McKinsey & Company

Strategy and Corporate Finance Practice

How to beat your biases and make better investment decisions

Cognitive biases can trip up even the most experienced executives but can be defeated using proven techniques.

August 2019

This episode of the *McKinsey Podcast* features highlights from McKinsey's Inside the Strategy Room podcast series. In recent episodes, the series has focused on insights from *McKinsey Quarterly's* Bias Busters collection, which delves into how executives can harness insights from behavioral science to make better business decisions.

Podcast transcript

Simon London: Hello, and welcome to this episode of the *McKinsey Podcast*, with me, Simon London. Today, we're revisiting the topic of cognitive bias—the wholly irrational but largely predictable patterns of thinking that get in the way of good decisions.

This episode is actually a "best of"—bringing you some highlights from our Inside the Strategy Room podcast series. If you like what you hear, please subscribe to that series.

Providing the insights are Tim Koller, a McKinsey partner, and Dan Lovallo, a former McKinsey consultant and professor of business strategy at the University of Sydney. Asking the questions is the host of the podcast series, Sean Brown.

To lead off, here's Dan Lovallo talking about how to overcome the biases inherent in what's known in behavioral science as "the inside view."

Dan Lovallo: Quite a while ago, I think about 1993, Daniel Kahneman and I came up with the terms the "inside view" and the "outside view" to explain different modes of thinking. The inside view is the way people usually think about decisions—in other words, they start on the problem and they build the case. In business, the case often starts with an Excel spreadsheet. You start putting in numbers, and they look into a crystal ball and try to see the future and plan that future out in advance. That's the natural way to think.

Another way to do it is to use lots of analogies or cases. When you use these cases, you take a more statistical view.

I can tell a story about the movie industry. I thought I was going to get rich—that's the important part of this story. I did a study, along with Colin Camerer. What we were able to do was forecast movies out of the domestic box office based simply on a poster and a one- or two-line description on Metacritic. What we got were the similarity ratings between the focal movie we were trying to predict and a reference class.

The reference class was formed by genre, actors, and story line. That's it. We took the intersection of those, and there were more than enough movies. When we added similarity ratings to the typical regression things they use like budget and whether it's an action movie or not, our mean average residual error was 25 percent, which is exceedingly low for something like that based on the fact that all they needed was a poster and a paragraph about the movie.

The reason I thought I was going to get rich was because I got the meeting with the head of the studio, and I walked him through our data. I said, "What's your error on average?" And he said, "One hundred percent. We think a movie is going to make $100 million, and it makes $50 million, or it makes $150 million. We're just not close." And I said, "Well, you know, I can cut that by 75 percent. Let's work together." And he thought for a while, and said, "No, I don't think I can do it." And I said, "Well, how many analogies do you use to make forecasts of movies at the beginning?" And he said, "Well, sometimes we use one." In other words, they use almost all the inside view.

And I said, "What's the most you've ever used?" And he said, "Two." I said, "Well, do you believe what I'm doing? If you don't believe me that we did this beforehand, give me your next slate of movies. If I'm helpful, you pay me, if not, it's free." I thought I can't go wrong there. And he said, and this is where he was very honest, he said, "No, it's not that. I get to pick 12 movies a year. And I'm only in this job for a few years. I don't want evidence out there that I should've done something different." Very candid. And there was no money for me.

Sean Brown: It seems that one of the key challenges is figuring out what those reference cases are. Are there any specific skill sets or techniques that help somebody who's trying to take the outside view to do it effectively?

Dan Lovallo: It's part art and part science. A way to think about it, for those who know how to do regressions—and I'm not going to assume everyone does—but they would be the main variables in your regression. In the case of movies, it's genre, actors, story line, and things like that. You can do it with almost anything.

Tim Koller: A lot of it will depend on the industry. It's not something you can pull from the financial statements. But if, for example, you're in the chemical industry, first, you've done these things before yourself, and, second, it's pretty easy to see what your peers have done, and what's worked, and what hasn't worked. Often these are big, visible projects. It's a combination of past projects, with projects inside the company, plus knowledge of what your peers have done. There will be people who will know what's happened in other companies and what their experiences were that you can bring into this as well.

Dan Lovallo: This is an important point. There are three types of learning that can go on. One, if you're the decision maker, you learn from your own past experience. Two, you can learn from the past experience of the company or organization that you're working for. And then the step that people rarely take is learning from others' experience.

There's an awful lot of information there. And I haven't been presented with a problem where I couldn't come up with a reference class.

Overcoming the anchoring bias

Simon London: With the inside view vanquished, the team turns its attention to a bias known as anchoring. For example, during the annual budget process, the numbers for next year typically end up looking pretty much like this year's numbers, plus or minus a few percent. How can companies overcome anchoring in order to reallocate capital more quickly and effectively? Again, leading us off is Dan Lovallo.

Dan Lovallo: Anchoring was something developed by Daniel Kahneman and Amos Tversky. It refers to the tendency to anchor on any number given. One of the best examples of it was people were asked to estimate Gandhi's age at death. This may be imprecise, but one group was given an age of ten and another was given an age of 140. These are obviously completely ridiculous numbers, but they strongly affected the predicted age of death. If something that you know is ridiculous impacts your forecast, then imagine how much more salient last year's budget is.

We call this endowed anchoring. You're endowed with this anchor, which is what you did last year. You're anchored to that in psychological ways.

Tim Koller: Another example that I think is interesting: we take a group of executives in a room and split them into two. We give one group a set of opportunities for different resource allocation for the coming year and what the payoffs would be. That's all you give them. You give the second group exactly the same information, plus we give them the information about how the resources were allocated across units the prior year. What you'll find is the answer is very different between the two groups because that second group automatically assumes this bias that maybe last year's information was right. There is an innate push toward what happened last year, and that is a major problem that has to be overcome in organizations. Organizations typically start with last year and then move up or down incrementally. Rarely do they start from scratch and think, "What's the best place to allocate our resources?"

Sean Brown: Clearly there's an unconscious bias that's going on with anchoring. But there might be other reasons why budgets don't change year to year. I'd imagine it's quite disruptive to reallocate people and capital dramatically on a year-to-year basis. How have you seen companies that have overcome the anchoring bias actually then overcome some of those other challenges associated with significant reallocation?

Tim Koller: I might challenge your premise about how difficult it is to do that. I think that's often an excuse. It's not that you're going to change a lot every year across different business units. In other words, if one business unit is on the upswing and it needs more resources, that's probably going to be the same for a couple of years, versus a business unit which is more mature. I don't think it's a good excuse to say it's hard to do, because you need to do it at some point or another. And better to do it sooner rather than later.

With respect to projects, though, it is much easier because you typically have resources that you can move. For example, some of the oil companies will just do a forced ranking of all the different projects. Then they can pick the projects. Initially they do it mechanically and pick projects that have the highest returns. You might make some modifications later on, but that is a technique you can do for business units as well. You start out with something which is a forced ranking, which ignores history, and is entirely forward looking. Therefore, your starting point in the discussion is an allocation that is sometimes very different than what you did last year. You might modify it a little bit, but at least you've got a starting point that pushes you away from the anchor of last year.

Dan Lovallo: The only way to fight an anchor is with a new anchor. This is a little bit like how you form a reference class and a little bit like how you do a regression. You pick some of the biggest indicators. In other words, let's say you're a consumer-goods company, and you've got a store in Hong Kong. You might look at the growth of the entire sector, the growth in Hong Kong, and the growth of your product line. This is a quick-and-dirty thing. It doesn't have to be fundamentally detailed. You take these attributes, and you run a quick regression model and predict what you think the sales would be in the Hong Kong store for this product line. And then you're going to compare that to your inside view, which means focusing on your own plans and actions and last year's budget, to set up a different anchor.

Tim Koller: The other thing you have to overcome is the internal dynamics or politics of the organization. For example, one technique people will use is to try to talk to everyone who's going be in a meeting before the meeting to get them to buy in so there's no debate at the meeting, and they get their point of view. Or people will go around the process directly to the CEO and say, "I really need those resources." Because of this, you really have to be strict about the processes that you go through in order to overcome those internal dynamics that will work against the more fact-based objective reallocation. People's prestige and their view of how they're perceived in the organization are all at stake here. Often, they're not looking at it from the perspective of the whole company. They're going to try to move things back to where they had them. This is particularly true for those who have a lot of resources. You have to take specific steps to prevent that from happening.

Sean Brown: Tim, does that lend itself more to an annual budget process, or one where it happens more regularly? Is there a way to lower the stakes so there's not as much pressure to try and save your resources?

Tim Koller: We did some research and found that when companies did their strategic planning at the beginning of the year, they put together a three- to five-year forecast. The first year of that forecast, ideally, should be the budget for the next year. But what we found was that when companies actually got around to budget time—let's say they did they did the strategic plan in June, and the budget in October—the budget reverted back to the prior year's allocation even though in their strategic plan they were going to have a big reallocation. One thing you can do is be more precise, or require that you don't make changes from that three- to five-year plan to the next year's budget through the process.

Another thing that you can do to give you more flexibility during the course of the year is to not allocate all of the resources and hold some back either for the CEO, or for the investment committee, or whatever body makes those decisions. That way when opportunities come along, you can allocate new resources.

The other thing you need to do is to make sure that even if you did allocate something to a project, if you're not going to start spending the money until next September, when September comes around, you actually revisit and say, "Do we still want to spend it?" You might take money back. You might have a budget that is not set in stone. You can add pieces in some cases and take it away as you gain more information.

The sunk-cost fallacy

Simon London: Next, the team turns its attention to the infamous sunk-cost fallacy—the bias that leads us to continue with an activity based on the time and money we've already put into it, as opposed to taking a clear-eyed view of the prospects. As we'll hear, the sunk-cost fallacy often gets in the way of decisions to close projects that should have been ended years ago, or indeed to divest perennially underperforming business units. Here's Dan Lovallo to start the discussion.

Dan Lovallo: We did some research that looked at projects that were both unprofitable and cash needy. In other words, their sales growth couldn't sustain the investment. You would think that somebody would kill them at some point. But the situation is those projects or business units had an 80 percent chance of surviving ten years.

This is really common. It's about as common as you get. Part of the difficulty is there's a lot of emotion tied to firing people. Naturally, people don't want to do it. And so they don't, and these things tend to keep living. If people don't have processes in place, it's very unlikely to happen.

Tim Koller: Some other research, as it applies to business units, for example, is that we found that companies often would not divest a business unit until a year or so after people had already started talking about it in the press.

The rule of thumb seems to be that it takes about two years after a company starts to have the conversation before they actually get around to doing anything about it. Or they come into the mode of, "Oh, we'll fix this thing before we sell it," which never happens, of course. You lose a lot of value over those two years by waiting. In addition to projects, divesting or pruning whole business units in the portfolio is also something that companies need to be better at in order to maximize the value creation.

Dan Lovallo: In the US economy, for stand-alone companies, that's about 30 percent of the companies that have at least one business segment that's not dying soon enough. And for multi-business-unit companies, that's about 43 percent of the companies. It's a ubiquitous problem.

Tim Koller: The reason they fall into this trap is typically because of the sunk-cost fallacy: "We've invested all this capital and time and energy into it, so let's keep it going."

We all know this does not make any sense at all. You need to have processes to eliminate that.

One approach is something that we sometimes call "stage gating." Every time you want to spend more money on a project, at a certain phase of that project's life to move onto the next phase, you have to go back and get permission to keep it going. There's not the automatic presumption that once a project gets started, it will go to completion. The presumption is at certain milestones that are predefined ahead of time, you have to get permission to keep going with the project.

Dan Lovallo: And there's a point in there: you have to define these milestones in advance, or else you're going to get slippage and you're going to fall into the sunk-cost fallacy. You have to have somebody who's in charge who's going to actually stop things if you don't meet those milestones.

Sean Brown: One of your suggested remedies is something you call the burden of proof. Can you say more about that, Tim?

Tim Koller: The burden of proof is the idea that instead of making the case for divesting, you have to change it to, "Why should we keep this business?" Obviously, you can't do it for the whole company every year, but one approach is, for 5 or 10 percent of the big projects or units every year, you ask that question, "Why should we hold onto this particular unit?" Or, "Why should we keep pursuing this particular project?" It's a way of changing the sequence around so that if you can't prove the case that you absolutely need to keep it, then the assumption is, "OK, we're going to get rid of it."

Sean Brown: In your article, you raised the notion of categorizing business investments.

Tim Koller: In this case, what you need to do is make sure that every unit or every investment is put into one of three or more categories.

There are those businesses or things that we want to accelerate. There are those that we want to maintain or defend. And then there are those that we need to dispose. And you force yourself to categorize each business into one of those so that, once again, if you can't make the case for maintaining or accelerating a project or a business, then it automatically goes into the dispose category using some criteria that you don't let the processes get in the way of.

Dan Lovallo: One thing that's absolutely important is this won't happen if it doesn't come from the top. It simply won't happen. The CEO has to make a call, and this has to be driven from the top.

Resisting 'glamour' projects

Simon London: And finally in this episode, from value-destroying projects that refuse to die, to value-creating projects that don't get the investment they need because they are not, well, very glamorous. What can executives use to ensure that the allure of "glamour" projects doesn't get in the way of optimal resource allocation? Dan Lovallo starts this segment with a good example.

Dan Lovallo: My favorite example is when you've got a big oil field there are a number of things that have to be done. One of them is drilling new resources, oil, or gas. And that's exciting. Now there's another thing that has to be done in a big oil field, and that's maintenance. You've got to maintain the wells. You've got to make sure that nothing untoward happens. That's not so glamorous.

Frequently, those two necessary activities are completely separated. The way to go forward in the company is to be a great explorer, not so much a great maintenance person. That creates a problem and has created some issues that we all know about.

This kind of problem isn't just limited to my favorite example in oil and gas. Tim can talk about how this generalizes to many more types of projects.

Tim Koller: When companies typically evaluate projects, like maintenance projects or the performance of a business unit, they typically make the assumption that if they don't make an investment, then things will go along business as usual. Therefore, they end up with the assumption that a maintenance project has a zero present value. Or a business unit's value is a steady-state kind of thing. The reality is that in many of these cases if you don't make the right decisions and make the right investments, that project or that unit will decline in performance. But it's often not acceptable to show a base case which is a decline. The base case is almost always a steady state. Therefore, the investments that keep it at a steady state don't show up as creating any value. And that's the source of a lot of the issue with respect to focusing on the new versus the maintenance of the old in some situations. That's a mind-set that has to change.

Sean Brown: What are some of the tips and tricks that you would offer our listeners in terms of better valuing these maintenance approaches?

Tim Koller: One of the techniques is to make sure that different parts of the organization aren't operating in silos. Decisions should be made where people who are doing maintenance things understand what's going on with the faster-growing businesses and vice versa so people can understand what it's all about.

Dan Lovallo: And one of the mechanisms from the oil and gas example earlier is really quite simple and I think somewhat ingenious: it's the committees in most oil and gas companies. There'll be an exploration committee and a maintenance committee, and different people oversee those. This leads to underinvestment in maintenance because there are more senior people on the exploration committee.

What one company did—and this company had a great safety record—was to have overlapping committee members. This allowed them to more easily balance out investments in maintenance and exploration.

The committee members that overlapped on the exploration committee but also sat on the maintenance committee could see the needs of maintenance and had an easier way to argue to the leader of the field and say, "Listen, they really need this money." Overlapping committees are a way to get more information to the top in a more direct manner.

Tim Koller: Another is the CEO has to take the lead in saying, "We're no longer going to accept the base case being business as usual." There are base cases where the business or the project is declining or we're going to destroy some value because we're not spending enough money on maintenance, for example. It is essential to have that kind of mind-set built into the organization, where we're always looking at, "What would really happen if we didn't spend this money?"

Sean Brown: How do you tie this to the depreciation charges that companies are taking on their capital base?

Tim Koller: The depreciation is a crude metric, but sometimes it can be very useful. If your spending is less than the depreciation, you might want to be worried. If you're spending much more but not growing, you want to be worried as well. It's a starting point but it's not sufficient. You really need to look at the business from the bottom up to make the right decisions. Once again, you don't want to be anchored in the past. You want to be thinking about the future. The people who are closest to those assets will know what needs to get done or when something needs to get done. And you need to have a mechanism where you can hear their voices so that you can take that into consideration as you're making your investment decisions.

Sean Brown: Any other suggestions for how executive teams can avoid this issue?

Dan Lovallo: One of the ways is to take personality out of the system. In one company, people are allowed to propose projects. The head of the business unit or division proposes the projects and puts them in writing. When the allocation decision is made, they're not in the room. The CEO, the CFO, and the head of technology make the decision without anyone else. And you don't get the head of the most glamorous division, or the head of the biggest division, or somebody who's been there longest arguing for their project face-to-face. That makes a big push toward objectivity. I would recommend that to all companies.

Simon London: So there you have it. A hatful of ideas for overcoming cognitive bias from Tim Koller, Dan Lovallo, and the host of the Inside the Strategy Room podcast series, Sean Brown. If you're hungry for more, please do subscribe to the series, which you'll find in all the same places you find the *McKinsey Podcast*. Also check out the most excellent series of articles entitled Bias Busters, which you can find on McKinsey.com or the McKinsey Insights app.

Sean Brown is McKinsey's global director of communications for strategy and corporate finance and is based in the Boston office, **Tim Koller** is a partner in the New York office, and **Dan Lovallo**, an alumnus of McKinsey's San Francisco office, is a professor of business strategy at the University of Sydney. **Simon London**, a member of McKinsey Publishing, is based in the Silicon Valley office.

Designed by Global Editorial Services