

Mental Model
The Map is not the Territory

“You don’t fully understand something until you can simplify it. But if all you do is simplify it, you don’t understand it.”

“That’s another thing we’ve learned from your Nation,” said Mein Herr, “map-making. But we’ve carried it much further than you. What do you consider the largest map that would be really useful?”
“About six inches to the mile.”
“Only six inches!” exclaimed Mein Herr.
“We very soon got to six yards to the mile. Then we tried a hundred yards to the mile. And then came the grandest idea of all! We actually made a map of the country, on the scale of a mile to the mile!”
“Have you used it much?” I enquired.
“It has never been spread out, yet,” said Mein Herr: “the farmers objected: they said it would cover the whole country, and shut out the sunlight! So we now use the country itself, as its own map, and I assure you it does nearly as well.”
— [Lewis Carroll, Sylvie and Bruno](#)
[Concluded](#)

In 1931, in New Orleans, Louisiana, mathematician Alfred Korzybski presented a paper on mathematical semantics. To the non-technical reader, most of the paper reads like an abstruse argument on the relationship of mathematics to human language, and of both to physical reality. Important stuff certainly, but not necessarily immediately useful for the layperson.

However, in his string of arguments on the structure of language, Korzybski introduced and

popularized the idea that the map is not the territory. In other words, the description of the thing is not the thing itself. The model is not reality. The abstraction is not the abstracted. This has enormous practical consequences.

In [Korzybski’s words](#):

- A.) A map may have a structure similar or dissimilar to the structure of the territory.
- B.) Two similar structures have similar ‘logical’ characteristics. Thus, if in a correct map, Dresden is given as between Paris and Warsaw, a similar relation is found in the actual territory.
- C.) A map is not the actual territory.
- D.) An ideal map would contain the map of the map, the map of the map of the map, etc., endlessly...We may call this characteristic self-reflexiveness.

Maps are necessary, but flawed. (By maps, we mean any abstraction of reality, including descriptions, theories, models, etc.) The problem with a map is not simply that it is an abstraction; we need abstraction. Lewis Carroll made that clear by having Mein Herr describe a map with the scale of one mile to one mile. Such a map would not have the problems that maps have, nor would it be helpful in any way.

(See Borges for [another take](#).)

“We are so reliant on abstraction that we will frequently use an incorrect model simply because we feel any model is preferable to no model.”

To solve this problem, the mind creates maps of reality in order to understand it, because the only way we can process the complexity of reality is through abstraction. But frequently, we don't understand our maps or their limits. In fact, we are so reliant on abstraction that we will frequently use an incorrect model simply because we feel any model is preferable to no model. (Reminding one of the drunk looking for his keys under the streetlight because “That’s where the light is!”).

Even the best and most useful maps suffer from limitations, and Korzybski gives us a few to explore: (A.) The map could be incorrect without us realizing it; (B.) The map is, by necessity, a reduction of the actual thing, a process in which you lose certain important information; and (C.) A map needs interpretation, a process that can cause major errors. (The only way to truly solve the last would be an endless chain of maps-of-maps, which he called self-reflexiveness.)

With the aid of modern psychology, we also see another issue: the human brain takes great leaps and shortcuts in order to make sense of its surroundings. As Charlie Munger has [pointed out](#), a good idea and the human mind act something like the sperm and the egg — after the first good idea gets in, the door closes. This makes the map-territory problem a close cousin of [man-with-a-hammer tendency](#).

This tendency is, obviously, problematic in our effort to simplify reality. When we see a powerful model work well, we tend to over-apply it, using it in non-analogous situations. We have trouble delimiting its usefulness, which causes errors.

Let’s check out an example.

By most accounts, Ron Johnson was one the most successful and desirable retail executives by the summer of 2011. Not only was he handpicked by Steve Jobs to build the Apple Stores, a venture which had itself come under major scrutiny – one [retort printed](#) in Bloomberg magazine: “I give them two years before they’re turning out the lights on a very painful and expensive mistake.” – but he had been credited with playing a major role in turning Target from a K-Mart look-alike into the trendy-but-cheap Tar-zhey by the late 90s and early 00s.

Johnson’s success at Apple was not immediate, but it was undeniable. By 2011, Apple stores were by far the most productive in the world on a per-square-foot basis, and had become the envy of the retail world. Their sales figures left [Tiffany’s in the dust](#). The gleaming glass cube on Fifth Avenue became a [more popular](#) tourist attraction than the Statue of Liberty. It was a lollapalooza,

something beyond ordinary success. And Johnson had led the charge.

With that success, in 2011 Johnson was hired by Bill Ackman, Steven Roth, and other luminaries of the financial world to turn around the dowdy old department store chain JCPenney. The situation of the department store was dour: Between 1992 and 2011, the retail market share held by department stores had declined from 57% to 31%.

Their core position was a no-brainer though. JCPenney had immensely valuable real estate, anchoring malls across the country. Johnson argued that their physical mall position was valuable if for no other reason that people often parked next to them and walked through them to get to the center of the mall. Foot traffic was a given. Because of contracts signed in the 50s, 60s, and 70s, the heyday of the mall building era, rent was also cheap, another major competitive advantage. And unlike some struggling retailers, JCPenney was making (some) money. There was cash in the register to help fund a transformation.

The idea was to take the best ideas from his experience at Apple; great customer service, consistent pricing with no markdowns and markups, immaculate displays, world-class products, and apply them to the department store. Johnson planned to turn the stores into little malls-within-malls. He went as far as comparing the ever-rotating stores-within-a-store to Apple's "apps." Such a model would keep the store

constantly fresh, and avoid the creeping staleness of retail.

Johnson [pitched his idea](#) to shareholders in a series of trendy New York City meetings reminiscent of Steve Jobs' annual "But wait, there's more!" product launches at Apple. He was persuasive: JCPenney's stock price went from \$26 in the summer of 2011 to \$42 in early 2012 on the strength of the pitch.

The idea failed almost immediately. His new pricing model (eliminating discounting) was a flop. The coupon-hunters rebelled. Much of his new product was deemed too trendy. His new store model was wildly expensive for a middling department store chain – including operating losses purposefully endured, he'd spent several billion dollars trying to effect the physical transformation of the stores. JCPenney customers had no idea what was going on, and by 2013, Johnson was sacked. The stock price sank into the single digits, where it remains two years later.

What went wrong in the quest to build [America's Favorite Store](#)?

It turned out that Johnson was using a map of Tulsa to navigate Tuscaloosa. Apple's products, customer, and history had far too little in common with JCPenney's. Apple had a rabid, young, affluent fan-base before they built stores; JCPenney's was not associated with youth or affluence. Apple had shiny products, and needed a shiny store; JCPenney was known

“It was exactly his competence and past success that got him in trouble.”

for its affordable sweaters. Apple had never relied on discounting in the first place; JCPenney was taking away discounts given prior, triggering massive [deprivation super-reaction](#).

In other words, the old map was not very useful. Even his success at Target, which seems like a closer analogue, was misleading in the context of JCPenney. Target had made small, incremental changes over many years, to which Ron had made a meaningful contribution. JCPenney was attempting to reinvent the concept of the department store in a year or two, leaving behind the core customer in an attempt to gain new ones. This was a much different proposition. (Another thing holding the company back was simply its [base odds](#): Can you name a retailer of great significance that has lost its position in the world and come back?)

The main issue was not that Johnson was incompetent. He wasn't. He wouldn't have gotten the job if he was. He was extremely competent. But it was exactly his competence and past success that got him into trouble. He was like a great swimmer that tried to tackle a grand rapid, and the model he used successfully in the past, the map that had navigated a lot of difficult terrain, was not the map he needed anymore. He had an excellent theory about retailing that applied in some circumstances, but not in others. The terrain had changed, but the old idea stuck.

One person who well understands this problem of the map and the territory is Nassim Taleb, author of the Incerto series--[Antifragile](#), [The Black Swan](#), [Fooled by Randomness](#), and [The Bed of Procrustes](#).

Taleb has been vocal about the misuse of models for many years, but the earliest and most vivid we can recall is his firm criticism of a financial model called Value-at-Risk, or VAR. The model, used in the banking community, is supposed to help manage risk by providing a maximum potential loss within a given confidence interval. In other words, it purports to allow risk managers to say that, within 95%, 99%, or 99.9% confidence, the firm will not lose more than \$X million dollars in a given day. The higher the interval, the less accurate the analysis becomes. It might be possible to say that the firm has \$100 million at risk at any time at a 99% confidence interval, but given the statistical properties of markets, a move to 99.9% confidence might mean the risk manager has to state the firm has \$1 billion at risk. 99.99% might mean \$10 billion. As rarer and rarer events are included in the distribution, the analysis gets less useful. So, by necessity, the “tails” are cut off somewhere and the analysis is deemed acceptable.

Elaborate statistical models are built to justify and use the VAR theory. On its face, it seems like

“A model might show you some risks, but not the risks of using it.”

a useful and powerful idea; if you know how much you can lose at any time, you can manage risk to the decimal. You can tell your board of directors and shareholders, with a straight face, that you’ve got your eye on the till.

The problem, in [Nassim’s words](#), is that:

A model might show you some risks, but not the risks of using it. Moreover, models are built on a finite set of parameters, while reality affords us infinite sources of risks.”

In order to come up with the VAR figure, the risk manager must take historical data and assume a statistical distribution in order to predict the future. For example, if I could take 100 million human beings and analyse their height and weight, I could then predict the distribution of heights and weights on a different 100 million, and there would be a microscopically small probability that I’d be wrong. That’s because I have a huge sample size and I’m analysing something with very small and predictable deviations from the average.

But finance does not follow this kind of distribution. There’s no such predictability. As Nassim has argued, the “tails” are fat in this domain, and the rarest, most unpredictable events have the largest consequences. Let’s say you deem a highly threatening event (for example, a 90% crash in the S&P 500) to have a 1 in 10,000 chance of occurring in a given year, and

your historical data set only has 300 years of data. How can you accurately state the probability of that event? You would need far more data.

Thus, financial events deemed to be 5, or 6, or 7 standard deviations from the norm tend to happen with a certain regularity that nowhere near matches their supposed statistical probability. Financial markets have no biological reality to tie them down: We can say with a useful amount of confidence that an elephant will not wake up as a monkey, but we can’t say anything with absolute confidence in an [Extremistan arena](#).

We see several issues with VAR as a “map,” then. The first that the model is itself a severe abstraction of reality, relying on historical data to predict the future. (As all financial models must, to a certain extent.) VAR does not say “The risk of losing X dollars is Y, within a confidence of Z.” (Although risk managers treat it that way). What VAR actually says is “the risk of losing X dollars is Y, based on the given parameters.” The problem is obvious even to the non-technician: The future is a strange and foreign place that we do not understand. Deviations of the past may not be the deviations of the future. Just because municipal bonds have never traded at such-and-such a spread to U.S. Treasury bonds does not mean that they won’t in the future. They just haven’t yet. Frequently, the models are blind to this fact.

“The models themselves influence the outcome they are trying to predict.”

In fact, one of Nassim’s most trenchant points is that on the day before whatever “worst case” event happened in the past, you would have not been using the coming “worst case” as your worst case, because it wouldn’t have happened yet.

Here’s an easy illustration. October 19, 1987, the stock market dropped by 22.61%, or 508 points on the Dow Jones Industrial Average. In percentage terms, it was then and remains the worst one-day market drop in U.S. history. It was dubbed “[Black Monday](#).” (Financial writers sometimes lack creativity — there are several other “Black Monday’s” in history.) But here we see Nassim’s point: On October 18, 1987, what would the models use as the worst possible case? We don’t know exactly, but we do know the previous worst case was 12.82%, which happened on October 28, 1929. A 22.61% drop would have been considered so many standard deviations from the average as to be near impossible.

But the tails are very fat in finance – improbable and consequential events seem to happen far more often than they should based on naive statistics. There is also a severe but often unrecognized recursiveness problem, which is that the models themselves influence the outcome they are trying to predict. (To understand this more fully, check out our post on [Complex Adaptive Systems](#).)

A second problem with VAR is that

even if we had a vastly more robust dataset, a statistical “confidence interval” does not do the job of financial risk management.

[Says Taleb:](#)

“There is an internal contradiction between measuring risk (i.e. standard deviation) and using a tool [VAR] with a higher standard error than that of the measure itself.

I find that those professional risk managers whom I heard recommend a “guarded” use of the VAR on grounds that it “generally works” or “it works on average” do not share my definition of risk management. The risk management objective function is survival, not profits and losses. A trader according to the Chicago legend, “made 8 million in eight years and lost 80 million in eight minutes”. According to the same standards, he would be, “in general”, and “on average” a good risk manager.”

This is like a GPS system that shows you where you are at all times, but doesn’t include cliffs. You’d be perfectly happy with your GPS until you drove off a mountain.

It was this type of naive trust of models that got a lot of people in trouble in the recent mortgage crisis. Backward-looking, trend-fitting models, the most common maps of the financial territory, failed by describing a territory that was only a mirage: A world where home prices only went up. (Lewis Carroll would have approved.)

This was navigating Tulsa with a

“In our march to simplify reality with useful models, of which Farnam Street is an advocate, we confuse the models with reality.”

map of [Tatooine](#).

The logical response to all this is, “So what?” If our maps fail us, how do we operate in an uncertain world? This is its own discussion for another time, and Taleb has gone to great pains to try and address the concern. Smart minds disagree on the solution. But one obvious key must be building systems that are robust to model error.

The practical problem with a model like VAR is that the banks use it to optimize. In other words, they take on as much exposure as the model deems OK. And when banks veer into managing to a highly detailed, highly confident model rather than to informed common sense, which happens frequently, they tend to build up hidden risks that will unhide themselves in time.

If one were to instead assume that there were no precisely accurate maps of the financial territory, they would have to fall back on much simpler heuristics. (If you assume detailed statistical models of the future will fail you, you don’t use them.)

In short, you would do what Warren Buffett has done with Berkshire Hathaway. Mr. Buffett, to our knowledge, has never used a computer model in his life, yet manages an institution half a trillion dollars in size by assets, a

large portion of which are financial assets. How?

The approach requires not only assuming a future worst case far more severe than the past, but also dictates building an institution with a robust set of [backup systems](#), and [margins-of-safety](#) operating at multiple levels. Extra cash, rather than extra leverage. Taking great pains to make sure the tails can’t kill you. Instead of optimizing to a model, accepting the limits of your clairvoyance.

The trade-off, of course, is short-run rewards much less great than those available under more optimized models. Speaking of this, Charlie Munger has noted:

“Berkshire’s past record has been almost ridiculous. If Berkshire had used even half the leverage of, say, Rupert Murdoch, it would be five times its current size.”

For Berkshire at least, the trade-off seems to have been worth it.

The salient point then is that in our march to simplify reality with useful models, of which Farnam Street is an advocate, we confuse the models with reality. For many people, the model creates its own reality. It is as if the spreadsheet comes to life. We forget that reality is a lot messier. The map isn’t the territory. The theory isn’t what it describes, it’s



simply a way we choose to interpret a certain set of information.

Maps can also be wrong, but even if they are essentially correct, they are an abstraction, and abstraction means that information is lost to save space. (Recall the mile-to-mile scale map.)

We must always be vigilant by stepping back to understand the context in which a map is useful, and where the cliffs might lie. Until we do that, we are [the turkey](#).

fs

How do we do better? This is fodder for another article, but the first step is to realize that *you do not understand a model, map, or reduction unless you understand and respect its limitations.*

This document is provided to [Farnam Street Members](#) and is intended for a single recipient. If someone passed it along to you, consider subscribing. If you would like to order multiple subscriptions for your team with a group discount (minimum 5), please [contact us directly](#).