



CAPITAL ALLOCATORS

How the world's elite money
managers lead and invest



TED SEIDES

Bestselling author and host of the
Capital Allocators podcast



About *Capital Allocators*

Who are these capital allocators and why have conversations with them on the *Capital Allocators* podcast been listened to four million times?

Capital allocation is the process of deciding where to invest limited resources. It occurs when corporate executives decide how to fund operations and initiatives. It occurs when portfolio managers at investment firms select and size positions. And it occurs when Chief Investment Officers (CIOs) for end owners of capital commit to investment products. The concept of capital allocation pervades finance and its real-world consequences are far reaching.

I created the *Capital Allocators* podcast in 2017 to focus on this last group, the people who sit at the top of the global food chain of capital. When I refer to “allocators” in this book, it is the CIOs and their teams that I have in mind.

At the bottom of the food chain, businesses are conceived by entrepreneurs. They raise capital to pursue their ideas by selling a piece of their business or borrowing money with the intent to pay it back. They make business decisions about which products or services to create, markets to enter, and forms of financing to fund new projects or return capital. These decisions are important ones in the success of their business. They are also different from the decisions faced by CIOs.

The middle of the food chain is populated by money managers. When overseeing their business, they make choices like those of business executives. When producing their product, they make capital allocation decisions like CIOs, picking investments and constructing portfolios.

CIOs sit at the top of the food chain. Endowments and foundations, high net worth individuals, family offices, corporate and public pension funds, and sovereign wealth funds are end owners of capital. CIOs lead their investment operations.

End owners of capital frequently staff their investment team with a small number of professionals. The team occasionally invests directly in securities or deals, and more frequently allocates capital to the products run by money managers.

Money managers often specialize in one style of investing. For example, Andreessen Horowitz, or a16z (Scott Kupor – *Capital Allocators*, First Meeting

episode 7), is a leading venture capital firm and Wellington Management (Jean Hynes – episode 82) is best known for public equity and fixed income management. Others are intermediaries, like Outsourced Chief Investment Officer (OCIO) Hirtle Callaghan (Jon Hirtle – episode 98), who raise capital from end owners and invest it in money management products.

There are so many specialized, well-resourced money managers in different strategies around the world that one sensible investment approach for CIOs is to identify the best of them across each area of expertise. CIOs partner with specialists, rather than hire a team to compete with them.

CIOs are students of the entire investment universe. They have a unique vantage point to look across assets, strategies, and geographies and build a portfolio of investments that best meets their objectives. They assess both the merits of the money management organization and the attractiveness of the underlying assets held by the manager. They have access and exposure to talented minds across the investment universe. Their seat is the broadest and most fascinating in all of investing, and their decisions significantly influence how capital flows throughout the world.

I started my career working for the most famous CIO in the world, David Swensen, who for 35 years has overseen the Yale University Investments Office. He is a brilliant investor and gifted teacher. A dozen of David's protégés manage some of the largest endowments and foundations.¹

In the 20 years after my apprenticeship at Yale, I continued to study money managers across public and private markets. Three and a half years ago, I created a podcast called *Capital Allocators* to share conversations with a range of CIOs around the world. I hoped to learn about different frameworks and best practices for how these holders of the keys to the kingdom allocate their time and money.

I never expected to write another book. Most first-time authors lose the naïve boon of seeing their name on a book jacket once they meet the reality that writing a book sucks up time and money, two of the most precious resources we

1 My colleagues at Yale included Andy Golden, CIO of PRINCO, Ellen Shuman, former CIO of Carnegie Corp, Donna Dean, former CIO of Rockefeller Foundation, Mary McLean, former CIO of Kauffman Foundation, Seth Alexander, CIO of MITIMCo, Lauren Meserve, CIO of Metropolitan Museum of Art, Paula Volent, CIO of Bowdoin College, and Casey Whalen, CIO of Truvvo Partners. Other notable alumni who came after my time include Robert Wallace, CEO of Stanford Management Company, Kim Sargent, CIO of Packard Foundation, Peter Ammon, CIO of University of Pennsylvania, Anne Martin, CIO of Wesleyan University and Randy Kim, CIO of Rainwater Charitable Foundation.

have. Yet, here I am, sitting in front of my computer, following the serendipity that called me to put pen to paper once more.

After so many interviews with allocators, managers, and thought leaders – where each had wisdom to share about how to enhance the investment process – I no longer could retain all the information in my head. David Swensen’s book remains the seminal tome in explaining the structure of capital allocation of institutional portfolios. This book is different. It endeavors to explain how these investors implement their craft.

My intent in writing is to help investors make better decisions. Investment success is increasingly challenging with every passing year, and the breadth of skills required to succeed expands commensurately. The lessons imparted by guests on the podcast speak to the needed skills and prerequisites for success going forward.

Introduction

“Did you find the podcast, or did the podcast find you?”

– *Michael Mervosh*

On a wintry day in February 2016, I sat on a chairlift at Haystack Mountain in Vermont with my friend Gregg Clark. Gregg is a warm-hearted guy who I’m quite sure could be a world champion on *Jeopardy*. He holds encyclopedic knowledge of both the useful and useless, equally facile in the arcane details of financial models and the wisdom of the Toltec culture. An engineer by training, Gregg saw one weekend that the full-sized electronic bowling alley in the ski lodge had stopped working. He heard the manufacturer couldn’t repair it until Monday, so he took the manual overnight and fixed it himself the next morning.

As we rode our way up the mountain, Gregg asked if I listened to podcasts. At the time, I had barely heard of a podcast and didn’t know the purple app on my iPhone existed. He directed me to *The Tim Ferriss Show*, and I was hooked.

Around the same time, my first book *So You Want to Start a Hedge Fund: Lessons for Managers and Allocators* released. After 14 years focused on sourcing, researching, and investing in early-stage hedge funds, I had acquired a body of knowledge about the ecosystem that few aspiring entrepreneurs in the industry understood. I noticed I was having the same conversation repeatedly and decided to put pen to paper so more people could learn from the past lessons of others.

The book also marked the end of a chapter in my professional life. I was keen to leave a narrowly siloed world of investing in small hedge funds and get back to a broader mandate akin to where I started my career.

In this window of transition, I received invitations to share the lessons of the book publicly. I appeared on television, radio, and lo and behold, a series of podcasts. One podcast conversation stood out. In October 2016, I sat down to speak to Patrick O’Shaughnessy in what became the seventh episode of his podcast, *Invest Like the Best*. I was blown away by his preparation and ability to ask just the right questions. We became fast friends.

The experience with Patrick demystified what happens behind the curtain of podcasting. We met in his office and had a conversation in front of two microphones. It seemed as simple as that.

A short time afterwards, I saw a Facebook post that my college friend Chris Douvos was a guest on a different podcast. Chris is the highly entertaining ‘Super LP’ of venture capital and founder of Ahoy Capital. I listened and was appalled. The questioning was so rote and scripted and the energy so low that the hosts managed to make Chris sound dull. If you listen to my conversation with Chris (*Capital Allocators*, Ep.14), you’ll understand that it’s quite a feat to crush his energy.

That day, a lightbulb went off in my head. I had time on my hands and wanted to reconnect with some of my old friends who I didn’t have time to see in the day-to-day grind of managing money. I aspired to interview with Patrick’s natural dexterity, and I figured I could do better than the folks who interviewed Chris.

I created a podcast interviewing people who are tasked with managing large pools of capital. I called Patrick to inquire about how it all worked. He sent me on a voyage to purchase a laundry list of equipment, record a few conversations, and connect with Mathew Passy to assist with the production. Off I went.

I recorded a fantastic conversation with my friend Steve Galbraith, who had written the foreword for *So You Want to Start a Hedge Fund*. I guided the conversation through Steve’s remarkable career and filtered in fun stories about his owning a local brewery and European football team. I couldn’t believe how well the conversation went and was excited to share it.

There was only one problem. I lost the recording.

For the next six hours I couldn’t find the recording on my H6 Handy Recorder. I searched the internet for a solution, called customer service at the manufacturer, and reached out to Mathew. All of it was to no avail. I was about to declare the end of my nascent podcasting venture when I thought of calling Gregg Clark, the one person who might have an answer when no one else did. Sure enough, Gregg hooked up the device to his computer, plugged in a few extra wires, hit play, and an hour and 11 minutes later downloaded a perfect digital recording of the episode out of thin air.

Before launching the podcast, I recorded two more conversations with dear friends. One was with André Perold, my former business school professor. The other was with Paula Volent, my former colleague and star endowment

manager at Bowdoin College. A week later Paula called to inform me that she couldn't get permission to put it on air. My plan to launch with three episodes got cut by one-third just like that.

I released the episodes with Steve and André and sent an email to friends in my contacts. Two hundred of them listened to one of the initial episodes. From there, I asked some friends to join me for an interview. A few conversations turned into a few more and three and a half years later, here we are.

* * *


That's the chronological story of how the podcast came together, but there's another side to the story that is harder to explain. Michael Mervosh described a different way of moving through the world during our podcast conversation. He posed the question, "Did I find the podcast, or did the podcast find me?"

The podcast journey was one of the first times in my life that I took on a professional challenge without a goal in mind. I thought it would be a fun and productive way to spend some time while searching for my next full-time engagement. If nothing else, I suspected the adventure would create some optionality, even if I had no idea what those options might be.

I was full of doubts when getting started. The only thing I knew for sure was that I didn't like the sound of my own voice and was fairly confident others wouldn't either. I didn't know if I would interview well or if I would be able to keep a conversation flowing. I wondered if I would run out of guests to interview, of subjects to discuss, or of time to keep going. And I never once thought of the podcast as a business, as even a cursory level of analysis would have suggested that offering content for free to a finite natural audience is not a road to riches. But I really didn't care about that, because I was ready to try something new without feeling tied to an outcome.

As a result, I can't explain what happened over the last three and a half years. I sent out one email, started posting episodes on my empty LinkedIn and Twitter accounts, and through word-of-mouth the 200 downloads that first week grew to 20,000× that number in a trajectory that looks like a stock chart presented in thousands of manager pitchbooks.

WEEKLY DOWNLOADS OF CAPITAL ALLOCATORS



I started receiving positive feedback from old friends and new. Some commented on my excellent “voice for radio,” rebuffing the only definitive prior I held. Others reflected an interview style of humility and curiosity that I would describe as a mirror of my best self. Most importantly, the podcast guests universally derived value from the experience, commenting on the platitudes and unexpected benefits they received after coming on the show. I loved being in the business of building goodwill.

As I released episodes one week after the next, I started to realize I had come across an educational platform unlike anything else available. I had the privilege of working with many top-flight investors in my two decades of experience, and 45 of the first 50 podcast guests were friends and peers from that time.

I hadn’t quite appreciated how little these insightful investors were in the public eye. One of my favorite examples came from Jim Williams, the long-time CIO of the Getty Trust. Following our conversation, Jim mentioned how much he had learned from the podcast with his peer, Scott Malpass, the recently retired CIO of Notre Dame University. I was surprised, knowing that the two had been friends for decades. Jim explained that although they speak a few times a

year, he never before had the chance to hear Scott’s story in its entirety. At that moment, I realized that I hadn’t either, nor had just about anyone else.

My guiding light in building the roster of guests has been answering the questions of who I am curious to speak to now, and what I would like to learn. I love hearing personal stories – investing is indeed a people business – so all the episodes start that way. Many guests fit into my original concept of diving into the people, philosophy, and process of the holders of the keys to the kingdom in asset management. I’ve also explored other disciplines that help allocators of capital think about improving their skill set.

Three years into this adventure, the podcast has gone beyond my expectations and been great fun. I’ve had wonderful conversations with brilliant thinkers and met incredible people along the way. It has proven a fruitful exercise in what Patrick has dubbed “growth without goals.”

Before we get to the lessons from the many guests on the show, I want to share some thoughts about the active vs. passive debate. The podcast and this book are entirely about the pursuit of active management through allocation to money managers. Critics of active management think this entire pursuit is a folly. I disagree, so let me tell you why.

Chapter 2

Decision-Making

“If Danny Kahneman still commits these behavioral errors himself, how is Drew Dickson not going to?”

– *Drew Dickson*

Back in 2007, a good friend of mine led the due diligence on a European hedge fund run by a former Soros Fund Management portfolio manager. My friend was the head of research at a large fund of funds and was incredibly thorough. In fact, his team’s work was so deep that managers referred to their process as a proctology exam.

My firm had seeded the hedge fund a few years before and after a strong stretch of performance, it crossed \$1 billion in assets. I watched my friend and his team gather more information about the fund than anyone else. They understood everything about the manager and his approach and confirmed every detail along the way. I don’t ever recall seeing more t-s crossed and i-s dotted.

During the process, he learned that the fund’s returns were volatile. Historically, his firm preferred managers that delivered a smooth return stream. When we discussed the potential mismatch, he shared that his investment committee was looking to add more juice to their portfolio.

With all the information available in their hands, the investment committee approved a substantial investment. Sure enough, the fund hit a rough patch shortly thereafter and within a year, the fund of funds redeemed. I have wondered what happened that day in the committee meeting ever since.

* * *

All the information gleaned from interviewing and conducting due diligence on managers can go for naught if an allocator does not follow an effective decision-making process.

And here’s the problem: we are hardwired to make poor decisions.

The field of behavioral economics sheds light on this unfortunate set of conditions. The human brain is designed to survive in the wild. Generations ago a fight or flight response was the difference between life and death. Daniel Kahneman referred to this instinct for action as System 1 thinking in his book, *Thinking, Fast and Slow*. Most decisions we make are based on instinctive reactions. Occasionally, we consciously shift to System 2 thinking in which we carefully process information.

A CIO seeks to follow a sound process and get as close as possible to discovering the truth with the information available at the time. It follows the tagline of the Alliance for Decision Education, a non-profit co-founded by Annie Duke, “what is true and what to do.”

We’ll start by developing an understanding of why it is so hard to make good decisions and then turn to how to improve the process.

Why this is so hard

Annie Duke, poker legend and author of *Thinking in Bets* and *How to Decide*, spent 20 years at a petri dish for decision-making: the poker table. In each hand of poker, she would make moves, obtain real-time feedback, and have an opportunity to assess her decision process.

Annie’s research in the field helped her understand why people make bad decisions. Decisions are like bets that arise from beliefs about the way the world works. When a CIO buys a stock or a fund, they are betting on their belief that the investment is better than the many other options of putting the money to work.

The way humans form beliefs is counterintuitive. One might think that we hear an idea, ponder it, decide if it’s true, and then form a hypothesis. But that is not how the brain works. Instead, we hear an idea, immediately decide it is true with our System 1 brain, and maybe later shift to System 2 thinking and vet the idea.

Beliefs, in turn, are based on facts and predictions. Our ability to process facts and make predictions is compromised. Annie points to confirmation bias, motivated reasoning, and tribe identification as three of the most pernicious ways the inputs that form our beliefs can be flawed. Confirmation bias is a well-known behavioral flaw that we notice and emphasize information that confirms our existing point of view. Michael Mauboussin said, “If you’re in the investing

business and you haven't fallen for confirmation bias, you're not doing your job." Motivated reasoning is confirmation bias on steroids. Investors work incredibly hard to discredit information that is contradictory to their hypothesis. Lastly, our affiliation with tribes influences our views. As an investor, we might be a card-carrying member of the value investing tribe, the trend following tribe, or the long-term investor tribe. We don't trust information from sources that arise from outside of our tribe.

Just becoming aware of our biases isn't enough to allay the pattern. We suffer from resulting, self-serving, and hindsight biases that inhibit our ability to improve our decision-making process. Annie refers to "resulting" as the inability to separate a good decision from a good outcome when assessing others. When allocators see good performance from a manager, they think the manager has made good decisions; bad performance must come from a bad decisions. As Paul Isaac of Arbiter Partners mocks managers, "volatility is only downside deviation. Upside deviation is called performance." All investors have a self-serving bias that wants to take credit when good things happen and defer blame when bad outcomes arise. It takes System 2 thinking to investigate when a bad outcome came out of a good decision process, and is even harder to recognize when dumb luck caused a good outcome to arise from a bad process. Hindsight bias gets in the way of accurately assessing past decisions. Investors subconsciously change the facts that led to decisions. Allocators are intelligent and thoughtful, which unfortunately only serves to strengthen their resolve. The smarter they are, the better they are at rationalizing existing beliefs, perfecting motivated reasoning, and sticking to opinions.

Investing in managers has less time pressure than a round of poker. While one might think allocators can mitigate these issues, even decisions without a timeline are biased.

Making good decisions as a team or with a committee compounds the challenge. Group dynamics can impinge independent thinking and truth discovery, two of the most important prerequisites for making good decisions.

CIOs generally work with teams, so let's turn to some steps to improve the decision-making process in a group setting.

Making better decisions

I have good news and bad news to share.

Let's start with the bad news. We are hardwired to make bad decisions even after we learn all the ways we are conditioned to muck it up. No matter how many times we read Kahneman's book, or how many times we listen to Annie on the podcast, we will still make the same mistakes. Drew Dickson of Albert Bridge Capital is keenly aware of this fact, cleverly articulated in the quote at the top of the chapter. But his knowledge of it alone won't change how he makes decisions.

Take a deep breath.

The good news is this awareness compels us to design systems that help mitigate the problem. We cannot change how our brain works, but we can think about information in a new way by striving to find an accurate representation of the truth and recognizing that we will only experience one of many outcomes that might occur. These two mental models open us to uncertainty, reduce overconfidence, and encourage contrary opinions.

CIOs can focus on two broad areas to improve their decision-making process. First, they can carefully design the structure of the team. Once this is done, they can add steps to their decision process to increase the probability of making good decisions.

1. Structure

Setting up the structure of the team lays the groundwork for a sound decision-making process. Teams can be more effective than individuals in unearthing and calibrating a range of possible outcomes in an investment decision. *Can be* is an important caveat, as a group that is constructed or managed poorly will exacerbate behavioral biases.

The size, continuity, and diversity of teams each have a substantial impact on a CIO's decision-making effectiveness.

a. Optimal size

Most CIOs oversee small teams internally. Boards can range in size. Michael Mauboussin cites research that the optimal size of a decision-making unit is

four to six people. Annie Duke adds that a team of as little as three can be highly effective.

b. Continuity

The longer a team stays together, the better a CIO can function as a decision maker. More repetitions allow a team to develop institutional knowledge and grow confidence in their decision-making process. CIOs that have the ability and wherewithal to retain key team members, like David Swensen and Scott Malpass, have been some of the most successful over time.

c. Cognitive diversity

An ideal team has cognitive diversity, benefitting from members who think differently. Cognitively diverse team members may come from different backgrounds, training, experiences, and personalities. Social diversity (race, gender, age, ethnicity) may offer a wide range of life experiences and perspectives. Interestingly, social diversity does not result in cognitive diversity all the time. The key to good decision-making is fostering a team whose members think differently, not who look different but think the same way.

CIOs walk a thin line between the salutary benefit of different thinkers and the potential challenge of team cohesion. Cognitive diversity necessarily brings together team members from different tribes, who may instinctively be biased against each other's opinions due to differing norms of conduct. Ash Fontana at Zetta Ventures tries to find people to bring onto the team that are as different as they can possibly manage culturally.

Bridging this gap between diversity and cohesiveness is a challenge. Michael Mauboussin suggests seeking individuals with a high RQ, or rationality quotient. Rational thinkers are more adept at accepting differences in order to gain from the benefits of cognitive diversity. Gathering a team of high RQ, cognitively diverse people gets the ball rolling. The way the group takes shape determines its success.

2. Conduct

Once constructed, the course of conduct significantly influences how teams perform. Sound decision processes encourage cognitive safety, independence of thought, and behavioral awareness.

a. Cognitive safety

Team members need to feel safe in their role to verbalize divergent opinions. Much of this safety comes from the leader of the team, who sets the tone through their words and actions. Ana Marshall at the William and Flora Hewlett Foundation opens her team to expressing independent views by reminding them to “argue as if you’re right, listen as if you’re wrong, and be willing to change your mind.”

The rest of the team consciously and subconsciously pays close attention. Effective leaders invite collaboration, encourage out-of-the-box thinking, and avoid blaming others for bad outcomes that follow good processes.

b. Independence of thought

The sequence of conversation in a group decision process can significantly alter the outcome. When a leader is respected or a perceived expert expresses an opinion, the rest of the team’s otherwise independent views get compromised. If a venture capital specialist weighs in on the outstanding opportunity to invest with a new manager, other members of the team are more likely to defer to the specialist’s opinion and withhold reservations they may have.

Effective leaders manage the potential to infect others with their beliefs through their orchestration of meetings. First, they allow the most junior or least informed team member to go first. Second, they steer the conversation away from more charismatic extroverts in the group, enabling introverts to have equal participation. Third, they withhold their view until others in the group have an opportunity to weigh in. Fourth, they communicate in probabilities to express uncertainty and invite input. Lastly, they use a ballot voting system to avoid social proof when it comes time to weigh in on a decision.

c. Behavioral awareness

Each team member brings biases to the table. Teams that understand each other’s individual biases are more effective at getting to the heart of issues without interference.

Kim Lew asks each of her team members to create a list of what gets in their way of objectively looking at investments. Where are they willing to take excessive risk, or why might they miss out on opportunities they should pursue? She also asks her team to come up with ideas to help them balance out

these biases. Someone who is overconfident in a certain asset class may need a devil’s advocate. Someone could be risk-averse and need encouragement in their convictions. When team members become aware of each other’s biases, they can work together to form better decisions.

Noting other common biases can help CIOs avoid decision-making roadblocks. CIOs combat the natural instinct to give more credibility within a tribe by separating the quality of an idea from the deliverer of the idea. They can monitor how team members instinctively handle bad outcomes in an effort to focus on facts and away from emotions.

3. Thought process

Once the structure and norms of conduct for the team is in place, CIOs can turn to a checklist of considerations to organize investment decisions. These include thinking probabilistically, incorporating base rates, conducting risk assessments, checking biases, and eliciting feedback.

a. Think in probabilities

Effective teams speak in probabilities. In *Thinking in Bets*, Annie Duke recommends the phrase “Wanna bet?” for occasions when someone states an opinion as a *fait accompli*. Challenging someone to a bet holds them accountable in a way that normally does not happen. They instinctively shift from System 1 to System 2 thinking and compel teammates to move from certainty to uncertainty. Bob Rubin had a rule on his famed risk arbitrage desk at Goldman Sachs that no one could speak about an investment without framing the discussion with a probability tree of expected values.²

In *How to Decide*, Annie recommends a six-step decision process for making important decisions between two options. For each possibility, create a decision matrix that includes 1) all possible reasonable outcomes, 2) the payoffs of each, and 3) the probability of each occurring. Following these six steps encourages counterfactual thinking, where decision-makers broaden their perspective beyond the most probabilistic outcome to consider a range of possibilities.

² The desk spawned a legacy of hedge fund titans, including Tom Steyer (Farallon Capital), Richard Perry (Perry Capital), Frank Brosens (Taconic Capital), Dan Och (Sculptor Management), Eric Mindich (Eton Park Capital), and Eddie Lampert (ESL Investments).

b. Base rates

Incorporating base rates, or using the “outside view,” is highly effective in reducing overconfidence.

Most of the time, we make investment decisions based on the “inside view.” We gather information about an investment opportunity, dive into due diligence, use our experience and judgement, and make an assessment about the future.

Alternately, the outside view considers a statistical sample of similar past situations. Rather than assign expected returns and risks to one opportunity, an allocator studies a large group of similar opportunities in the past to put the decision in perspective.

For example, consider an analyst forecast for Amazon stock. The analyst believes Amazon will grow revenues by 15% per year for the next ten years and creates an amazingly detailed model to describe how Amazon will get there. Amazon has grown faster than that for a long time, so anchoring to recent history makes the model appear sensible.

However, the base rate for Amazon’s future sales might look at every company since 1950 that has had \$100 billion or more in revenue, adjusted for inflation, and ask how those companies grew in the subsequent ten years. It turns out there are 313 companies in that sample set and exactly zero of them grew at 15% annually from that large a revenue base. Only seven grew more than 10% per annum, or 2% of the total in the reference class. It is certainly possible that Amazon will grow at 15% for the next ten years, but the base rate of 0% might inform the analyst that their model has a lower probability of hitting the revenue target than they initially projected from their inside view.

Base rates are helpful for allocators both in their own investment decisions and in assessing the decisions of managers. Allocators tend to put little weight on the outside view of managers, relying instead on their assessment of individuals in their analysis. Many CIOs are excited about the hedge funds in their portfolio yet are pessimistic about the average hedge fund return. Incorporating base rates into the research process can prove helpful in cooling overconfident expectations.

Additionally, allocators can consider managers’ use of base rates in their due diligence. Managers that consider both the outside view and inside view in their security selection are likely to make better decisions over time. Those unfamiliar with the concept may have a blind spot.

c. Risk assessments

When the investment case to underwrite a manager is complete, CIOs can take a step back and reconsider what might go wrong. Pre-mortem analysis, red teams, devil’s advocates, and pro-con lists all help unearth risks the CIO may not have considered when building their thesis.

1. Pre-mortem analysis

Gary Klein is a cognitive psychologist who studies how professionals become highly skilled at their craft. He created the pre-mortem analysis to help fighter pilots in the air force. A pre-mortem is a risk management tool that moves up a post-mortem analysis chronologically to uncover alternative possibilities that may not have been considered.

An example process for a pre-mortem is shown here.

a. “The plan has failed.”

The CIO kicks off the pre-mortem by telling the team they have looked in a crystal ball and seen that the plan has failed. Framing the conversation this way, as opposed to asking “What could go wrong?”, changes the mindset to see a disaster and own it. The counterfactual allows the brain to access thinking and speculation that it would not otherwise engage.

b. Take a pause

The CIO gives the team a few seconds to sit with this reality. The plan has failed.

c. Write down what happened

Each team member individually takes two minutes to write down all the reasons why the plan has failed. Two minutes affords them enough time to think, but not so much time that the energy in the room is lost. Writing down ideas encourages independent thought.

d. Share ideas one at a time, starting with the leader

The leader begins sharing ideas for why the plan has failed. In starting out, the leader purposely infects the group in taking the exercise seriously. The rest of the team will follow suit.

The leader calls on team members one at a time, requiring each team member to share one potential problem with the plan. The format encourages introverts to speak out and controls extroverts from dominating the discussion. It fosters cognitive safety and gives permission for everyone to bring up ideas.

The leader continues to orchestrate a flow around the table until the team has presented all of its ideas.

e. Write down actions to mitigate risk

Seeing everything that can go wrong, the team then takes two more minutes to write down all of the ways they individually can prevent these outcomes from occurring.

f. Share risk prevention actions

The CIO initiates a round of hearing what can be done to mitigate the risks presented.

g. End meeting

The CIO adjourns the meeting, steps back and thinks about the decision, now armed with an additional array of information. They can reassess the decision and make a new plan of action using all the material unearthed in the pre-mortem.

Pre-mortems reduce overconfidence of decisions more than any other technique, foster a culture of candor, and makes teams smarter.

2. Red teams

Andy Golden utilizes “red teams,” in which a distinct group is tasked with conducting research and poking holes in the recommendation made by the recommending “blue team.” Red teaming is more time-intensive than a pre-mortem and requires more resources.

Red teams also bring along a host of potential biases. Team leaders again need to ensure independence of thought and cognitive safety. Importantly, the red team must take ownership of their role. Red and blue team members are typically part of the same organization, a tribe of its own. The approach loses effectiveness if the red team holds back on their critique. Red teams are most

effective when the individuals come with different areas of expertise and are not personally invested in the recommendation.

3. Other risk management tools

Devil's advocate and *pro/con lists* are other tools that can be used to unearth contrary states of the world. A devil's advocate is tasked with poking holes in a thesis. Red teams generally are more effective than devil's advocates. The devil's advocate approach loses the cognitive diversity of the whole group and can get antagonistic. However, when teams are constrained by resources, devil's advocates can add a valuable step to the investment process.

Listing out positives and negatives of a decision – a pro/con list – is a reasonable starting point. This is better than not considering alternatives at all, but Gary Klein's research discovered that a list of cons is less effective than pre-mortems in tapping into the creativity of the team.

d. Bias checklist

CIOs can use a checklist at the point of decision to review common biases. Timing and overconfidence are two examples that tend to afflict investment decisions.

Ostensibly, all investment research concludes that individuals and institutions chase performance. We tend to invest in managers or buy securities after strong periods of returns and redeem or sell after periods of weakness. We can create a high-functioning team, consider base rates, conduct pre-mortems, and still make decisions that chase performance. Separating the decision about investing with a manager from the timing in doing so when permissible is a helpful trick to mitigate the bias. Getting excited about a manager need not equate to investing immediately when the enthusiasm takes hold.

CIOs and their teams are rarely more enthusiastic about a manager than the first day they put capital to work. Their team has put significant time and effort into research and is deeply vested in the new relationship. Optimism runs high and overconfidence can easily creep into their thinking. We broadly tend to overestimate a manager's edge and underestimate the future volatility of a strategy at the moment of decision. Adding a check on position sizing can help combat this bias. Overconfidence may lead to sizing a position higher than would be warranted, whereas starting off conservatively allows a

more realistic assessment of the future over time. Some investments offer the opportunity to cut losses as well, which can help prevent mistakes when we become emotionally unhinged.

e. Feedback loops

We can make very good decisions that do not work out. When I look back at my decision to bet with Warren Buffett, I still believe it followed a good process and the odds of winning were in my favor.

Alternatively, we can make terrible decisions that do work out. Annie Duke shared that her decision to quit academia to play poker came from one of the worst decision processes of her life.

Our natural tendency as individuals is to draw the wrong lessons from decisions: bad outcomes are not our fault and good outcomes come from good processes. Making better decisions requires feedback to evolve and enhance the investment process. Documenting decisions, creating a decision group, and considering the role of luck are useful tools to build feedback loops in the decision process.

1. Decision journal

Improving the investment process requires an accurate statement of the facts and beliefs known at the time the decision gets made. Hindsight bias will cloud our memory of what we knew. A decision journal ensures a CIO can learn and grow from a review of past decisions.

Part of my conviction in the quality of my decision to bet with Buffett comes from a white paper I wrote at the time of the bet. I based my assessment of the probability of success on an outside view of the wager – the S&P 500 had never before outperformed the hedge fund index over a ten-year period. However, I did not conduct a pre-mortem, which may have unearthed the possibility of markets shifting and lowered my conviction.

Annie had no such process when she moved from the academic world to poker. She looks back and thinks she barely even made a conscious decision to play poker in the first place. It was the only path that came up and she just went with it. Hence why she recalls it now as a terrible decision process.

2. Convene a decision group

A CIO can address *decision washing* by leaning on others in a decision group.

An effective decision group follows the structure and conduct principles above, ensuring open and frank communication. The group seeks to discover the truth about a decision rather than to determine who was right.

Annie uses the acronym CUDOS to describe the characteristics of a good decision group.

C: Communism – share all the data transparently, especially information that paints a view in a bad light.

U: Universalism – seek the objective truth. Pay attention to the bias to overweight or dismiss facts depending on who is communicating the message.

D: Disinterested – note our emotional conflicts of interest. Be mindful of motivated reasoning, confirmation bias, and hindsight bias.

OS: Objective Skepticism – approach the world asking why things aren't true.

3. Role of luck

As the skill of managers increased over time, luck has become more important in outcomes. The dynamic pointed out by Michael Mauboussin has permeated investing just as it did with Major League Baseball hitters over the last few decades.

Money managers tend not to see the importance of luck in their success. Among the many managers Brian Portnoy met, he found that the vast majority would not accept that luck was part of the equation when their returns were strong. Good performance demonstrated skill, bad performance was bad luck: a textbook example of self-serving bias.

Allocators need to be patient and gather lots of information to distinguish between a manager's luck and skill. Tim Recker at the James Irvine Foundation reminds his team that business cycles can last seven to eight years, and some managers need a full business cycle to show that their strategy works.

Summary

Making good decisions is hard. We are hardwired to do this poorly, so building processes to mitigate our pernicious biases can help make our decisions *less bad*.

Investment ideas are not black and white. Known and unknown risks are a part of every investment, and it is rare to have an honest consensus on a team. When functioning at a high level, teams unearth many scenarios and assign probabilities to potential outcomes that provide the full mosaic of information for a CIO to decide.

When a CIO has made the decision to hire a manager, they may find they are a price-taker, needing to accept the terms offered by a manager in high demand from allocators. Other times, the CIO may be in position to negotiate terms. The next chapter discusses how to make the most out of those opportunities.

To learn more

Podcasts

Capital Allocators: Annie Duke – Improving Decision-Making (Ep.39)

Capital Allocators: Annie Duke – How to Decide (Ep.156)

Capital Allocators: Gary Klein with Paul Johnson and Paul Sonkin – Conducting Pre-Mortem Analysis (Ep.109)

Capital Allocators: Michael Mauboussin – Active Challenges, Rational Decisions and Team Dynamics (Ep.36)

Books

Thinking, Fast and Slow, Daniel Kahneman

Thinking in Bets, Annie Duke

How to Decide, Annie Duke