

May 2014

NeuroLeadership Journal

Volume 5

Breaking Bias

Matthew D. Lieberman, David Rock and Christine L. Cox

www.neuroleadership.org



The NeuroLeadership Journal is for non-commercial research and education use only. Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third-party websites are prohibited.

In most cases authors are permitted to post a version of the article to their personal website or institutional repository. Authors requiring further information regarding the NeuroLeadership Journal's archiving and management policies are encouraged to send inquiries to: journal@neuroleadership.org

The views, opinions, conjectures, and conclusions provided by the authors of the articles in the NeuroLeadership Journal may not express the positions taken by the NeuroLeadership Journal, the NeuroLeadership Institute, the Institute's Board of Advisors, or the various constituencies with which the Institute works or otherwise affiliates or cooperates. It is a basic tenant of both the NeuroLeadership Institute and the NeuroLeadership Journal to encourage and stimulate creative thought and discourse in the emerging field of NeuroLeadership.

NeuroLeadership Journal (ISSN 2203-613X) Issue Five published in May 2014. We encourage readers to propose a paper for the next edition of this Journal. We recommend sending in a two-page proposal before submitting a finished paper and welcome pure science, as well as case studies and discussion pieces. For further information as to how to submit a paper for the next Journal go to www.neuroleadership.org

NeuroLeadershipJOURNAL

Cognitive biases, the unconscious drivers that influence our judgment and decision-making, are a pervasive issue in organizations. Individuals are notoriously poor at recognizing and controlling their own biases – in fact, our brains are wired to promote fast, efficient information processing. This can be highly adaptive, yet can also be detrimental in certain circumstances, such as making unbiased decisions about who to hire or promote, where to allocate resources, and what information to prioritize in decision-making. Here, we present a framework that lays out the three steps we believe organizations should follow to “break bias”:

1. **ACCEPT** that we are biased,
2. **LABEL** biases using four main categories detailed in our easy-to-remember, comprehensive **COST™** model (i.e., corner-cutting, objectivism, self-protection, and time and money biases), and
3. **MITIGATE** bias according to specific strategies targeting the neurobiological mechanisms associated with each category of bias. Specific examples illustrate each of the four major categories of bias according to the **COST™** model, highlighting the influence of cognitive bias in organizations and strategies we believe will be useful in combating their negative effects.

Breaking Bias

Matthew D. Lieberman, David Rock and Christine L. Cox

A mid-level manager in a financial services firm is trying to hire a new employee. While reviewing resumes, he unconsciously prefers candidates of a similar age and background to his team. The manager tells himself he is trying to build a cohesive team, unaware that he is biased or that this team will make worse decisions as a result.

A senior executive of a real estate firm once voted against investing in a significant new development project. Conditions have changed, and the project would now be an ideal fit for the business. When the project is presented again, she easily recalls the older data that led her to veto the idea, even though newer data, with which she is much less familiar, would suggest that it is now a good investment. She has no idea that she is biased or that a big business opportunity has been lost.

A sales representative in an industrial firm spends most of his time calling on clients in his home city, because he feels he knows the area best, even though there are significantly bigger clients in other cities in his territory. He has no idea that he is being biased and is costing himself and his firm significant revenues.

These are examples of common everyday biases. Biases are unconscious drivers that influence how we see the world. Biases are the invisible air we walk through – exerting their influence outside of conscious awareness, adaptive mechanisms evolved to help us make quick, efficient judgments and decisions with minimal cognitive effort. Thus, biases can impact every decision we make. We cannot go shopping, turn on a computer, or start a conversation without the potential for bias.

On the one hand, biases are helpful and adaptive. Biases help us use previous knowledge to inform new decisions, a kind of cognitive shorthand, as we do not have the cognitive resources to make every decision fresh.

However, many of our biases can also be unhelpful. Biases can blind us to new information, or inhibit us from considering a broad range of options when making an important decision and choosing our path carefully.

Biases are the invisible air we walk through – exerting their influence outside of conscious awareness...

Writ large, unhelpful biases were at the heart of the 2007 global financial crisis (and dozens of similar crises over the centuries) – i.e., ignoring evidence that current practices were going to have devastating long-term effects (known as the “confirmation bias”), and sacrificing long-term future outcomes for more immediate gains (known as “temporal discounting”).

They were at the core of why Japan suffered so much from their 2011 tsunami, and New York City from Hurricane Sandy in 2012 – i.e., overestimating the degree to which individuals, the government, etc., would be able to control the negative effects of these natural disasters (known as the “illusion of control”), and underestimating the time and effort it would take to prepare (known as the “planning fallacy”). And they are at the core of the dysfunction of many great companies and institutions, including the challenges faced by governments the world over.

In a hyper-connected world where poor decisions can multiply like a chain reaction, breaking free of unhelpful bias has never been more urgent or important...

In a hyper-connected world where poor decisions can multiply like a chain reaction, breaking free of unhelpful bias has never been more urgent or important – for individuals, teams, schools and institutions, organizations, and for whole societies. Various other pressures of the current world in which we live and work highlight the urgency of mitigating bias. Some examples include:

1. an increase in the complexity, ambiguity, and volatility of the problems we are facing, problems in which our default solutions are unlikely to work;
2. problems requiring slow thinking, cognitive effort, and the ability to approach potential solutions from multiple perspectives;
3. the parallel increase in the need for diversity of thought as our reliance on technology and social networks increases; and
4. reduced opportunities to slow down and engage cognitive effort in bias mitigation as demand for speedy decision-making increases.

In our large organizations, increasingly large sums of money and resources are spent educating people about biases. For example, U.S. companies spend an estimated \$200–300 million a year on diversity programs (Flynn, 1998; Vedantam, 2008). This spending is in the form of ‘diversity’ or ‘sensitivity’ training, where executives, managers, and all employees are told to watch out for biases, in particular around hiring and promotion decisions. These programs tend to be more narrowly focused on people-related bias in decision-making (e.g., the unconscious influence of negative evaluations of others who are dissimilar from oneself and positive evaluations of others who are similar to oneself [Ross, 2008; Lublin, 2014]). One motivation for the development of our model of “Breaking Bias” that we will present here is to expand the scope of bias awareness and mitigation strategies to include a wide range of biases, not only those related to “people decisions.”

While many executives are beginning to recognize that there is a real bias problem, and very specific, case-directed training has shown some promise in reducing bias (MacLean et al., 2013), there is little evidence that just educating people or raising awareness about bias currently does much to reduce those biases (e.g., Pronin et al., 2002; Kahneman et al., 2011). A recent global survey found that though diversity training improves the perceived value of having an inclusive work environment, less than 50% of surveyed employees report that their work environment is actually inclusive or values diverse perspectives (CEB, 2012).

...there is little evidence that just educating people or raising awareness about bias currently does much to reduce those biases...

One study, which explored the efficacy of diversity training over a 30-year period, found that efforts to improve diversity that focused specifically on training, educating, or providing managers with feedback around their biases were the least effective methods for improving levels of diversity (Kalev et al., 2006).

In some cases, this type of training actually reduced diversity rather than improved it. This is surprising, as it seems logical that bias education and feedback would improve an individual's ability to recognize when he or she is being biased and to use intervening strategies.

Why, then, are bias training programs so ineffective? Partly, this is because biases occur outside of conscious awareness. We literally are unaware of the fact that we are being biased at any moment. Not only does educating people about biases do little; there is a bigger challenge here: Currently, there is no theory in practical use for bias mitigation. To illustrate this point, a search for "bias mitigation" on Wikipedia states, "There is no coherent, comprehensive theory or practice of cognitive bias mitigation" (we note that Wikipedia is not an accepted academic source, but our intention is to highlight the information that is available to those searching online for information regarding "bias mitigation"). While there are commercial initiatives that offer cognitive bias modification services, "there is no evidence that this service is backed by peer-reviewed research results" (http://en.wikipedia.org/wiki/Cognitive_bias_mitigation), although we acknowledge the existence of many research-based education and training programs aimed at reducing bias.

We literally are unaware of the fact that we are being biased at any moment.

Could it be that billions of training dollars and countless employee hours are being wasted trying to educate employees to do something that just cannot be done with our current approaches? Is there a way to conceptualize and implement bias mitigation strategies that would yield a higher payoff?

The discussion above is based on what we know about bias in the *individual* – people are notoriously bad at knowing that their thoughts, beliefs, interactions, judgments, and decisions are affected by unconscious drivers. Recently, attention has shifted to bias mitigation strategies at the systems, or *organizational* level. This is reflected in research exploring the idea of "group intelligence," where groups of people make better decisions as a whole than each individual that comprises the group (Wooley et al., 2010).

The importance of implementing strategies for change on the organizational level is consistent with the work of Peter Senge on organizational learning and systems thinking – team learning, appreciating and valuing the perspectives of others, and leveraging the skills of a group of diverse individuals will improve the organization (1990). There is evidence that the efficacy of diversity training programs improves when organizations establish structures of responsibility (e.g., committees, staff positions, departments [Kalev et al., 2006]).

...groups of people make better decisions as a whole than each individual that comprises the group...

The idea we would like to propose is that organizations or teams may be able to become self-aware of bias in ways that individuals cannot, and that strategies can be put into place that will facilitate this awareness, foster an organizational culture that assumes and accepts bias inherent in human decision-making, and thus mitigate the wide-ranging effects of bias in this context.

In his recent book *The Righteous Mind: Why Good People Are Divided by Politics and Religion*, Jonathan Haidt (2012) sums up this idea nicely:

"... if you put individuals together in the right way, such that some individuals can use their reasoning powers to disconfirm the claims of others, and all individuals feel some common bond or shared fate that allows them to interact civilly, you can create a group that ends up producing good reasoning as an emergent property of the social system" (pg. 105).

Strategies that have been suggested for bias mitigation in organizations have thus far either been specific and limited in scope (e.g., Kahneman et al., 2011), or helpful though expansive, calling for radical organizational change (e.g., Lovallo & Sibony, 2010). This is likely due to the overwhelming number of cognitive biases that can currently be found with any Google search (~150) and the difficulty of putting these biases into a useful framework so that they can be easily remembered.

Short of going through a checklist of 150 possible biases that could influence major decisions, what is one to do, given that these biases exert unconscious influences and are so difficult to detect?

In this paper, we propose a model that provides an easy-to-remember framework for addressing bias at an organizational level. Our goal is to help systems – organizations, teams, or processes – to address bias in a whole new way that does not rely on individuals having to catch themselves being biased. The model involves three steps:

1. **Accept** that people and systems are deeply biased and do not know it;
2. **Label** the biases likely to occur in any given system or decision, based on the four major categories into which they fall. Our model condenses the ~150 biases into four overarching bins based on the common underlying biology driving a particular bias. We call this the **COST™** model of bias;
3. **Mitigate** bias by attacking bias with strategies that go directly to the core mechanisms underpinning that bias.

As we go into detail on each of these three steps, we will outline our **COST™** model of cognitive bias and provide real-world examples to illustrate how these major types of bias can be mitigated in an organizational setting. To begin, we must acknowledge and **accept** that we are biased in the first place.

Most of us recognize that we can fall prey to bias, but we almost never think we are biased in the current moment.

Step 1: Accept

People do not want to believe that they are biased. We all are quick to detect and criticize biased thinking and decision-making in others, but believe that we ourselves are far less susceptible to these same biases (a phenomenon that has been termed the “bias blind spot” [Pronin et al., 2002]).

In fact, even high cognitive ability does not protect someone from the effects of bias (West et al., 2012). Most of us recognize that we can fall prey to bias, but we almost never think we are biased in the current moment. As the oft-quoted saying goes, “I always think I’m right, but I don’t think I’m always right.” However, it does little good to recognize that somewhere in the last thousand decisions made there must have been bias if this does not help us to recognize it the next time it occurs, before it has influenced outcomes.

Why is it so difficult for people to accept that their beliefs, decisions, and actions can be influenced by the unconscious drivers of cognitive bias?

This resistance to evidence of our own susceptibility, paired with the often-unconscious nature of cognitive bias, creates a perfect storm in which bias is perpetuated and rarely adequately recognized or managed. The insidious nature of cognitive bias, and its effect on human judgment and decision-making, has led psychology researchers to propose that research on and efforts to educate the public against such bias should be a top priority in the field of psychology (Lilienfeld et al., 2009). More recently, prominent social psychology researchers highlighted the problem of persistent implicit social (e.g., gender and racial) biases in academia and recommended empirically validated diversity training interventions implemented at the organizational level to reduce the impact of bias (Moss-Racusin et al., 2014).

Why is it so difficult for people to accept that their beliefs, decisions, and actions can be influenced by the unconscious drivers of cognitive bias? A big part of the answer involves the unconscious nature of biases themselves, as indicated above. However, there is another, more potentially insidious and problematic answer to this question. In short, *it feels good to be right*.

For example, consider the following problem:

You have \$1.10, a bat, and a ball. The bat costs \$1.00 more than the ball. How much does the ball cost?

Likely, you were able to arrive at a solution fairly quickly, one that felt obvious and at least satisfying to you (i.e., the bat costs \$1.00 and the ball costs \$0.10). Arriving at a correct answer is associated with contentment and certainty. Being right is rewarding and activates the brain's reward circuitry. Even if people are completing a relatively uninteresting task for no money or other incentives, just a feeling that they are doing the task correctly leads to activation in the ventral striatum, a brain region consistently implicated in processing reward (Satterthwaite et al., 2012). This positive emotion – the enjoyment we experience from being right – is one of the main reasons that we are motivated to overlook our own biases and their contribution to the errors we make.

...the enjoyment we experience from being right – is one of the main reasons that we are motivated to overlook our own biases and their contribution to the errors we make.

Not only does it feel good to be right; *it feels bad to be wrong*. Making errors and mistakes is painful and distressing and activates brain regions associated with processing pain and negative emotion. In a task in which participants had to learn to classify shapes as belonging to one of two categories, making a mistake (misclassifying a shape) was associated with activation in the dorsal anterior cingulate and the anterior insula, brain regions that are part of the “pain matrix” (Daniel & Pollmann, 2012), even when there were no material consequences to being wrong. In addition, we often feel angry and frustrated when making errors. The amount of frustration and negative emotion a person feels after making an error is positively associated with activation in the dorsal anterior cingulate (Spunt et al., 2012).

These two related principles – that being right is rewarding, and being wrong is painful – are central to understanding how our judgments and decisions are so susceptible to unconscious cognitive bias and why it is so difficult to overcome the influence of bias. We are motivated to seek out reward, and we are motivated to avoid pain. The positive emotion and rewarding feeling of being right does not just occur when we are objectively right; it also occurs when we *believe* we are right, when we have a *feeling* of being right, regardless of the objective reality of whether or not we are actually right.

The reinforcement we get from believing that we are right (that we have answered a question correctly, that we have made the right decision, etc.) further motivates us to seek out situations in which we feel that we are right. Further, and perhaps more importantly, it motivates us *not* to seek out information suggesting we might be wrong, and even to ignore disconfirming information that is right in front of us.

These principles – seeking pleasure and avoiding pain – are two of the most important contributors to cognitive bias: **We make judgments and decisions based on what feels right, even though what feels right may be based on information that is irrelevant, faulty, or just plain wrong.**

Individuals in particular have a very difficult time with and may never accept that they are biased, but at the organizational level, a system of individuals may be able to (and need to) operate under the assumption that judgments and decisions are influenced by cognitive bias, and that they can put a structure in place to actively mitigate the effects of bias.

We make judgments and decisions based on what feels right, even though what feels right may be based on information that is irrelevant, faulty, or just plain wrong.

What would it look like for a firm to 'Accept' bias, at an organizational level? Here are several key ideas:

- Recognize that intelligence does not make people less biased. The brilliant researcher or genius engineer is just as likely to be biased as a mid-level employee.
- Recognize that experience or expertise does not necessarily make people less biased. Seasoned executives can be just as biased as new recruits, perhaps more so. Expertise may change the kind of bias present without eliminating it. Systems need to be put in place at all levels of an organization to mitigate bias, all the way from the shop floor to the boardroom.
- Recognize that educating people is not enough. Processes need to be built into organizational systems that mitigate bias.

We need to accept that intelligence, expertise, and education simply do not reduce bias in a meaningful way.

Above all, the goal of the 'Accept' stage is to educate executives and all employees that biases are a fact of life, and that it is normal to not know you are biased. We need to accept that intelligence, expertise, and education simply do not reduce bias in a meaningful way. With this acceptance, we can get to work on setting up systems that reduce bias at a systemic level. This brings us to the second step in this process, after accepting that our systems are biased, which is to **label** bias in a useful and effective manner.

Step 2: Label

To be effective at combating bias, we need an easy-to-remember framework that will allow executives, managers, team leaders, etc., to quickly identify the major *types* of biases that all too often affect major business decisions. To condense the overwhelming number of individual potential biases that have been described, we have developed the **COST™ Model of Bias**. This model was developed over a year of research by the authors by beginning to identify the core neurobiological correlates associated with the key biases and, through trial and error, organizing a framework that separated the biases into categories.

Several other models were proposed over the year, and, through a process of in-depth assessment of various theoretical orientations to cognitive bias and feedback from colleagues with expertise in this field of research, multiple iterations of the model were further developed and refined.

We now believe that biases can be divided into four main types:

1. **C**orner-Cutting
2. **O**bjectivism
3. **S**elf-Protection
4. **T**ime and Money

Each type is described in detail below, and extensive examples of each type of bias are provided in Appendix A. We do not mean to suggest that every bias fits into only one category or that every bias is accounted for by this model. But, we think there is value in simplifying the vast majority of errors and biases into these groupings. Each category has defining features as well as category-specific mitigation strategies that can be applied.

The COST™ Model of Bias

COST: Corner-Cutting

Corner-cutting biases can be described as mental shortcuts that help us make quick and efficient decisions. The downside to this efficiency is that those decisions may be based on incorrect judgments. Typically, when corner-cutting biases occur, the brain is using a fast, intuitive system and makes decisions based on what information is easily accessible and feels right. This has been labeled the brain's "System 1": The system that relies on fast, easy associations and intuition (Kahneman, 2011; Satpute & Lieberman, 2006).

Corner-cutting biases can be described as mental shortcuts that help us make quick and efficient decisions.

However, we often need to make decisions based on more objective information, which is often not so easily accessible and takes more mental effort to access and to use when making judgments.

The brain's "System 2" is the slower, more effortful overseer of the fast, more intuitive System 1 (Kahneman, 2011). System 2 is sometimes called the "lazy fact checker" since it can be called upon to correct System 1's mistakes, but it often is not since it requires more cognitive effort to engage.

Remember the "bat and ball" problem that was posed above?

You have \$1.10, a bat, and a ball. The bat costs \$1.00 more than the ball. How much does the ball cost?

Most people will answer quickly and confidently that the bat costs \$1.00 and the ball costs \$0.10. It is the fast, instinctive answer that comes to mind, and it makes intuitive sense and feels right (System 1). But it is wrong. If the bat costs \$1.00 more than the ball, then the ball must cost \$0.05, and the bat must cost \$1.05. However, arriving at this answer requires most people to engage their System 2 – you must do some mental algebra to come to the right answer. Engaging System 2, fact-checking and correcting System 1's mistake, is harder work, so if System 1's answer comes more easily and feels right, why think more about it? Most people do not. In fact, around 50% of students at Harvard, Princeton, and MIT also tend to get this problem wrong (Frederick, 2005; Kahneman, 2003).

We are being guided by an intuitive sense of right and wrong, rather than performing a logical analysis.

The following syllogisms provide another example of a corner-cutting bias at work.

If the premises are true, does the conclusion logically follow?

1. *Premise: If it rains, the game will be cancelled.*

Premise: It did not rain.

Conclusion: The game was cancelled.

2. *Premise: No addictive things are inexpensive.*

Premise: Some cigarettes are inexpensive.

Conclusion: Some cigarettes are not addictive.

Most people will (correctly) reject the conclusion from #1, but many people will also reject the conclusion from #2, even though they should technically accept it, because it does follow logically from the two premises (Evans et al., 2001). This is an example of the "belief bias" – when our belief that the conclusion is not true gets in the way of judging the logic of the syllogism.

We know and believe that cigarettes are addictive, so accepting the conclusion that says some cigarettes are not addictive is difficult, even though the truth of the premise has no relevance to the validity of the syllogism. When we reject the conclusion of #2, System 2 is not kicking in to correct the mistake of System 1. We are being guided by an intuitive sense of right and wrong, rather than performing a logical analysis. This is further evidenced by the fact that people who are required to respond within a 10-second time limit (i.e., have less time to engage their System 2) are more likely to make the mistake of rejecting the conclusion of #2 than those who are given more time to deliberate (Evans & Curtis-Holmes, 2005). This finding points to the importance of taking enough deliberative time to make a thoughtful and well-informed decision – something that can be very difficult in a workplace culture of urgency that puts a premium on a fast turnaround, expecting decisions to be made and answers to be given very quickly. The problem with impatience and urgency, as we illustrate with these examples, is the increased influence of cognitive bias and the likely sacrifice of the quality of a decision in favor of quantity.

A classic example of a corner-cutting bias is the availability bias, or our tendency to make a decision based on the information that's most readily accessible (i.e., the information that comes to mind most quickly) instead of on objective information (Tversky & Kahneman, 1981). The availability bias is a difficult problem for organizations because it prevents the consideration of all potentially relevant information, impeding objective and perhaps more adaptive decision-making. Remember the senior executive of a real estate firm mentioned in the opening paragraph of this paper?

A senior executive of a real estate firm once voted against investing in a significant new development project. Conditions have changed, and the project would now be an ideal fit for the business. When the project is presented again, she easily recalls the older data that led her to veto the idea, even though newer data, with which she is much less familiar, would suggest that it is now a good investment. She has no idea that she is biased or that a big business opportunity has been lost.

The executive in this scenario was subject to an availability bias and lost a business opportunity as a result.

In summary, corner-cutting biases can be detrimental to decision-making in organizations. If we make judgments based on our quick intuitions about what is right or what we *want* to be right instead of taking more time to deliberate, gather relevant information, question our initial assumptions, and make objective decisions, then we are likely to let irrelevant, incomplete, or flat-out *wrong* information guide our choices.

See Appendix A for a detailed list of corner-cutting biases.

COST: Objectivism

Objectivism biases are a result of our brains being built to experience the world as a direct and objective representation of what is really out there in the world. It is as if we have an implicit belief that our perceptions, beliefs, understanding, and experiences are objectively true. This assumption that our experience corresponds to reality is referred to as “naive realism” (Griffin & Ross, 1991). The problem with this implicit belief is that it overlooks the varying array of behind-the-scenes processes by which our experience of reality is constructed. Our expectations, past experiences, personality, and emotional state are just a handful of the factors that color our construal of what is happening “out there” in the world.

Objectivism biases are a result of our brains being built to experience the world as a direct and objective representation of what is really out there in the world.

There are two main reasons that objectivism biases are especially pernicious. First, they happen outside of conscious awareness, so it is nearly impossible to monitor for them. Second, because we hold a strong conviction that we are seeing reality as it is, we tend to believe that anyone else who sees things differently must either see things incorrectly or pretend to see them incorrectly for some other reason. If two people have different expectations, and thus experience two different “objective” realities, then each person is likely to think the other must be crazy, mean, stupid, biased, or lazy (Lieberman, 2013).

Thus, it is very difficult to convince someone who has an objectivism bias that, in fact, he or she might be the one who is mistaken. These biases are similar to visual illusions – even if you logically know that it is an illusion (i.e., two lines are the same length even though they *really* look like they are different lengths), in that it is practically impossible to change your experience of it.

It is very difficult to convince ourselves that our intuitive experience is incorrect, even when confronted with strong evidence to the contrary.

We have already introduced a key example of an objectivism bias, the “bias blind spot,” which describes the fact that it is relatively easy to identify biases in others but not in oneself. People rate themselves as less susceptible to biases than others and see their answers as less biased than the answers of others, even when given information about how biases could (and most likely do) affect them (Pronin et al., 2002). It appears that drawing individuals’ attention to this bias is not enough to make them aware of their own biases or to mitigate their effects.

Another objectivism bias is the “false consensus effect,” or overestimating the extent to which others agree with you or think the same way you do. For example, if you prefer vanilla to chocolate ice cream, you are likely to think that more people in the general population have the same preference (e.g., 75% of people prefer vanilla); someone who prefers chocolate to vanilla ice cream, however, will also think that 75% of the general population agrees with him and prefers chocolate. In an organizational setting, this assumption can lead to a variety of problems, especially if leaders assume that others agree with their preferences (e.g., for a certain development strategy) and make decisions without asking others’ opinions or seeking input regarding potentially superior alternatives.

...it is relatively easy to identify biases in others but not in oneself.

Objectivism biases may be the most difficult for individuals to accept, label, and mitigate, but are prime targets for an organizational systems approach. If leaders assume that objectivism biases will be present and are highly likely to affect decision-making, then strategies can be developed and checks put into place that will minimize their influence.

See Appendix A for a detailed list of objectivism biases.

COST: Self-Protection

People are highly motivated to feel good about themselves, to focus on information that allows them to see themselves in the best possible light, and to maintain that positive self-image. Sometimes these self-interested and self-sustaining motives can be in conflict with an objective perception of ourselves, others, and the world.

The “overconfidence effect” is one type of self-protection bias. We believe that we are better than average in a whole host of different ways (e.g., driving, attractiveness, leadership, teaching ability), and this tendency toward overconfidence may even have conferred an evolutionary advantage (van Veelen & Nowak, 2011; Johnson & Fowler, 2011). In situations where an honest and objective assessment of our abilities, strengths, and weaknesses is necessary, such as a leadership role in an organization or as a team member, the overconfidence effect can be detrimental. Difficulty with being realistic about our abilities and/or lack thereof, the accuracy of our judgments, and receiving criticism is likely to lead to difficulties in the workplace where honest assessment and feedback are necessary.

We believe that we are better than average in a whole host of different ways...

The “in-group bias” and the “out-group bias” are two self-protection biases linked to promoting and protecting one’s own group (e.g., your family, your team, your company), but are also associated with the development and perpetuation of stereotypes and prejudice. The in-group bias refers to the more *positive* perception of people who are more *similar* to you compared to those who are less similar to you. The out-group bias refers to the more *negative* perception of people who are more *different* than you compared to those who are more similar to you. These biases are reflected not only in the perception of in-group and out-group members, but in one’s behavior toward them – e.g., more resources are allocated to in-group (vs. out-group) members. As such, if left unchecked and unaddressed, these can be particularly harmful in organizations. Remember the mid-level manager in a financial services firm mentioned in the opening paragraph of this paper?

A mid-level manager in a financial services firm is trying to hire a new employee. While reviewing resumes, he unconsciously prefers candidates of a similar age and background to his team. The manager tells himself he is trying to build a coherent team, unaware that he is biased, or that this team will make worse decisions as a result.

By only hiring people similar to his existing team, this manager was exhibiting an in-group bias.

Social neuroscience research has shown that we perceive and relate with in-group and out-group members very differently (Banaji & Greenwald, 2013). In fact, the “Relatedness” component of the SCARF model (Rock, 2008; Rock & Cox, 2012) deals with this topic in detail. Research has shown that merely assigning people to arbitrary teams creates affinity for their own team members, relative dislike of members of the other team, and greater activity in several brain regions (e.g., amygdala, orbitofrontal cortex, striatum) for in-group vs. out-group faces. This effect of team membership was seen regardless of other group differences like race (Van Bavel et al., 2008). People like members of their in-group more than out-group members, and they are also more empathic toward members of their own group. Adams et al. (2009) asked Japanese and American participants to perform the “mind in the eyes” test, in which they had to choose the correct emotion expressed by seeing only images of different individuals’ eyes; crucially, they showed images of both Japanese and American eyes. Japanese participants were much better at correctly identifying the emotions expressed in images of Japanese eyes, and Americans were better for images of American eyes. Not only were participants more accurate in judging the correct emotion for their own culture (their in-group), but a region of the brain important for making these social perception judgments (the superior temporal sulcus) was significantly more active when participants saw images of their own vs. the other culture. There are significant behavioral and neural differences associated with processing information about in-group and out-group members which can impact the way in which we interact with and interpret the people around us.

People like members of their in-group more than out-group members...

Increasing awareness of self-protection biases, and in-group/out-group biases in particular, and implementing strategies to foster unbiased hiring strategies; team assignment; intergroup interaction across race, gender, age and other factors; and enhancing communication, conflict resolution, and perspective taking, are all critical to mitigating the detrimental effects of the influence of these largely unconscious biases.

See Appendix A for a detailed list of self-protection biases.

COST: Time and Money

Humans are very strongly driven by two asymmetries:

1. We are more driven by *negative* than by positive, and
2. We value things more that are in greater *proximity* to us in terms of ownership, spatially and temporally, than those that are farther away.

The fact that negative information tends to be more salient and motivating than positive information is evolutionarily adaptive – you will stay alive longer if you remember more quickly that the snake will kill you than that the bunny is cute. Put another way, losing \$20 feels worse than finding \$20 feels good. This principle, termed the “negativity bias” (e.g., Skowronski & Carlston, 1989), is also manifest in several time and money biases. “Loss aversion” and the “framing effect” both refer to the fact that humans are highly sensitive to information about whether we expect to lose something or gain something, and that that information changes our decisions.

People will choose to avoid a risky decision if the outcome is positive (i.e., if you expect to win money), but will seek more risks in order to avoid a negative outcome (i.e., if you expect to lose money). In both cases, people are exhibiting loss aversion (Kahneman & Tversky, 1984), and the negative information is what is salient – minimize the risk of not winning and increase the chances of avoiding *losing*.

The fact that negative information tends to be more salient and motivating than positive information is evolutionarily adaptive...

Similarly, with the framing effect (Tversky & Kahneman, 1981), if information is presented, or framed, as a *gain*, people choose to avoid a risky decision (e.g., do not take the risky bet for a 60% probability of winning \$20). However, if information is framed as a *loss*, then people choose to take the risk to avoid the loss (e.g., take the risky bet for a 40% probability of losing \$20). This is true even though the objective information is the same in both cases (i.e., 60% chance of winning, 40% chance of losing).

Proximity is also a salient driver of decision-making. It appears there is one network in the brain for all types of proximity – the proximity of owning vs. not owning an object, as well as proximity in space and in time (Tamir & Mitchell, 2011). Examples of this include the “endowment effect” and “temporal discounting.” We value things more if we own them than if we do not. For example, someone may say that she is willing to pay \$1 for a neutral object, for example, a bottle of water. However, if you give her a bottle of water (i.e., endow her with it), and ask how much she would be willing to accept as payment for this bottle of water that she now owns, she may say \$2. This phenomenon is known as the endowment effect (Kahneman et al., 1990) – we expect others to pay more for something we own than we would be willing to pay for it independently.

We value things more if we own them than if we do not.

Remember the sales representative in an industrial firm mentioned in the opening paragraph of this paper?

A sales representative in an industrial firm spends most of his time calling on clients in his home city, because he feels he knows the area best, even though there are significantly bigger clients in other cities in his territory. He has no idea that he is being biased and is costing himself and his firm significant revenues.

By not focusing on more valuable clients in others cities, he was subject to a proximity bias.

Temporal discounting (Kirby & Marakovic, 1995) is a similar phenomenon, but concerns proximity in the temporal as opposed to the physical or spatial domain. People tend to value things differently depending on whether they get them now vs. later. For instance, given a choice between \$10 right now and \$20 paid out in a month, most people will choose the \$10, even though no reliable investment strategy will make the \$10 worth more than \$20 in such a short period of time. In other words, the \$20 is rationally worth more, but we devalue or discount this future worth because it is off in the distance and less tangible than the money we can receive right now. In our evolutionary past, survival may have benefited more from focusing on current needs, but in the modern world, this overdependence on immediate outcomes is often less beneficial in the long-term.

...many types of decisions made in the modern world, and especially in the organizational world, largely depend on objective and rational judgments based on data instead of gut feeling.

An awareness of these time and money biases are crucial to effective decision-making in organizations. Evolutionarily, these biases may have been adaptive, and they may still feel intuitive since our brains developed in a context in which they served an important purpose.

Though some research suggests that intuitive (vs. analytical) decisions are very effective when an individual has a high level of expertise in a specific domain (Dane et al., 2012), many types of decisions made in the modern world, and especially in the organizational world, largely depend on objective and rational judgments based on data instead of gut feeling.

See Appendix A for a detailed list of time and money biases.

We have outlined the importance of accepting that many of our judgments and decisions are subject to unconscious cognitive biases, and we have provided a framework, the **COST™** model of bias, that characterizes the four major types of bias (see Table 1 for an overview). The final step is to delineate strategies that individuals and organizations can use to **mitigate** the negative consequences of bias in judgment and decision-making.

Step 3: Mitigate

For each major type of bias outlined in the **COST™** model, we will now present example scenarios in which these biases are evident in an organizational setting and provide mitigation strategies designed to address the effects of these biases by targeting the root cause of each type of bias. These strategies are practical ways of helping people activate their brain's "braking system" (Lieberman, 2009) and inhibit biased responses.

Bias	Definition	Examples	Mitigation
Corner-Cutting	Mental shortcuts that help us make quick and efficient decisions	Availability Bias – making a decision based on the information that comes to mind most quickly	Increase motivation to engage cognitive effort Consider all information and alternatives, not just what's most readily accessible
Objectivism	The implicit belief that our perceptions, beliefs, understanding, and experiences are objectively true	Bias Blind Spot – identifying biases in others but not in ourselves	Request objective, outside opinions from others Try to see the information from the perspective of an outside observer
Self-Protection	We are motivated to feel good about ourselves and the group(s) we belong to	Ingroup/Outgroup Biases – perceiving and behaving more favorably toward others who are more similar to you (vs. more dissimilar)	Reduce self-defensiveness through self-affirmation Reduce the threat of outside others through thinking of shared goals and values
Time & Money	Negative information is more salient than positive; we value things more that are in greater proximity to us	Temporal Discounting – preferring a smaller reward now to a larger reward in the future	Make decisions less attached to the self Imagine that you are making the decision for someone else or that it has already been made in the past

Table 1

Mitigating Corner-Cutting Biases

Corner-cutting biases might occur in everyday decisions that involve complex calculations, analysis, evaluation, or identifying conclusions out of data, for example, building a spreadsheet to analyze a project, working out the cause of a problem with a machine, or identifying the right solution for a client's needs. Let's take this last example: Identifying the right solution for a client's needs. If a sales rep is busy and tends to prefer one solution more than others, he might suffer from an availability bias when he hears a client's problem and automatically thinks of his favorite solution, instead of really listening to the client and considering other potential solutions that might better fit the client's needs. Or, think of a doctor who has recently seen many patients with a particular virus; she might automatically assume a new patient with similar symptoms also had the same virus without more carefully analyzing the details of his condition.

If people do think more deeply, they may be able to make better decisions.

Corner-cutting biases will be especially likely when people are in a hurry or are cognitively depleted, something very common in many organizations. The key issue with corner-cutting biases is that people take the easy path. There is no incentive to think more deeply and search for a wider possible set of solutions. If people do think more deeply, they may be able to make better decisions. In this instance, the goal is to create some kind of incentive for individuals to identify their own mistakes. In other words, we suggest a mitigation strategy of increasing the motivation to engage our "System 2," activating the brain's braking system (i.e., ventrolateral prefrontal cortex [VLPFC] [Lieberman, 2009]), resulting in more deliberative and thoughtful decision-making, and inhibiting our quick, reflexive "System 1" – our brain's tendency to engage easy, less-cognitively effortful habits and default responses.

In the example above, the sales rep is experiencing unconscious corner-cutting biases, technically known as availability bias (making a decision based on the information that comes to mind most quickly instead of on objective information) and anchoring bias (relying too heavily on the first piece of information offered when making a decision).

A process that might work here would be for the sales rep to lay out the logic of his decision step-by-step and be encouraged to find any potential flaw in his logic, with his manager providing clear motivation and positive reinforcement that finding flaws in one's thinking is a sign of strength.

Other strategies that may work to mitigate corner-cutting biases include developing step-by-step approaches that encourage breaking a problem into its component parts. It may also help to involve other people and get outside opinions as part of the typical decision process, as well as implementing a mandatory "cooling off" period (e.g. 10 minutes of relaxation exercises or a walk outdoors) before making decisions likely to be impacted by corner-cutting biases. Using a human-centered design process (Brown, 2008), allowing for the opportunity to seek out and evaluate opposing or conflicting views would also be very useful in these cases.

Mitigating Objectivism Biases

Objectivism biases can happen anytime that you fail to appreciate that the way you see things may not be the way they actually are, and in any situation where you fail to appreciate other people's perspectives. Objectivism biases can occur anywhere, as they are about your perception of situations. In the workplace, they might commonly occur in any process where you are looking to influence others or sell an idea. A salesperson can easily gloss over that people are not as excited by a product as he is. A presenter to an audience can easily forget that others do not know what the presenter knows (i.e., "curse of knowledge" bias). An executive can easily miss the fact that not everyone is as on board with a big organizational change as she is.

Objectivism biases can occur anywhere, as they are about your perception of situations.

While corner-cutting biases can be mitigated by encouraging more cognitive effort from employees, the same cannot be said for objectivism biases in most cases. Objectivism biases occur because of invisible processes, cognitive machinery at work outside of our conscious awareness.

Putting in more effort typically does not resolve the problem here (and would be most likely to reinforce one's initial bias).

Instead, when you think an objectivism bias might be occurring, what is more likely to help is to get objective, outside opinions from others not on the team or project. Another technique to mitigate objectivism biases is to revisit ideas after a break to see them in a fresher, more objective light, and in particular trying to look at yourself and your message through other people's eyes. In these mitigation strategies, the brain's braking system (VLPFC) also plays a role in our ability to exercise cognitive control and disengage from our own, self-specific viewpoint. Taking a step backward and seeing ourselves and our decisions from a more objective perspective, putting ourselves in the mind of someone else, is also associated with brain regions associated with mentalizing (e.g., temporoparietal junction, medial prefrontal cortex). It is likely that the most effective strategies for mitigating objectivism biases will engage this neural circuitry, promoting perspective taking and self-evaluation.

...putting ourselves in the mind of someone else, is also associated with brain regions associated with mentalizing...

Mitigating Self-Protection Biases

Self-protection biases involve protecting your sense of self and retreating to a sense of safety with people you feel are similar to you or who share similar goals.

These kinds of biases will be common in "people decisions." Self-protection biases might occur in hiring decisions, in how teams are formed, in who is selected to be promoted, in deciding what kind of clients to work with, or in deciding who to have in a social network. Think of a recruiter who hires a person because he or she resembles others who have succeeded previously without paying enough attention to that individual's history or skill set. Or, consider a purchasing manager who feels more comfortable buying from someone who grew up in their hometown, just because it "feels safer." Or, consider a board deciding to give a key role to someone who most looks the part vs. someone who can do the best job.

Where corner-cutting biases can be mitigated with cognitive effort, and objectivism biases can be mitigated by introducing alternate perspectives, self-protection biases need to be mitigated in another way entirely. Here, a motivation to defend one's identity and worth leads to poor decisions. The solution is to diminish this motivation.

...a motivation to defend one's identity and worth leads to poor decisions.

There are at least two routes in accomplishing this. First, engaging in self-affirmation (thinking about things you value or people who are important in your life) affirms our sense of who we are and makes us less likely to "act out" in defense of the self (Sherman & Cohen, 2006; Reed & Aspinwall, 1998; Sherman & Cohen, 2002). Second, we can find ways to think of those who are different from us, and potentially a threat to the self, as more similar to us. One example of such a strategy is the jigsaw classroom pioneered by Elliot Aronson (e.g., Aronson, 2000), which promotes cooperative learning, engagement, and empathy in an educational setting. Thinking of ways that we and dissimilar (potentially threatening) others share goals, values, or preferences can help us think of ourselves as part of a larger group, in which case the threat to self is diminished.

These strategies for mitigating self-protection biases are based on promoting a positive and affiliative interpretation of ourselves and others and are associated with engaging the VLPFC (controlling our impulse to see different others negatively), disengaging the amygdala (reducing our arousal and threat response), associating dissimilar others with reward instead of threat (i.e., engaging reward neural circuitry including the ventral striatum and orbitofrontal cortex), and bringing online brain regions associated with mentalizing and promoting an identification between ourselves and others (e.g., medial PFC, temporoparietal junction).

For "people decisions" such as hiring or promoting, organizations could make it a policy to remove any identifying and potentially biasing information or features (e.g., name, gender, or ethnicity) from materials. This would be one way to prevent or mitigate self-protection biases at the outset of these types of decisions, but there are limitations to these strategies when face-to-face interactions are necessary (e.g., interviews, etc.).

Mitigating Time and Money Biases

Time and money biases can happen any time you are making decisions about the probability of risk or return, where to allocate money, or how to allocate resources including time, people and other assets. These might occur in financial decisions, investment decisions, resource allocation, strategy development, or planning for strategy execution. Examples include an executive not being able to let go of a business unit because of resources already invested in a project, or a CEO who is not willing to innovate in a new direction because it would compete with the company's existing business.

Time and money biases can happen any time you are making decisions about the probability of risk or return, where to allocate money, or how to allocate resources...

With time and money biases, again we need to get directly at the core biology driving the bias. Strategies that can work for time and money biases include imagining that you are making the decision for someone else; there is evidence that, when making decisions for others, you can be less biased, because the decision is less attached to the self (Gilbert et al., 2009; Hershfield et al., 2011). Getting greater distance between you and a decision is one strategy that might help. For example, you can imagine that the decision has been already been made in the past, and you are seeing it from a later, more objective and distanced point in time. In fact, studies suggest that recalling yourself in past events, as well as imagining yourself in future events, from a more objective, third-person perspective makes those events less emotional and less tied to the self (Libby et al., 2005; Pronin & Ross, 2006).

These strategies also rely heavily on the brain's braking system (VLPFC), which allows us to exercise the cognitive control needed to take a more objective, outside perspective and to engage our unique ability to project ourselves into a hypothetical future or past.

The Brain's Braking System for Breaking Bias

We have presented evidence of our behavioral and neural tendencies to seek reward and avoid pain and how these tendencies contribute to our susceptibility to the effects of bias. Though research on the neuroscience of breaking bias is in its infancy, there is evidence that specific brain regions are involved and show greater activation when people are behaving in a less-biased manner – specifically, the ventrolateral prefrontal cortex (VLPFC), the brain's "braking system" (Lieberman, 2009), as mentioned above.

Activation in the VLPFC has been associated with a reduced susceptibility to cognitive bias. Individuals who were able to correctly solve the "bat and ball" problem showed increased activity in the VLPFC (Spunt & Lieberman, in prep). Furthermore, individuals who were less susceptible to the framing effect (De Martino et al., 2006), temporal discounting (Boettiger et al., 2007), overconfidence effect (Beer & Hughes, 2010), and belief bias (Goel & Dolan, 2003) all showed increased activation in the VLPFC. When individuals are able to engage their mental brakes, inhibit their initial automatic responses, take a more objective view, and engage their System 2, then it appears that the influence of cognitive bias can be mitigated. This same system is central in regulating emotions and minimizing the impact of threat or reward (Lieberman, 2009). It seems there is one system for both managing emotions and managing biases, some of which may be driven by emotional (threat/reward) responses.

Activation in the VLPFC has been associated with a reduced susceptibility to cognitive bias.

Interestingly, it also appears that the more mindful a person is (i.e., how aware and receptive to his experiences in the present moment), the more active his VLPFC is during the labeling of negative emotions, which is a common emotion regulation strategy (Creswell et al., 2007).

... mindfulness training can increase positive judgments and reduce the negativity bias, as well as improve affective forecasting...

This study also showed that people high in mindfulness were able to neutralize the threat response when labeling an emotion (Creswell et al., 2007). Other studies show that mindfulness training can increase positive judgments and reduce the negativity bias (Kiken & Shook, 2011), as well as improve affective forecasting, which is another example of a time and money bias [see Appendix A; Emanuel et al., 2010]. One article recently published in this journal reviewed neuroimaging evidence that mindfulness training enhances self-awareness and reduces susceptibility to unconscious bias and emotional reactivity, highlighting the potential usefulness for mindfulness training in corporate leadership (Kirk, 2012).

...research supports the beneficial role of mindfulness in mitigating bias at an individual level...

If increased VLPFC activity is associated with reduced susceptibility to many types of cognitive bias and is also associated with greater mindfulness, individuals as well as organizations can foster mindfulness as a means of mitigating susceptibility to bias across the board. Currently, research supports the beneficial role of mindfulness in mitigating bias at an individual level, but more research needs to be done on strategies for increasing mindfulness in ways suitable to “people managers” at scale.

Issues include the “packaging” of the idea to be more acceptable to organizations, as well as finding minimum times for practice that can be effective. We are not advising this as a one-size-fits-all approach to any bias, simply as an area for more research. The ideal strategy may be increasing mindfulness of leaders, combined with applying processes that mitigate biases according to the underpinning neural issue driving that bias, according to the **COST™** model.

Other Current Models of Bias Mitigation

There are quite a few models and leadership education and training programs that are relevant to cognitive bias mitigation. As mentioned previously, diversity or sensitivity training programs tend to focus more narrowly on people-related decisions and encompass a subset of the self-protection and corner-cutting biases we discuss in this article (e.g., in-group and out-group biases, hot hand fallacy, halo effect; Kaley et al., 2006; Ross, 2008; Lublin, 2014; Babcock, 2006).

Other models and strategies for bias mitigation, like ours, have been more expansive and inclusive of a wider variety of bias. For example, Campbell et al. (2010) identify three “red-flag conditions” that are likely to lead to biased decision-making (i.e., presence of inappropriate self-interest, distorting attachments, and misleading memories) and outline a seven-step process for identifying those red flags. Though an alternative approach that is useful and concise, we believe that our model provides more detail and structure surrounding the neuroscience of bias – not only its cause, but potential mitigation strategies – and why different targeted strategies could work for specific categories of bias.

Other comprehensive models of leadership development less directly target unconscious bias, but are nonetheless aimed at reducing the impact of bias in decision-making. Work by William Torbert focuses on “action logic” profiles, or leadership styles characterized by how one interprets his/her environment and how s/he reacts to perceived challenges (Rooke & Torbert, 2005). Though never explicitly discussed as unconscious biases, the descriptions of the problems underlying each style are consistent with the biases we discuss here (e.g., “Opportunists tend to regard their bad behavior as legitimate [...] reject feedback, externalize blame” [Rooke & Torbert, 2005, pg. 68] – similar to fundamental attribution error, self-serving bias, and egocentric bias). Mitigation strategies in this program stress awareness, learning about the drivers of our behavior, and perspective taking.

Similarly, work by Kegan and Lahey at Harvard highlights the underlying, largely unconscious motivations behind human behavior that can hold people back or even be in direct conflict with their values and goals, resulting in an “immunity to change.”

Though again not explicitly characterized as unconscious bias, their model focuses on “competing commitments” and “big assumptions” that people may be unaware of, and their mitigation strategies stress the need for self-reflection, promoting self-awareness, and taking a more objective perspective (Kegan & Lahey, 2001).

We believe that our model complements models and programs such as these – bringing the ‘unconscious’ part of bias to the forefront and presenting a model where people can readily identify different categories of bias (regardless of leadership style or position in the company), understand the neural underpinnings of bias, and highlight specific mitigation strategies and their neural underpinnings as well. Understanding the neural basis of bias and its mitigation can promote awareness and pave the way for the acceptance of strategies to prevent bias in decision-making at the organizational level.

Summary

Biases are a significant issue in organizations. In this article, we have presented a serious attempt to organize the roughly 150 biases into a model that can be applied in organizational settings. Our goal is to help individuals in a wide range of positions, from business leaders all the way to front-line staff, and, more broadly, organizations as a whole identify and then mitigate biases based on the underlying issues associated with each broad category of bias. For example, to identify and address corner-cutting biases, we must appreciate our neural predisposition to make fast and efficient judgments, identify situations in which more deliberative thought and strategies are necessary to avoid bias, and encourage processes that place a premium on engaging cognitive effort instead of going with intuition or gut instinct in these situations. Alternatively, for objectivism biases, cognitive effort is generally not the main issue – instead, we must appreciate that our brain’s default setting is an egocentric one which assumes that our experience and perception of reality is the objective truth. In order to identify and address these biases, implementing strategies that encourage actively seeking out more objective perspectives and others’ viewpoints will be most helpful.

Biases are a significant issue in organizations.

While this model is just being released, and considerably more research and development needs to be done around both theory and practice, we believe this model may be a useful step in reducing the unhelpful biases that are at the heart of many organizational challenges today. We appreciate feedback and input about the model from academics and practitioners alike.

Appendix A: The COST™ Model of Bias

Examples of Corner-Cutting Biases

Belief bias – Deciding whether an argument is strong or weak based on whether or not one agrees with its conclusion. Like the example in the text about the addictiveness of cigarettes, this bias entails letting one’s beliefs influence how one evaluates information.

Confirmation bias – Seeking and finding evidence that confirms one’s beliefs and ignoring evidence that does not support those beliefs. Some examples include selectively reading studies and articles that support your views or theories and ignoring those offering conflicting information; and only reading news sources that support your political beliefs.

Availability bias – Making a decision based on the information that’s most readily accessible (comes to mind most quickly) instead of on objective information. For example, you might think it’s more likely to die from a shark attack than from falling airplane parts because shark attacks are more widely publicized but happen less often, and deaths from falling airplane parts are less widely reported but happen more often (Read, 1995).

Anchoring bias – A tendency to rely too heavily on the first piece of information offered when making a decision. This piece of information is the “anchor”, and other information is interpreted around this anchor. For example, the initial price set for buying a car is the price that subsequent negotiations will follow from.

Base rate fallacy – The tendency, when judging how probable something is, to ignore the base rate (the rate that it occurs in general) and to focus on other information. For example, only 5% of applicants are interviewed for a certain job, but you know that you are perfect for the job and are convinced that the probability of your getting an interview is higher than 5%.

Planning fallacy – The tendency to underestimate how long it will take to complete a task, how much it will cost, and its risks, while at the same time overestimating its benefits.

Representativeness bias – Misjudging that something that is more representative means that it is more likely. For example, if given a choice between teacher and yoga instructor, we’re more likely to think that someone who is described as being very spiritual, doing yoga, and meditating every day is a yoga instructor because s/he is representative of that group. But in reality, teacher is more probable because there are more teachers than yoga instructors.

Hot hand fallacy – Believing that someone who was successful in the past has a greater chance of achieving further success.

One example is expecting a gambler who has had a winning streak to be more likely to continue winning, even though the probability of winning has not changed.

Halo effect – Letting someone’s positive qualities in one area (e.g., attractiveness, optimistic personality) influence one’s perception of him/her in other areas (e.g., job performance, leadership ability).

Examples of Objectivism Biases

Bias blind spot – Identifying biases in other people but not in oneself.

Curse of knowledge – The more one knows, the harder it is to appreciate the perspective of those who know less.

False consensus effect – Overestimating the extent to which others agree with you; the tendency to assume that your beliefs, habits, and opinions are “normal” and that others think the same way.

Fundamental attribution error – Believing that one’s own errors or failures are justifiable due to external circumstances, but others’ errors are due to their character, or internal factors and are cause for greater concern. For example, “I made a mistake because I was having a bad day; you made a mistake because you’re not a very intelligent person.”

Hindsight bias – Seeing past events as having been predictable even though they may not have been; the feeling of “I knew it all along” even though the outcome was mostly likely unforeseeable.

Illusion of control – Overestimating the degree of control one has over external events. For example, believing that if you had just left the house 2 minutes earlier, you would have avoided getting caught at every traffic light is an illusion of control.

Illusion of transparency – Overestimating the degree to which your mental state is accessible to others. For example, public speakers believe their nervousness and stage-fright were obvious to the audience, but were really not.

Egocentric bias – The general tendency for information about oneself to have a disproportionate effect on judgments and decisions. For example, overestimating our ability to communicate with others, assuming that others understand what we understand.

Examples of Self-Protection Biases

Overconfidence effect – The tendency to believe that we are better than average (e.g., better drivers, more attractive, better leaders, better teachers, better friends, than others) and to be overconfident in the accuracy of our judgments.

In-group bias – Perceiving people who are similar to you (e.g., are of the same ethnicity, practice the same religion, are from the same hometown) more positively than people who are more different from you.

Out-group bias – Perceiving people who are different from you (e.g., are of a different ethnicity, practice a different religion, are of a lower or higher socioeconomic status) more negatively than people who are more similar to you.

Self-serving bias – Believing that one’s own successes are due to internal, personal factors but one’s own failures are due to external circumstances. For example, “I got my new job because I am intelligent and talented; I lost my old job because the economy is bad.”

Self-evaluation maintenance – Undermining others who are likely to be the people you are socially compared with (friends, close colleagues), even if they are close to you, in order to make yourself look better by comparison. One example of this is when a person is more likely to help a stranger win a competition than a friend because he is more likely to be compared with his friend – his friend’s success makes him look less successful by comparison.

Examples of Time and Money Biases

Loss aversion – Making a risk-averse choice if the expected outcome is positive, but making a risk-seeking choice in order to avoid negative outcomes. For example, if a person is going to win money, s/he is more likely to take a less-risky bet to minimize the chances of losing; if a person is going to lose money, s/he is more likely to take a more risky bet to increase the chances of avoiding the loss.

Framing effect – Making a different judgment based on whether the decision is presented as a gain or as a loss, despite having the same objective information. For example, choosing to take a risk to avoid a 40% probable loss, but choosing to avoid a risky decision for a 60% probable gain.

Sunk costs – Having a hard time giving up on something (e.g., a strategy, an employee, a process) after investing in it (e.g., time, money, training), even though the investment has already been made and can’t be recovered.

Endowment effect – Expecting others to pay more for something that we own than we would be willing to pay for the same thing that someone else owns.

Affective forecasting – The fact that people are surprisingly poor judges of their future emotional states.

Temporal discounting – The tendency to devalue rewards as they move farther into the future. For example, given a choice between \$5 now and \$10 tomorrow, people choose \$10 tomorrow. But, given a choice between \$5 now and \$10 in six months, people choose \$5 now.

References

- Aronson, E. (2000). Nobody left to hate. *The Humanist*, 60, (3), 17–21.
- Babcock, P. (2006). Detecting hidden bias. *HR Magazine*, 51, (2), retrieved from: <http://www.shrm.org/Publications/hrmagazine/EditorialContent/Pages/0206cover.aspx>.
- Banaji, M.R., & Greenwald, A.G. (2013). *Blindspot: Hidden Biases of Good People*. New York, NY: Delacorte Press.
- Beer, J.S., & Hughes, B.L. (2010). Neural systems of social comparison and the “above-average” effect. *NeuroImage*, 49, 2671–2679.
- Boettiger, C.A., Mitchell, J.M., Tavares, V.C., Robertson, M., Joslyn, G., D’Esposito, M., & Fields, H.L. (2007). Immediate reward bias in humans: Fronto-parietal networks and a role for the catechol-o-methyltransferase 158^{Val/Val} genotype. *The Journal of Neuroscience*, 27, (52), 14383–14391.
- Brown, T. (2008). Design Thinking. *Harvard Business Review*, 86, (6), 84–92.
- Campbell, A., Whitehead, J., & Finkelstein, S. (2010). Why good leaders make bad decisions. *Harvard Business Review*, 87, (2), 60–66.
- CEB (The Corporate Executive Board Company). (2012). CEB 2012 Global Labor Market Survey.
- Creswell, J.D., Way, B.M., Eisenberger, N.I., & Lieberman, M.D. (2007). Neural correlates of dispositional mindfulness during affect labeling. *Psychosomatic Medicine*, 69, 560–565.
- Dane, E., Rockmann, K.W., & Pratt, M.G. (2012). When should I trust my gut? Linking domain expertise to intuitive decision-making effectiveness. *Organizational Behavior and Human Decision Processes*, 119, 187–194.
- Daniel, R., & Pollmann, S. (2012). Striatal activations signal prediction errors on confidence in the absence of external feedback. *NeuroImage*, 59, (4), 3456–3467.
- De Martino, B., Kumaran, D., Seymour, B., & Dolan, R.J. (2006). Frames, biases, and rational decision-making in the human brain. *Science*, 313, (5787), 684–687.
- Emanuel, A.S., Updegraff, J.A., Kalmbach, D.A., & Ciesla, J.A. (2010). The role of mindfulness facets in affective forecasting. *Personality and Individual Differences*, 49, (7), 815–818.
- Evans, J.St.B.T., Handley, S.J., & Harper, C.N.J. (2001). Necessity, possibility, and belief: A study of syllogistic reasoning. *The Quarterly Journal of Experimental Psychology*, 54A, (3), 935–958.
- Evans, J.St.B.T., & Curtis-Holmes, J. (2005). Rapid responding increases belief bias: Evidence for the dual-process theory of reasoning. *Thinking & Reasoning*, 11, (4), 382–389.
- Flynn, G. (1998). The Harsh Reality of Diversity Programs. *Workforce*, 77, (12), 26–28.
- Frederick, S. (2005). Cognitive Reflection and Decision Making. *The Journal of Economic Perspectives*, 19, (4), 25–42.
- Gilbert, D.T., Killingsworth, M.A., Eyre, R.N., & Wilson, T.D. (2009). The surprising power of neighborly advice. *Science*, 323(5921), 1617–1619.
- Goel, V., & Dolan, R.J. (2003). Explaining modulation of reasoning by belief. *Cognition*, 87(1), B11–B22.
- Griffin, D.W., & Ross, L. (1991). Subjective construal, social inference, and human misunderstanding. *Advances in experimental social psychology*, 24, 319–359.
- Haidt, J. (2012). *The Righteous Mind: Why Good People Are Divided by Politics and Religion*. New York, NY: Pantheon.
- Hershfield, H.E., Goldstein, D.G., Sharpe, W.F., Fox, J., Yeykelis, L. Carstensen, L.L., & Bailenson, J.N. (2011). Increasing saving behavior through age-progressed renderings of the future self. *Journal of Marketing Research*, 48, (SPL), S23–S37.
- Johnson, D.D.P., & Fowler, J.H. (2011). The evolution of overconfidence. *Nature*, 477, 317–320.
- Kahneman, D. (2003). A perspective on judgment and choice: Mapping bounded rationality. *American Psychologist*, 58, (9), 697–720.
- Kahneman, D. (2011). *Thinking, Fast and Slow*. New York, NY: Farrar, Straus, and Giroux.
- Kahneman, D., Knetsch, J.L., & Thaler, R.H. (1990). Experimental tests of the endowment effect and the coase theorem. *Journal of Political Economy*, 98, (6), 1325–1348.
- Kahneman, D., Lovallo, D., & Sibony, O. (2011). The Big Idea: Before You Make That Big Decision.... *Harvard Business Review*, 89, (6), 50–60.
- Kahneman, D., & Tversky, A. (1984). Choices, Values, and Frames. *American Psychologist*, 39, (4), 341–350.
- Kalev, A., Dobbin, F., & Kelly, E. (2006). Best Practices or Best Guesses? Assessing the Efficacy of Corporate Affirmative Action and Diversity Policies. *American Sociological Review*, 71, 589–617.
- Kegan, R. & Lahey, L.L. (2001). The real reason people won’t change. *Harvard Business Review*, 79, (10), 84–92
- Kiken, L.G., & Shook, N.J. (2011). Looking up: Mindfulness increases positive judgments and reduces negativity bias. *Social Psychological and Personality Science*, 2, (4), 425–431.

- Kirby, K.N., & Marakovic, N.N. (1995). Modeling myopic decisions: Evidence for hyperbolic delay discounting within subjects and amounts. *Organizational Behavior and Human Decision Processes*, 64, 22–30.
- Kirk, U. (2012). Neural substrates of corporate decision-making. *NeuroLeadership Journal*, 4, 70–80.
- Libby, L.K., Eibach, R.P., & Gilovich, T. (2005). Here's looking at me: The effect of memory perspective on assessments of personal change. *Journal of Personality and Social Psychology*, 88, (1), 50–62.
- Lieberman, M.D. (2009). The brain's braking system (and how to 'use your words' to tap into it). *NeuroLeadership Journal*, 2, 9–14.
- Lieberman, M.D. (2013). *Social: Why Our Brains Are Wired to Connect*. New York, NY: Crown.
- Lilienfeld, S.O., Ammirati, R., & Landfield, K. (2009). Giving Debiasing Away: Can Psychological Research on Correcting Cognitive Errors Promote Human Welfare? *Perspectives on Psychological Science*, 4, (4), 390–398.
- Lovullo, D. & Sibony, O. (2010). The case for behavioral strategy. *McKinsey Quarterly*, 2, 30–43.
- Lublin, J.S. (2014). Bringing hidden biases into the light: Big businesses teach staffers how 'unconscious bias' impacts decisions. *Wall Street Journal*, retrieved from: <http://online.wsj.com/news/articles/SB10001424052702303754404579308562690896896>.
- MacLean, C.L., Brimacombe, C.A.E., & Lindsey, D.S. (2013) Investigating Industrial Investigation: Examining the Impact of A Priori Knowledge and Tunnel Vision Education. *Law and Human Behavior*, 37, (6) , 441–453.
- Moss-Racusin, C.A., van der Toorn, J., Dovidio, J.F., Brescoll, V.L., Graham, M.J., & Handelsman, J. (2014). Scientific Diversity Interventions. *Science*, 343, 615–616.
- Pronin, E., Lin, D.Y., & Ross, L. (2002). The Bias Blind Spot: Perceptions of Bias in Self Versus Others. *Personality and Social Psychology Bulletin*, 28, 369–381.
- Pronin, E., & Ross, L. (2006). Temporal differences in trait self-ascription: When the self is seen as an other. *Journal of Personality and Social Psychology*, 90, (2), 197–209.
- Reed, M.B., & Aspinwall, L.G. (1998). Self-affirmation reduces biased processing of health-risk information. *Motivation and Emotion*, 22, (2), 99–132.
- Rock, D. (2008). SCARF: A brain-based model for collaborating with and influencing others. *NeuroLeadership Journal*, 1, 78–87.
- Rock, D., & Cox, C. (2012). SCARF in 2012: Updating the social neuroscience of collaborating with others. *NeuroLeadership Journal*, 4, 129–142.
- Rooke, D. & Torbert, W.R. (2005). Seven transformations of leadership. *Harvard Business Review*, 83, (4), 66–76
- Ross, H. (2008). Proven strategies for addressing unconscious bias in the workplace. *CDO Insights*, 2, (5), 1–18.
- Satpute, A.B., & Lieberman, M.D. (2006). Integrating automatic and controlled processes into neurocognitive models of social cognition. *Brain Research*, 1079, 86–97.
- Satterthwaite, T.D., Ruparel, K., Loughhead, J., Elliott, M.A., Gerraty, R.T., Calkins, M.E., Hakonarson, H., Gur, R.C., Gur, R.E., & Wolf, D.H. (2012). Being right is its own reward: Load and performance related ventral striatum activation to correct responses during a working memory task in youth. *NeuroImage*, 61, (3), 723–729.
- Senge, P.M. (1990). *The Fifth Discipline*, New York: Doubleday/Currency.
- Sherman, D.K., & Cohen, G.L. (2002). Accepting threatening information: Self-Affirmation and the reduction of defensive biases. *Current Directions in Psychological Science*, 11, (4), 119–123.
- Sherman, D.K., & Cohen, G.L. (2006). The psychology of self-defense: Self-affirmation theory. *Advances in experimental social psychology*, 38, 183.
- Skowronski, J.J., & Carlston, D.E. (1989). Negativity and extremity biases in impression formation: A review of explanations. *Psychological Bulletin*, 105, 131–142.
- Spunt, R.P., Lieberman, M.D., Cohen, J.R., Eisenberger, N.I. (2012). The phenomenology of error processing: The dorsal ACC response to stop-signal errors tracks reports of negative affect. *Journal of Cognitive Neuroscience*, 24, (8), 1753–1765.
- Tamir, D.I., & Mitchell, J.P. (2011). The default network distinguishes construals of proximal versus distal events. *Journal of Cognitive Neuroscience*, 23, (10), 2945–2955.
- Tversky, A., & Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science*, 211, 453–458.
- van Veelen, M., & Nowak, M.A. (2011). Evolution: Selection for positive illusions. *Nature*, 477, 282–283.
- Vedantam, S. (2008, January 20). Most Diversity Training Ineffective, Study Finds. *The Washington Post*, Retrieved from <http://www.washingtonpost.com/wp-dyn/content/article/2008/01/19/AR2008011901899.html>
- West, R.F., Meserve, R.J., & Stanovich K.E. (2012). Cognitive sophistication does not attenuate the bias blind spot. *Journal of Personality and Social Psychology*, 103, (3), 506–519.
- Woolley, A.W., Chabris, C.F., Pentland, A., Hashmi, N., & Malone, T.W. (2010). Evidence for a Collective Intelligence Factor in the Performance of Human Groups, *Science*, 330, (6004), 686–688.

