

## Innovating Government Decision Making Through Analytics

Steve Bennett, Ph.D., SAS Institute Inc.

### ABSTRACT

Everything that a government agency does revolves around decision making. While the types of decisions can be highly varied, from strategic decisions about agency budget priorities and delivery of social benefits, to split-second operational decisions on the battlefield, decisions are at the center of what governments do. But how do we make decisions? What evidence and information gets considered? And ultimately, since humans either alone or in combination are the ultimate makers of decisions, how does that evidence and information get interpreted, processed, synthesized, and assessed in our brains to ultimately arrive at a particular alternative, given an array of possible choices? Understanding how we as humans make decisions, and even more importantly, how we can make them better – is critical for successful execution of government agency missions.

The challenge is that over the last 50 years, the biological, social, and behavioral economic fields of scientific study have started to show that humans are not very good at making rational decisions, at least from a mathematical sense. When we receive and process information, and then render judgments, we are subject to a host of biases, ranging from impacts arising from the order in which we receive inputs, to subtle and often subconscious assessments of input quality because of something we already think or believe. These biases can have a significant impact on how information is processed and assessed, as well as how judgments are rendered. In addition, despite the fact that modern governments are awash in more quantitative information than at any point in history, the fact that humans are generally poor at creating an accurate mental representation of quantitative information for very large or very small numbers remains. Ranging from budget discussions, to risks from rare low-frequency events, government decisions makers must make decisions – sometimes life or death – with built-in cognitive limits to how well quantitative information can be internalized and processed. And finally, humans are limited in our decision making rationality and consistency when those decisions or the input data upon which they depend, include uncertainty. Given that many if not all critical government decisions are fraught with significant, if not high degrees of uncertainty, this creates difficulty in making sure that decisions are made in the best way possible, and securing the best possible outcomes for missions and for citizens. From how we understand uncertainty, to how we intuitively understand numbers, to inherent judgment biases; we make decisions that much of the time, are not rational.

In this paper, we will motivate the need for innovation in decision making based on these cognitive decision making challenges, drawing from historical examples, as well as examples from current government challenges. We'll explore aspects of cognitive decision science focusing on the three limitations for human decision making outlined above – judgment bias, quantitative intuition, and difficulty under uncertainty, highlighting our decision making irrationality. We'll then highlight the impacts and challenges of irrational decision making using real-world examples. Finally, analytics will be discussed as an important source of evidence for helping to improve how all of us – and government in particular – makes decisions. We'll talk specifically about analytics from an organization-of-information perspective in helping us circumvent bias, improve intuition for quantitative information, and help us do better and effectively handling uncertainty in the decision making process. Limitations and descriptions of how analytics can help overcome them will be presented in the context of government decision making, but are applicable broadly in range of organizations and industries.

### INTRODUCTION

In 1995, two friends from Saudi Arabia left their homes to travel to Eastern Europe to fight alongside the Muslims in the Bosnian War. Returning home with significant combat experience, they traveled to Afghanistan in 1999 to fight alongside the Taliban against the Afghan Northern Alliance. By this time, Nawaf al-Hazmi and Khalid al-mihdhar had gained the attention of the senior leadership of al Qaeda, to

include Osama bin Laden, who was looking to recruit loyal soldiers for a bold attack against the United States on September 11<sup>th</sup>, 2001.

In January of 2000, Hazmi and Mihdhar attended an al Qaeda summit in Malaysia at which attack planning was discussed for the September 11<sup>th</sup> attacks. United States intelligence services were aware of the summit through an intercepted phone call from Midhar to Yemen, and asked Malaysian intelligence services to infiltrate the meeting and gather video and audio information. While the Malaysians were not able to gather the attack details, U.S. intelligence nevertheless had information in video form. While Hazmi was videotaped, no connection was made to his name – and authorities already knew he was in the country from an overstayed visa.

While in the United States, Hazmi and Mihdhar applied for and received drivers licenses in multiple states as well as other forms of identification. No connections were made to national intelligence data. In addition, they interacted frequently with state and local law enforcement, including an April 2001 speeding stop in Oklahoma in which a driver's license check turned up no warrants even though their names were known by both the CIA and NSA as suspected terrorists. On August 23<sup>rd</sup>, 2001, just a few weeks before 9/11, the CIA received a list of 19 names from a foreign intelligence agency, indicating people suspected of imminent attack planning inside the U.S. Four of the names were known for certain to be involved, and Hazmi was one of those four names.

And finally, about three weeks before the attacks of September 11<sup>th</sup>, Zacarias Moussaoui was arrested in Minnesota for immigration-related issues; in his possession were a laptop, two knives, Boeing 747 training manuals, and flight simulator software. U.S. authorities were unable to search his laptop, and did not make a crucial connection between him and another detained al Qaeda operative who after 9/11 said he trained with Moussaoui in an al Qaeda training camp in Afghanistan. The independent 9/11 Commission established to investigate the law enforcement, security, and intelligence gaps that contributed to 9/11 believes that connecting this information together would have disrupted if not derailed the 9/11 attacks.

Ultimately the events of September 11<sup>th</sup>, 2001 are an example of the inherent difficulties and challenges government have in synthesizing and connecting many different pieces of information together to make what could be critical security decisions for the safety and security of citizens. Specifically, the 9/11 Commission said it this way:

*“The U.S. government has access to a vast amount of information.... the storehouse is immense...But the U.S. government has a weak system for processing and using what it has.”*  
(Commission, 2004)

The challenges of organizing, synthesizing, interpreting, analyzing, and presenting data and information in a context that supports timely and effective decision making is not just unique to the run-up to the events of September 11<sup>th</sup>, but persists throughout almost all aspects of Government decision making. Yet this comes at a time when (as the 9/11 Commission notes) we are awash in unprecedented amounts of data and information in almost every agency in government. If the most important thing government agencies do is make timely, effective, and efficient decisions in their missions, how can we improve our ability leverage this vast amount of information to make sure we're doing our best to serve, support, and protect citizens? The answer lies in first understanding the limits of our own cognitive abilities with respect to decision making.

## **COGNITIVE CHALLENGES IN HUMAN DECISION MAKING**

Humans are not very good decision makers. How we define a “good decision” is certainly up for debate, but in terms of adherence to rationality and how well information and uncertainty is successfully and appropriately integrated into someone's decision making logic and processes, there are years of experience that show our susceptibility to a whole host of biases, inconsistencies, influences, limitations, and other challenges. While cognitive psychologists, behavioral economists, and neurobiologists have studied the biology and psychology behind human decision making for quite some time, recent research over the last 50 years has crystallized and quantified the impacts of these challenges on our decision making in new ways. Although there are many, I'll focus on three particularly key cognitive limitations that

impact our decision making: biases in judgment, difficulties in intuitively handling quantitative information with a high dynamic range, and irrationality when decisions involve uncertainty.

## JUDGMENT BIASES

When the blockbuster movie “Jaws” debuted in 1975, Americans flocked to theaters to see Richard Dreyfus spend several hours fighting off a large robotic shark. But if we imagine assessing judgments by the general public regarding an estimate of the frequency of shark attacks, following the release of the movie, Americans’ estimates of the frequency of shark attacks went up and were much higher than the real risk. Why is this? Was there some new evidence provided to Americans regarding these events? Was new data released to which Americans had access, when making these judgment that they did not have previously? Almost certainly not. The vivid 2-hour depiction of a battle with a fearsome great white shark created a salient and vivid... let’s say “available” impression in our minds, that impacted our assessment of the frequency of what is an incredibly rare event (we’re 30 times more likely to die from a lightning strike than from a shark attack). This classic example of what economist Daniel Kahneman and Amos Tversky called the “availability heuristic,” or availability bias demonstrates how easy it is for our judgments – our decisions – to be impacted by something recent, vivid, and/or available to us in our minds (Tversky 1973). There are a litany of other examples of availability bias. Consider the following statement: “Smoking may not really have all that many health impacts because I know someone who smoked every day and lived to be 100.” But this and other biases show up in real problems in government decision making as well, not just in movies and around personal assessments of health risks.

Availability bias, and other judgment biases we will address below aren’t just things that pop up from scary movies and assessments of personal health. They show up in critical decisions governments must make regarding the safety and security of their citizens. During the twelve years I worked for the Department of Homeland Security following the events of September 11<sup>th</sup>, I had the opportunity to lead a number of large scale risk analyses seeking to support decision making around how to prioritize areas of biological weapons counterterrorism for investment. The U.S. has just been attacked with weaponized Anthrax transmitted through the U.S. mail in the fall of 2001, and were we to try and develop countermeasures for every possible manner of bioterrorism, we would bankrupt the nation many times over. Coming up with evidence-driven priorities was essential. As we set about designing the study, we needed two things to calculate quantitative estimates of risk as expected loss. First, we needed consequence estimates for different types of attack scenarios; e.g., how many deaths and illnesses would we expect for different types of attack scenarios with different types of biological weapons. The consequences were fairly easy to quantitatively calculate, as there were many high-fidelity models that would take inputs about an attack and generate high-quality estimates of the expected consequences. The second thing that we needed was much harder to get, and required human judgment – probabilities of the occurrence of different types of attacks. We needed experts from the government to tell us what the relative frequencies were estimated to be for different types of attacks based on their understanding of the difficulty of those attacks, preferences of terrorism adversaries, and a range of other criteria. Our team developed a sophisticated expert elicitation protocol to be used in interviewing these experts to carefully and accurately assess their judgments on these attack frequencies. Then, about a week before we were ready to sit down with these experts to capture their judgments, a man in Las Vegas hotel room attempted to kill his wife with a particular poison, and it was all over the news. This event had nothing to do with biological terrorism, nor did it provide any real evidence that should have informed expert judgments about terrorism frequencies. Yet, guess what these government experts wanted to talk about? They wanted to talk about this particular poison, almost regardless of the independent information they had from their experience and data sources about bioterrorism. They had just seen “Jaws” and were looking for sharks. This example and many others like it demonstrate that it’s not a matter of a lack of training or expertise that makes us susceptible to availability bias – rather it’s the fact that we’re *human*.

Kahneman and Tversky first wrote about availability bias and a raft of other judgment biases in a seminal 1974 paper which outlined for the first time the impact of how and when we receive information, and in what form, on our judgments and decisions. Table 1 lists those initial heuristics and biases described in 1974 – and although many others have been articulated since then, these form a good place to start. I won’t address all of them, but will hone in on several that are of particular note for decision making in government.

<b>Availability Heuristic:</b> Estimation of an event's likelihood is based on the ease with which instances or associations come to mind. (Tversky 1973, Tversky 1974)	<b>Overconfidence:</b> Excessive confidence in one's own answers to questions. (Hilbert 2012)
<b>Anchoring:</b> The tendency to rely too heavily, or "anchor", on one trait or piece of information when making decisions (usually the first piece of information acquired on that subject) (Tversky 1974)	<b>Confirmation Bias:</b> The tendency to search for, interpret, focus on and remember information in a way that confirms one's preconceptions. (Oswald 2004).
<b>Representativeness:</b> When determining if something belongs in a particular category, this heuristic evaluates the degree to which that something is representative of the category (Tversky 1974)	<b>Conservatism Bias:</b> The tendency to revise one's belief insufficiently when presented with new evidence. (Hilbert 2012).
<b>Outcome Bias:</b> An error made in evaluating the quality of a decision when the outcome of that decision is already known (Baron 1988)	<b>Survivorship Bias:</b> Concentrating on the people or things that "survived" some process and inadvertently overlooking those that didn't because of their lack of visibility (Shermer 2014).

**Table 1. Select judgment biases and heuristics. Shaded cells are drawn from Tversky 1973/1974; remaining cells from cited references**

Another important judgment bias that impacts decision making for all of us is something that Kahneman and Tversky called “adjustment and anchoring,” highlighting the fact that humans have an inherent tendency to give more weight to the first piece of evidence or information that they encounter. They performed a classic experiment demonstrating this in which they gave two groups of high school students two math problems, and then gave them 5 seconds to estimate (not calculate) an answer. The first group of students was given the problem below in expression (1);

$$1 \times 2 \times 3 \times 4 \times 5 \times 6 \times 7 \times 8 \tag{1}$$

The median estimate across the sample of high school students for this fairly simple multiplication problem was 512. A second group of students was given the mathematically-equivalent problem below in expression (2):

$$8 \times 7 \times 6 \times 5 \times 4 \times 3 \times 2 \times 1 \tag{2}$$

The median estimate for the group of students asked to estimate the solution for expression (2) was 2,250; well over four times greater than the median estimate of the students estimating expression (1). Why were the estimates so different for what is essentially the exact same math problem? The answer is anchoring bias. Because the written language of English proceeds from left to right, the expression (2) estimators encountered a large number as the first piece of information they saw, whereas the expression (1) estimators encountered a small number as their first piece of information. Each group of students “anchored” on this early information, and it biased their judgments in quantifiably measurable ways.

(By the way, the correct answer is 40,320).

But while this clever math problem is interesting and effectively highlights and measures anchoring bias, this important form of judgment bias has much more important consequences in government decision making contexts. Anchoring bias is something to which people are inherently susceptible, and in professions that require assessment of evidence for law enforcement or security decision making, an over-reliance on early information can be to the detriment of the quality of the decisions that are made later. Imagine intelligence analysts who receive early information about a particular terrorist organization, and what their plans might entail. This information forms the first mental picture in the mind of the analyst about what the state of reality may be regarding this particular group. But what if the next several pieces of evidence encountered by the analyst contradict or challenge the initial information? What if these later

pieces of information are of higher quality or are better sourced? Anchoring bias means that as humans we have a harder time “coming off” early information even if subsequent information is higher quality and/or contradictory. The decision challenges presented by anchoring bias manifests everywhere in government decision making, from intelligence, to law enforcement, to foreign policy. And anyone with multiple children living in their home understands this bias very well – when my wife or I hear the crash of something breaking in the next room, the kids in that room understand full well the value of being the first to race into the room where we are to get their version of the story out first!

Another important judgment bias related to anchoring bias, is confirmation bias. Confirmation bias refers to the human tendency to ascribe higher value or evidentiary weight to information that confirms something we already believe. Put another way, we have an inherent preference for information that does not challenge our *status quo*. In popular life, we can see this in political discussions and arguments in which a people tend grasp onto information that confirms their beliefs or values, instead of wrestling with the evidentiary value of information that could challenge their beliefs.

But again, political discussions aside, confirmation bias is something that can greatly limit government decision making in that it makes it difficult for information to challenge current ways of thinking. The missed intelligence and law enforcement signals leading up to the events of September 11<sup>th</sup>, 2001 are at least in part an example of confirmation bias, in that the failure to detect and disrupt those events has been described as a “failure of imagination.” (Commission, 2004).

A more personal example of this came in a large-scale risk analysis project my team and I had been working on for a number of years inside the Department of Homeland Security. We had been working on producing analyses that integrated key homeland security risks with the effectiveness of current (and possible future) government programs that could address or mitigate those risks. In short, the analysis we produced could, in limited fashion, answer questions about which programs reduced the most risk, and for the first time, could inform trade-off decisions between different programs. This had never been done before, and (theoretically) afforded senior leadership with a mechanism for informing where the next marginal dollar was best spent, within a particular mission scope.

After two years of work, we were ready to present results to an audience of very senior, seasoned government leadership from across the Department. After concluding what we thought was a very compelling presentation about the impacts of our analysis and where potential future investments could be made to best reduce risk for Americans, one of the senior leaders in the room leaned back in his chair, and informed us: “This is a great presentation. Outstanding work. But here in [our agency], when we need to make a decision, we just trust our gut.” Now to be sure, this event was more than just an example of confirmation bias – our analyst team clearly did not do a great job managing this stakeholder and bringing him along over the two years of our work – a lesson I learned well here as a young analyst. But certainly, this senior leader was clear in communicating that whatever information he received was to be filtered through his intuition and currently-held beliefs. Experience and a well-seasoned “gut” is an essential tool for any senior leader, and we dare not reject experience-informed intuition as important in our agency leaders. But in complement to this intuition, what if there really was new information that *should* challenge the way decisions have been made previously? What mechanism should exist to effectively inform and supplement experiential judgment with new information? Confirmation bias makes this difficult.

There are a number of these biases that we haven’t touched on in high detail, but are nevertheless important. One example is *overconfidence*. This likely needs little explanation; it represents the fact that humans tend to wildly overestimate the certainty of their own judgments. There have been lots of examples of infrastructure failures and other disasters resulting from overconfidence in judgments, ranging from the sinking of the Titanic to the 1976 failure of the Teton dam, to fire at the Browns Ferry Nuclear Power Plant. Another additional bias is *outcome bias* in which we tend to assess the quality of a decision based on the outcome as opposed to whether the decision effectively synthesized all available information. This is captured by the observation that quite terrible decisions can still be “lucky” and produce a good outcome. Another related bias is *survivorship bias* in that we overweight information because it “survives,” while we ignore the information that did not. For example, consider the foolish statement, “I know that Apple and Hewlett-Packard both started as small companies in their founders’ garages, and have now grown to be multi-billion dollar corporations. Therefore it must be pretty easy to


start a company in my garage and become wildly successful.” Here, we are of course overweighting the “survivors” and not considering the thousands upon thousands of companies that failed, and so we never heard about them.

Hopefully this section has highlighted some of the many ways that judgment bias can and does impact decision making in important government missions areas. Next, I will turn to our second inherent limitation in decision making – the inability of humans to intuitively handle quantitative information that has a large dynamic range.

### **HUMAN DIFFICULTIES REASONING ABOUT QUANTITATIVE INFORMATION WITH A LARGE DYNAMIC RANGE**

Beyond the judgment biases inherent in human decision making, there is another important limitation to human decision making that’s particularly relevant as we have access to more and more data. This limitation is specifically that humans have an increasingly hard time reasoning about quantitative information that has a large dynamic range. In short, when numbers get very big or very small, we begin to lose our ability to intuitively reason about them. When that happens, the impact that those numbers could (or should) have on us when we make decisions, gets diminished.

A great example of this “quenching affect” as quantities increase is highlighted in the excellent book on this subject, “Numbers and Nerves” by Scott and Paul Slovic. They highlight an experiment from a University of Oregon risk perception experiment in which perceived brightness of a visual scene is demonstrated against a constant increase in luminance. Figure 1A-D highlights the experiment.

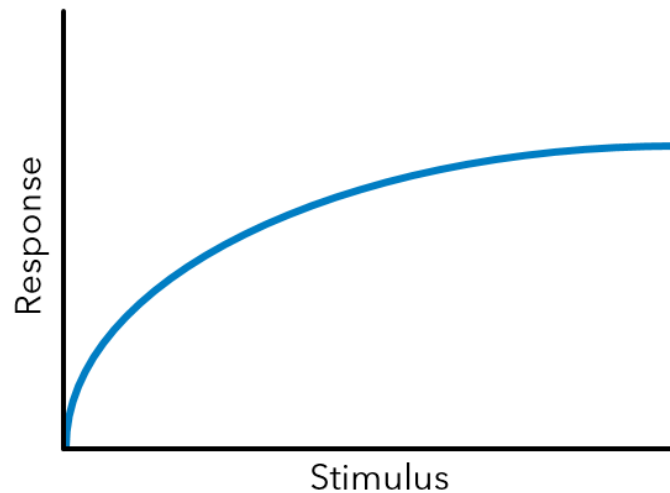


**Figure 1 (A-D) Stimulus-response demonstration, human psychophysics of perceived brightness.**

In Figure 1A, a single candle illuminates a multi-colored scene. The perceived brightness of the scene is commensurate with a single candle. In Figure 1B, an additional candle is added to the scene. Specifically, a constant amount of luminance, or light energy has been added to the scene. Our eyes and

the visual cortex in our brain see a significant increase in perceived brightness of the scene, commensurate with the increased luminance. More of the scene is now perceived. In Figure 1C however, there are 20 candles – the scene is quite bright and well-illuminated. In Figure 1D, a 21<sup>st</sup> candle is added. The perceived difference in brightness, or scene illumination between 1C and 1D is for the most part, imperceptible. Many of us might have a hard time choosing whether 1C or 1D contained more candles without explicitly counting them. As in the transition from 1A to 1B, the same amount of light energy – one candle – is added to the scene, yet there is a far greater perceived change in brightness from 1A to 1B than there is from 1C to 1D.

Why is this? The psychophysics of our perception for many different quantities is very much non-linear. That is, with constant increases in the magnitude of a stimulus (here, light energy from adding one candle), the stimulus evokes a smaller and smaller response (Figure 2).



**Figure 2. Stimulus-response curve for a range of perceived human input**

This nonlinear stimulus-response behavior is observed for a number of perceived quantities besides visual illumination – the same behavior is observed for the perception of loudness and heaviness, and even for non-sensory stimulus such as wealth.

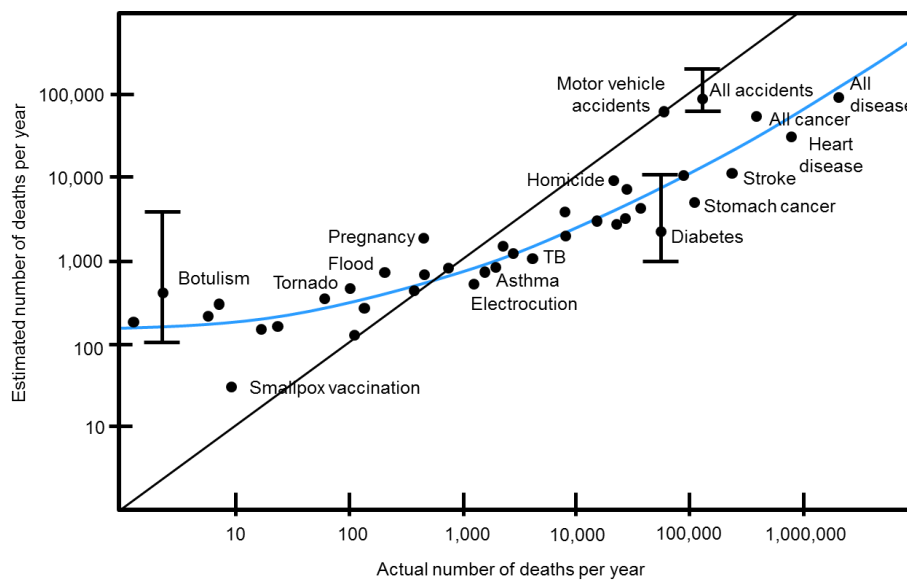
So how does this relate to decision making? It turns out that this same stimulus-response characteristic is observed in our response when the stimulus is reasoning about large numbers. Slovic and Slovic call this psychophysical numbing, citing numerous examples such as the salience, or “response” when we hear about the accidental death of a single child, as opposed to the reduced response when we hear about hundreds of thousands of people killed in genocide in an African nation. While we can certainly “do the math” with these large numbers, and logically calculate and communicate how much greater they are than a smaller set of numbers, there is a much reduced “response” to the stimulus of large numbers – we don’t intuitively reason about them as effectively, in proportion to their magnitudes, as we do smaller numbers. In government decision making about policy, strategy, and operations, this effect can quench or limit the ability of large numbers to have the intuitive impact that they “should” given their size.

We don’t need to look far to find government examples in which this psychophysical numbing can occur. Imagine U.S. government budget discussion in which quantities like 10 billion and 1 trillion are compared. We can certainly do the math and communicate that 1 trillion ( $10^{12}$ ) dollars is 100 times greater than 10 billion ( $10^{10}$ ) dollars, but do we really have an intuition for the relative impacts of these numbers in the same way we do in the case of 1 dollar versus 100 dollars? Likely not, and this impacts decision making when numbers become very large.

One of the numerous incidents that occurred while I served in the Department of Homeland Security was the Deepwater Horizon accident. Following that incident, numbers describing the incident arose - like 210 million gallons of oil spilled, 176,100 square kilometers of ocean surface impacted, 23 billion dollars in lost tourism, and 22,000 jobs lost. Many important issues require government to make use of the best possible statistical information – and decision making in terms of situational assessment and selection of

policy or operational alternatives are driven by this kind of information. But as the numbers get larger, human ability to intuitively reason about them commensurate with their true size begins to diminish. And this challenge isn't an issue only for large numbers.

In the 1960s and 1970s, nuclear power plants were being planned and built at increasing rates, and an important part of the public conversation was the level of risk from catastrophic accidents. Sophisticated probabilistic risk assessments were performed by the Nuclear Regulatory Commission and others, which came up with frequencies of catastrophic accidents in the neighborhood of  $10^{-4}$  and  $10^{-5}$ . But what would it mean to the public if the risk were  $10^{-5}$  versus, say,  $10^{-3}$ ? An understanding of risk perception was necessary to understand how these risks were perceived. Groundbreaking work from the 1970s by Slovic and Fischhoff demonstrated that our ability to intuitively reason about numbers on either end of the spectrum – large and small – is similarly limited. Figure 3, reproduced from a key paper from Slovic and Fischhoff show the results of an experiment in which humans were asked to estimate risks from a range of hazards, and then those estimation results were compared with the actual risks from those hazards.



**Figure 3. Actual versus human-estimated hazard risks. Hazard risk dynamic range compressed in human risk estimates as compared to actual risks. Adapted from Lichtenstein 1978, Figures 10, 11; and Slovic 1979, Figure 4.**

Each point in Figure 3 represents the actual number of deaths per year from a particular hazard (horizontal axis) and the human-estimated number of deaths per year from that hazard (vertical axis). If humans were perfect estimators of the actual risks from these hazards, all of the points would line up along the horizontal black line. However, this is not what was observed. In the neighborhood of around 100 to 1,000 actual deaths per year (a dynamic range of about a factor of 10) there is decent agreement between estimates and actual values for those hazards (that is, the points are near the diagonal line). But for actual death rates higher and lower than this neighborhood, the estimates deviate further and further from the actuals the further away we get, either smaller or larger. Fundamentally, this data showed that people are not very good at estimating or working with numbers that are very big or very small. Specifically, people tend to underestimate large risks, and overestimate small ones. This has the effect of taking what in actuality is quite a large dynamic range of deaths per year (ranging between about 1 and about 1,000,000 – a factor of a million) and compressing that into a much smaller range of between 100 and 100,000 – a factor of only 1,000. In their estimates, people reduced the dynamic range of the true numbers by a factor of 1,000.



There are a number of other studies that further bear out this phenomenon, but for the purposes of this discussion, we'll leave it with these examples, having briefly highlighted the first two of our three cognitive decision making limitations – judgment biases, and now dynamic range compression of quantitative information. I'll now turn to the last of the three areas of decision making limitations – irrationality under uncertainty.

## **HUMAN IRRATIONALITY UNDER UNCERTAINTY**

### **Expected Value Theory**

Prior to the early 1700s, it was believed for the most part, that rational decision making meant choosing the alternative, from amongst a set of alternatives, that had the highest *expected value*, where the expected value of a decision alternative is the statistical expectation of that alternative; that is, the probability of its outcome multiplied by the valuation of the outcome. A simple example might be the decision alternative that had a 25% chance of a \$100 payout. The expected value would then be  $0.25 \times 100$ , or \$25. Compared with a separate decision alternative having a 50% chance of an \$80 payout (an expectation value of  $0.5 \times \$80 = \$40$ ), expected value theory would dictate that rational choice would be the latter choice (50% chance of \$80), since it has the higher expectation value of \$40.

A number of mathematicians in the early 1700s observed however, that expected value theory did not fully account for the fact that in some cases, observed choices seemed to contradict expected value theory. This is due to the fact that many people are risk averse and will choose a more certain outcome even if it has a lower expected value. For example, consider a decision alternative with a 100% (certain) chance of delivering a \$100 payout (an expectation value of \$100). Then consider an alternative with a 12% chance of delivering a \$1000 payout (expectation value of \$120). While the latter alternative is the “better” choice according to expected value theory, since  $\$120 > \$100$ , many people are “risk averse” and choose the former.

### **Expected Utility Theory**

Bernoulli and others proposed that a mathematic correction should be made to the expected value to account for risk aversion – specifically a *utility function* that captures the utility of the decision maker, incorporating their level of risk aversion. With this correction made to expected value, it was then *expected utility* that was thought to be the deciding criteria for decision amongst alternatives when there is uncertainty. But some key things to note about expected utility theory, as it was called, was that it implied – like expected value theory – internally consistent decision making under the expected utility calculations. For particular decision makers with particular risk tolerances, decisions should be consistent. One specific form of this consistency was the (fairly obvious) substitution axiom that under expected utility theory, if alternative A is always preferred to B then any probability applied to those alternatives should not violate that preference. Said another way, if A is preferred to B, then a 25% chance of A should be preferred to a 25% chance of B. For about 250 years, expected utility theory was the dominant theory thought to best describe how people should (and did) make decisions.

### **Prospect Theory**

Despite the expected utility adjustments to expected value theory to account for risk tolerance, behavioral economists began to observe that the consistency axioms of expected utility theory did not always hold – specifically that preferences could be induced to reverse depending on how a problem was framed. First articulated by Allais in 1953, these violations of expected utility theory were difficult to explain until Daniel Kahneman and Amos Tversky in 1979 described a new theory of choice for which they later received a Nobel prize. The articulation of the problem by Allais, Kahneman, and Tversky is highlighted by the following two problem statements (adapted from Kahneman, 1979). Each problem states two choice alternatives, and about 100 people were asked which they would prefer. The number that appears after each choice alternative in [brackets] is the number of respondents that chose that alternative.

PROBLEM 1:

CHOICE A: 80% chance of \$4,000 [20]  
CHOICE B: \$3,000 for certain [80]

PROBLEM 2:

CHOICE C: 20% chance of \$4,000 [65]  
CHOICE D: 25% chance of \$3,000 [35]

Here, we don't need to know the risk tolerance of the respondents, or their utility functions – we've measured their actual choices so we know the outputs of those things for these problems. Looking at Problem 1, the majority preferred choice B. So, from expected utility theory, we can say that for these respondents, the utility of choice B compared to the utility of choice A must be described by the following inequality, otherwise respondents would not have chosen B:

$$u(3,000) > 0.8u(4,000) \tag{3}$$

Now, looking at Problem 2, we see that most respondents preferred Choice C over Choice D. So we can say that for this majority of respondents:

$$0.2u(4,000) > 0.25u(3,000) \tag{4}$$

Multiplying everything in Inequality (4) by four to facilitate comparison, we get the following inequality, which is equivalent to Inequality (4):

$$0.8u(4,000) > u(3,000) \tag{5}$$

Comparing Inequalities (3) and (5) we can see that for the same respondents with the same inherent utility evaluations of \$3,000 and \$4,000 we get a complete contradiction of preferences under the same utility functions, just from the difference in how problems 1 and 2 are framed. This preference inconsistency violates expected utility theory – it's easy to see that Inequalities (3) and (5) just don't make sense together.

Without going into as much explanatory detail, let's quickly look at the reverse problem – when we are talking about losses instead of gains:

PROBLEM 1' :

CHOICE A: 80% chance of losing \$4,000 [92]  
CHOICE B: Lose \$3,000 for certain [8]

PROBLEM 2' :

CHOICE C: 20% chance of losing \$4,000 [42]  
CHOICE D: 25% chance of losing \$3,000 [58]

Without working out the inequalities, we can see the same preference reversal with losses, but in the opposite directions - while we are risk averse when looking at gains (we'll take the sure bet over the gamble, even at higher expected utility), we are risk seeking when we are facing losses (we'll gamble to avoid the loss rather than the certainty of loss).

Kahneman's and Tversky's invalidation of expected utility theory as a descriptive model of decision making in favor of prospect theory does a much better job at describing how people actually evaluate choices and make decisions.

But what does this have to do with government decision making? It turns out that myriad decisions from geopolitics to health care to budget and legislative deliberations are well-described by prospect theory – risk averse with gains, and risk-seeking to avoid losses.

But while prospect theory does better at describing how decisions are actually made, for many decisions, we may argue that expected utility- or even expected value-based decisions may in fact be *better* decisions over the long run, since they quantitatively capture the known input information about probability and outcomes without the human decision biases and risk tolerance inconsistencies.

Having briefly reviewed three sources of limitations in human decision making – judgment biases, inability to handle large dynamic range quantitative information, and inconsistencies in how we make decisions uncertainty – in the final section, I will turn to a discussion of analytics as an important component of the decision making process to help us do better.

## **ANALYTICS AS A SOURCE OF EVIDENCE FOR IMPROVED DECISION MAKING**

To this point, I've devoted a significant amount of space to the survey of select aspects of cognitive limitations to our rationality as human decision makers. Even if we stopped here, this is already a valuable discussion for government decision makers, since simply *knowing* about these cognitive limitations can move decision makers closer to being able to combat them. Intelligent analysts might intentionally consider things like availability bias and confirmation bias as they receive information about threats to national security and synthesize that information into a model of reality. Law enforcement officers might think about their own propensity for anchoring bias, by seeking to avoid over-weighting the early information in an investigation at the expense of higher-quality information that could come in later as the investigation develops. When creating policy or planning for mission execution, senior government agency leaders dealing with very large numbers or very small risks might think about asking their staff for different representations for those numbers, scaling them for better intuitive comparison, or looking for analogies that facilitate better intuitive understanding of the numbers. When government decisions involve uncertainty, agency managers can frame problems in terms of both gains and losses to try and combat decision inconsistency due to problem framing.

But beyond these benefits that come from simply knowing about these decision making challenges, analytics can provide a significant amount of benefit in the form of evidence that can vastly improve the government's ability to make sound defensible decisions for good mission outcomes and efficient operations. But first, we need a good working definition of analytics. With analytics being one of many technology buzzwords floating around the field of information technology, there are quite a few definitions to choose from. For this discussion, I will choose the definition from the Institute For Operations Research and Management Science (INFORMS) – the international trade association for the field of analytics:

*"Analytics is defined as the scientific process of transforming data into insights for making better decisions."* (INFORMS, 2018)

I choose this definition first and foremost because it focuses on the most important job of any leader in government – decision making. Analytics is at its core, a decision support activity. The science must be right to effectively transform information into a form that can be used for decision making. But if not coupled to helping someone make a decision better, faster, or at lower cost, the best science conceivable will not have any real impact (we could spend a lot of time talking about how this fact – starting with the science and technology instead of the decision requirements explains why a number of government agencies find their analytics programs to have only limited value, but this is a topic for another paper).

Let's now look at four broad aspects of modern analytics, and briefly discuss how they can help serve as sources of evidence to improve decision making.

### **DATA MANAGEMENT**

In any organization, information can be fragmented and siloed in different departments, divisions, or offices. But because of the unique organizational structure of Government, with dilute and decentralized authority and decision making, this can be particularly true. Agencies may have data not only on different servers and in different parts of the organization; often these critical information sources are on individual computers as spreadsheets, or even on paper.

Thinking about the judgment biases that Kahneman and Tversky first outlined in 1973 and 1974, and the many others that have been articulated since, one of the attributes that a number of them have in common is the sequence or time-order in which information arrives. For example, anchoring bias occurs when someone provides additional weight to information that they receive or encounter first, even if challenged by higher-quality information later. Related to anchoring bias, confirmation bias means that we are less likely to incorporate or value information that challenges what we may already believe to be true – which can be a result of receiving information in a particular sequence in time.

Simply managing government data and information in a better way – even absent any real manipulation of that data at all – can make a significant impact on the quality of decision making by helping to manage the impacts of some of these biases. Relevant information for a problem or mission organized and presented in a relevant and comparison way can help people who have to make decisions in those missions areas properly weight and compare that data. Imagine a law enforcement officer or intelligence analyst receiving organized information about an investigation or threat integrated in one place, facilitating interaction with multiple sources of relevant data simultaneously. Objective indications of the information source quality can also be incorporated so that an analyst or decision maker not only has all the relevant source of information in one place, but has a built-in assessment of each source's uncertainty and quality.

A great example of the impact that simple data management can have on decision making for a government mission is North Carolina's Criminal Justice Law Enforcement Analytics Data Service or CJLEADS. Prior to CJLEADS, law enforcement and corrections in North Carolina had no easy way to see across the 90+ sources of information about suspects and offenders in the criminal justice system. Beyond just the time impacts to officials as they tried to search and synthesize all of this information, the consequences could have life-and-death impacts. In 2008, a North Carolina college student was murdered by two suspects who had appeared before a court several weeks prior. Had all relevant criminal justice information about these men been available to officials, it would have been clear that they should not have been released and able to commit additional crimes.

CJLEADS integrates more than 90 different data sources across the criminal justice system in North Carolina, ranging from sex offender data to department of motor vehicles information, even fishing and hunting license data. This system is available to every member of the law enforcement community in real time in the field, providing an integrated single view of an offender. Even with no "analytics" in terms of manipulating the data in any significant way, simply integrating the relevant information together in one place enhances decision making, reduced uncertainty, and helps avoid time-sequence judgment biases. North Carolina has measured these impacts, and has found that CJLEADS saves the state approximately \$13 million annually, and estimates that the improved knowledge field officers have about the people they are interacting with results in, on average, four officers' lives saved each year (McCoy 2013, page 33).


Returning to the opening example from this paper, the information integration failures prior to September 11<sup>th</sup>, 2001 are yet another historical example of the impacts of separate, siloed government data. An effective and integrated data management strategy, even absent any advanced analytics, can make a big difference for government agencies seeking to improve the quality of their decisions for better mission effectiveness.

## **DATA VISUALIZATION**

Assuming data is effectively managed and gathered into the same place for integrated access by decision makers, an additional way to help improve decisions is to effectively represent that data in a way that supports effective interpretation and meaning from the data for use in mission decision making. This paper will not seek to catalog good and bad methods for visualizing data – many resources have been produced that do an excellent job laying out when certain visualizations should be used as opposed to others. Rather, the focus here will be on the impact that effective visualization – whatever its form – can have on improving decision making.


As we've seen, one of the important challenges we have in decision making is our difficulty in interpreting very large and very small numbers. Effective data visualization can go a long way towards helping us do better at effectively internalizing and intuitively understanding these numbers in comparison to one another. The visualizations don't even necessarily need to be overly complex. Humans are very good at comparing linear distances, so simple bar charts can help. As a simple example, consider the problem

when a decision maker or analyst must compare, synthesize, or gain an intuitive understanding of the difference between a thousand and a million –  $10^3$  and  $10^6$  respectively. We can all plug those numbers into our computers or calculators and work with them easily, understanding how they compare mathematically, and even knowing that they differ by a factor of a thousand. But this isn't the same as an intuitive sense for what the numbers mean. Figure 4 provides a simple visualization on a linear scale of what these two numbers “look like.” Because our brains are good at comparing linear distances, we can quickly (and now *intuitively*) see that depending on the decision context, one number “matters,” and one number may not matter at all. While the mathematical computations we may have done in the spreadsheet are still as accurate as they were before, the additional value from visualization can help a great deal at combating our difficulties in handling large and small numbers.



**Figure 4. Simple bar chart, visualizing a large dynamic range;  $10^3$  and  $10^6$ .**

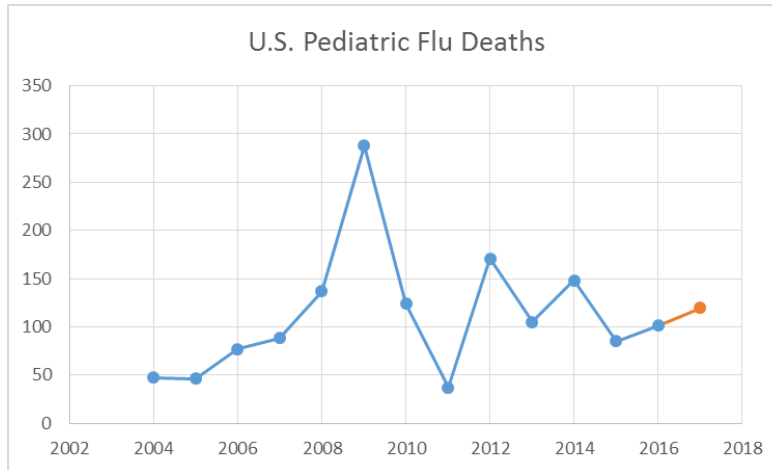
Another important way that data visualization can help improve decision making is by helping to resolve, or limit the impact of our inherent judgment biases. Recall that availability bias is the judgment bias in which something that is available or vivid, may receive additional weight in our decision making than it possibly should. Data visualization can help us properly ground and compare this “available” or vivid information to other evidence to help us keep it properly in perspective. Let’s take two examples to demonstrate the value of data visualization for managing availability bias. First, consider the very vivid example of fatalities in the United States from terrorism. Most of us remember September 11<sup>th</sup>, 2001, and we have pictures in our minds of those events. But how should we weight or balance terrorism against other hazards in terms of loss of life? Visualization can help by fixing the “available” and vivid information to other information that can help us weight it more properly. Figure 5 visualizes U.S. terrorism fatalities between 2000 and 2014 alongside annual deaths in the U.S. from automobile crashes and influenza.



**Figure 5. U.S. Fatalities: Automobile Crashes, Influenza, Terrorism, 2000-2014 (IIHS 2018, Statista 2018, START 2015)**

I am not trying to make the case here that all deaths should be counted with equal weight in policy making – society has shown that it tolerates different levels of fatalities from different causes, differently. Deaths from disease and other “natural” causes like cancer, are accepted at higher rates than supposedly “avoidable” deaths, like homicide and terrorism. However, attaching that “available” and vivid information – here, terrorism data – alongside other information can go a long way towards helping an analyst or decision maker more effectively weight that information, limiting the impact of availability bias.

Let’s look at one more example, this one more recent from the ongoing U.S. 2017-2018 influenza season. Watching the news, both locally and nationally, one may have thought that children were dying at a unprecedented rate from influenza. News report after news report seemed to feature families who had lost a child due to influenza. And this reporting was unique compared to previous years. This vivid information presented to us creates the risk of availability bias – we remember these terrible stories of children who have lost their lives, and it can cause us to potentially overweight this information if it’s not accurate, as governments seek to make sound vaccine and public health decisions. Figure 6 shows the influenza deaths for children over the past 14 years, including this year’s 2017-2018 season which appears to have peaked and is winding down.



**Figure 6. U.S. Influenza deaths, children, 2004-2005 through – 2017-2018 seasons (CDC FLUVIEW 2018).**

We can see that while the pediatric deaths from influenza from the 2017-2018 season (which still isn't complete) are higher than the most recent flu seasons immediate preceding this year, with the season winding down it looks likely that pediatric deaths from influenza appear to be most similar to the 2012-2013 and 2014-2015 seasons; certainly nowhere near the 2009-2010 season just 8 years ago. Again, I am in no way minimizing any tragic death of a child from influenza; rather that good visualization of the data can help us avoid judgment biases based on available or vivid information we may see in the media, absent contextual information.

Finally, in addition to helping limit impacts from judgment bias and improve how we interpret and internalize large and small numbers, visualizations can help us make better decisions under uncertainty as well. From our discussion about utility theory, recall that for a given decision maker and risk tolerance, depending on how a problem was framed, the preferences could be influenced to reverse, depending on whether or not we were talking about seeking gains or avoiding losses. Effective visualization of outcomes and probabilities can help analysts and decision makers maintain consistent problem framing, and “apples to apples” comparisons to help preferences remain close to what the preferences and utility functions (risk tolerances) of decision makers would dictate.

So we can see in this brief discussion that simple visualization of information can be of immense benefit in helping decision makers create improved intuition about large and small numbers, limit the impact of judgment biases, and improve consistency of preferences under uncertainty. Let's now turn to some aspects of analytics that manipulate data in more sophisticated ways.

## **ANOMALY DETECTION**

One of the most important things that any government organization needs to do with its data is to be able to mine it for anomalies. Finding the thing or things that “stick out” against a “normal” background of information, or the analogous task of being able to match government data against patterns of known anomalies is a critical element of many government missions across the entire spectrum of civilian and security agencies. With billions of dollars improperly leaving agencies in governments around the world, civilian agencies responsible for benefits delivery must make sure that benefits and services are delivered to citizens in a timely fashion, while at the same time detecting and preventing intentional fraud and accidental improper payment before funds leave the agency. With a multibillion dollar tax gap, revenue collection agencies at all levels of government must ensure that taxes are collected accurately, detecting and flagging potentially anomalous payments and returns, while at the same time not impeding law-abiding citizens from processing their accounts quickly. With the risk of collusion between external entities and internal agency staff, detecting out-of-the-ordinary behavior amongst agency staff is a critical part of maintaining government program integrity. Cybersecurity efforts in all agencies hinge on the ability to tell when something is happening to a network or information technology asset that is out of the ordinary. And finally law enforcement and national security personnel must be able to effectively detect outliers and

anomalies in their data that could reflect criminal or national security threats to our communities. And we could go on.

But how can analytics applied to the problem of detecting anomalies help make better decisions, and combat some of the limitations that we've articulated earlier in this paper? The answer lies with the way in which humans tend to look for information that differs from a background. In simple situations, anomaly detection can be pretty easy. In the case of a physical system, say a temperature reading on a nuclear reactor, the anomaly detection is easy, and can (theoretically) be done to high quality with good data management and data visualization – when the temperature exceeds a known threshold that is anomalous, decisions can be made given that the anomaly has occurred. (I say “theoretically” because even this simple case is not immune from judgment bias. There are many examples of nuclear power plant accidents, including Three Mile Island, which happened because the human didn't *believe* the anomaly detection systems when they said something was catastrophically wrong – classic cases of confirmation and other biases.)

But many, if not all of the anomalies our government agencies need to find are difficult, and not obvious. They are in highly multidimensional spaces, not just “temperature.” They are moving targets in that the anomalies are originating from thinking, adaptive adversaries like criminals, fraudsters, and hackers. As such, they are not amenable to detection via simple rules and thresholds alone. As humans, we tend to “look for our keys under the streetlight” – we search for anomalies and compare signals against data backgrounds in ways that make sense to us and which fit with our experience and background. This is certainly valuable – people with expertise and intuition about data sets are essential and necessary to finding the most important anomalies. But this is not sufficient. Looking for anomalies that only fit into what we can conceive of or understand, is not effectively making use of all available information, and subjects the detection and investigation tasks to significant confirmation bias and overconfidence, in which our detection and investigation units can sometimes convince themselves that: 1) they already understand most of the different ways mission-impacting anomalies can impact their agencies, and 2) that they are fairly certain they are capturing most of the anomalous behavior.

It is impossible to know for sure whether all fraudulent activity is being detected based on rules and thresholds set from experience. But new techniques, particularly machine learning and deep learning more specifically can provide a significant boost to agencies' abilities to increase their certainty that they are avoiding judgment bias, and truly assimilating and evaluating all possible evidence as it pertains to finding anomalies in highly complex agency data sets.

There are many excellent resources that explain the value of deep learning and machine learning for anomaly detection and pattern recognition; for the purposes of this paper, I will simply state that these methods greatly expand the aperture of the agency in its ability to see things that are out of the ordinary in their data. Conceptually, machine learning inverts the process of detecting anomalies from what has historically been done. Rather than starting with the humans, and establishing rules based on knowledge and experience, and then applying those rules to the data to look for signals, machine learning instead starts with the data, and uses specialized algorithms to ask the *data* – in all of its complexity – what the best rules should be for optimally discovering what sticks out from what's normal. This expands the ability of the agency to learn things it didn't know before, and possibly see things that just weren't possible without advanced methods looking across tens or even hundreds of dimensions of information.

While the ability for things like machine learning to help agencies do a better job avoiding judgment bias, and uncertainty challenges in our anomaly detection and decision making, it's important to note that there are important organizational and cultural implications to implementing this kind of approach. First, agency experts must understand that the tools and algorithms presented by machine learning and deep learning approaches do not purport to *replace* agency experts in the detection and investigation mission. Rather, they serve to arm these experts with even better information that further informs, sharpens, and augments their experience. When the first real explosion of deep learning came into the public conversation in 2015 with Google's deep learning system, AlphaGo, defeating a human player in the game of “Go,” the best players in the world did not feel threatened by the computer system. Rather, they studied the non-traditional and unexpected strategies and moves that AlphaGo made, and learned to be better players as a result themselves.



Second, it's important that an agency permits itself to have the status quo challenged. Machine learning and deep learning results may indicate new avenues for anomaly detection and investigation that an agency has not previously conceived or pursued. These potential new directions must of course be checked and validated, and not blindly followed, but if an agency is not willing to learn new methods or patterns, then any investment in technology and analytics to provide evidence, will amount to wasted resources. The technology is important, but of equal (or even greater) importance are the agency's people, processes, organization, and culture.

## OPTIMIZATION

The last area of analytics we'll address that can provide evidence in decision making and help overcome judgment biases and other limits is optimization. Optimization applied to government problems is not new, as it first was used as part of Operations Research in World War II to make effective evidence-informed tactical decisions on the battlefield. But here, I want to briefly point out that this technique – designed for problems that have constrained resources (which is pretty much every single government problem) – can greatly help avoid the decision limitations that come from all three of the challenge areas we've discussed - judgment biases, large dynamic range quantitative information, and preference reversals based on problem framing.

When a decision maker faces a portfolio of possible decision alternatives, and constrained resources are involved – time, money, personnel, equipment, and/or others – optimization can help the decision maker frame their priorities across those constraints, and then provide the decision maker with analytically-derived “optimized” solutions in priority order. This can be an invaluable source of evidence for the decision maker alongside other information she must consider, in that it mixes the data together for the decision maker free from comparative bias, effectively manages the magnitudes of the input quantities, regardless of their dynamic range, and maintains consistent problem framing as it integrates constraints and problem features on the way to proposing a ranked list of optimal solutions or decisions.

While optimization is part and parcel in many industries, such as manufacturing or retail, it is something of an unsung hero in government. Given the high-dimensional constraints in government decisions, and the difficulty in marshaling evidence bases for them, it is an important tool alongside the others we've mentioned in this paper to help improve the quality of decisions.

## CONCLUSION

In this paper, I have explored three important cognitive challenges in human decision making – judgment biases, difficulties intuitively understanding large dynamic range quantitative information, and difficulties making consistent decisions under uncertainty. We've looked specifically at government examples for each of these, and shown that these decision making challenges have real impacts in government decision making today. We then turned to specific examples about how analytics can serve as a critical source of evidence to limit the impact of these cognitive decision challenges, and can go a long way towards helping our government colleagues make the best possible decisions to give the best possible outcomes for the citizens they serve. The few limited techniques I've discussed here are by no means an exhaustive evaluation of all possible analytics approaches that could improve decision making in government; however they are an excellent place for any government agency leader to begin to start innovating government decision making with better evidence.

## REFERENCES

1. National Commission on Terrorist Attacks Upon the United States. 2004. *The 9/11 Commission Report*. Available at: <http://govinfo.library.unt.edu/911/report/911Report.pdf>.
2. Schwartz, N, et al. 1991. “Ease of Retrieval as Information: Another Look at the Availability Heuristic.” *Journal of Personality and Social Psychology*, Vol. 61, No. 2, 195-202.
3. Slovic, S., and Slovic P. (eds). 2015. *Numbers and Nerves: Information, Emotion, and Meaning in a World of Data*. Corvallis, OR. Oregon State University Press.

4. Tversky, A, and Kahneman, D. 1973. "Availability: A heuristic for judging frequency and probability." *Cognitive Psychology*, Vol. 5, No. 2, 207-232.
5. Tversky, A, and Kahneman, D. 1974. "Judgment under Uncertainty: Heuristics and Biases." *Science*, Vol. 185, No. 4157, 1124-1131.
6. Baron, J. *et al.* 1988. "Outcome Bias in Decision Evaluation." *Journal of Personality and Social Psychology*, Vol. 54, No. 4, 569-579.
7. Hilbert, M. 2012. "Toward a synthesis of cognitive biases: How noisy information processing can bias human decision making." *Psychological Bulletin*, Vol. 138, No.2, 211-237.
8. Oswald, M. E., Grosjean, S. 2004. "Confirmation Bias". In Pohl, Rüdiger F. *Cognitive Illusions: A Handbook on Fallacies and Biases in Thinking, Judgement and Memory*. Hove, UK: Psychology Press. pp. 79–96.
9. Shermer M. 2014. "How the Survivor Bias Distorts Reality." *Scientific American*, Sept 1 2014.
10. Slovic, P., *et al.* 1979. "Rating the Risks." *Environment*, Vol. 21, No. 3, 14-39.
11. Lichtenstien, S., *et al.* 1978. "Judged Frequency of Lethal Events." *Journal of Experimental Psychology: Human Learning and Memory*, No. 4, 551-578.
12. Allais, M. 1953. "Le Comportement de l'Homme Rationnel devant le Risque, Critique des Postulats et Axiomes de l'Ecole Americaine." *Econometrica*, No. 21, 503-546.
13. Kahneman, D., and Tversky, A. 1979. "Prospect Theory: An Analysis of Decision Under Risk." *Econometrica*, Vol. 47, No. 2.
14. Tversky, A., and Kahneman, D. 1992. "Advances in prospect theory: Cumulative representation of uncertainty." *Journal of Risk and Uncertainty*, Vol. 5, No 4, 297-323.
15. INFORMS. 2018. "Operations Research and Analytics." Institute for Operations Research and Management Science (INFORMS). 2018. Available at: <https://www.informs.org/Explore/Operations-Research-Analytics>.
16. McCoy, David T. "North Carolina Government Data Analytics Center Program." North Carolina Office of the State Controller. 2013. Available at: <https://files.nc.gov/ncdit/documents/files/GDAC-Legislative-Report-October-2013.pdf>.
17. "Overview of Fatality Facts." International Institute for Highway Safety. 2018. Available at: <http://www.iihs.org/iihs/topics/t/general-statistics/fatalityfacts/overview-of-fatality-facts>.
18. "Deaths by influenza and pneumonia in the U.S. from 1950 to 2015." Statista. 2018. Available at: <https://www.statista.com/statistics/184574/deaths-by-influenza-and-pneumonia-in-the-us-since-1950/>.
19. "American Deaths in Terrorist Attacks." National Consortium for the Study of Terrorism and Responses to Terrorism, University of Maryland. 2015. Available at: [https://www.start.umd.edu/pubs/START\\_AmericanTerrorismDeaths\\_FactSheet\\_Oct2015.pdf](https://www.start.umd.edu/pubs/START_AmericanTerrorismDeaths_FactSheet_Oct2015.pdf).
20. "Influenza-associated Pediatric Mortality." Centers for Disease Control and Prevention, FLUVIEW. 2018. Available at: <https://gis.cdc.gov/GRASP/Fluview/PedFluDeath.html>.

## CONTACT INFORMATION

Your comments and questions are valued and encouraged. Contact the author at:

Steve Bennett, Ph.D., Director, Global Government Practice  
 100 SAS Campus Drive  
 Cary, NC 27513  
[steve.bennett@sas.com](mailto:steve.bennett@sas.com)  
 Twitter: @AnalyticSteve  
 LinkedIn: [www.linkedin.com/in/AnalyticSteve/](http://www.linkedin.com/in/AnalyticSteve/)

SAS and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries. ® indicates USA registration.

Other brand and product names are trademarks of their respective companies.