MAKE RULES WITHOUT MAKING ENEMIES

how to draw the lines without drawing fire

MAKE RULES WITHOUT MAKING ENEMIES

how to draw the lines without drawing fire

by
Shay Hill

First Edition: March 2023

Dedication

At every audit, after every investigation, during every “roll out” of a new procedure, friends in operations would assure me, “Just tell us exactly what you need, and we’ll make it happen.”

This book is for you.

Table of Contents

The Path

A father and his young daughter walk together along a path through the forest. From time to time, the daughter wanders away from the path—as children are inclined to do.

The father gently reprimands his daughter, reminding her repeatedly to stay on the path as instructed.

This has no effect, but the father is patient. Again gently, he reminds his daughter of the warnings he's given her. "Do you remember, daughter, what I said would happen if you left the path?"

The daughter dutifully replies, "If I leave the path, I could get lost, and you would miss me very much."

The father, hearing his own words repeated back to him, is confident that his daughter has listened and will now stay on the path.

But again, the daughter wanders away from the path.

Fortunately, the father has had some communication training. He starts by encouraging his daughter, "You've done very well on this walk. I know we are farther from home than we've ever been, but you're being brave and strong, and I am proud of you."

… before harshly—but unavoidably—leveraging her own fear against her.

> Father: What would happen, daughter, if you lost your way?
>
> Daughter: I wouldn't know where to go.
>
> Father: And how would you feel?

Daughter: I would be afraid.

Father: What would you be afraid of?

Daughter: That I couldn't find my Daddy.

Father: What would happen if you never found your Daddy?

Daughter: I would miss him, and I wouldn't know where to go.

Father: Where would you sleep?

Daughter: I don't know.

Father: What could happen to a child in the forest after dark?

Daughter: …

Father: What could happen to me when I was looking for you after dark?

The daughter begins to sob, and the father stops, satisfied that his daughter, now adequately chastened, will stay on the path.

But again, the daughter wanders away from the path.

Now the father is gripped by *his own* fear. He is terrified that his daughter will be lost, possibly worse, if he cannot convince her to stay on the path.

He ~~offers her a bribe~~ establishes an incentive program.

Father: Here's the deal; if you stay on the path, I will buy you a new toy. Would you like to have a new toy?

Daughter: I would love to have a new toy, Daddy!

Father: Then we agree, you stay on the path, and Daddy will buy you a new toy.

The father holds out his large hand, and the daughter shakes it, feeling quite grown up and responsible.

But again, the daughter wanders away from the path.

The father panics. He ~~resorts to threats~~ develops a performance improvement plan.

Father: If you leave the path again, I will take away your toys and throw them in the fireplace! Do you understand?"

Daughter: Yes, Daddy.

Father: Then we're clear? If you wander away from the path again, Daddy will throw all your toys in the fireplace.

To make sure, the father drafts an agreement. He removes a pen and paper from his backpack, writes down the arrangement with two signature lines at the bottom, signs the top line, then offers the paper to his daughter.

He *encourages* the daughter to sign, but not before *encouraging* her to read the document aloud and answer a few follow-up questions to make sure she understands the agreement.

The daughter reads the document aloud, answers the questions, and signs, no longer feeling at all grown up.

The father didn't want to go this far, but he's confident that *if* his daughter can feel the same fear he's feeling, she will "stick with the program" and make it home safe.

The daughter is indeed afraid and does not want her toys thrown in the fireplace, but she continues to stray from the path.

The father loses his temper. "What do I have to do to convince you? Don't you care about yourself? Don't you care about me? Do you *want* to be eaten by wolves? WHY WON'T YOU STAY ON THE PATH?!"

The daughter, terrified, replies tearfully, "Daddy, what is a path?"

What is a Path?

The daughter wanted to please her father and she wanted to keep her toys and she *didn't* want to be eaten by wolves, but she couldn't see the path.

It is the same with our front-line employees. They want to get paid, they want to be safe, they want to succeed, and they want the organization to succeed, but they still have to see the path. On too many unfortunate occasions, industry workers have literally died to succeed, to meet a deadline, or to please a mentor. So why do we have to work so hard for compliance with our management system?

The hard thing about following rules is that there so many of them:

- the rules in the management system
- the "rules" in the org chart
- the "rules" in the job description
- the "rules" in the training
- the "rules" of a trade
- the "rules" of social interaction and hierarchy
- the "rules" of professional advancement and reward
- the "rules" of productivity
- the "rules" of biology, focus, and mental health
- the "rules" of an 8-hour or 12-hour workday

It's a lot to keep track of, and if you want the *opportunity* to follow a management system, you have to have a job, which means *none* of these rules can be ignored. As rule *makers*, we might believe that some of these rules are more important than others, but we don't write rules for the unemployed.

If you read the introduction with the father and daughter walking through the forest, I will forgive you for believing "the path" in my analogy was the first rule set: the rules in the management system. That is our principal job, and a lot of this book is dedicated to it, but the rules are not the path.

Do or Die

One of my least favorite tropes in movies is the Gettin' Shit Done Guy.

We're shown a "situation room" in panic. All is lost. Most of the poor souls in the room are bent over machinery or computer desks, desperately working to keep catastrophe at bay. Those are the lucky ones. They occasionally look up from what they're doing just long enough to deliver increasingly bad news to the one person having a worse time than they are, a Beleaguered Person in Charge who is obviously in over their head and left with no alternative but to nobly "go down with the ship."

Then in walks Gettin' Shit Done Guy. GSDG asks for the "sit-rep," and the beleaguered person in charge lays it out:

> **Beleaguered Person in Charge**: We struck an iceberg; there's a twelve-foot diameter hole in the starboard hull; we're taking on water fast; most of the lifeboats are on fire; man-eating sharks are circling the ship, and just before we lost power, we spotted enemy submarines approaching from the South.
>
> **Gettin' Shit Done Guy**: Who is the best carpenter on the ship?
>
> **Beleaguered Person in Charge**: That would be Larry, sir.
>
> **Gettin' Shit Done Guy**: Bring me Larry!
>
> *Larry arrives, carrying a hammer*.

> **Gettin' Shit Done Guy**: How long would it take to gather driftwood and build new lifeboats?
>
> **Larry the Carpenter**: Six months, sir, *if* we had nails.
>
> **Gettin' Shit Done Guy**: … You have 60 minutes.

When we ask for big things without considering (or even realizing) what it might take to deliver those things, we are playing Gettin' Shit Done Guy. Sometimes it can be hard to sympathize with Larry the Carpenter because we didn't get as far as we have by saying "No" when someone asked *us* to do impossible things. We got them done! We figured them out! We made the impossible possible by working for free on nights and weekends! Sometimes, at least.

In the movies, this attitude often saves the day. In reality, it's almost silly. In reality, Gettin' Shit Done guy gets fed to the sharks. We like to believe that *our* purview is so important that it doesn't need justification or explanation. *Our* purview is not a priority, it's *the* priority. How's that working out for you?

If "do or die" actually got it done, we'd move up all our deadlines by a week or two, cut all our budgets by eighty percent, lay off half our staff, tell everyone they have no alternatives, then sit back and wait for the miracles to happen.

We aren't "the problem," because we're suffering under the same pressures as everyone else. We have our own rules to follow and our own corners we end up cutting. As a professional, you only work 50 or 60 hours a week, and you probably have a few hundred hours of obligations, so you get done what you can and try to keep the important plates spinning. But at what cost?

Is Fear a Good Motivator?

The Fourth R

When I was in school, my teachers taught me "the three Rs":

Reading: The ability to understand written language, the foundation of all learning that allowed me to acquire knowledge and engage with the world around me.

wRiting: The ability to express myself through written language, an essential skill for communication, problem-solving, and critical thinking.

aRithmetic: The ability to perform mathematical operations, including basic calculations and problem-solving. The keys to the physical universe.

Those are significant, important concepts. It's easy to look at the Three Rs and conclude they are relatively complete, that nothing academic could be added without taking something away. But one academic R is missing, taught to Roman children but now long-since relegated to self-help books and management seminars: Rhetoric.

Pathos (In the trade, we get here by "preaching safety"): A rhetorical device that aims to evoke an emotional response from the audience. Pathos is often used in speeches, advertisements, and other forms of communication to appeal to the audience's emotions and create a connection with the audience. As professionals, we may use pathos in a speech by sharing a personal story to evoke empathy and connect with the audience on an emotional level.

An organization may use pathos in various media by showing images of crying children to evoke sadness and a desire to avoid regret.

Pathos can be effective in persuasion because it appeals to people's emotions and values, which can motivate them to take action.

Logos: (In the trade, we get here by "teaching safety"): A rhetorical device that appeals to logic and reason. Logos is often used in arguments and persuasive writing to appeal to the audience's sense of logic and reason. Logos relies on facts, evidence, and reasoning to make a case. Using logos, we present logical and rational arguments to support our claims. We may use statistics and other evidence to support our point of view. A professional might use logos in a presentation by providing data to demonstrate the value of a new corporate initiative.

Logos is an effective way to persuade an audience because it appeals to our sense of logic and reason. However, it takes a lot of preparation and, for an authority figure, can feel like a retreat from the (unassailable) moral high ground.

Ethos: (In the trade, we get here by building a culture): A rhetorical device that appeals to the speaker or writer's credibility, trustworthiness, or charisma. Ethos is often used in arguments and persuasive writing to establish the credibility or authority of the speaker or writer.

Ethos is the core or familiar in-crowd-out-crowd appeals like, "I don't know how you did it on your old crew, but on *this* crew, we put our tools away after we finish using them," or "I don't think this is the best crew, I know it is. And if you want to be a part of *my* crew, you will have to show up on time and act like you give a crap."

The Levanthal Tetanus Study

In 1965, psychologist Howard Levanthal engaged a group of Yale University students to evaluate the persuasive strength of a pamphlet on the dangers of tetanus. To encourage participation, the university offered a free tetanus shot to any student who participated.

Each student was given a pamphlet to review, but not all pamphlets were the same. Some pamphlets were "low fear," others "high fear."

The "high fear" pamphlets graphically described (with photographs of infected children) the negative effects of tetanus, and these are frightening. Tetanus starts with muscle stiffness (whence the common name "lockjaw" is derived). This progresses to violent whole-body seizures and possible death. An intermediate symptom, "rictus grin" (whose nature horror-movie fans will have no trouble inferring from the other symptoms), may have been the inspiration for DC's Joker.

Images are every bit as disconcerting as you may imagine. I will spare you, but it is easy to see why the "high fear" pamphlet group replied overwhelmingly to the study's meta question: "Will you get a tetanus shot?"

Many students, particularly the "high fear" group, committed in writing to get a tetanus shot. Levanthal had successfully motivated students to "check a box."

As we too frequently evaluate seminars, interventions, quarterly reviews, stand-downs, and "this is your last chance" meetings, the Levanthal Tetanus Experiment was a tremendous success. Handshakes and back claps all around. I cannot find a reference, but

I wouldn't be surprised if the most fervently fearful were offered positions in Levanthal's research department.

But … what happened *after* the handshakes is less encouraging. Only 3% of participants actually went to get a tetanus shot.

WHAT WENT WRONG?

Malcolm Gladwell later recreated the Levanthal Tetanus Experiment.

Again, students were asked to evaluate a pamphlet and offered a free tetanus shot in exchange for their participation.

Again, each student was given a "low fear" or "high fear" pamphlet. But this time, the pamphlets had a small addition: a map to the student health center along with times when tetanus shots would be available.

Again with the reviews, again with the commitments, and again these were all pretext. The *real* data was how many students actually went to get a vaccination. The result? 28% of students went to get a tetanus shot after reviewing the pamphlets, far higher than the 3% from Levanthal's study.

Another perhaps surprising fact: students who received the low-fear pamphlets and students who received the high-fear pamphlets both got shots at the same 28% rate. The impact of information was overwhelming. The impact of fear wasn't even measurable.

BACK TO RHETORIC

So what happened? Were the students in the Levanthal Tetanus Experiment immune to fear? Should the study have done *more* to scare the students? Is fear a terrible motivator in general?

We can only speculate, but I believe fear *is* a powerful motivator, and I suspect the pamphlets were scary enough. The problem may be that the students were already scared when they walked into the study. Most of us have heard of tetanus and recognize it as a thing we want to avoid. Even if we don't know the specifics, we know to be afraid. As much as fear works, there was little to be gained by making the students *more* afraid.

While fear messaging can effectively promote behavior change or drive action, it can also have negative consequences. Over-reliance on fear messaging leads to anxiety, stress, and helplessness, undermining motivation and producing adverse outcomes such as learned helplessness or depression.

Fear messaging is particularly and clearly problematic, but all emotional appeals have similar limitations.

If you love your family, you'll follow the rules.

Many employers have a 100% glove policy for tradespersons. This means that if you are on the job site, you are required to wear gloves, whether you are swinging a hammer or scribbling on a clipboard. But this wasn't always the case.

Before gloves were mandatory, both gloved and non-gloved personnel were assured that we were good people who cared about our safety, our families, and our environment. We were "good people who cared about our safety, our families, and our environment" because we followed the safety rules in place at the time.

Safety messaging was everywhere. Particularly poignant were pictures of crying children, ostensibly orphaned by the unsafe behaviors of their working parents, in the meeting rooms and group areas. Captions on the pictures read, "Please come home safe, Daddy," or "What will you say to their children?"

But we didn't have to worry about *our* children, because we were the good guys who labeled our chemicals, tied off our tools, and filled out our paperwork. *Our* children were safe.

Then the rules changed. Sometimes temporarily for a client and sometimes permanently because someone had a better idea. Gloves were a big change, then Fire-Resistant Clothing, then new processes, then tablets. It was a common occurrence to leave your job a hero and come back a "dinosaur", "bubba", or "good ol' boy". Your safe behavior from the day before wasn't safe anymore. Yesterday, you might have been disciplined for wearing gloves in some work areas; today, you're an uncaring, unsafe "cowboy" if you don't. Yesterday,

you might have been safe at six feet; today, it's four; tomorrow, it might be six again.

Too often, the messaging didn't change. We were expected to feel fear and shame on queue when the rules changed. And, because progress isn't linear, we were often expected to turn off our fear and shame on queue when the rules changed back.

This kind of change can be tough on experienced employees, who are natural workplace leaders *until* they find themselves in a context where their experience and knowledge have no value, or even negative value. This is the condition we create when we reduce compliance to moral absolutes, when we solely rely on pathos.

But too often, we corporate types *like* this condition because if experience and knowledge aren't allowed into the conversation, our *lack* of experience or knowledge aren't allowed into the conversation. We don't have to lower ourselves to talk about *how* something might be accomplished; "shall" is enough.

*See, Fear **is** a Motivator*

This situation is partly created by *our* fear of engaging with other people as equals. I can deliver the same pulpit-pounding safety *sermon* across multiple industries and as a response to nearly any situation. But it can be difficult to reason and teach. Teaching requires preparation, cooperation, adaptation, learning from others, and a real potential for failure. It is a frightening thing, but that's how we get people to line up for a tetanus shot.

We have to loosen our grip on the idea that people don't work safely because they aren't adequately motivated or virtuous. We have to work on our own virtues. The Levanthal Tetanus Experiment suffers from the same deficiencies as many other behavioral studies (lack of rigor, low sample size, etc.). But if you've found a way to motivate people solely by showing them morgue slides and promulgating "shalls" at them, then I need to be reading *your* book.

Keep preaching. Keep motivating. It's critical, but don't forget to make *yourself* a little uncomfortable too.

This book is about some of the worst ways I've been wrong in my career and the tools I found to become less wrong. Some of them are so unusual to our industry that trying them (not hearing them, not reading them, *trying* them) just once has the potential to change your perspective forever. There are a few chapters on mathematical intuition. That's another thing that can be frightening. If you feel uncomfortable, embrace it. Slow down, maybe reread a paragraph or two. Numeracy is another thing they forget to teach us in school. They were too busy teaching us equations.

The Treadwell Effect

You can't cure acrophobia (fear of heights) by jumping out of a plane—at least, that's what the psychiatrists tell us—so let's put clinical phobias aside and constrain our conversation to everyday fear: fear of public speaking, fear of rejection, fear of walking under suspended loads.

Timothy Treadwell was an American bear enthusiast and environmentalist who lived among grizzly bears in Alaska for thirteen summers. He was born in New York City in 1957 and had a tumultuous childhood, marked by drug addiction and a series of arrests. But in the 1980s, he turned his life around and moved to California, where he began working as a guide for bear-watching tours.

In the early 1990s, Treadwell became increasingly interested in living among bears and began traveling to Alaska each summer to spend several months in the wilderness. He quickly became known for his close interactions with the bears and for his ability to seemingly coexist with them without incident. Treadwell would often film his experiences and use the footage to raise awareness about the importance of protecting grizzly bears and their habitat.

Treadwell's approach to living with bears was controversial, with many experts in the field criticizing him for getting too close to the animals and potentially putting both himself and the bears in danger. Despite these criticisms, Treadwell persisted in his efforts and even wrote a book, "Among Grizzlies: Living with Wild Bears in Alaska", in which he detailed his experiences and advocated for the protection of grizzly bears.

But he still had limits. Treadwell knew to stay away from the bears in the Fall, when the bears are fattening up to hibernate. Whatever else you might say about him, Timothy Treadwell was an expert on bears.

In September 2003, he left the camp he'd been sharing with his girlfriend, Amie Huguenard. It was getting late into the year, food was scarce, and the bears were acting more aggressive than usual.

Amie wrote in her diary that she was terrified of the bears and wanted to leave.

But, on the way out of town, they got into an argument at the airport over ticket prices. Timothy and Aime stormed out of the airport and returned to their camp in Katmai National Park (Alaska) to escape the madness of civilization. It seems they were more mad than they were scared.

This is a man who arguably knew more about bear behavior than any other person on Earth, but he brought the love of his life to sleep beside grizzly bears in a tent. Of course, they were attacked, killed, and partially eaten by the bears.

Tumultuous adolescence aside, Treadwell was a sane and capable man in most respects. He had a profound impact on the way people think about bears and their conservation. Through his films and writings, Treadwell gave people a glimpse into the lives of these magnificent animals and helped to raise awareness about the importance of protecting bears and their habitat. He was productive, articulate, and intelligent. The trap that killed him is waiting there for all of us: with repeated exposure, memory replaces imagination.

Treadwell started with a healthy fear of bears, but with each encounter—until the last—he established a pattern: bears don't kill. Over several years, the pattern overtook education, imagination, and common sense. Memory replaces imagination.

Ironically, that pattern was more or less accurate. Timothy and Aimee were the first humans attacked in the park's history. But Timothy wasn't the only one losing his fear. It was later determined that the bear had been habituated to humans due to Treadwell's close interactions with the animal. Memory replaces imagination.

It can happen with bears, and it can happen with equipment. Comfort can be one of our most tragic contributing factors.

But there's an opportunity here for us, because there are fears we *don't* want in our organization. We don't want personnel to fear change, we don't want personnel to fear management, and we don't want personnel to fear "standing up" for our values.

We can overcome these fears with the Treadwell Effect. It only takes patience and time. We have to let personnel experience and remember consistent, positive change, not "this time it will be different, I swear!" trainwrecks.

Frequency

I worked for a while auditing land-rig operations for a drilling contractor. I would visit our entire fleet of land rigs and our full, diverse assortment of land-rig clients. Each client, and most well sites, had their own safety orientation, and I had to complete each to access and inspect the rigs.

These orientations typically involved watching a video and PowerPoint presentation … on a laptop … alone in a room … behind a closed door

… because no one else wanted to see it.

There seemed to be a competition between clients and well sites to find and present the most gruesome scenes of death and dismemberment, meant to encourage awareness of safety hazards in the workplace. One thing I know for certain, I will never stop to wave to my family on the way out of a helicopter with a spinning rotor … never.

In safety, we like to use images, images of cooperation, images of inspiration, images of organization, images of process, images of trend lines, and images of severed limbs. Some prefer friendly images; others want to shock you.

I am going to use images in this chapter to help you unlearn some popular ideas about frequency, trends, and analytics. You might want to shut the door and turn down the volume, because some of this is going to be a little hard to take.

A Likely Story

We live and think—and spend—under a terrible misconception: that things generally happen where they are likely to happen. Rain happens in rainy places; safety incidents happen on bad crews; if you come across a camel, you're probably in the desert.

What do you mean misconception? All of those things are true!

All of those things are true … *under certain impossible conditions*.

Let's start with the first two.

HOUSTON RAIN

I live in Houston. At 53" average annual rainfall, we are a rainy city. Here is Houston compared to the five driest cities in Texas:

- Houston (53.34" annual rainfall)
- El Paso (8.19" annual rainfall)
- Lubbock (20.39" annual rainfall)
- Midland (20.91" annual rainfall)
- Amarillo (22.06" annual rainfall)
- San Angelo (22.65" annual rainfall)

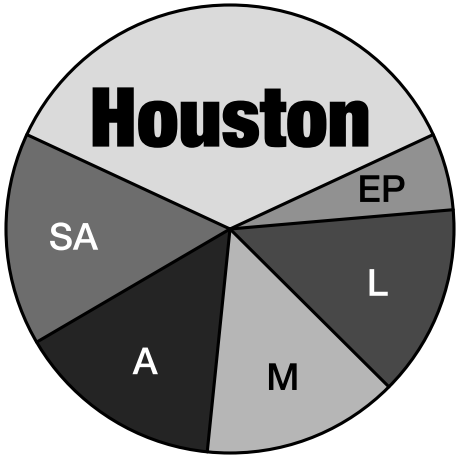

Houston is a rainy city

In a universe where I could only wake up in one of these cities, I might follow this heuristic: if I wake up and it's raining, assume I am in Houston.

If you hadn't seen the chart, you might have gone along with "If it's raining, I'm in Houston," but even *this* chart, even with "Houston" in big, bold letters, is clear enough to show you the problem with my heuristic. Houston may have more rain that *any* other city on the chart, but it doesn't have more rain than *all* other cities on the chart. In other words, when it's raining, I'm probably *not* in Houston.

Maybe you caught it this time, but how many times have we followed this exact heuristic when apportioning resources (and blame) after a safety incident? Chart or no chart, the idea persists that even one data sample (an incident) contains *some* truth, and that, if we chase that truth, we will over time apply our resources

(and blame) proportionately to the incident risk (and blame) present across our organization.

The idea that a small data sample still contains some truth goes by many names: availability heuristic, law of small numbers, illusion of validity, and more. These are common cognitive biases[1][^5] that frequently lead to incorrect judgments and decisions. They are all related to the tendency of people to make assumptions or draw conclusions based on limited or unrepresentative data.

The *availability heuristic* describes our tendency to overestimate the likelihood of events based on how easily they come to mind, rather than on a representative sample of information. A classic example is how the beach feels more dangerous after a shark attack, even if the statistics haven't changed. Even a statistician will have problems with this one. Statistician or no, your pulse is going to speed up once you wade in above your knees.

The *law of small numbers* describes our belief that a small sample is representative of a larger population, even if it is not. An example here is how easily our minds register one rumor from one person as "now people are saying X." This leads us to create systems that "nobody" asked for.

The *illusion of validity* describes our belief in the accuracy of predictions or decisions despite limited information or evidence. This is getting closer to the problem, but we're still not quite there.

1 The term "cognitive bias" is often used to mean "things other people have problems understanding." The true meaning is "things human beings (all of us) have trouble understanding." THIS MEANS YOU—AND ME.

We're not quite there because even as you read these, you're probably still thinking,

> "Those problems occur when we give information too much weight, but all information still has *some* weight, right?"

or,

> "The people in the examples weren't wrong, just less right than they thought they were,"

or,

> "Where there's smoke, there isn't *always* fire, but it's *always* a good idea to check."

Even though I know better—and even though I'm the one who wrote this chapter—it feels true as I type it out. It feels like there is *some* signal in the noise, and that we are better off paying attention to it. In a binary choice like Houston vs. El Paso, it may be true *enough*. It's nearly always true *enough* if we're omniscient. If you only ever deal with binary choices, or if you are in fact omniscient, carry on. The rest, keep reading.

OVER TIME

The pie chart looks enough like a dartboard that we can imagine the universe throwing darts at it. Each dart might be completely random, but *over time*, more darts will land in Houston than in any of the other cities. “If it’s raining, I’m in Houston” might not work every time, or even most of the time, but it will be right more than any other guess *over time*. All we have to do is keep betting on Houston, because the odds are clearly in our favor. Right?

Let’s look at the Texas rain chart again through the lens of our heuristic, “If it’s raining, I’m in Houston” or “Whenever it rains anywhere, let’s wager we’re in Houston.

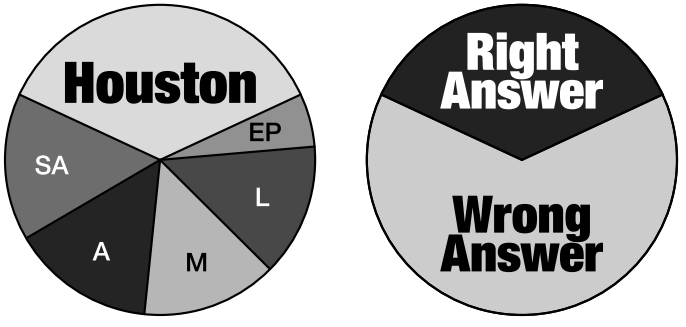

Our bet isn’t looking so good. Houston is the best choice, but the odds are still against us.

I suspect you caught this one when you saw the first chart. That’s good. We got to walk together through familiar territory. Now …

Small Data

We agree that it’s tough to make an inference from one incident. But what if we had a year of data, or ten years? Then the little bit of “rightness” in every data point would add up. Right?

We get the impression, as we watch data trickle in, that each observation brings us closer to the truth. This is rarely the case unless you’re dealing with hundreds (at least) of data points.

Let’s look at some data: five-thousand tosses of a fair coin. We’ll break these five-thousand tosses into one-hundred groups of fifty. In experimental terms, each of these fifty tosses is a “trial.”

Here are the results. Each wavy line is a record of fifty coin tosses. When the coin landed on heads, the line moved up. When the coin landed on tails, the line moved down.

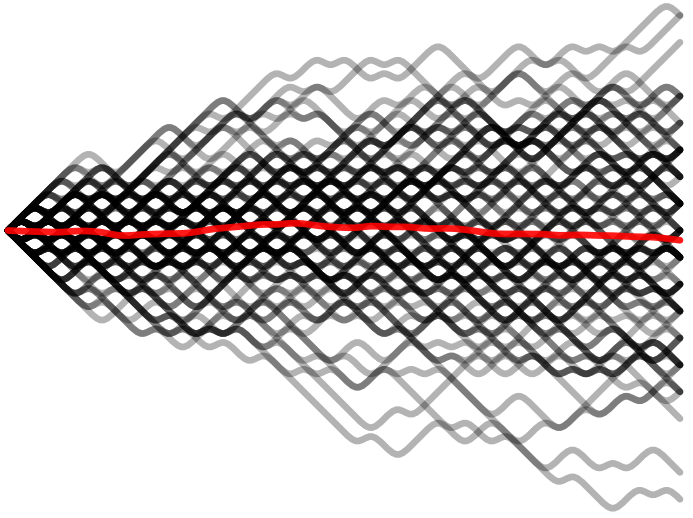

100 trials, 50 coin tosses each

The red line is the average of all one-hundred trials. This line points pretty squarely down the middle, which is what we'd expect from a fair coin. This line represents how we *think* we see data over time.

We have names for this point of view. You've likely heard of *regression to the mean*. This is the statistical rule that an uncharacteristic observation will likely be followed by a more representative observation. This checks out. This is the first formula in the book, so let's start with an easy one.

$$(1 + 2 + 3 + 4 + 5 + 6) \ / \ 6 = 3.5$$

the average roll on a 6-sided die

The average (mean) roll of a six-sided die is 3.5. This is the first equation in the book, so let's start with something simple.

If you roll a die and see a point far away from this average, let's say 6, you have

- 0 ways to see something farther from average on the next roll
- 2 ways to see something equally far from average (1 or 6)
- 4 ways to see something closer to average (2, 3, 4, or 5)

This gives you a 2:3 chance to see something more representative than 6 on the next roll. Over time, this *regression* means that the average of random observations tends to move toward the expected value (theoretical average) over time.

All true, but the situation we're describing is statistical "God Mode." In reality, we aren't observing one-hundred trials, we are traveling along the path of *one* trial, and we don't know which one. That last sentence is worthy of a Post-It on the corner of your monitor. If I went too quickly, here's another picture to clarify. There are no other paths to average; we are on *one* of these paths.

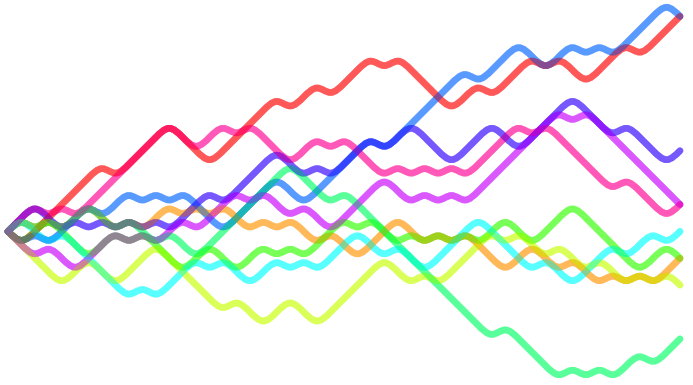

we only see one trial

At any point on a path, the expected value (expected average in the future) is equal to where we are currently standing on the path. We have a name for this too, *normalization of deviance*, it applies to rule-breaking, and it applies to random observations. There are still ways to estimate how far we may be from average, but they're trickier and require a *lot* more data. More in the chapter on statistical tests, but before that, just a little more bad news.

Deviance is normal. In the coin-flip images above, trials with an above-average ratio of heads combine with trials with a below-average ratio of heads to give us an average ratio of heads. This could give the impression that the *average* trial is, well, average. I', afraid not. I will show the trials again, this time averaging not by value but by absolute difference from the truth.

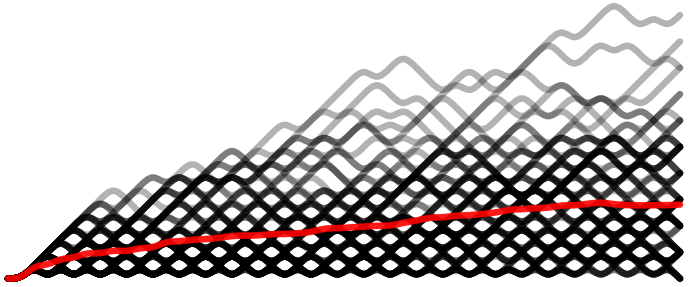

deviance is normal

The average path (the path we're probably standing on) is wrong.

If we had all the information the chart gives us, this might still work out. We'd just have to make sure the payout for "right answer" was high enough to cover the odds. Unfortunately, we don't know that Houston has 2.4 times the average annual rainfall of Amarillo. We only know where the darts hit.

Doubling Down, Tripling Down, etc.

Each of the pictures in the previous chapter was created to highlight a specific idea. To accomplish this, each picture gave you a lot more information than you would have in real life. Each gave you a taste of omniscience. These pictures represent data as we encounter it in a classroom. In work and life, we have a lot less to work with.

Let's look at five facilities in an imaginary company. We'll call them Kish, Eridu, Mari, Uruk, and The Pit of Death. The Pit of Death is TWICE as dangerous as any of the other four, but you don't know that. Let's look at our incident rate and try to figure it out.

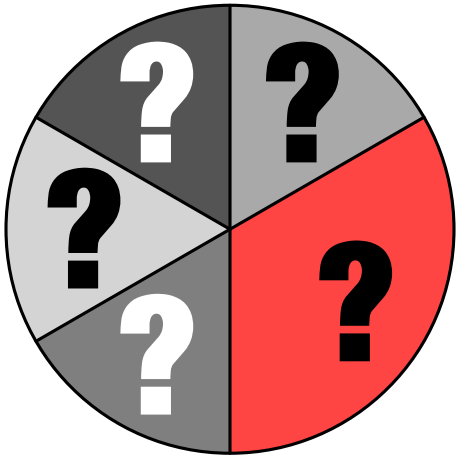

We're going to simulate incident distribution with a six-sided die representing our five facilities. The names of Kish, Eridu, Mari, and Uruk, your safer facilities, are each painted on one side of the die. Pit of Death is painted on *both* of the remaining sides. This means

that, on average, every third roll will land on Pit of Death. This is an absurd tilt. Pit of Death has worse odds than any game in Vegas, Monte Carlo, Macau, or Canto Bight. Pit of Death is the reason you pay the license for Power BI. Pit of Death is the subject of your next quarterly review. Pit of Death is a *sure* bet to lose. So let's roll …

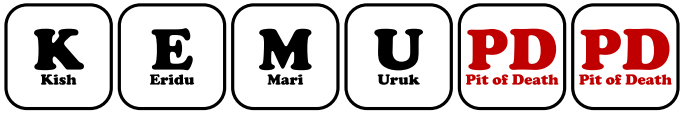

There's some math here, but I'm going to do it for you. We'll walk through three rolls. That's three incidents across five facilities, a 60% per-facility incident rate. Our imaginary company is certainly out of business. But we have data!

There are 216 ways three dice can land. If we saw every roll, we would see 516 candidates for "horribly dangerous facility." Sometimes these candidates would appear as singles (all dice land on different sides), sometimes as doubles (two dice the same, one different), and occasionally as triples (all dice the same). By "the same," I mean the same facility, even if that facility is painted on two sides. I'll repeat this again because it is easy to forget:

> We don't get to see all the rolls. Just one.

The challenge is to look at *one* roll and make a data-driven decision. Let's play analyst for a moment and define what Hold 'Em players would call the "Nash Equilibrium" for this challenge.

In Texas Hold 'Em, "Nash Equilibrium" refers to a state in which no player can improve their outcome by changing their strategy, given

the strategies being employed by the other players. In other words, perfect play. In our case, the only other players are risk and luck, so our perfect play is simple.

- If any side shows three times, we wager that is our dangerous facility.
- If any side shows twice, we wager the side showing twice is our dangerous facility.
- If all dice land on different sides, we assume our dangerous facility is one of them and pick randomly.

Following this, we would see

indicator		attempts		result
96 times PD appears as a single	%	396 singles	=	We are misdirected 76% of the time.
48 times PD appears as a double	%	108 doubles	=	We are misdirected 56% of the time.
8 times PD appears as a triple	%	12 triples	=	We are misdirected 33% of the time.

Of the 216 possible combinations, only the 12 triples would have a fair chance of steering our algorithm correctly. The trace is there, but we'll probably never see it. Our incident has already put us out of business, and we *still* probably don't know that Pit of Death is worse than the other four.

As our company grows, our chances get dramatically worse. With a ten-facility scope, we'd be waiting for quadruples or better (1 in 50) to see a fair chance of steering our algorithm correctly.

And our company isn't waiting with us! Each facility, supervisor, or team is another side on our die. With rare-event data, there is no such thing as "over time."

ONE right answer competes with ALL wrong answers for a chance to appear in our data.

Brevity

Here Bygynneth the [Chapter on Brevity]
Whan that Aprille with his shoures soote,
The droghte of March hath perced to the roote,
And bathed every veyne in swich licóur
Of which vertú engendred is the flour;
Whan Zephirus eek with his swete breeth
Inspired hath in every holt and heeth
The tendre croppes, and the yonge sonne
Hath in the Ram his halfe cours y-ronne,
And smale foweles maken melodye,
That slepen al the nyght with open ye,
So priketh hem Natúre in hir corages,
Thanne longen folk to goon on pilgrimages,
And palmeres for to seken straunge strondes,
To ferne halwes, kowthe in sondry londes;
And specially, from every shires ende
Of Engelond, to Caunterbury they wende,
The hooly blisful martir for to seke,
That hem hath holpen whan that they were seeke.[2]

That's a taste of Chaucer to start us off. And believe it or not, it's in English. Or at least it used to be. This is Old English and it's a good approximation of how "New English" might look to a non-native speaker. I read a little Spanish, and this is how Spanish looks to me. If I take my time, concentrate, and re-read a few times, I can kind of, sort of, almost follow along.

2 Chaucer, 1400

Gatsbys

The Great Gatsby (we all read it in school) has 47,094 words. Each 50-thousand words of documentation you want your personnel or contractors to read is one Gatsby.

You are most likely reading this at somewhere between 200 and 300 words per minute because I've worked to keep it breezy. Now, quick, how many cities can you name from the Pit of Death exercise? The kind of reading you are doing now is great for a general impression, but you wouldn't want to be tested over it.

The average *technical* reading speed is only 100 words per minute, slower still if you expect to be tested against it. So, each Gatsby in your management system is about eight hours of dedicated reading time multiplied by the number of people you expect to read it.

This can be difficult to appreciate because our manuals are not technical reading **to us**. Let's try it. Here's a puzzle (that means there is a test at the end) a third-grader could work out:

> Each of nine professionals (a fireman, a police officer, an accountant, a dentist, a lawyer, a medic, a teacher, a plumber, and a farmer) is responsible for exactly one of nine responsibilities (broken legs, crops, toothaches, leaky faucets, taxes, criminals, fires, lawsuits or students). No two professions share a responsibility.
>
> - Criminals are the responsibility of the police officer or the medic.
> - The accountant is responsible for broken legs, taxes, or fires.

> - Lawsuits are the responsibility of the lawyer or the teacher.
> - The teacher is not responsible for lawsuits.
> - Leaky faucets are the responsibility of the dentist, the plumber or the police officer.
> - Students are the responsibility of the teacher, the police officer, or the farmer.
> - Broken legs are the responsibility of the farmer or the medic.
> - Fires are the responsibility of the teacher, the accountant or the fireman.
> - The dentist is responsible for toothaches, taxes, or crops.
> - The farmer is responsible for crops or toothaches.
> - Crops are not the responsibility of the fireman, the police officer or the dentist.
> - The fireman is not responsible for toothaches or taxes.
>
> Who is responsible for taxes?

You probably didn't even try to work it out. And you wouldn't have to. You are an expert on this material. You can quickly scan this material and identify any mistakes. You knew before you started which of the nine professionals was responsible for taxes.

Let's try it again as a novice.

> Each of nine professionals (a hayward, a hostler, a weirkeeper, a rubricator, an ackerman, a webster, a

wainwright, a peruker, and a knacker) is responsible for exactly one of nine responsibilities (fish traps, titles, looms, carcasses, wagons, plows, fence posts, mules or wigs). No two professions share a responsibility.

- The ackerman is responsible for looms, mules, or plows.
- Looms are the responsibility of the rubricator or the webster.
- The wainwright is not responsible for mules, titles, or wigs.
- The wainwright is responsible for wagons or fence posts.
- Mules are the responsibility of the knacker, the hostler, or the webster.
- Carcasses are not the responsibility of the hayward, the hostler, or the peruker.
- The rubricator is responsible for wagons, titles, or wigs.
- The hayward is responsible for fence posts or fish traps.
- Fish traps are the responsibility of the knacker, the weirkeeper, or the peruker.
- Wigs are the responsibility of the weirkeeper, the peruker, or the hostler.
- The peruker is not responsible for fish traps.
- The knacker is not responsible for fish traps.

Who is responsible for carcasses?

That's what a strange manual, even clearly written, looks like. And that's the kind of speed and retention you can expect with new material.

You probably won't try to work this one out either. This is a puzzle a third grader could solve, but looking at it is disheartening because we know even this simple puzzle will be a chore. When the brain is constantly engaged in cognitively demanding tasks, even basic tasks, it is hard to focus on the details of a process. It can be difficult to remember important information or follow the steps of the process in the correct order. If you're not convinced, try the puzzle. You'll find the solution, but you're sure make at least one mistake on the way.

You know about these issues already, because if you're reading this book, you've seen or maybe shown the gorilla-suit video, the elevator video, or something similar to enlighten yourself or others about our capacity for distraction. We recognize the problem. So why do we still have 10-Gatsby safety manuals?

And still, I haven't used any difficult technical language (I've indented this next passage to make it easy to skip):

> A comprehensive assessment of the food-preparation area must be executed to ensure compliance with established industry standards and regulations. This assessment must include, but is not limited to, examination of equipment, utensils, and surfaces for cleanliness and proper functioning; evaluations of food storage and handling protocols; and inspections of the physical structure of the food-preparation area for code compliance. Any non-conformities must be documented and reported in accordance with established protocols. The assessment

> must be performed utilizing appropriate industry standards and guidelines. Procurement of bread must be initiated, with specific attention paid to ensure that the bread is of the appropriate texture and consistency. A thorough examination of the bread must be conducted to detect the presence of mold and/or indicators of non-freshness. This examination must include, but is not limited to, visual inspections, olfactory evaluations, and tactile assessments. Any deviations from the expected norm must be documented and reported accordingly. The examination must be performed in accordance with established protocol. Utilizing one (1) legally obtained and responsibly sourced spreading implement, such as a butter knife, peanut butter must be distributed on one (1) slice of the bread in a manner consistent with desired coverage. Inspect such implement before use utilizing appropriate industry standards and guidelines. Any deviations from the expected norm must be documented and reported accordingly. One (1) banana, previously obtained through authorized means, must be placed on a second slice of bread. The two (2) slices of bread must be joined together, creating a sandwich, which must then be cut in half with an approved cutting device which must then be cleaned, dried, inspected and returned to its proper storage cabinet and/or drawer and/or other appropriate and legally compliant storage location, with each half constituting one (1) unit serving.

One paragraph of that nonsense, and I'll bet you thought of putting this book down. Enough fun. Now, how is this hurting us?

The Expert Problem

Eventually, we process professionals become fluent with any procedure, no matter how long or complex.

We develop fluency over time through a process known as "acquisition." This process involves repeated exposure to and engagement with a complex thing, which helps us to gradually build up our knowledge and understanding of it. As we acquire more information and skills related to the complex thing, our ability to understand and interact with it improves.

With repeat exposure and engagement, the complex thing becomes more familiar and less intimidating. Our brains begin to recognize patterns and connections between different aspects of the complex thing, which helps us to understand it more easily. We start to develop a set of mental shortcuts, or heuristics, that help us quickly process and understand new information related to the complex thing.

We no longer have to consciously think about how to interact with or understand the complex thing; it becomes an automatic process. We may no longer even realize that something is complex once we become familiar with it.

The process of acquisition is not limited to technical or scientific subjects, acquisition can be applied to any complex thing, like a wordy safety manual, postmodern literature, or the layout of a city. Have you ever asked for directions in a strange city? In five seconds, you're ignoring whomever you asked, just thinking, "I'll wait this out and then put it in my GPS." Ever taken advice from a hobbyist friend to try Vim, cook a brisket, flash your bios, or rebuild

a transmission … you'll know what it's like to be on the other end of "acquisition."

Sadly, we're the old codger at the gas station saying,

> "Now, this is where it gets tricky. You're gonna wanna take the left fork, but only if you see a big red barn on your left. If you don't see no barn, then you done took the wrong turn, and you're gonna have to go back and try again. If you do see the barn, though, keep going down that road until you see a big old oak tree on your right."

And the person hearing the directions will *never* be able to contribute to our "continual improvement." The theory goes that we experts write a process, non-experts execute that process, then everyone gets together to improve the process. You'll know this as Plan > Do > Check > Act. The concept is simple. So, why does it almost never work?

- **Time Constraints**: Complex problems often require a significant amount of time and effort to understand and solve. This can be difficult for contributors who have limited time and resources, especially if they are already committed to other projects.
- **Technical Barriers**: Complex problems may require specialized knowledge and skills that not all contributors possess. This can create technical barriers that prevent some individuals from contributing to the solution.
- **Communication Breakdowns**: When a problem is complex, it can be difficult to communicate effectively about it. This can lead to misunderstandings and miscommunication, which can reduce the number of contributors involved in the solution.

- **Frustration and Burnout**: Working on a complex problem is challenging. Contributors may become frustrated with the lack of progress and may eventually burn out, reducing the number of individuals involved in the solution.

All of these can be summarized under "Cost of Entry." If it takes reading and understanding eight hours of material just to have a productive conversation, you're going to end up with a very exclusive conversation.

And not only users get excluded. Management is often excluded as well. No management participation can mean no management engagement can mean our process slides to the bottom of the priority list.

We all Pay

Potential contributors aren't the only ones paying this cost of entry. Most of our users will never submit process improvements or even ask questions. They'll just show up and go to work and for the most part make a fair effort to conform to expectations as they understand them. That can be difficult even when the rules are simple, because of the number of rule *sets* I mentioned in the first chapter. I'll repeat those here.

- the rules in the management system
- the "rules" in the org chart
- the "rules" in the job description

- the “rules” in the training
- the “rules” of a trade
- the “rules” of social interaction and hierarchy
- the “rules” of professional advancement and reward
- the “rules” of productivity
- the “rules” of biology, focus, and mental health
- the “rules” of an 8-hour or 12-hour workday

It’s almost impossible for one person to keep all these rules straight, but fortunately, coworkers are there to help. If you forget to put away your tools, someone will probably remind you. Mix the paint wrong? They’ve got you covered. Forget who the boss is? The boss will refresh your memory. Put your shirt on backward…, leave your zipper down …, support the wrong football team … you’ll hear about it.

This applies to our management system too, but only if everyone knows the management system as well as they know how to mix paint. The directions on the side of a paint can fit, conveniently, on the side of a paint can. If your management system is long and difficult, workers won’t know if they’re following it or not. And they certainly won’t know if the person next to them is following it. You’ve lost your first line of defense. And you’ve lost your head start. If you can only manage to be compliant *after* consideration and review, then you’re looking at do-overs, re-thinking, double work … and a lot of people waiting around for all of that to happen.

Technical Debt

As hard as your manual may be to read, it’s much harder to write and maintain.

Technical debt refers to the cost of maintaining and updating the systems used in a business or organization. It is the trade-off between immediate benefits and future costs, which results from a lack of attention to technical design and quality.

Technical debt is a by-product of a fast-paced and rapidly changing environment, where businesses need to constantly keep up with changing client expectations and react to evolving tools, practices, and standards. In order to meet deadlines and release material quickly, system designers may take shortcuts and make compromises in the design and quality of their systems, leading to technical debt.

This debt accumulates over time as the system is used and maintained. The more the system evolves, the greater the costs associated with fixing problems and ensuring its reliability. This results in additional personnel, additional hours, additional rework, additional layers, additional pages, and additional misery.

There are two main types of technical debt: intentional and unintentional. Intentional technical debt refers to decisions made by management to take shortcuts in order to meet deadlines. This type of debt is usually incurred as a result of a conscious choice and is considered manageable if the company is aware of the costs involved and is able to plan for them in the future.

Unintentional technical debt, on the other hand, is not a deliberate choice but rather a result of a lack of technical knowledge or

experience. This type of debt occurs when developers are unaware of the potential consequences of their actions or when the company has no processes in place to manage and mitigate technical debt. This type of debt is usually more difficult to manage, as it is not easily recognizable, and it may accumulate faster than intentional debt.

Regardless of whether technical debt is intentional or unintentional, it has a number of negative consequences for businesses. These include increased costs, decreased efficiency, and reduced productivity. When a company incurs technical debt, it must spend more time and money on maintenance and updates, which results in decreased productivity and efficiency. This can result in a vicious cycle, where the company is forced to take on more technical debt in order to keep up with the demands of its business, leading to even greater costs in the future.

Hiding Places

Brutalism emerged in the mid-twentieth century as a response to the perceived excesses of pre-war European culture and the desire to create a more honest, straightforward architecture that expressed the structure and materials of the building. The style was often used for public buildings, including government offices, as it was seen as a way to convey a sense of solidity, permanence, and authority.

While some appreciated the style's honesty, more found it austere and unwelcoming. Many of these buildings were later torn down because the sharp angles and bare surfaces made the public uneasy. There is no correlation between brutalism and totalitarian governments, but

brutalism nonetheless conveys a sense of totalitarianism to many. This is the aesthetic of "A Clockwork Orange," and most consider it *too* effective at conveying authority.

It is therefore ironic that public buildings built in the brutalist style have a pattern of encouraging lawless behavior. The wide columns and beams that give brutalist buildings their characteristic imposing air also provide a lot of hiding places. The limited visibility and dark corners give criminals (real and imagined) a place to hide. This is a pattern with complexity.

I had a client with several sites around the US. Each site was required to have a series of site-specific plans: security plans, emergency response plans, waste management plans, etc. The core requirements for these plans came from well document standards, but the client had identified several opportunities to leverage these plans in new and useful ways. Unfortunately, these improvements were sprinkled across the management system. Page 17, paragraph 3 of no chapter in particular might have a requirement, "Document [something] in your site-specific chemical-management plan" or "Update your site-specific emergency contacts list with [something]."

This made the plans difficult to create and audit.

I worked with the client's safety leadership to consolidate these requirements into a separate chapter for each. After this consolidation, all requirements for site-specific emergency response plans could be found under one heading, "Emergency Response Plans." All requirements for waste management plans could be found under "Waste Management Plans," etc.

This made things simpler, but creating and auditing a site-specific plan still wasn't entirely straightforward. Some of the requirements

were references to standards with an instruction to "read this and create a conformant plan." Other instructions were copied and pasted from those same standards and contained requirements that did not apply legally, ethically, or logically to the work my client was doing.

We simplified and clarified the instructions, but we didn't stop there. We created a template for each site-specific plan. Where we were unable to find a conformant example, we created one. Each site-specific plan now had

- clear instructions;
- an editable template; and
- a conformant example.

So not only had we made the instructions easier to read and follow, we'd made it so you didn't have to read them at all!

But we didn't stop there. After clarifying all of this *for ourselves*, we were able to identify and remove redundancies. These were indeed "site-specific plans," but existing document-control, material-management, and authorization procedures in the company satisfied many of the requirements without a need for any site-specific specification.

After our overhaul, if you had the right people in the room, you could download a template and draft one of these plans in 15 minutes. The mystery was gone. By every conceivable measure, creating a site-specific plan was far easier than it had ever been.

BUT THAT ONLY MATTERED IF YOU WERE, IN FACT, CREATING A PLAN.

The *official* process we'd replaced was

- review the entire safety manual (multiple Gatsbys);
- review the standards;
- open up an empty document on company letterhead; then
- write a plan.

But the *actual* process we'd replaced was

- ask for a copy of another site's plans;
- cross out anything that doesn't look right;
- add a few pages to cover actual site-specific requirements.

The set of plans making the rounds had been consolidated into one seventy-five-page document. The client, the contractors, and the site leadership (presumably having each spent several hours reviewing this document) signed to follow the plans inside. And the document, with accumulating additions, was eventually forwarded to the next project for another round of additions and review. Lather. Rinse. Repeat.

The surprising thing—this won't be a surprise if you're in industry—the "actual" process *actually* worked. A few bright and knowledgeable people became experts on this document and were able to surgically alter it without too much effort. These bright and knowledgeable people also identified redundancies, but they addressed them by all-but ignoring portions of the plan set. They knew where to find the meaningful content.

So it worked, with some extra time and effort, but it more-or-less worked. However, over time, this document became a safety manual in itself. The safety manual was "the safety manual plus the plan set," and these didn't always agree. Site-specific policy at one site became company-wide policy as this plan set was reused across the organization. And the plan set faced the same challenges as the full safety manual: complexity, heedless growth, and technical debt.

While almost everyone appreciated the new process, a vocal few preferred the old way … and not without reason. When you "spell out" responsibility and accountability, everyone gets a clear picture of what is actually required, and it's not always a pretty picture.

The Junk Drawer

Most of us have a drawer in our house where extra pens, appliance manuals, and half-used tubes of Super Glue get thrown when we don't know where else to put them. If it's a drawer, that tends to work out OK. Every so often, you open the drawer, see a mess, wonder why you ever kept any of that in the first place, and dump the drawer into the trashcan. This works because we can see into the drawer. If, instead of a drawer, we used a barrel with a slit on top, we'd have a pretty big mess on our hands by the time we knew it was full.

This is the kind of workload that may be hiding in our management system. When we have an idea for a new form, process, checklist, redundancy, authorization, or key performance indicator, we identify the most likely "owner" of this new responsibility and "shall" it over to that person's desk. Rarely is much consideration put toward

whether that person has any capacity for the new responsibility because quantifying that person's existing responsibilities is too big of a job.

If you have a Gatsbys-long management system with "the Quartermaster will … ," "the Quartermaster shall … ," "the Quartermaster ensures … ," and "… is the responsibility of the Quartermaster" scattered throughout, then you probably also have additional responsibilities assigned to the Quartermaster through aliases. Your Quartermaster might be "signatory" in one procedure, "authorizer" in another, "controller" in the one after that, and "qualified person" in a dozen more.

Not even the Quartermaster knows everything the Quartermaster is supposed to do, and the Quartermaster doesn't dare complain because someone might take a close look and notice one of the things the Quartermaster *isn't* doing. So the Quartermaster does the best he can. He finds a way to deliver everything he is asked for, promising himself that he'll get to the rest should he ever find the time. And it all works out in the short term because everyone except the rule-makers gets what they ask for.

But then we rule-makers want to change something.

When we want things to be different, we change the script. We edit the management system with the expectation that this will edit behavior on the worksite. That presumes the worksite was following the management system in the first place, which we know they weren't.

Why weren't they following the management system? We didn't give them a path. We didn't tell them where to find the time, because we didn't bother to find the time ourselves. Not only did we not get

what we asked for, we weren't even sure *what* we asked for. We knew we wanted some process to be followed or form to be filled out, but we didn't have any idea how much work that might take. We didn't know if we were asking for five minutes or five hours. We were playing Gettin' Shit Done Guy.

Pay the Cost to be the Boss

All confusion, loss of control, and debt created by complexity will eventually be paid back. Hopefully in hours or dollars, but too often we pay back in blood.

So, what is the answer?

Is it clarity? I've included a few chapters on clarity. This will help make things as simple as they are supposed to be, but sometimes things aren't going to be simple. Some processes have a lot of steps. Some clients have complex rules. Sometimes we have to use "50¢ words". Sometimes we make mistakes. Sometimes our links are broken.

Is it expertise? It would be nice if we knew everything going on everywhere, so we never had to ask for creativity or research from our personnel. Those are big "asks" that can generate a lot of frustration and resentment. Unfortunately, sometimes "do or die" is all we'll have to offer.

Is it quality? Quality is important, but perfection isn't the answer. Change is the enemy of perfect design. You cannot be agile, adaptable, and responsive if perfection is required. Sometimes, intentional technical debt is reasonable.

The big answer, the way to give us everything we want, is less. "Less is more" means more time, more consideration, more cooperation. Less brings clarity, expertise, and quality into actual reach. But it isn't easy.

First, stop digging.

We're stuck in a pattern. We add to our own responsibilities in good faith until we can't possibly keep up, and then, like a drowning swimmer, grab onto *anything* that floats past, often dragging it down with us. Writing and maintaining a management system can become so overwhelming that we "hand off" responsibilities (and authority) to any random person in the office who seems bookish and looks like they need something to do. Express an opinion (pro or con) on Oxford commas in the break room, and you'll likely be given something to write. These well-intentioned helpers often don't have the organizational discretion to order lunch on the company credit card, but we let them add expensive hours and complexity to our systems. We lock up the toner cartridges because printing is expensive, but we let "service providers" add slides and musical interludes to our training programs like we aren't paying for the time our people will spend watching them.

This is one time we need bureaucracy! Find the longest, most complex procedure in your management system. Maybe it's confined-space entry or energy isolation. Any change to your management system should take at least five times as much effort as completing your longest, most complex procedure. Changes should be submitted to gatekeepers with a budget in frequency and hours. All stakeholders should be aware from the beginning what these changes will cost, and we should have an informed conversation about how those costs in money and time will be paid.

If you're a gatekeeper, you are responsible for representing everyone who isn't in the room. Most initiatives are championed by a few people who believe their idea is so brilliant and important that cost shouldn't even be in the conversation. Sometimes they're right. But too often we use additions to the management system as rewards

for good effort, demonstrations of our faith in favorite employees, peace offerings to adversaries in the office, *yubitsume* after an incident, or because we're intimidated by the person who requested them.

And every word we add loosens our control over what actually happens in the workplace. All that refinement, all that nuance, all those pages … are *nothing* if they are not followed. Every additional page is an abdication of authority.

This is where it starts to hurt.

I am going to say some things in this chapter that will, as the saying goes, hurt me a lot more than they're going to hurt you. If you don't agree with what I say here, you will *never* want to hire me to coach your personnel or update your management system. This is the chapter with the greatest potential to destroy any chance of a professional relationship between me, the author, and you, the reader.

I am going to criticize *everyone*, because we all (have to) do the things I'm going to criticize. In fact, I won't even insist that you stop doing them. The purposes of simplicity are to gain or improve control (to improve compliance) and to eliminate *hidden* costs. So let me offer some provisos before we get started:

- If you already have control, if your personnel and contractors already understand and follow your management system, then don't change what you're doing!

- If all of your management-system deliverables (forms) are, in fact, delivered, then you have no *hidden* costs. You are funding all of your mandates. Don't change what you're doing!
- If you have more control today than you did yesterday, then you are leading. There might be something here to help get the rest of the organization to follow *faster*, but don't change what you're doing!

If you want more control, or if you are losing control as time goes by, then these are some of the areas you can look to cut. This might hurt me more than it hurts you, but it's going to hurt us all just a little.

From Aspirational Chaos to Authentic Consistency

I learned a big lesson the first time I went through a digital transformation. We went from a paper Control of Work system to a digital Control of Work system. I spoke both Control of Work and Information Technology, so I played a large role in transforming our analog management system into something the digital system could understand and manifest.

A digital system brings a lot of opportunities and guarantees, but also a few limitations:

- Digital systems think linearly: this happens, then this, then this. Generally, the system will not progress to step three until step two is complete.

- Digital systems are literal: we find out when "going digital" how much we were relying on assumptions and imagination in our paper systems.

So not everything could be translated with 100% fidelity. Some of the hardest to digitize were "dog pile" authorization schemes where, in theory, everyone including the janitor had to sign a piece of paper before work could begin. In practice, the janitor didn't work Thursdays, and some of the other gatekeepers didn't get to the job site before 10 am. The authorization scheme described a "best case" scenario, but if you asked anyone, they'd promise it was done that way every time.

And it *kind of* was done that way every time, if you didn't look too closely. When the janitor was absent, we'd make a good-faith effort to find the assistant janitor. When one of the other gatekeepers was gone, we'd maybe call and ask … *if* the job were phone-call-worthy. We followed *the spirit of* the procedure.

There are three big showstoppers when you put something like that into a computer:

- The computer isn't going to let you proceed until the Janitor gets there.
- The computer program probably isn't built to take these signatures out of order (nor the checklists, nor the myriad other "nice to have's" included in the procedure over time).
- The computer doesn't have a concept of what jobs are or are not "worthy" of absolute conformance.

So the procedure had to change. The “dog pile” of checklists and signatures was sometimes reduced, sometimes relegated to photo attachments. These photo attachments turned out to be purgatory for doomed “nice to have’s,” as it was only a matter of time before many of these were removed as well.

This was a “tough sell” because a lot of the dog pile hadn’t come from us, it had come from clients, big clients who were used to getting their way.

You can’t say no when your client wants to “exceed” your safety standards! To the client, six signatures were twice as good as three. The client was Gettin’ Shit Done Guy.

But we gave up a lot of complexity, complexity we’d fought about for years, complexity some very important people had sworn we could never live without. We gave it up because we were excited about what we were getting in return. The promise of our new digital system crashed like a bulldozer through our rigid mindsets.

There is promise in simplicity. You can get control back. You can bring your personnel back into the conversation. The cycle of continual improvement is within reach. You just have to get excited enough to make a change.

What Gets Measured Gets Managed

In programming, we don’t take it for granted that something works. We test it. It slows things down a bit over the short term but speeds things up over the long term.

For instance, I might write a function

```
function add(a, b) -> a + b
```

At this point, I *think* that function works, but I don’t **know** that function works. So I write tests.

```
# test commutative
assert add(a, b) = add(b, a)

# test associative
ab_c = add(add(a, b), c)
a_bc = add(a, add(b, c))
assert ab_c = a_bc

# test identity
assert add(a, 0) = a

# test inverse
assert add(a, -a) = 0

# test comparison
assert add(a, b) = a + b

# test fixed
assert add(2, 3) = 5
```

If I really want to be careful, I'll write the tests *first*.

It's good practice for software, but it didn't start with software.

We've had tests for a long time.

In programming, a test is an assertion. In the "real world," a test is a deliverable. When we give an instruction, we require a deliverable to **know** that the instruction has been followed. Before digital systems, that would look something like this.

> Requirement: test the horn on the forklift.
>
> Deliverable: check a box.

We might organize similar checkboxes into check*lists*.

> Requirement: inspect the forklift.
>
> Deliverable: submit a completed forklift checklist.

Anyone inspecting a forklift will fill out a checklist. The checklist will catalog everything that was inspected and any issues that were found.

We might have different checklists for monthly inspections, daily inspections, and pre-use inspections. We can end up with a lot of checklists. So, what do we do with the checklists?

There's a branch here, one path for "issues found" checklists and another for "no issues found" checklists. We might direct "issues-found" checklists into a corrective action process and "no-issues-found" checklists directly into a document-retention process. A

pre-digital management system would likely give that responsibility to the person filling out the checklist. At this point, that person is probably the only one who has seen the checklist, so this makes sense.

Now that we've instructed the person filling out the checklist to initiate a corrective action process, we need a deliverable for us to **know** this is taking place.

Remember that we're walking through this process in the pre-digital world.

> Requirement: review checklists and initiate the corrective action process if …
>
> Deliverable: submitted corrective-action forms.

We might later have a review (audit) of this review, later still, a review of the review of the review, and so forth. This is manageable because each review is higher level.

- The first review, performed by the person filling out the checklist, is performed on every checklist.
- The second review might involve some percentage of checklists per week.
- The review after that, a smaller percentage of checklists per project.
- The last review might itself be a single checkbox: "(yes or no) Is the project generating corrective-action forms?"

Eventually, we end up with a stack of fileboxes of checklists that will probably never be opened or reviewed again. The filled fileboxes at this point *are* a deliverable.

> Requirement: use inspection checklists.
>
> Deliverable: boxes of completed checklists.

Almost certainly, there are checklists in these boxes that have never been seen by anyone except the person who filled them out. There is a chance we may want to open these boxes for additional review in the future—else why make checklists instead of bullet lists? So we make a requirement to retain the boxes of checklists for some period of time.

> Requirement: retain inspection checklists.
>
> Deliverable: some physical location full of fileboxes.
>
> Requirement: define an expiration date.
>
> Deliverable: log boxes of checklists into an inventory.
>
> …

I didn't go through all the steps, but there's enough here to see a pattern: the responsibility is pyramid shaped. One person working at a project, inspecting one forklift, might spend ten minutes inspecting that forklift and completing a checklist. One person working in a warehouse, managing one palette of fileboxes of checklists, might spend ten minutes inspecting that palette of fileboxes of checklists. One person *managing* a warehouse full of palettes of fileboxes of checklists might spend ten minutes inspecting the full inventory of fileboxes. That is, the more forklifts in an employee's purview, the less work that employee is assigned *per forklift*.

BUT NOW WE USE DIGITAL SYSTEMS.

Not quite. Yes, we use computers, but we still work to expectations and legal standards left over from the analog world. The *way* we use computers brings to mind early attempts at airplanes with flapping, bird-like wings. As individuals, we may be ahead of the curve, but we *will* have to wait for for the rest of the world to catch up.

One thing we have not compensated for as a species is the loss of the pyramid-shaped allocation of responsibility. In the analog world, if we doubled the number of checklists, we would double the workload for every person who handles those checklists at each level. We'd also double our storage and printing costs. That is no longer the case. If the checklist is digital, a computer handles all the cost and all the work *except* filling out the checklist.

That's a massive exception! No one except the person filling out the checklist has any feedback on how much work doubling or tripling the number of checklists creates. Conceptually, we understand that there's *some* cost *some* where, but we can't sense it. It doesn't feel real to the person assigning the work.

And we don't just assign more checklists. We see all the "free" upgrades the digital systems can offer us, and we start asking for pictures, scans, additional authorizations, and anything else we can imagine—now that the cost to *us* has been eliminated. We end up generating hundreds of hours of other people's work just to create a little content for our ProMapp dashboards.

If you don't want to reduce your forms or your number of forms, look for ways to re-distribute the work back into a pyramid.

Do We Need Elaborate Incentive Programs?

Several years ago, I took part in a birthday celebration for a popular friend. Around forty of us went through the Body Worlds museum exhibit and had dinner at a nice restaurant afterward. It was my friend's birthday, so we wouldn't have let her pay for anything anyway, but we were especially generous because she was going through some financial difficulties at the time.

Not everyone had a huge appetite after looking at flayed human bodies for an hour and a half, but my friend had an amazing time. After dinner, she went home pleasantly tired, well fed, feeling loved … and still unemployed.

It occurred to me later that we had collectively spent a few thousand dollars. Everyone had a great time, and we made great memories, but that money could have paid two months' rent for the birthday girl. I won't say that we definitely should have given her the money instead of going out, but that idea should have been in the conversation.

If anyone gets hurt, no one gets a ball cap. If no one gets hurt, everyone gets a ball cap. It can be that simple.

I've worked with people whose full-time job was counting points for incentive programs. I have a pretty good idea of their salaries and a pretty good idea what they required from field personnel to complete their reports. It would buy a lot of ball caps.

Key Performance Indicators

People who work "graveyard shifts" in vulnerable glass buildings like gas stations or fast-food restaurants are trained to keep the buildings clean to discourage criminals. The theory goes that criminals prefer to operate in environments where they can blend in or hide, and a clean and well-lit glass building will make the criminals feel exposed and vulnerable.

This feels true, and even if it isn't, cleaning the glass building keeps everyone awake and aware, which is also likely to deter criminals. A staff that keeps the glass building clean is more likely to do other conscientious things like change the light bulbs, lock the door before they stock the freezer, keep the cameras working, deposit money into the drop safe, and pay attention to customers.

It's probably all true, but we'll never really know. Correlating common occurrences like mopping the floor with rare occurrences like robbery at gunpoint is nearly impossible. You can prove this with math, but the best argument may be that we're still using most of the KPIs we used in the twentieth century. The AIs haven't found anything new for us.

So we have to accept on faith that clean floors are a good metric for robbery prevention in general. The same thing goes for "key performance indicators" and "leading indicators." We have to accept them on faith and good sense.

Don't let that discourage you. If you take one thing from the math chapters in this book, let that one thing be that we take a lot more on faith than we think we do. And that's fine!

HOWEVER …

If we're accepting on faith that *clean floors* are a good metric for conscientious, robbery-preventing behaviors in general, can we accept *changing the light bulbs* instead?

What if it took a *lot* more work to measure clean floors than to measure new light bulbs—a *lot* more? What if, to keep track of clean floors, we used checklists, inspections, photographs, audits, and customer surveys, whereas, to keep track of light bulbs, we tracked how often we ordered light bulbs?

We assume there is some value to these cleanliness "indicators," but we don't (literally can't) know exactly how much that value is. Doesn't it make sense to use the least expensive indicators we can find?

Don't overspend on your KPIs. If there's something you can get from

- a computer system;
- data you're collecting anyway for some other reason,

start there. You can build the same reports, gain the same insights, and save the company a lot of complexity and hours.

The Rule Above All Rules

As much as I've had to say about the fallacy of indicators and analytics in this book, there is one qualifier I have to offer. Here is as good a place as any. The Platinum Rule, The Rule Above All Rules, The Cardinal Principle:

NOTHING IS EVER A WASTE OF TIME IF THE CUSTOMER IS WILLING TO PAY FOR IT

OK, back to it.

Don't Get Lost in the Mays

If your personnel are following all of your instructions, they might be receptive to some of your suggestions. If your personnel are *not* following all of your instructions, then the suggestions might be part of the reason why.

> Procedure: Water The Grass.
>
> Step 1: Employee will attach the hose coupling to the hose bib.
>
> Employee may use a small pair of channel locks or an 8-inch crescent wrench to "snug up" the hose connection.
>
> Ensure you do not over-tighten the hose connection.
>
> Some hoses have a hexagonal coupling that allows the hose to be "snugged up" or loosened with a brand-name

> 8-inch crescent wrench. Some cheaper 8-inch crescent wrenches will *almost* fit, but not quite. Personnel may order a Bahco, Craftsman, Crescent-brand, or similar quality 8-inch crescent wrench and submit an expense-report form for reimbursement.
>
> Personnel purchasing an 8-inch crescent wrench may additionally purchase a hose with a hexagonal coupling if the hoses on site do not have hexagonal couplings.
>
> Site leadership may purchase a brand-name 8-inch crescent wrench and secure it with a cable to the hose bib in order to save time locating a crescent wrench. Do not attach a 10-inch, 12-inch, or larger crescent wrench to the hose bib, as a larger wrench might encourage over-tightening.

Some will disagree with using *any* tool on a hose coupling. Others will insist on it. We've compromised with "may" so each supervisor can set their own policy.

In the discussion that led up to this "may" compromise, the anti-tool contingent brought up a lot of objections about over-tightening. The pro-tool contingent argued that this would not be a problem if an appropriate-sized tool were used. The group as a whole couldn't quite decide what tool was appropriately sized because they'd all had different experiences with 8-inch crescent wrenches. The safety leadership pushed back on requiring a specific tool size at all because they weren't sure how to measure or audit that requirement. This gave us a few more mays.

The pro-tool contingent was savvy. By leading the conversation into minutia about tool size, they'd maneuvered the anti-tool contingent away from an absolute "no tool ever" debate into a "how can we

make tools work?" conversation. The anti-tool contingent didn't realize this until it was too late, but they did eventually realize it, and they weren't happy about being outmaneuvered. The anti-tool contingent insisted that no one was going to be walking around with a just-right-sized crescent wrench and that the pro-tool contingent didn't have the common sense to realize this. The pro-tool contingent sketched out this idea of attaching an appropriate-sized crescent wrench to the hose bib. It wasn't a great idea, certainly not policy worthy, but it didn't have to be. As I said, the pro-tool contingent had already won.

The safety leadership didn't want the anti-tool contingent to go away mad. They assured the anti-tool contingent that their objections about tool size and availability were *so* important that they would have to be specifically addressed in the company-wide hose-attachment procedure. And here we are with this mess.

Does this sound familiar?

The Pdf Cold War

Do you have a Microsoft SharePoint site full of stale Adobe pdf files? Do you have to track down "the Word file" when you want to make a change? Do you often have to re-create "the Word file" when someone has lost it or left the company? Do you have an online document converter bookmarked so you can convert slowly but inexorably more corrupt pdf versions of your documents from *.pdf to *.docx, only to open the result in Word, make changes between a mess of broken headers, and again save the file as a pdf?

Do you have errors or old links in your document that you don't correct because fixing them would take too much time? Do you have page numbers that count 1, 2, 3, 7, 10? Do you have documents that print with four-inch margins on the left and truncated text on the right? Do you have checkboxes that won't uncheck? Do you have paragraphs that can only be read by copying the entire document and pasting it into Notepad? Has someone resorted to pasting *images* of text onto your org chart to make changes?

Welcome to the cold war.

Microsoft *.docx and Adobe *.pdf files both show documents, but the similarity pretty much ends there.

Docx is an open, XML-based format. Most of the information is actually in plain text. Try this:

- Rename some_document.docx to some_document.zip.
- Right-click some_document.zip and "extract all."

You'll end up with a folder, "some_document," full of plain-text XML files. You can read these in Notepad. If you know what you're doing, you can even edit them, zip the whole thing back up, rename it back to some_document.docx, and open the result in Word.

Pdf is an open, *binary* format. There isn't much "plain text" inside it, zipped or unzipped. To read a pdf, you'll have to know the size and layout of each data field in the file, including the number of bytes used to represent each field, the order in which the fields are stored, and any padding used to align fields on byte boundaries.

Both have pros and cons, but translating between them is all con.

Pdf has always been easier to view and harder to edit, even when Microsoft and Adobe both used binary formats, so pdf has gotten a reputation as being more universal and secure. A binary file requires less parsing than an XML file, so the translation from document to screen to paper is simpler. This is why your browser and printer like pdf files.

For these reasons, a lot of us treat Adobe's "Portable Document Format" as a "Publish my Docx Format." We create the files in Word, review the files in Word, edit the files in Word, then, when we're ready to send out the "official official" version, convert them to pdf.

Most of us are just innocently following habits or assumptions, but not everyone is innocent.

There may be a cold war going on in your organization, and it's causing you a lot of extra work. Some (only a few, but it doesn't take many) people cannot accept the fact that paid work product is company property. When they create something, they make it intentionally harder to edit by saving it in pdf format. This is often referred to as "locking it down." Locking it down discourages others from updating their creation (and stealing some of their limelight). Sometimes, these files are even marked with an ersatz copyright notice, something along the lines of "Tarantino 2025."

They're out there. But not in your organization, I'm sure. Everyone in your organization is well-intentioned and forthright. It's still a good idea to ban pdf conversion, just to prevent a lot of *unintentional* double work.

And, just in case you didn't know, Word files can be "locked down" with a password. They can even be partially locked down to make editable forms. But the *best* practice is to secure them in SharePoint or some other document-control system.

Risk Assessment

Be Careful

When my son, Oliver, was two years old, I caught him standing on a tall stool.

After positioning myself where I could catch him if the worst happened, I tried to use my own communication training. “What would happen if you fell off that stool, Oliver?”

He dismissed the question—along with any concerns I may have had—with “It’s OK, I’m being careful.”

His confidence *and vocabulary* surprised me. I’d never heard him use that word.

Curious, I asked, “Son, what exactly does ‘careful’ mean?”

He didn’t know. He was sure that word somehow had the power to keep him from falling off the stool, but he had no idea what it meant.

I considered this for a moment, then realized I hadn’t thought much about it either, at least not enough to clearly explain it to a two-year-old.

I had to think for an uncomfortable moment or two before explaining to Oliver that “being careful” meant “making good decisions.”

That was the first time in my life I’d said or thought anything cogent about “being careful,” but it made sense to both of us, so when I then asked him, “Is standing on the stool a good decision?” he had to hang his head and concede that it was not.

I'd struck gold there, but I didn't realize it until the next time I caught *myself* saying, "I'll be careful." Of course, I had every intention of being careful (whatever that meant), but I'd put ZERO thought into what that might look like. I had to stop myself and ask, "What good decisions can I make?"

I've tried to make that a habit, but I still slip. Oliver helps me out by challenging me with, "Is that a good decision, Dad?" Now that he's old enough to *enjoy* catching me, these reminders come maybe more often than I'd like.

It was uncomfortable and instructive to wonder how many times I said, "be careful," without ever stopping to consider what it meant. It was more uncomfortable and more instructive to remember how much value I'd given to "be careful." Not just with my own safety, as in "I'll be careful," but in exchanges like "You be careful," then later, "What happened? I *told* you to be careful?!"

How many times had I stood up at a safety meeting and had nothing else to offer but some elegant rephrasing of "be careful"? And how had I ever thought that was enough?

Snowblind

There's an old trope that Eskimos have fifty words for "snow." In grade school, they told us this to stress how important snow is to an Eskimo.

I don't speak Yupik, Innuit, or any other "Eskimo" language, so I'll have to take my teacher's word for it. What I do know is that English has a lot more than fifty ways to say "be careful." Here are fifty. I'll bet you've heard every one of them in a safety meeting. And I'll bet again that you could add a few to the list:

> "Watch out", "Take care", "Stay alert", "Mind yourself", "Don't do anything stupid", "Watch your back", "Keep your eyes peeled", "Stay on your toes", "Be on the lookout", "Use common sense", "Remain vigilant", "Pay attention", "Be on your guard", "Stay sharp", "Keep an eye out", "Be circumspect in your actions", "Be prudent", "Be sensible", "Work conscientiously", "Be cautious of your surroundings", "Be mindful of potential hazards", "Stay safe", "Expect the unexpected", "Exercise caution", "Stay out of harm's way", "Watch out for the person next to you", "Don't take unnecessary risks", "Use the good sense your mother gave you", "Don't be reckless", "Pull your head out of your backside", "Take it easy", "Handle with care", "Proceed with caution", "Work safe", "Don't get in a hurry", "Take heed", "Keep your wits about you", "Don't let your guard down", "Maintain high alert", "Practice situational awareness", "Stay focused", "Don't do anything you wouldn't tell your kids to do", "Don't put your hands where you wouldn't put your [something important]", "Avoid pinch points", "Watch your step", "Stay on the ball", "Keep out of harm's way", "Look before you leap", "Keep your mind on the job", "Remember why you're out here."

We say, "Be careful," a lot and in a lot of ways. Read through risk assessments in your own organization, and I'm sure you'll see liberal use of these phrases and others like them. We even have acronyms

for a few ways to say "be careful" so we can write or type them faster. Have you seen ABBI for "Above Behind Below Inside"? How about "PLT" for "Proper Lifting Technique"? These are worse than meaningless because they take time and attention away from any valuable content on your risk assessment.

Equally worse than meaningless are instructions to "stay hydrated," "avoid overheating," or "remember [whatever]."

There is a big problem with this kind of language, and then there is a bigger problem. Both are easily fixed.

The big problem is that it's hard for *us* to know our own state of mind. Anger, fear, sadness, elation, and even focus (especially focus) can make it difficult to recognize when we're being distracted. We've got a lot of ways to say "be careful," but maybe even more to say, "I didn't mean to."

We operate under the assumption that we are in control of ourselves. That's why when someone gets hurt, it's common to hear, "Why would [the injured party] *intentionally* put their hand in that spot?" as if intention and conscious decisions were the way most actions occur. But the truth is we're terrible at remembering to remember.

That's the big problem. The *bigger* problem is that other people can't read our minds. If you walk onto a job site without a hard hat, a dozen people will whistle, yell, wave their arms, and tap their heads to let you know. It takes maybe one second for the people around you to realize something is wrong and another second to identify what it is.

That can be even faster if we've just discussed hard hats in a safety meeting.

But what if, instead of hard hats, our safety meeting had focused on mindfulness? Maybe we all walked out to the job site committed to being more "present" in the workplace.

What if your mindfulness starts to wear off around lunchtime? One minute, you are engaging your senses, focusing on the sequence of job steps, and maintaining awareness of the people around you; the next minute … tacos.

Someone around you might eventually notice, but it won't happen right away. It's not something you're going to notice in others unless you're remembering to remember to remember to remember to be mindful yourself.

The Fifty-Foot Rule

So, you want mindfulness?

Every hobby is the same. You see something that holds your interest *just* long enough to see something else. When you go back to look at that something else, you see a third thing. Repeat this enough times, and these things you see start to form a narrative. The things have relationships to each other. There is a story, and it is fascinating. As you investigate this story, you meet other people on the same path. You learn from and teach each other.

These stories are like living things. They have a way to grow and even to defend themselves. You reach a point of understanding and enthusiasm that isn't available to people outside the hobby. You are

driven to continue learning and teaching, so you seek others who are advanced enough in the hobby to learn from and teach to. You are willing to spend increasing amounts of time processing increasingly complex lessons.

And you begin to *exclude* people who distract from this hobby. I don't watch baseball, but the local team has been in the World Series twice in the past few years. When your local team is in the World Series, strangers will approach you to engage in conversation about the World Series. You can't step into an elevator without making a new acquaintance. And if you *aren't* invested in the World Series, *friends* will avoid you—they're all off talking to strangers. Don't worry, they'll come back after the game.

This whole thing can be "kicked off" by watching a game with a friend, walking past a store window, receiving an interesting heirloom from a relative, hearing something contentious on the Internet, or intentionally feigning interest in your immediate environment. Mindfulness starts with someone reminding you to listen to your breathing and in the best cases ends with your finding a fascinating story in your own job or situation. Build enough of a narrative around the world within arms reach, and you will start excluding contentious things you hear on the Internet. It can be transformative.

WHAT DOES THIS HAVE TO DO WITH RISK ASSESSMENT?

I selected mindfulness because it's a bit squishy. It's still kind of a nebulous concept. If I challenged you to be tough, be humble, be a leader, take pride in your work, or take ownership, you'd have a pretty good idea what I meant. But if challenged you to be *mindful*, you would probably have a few follow-up questions. If you are familiar with mindfulness already, you doubtless have quibbles with the simple explanation I gave above. Mindfulness can be a hard thing to pin down, and anyone who tries usually resorts to other squishy words like "presentism."

This squishiness makes mindfulness an excellent example of a principle. Most concepts have two parts:

- an aspirational component; and
- a methodology

If you want mindfulness in your job plan, stick to methodology. If you can make something as aspirational and squishy as mindfulness visible from fifty feet, then you can "methodologize" pretty much anything into a series of actions that can be seen from fifty feet. Let's get started.

Instead of *encouraging* personnel to be present and engaged, *instruct* personnel to stop every [some number] of hours or iterations and review the job plan. Be specific and define what "review the job plan" looks like. Maybe that's one person reading aloud and the rest of the crew facing that person. Whatever it is, make it something you can see from fifty feet. Use big, objective actions. You should

know when you're doing them, and so should everyone else around you.

Now, proceed to easier problems than mindfulness. If the aspiration is to use proper lifting technique, stop and demonstrate proper lifting technique before the job. Can you see that from 50 feet? Absolutely! If the aspiration is to say hydrated, check and refill the water cooler every hour or stop for a five-minute water break every half hour. I can't measure your hydration from fifty feet, but I can see you taking a break. If the aspiration is to watch out for crush points, identify and mark these crush points before starting the job.

More on Being Present

I often start off my risk assessment class with three paper cups, marked "cup A," "cup B", and "cup C." Cup A is the cup full of coffee that I would have walked in with anyway. Cups B and C are empty.

I take a few sips out of cup A while I introduce myself. Sometimes, we "go around the room" and all introduce ourselves while I continue to enjoy the coffee in cup A. After introductions, I immediately start with the first exercise.

> Exercise One: Fluid Transfer
>
> We will transfer the remaining coffee in this cup, cup A, into cup B and then into cup C. Start by identifying the risks.

This is a verbal exercise, so students either raise their hands or just shout out answers. Invariably, someone says, "burns!"

When I hear this, I take a sip of coffee and say, "Sorry, I didn't catch that. What was it you said?"

"Burns," the student will repeat.

I'll take another sip of coffee, sometimes with an ostentatious slurping sound. "One more time, please; I didn't quite understand you."

By this time, everyone gets my point. Obviously, this coffee is not hot enough to burn me. The student who said "burns" looks like the butt of the joke. That student was using imagination instead of investigation to determine what the risks were. That student had treated our fluid transfer risk assessment as an academic exercise, not an assessment of actual risk.

But I point out to the rest of the class that they'd pretty much all fallen for it. Rarely does anyone walk up to inspect the cups, and even when some savvy person does, the rest of the class doesn't follow their example.

These are often experienced personnel who have filled out hundreds of risk assessments, and they do the same thing in class that they do in the field: they try to please the teacher. They assume I've got a list of words in my head and that if they repeat enough of those words, the teacher will be satisfied and leave them alone. They've been conditioned to play this gotcha game instead of doing an actual risk assessment. The whole thing becomes an exercise in manual writing instead of job planning.

I've seen many hundred risk assessments that use the word "if."

- "If the load won't fit down the hatch …"
- "If the board is too heavy …"
- "If the ground is wet …"
- "If the pipes are too hot …"
- "If the weather is too cold …"
- "If visibility is low …"
- "If the bucket is too full …"
- "If the shelf is too high …"

Meanwhile, the load, board, ground, pipes, weather, visibility, bucket, shelf, etc., are RIGHT THERE. We don't have to deal with "what if" when we already know. If something requires two people or a different tool or an additional step, why don't we just go get it now instead of writing out this bowl of conditional spaghetti?

A lot of this can be addressed with training. I lead students through contrived examples. Then the students tell me it isn't possible to be both brief and clear in *their* jobs. Then the students offer real examples. Then I express the real examples briefly and clearly. Then everyone makes a little progress. A few hours of this goes a long way.

But what can't be fixed in class is that too much responsibility has rolled down the proverbial hill to the front-line workers doing the risk assessment.

Responsibility is first assigned by company executives or legal standards. Their instructions are understandably broad, as they have

to apply to an entire industry or organization. Let's start with an OSHA requirement.

> To facilitate cleaning, every floor, working place, and passageway shall be kept free from protruding nails, splinters, loose boards …[3]

Every organization is different, but a nice example of distributed responsibility might go something like this:

- Before the contract is signed, the Corporate Safety Department makes an arrangement with the client to provide dumpsters or large-scale disposal of demolition waste.
- The Project Team has a good idea of how much demolition waste will be produced. The Project Team arranges for temporary storage and on-site transportation of demolition waste.
- The Site Safety Personnel create a waste management plan that describes the arrangements made by the Project Team and assigns responsibilities where required.
- The Site Safety Personnel acquire signed authorization from the client to have a daily or weekly scheduled deep cleaning.
- The Work Crew creates a risk assessment, identifies demolition waste, and instructs personnel to proceed as detailed in the site waste-management plan.

That's the best case. The worst case is easier for *almost* everyone:

- The Corporate Safety Department pastes the requirement into the management system.

3 OSHA 1910.141 Sanitation

- Responsibility rolls down to the Project Team, who attach a copy of the entire management system to some site-specific document just to make sure the client signs it.
- Responsibility rolls down to the Site Safety Personnel, who copy this portion of the management system onto a smaller document, save it as "Waste Management Plan," and upload it to the company SharePoint site.
- Responsibility rolls down to the Work Crew, who are given, at best, fifteen minutes to decide and document exactly what they're going to do with a pile of old boards with nails sticking out of them—along with every other hazard they might identify in their risk assessment.
- The next Work Crew pretty much starts over …

Each level looks at the previous step and thinks, "This is what the boss put down, so I guess this is what the boss wants back."

We can't have the best-case scenario every time. We can't anticipate every decision that will have to be made. We can't front load every possible project or contract with a closet full of paperwork. We will always require people to react. Work crews have the most training and the best view of the work, and we will forever rely on them as the last, best chance to anticipate risk.

But we cannot rely on them as the *only* chance to anticipate risk. Every level of the organization should digest their share of the complexity and complete or distribute work product that would be redundant at a lower level.

Incident Investigation Reports

Asking Questions

Closed-ended questions are questions that can be answered with "yes" or "no." We can get away with closed-ended questions when everyone is relaxed and receptive. For instance, imagine approaching a stranger and saying, "Hello, is your name Jennifer?" If the stranger is relaxed and receptive, she might say, "No, I'm Sharon."

But if she *weren't* relaxed and receptive. The conversation could get ridiculous.

Hello, is your name Jennifer?

- No.

OK, how about Mary?

- No.

Deborah?

- No!

Patricia? Michelle?

- No and No!!

Carol? Diane? Barbara?

- No! No! No!

Sharon?

- Yes, my name is Sharon.

Hello, Sharon, do you like baseball?

- No.

Tennis?

- No!

…

It feels awkward just typing that out.

This is old news; it has been common advice for a long time to ask "open-ended questions." Open-ended questions are questions that *cannot* be answered with yes or no.

- What is your name?
- What sports do you like?

This leaves an unreceptive person with the choice of answering or telling you to go away. It takes away "no," which was the easiest choice when you asked a closed-ended question. Open-ended questions also help prevent unintentional leading.

- Do you like your new boss?

is a different question than

- How do you feel about your new boss?

So when we train for investigation interviews, we insist on open-ended questions.

- What did you see?
- Where were you standing?
- Why? Why? Why?

That's better.

In class, I often have the students take turns asking each other open-ended questions. It's sometimes tricky at first, but everyone quickly figures out you just need to start with who, what, when, where, why, or how. It's great for nervous interviews and moody kids.

There's more to learn, of course.

- What on Earth were you thinking?!
- Why shouldn't I fire you?! and
- How could you let this happen?!

…are still technically open-ended questions. The "five whys" are a big improvement, but they aren't a panacea. Some coaching is still required.

Answering Questions

The harder part is preparing to *answer* open-ended questions. After an investigation, prepare to be asked most of these:

- How do we know this won't happen again?
- What have you done to fix it?
- Did you look at [potential causal factor]?
- Why didn't you [potential corrective action]?
- Why didn't you identify this sooner?

In addition to our favorites from the last section:

- What on Earth were you thinking?!
- Why shouldn't I fire your company?!
- How could you let this happen?!

You're embarking down a particularly painful path if you rely on firing personnel as a corrective action:

- How do we know the next person will be any better?
- How did you hire that person in the first place?
- How did this person make it through your training?
- Why didn't you fire the supervisor too?
- What is going on in your HR department? Did you fix it?

Your first opportunity to answer these questions is your incident investigation report.

Find Solutions, Not Just Problems

When I teach incident investigation, I ask the class to tell me the purpose of an incident investigation *report*.

A few answers almost invariably come up:

- to discover the root cause of an incident
- to prevent future incidents
- to identify corrective actions

I ask again, "These are the purposes of an incident investigation; what is the purpose of an incident investigation *report*?"

I've gotten some great responses, but I always insist on one narrow definition. The purpose of an incident investigation report, I tell the class, is "to convince anyone who reads it that your corrective actions are necessary and sufficient."

This is an example of a logical appeal. This is for the student who doesn't walk into class convinced of the value of spending two days examining why someone forgot to put a gas cap back on—especially when that someone has already been *relieved of their responsibilities*.

It saves us a lot of class time that would have been lost to fielding cynical questions like "What's the point of all this when we already know what happened?" or "What's the point when *they* are just going to edit our report anyway?"

"Yes!" I say to the cynics, "maybe it is all a show. Now here's how and why to put on a Hell of a show!"

Of course, I don't believe it's all a show. In fact, I believe incident investigation may be our *best* opportunity to identify and address opportunities for improvement. But you can argue that I am being cynical too. I used my own training to consciously design this dialog as a way to engage ship captains, site leadership, executives, and other VIPs who would prefer to "skip the BS" and proceed directly to the nearest available solution. Create a great report, I assure them, and the client won't say, "… and then what" after you've fixed the problem to your own satisfaction.

From the first time I tried it, the dialog worked for that purpose. In the pursuit of "convincing anyone who reads [our report] that our corrective actions are necessary and sufficient," even VIPs are willing to dig in and explore causal factors they might otherwise have dismissed without consideration.

But the frame "convincing anyone who reads it …" is back-of-the-book worthy because it works on the bloodhounds too. Bloodhounds are the type, usually safety professionals, who love to go hunting for causal factors. They look under every rock, check every dial, and find every potential incident cause … except one: change.

Bloodhounds sometimes fail to recognize change and disruption as potential incident causal factors, especially the change and disruption they might cause by treating an incident investigation as a comprehensive site audit and corrective actions as a crusade against imperfection. Throughout the investigation, we need to ask *ourselves*

- Why are my changes necessary?

and

- Why are my changes sufficient?

That second question is often best answered by a detailed explanation of what went *right*. Anyone reading your report should come away with a reasonable understanding of how your company succeeds *almost* always at preventing incidents. By writing this report, you should acquire that same understanding.

We want to find and seize opportunities to improve, but we don't want to disrupt. One step forward, zero steps back.

Statistics

In previous chapters, I've gone over some of the math that unfortunately doesn't work like we think it does. Here's a taste of what does work … if you have enough data.

Red Fish, Blue Fish: An Introduction to Statistical Tests

One year ago, you stocked a pond with exactly 50% red fish and 50% blue fish. Today, you net some fish out of the pond, count the number of red and blue fish you net, and attempt to decide if one color of fish has survived better than the other.

What does it tell you if every fish you net is red?

THE ANSWER ISN'T OBVIOUS

What if you only net one fish?

You stocked the pond with an equal number of red and blue fish, so you have no reason to expect your catch to be anything other than half red and half blue. One fish can't be half red and half blue—neither can three fish or five fish for that matter. At some point 3/7 or 4/9 or 5/11 is going to be close enough to half, but what is that point?

How close can you expect to be to 50:50?

If netting one fish doesn't prove anything, does netting two, even if they're one of each color?

Does netting 6 red and 4 blue prove anything? Does it prove half as much as netting 7 red and 3 blue? How well does your net reflect the population of the pond?

A STATISTICAL MODEL

A statistical model is a map of the way you *expect* things to be. If you're reading this blog, you are probably attached to industry in some way, which means a statistical model is a great map, because we aspire to eliminate all factors *except* random variation, and the models describe randomness perfectly.

If performance is affected by anything besides chance (if we're deviating from the model), it's something we want to know about and deal with.

Note: The models have parameters (randomness does not require a fair coin).

THE BELL CURVE IS THE MOST FAMOUS MODEL …

But there are many. Here are a few so we can see what they have in common.

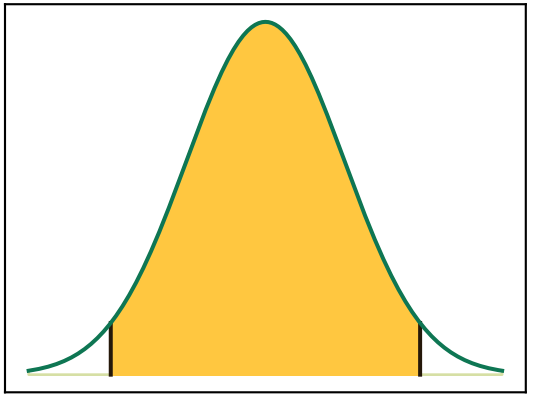

gaussian distribution

The Gaussian distribution, also known as the normal distribution, is suitable for modeling things that accumulate or contribute to a result:

- Physical measurements: Many physical phenomena, such as the height and weight of people, the length of objects, and the time taken to complete a task, can be modeled using a Gaussian distribution.
- Natural processes: Many natural processes, such as the distribution of rainfall, the velocity of wind, and the distribution of radioactivity, can be approximated using a Gaussian distribution.
- Human behavior: Many aspects of human behavior, such as intelligence, reaction time, and test scores, are distributed according to a Gaussian distribution.
- Errors and noise: Many types of errors and noise in data, such as measurement errors and sensor noise, can be modeled using a Gaussian distribution.

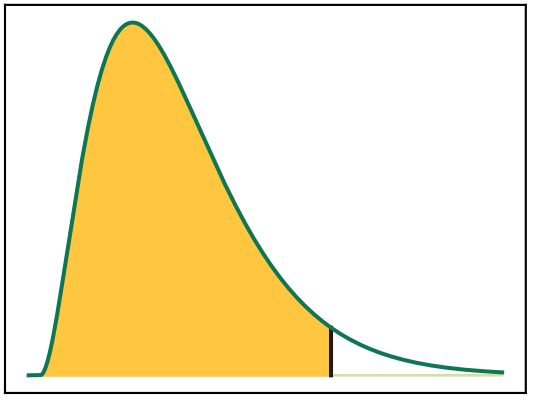

chi-squared distribution

The chi-squared distribution is suitable for comparing ratios:

- Goodness of fit tests: In goodness of fit tests, the chi-squared distribution is used to test whether observed data fits a theoretical distribution.
- Confidence intervals: The chi-squared distribution is used to construct confidence intervals for the variance of a normal distribution.
- Genetics: In genetics, the chi-squared distribution is used to test whether observed genetic ratios conform to expected ratios.
- Quality control: In quality control, the chi-squared distribution is used to test whether observed defects in a production process are consistent with expected defect rates.

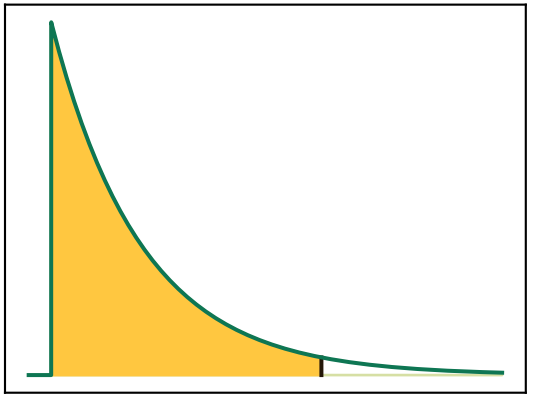

exponential distribution

A classic example of the exponential distribution is a box of light bulbs. If each averages 500 hours of life, what is the expectation for the lifespan of three bulbs used consecutively? In general, the exponential distribution is a good fit for any “memoryless” process:

Reliability analysis: The exponential distribution is commonly used to model the time to failure of a system or component, as it assumes that the failure rate of the system is constant over time.

Queueing theory: In queueing theory, the exponential distribution is used to model the inter-arrival times between customers waiting in a line or the service times of a computer server.

Radioactive decay: The exponential distribution is often used to model the time between decays of a radioactive substance, as it assumes that the probability of decay is constant over time.

Insurance risk: The exponential distribution can be used to model the time between insurance claims, which can help insurers to estimate the risk associated with a particular policy.

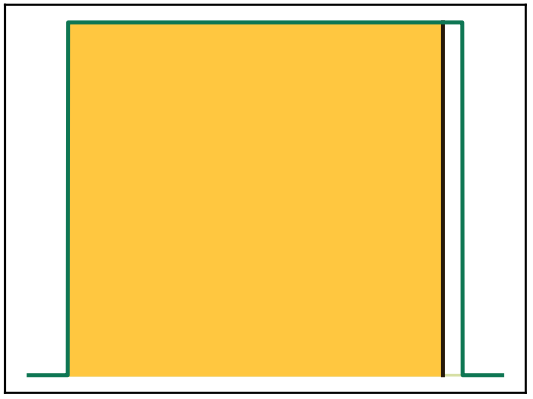

uniform distribution

The uniform distribution models situations where every outcome has an equal chance to occur:

Random number generation: In computer science, the uniform distribution is often used to generate random numbers between a specified range.

Probability modeling: The uniform distribution can be used as a simple and general model for probability distributions, especially when there is no prior knowledge or assumption about the shape of the distribution.

Physics: The uniform distribution can be used to model the position or velocity of a particle in a given range or the amount of energy absorbed by a material over a given range of wavelengths.

Games of chance: In games of chance, such as rolling dice or spinning roulette wheels, the uniform distribution can be used to model the probability of each possible outcome.

THESE MODELS KNOW HOW THINGS WOULD LOOK IF WE NETTED FISH (OR TIMED TRAINS, OR PULLED MARBLES OUT OF A BAG) FOREVER.

This is no small thing. You might or might not know that it is *twice* as easy to net 3 red fish out of 5 than it is to net 4 red fish out of 5. That's because there are 10 ways to net 3 and only 5 ways to net 4. That's why there's a curve in the bell curve: each draw has the same chance, but totals in the middle are more likely. For five fish, this is easy enough to calculate, but even a computer would be quickly overwhelmed trying to count its way to the difference in likelihood at commercial scales.

Here is the curve mapping how many red fish you'd expect to net if you netted 16 fish—IF your pond were still stocked 50:50.

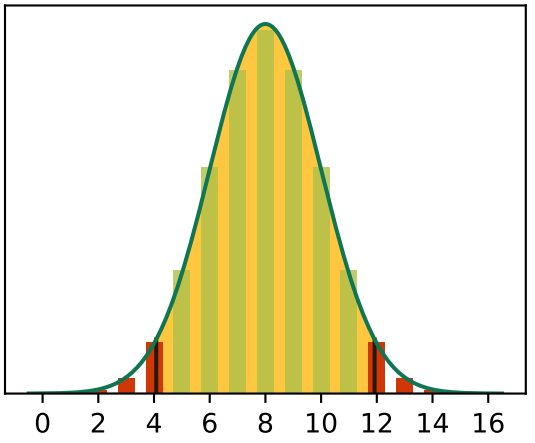

chance of netting some number of red fish out of 16

This model, like the ones above, has a yellow area covering 95% of conceivable outcomes. 95% is a typical statistical test boundary.

Did you notice the red bars outside the yellow area? Those red bars are what we shouldn't see.

We've modeled random chance, and random chance shouldn't produce those outcomes (not more than once in twenty anyway). If we get a red outcome, one of two things has probably happened:

- Our model parameters have always been wrong (have we forgotten that we stocked the pond at 70:30, not 50:50?).
- Something besides chance is affecting our result.

So our model is probably wrong. In statistical terms, we "reject" the assumption that there are currently 50% red and blue fish in our

pond. 95% may seem like a high burden of proof, but remember that in practice, you'd be performing such tests on something every day. With a lower boundary, you'd be rejecting all the time.

THE STATISTICAL TEST

Your model is a set of assumptions. Your statistical test looks at the likelihood of your outcome under those assumptions. If the likelihood is low, you reject your assumptions. This beats the heck out of the way humans tend to reject assumptions (usually recency bias).

The test is a formula, which—don't miss this—looks at not only the ratio of fish in your net, but the number as well.

Looking back to the last image (netting 16 fish), we see that you might change your assumptions if 2/3 of your fish were red or blue. If you only net 8 fish, however …

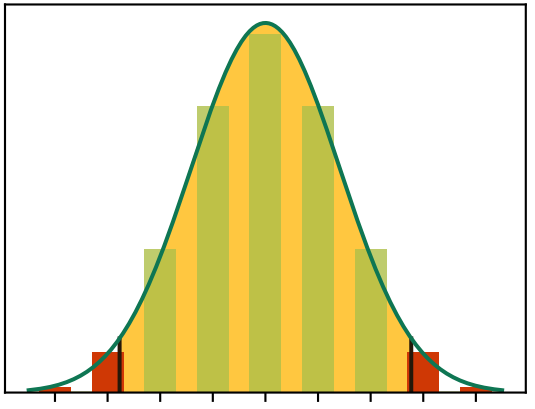

chance of netting some number of red fish out of 8

… you'll need more than 3/4 of one color to change your assumptions.

And if you only net 3 …

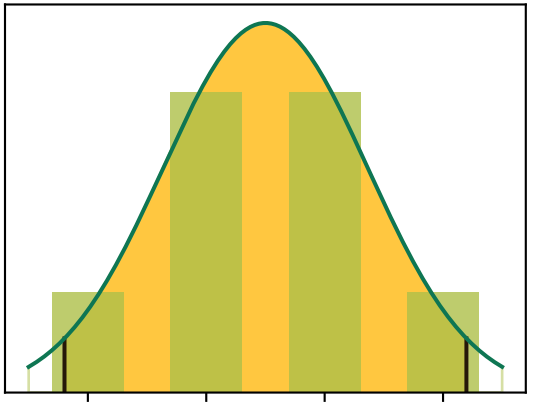

chance of netting some number of red fish out of 3

… you couldn't justify changing your assumptions no matter what was in your net. This is where it gets tough in safety, because in safety, we count BAD things, and fortunately, we don't have many BAD things to count. The general "rule of thumb" is that you need about thirty samples for this kind of test to be valid.

It's tough to get the numbers we need for proper analysis. In safety, we more often have to live by "It's not IF, it's WHAT IF?" and make our decisions that way.

Bar Chart Analytics

Even if you aren't familiar with statistics, you've probably spent a fair amount of time looking at and discussing distributions. Casual distributions are usually shown as pie or bar charts. When casually reviewing these pie or bar charts, we have casual metrics of significance. Twenty years ago, the typical metric might have been, "Is one of these bars conspicuously higher or lower than the others?"

Now, most professionals have some contact with IT personnel or analysts, so concepts like "expected value" and "standard deviation" have raised our collective mathematical sophistication. Now, we find ways to quantify "conspicuously higher."

A quick review:

> **Expected value** is the value you would expect to see *if* your model were correct.
>
> A **Model** is a set of assumptions, quantified by formulas, about how values *should* be distributed.

A model can be simple. Most companies use the model "each facility should have about the same number of incidents." That number is the expected value. For instance, if five facilities collectively have ten incidents, the expected value of *incidents per facility* would be two—that's the number of incidents divided by the number of facilities. You might use a more sophisticated model that takes facility size and other variables into account. The point is that you *do* have a model and you *do* have an expected value for indicators at each facility. That's an improvement from where you would have been twenty years ago—at least it's an improvement in vocabulary—but it's only a very short step in the right direction.

Let's specify our model and then try to use it to interpret a bar chart.

- We have five facilities of equal size
- We expect each facility to have 1/5 of the incidents. This is our expected value of incidents per facility.
- We will consider it significant if any facility has at least **twice** the expected value of incidents.

LET'S SEE HOW THIS WORKS OUT.

This obviously won't work with a very low number of incidents. If we have 2 incidents over 5 facilities, the expected number of incidents per facility is 2 / 5 = .4. Twice the expected value (our threshold for significance) is .8. One incident is more than .8, so every incident would be significant, and any single incident would be consistent with random chance.

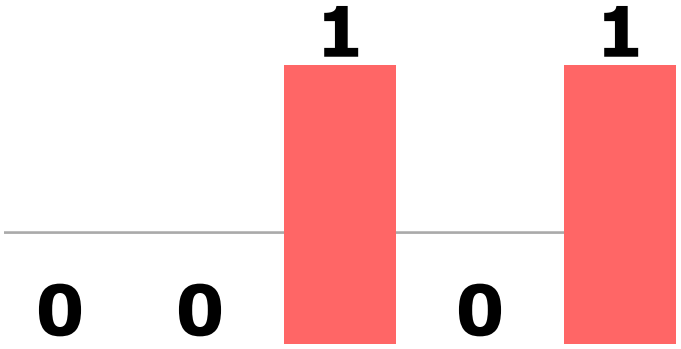

2 incidents across five installations

We are mathematically sophisticated enough as a population to see this graph and recognize that our expected-value calculation doesn't have enough data to work. Still, if I told you nothing besides, "Two of our facilities have 100% of our incidents" or "Our two worst facilities have two-and-a-half times the average incident rate," you'd likely raise an eyebrow. Even this graph is enough to emotionally support a "heroes and villains" narrative.

So how many incidents do we need for our casual metric to work? Let's have a look at five. The gray horizontal line is the expected value of incidents per facility.

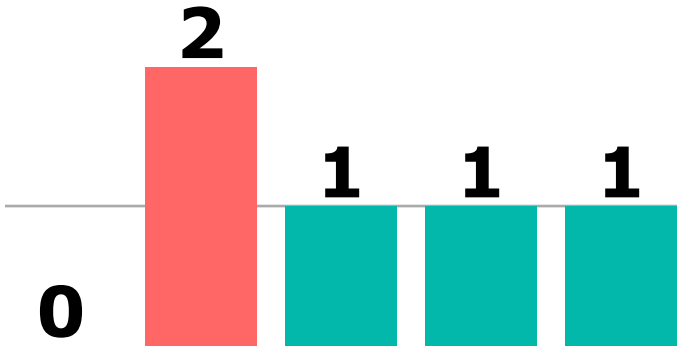

5 incidents across five installations

The expected value for five incidents is one per facility. 5 incidents / 5 facilities = 1. Twice the expected value (our threshold for significance) is two incidents at one facility. This is the least-worst (lowest standard deviation) possible distribution with at least two incidents at one facility.

I will follow that convention as we go: Every bar-chart distribution shown will be the least-worst with twice the expected value. With

five incidents, it is possible that one facility could have three, four, or five incidents. But I'm not showing those distributions. I'm showing the least-worst.

With five incidents across five facilities, 96% of random distributions would be as bad **or worse** as the bar chart shown. The chart tells a story, with heroes and villains, but the overwhelming probability is that this distribution is the result of pure random chance.

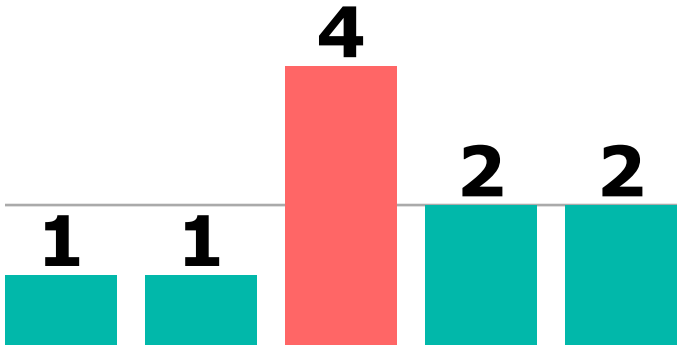

10 incidents across five installations

With ten incidents *randomly distributed* over five facilities, there is still a better than even chance (57%) of one facility's having at least twice the expected value of incidents.

When I say "randomly distributed" or "the result of pure random chance," I do not imply that incidents are *not* the result of causal factors. I imply that incidents may be the result of *uniformly distributed (or company-wide)* causal factors. If the distribution is likely the result of uniformly distributed or company-wide causal factors, then the corrective actions should be uniformly-distributed or company wide.

We will, of course, perform an incident investigation to identify causal factors. However, due to the nature of random distribution, it may not be productive to assume the causal factors are specific to the incident site. We do not have a "culture of blame," but we do have "a culture of monetary incentives." Maybe that's worth a conversation.

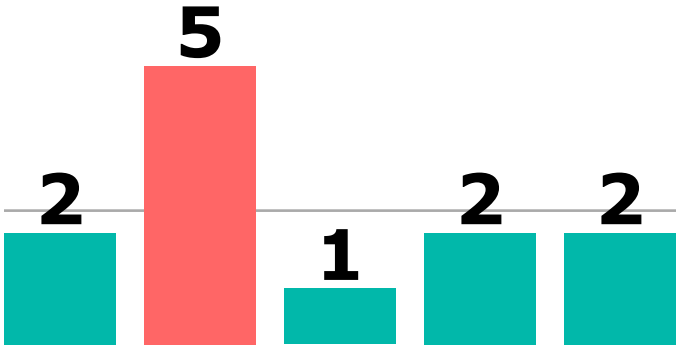

12 incidents across five installations

With twelve randomly-distributed incidents over five facilities, the chance of one facility with twice the expected value of incidents per facility is down to 36%. This graph shows the most even distribution possible with

- twelve incidents;
- five facilities;
- at least one facility with at least twice the expected number of incidents.

Still, this isn't a graph we want to bring into a staff meeting. Experience tells us how this information will be received, even

though random chance will give us this distribution *or worse* 36% of the time. Worse than that, If we collect the same information next quarter, we are 59% likely to see this distribution *or worse* in one of the two quarters.

This is like rolling a die.

- Your chance of seeing a six after one roll is about 17%;
- your chance of seeing at least one six after two rolls is 31%;
- after three rolls, 42%;
- after four rolls, 52%;
- after five rolls, 60%;
- etc.

So, a 36% chance is still extraordinarily likely if we repeat the same experiment multiple times.

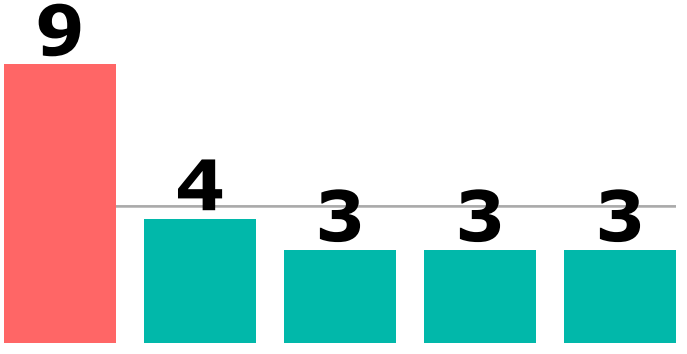

22 incidents across five installations

At twenty-two randomly-distributed incidents over five facilities, the chance of one facility with twice the expected value of incidents per facility is still 10%.

Ten percent is pretty low, but this is *still* beneath the typical threshold of statistical significance. This 10% would not be enough to reject the hypothesis that causal factors are uniformly distributed (or company wide).

What Are the Chances?

When you write about math, your friends and family very often want to ask about "The Chances." The Chances are the dark side of probability. "What are The Chances?" is only a question in the rhetorical sense. "What are The Chances?" is, in fact, an accusation. "What are The Chances?" means "I *know* someone cheated" or "I *know* something is wrong."

- What are The Chances the lottery winner would work in the state capitol building?
- What are The Chances my coworker's house was the only one to burn down … *in flood*?
- What are The Chances the neighbor would break his leg in a five-mile-per-hour collision?
- What are The Chances Bridget would win all three games of Bingo last Saturday night?

THIS CHAPTER IS ABOUT HUMILITY, SO LET'S TAKE A LOOK AT THAT LAST ONE

The pedantic answer is easy. With 100 people playing Bingo,

- 1/100 chance Bridget wins the first game.
- 1/100 chance Bridget wins the second game.
- 1/100 chance Bridget wins the third game.

$$\frac{1}{100} * \frac{1}{100} * \frac{1}{100} = \frac{1}{1,000,000}$$

one-in-a-million chance

The Chances Bridget won all three games of Bingo last Saturday are **one in a million**. Bridget is obviously a witch!

DOES IT HAVE TO BE BRIDGET?

One in a million is the answer to the question *as asked.* But the intended question might have been "What are The Chances any one person would win all three games of Bingo last Saturday." Those odds are much better than one in a million.

- 100% chance *someone* wins the first game.
- 1/100 chance the same someone wins the second game.
- 1/100 chance the same someone wins the third game.

$$\frac{1}{1} * \frac{1}{100} * \frac{1}{100} = \frac{1}{10,000}$$

one-in-ten-thousand chance

Things are already looking a little less nefarious. But 1:10,000 is still a pretty small chance. You'll probably be talking about this for years, which invites the question …

DOES IT HAVE TO BE LAST SATURDAY?

Will we still be talking about Bridget's "impossible" streak in 100 Saturdays? If so, then *which* of the 100 Saturdays isn't important.

The math for this one is a bit trickier.

Where:

- **p** is the chance of success *per trial*
- **n** is the number of tries
- **k** is the number of successes

$$\binom{n}{k} p^k (1-p)^{n-k}$$

binomial pmf

That is the *probability mass function of a binomial distribution*. I'm only writing that out so you can find it later on Wikipedia. A more appropriate description is "What are The Chances of winning **k** times in **n** tries?"

It's actually a little more formula than we need here, but keep it handy, and you'll usually be able to figure out The Chances even when there are no math writers around.

Let's set our variables.

- **p = 1/10000** is the chance of any one person's winning all three games on a particular Saturday.
- **n = 100** is the number of Saturdays.
- **k = 0** is the number of successes. Why **0**? That leads to the ungainly, but pretty much necessary explanation found in every probability article.

> Ungainly but necessary explanation:
>
> We're not talking about it, but there is a small chance we could have *more than one* "impossible" Saturday. In that case, we'd have to figure out The Chances for exactly 1 impossible Saturday + The Chances for exactly 2 impossible Saturdays + The Chances for exactly 3 impossible Saturdays, … all the way up to 100. Instead, we'll start with a 100% chance and subtract the chance of 0 "impossible" Saturdays.

$$1 - \binom{n}{k} p^k (1-p)^{n-k} \approx \frac{1}{100}$$

one-in-a-hunded chance

We're all the way down to a reasonable (1%) chance Bridget's "impossible" Saturday happened without forged cards, sleight of hand, bribing the Bingo caller, or other such witchcraft. Were we a little quick to judge Bridget?

BEFORE YOU ANSWER …

We've already decided that it didn't *have* to be Bridget and it didn't *have* to be last Saturday. Did it *have* to be Bingo?

Every time we roll dice, flip coins, deal cards, buy stocks, draw Scrabble letters, kick a ball, name a pet, skip a rock, change the channel, open the curtains, answer an anonymous phone call, guess a number, hold a raffle, drop a glass bottle, throw trash into a far-away basket, or close our eyes and spin around, there is a small-but-probably-larger-than-we-suspect chance to see something "impossible."

Here's one more time through our formula

- **p = 1/100** chance of "impossible" event *per game* (a game's being any witnessed event)
- **n = 100** games witnessed 100 times each
- **k = 0** "impossible" results

$$1 - \binom{n}{k} p^k (1-p)^{n-k} \approx 63\%$$

a great chance

That's a 63% chance of seeing something "impossible" if you watch 100 games 100 times.

GO APOLOGIZE TO BRIDGET.

On television, the smart detective never believes in coincidences. Things don't "just happen."

When you look hard enough, things "just happen" every day. Before reading the math chapters, you might have followed the "alligator metric"[4]: if indicators at site A are greater than indicators at site B, then site A is a problem. Site A gets the scrutiny; site A gets the intervention; site A gets a red bar on the monthly report bar chart; the leadership of site A gets to hang their heads while everyone else softly clucks their tongues in disapproval.

The boundary for even casual statistics is much higher. In order to reject the hypothesis that indicators are evenly distributed, you have to see something highly improbable (only a 5% chance) if that hypothesis were true. 5% is about two standard deviations away from the expected result under your hypothesis.

This raises the (red) bar considerably, but you're still going to see it. If something is 5% likely when you check once, it's 9.75% likely when you check twice. That same something is *more likely than not* (51%) when you check for it 14 times. Do you have 14 indicators? If you track a lot of indicators, you have to start believing in coincidences.

4 The less than < and greater than > operators are the "alligators" that eat the larger number.

Why Hasn't Machine Learning Changed the (Entire) World?

In a casual game of craps, there are usually two bets: win or lose. Payout on either is 2:1. Betting takes place before the first roll, at which time the player holding the dice has a consistent, slight disadvantage versus the players betting against him. In a casino, there are many possible bets on a craps table. Payouts vary from 2:1 to 19:2, and house advantage varies from "sucker bet" to "keep 'em there long enough to buy dinner."

But every bet, even the sucker bets, has an identical, simple design: the minimum house advantage payable in whole chips. Variation in odds is not due to some elaborate strategy on the part of the casino, but to the limited number of chip denominations and characteristics of the infinite-but-partitionable potential sequences of die rolls.

Why the minimum house advantage? That's a part of the other game the casino is playing: don't just make the player lose; keep the player losing.

Everything on the table is controlled by just two dice. To win the second game, the casino has to manage light, temperature, sound, color, the portion size of complimentary drinks, and 100 more.

The game on the table hasn't changed with the digital age. The game in the casino has, but the critical difference isn't the number of factors or even the complexity.

I'll come back to that critical difference. But first … we'll have to loosely define a few math terms.

Continuous vs. Discrete

CONTINUOUS

comprised of uncountable parts, infinitely divisible.

If I pour water into a glass, the volume is continuous. In other words, there are infinite possible volumes between 0 and full:

- 1/3 of a cup
- 1/3 of 1/3 of 1/3 of 1/3 of a cup
- 99.9999999% of a cup
- and so on …… forever.

DISCRETE

comprised of countable parts.

If I put marbles into a glass, the volume is discrete. In other words, there are finite possible volumes between 0 and full:

- 0 marbles
- 1 marble
- 2 marbles
- up to maybe 30 marbles depending on the size of my glass.

If we're being pedantic, the volume of water isn't strictly continuous (it's made of countable molecules). Only things like acceleration,

bias, and risk are strictly continuous. But water of sufficient volume is *effectively* continuous, so we treat it that way—even in equations.

And that's the way I'm using continuous and discrete here: effectively continuous and effectively discrete. On the King Ranch (the world's largest ranch), the cattle population is effectively continuous. In my backyard, the cattle population is discrete.

I've paused to highlight the "effective" distinction because I'll be using the particularly non-mathematical term "very discrete." The "more discrete" a quantity is, the less effectively it can approximate a continuous value. For instance, you can stand on a stack of books to change any light bulb, but you can't reach most light bulbs from a stack of libraries.

CONTINUOUS AND DISCRETE MEASUREMENTS CAN "PREDICT" EACH OTHER.

Given the volume of barking in my driveway (continuous), we could estimate the number of dogs (discrete) living in my neighborhood.

Given the number of barks heard in 24 hours (discrete), we could estimate the amount of time (continuous) those dogs spend outdoors.

THE CONTINUOUS ADAPTS WELL TO THE DISCRETE.

A catering service can estimate how much food (continuous) to cook for five people (discrete).

BUT THE VERY DISCRETE DOES NOT ADAPT WELL TO THE CONTINUOUS.

How many cows are needed to yield at least one pound of meat?

In short, we can tune to a station a whole lot better than we can station to a tuning. And that's where we get in trouble.

A Friend's Wedding

You make a plan to drive 100 miles to a friend's wedding. You check the price of gasoline and budget for the current price per gallon. As the wedding approaches, the price of gasoline rises 10%.

If the trip were continuous, you'd simply take less trip. You would drive only 91 miles and still make your budget. You'd miss the wedding anyway, but you'd be comforted by the fact that you'd made a data-driven decision.

Your real choices are discrete:

- stay home
- increase your budget

"Increase your budget" is a tempting solution because you've found a continuous piece to this puzzle. You could take just the right amount of continuous (and fungible!) money from somewhere else and leave the minimum possible continuous hole.

But what if money were discrete? What if you were paying in gold bricks? What if a 10% increase in gasoline weren't a few more dollars, but *another* gold brick?

That's the situation we often face when trying to alter operations and move personnel (a discrete resource) in response to single-digit analytical revelations.

Back to the Craps Table

"Single-digit analytical revelations" are what machine learning provides. Some combination of factors correlates with a small effect somewhere. If the effect were large, we'd have known about it without machine learning.

Once we filter out the factors outside our control, the effect is even smaller. (It is important to note that where these effects are on a boundary, e.g., "has cancer?", the significance is tremendous).

The casino is definitely in control of the environment inside the casino. The casino can turn continuous knobs, tweak continuous jiggers, and continuously shift things around to maximize even a small effect.

The casino is also in control of the game (craps), but the game doesn't have any continuous knobs. Machine learning might reveal (might have already revealed) that players are more likely to leave and never return after three losses in two hours on "hard 8." But what can the casino do with that information?

They can't provide less loss. They can't shift the odds toward the player's favor without shifting the odds *to* the player's favor. The casino has to pay the player in (discrete) chips. To return to a previous metaphor, they're either too high or too low to reach the light bulb.

THE TAKEAWAY

Machine learning and advanced probabilistic models require a lot of data. Applying ML requires a lot of (granular) control. The second is not a problem for investment markets for reasons that should be obvious by now, but it can definitely be a problem for those of us who more often have to decide "how many?" than "how much?"

For the discrete choices we face, there will never be an "aim small." Small revelations should not compel us to swift action. In these domains, what we already know may be more important than what we will ever learn.

How to Write a Procedure

META

If you're familiar with procedures, you've probably seen headings like "purpose," "scope," and "process." Most likely, you've never seen a "meta" heading, but you've certainly seen meta content.

Meta Content

Meta content is the tone, background, and justification for your procedure. Arguably, this content has a place, but that place is not in the body of your procedure. Put meta content in its own block or in a separate document.

Meta Content Example

Writing a procedure is difficult. Your job as a technical writer is to codify hundreds of pages of legislation, guidance documents, best practices, incident findings, technical knowledge, and company expectations into a set of clear, conformant instructions.

Your customer will have guidelines on the form and style of your document. These are not covered here. This chapter codifies hundreds of pages of loose guidance into a set of clear instructions.

USE BIG TEXT FOR BROAD STEPS.

Instead of using paragraphs to create a narrative:

There are two primary goals when making a sandwich: avoid food allergies and avoid food contamination.

Clean the food preparation area with anti-microbial cleaner. Wipe down the fans and lights above the food preparation area.

To avoid food allergies when selecting food ingredients, avoid peanuts, bananas, and soy; also, do not place bread, pasta, or ketchup food in the gluten-free preparation area. To prevent food contamination when selecting ingredients, only use food ingredients with a current inspection tag.

Use Title, Purpose, Overview, and subheadings to provide context, guide readers, and highlight important content. Only use paragraphs for the small steps:

Heading: Make a Sandwich

Subheading: Purpose

Define safe practices for avoiding food allergies and contamination during sandwich preparation.

Subheading: Overview

- Part 1: Prepare the work area.

- Part 2: Select food items.

Subheading: Part 1 - Prepare the work area

Clean the food preparation area with anti-microbial cleaner. Wipe down the fans and lights above the food preparation area.

Subheading: Part 2 - Select food items

Avoid food allergies

Avoid peanuts, bananas, and soy. Do not place bread, pasta, or ketchup in the gluten-free preparation area.

Avoid food contamination

Examine each ingredient. Verify each inspection tag is current.

USE A PERSON AS THE SUBJECT.

Instead of:

A flotation device must be worn.

Address the reader:

Wear a flotation device.

Where required, address a position:

Gardeners will wear flotation devices.

WRITE INSTRUCTIONS, NOT REQUIREMENTS.

Instead of defining a requirement:

Personnel shall wear a flotation device.

Give an instruction:

> Wear a flotation device.

AVOID "WHICH" CLAUSES.

Instead of combining two instructions, or an instruction and a long description:

> Wear glasses, which must have side shields.

Remove verbs (here: "have") from clauses:

> Wear glasses with side shields.

Replace clauses with adjectives:

> Wear side-shield-equipped glasses.

Create two sentences:

> Wear glasses. Glasses will have side shields.

ELIMINATE "ALL," "ALWAYS," "EVERY," AND "AT ALL TIMES."

Instead of:

> Always inspect every load before any lift at all times.

Eliminate redundant qualifiers:

> Inspect loads before a lift.

ELIMINATE REDUNDANT "DON'TS."

Instead of saying the same thing two ways:

> Wear long-sleeved coveralls. Do not wear short-sleeved coveralls.

Eliminate the meaningless inversion:

> Wear long-sleeved coveralls.

COVER 99% OF SITUATIONS.

Instead of:

> Where this is not practicable, the Crane Operator may allow the use of substitute tools under the following 27 conditions …

The management system should have a deviation process. Allow this process to work as designed:

> Use this tool.

KEEP MODIFIERS AT 5 WORDS OR LESS.

Instead of burying the instruction:

> In such conditions (high-velocity winds, heavy rain, compromised hull integrity, nuclear attack, etc.) that stability may not be guaranteed, …

Use a short modifier:

> In unstable conditions, …

PUT THE LONGEST ITEM AT THE END OF A SERIES.

Instead of:

> Count the horses, only those chickens that lay eggs in the yellow barn beside the old cistern, cows, and pigs.

Put the long item last:

> Count the horses, cows, pigs, and only those chickens that lay eggs in the yellow barn beside the old cistern.

FORMAT LONG SERIES AS LISTS.

Instead of:

> Personnel with corrective lenses will wear prescription safety glasses; safety bifocals with side shields; or safety glasses or impact-type safety goggles over prescription glasses.

Create a list:

> Personnel with corrective lenses will wear one of:
>
> - prescription safety glasses
> - safety bifocals with side shields
> - safety glasses or impact-type safety goggles over prescription glasses

ELIMINATE “WHYS.”

Instead of “selling” the instruction:

> We use LOTO as a means of protecting personnel from danger. It is critical that all equipment be locked out, tagged out, and tried out before attempting even routine maintenance.

Give a clear instruction:

> Isolate equipment per MS-ABC-0000.

DO NOT PARAPHRASE OR QUOTE OTHER PROCEDURES.

Instead of:

> Per MS-ABC-0000, seek authorization from the process owner and Installation Manager. Perform a risk assessment and confirm that all risks have been reduced to As Low As Reasonably Practicable (ALARP) before requesting this authorization.

Refer to the existing description:

> Request authorization per MS-ABC-0000.

DO NOT RE-WRITE OR RE-INVENT THE CONTROL-OF-WORK PROCESS.

Instead of (outside the control-of-work procedures):

> Before using the hammer, identify potential conflicting operations. Ask yourself, “How could I get hurt?” Look above, below, behind, and inside.

At most, instruct the reader to perform a risk assessment:

> Perform a documented risk assessment.

DO NOT QUOTE DIRECTLY FROM STANDARDS.

Instead of:

> Protection against the effects of noise exposure shall be provided when the sound levels exceed those shown in Table 1. For guidance on measuring sound levels, see ANSI S12.19-1996: Measurement of Occupational Noise Exposure; ANSI S1.13-1995: Measurement of Sound Pressure Levels in Air; and ANSI S12.36-R1997 Table 1. For guidance on measuring sound levels, see ANSI S12.19-1996.
>
> **Table 1—Maximum Permissible Noise Exposures**
>
> | Duration per day, hours | Sound level dBA, slow response |
> |---|---|
> | 12 | 85 |
> | 8 | 90 |
> | 6 | 92 |
> | 4 | 95 |
> | 3 | 97 |
> | 2 | 100 |
> | 1 1/2 | 102 |
> | 1 | 105 |
> | 1/2 | 110 |
> | 1/4 or less | 115 |

Create an instruction that meets or exceeds the requirement:

> The Company will provide hearing protection. Wear earplugs when outside the living quarters.

PUT THINGS IN ORDER.

Instead of:

> Hammer the nail into the wall in the approximate location you have selected for the picture. Wear gloves and safety glasses during this step. Inspect the hammer before use for cracks or other damage.

Deliver the instructions in the order they are required

> 1. Select and mark a location for the picture.
> 2. Inspect the hammer for cracks or other damage.
> 3. Put on gloves and safety glasses.
> 4. Hammer the nail into the wall.

Use Clarity to Achieve Conformity

The Code

Rewriting a procedure can feel overwhelming because it seems like there are hundreds of alternative ways to write each sentence. I'll make this easier. I'll narrow that down to **five**. If you can't write 90% of your procedure with these five templates, visit or re-visit the chapter "How to Write a Procedure" and then come back.

Five is generous. You are probably familiar with the *lex talionis* doctrine in the Code of Hammurabi: "an eye for an eye and a tooth for a tooth." Hammurabi promulgated 282 codes, and he wasn't messing around. If you violated one of the codes, you could be put to death or, for minor offenses, leave the court without an eye, tooth, hand, tongue, or breast.

Hammurabi's scribes had to make sure everyone understood the laws, lest all of Babylon end up dead or disfigured. So every one of the 282 codes follows **one** template. In Accadian, of course, but it goes something like this:

"If a person …"
"If a doctor …"
"If a merchant …"
"If a husband …"
"If a wife …"

Plenty of ways to go wrong in ancient Babylon, but uncertainty about the law wasn't one of them.

IMPERATIVE MOOD

Imperative is a grammatical mood used to express commands, requests, and advice. The imperative voice is suited for expressing urgency and clarity. This is the voice you use when you speak to pets, children, and "virtual assistants." Even then, it can sound a little harsh, so it's common to put a "please" in front.

We get away from the imperative mood for various reasons, usually to avoid "commanding" someone we can't or would rather not command. Sometimes we get our point across anyway, but we've all had trouble with

- **Subjunctive:** "If I were you, I'd feed the chickens."
- **Conditional:** "If you feed the chickens, they won't starve to death."
- **Interrogative:** "Have you fed the chickens?"
- **Infinitive:** "To feed the chickens is a daily task."
- **Gerund:** "Feeding the chickens every day is your job."
- **Indicative:** "You will feed the chickens."
- **Optative:** "May you kindly feed the chickens."
- **Jussive:** "Let it be that you feed the chickens every morning."
- **Potential:** "You can feed the chickens if you want to."
- **Negative imperative:** "Do not forget to feed the chickens."
- **Legal**: "I hereby instruct and direct you to provide sustenance to the avian livestock commonly referred to as 'chickens.'"
- **Regulatory**: "The owner shall ensure that each affected chicken is provided with adequate protection from starvation, undue hunger, and injurious weight loss while under the care of the farm."
- **Parental**: "Don't make *me* go out there and feed the chickens!"

Some of these work better than others, but the imperative "feed the chickens" is a clear winner. You might have spotted its cousin, negative imperative, in the list. I'll come back to that.

Imperative sentences are easy to put together if you understand how they work:

- No explicit subject: Imperative sentences do not have an explicit subject, as the subject is usually implied to be the listener or the person being addressed. For example, "Close the door" does not explicitly state who should close the door, but it is implied that the listener is the one who should do it.
- Use of base verb forms: Imperative sentences use the base form of the verb, without any inflections such as "-s" or "-es" at the end of the verb. For example, "Come here" uses the base form "come" rather than "comes."

The basic recipe is [verb][article][noun]:

> "Shut the door," "Pass the salt," "Water the plants," "Pick a color."

Add an indirect object if necessary:

> "Give **me** the candle," "Tell **her** the secret," "Give the food **to the chickens**."

Find an alternative for adverbs. It's pretty clear to say, "Only use the red hammer," but "Use the red hammer" is better.

Try to find an alternative for the negative imperative. Avoid "**Don't** feed the ducks." Instead, assign someone to lock up the duck food.

If all we had to do was give instructions to *anyone*, imperative mood would be the obvious choice.

- Insert tab A into slot B
- …

But our job is more complicated because we often need to specify personnel and positions for each step. You can stick a subject in front of or behind an imperative sentence, but the result isn't always as clear or polite as we'd like.

> Electrician, get a ladder.
>
> or
>
> Get a ladder, electrician.

There's a case to be made for sticking with imperative here

> Get a hammer from the storage closet.
>
> Take before pictures.
>
> Electrician, get a ladder.

We could make things even clearer by assigning *every* step

> Electrician, appoint a helper.
>
> Helper, get a hammer from the storage closet.
>
> Helper, take before pictures.
>
> Electrician, get a ladder.

But a lot of us don't want to work from something that reads so much like a script. And that wouldn't get us all the way through a management system anyway because we don't just make rules for people. It would look pretty silly to say

> Electrician, get a ladder.
>
> Ladder, be IA duty rated for at least 300 pounds.

INDICATIVE MOOD

Our second-best choice is the indicative mood.

> Imperative: Get a hammer from the storage closet.
>
> Imperative: Take before pictures.
>
> **Indicative**: The electrician will get a ladder.
>
> Imperative: Use a ladder rated for at least 300 pounds.

This isn't perfect, but it's the compromise most of us have settled on. And there's a *mixed* blessing in disguise. The fact that mixing imperative and indicative looks a little funny has encouraged us to specify a subject for *every* step in a procedure. That's a reasonable impulse, but don't take it to extremes unless you want

to start a detailed conversation about qualification, competency, and documentation thereof for getting something out of a storage closet.

The basic indicative-mood recipe (for instruction) is [subject] [auxiliary verb][verb]:

> "Jack will clean," "Jill will cook," "Jimmy will set the table."

If we had to, we could use the Indicative mood auxiliary verbs:

> The Electrician gets the ladder.

But that sounds a bit *Silence of the Lambs*. "Subject will" is the way to go here.

We've even gotten used to

> Ladders used on this site will have an IA duty rating.

As an alternative to the imperative "Use ladders with an IA duty rating." Both can get wordy when you have a lot to say about inanimate objects. Don't be afraid to use bullets. This is the one place I will forgive the use of "be" or "have."

NO MOOD

If you can't find the right verb, you may be able to communicate your instructions better without a verb.

> Requirements for ladders:

- IA duty rating
- Inspected and tagged within the last 31 days
- Clean and undamaged
- …

The "right verbs" are active, base verb forms: open, turn, lift, inspect

The *wrong* verbs are "be" verbs: be, have, is (as in "is responsible for" or, worse, "is the responsibility of").

The Five Templates

For maximum clarity and brevity (and therefore compliance), try to make every sentence into one of these:

IMPERATIVE

1. "Feed the chickens."
2. "Feed *me* the chickens."

INDICATIVE

3. "The farmer will feed the chickens."

4. “The chef will feed the chickens to me.”

DATA

5. “Menu: Chicken”

That’s it! If you can’t fit your instruction into one of these five, then you may not have an instruction to offer. I’ve seen a lot of

> The employer is responsible for ensuring that any employee who may be at risk of foot injuries from falling or rolling objects or objects piercing the sole, as well as those who may be exposed to electrical hazards like static discharge or electric shock even after other protective measures have been taken, wears appropriate protective footwear while working in the designated areas.

Which is just a re-worded OSHA requirement[5]. At least someone reworded it, but that doesn’t get our job done. Our job is not to repeat obligations; our job is to develop a sequence of actions that meet or exceed those obligations. If your procedure can’t be expressed in procedural language, then you might not have a procedure at all.

5 OSHA 1910.136(a) Foot Proction, General requirements

Right More or Less Wrong?

This starts with a simple puzzle. The riddles in this book can all be solved, but the rest are more rhetorical. This one is here for you to solve. You won't need a pencil, much less a calculator, and it will help strengthen your probabilistic intuition. Going forward, I recommend using this as an interview question for anyone pitching themselves as rational or technical. You'll want to tell them *you* figured it out yourself, so do try to work this one out.

The Premises

PREMISE 1

I have a bag of at least 1 but no more than 100 numbered tiles.

PREMISE 2

The tiles are sequentially numbered. Numbers start at 1, and no numbers are skipped, so a bag of 72 such tiles would contain tiles numbered…

1	2	3	4	5	6	7	8	9	10
11	12	13	14	15	16	17	18	19	20
21	22	23	24	25	26	27	28	29	30
31	32	33	34	35	36	37	38	39	40
41	42	43	44	45	46	47	48	49	50
51	52	53	54	55	56	57	58	59	60
61	62	63	64	65	66	67	68	69	70
71	72								

tiles_grid.svg

PREMISE 3

I reach into the bag, randomly select one of the numbered tiles, show it to you, and place it back in the bag. The tile is number **5**.

THE QUESTION

How many tiles are in my bag?

If you don't know the answer, take just a minute to think it through. While you're thinking (and to put some space between the question and the answer), here's a brief eulogy for the Big Foot.

A Brief Eulogy for the Big Foot

In 1980, Mt. St. Helens in Washington State erupted for the first time in 120 years, blowing volcanic gasses and super-heated ash 15 miles into the atmosphere (commercial airliners fly at around 5 miles up). The eruption destroyed hundreds of square miles of wilderness, killing 57 human beings and many thousands of animals.

I was 6 years old.

Perhaps five years later, I came across a magazine article on Big Foot, the mysterious, hirsute hominin rumored at the time to be hiding in untamed portions of the US Pacific Northwest. Now, the idea is ridiculous, but at the time, Big Foot was at least a lot of fun to consider. Satellite photography was poor, and portable cameras were large, fragile, expensive, and rare. There was more room then for the unknown.

But not all of us believed. Positions on Big Foot were cleanly divided between

- the incidentally-true-but-meaningless, "There is no proof"; and
- the equally true, equally meaningless, "Yeah, but absence of evidence is not evidence of absence."

Like most kids, I took the latter position. I believed in Big Foot the way I believed in giant squid—and (remember, this was the 80s) I'd never seen a clear photo of either.

But the article made a compelling point: no dead Big Feet were found after the Mt. St. Helens eruption. That fact was not easy to dismiss. If Big Feet lived on Mt. St. Helens, they should have died

on Mt. St. Helens. This was evidence that *shouldn't* have been absent. This is when I lost my belief. Big Foot died in 1980 (RIP), even if it took me five years to figure it out.

It was a sad day for mythical hominini, but an important early lesson in probability: When you can't answer a question, question an answer.

Was that a hint?

To review, I have a bag with somewhere between 1 and 100 consecutively numbered tiles. I reach into the bag and randomly select one of the tiles. The tile is number 5. How many tiles are in the bag?

We know there are at least 5, but we can't "rule out" any number between 5 and 100. So where do we go from here?

Let's start with *how* we know there are at least 5. We know there are at least 5 tiles in the bag because *if* there were only 1, 2, 3, or 4, there would not be a tile number 5. See what we did? We *questioned an answer*—four answers, in fact.

Here it is on a particularly boring graph.

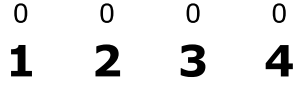

chance to draw 5 from a bag of n tiles: 1 <= n <= 4

Now let's look at one of the *possible* answers. *If* there were 5 tiles in the bag, we'd have a 1-in-5 chance of drawing the number 5 tile. Let's add that to our graph.

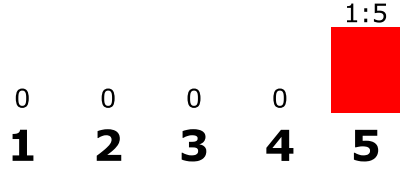

chance to draw 5 from a bag of n tiles: 1 <= n <= 5

5 is clearly a better candidate than 1, 2, 3, or 4, but is it the best candidate? Let's try a few more.

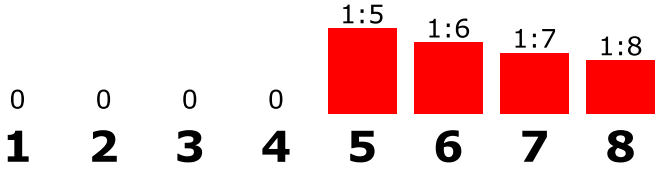

chance to draw 5 from a bag of n tiles: 1 <= n <= 8

- If there were 6 tiles in the bag, the chance of drawing tile 5 would be 1:6
- If there were 7 tiles in the bag, the chance of drawing tile 5 would be 1:7
- If there were 8 tiles in the bag, the chance of drawing tile 5 would be 1:8
- If there were [any n greater > 4] tiles in the bag, the chance of drawing tile 5 would be 1:n

The trend here is obvious. We will never see a better candidate than 5.

So the best answer is…

> When you select the **n**th tile from a bag of tiles, the best guess for the number of tiles in the bag is **n**.

The percentages will change, but **n** will always be the most likely choice, whatever tile I show you and however many tiles *might* be in the bag.

DOES THAT SOUND WRONG?

That's because it usually is.

RIGHT MORE

n is the right answer when you have to have the right answer. If you guess **n** (with a maximum of 100 tiles), you'll be exactly right around 5.2% of the time and exactly wrong the rest. If you want to be right more often, guess **n**.

LESS WRONG

But how often do we have to be exactly right? And how often do we get to be exactly wrong for free?

Instead of numbered tiles, let's imagine sequentially numbered military tanks, somewhere between 1 and 100 of them. You see one tank, number 5, and you now have to organize an opposing force with limited resources. Too little opposition is bad. Too much opposition is bad. Will you still assume there are only 5 tanks?

As far as I know, only magicians and leprechauns will ever approach you with a bag of mysterious numbered tiles, but the tank scenario actually occurred. This is known as the "German Tank Problem," referring to Allied attempts to estimate the monthly production of German tanks based on infrequent samples of ascending serial numbers.

Let's model the tank scenario

We'll model the tank scenario with our bag of tiles. Everything else remains the same, except that now, you'll have to pay if you're wrong, and the more wrong you are, the more you'll have to pay. For example, if you guess 5 and the answer is 7 or 3, you'll pay two dollars for missing by two. If the answer is 8 or 2, you'll pay three dollars for missing by three.

This is a little more math. Starting from the chart we've already seen, we have to look at every combination of guess and truth.

If we guess 5:

- There is a 6.4% chance we are correct, so the penalty would be 0.

- There is a 5.2% chance of 6 tiles in the bag, so a 5.2% chance penalty = 1
- There is a 4.6% chance of 7 tiles in the bag, so a 4.6% chance penalty = 2
- …
- There is a 0.3% chance of 100 tiles in the bag, so a 0.3% chance penalty = 95

Repeat that for every guess, and we end up here.

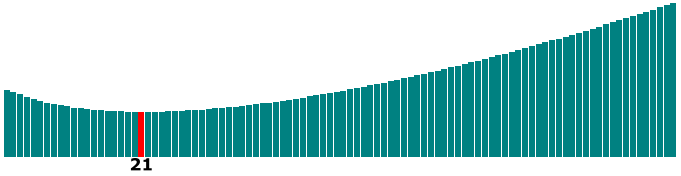

expected penalty having seen tile 5

Look closely. There is a probability mini-lesson in the fact that the sure losers (1 through 4) have a better expected return than many of the possible winners.

Following the "Less Wrong" strategy, the best guess (having seen tile 5) is 21.

Do you want to be right more or less wrong?

We have two strategies, "right more" and "less wrong". Right more is right more than less wrong (5.2% for right more vs. 1.9% for less wrong), and less wrong is less wrong than right more ($16.81 for less wrong vs. $24.75 for right more).

But that's not all you know, or need to know! Those numbers change once the random tile is revealed.

You could devise many hybrid strategies to balance right and wrong. But don't get in a hurry. Remember that you'll see some tiles more than others. If there is a uniform chance each bag will contain 1, 2, 3, 4, … 100 tiles:

> There's a 1 tile in every bag, a 2 tile in almost every bag, a 30 tile in most bags, etc.

If you get that far, devise a strategy for having seen two tiles. Then three. There's more than one answer for those too.

Whether you want to be more right or less wrong, the numbers are on your side. But they're on my side too!

THE TAKEAWAY

Math will give you *the* answer to an equation, but it will rarely give you *the* answer to a problem. And the more math you know, the more competing answers you'll find. Winning doesn't always mean you're not losing, and losing doesn't always mean you're wrong. Be wary of certainty, even when you have "the evidence".

Back to That Camel

Mawsynram, a village on the Khasi Hills in Meghalaya, India, may be the rainiest place in the world, with an average annual rainfall of around 467 inches. This makes Houston's 53 inches of annual rainfall arid by comparison. If I built a rainfall pie chart with Mawsynram and Houston, no about of Right Answer / Wrong Answer whataboutism or combinatorial demonstration would break the heuristic from the beginning of the Frequency chapter: "If it's raining, I'm in Mawsynram."

But the heuristic is still broken.

I introduced the Frequency chapter with three wrong assumptions:

1. rain happens in rainy places;
2. safety incidents happen on bad crews; and
3. if you come across a camel, you're probably in the desert.

It's time to talk about the camel.

Camels can be found in several deserts around the world:

- The Sahara Desert in Africa
- The Arabian Desert in the Middle East
- The Gobi Desert in China and Mongolia
- The Taklamakan Desert in China
- The deserts of Iran, Pakistan, and Afghanistan
- The Thar Desert in India
- The deserts of Australia, particularly in the northern regions.

Not in the United States, where we haven't had wild camels for ten-thousand years, and domesticated camels are few and far between.

I've seen two camels in the US, one at a zoo and one at a renaissance fair, neither in the desert. Those are the only two camels I have ever seen, I saw them both in Texas, and the next camel I see will most likely also be in Texas, and the next elephant, and the next penguin, and the next panda, and the next kangaroo, and—who knows?—maybe the next Tazmanian tiger.

And if I ever meet a Martian, that will probably be in Texas too. It certainly will *not* be on Mars.

Mawsynram is a small village on the other side of the world. I couldn't find it on a map. Whether it's storming, drizzling, showering, hailing, sleeting, misting, or raining frogs, you can be sure that I am not in Mawsynram. If I am standing in the rain under a sign that says, "Welcome to Mawsynram," don't bet your mortgage that I've made it to Mawsynram. If you see a picture of me standing in the rain, under the sign, shaking hands with the mayor of Mawsynram … that picture was not taken in Mawsynram.

You find camels where you look for camels.