Chapter 12

Judgment Heuristics, Deduction, & Induction

By Dr. Charles Wallis

Last Revision: 4/28/2020

Chapter Outline

12.1 Review of Basics of Human Inference Abilities

Chapter 5 introduces a general characterization of human inference abilities and their typical deployment. After that general introduction Chapter 5 focuses primarily upon deductive inferences and human deductive inference abilities. This chapter's focus revolves around inductive inferences and a few of the System 1 inference heuristics thought to shape human inference abilities with an eye to understanding the inferential strategies of these heuristics as well as their potential biases and limitations. However, before narrowing its focus to inductive inferences this chapter reviews some of the main ideas from the first inference chapter—Chapter 5. The first eight sections of this chapter review the basic framework introduced in the Chapter 5. Students may feel tempted to skip the review sections. Nevertheless, I encourage students to review this material. A thorough understanding of human inference abilities, inabilities, and biases it will facilitate comprehension of new material. Moreover, a thorough understanding of human inference abilities, inabilities, and biases is a central pillar upon which students can build more complete, consistent, and systematic worldviews as well as utilize those worldviews to craft more optimal inferences and decisions.

12.2 Characterizing Inferences

Beginning with Chapter 1 and continuing through the current chapter this text and course emphasize the distinction between innate human inference and decision-making processes and the formal artifacts created by thinkers throughout the last three millennia. These inferential and decision-making artifacts often function as bug fixes and workarounds for those problems and/or those circumstances for which native human reasoning processes prove less than optimal. Understanding the difference in strategies as well as costs and benefits between native reasoning processes and those formal artifacts gives students a chance to anticipate potential pitfalls and to adopt, at least sometimes, more optimal approaches to inference and decision problems. Therefore, this chapter begins by reviewing that basic distinction between inferences and formal inference systems.

12.2.a The Pervasiveness of Inferences

Chapter 5 argues that inferences pervade every aspect of every moment of every day. For instance, when people play basketball their visual systems processes reflected and projected light to identify the basket and its position in three-dimensional space. These visual inferences occur completely outside of conscious awareness; only the results of these inferences can enter consciousness through visual working memory. When people shoot a basket, in contrast, they engage in semi-conscious inferences. They are conscious of their intent to shoot a basket and the position of the basket. They might even think about aspects of their movements consciously. But, their brains coordinate their movements using vision without their ever consciously planning out those movements. Indeed, the extent of our unconscious motor planning becomes evident in the video (left) when basketball players don prismatic goggles that shift their vision thirty degrees to the right. The players suddenly can't sink a basket--their unconscious inferences breakdown. These players must shift their movement more strongly towards conscious control in order to override their unconscious visual processing until their brain engages in visual adaptation. When the goggles come off, the problem emerges all over again. Of course, we also engage in conscious inferences. For instance, we may consciously plan our day ahead of time to allow for some time on the basketball court. Despite the ubiquitous nature of inferences, few people can readily answer three basic questions regarding inferences. (1) What are inferences? (2) What are the functions of inferences? (3) What are

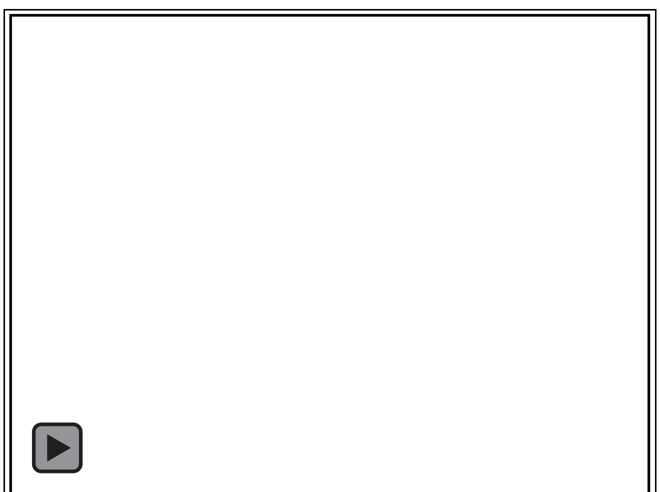

Dan Simons gives NYC basketball league players prismatic goggles that shift the visual scene thirty degrees to the right.

the goals of inferences? Hopefully, this text and course provides answers that have found their way into your worldview. The next sections ask and answer these three fundamental questions.

12.2.b What Are Inferences?

Recall that the inferences chapter characterizes inferences as psychological processes that transform information. Specifically, inferences transform available and explicit information to create new available and explicit information. When information gets encoded in a fashion that allows our brains to use that information for the task at hand, the information becomes explicit. When an inferential process has access to information, the information is available. An analogy might help make these ideas clearer. Imagine that you receive a check from your grandmother and you go to deposit it in the bank. The bank informs you that your funds will be available in two days. In other words, it may be your money, but you don't have access to it at the moment. When an inferential process cannot access information, say you can't remember the conjunction rule for to a quiz question, that information is unavailable in much the same way your deposit is unavailable. Now, suppose that you go to the bank after two days and withdraw half of the money from your grandmother's check. However, the bank teller gives you euros instead of dollars. You have your money, but it is not in a form you can spend. Your money is the financial equivalent of implicit information. When you complain, the teller exchanges the euros for an equivalent sum in dollars. Now your money has become financially explicit, you are free to spend it at the store. When information gets encoded in a manner that an inference process cannot utilize, that information is implicit information. To illustrate the difference between implicit and explicit information consider the following puzzle:

> To get to school Olivia drives west 5 miles turns left and drives 15 miles turns left and drives 25 miles then turns right and drives 5 miles. What's the shortest distance between Olivia and her school and in what direction is her school relative to her house?

The Olivia vignette provides all the information necessary to discover the answer. However, the information remains implicit in the vignette. The movie (below) presents the information in a fashion that explicitly represents the information needed to determine the solution. The graphic representation of the Olivia problem, in contrast to the vignette representation, makes the solution explicit. If you know the Pythagorean Theorem and use a drawing to make the route explicit, you should be able to answer this question.

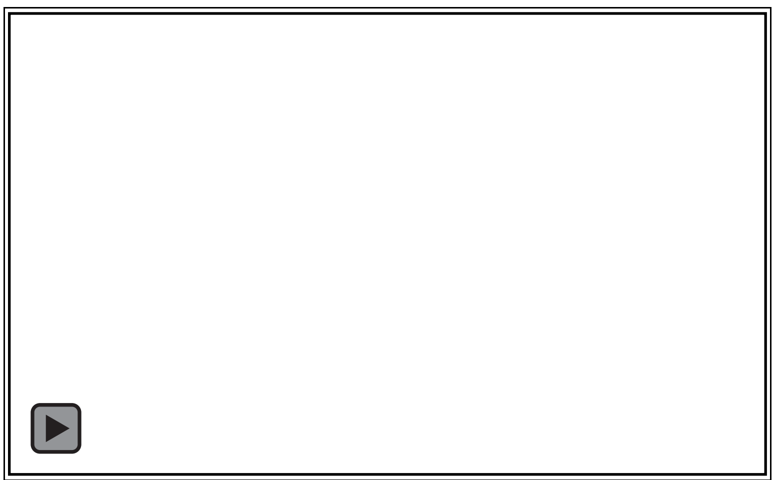

Movie showing how making the information in the vignette explicit in a diagram makes solving the problem much easier.

Similarly, the two different ways of writing *pi*, π and 3.14159…, make different information explicit. The symbol π refers to the mathematical constant, the value of which is determined by the ratio of a circle's circumference to its diameter. Using the symbol, π, allows one to explicitly refer to that constant in a manner that allows one to, for example, write the equation for the area of a circle: $A = \pi r^2$. However, the symbol does not make the value of that constant explicit. As a result, one cannot calculate the area of a circle unless one uses the decimal approximation of *pi*. The decimal approximation makes the value of the constant (partially) explicit. The table below illustrates the relationship between explicit and available information:

	Type of Information			
	Explicitly Encoded		Implicitly Encoded	Not Encoded
Type of Accessibility	Immediate	Inaccessible	Inaccessible	Inaccessible
Availability Status	Available	Unavailable	Unavailable	Unavailable

One can usefully characterize inferences as processes transforming information encoded in one's brain. In order to make an inference that information must be explicit and available to the inferential process. Neither the inferential process nor the information need enter conscious awareness. You remain unaware of the inferences made by, for instance, your vestibular system. Rather, the vestibular system needs access to the information from the inner ear (availability) and the information must be encoded in a manner that allows it to exploit that information (explicit) to determine your spatial orientation and bodily balance. Thus, inferences occur at many levels; unconsciously as with visual recognition, semi-consciously as with inferences about shooting baskets, as well as explicitly and consciously as when solving math problems or planning one's day. Available and explicit information for these inferences consists in accessible information encoded in a usable manner. In short, inferences of all varieties consist in information transforming psychological processes that take explicit, available information--accessible information encoded so that a process can utilize it—and transform it into new explicit and available information.

12.2.c What are the Functions of Inferences?

Why are inferences so important? Simply put inferences provide a mechanism whereby cognizers can adapt in a highly complex and variable world. **Inferences generate new explicit and available information or adapt existing information to novel circumstances. Inferences allow us to evaluate and integrate new information into our belief systems and worldview. They help people to adapt to the world by transforming information into useful formats--making implicit information explicit. Inferences even allow one to discover inaccurate information, poorly evinced information, and inconsistent information so that one can discard that bad information and/or render one's worldview more consistent.** The positional multiplication technique provides a simple illustration of inference as a process transforming explicit and available information. For instance, one might have explicit and available information for each multiplicand. One simply transforms that information to create an explicit and available representation of their product. Thus, a process whereby one transforms Hindu–Arabic numerals representing different multiplicands allows school children to solve multiplication problems with ease. Indeed, people use the positional technique to solve a variety of problems, even to detect and correct inconsistencies in, for example, their check book. Notice that the positional technique works well for Hindu–Arabic numerals; Roman numerals have fewer primitives and a different positional system [LV = 55, VL = 45] so that the relevant quantitative information, while available is not explicit [in usable form] for the positional technique. Thus, the positional technique does not work for Roman numerals.

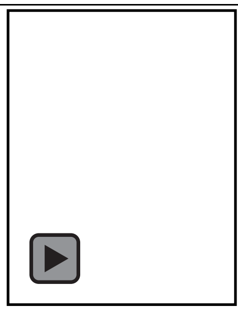

Multiplication as an inference that transforms the explicit and available information about each multiplicand into explicit and available information about the product. Click on the diagram to see the animation.

12.2.d What are the Goals of Inference

One might suppose that inferences have a single, obvious goal—truth. True, accurate, or veridical information can guide one's interactions with the world in a manner that respects the world's actual structure. If you infer that you should urinate on your friend's jellyfish sting or you should take off your red coat so that the bull won't charge, things will not go well. Experts recommend removing the jellyfish stingers and washing the area with vinegar or sea water. Urine is mostly water and can cause the stingers to inject their poison. Bulls are color blind; the bull will likely perceive the motion of taking off your jacket as a threat. Consistent with the focus of this chapter, inferring true beliefs in a complex and highly variable world often represents an unattainable goal. In many, many cases inferring a highly probable belief on the basis of good evidence represents the only achievable goal. Thus, if you infer that hydroxychloroquine plus azithromycin "have a real chance to be one of the biggest game changers in the history of medicine," you may find that you have diverted resources and risked lives during a pandemic based upon anecdotal reports and a French study described by experts "as a complete failure."[1, 2]

Does truth or high likelihood exhaust the potential goals of inferences? No. A much more diverse set of inference properties affect the relative merit of potential inference strategies. For instance, inferential power proves very desirable in inferences. Fans of the works of Sir Arthur Conan Doyle might recall Sherlock Holmes' amazing ability to make unobvious inferential extrapolations from seemingly irrelevant or miniscule bits of information. Such powerful inferences--inferences greatly extending one's initial knowledge--prove both necessary and desirable in everyday life. Every time one utilizes information from past experiences to guide present actions, one uses powerful inferences to extend one's knowledge of the past into the future.

Similarly, one's brain makes visual inferences and recognizes objects in about 300ms (1/3 of a second). This incredible speed proves important since it allows, for instance, the quick identification of threats. Speed, therefore, represents another feature of inferences that one might want to optimize. Speed, power, and truth…all have value as features of a given inference strategy. However, in practice one must usually trade-off strength in one of these features for gains in another feature. Processes like vision--fast, reliable, and powerful—prove the exception in human inference rather than the rule.

For example, suppose that a computer science student wants to create a chess playing program that always ties or wins the games it plays. One inference strategy that might seem initially promising would generate every possible permutation of every possible move after the initial move. At each turn, the computer program would then choose its move from all those possible games. Since the computer now has generated explicit and available representations of how all the possible games will end, it can choose only those moves that would end in a win or tie. Such a program would represent a powerful and highly reliable inference strategy. However, no computer yet built would prove sufficiently fast to execute such a program. Thus, the "generate-all-possible games" strategy represents a non-viable strategy for accomplishing the computer student's chess-playing goals. Specifically, the average chess game has approximately 40 moves per player. For each player's turn, the number of possible moves equals all of the moves that the rules of chess allow. Each move, likewise, allows for a large number of possible counter-moves —especially at the beginning of the game. An American computer scientist and cryptographer named Claude Shannon[3] (1916-2001) proved that in a single chess game, the average number of possible combinations of moves involves 10^{120} possible moves. This number of possible moves, and hence possible games, now bears the name the Shannon Number.[4] The Shannon number poses a problem for the computer science student; 10^{120} moves means that the number of possible moves in every permutation of an average chess game exceeds the number of seconds since the big bang. The student's program plays wonderful theoretical chess, but would prove impossibly slow for real use.[5-7]

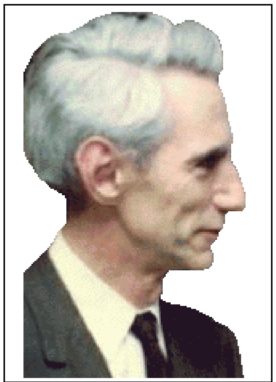

Claude Shannon (1916-2001)
Adapted from:
netzspannung.org

Thus, the fourth important potential property of an inference strategy is tractability—the potential to complete the inference in a reasonable amount of time (or even at all) utilizing only the available resources. In order to survive and especially to thrive humans need to solve the problems that confront them. As the discussions of various inferences and inference strategies unfold in the chapters, one theme that appears time and time again is that inference strategies almost always represent some trade-off between truth preservation, inferential power, speed, and/or tractability. As a result, all inference strategies have strengths and weaknesses—costs and benefits.

12.3 Innate Reasoning Abilities: Origins and Elements

To understand human inference abilities one must first understand the origins of those abilities. Indeed, human and proto-human evolution determined two central elements of human inference abilities—the human brain and the innate strategies the brain employs to make the vast majority of inferences. By understanding the origins of human inference abilities one can understand the forces that shaped both the brain's inference capacities and the innate strategies that drive the majority of human inferences. Such an understanding of the human brain's inference capacities and strategies allows one to recognize the strengths and weaknesses of innate human inference abilities. I begin this section by discussing the origins of humans and proto-humans (called Hominini by scientists[8]). The general, long-term environmental features and inferential challenges during Hominini evolution have shaped both the human brain and the inference strategies that modern humans employ to solve problems in the contemporary world.

12.3.a The Evolutionary Origins of Native Human Inference Abilities

So, how did human inference abilities evolve and what assumptions do they embody? Most inference abilities probably evolved during the hunter-gatherer phase of Hominini (human and proto-human) evolution. Scientists now theorize that this period of Hominini existence lasted for approximately 4.4 million years to 7 million years. The exact period depends upon which of the candidate fossil species one includes as Hominini and which species fall into the common human/ape lineage. If one includes proto-humans like the hominoid recently discovered in current-day Ethiopia named "Ardi" (*Ardipithecus ramidus*), the period extends to about 4.4 million years.[9-23] If one includes the fossil species *Sahelanthropus tchadensis*, then the period extends to over 7 million years.[19] The fossil record of Hominini provides good evidence that they engaged in subsistence foraging and hunting. For instance, ample evidence exists characterizing the lives of *Homo erctus* 1.8 million years ago as well as *Homo sapiens* starting around 200,000 years ago as surviving by subsistence foraging and hunting.[20, 21] The hunter-gatherer existence represents the exclusive mode of human existence until a mere 10,000 to 12,000 years ago when the Mesolithic era ended.[17] The Neolithic Revolution marks the end of the Mesolithic era and signals the advent and slow spread of anatomically modern humans who domesticate animals, develop agriculture, and live in larger, relatively permanent groups.[22, 23] In other words, scientists suppose that beginning approximately 10,000-12,000 years ago human existence begins to undergo a metamorphosis during which increasing numbers of humans stop living primarily by subsistence foraging and hunting and start living by growing food as well as breeding and raising animals. This change in how humans go about getting food brings with it important changes in all aspects of how humans live. For instance, the introduction of agriculture results in humans living in larger groups and creating relatively permanent settlements.

Scientists currently hypothesize that human languages develop during the Paleolithic Era approximately 100,000 to 50,000 years ago.[24-30] Proto-written language does not develop until approximately 8600 years ago. Alphabetic languages date to approximately 3100 years ago. Thus, the advantages of language—the ability to externalize memory and to share relatively complex and large amounts of information between individuals and across time--likely do not play a major role in shaping the human brain and inference abilities. Written language dramatically impacts human thriving, but it emerges far too recently to affect human evolution. One might find this conclusion unintuitive given the integral role that language—spoken and written—plays in contemporary life. Nevertheless, scientific research seems relatively homogeneous and substantial in support of the limited period during which language shapes human thinking.

In the hunter-gatherer era humans make inferences about, for instance, the likelihood and/or relative incidence of objects, properties, and events just as we do today. However, their environment differs from our own. Hunter-gatherers have short lives and few tools or other artifacts. Hunter-gathers live in small groups relatively isolated from most other proto-humans. As a result, the environment in which they solve problems proves relatively small. With no means of travel besides walking, most Hominini likely travel only 30 miles or so from their birthplace. Though major changes occur during the 4.4 million years of Hominini hunter-gather existence—ice ages, for example—most humans do not live long enough to experience much change. The mean hunter-gatherer lifespan is probably 21-37 years.[31] Approximately 60% live past 15, and of those who live past 15 approximately 60% live to 45 (between 23% and 43% total). Since Hominini have little technological development and short lives their environment proves pretty stable during their lifetime. In similar fashion, the combination of a small environment and a stable environment means that a hunter-gatherer solves problems in a relatively homogeneous environment. That is, things do not vary much from one part of their environment to another or even during the course of their relatively short lives. Thus, researchers characterize the environment in which individual humans and proto-humans solve problems for something like 7 to 4.4 million years as relatively small, stable, and homogenous. Since the environment remains stable and homogenous

		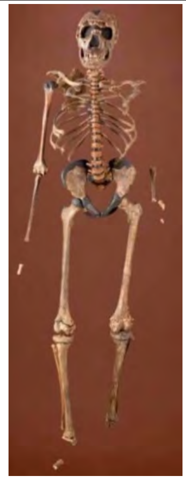	
Recovered skeleton of *Ardipithecus ramidus*. Image from: [Wikipedia]()	Artist rendering of *Ardipithecus ramidus*. Image from: [Wikipedia]()	*Homo erctus* skeleton. Image from: [Smithsonian Science]()	*Homo erctus* statue at the Smithsonian museum. Image from: [Smithsonian Science]()

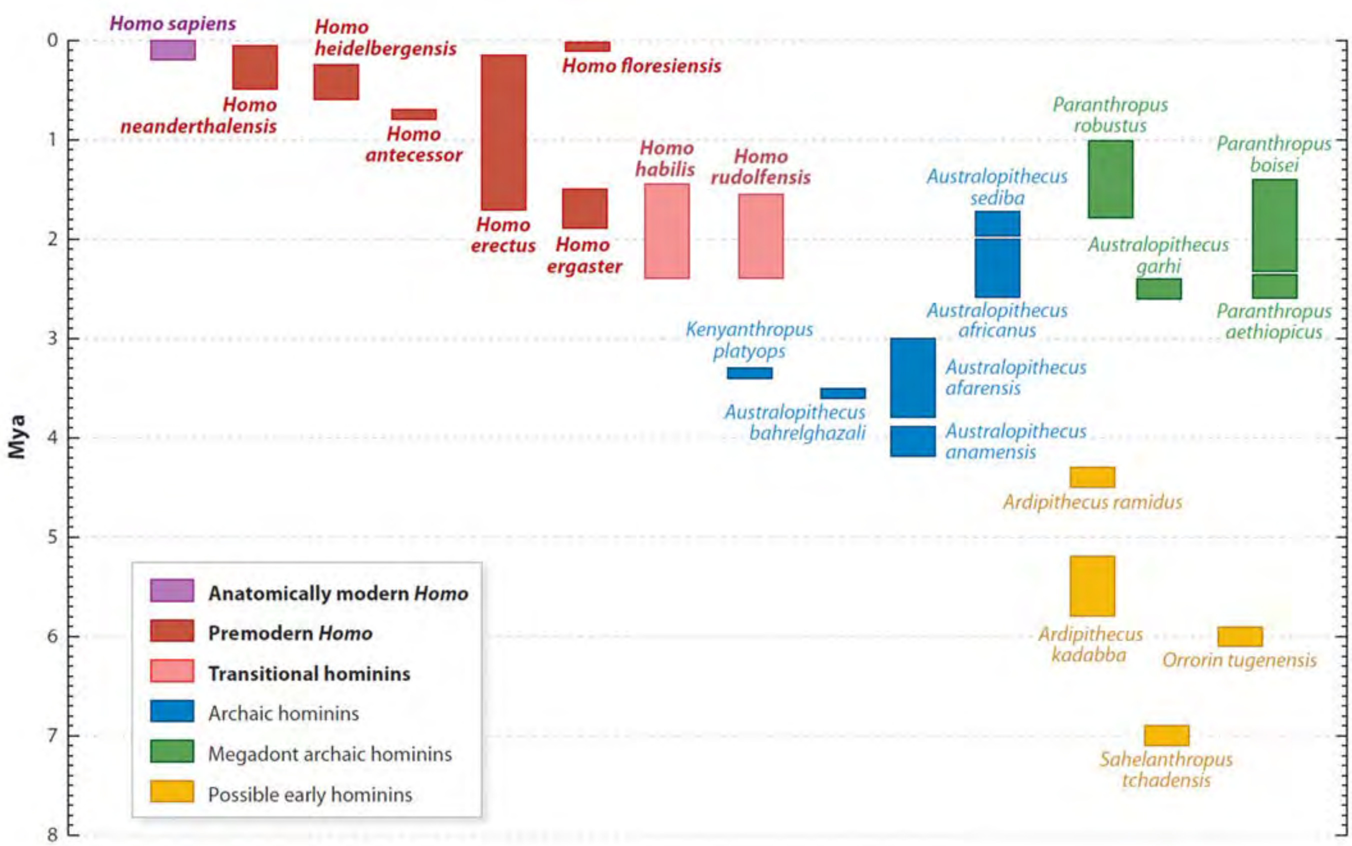

Diagram illustrating the various fossil specimens classified by their relationship to/within the Hominin taxonomy and the era during which scientists suppose that they flourished. From: Wood's and Baker's *Evolution in the Genus Homo*.[11]

during an individual hunter-gatherer's lifetime, their inferences are largely reactions to specific problems at hand. This combination of problem-solving environment and problem types limits hunter-gather problem-solving to largely reactive, relatively simple, and concrete problem-solving linked to specific contents (problems) and contexts (situations). Hunter-gatherers spend a great deal of time identifying features of their environment and responding to those features all within a relatively short time-frame. Additionally, hunter-gatherers also live in small social groups. As social creatures hunter-gatherers must solve problems having to do with whether their social hierarchy permits certain actions and whether their social hierarchy obligates them to perform (or refrain from performing) certain actions. Finally, in their more abstract and long-term inferences, like estimations of the likelihood of events, hunter-gatherers can rely upon their experiences as representative samples of their environment. In a relatively small, stable, and homogenous environment an individual human's experiences provide pretty accurate samplings of the overall environment.

12.4 Two Elements of Inference Ability: The Brain's Architecture and Its Inference Strategies

Though it might seem counterintuitive to contemporary humans, the human brain—like the brains of vertebrates generally—evolved to optimize problem-solving and decision-making of a reactive and rather immediate nature. One design choice selected by evolution to optimize performance in such circumstances utilizes specialized brain systems to quickly and reliably gather information from the environment with relatively little conscious input. Thus, humans can recognize an object very quickly without much conscious effort. However, the brain also employs a second design choice of a quite different nature—conscious inferences employing working memory. The remainder of this chapter and lecture will discuss these two strategies, when they succeed, when they fail, how and when they collaborate, and when they fail to interact with each other.

12.5 The Brain: Conscious vs Unconscious Inference

Each of the brain's native strategies has its strengths and weaknesses. Unconscious inference strategies operate automatically, quickly, often handling larger and more complex bodies of information, and they tend to be robust. However, these strategies can also prove inflexible, especially when the problem violates one or more of the assumptions implicitly driving the inferences. In contrast, conscious inferences tend to be much more flexible and adaptive. But conscious inferences prove resource intensive and can only handle a very limited amount of relatively simple information. Unlike unconscious inferences, which occur throughout the brain, conscious inferences as well as conscious components of semi-conscious inferences occur in working memory. It makes sense, therefore, to discuss what psychologists and neuroscientists currently know about working memory. However, before turning to working memory, the chapter and lecture discuss the relationship between conscious and unconscious inference. The discussion highlights the relative numbers and complexity of inferences performed unconsciously vs consciously. Students likely exhibit a common bias towards supposing that most inferences occur consciously. However, as the next section indicates, most inferences—especially complex inferences occur outside out conscious awareness. Working memory provides humans with conscious access to the final products of many of these unconscious processes, but it rarely captures more than a tiny portion of the inference or the information involved in the inference.

12.5.a Most Inferences are Made Unconsciously

The idea that conscious inferences constitute a miniscule portion of the inferential life of the brain and the information processed by the brain, strikes many students as contradicting their lived experiences. So, some illustrations seem in order. By now early vision is a familiar example in the text and lectures. Let us start there. None of the information or inferences discussed above in the processing of early vision enters working memory or consciousness until a small portion of the final products become accessible through working memory. For instance, humans have absolutely no conscious access to the initial light-intensity information collected by 120 million photosensitive receptors in each eye, nor can working memory access the inferences and information that occur in the eye, the lateral geniculate nucleus, and the striate cortex. Only when visual information enters into the parietal and temporal cortexes can elements of the visual scene potentially enter into consciousness; even then, only a very, very small percentage of that information actually enters into working memory at any given moment. In

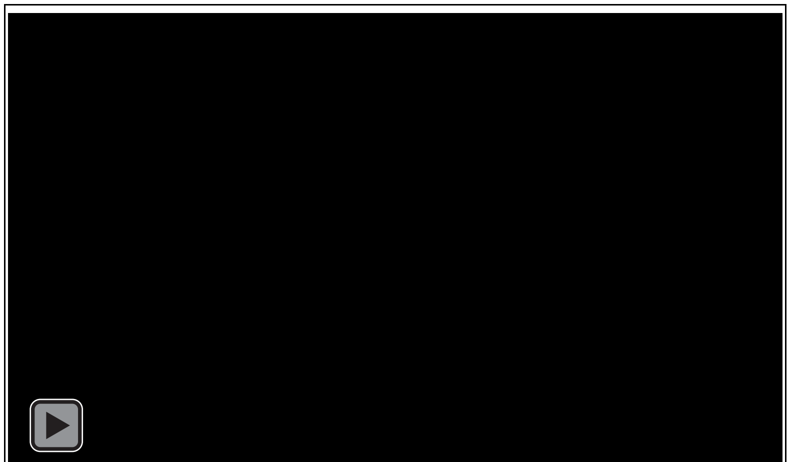

Movie illustrating how our conscious visual awareness is dependent upon our limited visual attention. From: YouTube order for even that small bit of the processed visual information to enter consciousness a person must focus their attention upon it. For instance, the movie (right above) illustrates the dependence of conscious visual awareness upon our limited visual attention. Trying to process images of movie stars without attending to them results in extremely distorted images. Lest students think that visual processing proves the exception to the rule, consider the words of researchers John Bargh and Tanya Chartrand from their aptly named article, *The Unbearable Automaticity of Being*:[32]

> Our thesis here—that most of a person's everyday life is determined not by their conscious intentions and deliberate choices but by mental processes that are put into motion by features of the environment and that operate outside of conscious awareness and guidance—is a difficult one for people to accept. (p.462)

For instance, facial characteristics like pupil dilation, averageness (mean values) of features, symmetry of features, skin color, skin texture, as well as gender-specific dimorphisms (two forms distinct in structure within a single species) heavily influence judgments of attractiveness despite typically playing no role in conscious explanations of facial attractiveness.[33-42] Additionally, situational and idiosyncratic factors like familiarity, dissimilarities (the degree of dissimilarities in immune responses that can prove compatible in an individual imprinting during development, hormone levels, fertility cycles in women, major histocompatibility complex resulting from reproduction), peer evaluations, self-perceptions (of attractiveness and personality characteristics), social status, and social learning all modulate impact of physical facial features without being included in people's conscious explanations of facial attractiveness.[38, 39, 43-48]

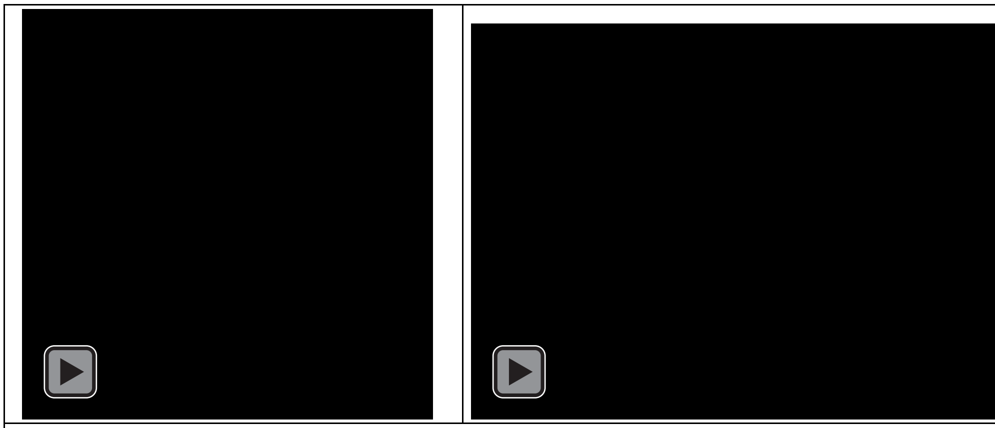

Many people hear a sound when the pylon **(above, left)** jumps. This is called visually evoked auditory response. From: [YouTube] The dot **(above, right)** seems to be climbing endless stairs just as the tone seems to climb endlessly higher and higher in the shepard tone illusion. From: [YouTube]

Consider jumping pylon (far left), it creates the illusion of a sound in about 20% of people. These people have a form vision-sound synaesthesia in which the visual movements (often something pulsing, or moving rhythmically) can trigger sound associations. Everyone has had the experience and identifying some object to grab, looking away while grabbing it, and fumbling the pick-up. The video (left) shows just how dissociated conscious perceptions of our body are from our brain's inferences about what body parts are where.

12.5.b Conscious Inference Requires Working Memory

So, the examples in the last section of this chapter illustrate the enormous volume of information and the complexity of information that the gets processed unconsciously by the human brain. What about conscious inferences and the conscious aspects of semi-conscious inferences? All such inferences utilize working memory. What is working memory and what do psychologists and neuroscientists currently know about working memory? Psychologists and neuroscientists currently know quite a lot about working memory. But, as with unconscious inferences, the answers scientists offer differ quite significantly from what students might expect. To start, one might suppose that working memory is a simple container in which shorter-term memories are stored. This is, in fact, not the most common model of working memory. Most psychologists and neuroscientists have adopted the "multi-component model" of working memory. So, how does the multi-component model differ from the little chalkboard used in the analogy earlier in the chapter? Are there just multiple little chalkboards?

The origins of the multi-component model of working memory dates back to a 1974 paper by Allen Baddeley and Graham Hitch.[49] In that paper Baddeley and Hitch tell readers that their model conceives of working memory as single common system composed of multiple sub-systems. That linked collection of subsystems is "limited in capacity and operates across a range of tasks involving different processing codes and different input modalities."[49] (p.35) By 2003 Baddeley refines his initial model to the one depicted in the diagram (above) and researchers start to determine what areas of the brain are responsible for the various components and operations depicted in the model. Baddeley's model includes three different memory stores; the visuospatial sketch pad, the phonological loop, and the episodic buffer. Each of these memory stores holds a specific kind of information represented in a specific manner. The visuospatial sketch pad (VSP) stores visual and spatial information in a non-verbal format that encodes features and objects which it

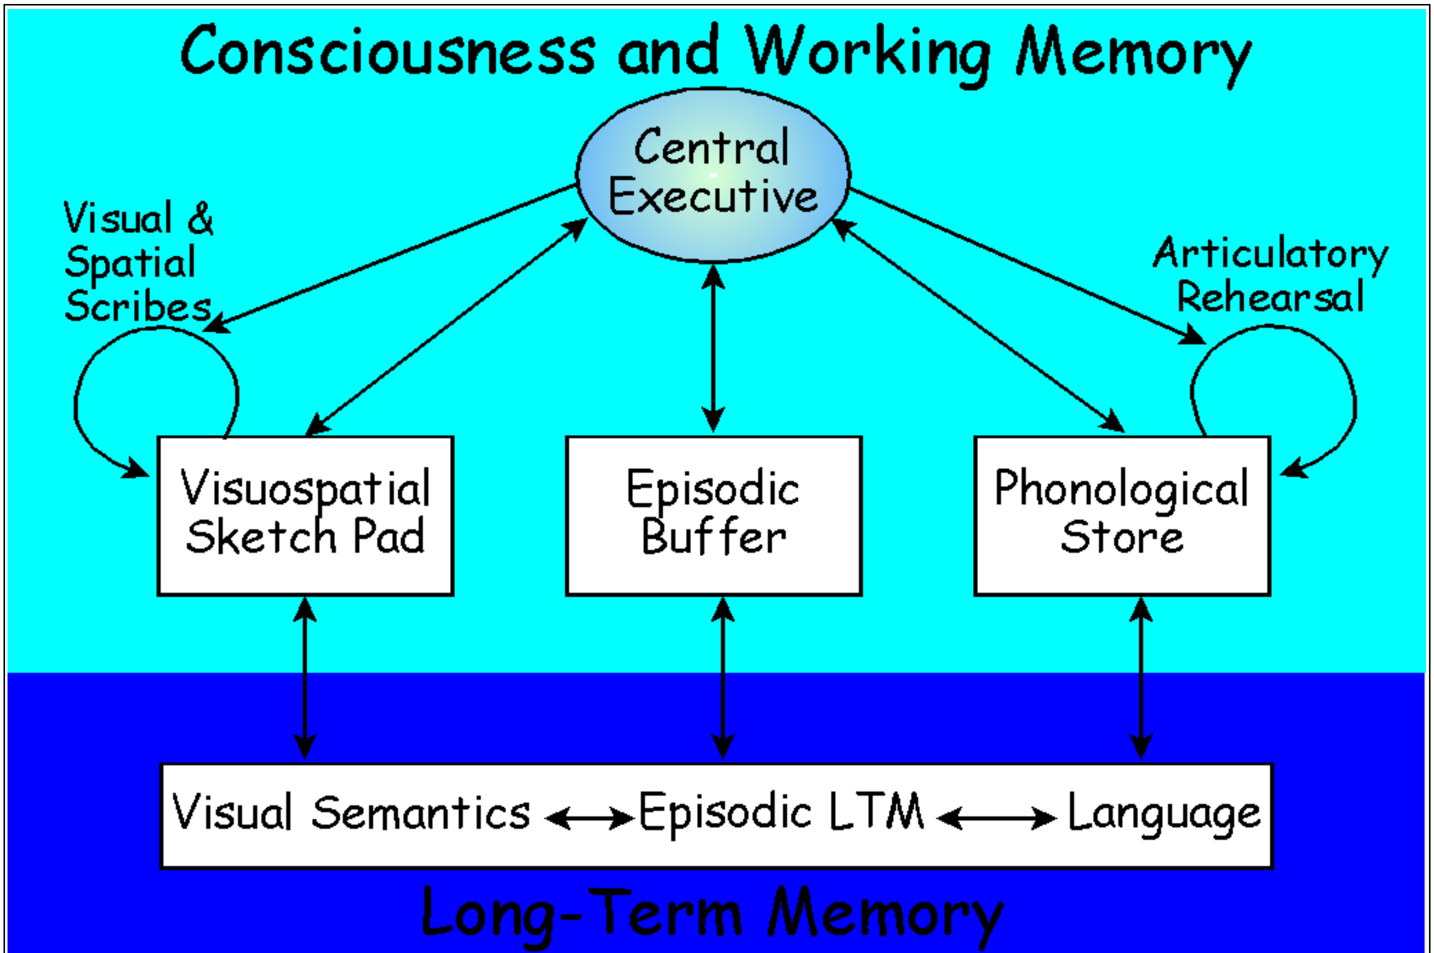

Diagram depicting the Baddeley multi-component model of working memory adapted from his 2003 diagram.[50]

can bind together into visual objects. For instance, the VSP would encode a red triangle by binding its representation of redness and its representation of triangularity. Information enters the visuospatial sketch pad (VSP) when the visual system attends to it. Once in the visuospatial sketch pad, information will degrade if not maintained by processes called the visual and spatial scribes, which are intimately related to attention.[50, 51] The phonological loop stores acoustic and/or phonological and order information. The phonological loop is implicated in human language learning. Once information enters into the phonological loop it will degrade relatively quickly unless it is maintained by the processes of articulatory rehearsal. Thus, once you hear a series of numbers you must rehearse those numbers to maintain them in the phonological loop. The final store, the episodic buffer, stores information in a complex multi-model format as scenes or episodes. The episodic buffer, under the control of the central executive, transfers and translates information between the phonological loop and the visuospatial sketch pad. It combines information into complex scene and episode representations that it can manipulate to consciously solve problems in parallel and serial fashion. Additionally, the episodic buffer facilitates information transfer between LTM and working memory.[50, 51]

Finally, the central executive directs information flow among the component stores within working memory and between working memory and long-term memory when such transfers are not habitual. The central executive, directs attention to specific information, suppresses distractions, inhibits inappropriate actions, information and inappropriate actions, coordinates processing for a task, and coordinates between tasks when multi- tasking.[50, 51]

12.5.b.1 Working Memory is Relatively Small Measured in Both Size and Complexity of Information
As has been mentioned in earlier chapters, working memory has somewhat severe limitations on the amount of information and the complexity of information it can store and/or process. The specific limitations depend upon the specific memory stores within working memory. The number of individual items available in the phonological loop of working memory ranges between three and eight items of rather limited complexity. In contrast, the iconic memory of the visual system contains and briefly stores, for instance, information about the entire visual scene in the visual cortex. It makes massive, highly complex inferences with this initial data even before any information leaves the eye via the optic nerve.

12.5.b.2 Limits on Information in Working Memory Limits Conscious Reasoning Abilities
Measures of working memory indicating a capacity ranging between five and nine items predate the concept of working memory itself.[52] Probably the most famous measure of working memory capacity appears in George Miller's 1956 paper, "The Magic Number Seven, Plus or Minus Two."[53, 54] Contemporary researchers tie capacity estimates to the specific component of working memory as well as the complexity of information. For instance, researchers estimate the capacity of the phonological loop to store words ranges from three elements to eight. However, the number of items varies with their length in that the number of words one can store decreases as a function complexity. The number of items decreases proportional to the time it takes to speak those; long words mean fewer items. Likewise, the capacity of the phonological loop for stored words decreases for very similar-sounding words, and increases for dissimilar-sounding words. In short, more complex items exhaust capacity sooner. Chunking items together can mitigate these limits somewhat. For instance, remembering sequences of three-digit chunks often allows one to remember more digits than remembering each digit individually.[51, 53, 55] Measures of the capacity of the visual component of the visuospatial sketch pad currently place the number of items between three to four items having one to four kinds of features. Within this framework complexity does not seem to affect capacity. However, three to four items appears to be a hard limit on visual working memory capacity.[51, 56-58]

Students who wonder if the limits in information capacities just discussed might be overcome by, for instance, brain training will find little support in current scientific research. Most research suggests that the amount and complexity of information one can store in working memory has strong genetic determinants.[59-61] Training on specific tasks often improves performance on that task. However, improvements in a specific task do not appear to transfer to improved performance overall. Nor do task-specific performance improvements tend to last after such training stops. Moreover, like many cognitive functions, the capacity of working memory appears to decrease with age.[62, 63] Some evidence does seem to suggest that brain training (and generally having an active intellect) might mitigate age-related declines in working memory.

Finally, measures of working memory capacity are strongly related to fluid intelligence—the ability engage in adaptive problem solving and decision-making as well as spotting patterns in experience, particularly in novel, uncertain, and low-information contexts.[64] In psychological parlance, working memory capacity explains most of the variance between individual levels of fluid intelligence. Roughly speaking, the greater the capacity of various components of a person's working memory, the greater the level of fluid intelligence the person exhibits in tasks related to that capacity.

Alternatively, working memory generally serves as a bottleneck in cognitive processing, limiting the amount and complexity of information an individual can utilize in conscious inference and problem-solving.

12.6 Human Inference Strategies and their Typical Deployment

So far the discussion in this chapter characterizes inferences and the properties that can distinguish good inferences from less useful inferences. It likewise notes that the human brain make inferences consciously, semi-consciously, and unconsciously. Important and interesting questions thus arise: "What human inferences tend to have these properties?" "Which strategy proves better?" To answer these questions, one needs to differentiate (divide or classify) human inference strategies into three different classes--three tiers of human reasoning abilities. Psychologists further categorize these classes of human inference strategies into two relatively independent systems for human inference.

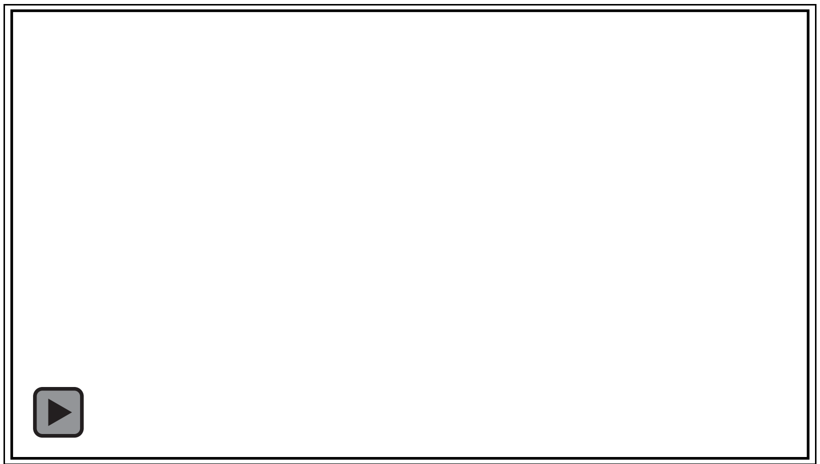

Diagram depicting the three kinds of inference strategies deployed by human beings, the relative likelihood of each being used in a given circumstance, and the relative general reliability of each kind of strategy. The two tiers (classes of inference) in blue collectively form what many psychologists call System 1. System 1 strategies share the properties of (a) automaticity (they work automatically without having to think about or choose them) (b) contextualization (i.e., System 1 inference strategies operate by bringing contextual and content-relevant information to bear on the problem), as well as operating associatively. System 1 strategies exhibit (c) autonomy, meaning that they tend not to draw heavily on working memory. As a result, people exhibit limited conscious awareness, oversight, and insight. In contrast to the inference strategies in System 1, System 2 inference strategies include only the third tier or class of inference strategies, learned rules, depicted in green. Choosing a System 2 process requires conscious awareness and conscious attention to execute. Click on diagram to display animated version.

The first two tiers or classes of inference strategies encompass reasoning strategies that operate relatively automatically with little oversight from consciousness. For this reason, psychologists tend to group them together into a single system, often called "System 1."[65-70] Thus, System 1 includes context-dependent reasoning strategies as well as general heuristics in the diagram above. The term "system" is somewhat misleading in that the categories do not actually pick out determinate, fixed brain systems like the category, "primary visual pathway." Rather, these two inference categories represent different strategies for making inferences and decisions, though each system tends to rely upon certain brain systems more heavily.[71-73] Vinod Goel suggests the following characterization of the data from neuroscience:[72]

> In particular, we need to confront the possibility that there might be no unitary reasoning system in the brain. Rather, the evidence points to a fractionated system that is dynamically configured in response to certain task and environmental cues. The three lines of demarcation reviewed above include (i) systems for heuristic and formal processes (with evidence for some degree of content specificity in the heuristic system); (ii) conflict detection/resolution systems; and (iii) systems for dealing with certain and uncertain inferences. There are undoubtedly others. (p.440)

The misleading connotations of these categories led researchers to propose alternative names. Daniel Kahneman often uses the terms "fast" and "slow."[74] Jonathan Evans and Keith Stanovich have adopted the categories "Type1" and "Type2."[75] Other researchers like Adam Darlow and Steven Sloman adopt the categories "intuitive" and "deliberative."[76]

12.6.a What are General Heuristics?

Before discussing System 1 and System 2 in greater detail, let's take a moment to better understand the term "heuristics." In practice, psychologists call replicable methods or practices directing one's attention in learning, discovery, or problem-solving "heuristics." Pappus of Alexandria, an Greek Mathematician, first introduced the term, which comes from the Greek "heurisko", meaning "I find."[77] Computer scientists both call simple, efficient rules of thumb "heuristics" or "heuristic knowledge." In practice, psychologists call replicable methods or practices directing one's attention in learning, discovery, or problem-solving "heuristics." One employs a heuristic when confronted with a complex problem or when one has incomplete or inaccurate information. In other words, heuristics represent strategies that trade a degree of truth-preservation in one's inference in order to gain the inferential power, speed, or tractability necessary to generate an answer to a given problem. Heuristics implicitly presuppose certain facts about the world and/or the problem in order to facilitate a solution. Hopefully, these implicit presuppositions hold true most of the time, though such presuppositions often have significant exceptions. As a result, heuristics work well under most circumstances, but in certain cases lead to systematic errors in reasoning. Errors arise most often when the conditions under which one employs a heuristic vary dramatically from the conditions under which the heuristic evolved. That is, heuristics implicitly make assumptions designed to facilitate problem-solving in the environment that leads to their selection. Whenever the conditions or current use violate those assumptions, one can expect to see systematic errors result from the use of judgment heuristics.

Thus, the first tier of inference strategies, general heuristics, consists of inference strategies one utilizes in general problem solving (that's the general part) and which involve the implicit presupposition of various facts about the problem or the world in order to generate solutions in a timely fashion given the information available (that's the heuristic part). For example, [Amos Tversky]() and [Daniel Kahneman]() are famous for formulating the judgment heuristics humans seem to employ for estimating probability simpliciter.[78-83] Like all System 1 inference strategies, one does not choose or monitor judgment heuristics consciously. Indeed, one exhibits extremely limited conscious awareness of their use, much less insight into, or oversight over their functioning. Finally, judgment heuristics implicitly rely upon assumptions regarding the nature of the world to facilitate their functioning. As a result, though heuristics often prove useful, they sometimes they lead to systematic errors. Errors arise most often when the conditions under which one employs a heuristic vary dramatically from the conditions under which the heuristic evolved. That is, these heuristics implicitly make assumptions designed to facilitate problem-solving in the environment that leads to their selection. Whenever the conditions or current use violate those assumptions, one can expect to see systematic errors result from the use of judgment heuristics and limited conscious awareness of their use, appropriateness, or sub-optimal performance.

12.6.b System 1

As noted earlier, the first two tiers or classes of inference strategies encompass reasoning strategies represent part of the human innate brain architecture and functioning. In other words, many of these inference processes are innate tendencies according to which humans process information. That is, humans develop many System 1 inference processes without any explicit instruction. These inference processes also operate relatively automatically with little oversight from consciousness. The first two tiers or classes of inference strategies encompass reasoning strategies that. For this reason, psychologists tend to group them together into a single "system," under various monikers like "System 1," "Type 1," "Fast Thinking," or "Intuitive Judgment System."[65, 68, 74-76] For simplicity I will adopt the original name—System 1. Again, I should emphasize that the term "system" misleads in that the categories do not actually pick out determinate, fixed brain systems like the category, "primary visual pathway," though each system tends to rely upon certain brain systems more heavily.[71-73] Rather, these two inference categories group together inference processes that represent two different strategies for making inferences and decisions.

System 1 includes context-dependent reasoning strategies as well as general heuristics in the diagram above. System 1 processes tend to share several properties, such as, (1) automaticity (they work automatically without having to think about or choose them). In fact, (1a) many of these inference patterns are innate, emerging as part of normal development. In some case learned strategies become consolidated and automated by the brain over time thereby reducing or eliminating the need for attention.[84] System 1 processes also tend to exhibit high levels of (1b) contextualization and function associatively. That is, these processes tend to rely heavily information regarding the specific objects, properties, etc., involved in the current situation and the manner in which that situation presents those objects, properties, and etc.. Likewise, System 1 processes often operate by associating problem elements (for example, associating similar items or the past with the present). System 1 processes exhibit (2) autonomy in that they operate largely outside of working memory. As a result, people tend to exhibit (2a) limited conscious awareness, (2b) oversight, and (2c) insight into the operation of System 1 processes. That is, (a) one employs a System 1 inference as a natural reaction to a situation (1) and (2) without having a conscious awareness of doing so. One has very little ability to affect the operation of a heuristic and very little insight into how one actually solves the problem.

For example, suppose that you need to buy a birthday present for your mom. You might look through a webpage from a store and make judgments about whether she would like various items. You might well make these judgments by employing the representativeness heuristic discussed below. That is, you judge the likelihood that she will like an item by unconsciously comparing it to your concept (understanding) of your mom's taste. You do this (1) as a natural inferential disposition that automatically activates (2) without any awareness that you have reacted to the task by automatically employing the representativeness heuristic. The representativeness heuristic generates these judgments by drawing upon information that you would probably have great difficulty articulating explicitly and overtly, and that you would likely not list as your reasons for your judgment. Moreover, you likely would have great difficulty altering your innate disposition to use the representativeness heuristic in such cases and little to no control of the information upon which the heuristic draws. Finally, since the representativeness heuristic relies heavily upon the content and context of an inferential situation (1b), your shopping inference would prove quite different were you shopping for someone else, like your father. Your search through potential gifts would also very likely go differently if you were in a mall as opposed to sitting at home. You likely will not consider possible gifts, for instance, that are not explicitly presented for your consideration. So, both the content (in the form of the nature of the objects about which you make the inference and the person for whom you are shopping) and the context (in the form of the shopping venue) influence your inferences. Likewise, the context in the form of the features of the situation in which you make the inference will influence the inference. For example, you might think differently in the context of Christmas shopping as opposed to birthday shopping or Mother's Day shopping. Likewise, if you just paid a big bill you might gravitate towards lower priced gifts, while you might spend more if you just got a big bonus. Indeed, the range of prices for those potential gifts and the order in which you consider those potential gifts will likely affect your choice as well.

In summary, both one's general heuristics and one's context-dependent strategies generally consist of innate inferential dispositions that operate automatically in reaction to problems that one encounters. These strategies exhibit very limited conscious awareness, oversight, and insight in their operations because they operate largely outside of working memory. Thus, psychologists often characterize them a forming one system of inference strategies—System 1. One can think of the inference strategies characteristic of System 1 by analogy with the development of search engines and personalization algorithms for the internet. Both System 1 inferences and search engines represent a strategy to quickly and efficiently process large amounts of often complex information. Both accomplish their tasks largely by relying upon heuristic assumptions and specialized systems that operate largely outside the awareness of end users.

12.6.c System 2

In contrast to System 1, psychologists differentiate a second human inference strategy—System 2. System 2 encompasses the third and final tier or class of human inference strategies--one's learned, consciously executed inference strategies. Unlike System 1 inference strategies, System 2 inference strategies tend to require conscious effort—both in deciding to use the strategy and in using the strategy. For instance, towards the end of the term students will learn how to use Bayes' Theorem to infer how a new piece of information affects the probability of an event. Naturally, since one tends to deploy these learned inference strategies consciously one has much more ready access to their functioning when one uses them in problem solving. Thus, psychologists often think of these strategies as composing a separate system for solving problems—System 2.

12.6.d The Relationship Between System 1 and System 2

Together System 1 and System 2 comprise two quite different sets of inference and decision-making processes embodying quite different strategies. These inference strategies operate much like and train and an engineer. System 1 operates like a train rushing towards its destination constrained by the narrow limits of its tracks. The combination of automaticity and autonomy make for a speeding inference system blindly guided by its implicit assumptions. In contrast, System 2 operates like an old, distracted engineer. System 2 will notice if disaster strikes; if the System 1 train jumps of the track, then System 2 can hit the brakes or divert to a different track. However, System 2 much like the engineer operates with only partial awareness of System 1. The diagram below illustrates two important points about these two inferential strategies and how inference strategies from each system function in human reasoning. Specifically, the probability that one will employ a System 1 process to solve a given problem far exceeds the likelihood that one will employ a System 2 strategy.

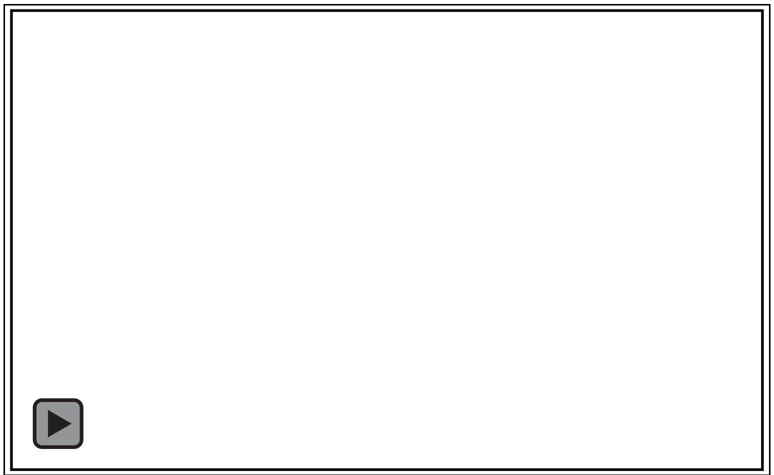

Diagram depicting the two human inference categories and their respective properties. System 1 includes general heuristics and context dependent inference strategies and evolved to automatically engage native inference dispositions whenever humans face a problem. System 1 processes tend to contextualize problems by relying upon the specific context and content of the problem. System 1 inferences require little conscious awareness and oversight to operate. As a result, these strategies allow for very little conscious access into their functioning and very little conscious oversight of their operations. System 2 inference processes, in contrast, draw more heavily on working memory and are learned. They require conscious awareness and oversight to operate. System 2 inferences are not automatically engaged. Indeed, they often prove difficult to engage. However, they tend to compensate for weaknesses inherent in System 1 processes and prove more generally reliable because they tend embody more decontextualized solution strategies. System 2 inference strategies also provide humans with greater conscious insight and oversight into their functioning. Click diagram for animation depicting the potential roles of each category.

However, if one looks at the general reliability of these processes the reverse relationship holds—System 2 strategies (learned rules) tend to have a higher general reliability than System 1 strategies. In short, the inconvenient truth of human reasoning consists in the fact that one is more likely to use a less generally reliable inference strategy to solve a given problem! Worse still, as mentioned in the discussion of critical thinking, innate, genetically determined features of one's brain create this disposition toward less generally reliable strategies. As a result, one cannot significantly temper one's predilection to employ less generally reliable inference strategies since one cannot significantly alter the genetically determined architecture and dispositional functioning of one's brain. The next sections, discuss each of the two tiers or classes of inference strategies in System 1, giving several illustrative examples of strategies from each tier.

12.7 Innate Reasoning Abilities, Inabilities, & Biases: Two Types of Inferences

The last section suggests that one can distinguish System 1 inference strategies from System 2 inference strategies by noticing that System 1 inference strategies represent genetically encoded brain functioning and architecture solutions originating in evolutionary selection in response to a specific environment. In contrast, many of the most important and widespread System 2 inference processes have their origins in the cultural heritage of the last approximately 10,000-12,000 years. People must learn many System 2 inference processes from other people, and people generally must consciously choose to employ those learned System 2 inference processes. In addition to distinguishing between two different sources whereby humans acquire their reasoning abilities, one can also distinguish between the two major classes (kinds) of inferences that humans make. One can base this second distinction on the relationship between the truth of initial information for an inference and the truth or likely truth of the information resulting from the transformation of that information through the inferential process. Logicians call the two classes of inferences deductive inference and inductive inference.

12.8 Deduction: Form, Content, Contextualization

This chapter focuses upon inductive inferences and native human inductive abilities. But, before focusing on induction, let's review the conclusions of Chapter 5 regarding deductive inferences. Deductive inferences serve several important functions: Deductive inferences allow people to transform initial information into new information in a manner that

ALWAYS PRESERVES TRUTH. In other words, if one begins with true beliefs and makes a good (valid) deductive inference, the resulting conclusion will always also be true. Deductive inferences can help to render one's beliefs and worldview systematic and consistent. When one's worldview contains inconsistent beliefs, it contains beliefs that cannot all be true. When one's worldview contains contradictions, then one's worldview contains beliefs the truth of which would imply the falsity of other beliefs in their worldview. Deductive inferences can help an individual to reveal any inconsistencies or contradictions in their worldviews by revealing--making explicit--that some of that individual's beliefs either imply a contradiction or directly contradict other beliefs one also holds. In short, these individually inconsistent beliefs together result in a statement that is necessarily false—a contradiction. Likewise, deduction facilitates the formation of a systematic belief system or worldview by providing a means of assessing whether a belief or a collection of beliefs in the system guarantees the truth of another belief or collection of beliefs. In other words, deduction can help to illuminate the gaps in one's belief system as well as reveal the logical difficulties within one's belief systems and worldview.

12.8.a Deduction Preserves Truth and is Non-Ampliative

Deductive inferences work to preserve the truth of the initial information across the inferential transformation. Thus, good deductive inferences (valid deductive inferences) have a structure such that if one begins with true initial information, the inferential transformation generates necessarily true information as the conclusion. Deductive inferences, as a result, only reveal what must be true given the truth of one's initial information. In one sense, then, deductive inferences do not increase a reasoner's stock of truths. But, in another sense, deductive systems do increase the reasoner's stock of truths. Specifically, deductive inferences transform inexplicit and unavailable truths in the reasoner's stock of information into explicit and available truths. So, deductive inferences serve a very useful purpose despite only revealing what must already be true given the truth of one's current information. Logicians and philosophers call such inferences non-ampliative inference in that these inferences do not increase (amplify) the number of potential truths (explicit and inexplicit information) that the reasoner possesses.

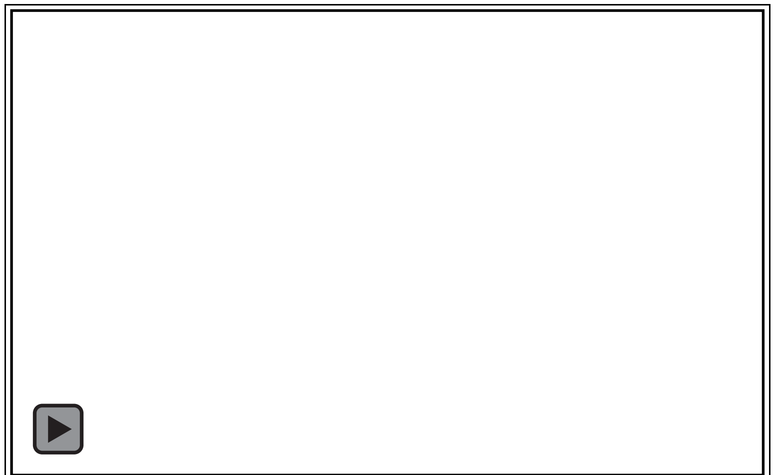

Diagram depicting deductive inferences in terms of the relationship between the initial information explicit and available to the reasoner and the information generated by the inferential transformation. Click on image to see animation.

12.8.b Native Human Deductive Abilities Often Lack Working Memory Resources

Chapter 5 notices that deductive inferences seem more dependent upon language and hence more closely tied to working memory and working memory limitations. As a result, normal human formulations of deductive arguments and evaluation of deductive arguments quickly run into the very real capacity limitations of working memory.

For example, consider another argument taken from Charles Lutwidge Dodgson, better known as Lewis Carroll:[85]

> Animals are always mortally offended if I fail of notice them;
> The only animals that belong to me are in that field;
> No animal can guess a conundrum, unless it has been properly trained in a Board-School;
> None of the animals in that field are badgers;
> When an animal is mortally offended, it always rushes about wildly and howls;
> I never notice any animal, unless it belongs to me;
> No animal, that has been properly trained in a Board-School, ever rushes about wildly and howls.
> --
> Therefore, No badger can guess a conundrum. (p.123)

Is the above argument a good deductive argument? Most people have almost no idea. Carroll's arguments have the character of rambling, unconnected sentences. However, careful analysis reveals the argument's validity. Put simply, the argument proves too complex for intuitive evaluation. The ability of humans to effectively reason, particularly reasoning employing working memory, varies inversely with the amount and complexity of information involved in the inference. For example, clinicians, (doctors, psychologists) perform no better—often worse--on a wide range of clinical judgment tasks when given access to more information (though their subjective confidence in their judgments increases).[86-94] In short, information--even when highly predictive--only proves useful to the extent that the reasoner can exploit the information for the purposes of the inference. Utilizing large amounts of complex information has benefits, but the human ability to utilize such information proves quite finite. As a result, deductive inferences rather quickly become too complex and involve too much information for native human reasoning abilities reliant upon working memory capacity. We saw a similar problem emerge in the informal fallacy of false cause via over-simplified cause. In such cases the reasoner focuses exclusively on a single causal factor and ignoring the true, but much less managable complexity of the causal relationships in the situation.

12.8.c Native Human Deductive Abilities Rely Upon Contextualization

Chapter 5 suggests that the amount of information as well as the complexity of information can quickly and adversely impact intuitive evaluations and formulations of deductive inferences. Information also enters into deductive inference abilities more directly through the salience of content in formulating and evaluating arguments. Indeed, research demonstrates a strong dependence upon content and context in the formulation and in the evaluation of deductive inferences by human subjects. As a result, researchers can present a clear and detailed hierarchy of difficulty of argument types for human formulation and intuitive evaluation.

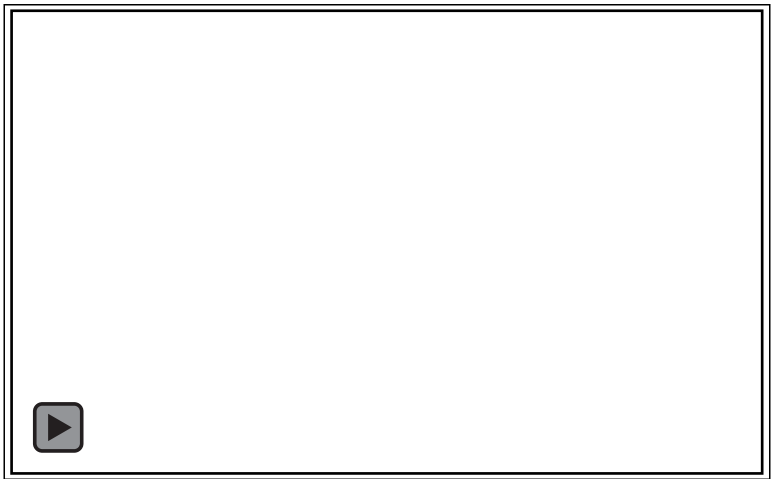

Diagram indicating the relative difficulty of making or evaluating deductive inferences with various kinds of content. The easiest types of inferences to make or evaluate involve familiar content (i.e., are about, familiar objects, properties, events, or relations) where the premises and conclusion are true (valid) or false (invalid). Replacing familiar with abstract content (i.e., like symbols) makes inferences and their evaluation more difficult. Replacing familiar content with nonsense words (i.e., pseudo-content that the brain tries to use) increases the difficulty. Finally, replacing familiar content where the premises and conclusion are true (valid) or false (invalid) with familiar content where the truth-values of the premises and conclusion vary from all true or all false makes inferences the hardest to correctly perform or evaluate. Click on diagram to view animation.

In general, researchers report that people have the least difficulty in evaluating deductive arguments when those arguments involve content with which the person has familiarity. Similarly, people perform better when argument content mirrors the underlying logical structure of the argument (e.x., true premises, true conclusion—valid; false premises, false conclusion—invalid). People tend to find arguments lacking content, like the abstractly symbolized argument in the second box, more difficult. In fact, performance on argument evaluation tasks drops significantly.[95-97] Arguments employing pseudo-content (meaningless word-like content) prove even more difficult for most people to evaluate. Finally, arguments in which the content seems inconsistent with one's beliefs or in which the argument's content fails to mirror the underlying argument structure prove the most troublesome for people (i.e., false premises, false conclusion—valid; true premises, true conclusion—invalid).

12.8.c Strengths and Limitations of Native Deductive Reasoning Abilities

As Chapter 5 observes, contextualization and limited working memory capacity make complex deductive inference problems extremely challenging. Lewis Carroll above nicely illustrates one general manifestation of this limitation. Evans' work on negative conclusion bias (also called double negation bias) illustrates another such manifestation.[98-100] Consider the following two conditional inferences:

If the number is 5, then the letter is Q.
The letter is not Q.

The number is not 5.

and

If the circle is not red, then the square is blue.
The square is not blue.

The circle is red.

Evans reports that though both inferences are correct (valid) on average the responses regarding the correctness of the first inference usually reach 75% while the responses identifying the second inference as correct usually only reach about 45%. Evan suggests that the additional negation in the conclusion, "The circle is not, not red," causes the decrease in correct responses. Wason selection tasks illustrate how contextualization, this time in the form of a seeming context dependent conditional inference ability, likewise limits native human deductive inference abilities. In general, humans appear to find conditional reasoning tasks with deontic conditionals (conditionals expressing permission and obligation) easier than any other conditional-based reasoning tasks. Familiar content tends to make all conditional reasoning tasks somewhat easier than those with unfamiliar content as well.[104-107]

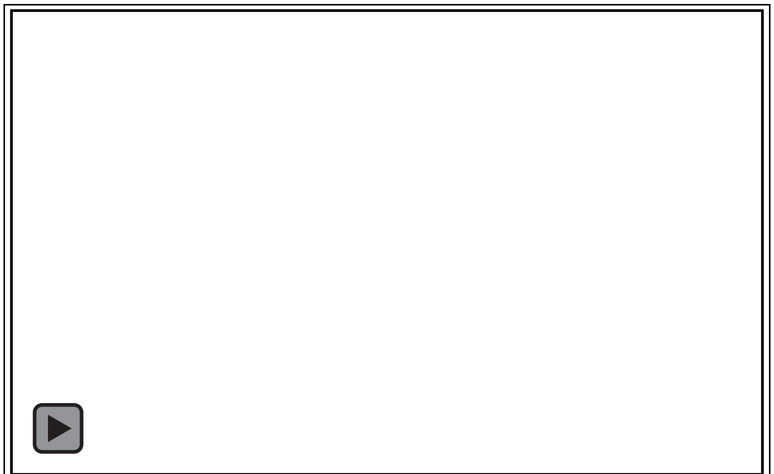

Example of two conditional reasoning tasks Wason explores in adults and that Cummins, Chao, and Cheng[101-103] in explore in development. Such studies seem to support the hypothesis that people have difficulty making conditional inferences as well as a difficulty evaluating the truth of conditionals in many cases. Click on the diagram to play the video

Finally, content-reliance introduces one systematic bias in innate human deductive reasoning resulting from the tendency to contextualize (i.e., rely heavily on content and context) deductive reasoning. Specifically, people tend to judge as good (valid) arguments with believable or believed conclusions; people tend to judge as bad (invalid) arguments with unbelievable or not believed conclusions. Researchers call this tendency "Belief Bias".[108-112] Belief bias arises because conclusion believability can prove logically irrelevant, but psychologically relevant to humans. The graphic below illustrates the relationships between an argument's content and the difficulty it presents to typical humans when the try to formulate or evaluate the argument. Importantly, familiarity of the content, the type of content, and the relationship between content and underlying logical structure affect human performance on deductive reasoning tasks.

12.8.d Formal Deductive Logics Separate Form & Content

Beginning in the Arguments Chapter and continuing in the material on deductive logics, probability, and statistics the

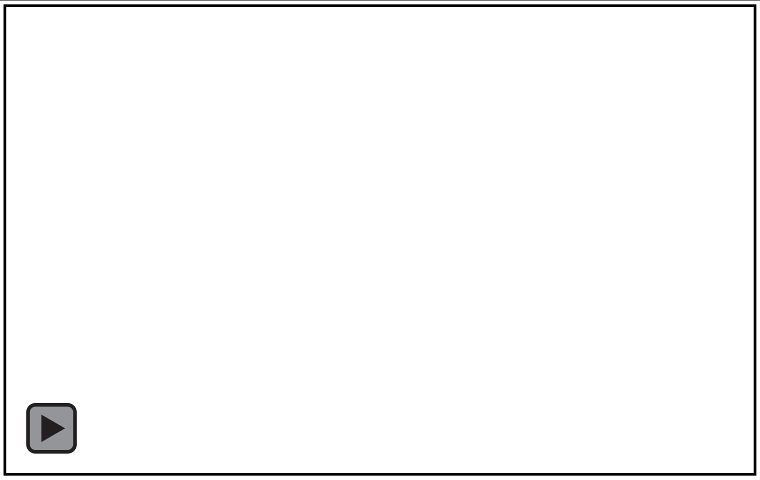

Video illustrating the relationship between content and logical for in both bad (invalid) and good (valid) deductive logic. Click diagram to play video.

text emphasizes that formal inference systems provide an alternative strategy for making inferences. The lectures on categorical and propositional logics introduce the distinction between logical form and content. Logical form designates the relationships between individual content elements in a given argument. Aristotle as well as the logicians and mathematicians that follow him, recognize that arguments having different content, i.e., arguments about different things, can nevertheless have identical logical form—underlying structure. In recognizing that the quality of arguments results from both the quality of their contents (i.e. true or false premises) and the quality of their logical form (i.e., valid--truth-preserving structural relationships, or invalid—structural relationships that fail to preserve truth), Aristotle begins a tradition of developing formal systems, i.e., symbolic representations of the logical structure of arguments.

One creates a logical system to make the inexplicit logical structure of a class of arguments/inferences/decisions explicit, i.e., to formalize them so as to evaluate the logical form of arguments free of the potential biases of content and context that the text has observed so often. This formalization of a class of arguments/inferences/decisions involves three elements: First, one must create an explicit representational system (a language or symbol system) that makes the important logical structure of the argument class explicit. Logicians, philosophers and mathematicians call this representational system by names such as the language, the symbol system or, as in this text, the formal system. For instance, numerals symbolize quantities in mathematics and symbols like "+" and "=" symbolize operations and relationships. The logical operators of propositional logic--●, v, ~, ⊃, and ≡--symbolize truth-functional relationships between statements.

In addition to making the underlying logical structure explicit, one formalizes a given class of arguments with two goals in mind. First, one wants to capture all the arguments of that class in one's formalization, i.e., you want power. Second, one also wants to make one's formalization as simple to learn and use as possible, i.e. to maximize speed and tractability. These two goals are somewhat orthogonal, so one tries to balance completeness and power against ease of learning and use. Thus, one creates an artifact, a simplified system representing the underlying structure of that class of arguments.

Second, one must create a decision procedure that can take the new information about the underlying structure of the problem to evaluate that structure. One must find techniques to evaluate the logical structure of a given class of arguments so that users can decide if any given argument has good logical form. These decision procedures prove crucial since humans have greater difficulty in evaluating arguments when stripped of their content. In deductive logics the decision procedure is effective because it will tell you for all such arguments whether the arguments are valid or invalid. Logicians and mathematicians prize effective decision procedures because such procedures will always give the user a solution. Unfortunately, not all decision procedures prove effective. For instance, the counterexample technique from the arguments chapter could only show arguments invalid--and only if students discovered a substitution instance. Likewise, in most areas of math we have only solution procedures.

One rigorously specifies the formal system and the decision procedures. However, one must pay for the tradeoff between power and speed/ease of use in one's specification of the third element of one's formal system--the translation heuristics. Specifically, one must find general, but fallible, strategies for taking the more complex arguments and inferences in the target class and adequately representing their logical form in the simplified formal system. For instance, in math one translates by counting and measuring objects.

Formalized inference and decision-making systems operate by making explicit and utilizing the underlying structural features common to classes of problems. Their operation, therefore, is highly decontextualized. They tend to require both working memory capacity and externalized memory (such as writing problems and steps down). The advantages of formal systems carry a number of costs as well. Humans have a difficult time stepping out of their normal, automatic reasoning patterns. Thus, even people with significant training in logic and statistical training often fail to recognize that particular inferences and arguments are more optimally addressed with a formal system, failing as a result to utilize the relevant formal system and decision procedure. Even if one recognizes the applicability of formal inference systems, their actual application can prove time consuming and cumbersome. In other words, these formal systems tend to optimize the exploitation of underlying structural features in inference and decision-making, making them unintuitive and difficult for people to use. For example, formal logics utilize underlying logical form to better evaluate an arguments or inference's ability to preserve truth. But the use of such systems costs speed, informational content, cognitive resources. To minimize the costs in cognitive resources and speed, logicians deliberately craft formal systems to balance simplicity against informational content. Thus, the deductive lectures present a categorical logic having only three quantifiers and four standard forms. The propositional logic has one unary connective (negation) and four of a possible sixteen binary connectives. These simpler systems facilitate ease of learning and use, but the cost appears when one translates the richer and more complicated representations of ordinary cognition and language into these formal systems. For this reason, translating real-world problems into the formal system remains heuristics in nature and dependent upon judgment.

Hopefully, students have come to appreciate the relative strengths and weaknesses of native human inference abilities and formal deductive systems during the course. The sections of the course covering probability and statistics likewise gave students some familiarity with formal inductive inference systems. But until now the text has said significantly less about native human inductive inference abilities, disabilities, and biases. The remainder of this chapter will focus on native human inductive inference processes, their strengths, presuppositions and their weaknesses.

12.9 Induction: Ampliative, Presumptive, and Risky

Deductive inferences trade inferential power for truth preservation. The other major class (kind) of inferences--inductive inferences--trade little bit of the inference's guarantee of truth for increases in power, speed, and/or tractability. As a result, inductive inferences can extend what one knows beyond what is guaranteed true given one's current knowledge. As a result, inductive inferences prove essential to surviving and thriving. Whenever you make inferences about the future, like whether you will get paid next Friday, you use induction. When you bring your past experiences about an object or context to your current situation, like when you infer that the 405 will be too crowded coming home from work, you use induction. When you make inferences about that you cannot see based upon what you can see, like when you identify a partially obscured object, you use induction. Indeed, you make explicit or implicit inductive inferences almost every minute of your conscious day.

12.9.a Induction is Ampliative, Adding to a Reasoner's Stock of Truths

Inductive inferences seek to stretch the information available to the reasoner to cover new and possibly different situations. Thus, inductive inference is ampliative. That is, inductive inference attempts to add information to a reasoner's stock of truths. All inductive inferences, as a result, transform one's initial information in accordance with one or more implicit assumptions about the structure of the world or about a regularity in the way the world changes. The implicit assumptions driving ampliative inferences take the form of the inference strategies or rules themselves--the actual mechanisms of information transformation. As a rule, inductive inferences suppose (at least) that new situations will resemble old situations in some respect and to some degree.

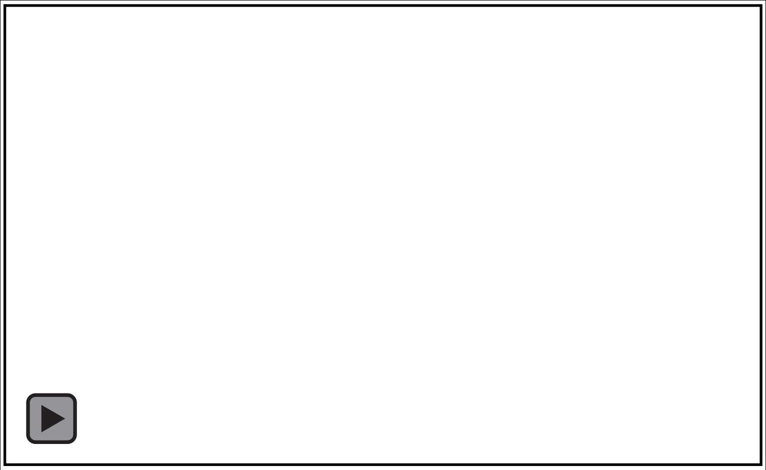

Diagram depicting inductive inferences in relation to both deductive inferences and relationship between the initial information explicit and available to the reasoner and the information generated by the inferential transformation. Click on image to see animation.

12.9.b Ampliative Inferences Presuppose Structure in the World

Ampliative inferences extend the reach of inference beyond what must be true given what a reasoner already knows. In order to extend beyond what must be true, ampliative inference strategies must leverage features of the world. For instance, the British Empiricist philosopher David Hume characterized inductive inference as supposing that the future must resemble the past in certain respects and to certain degrees.[113] All inductive inferences, as a result, transform one's initial information in accordance with one or more implicit assumptions about the structure of the world or about a regularity in the way the world unfolds. The implicit assumptions driving ampliative inferences take the form of the inference strategies or rules themselves--they are generally the actual mechanisms of information transformation. Ampilative inferences suppose (at least) that new situations will resemble old situations in some respect and to some degree. An argument from analogy nicely illustrates that feature of ampliative inferences: Suppose that you watch as movie starring Alicia Vikander and find it very good. You might reason that other movies starring Alicia Vikander will also be of high quality. In general, the sorts of implicit assumptions that drive ampliative inferences don't always hold or hold to the degree one supposes. For example, you may have just watched a major feature based upon classic novel with great screenplay and high production values, like Anna Karenina. When you watch Seventh Son, you may find that it lacks the polish and intelligence of Anna Karenina.

12.9.c Ampliative Inferences Trade Truth for Power Thereby Introducing Epistemic Risk

The imperfect relationship between the truth of one's initial information and the truth of the resulting inferentially generated information means that ampliative inferences trade truth for inferential power. The truth of one's initial information does not guarantee the truth of the conclusion. Indeed, your disappointment with the Seventh Son results from the imperfect relationship between the overall quality of a movie and the actors in that movie. The sorts of implicit assumptions that drive ampliative inferences don't always hold or hold to the degree one supposes. Hume searched for and ultimately despaired of finding a means of showing that his principle of inductive inference must be true given what he already knew. The fundamental truth, therefore, is that ampliative inferences trade some of that guarantee for truth in order to purchase inferential power, speed, and/or tractability. In other words, even a very good inductive inference can result in a conclusion that proves false. As a result, all inductive inferences introduce a certain degree of epistemic risk—risk that your inference will not yield knowledge.

12.9.d Good Ampliative Inferences Manage Their Risk

The virtue of inductive inference, then, does not lie in the perfect preservation of truth from initial information to the conclusion. One should not despair at the ineliminability of inductive epistemic risk. The real inferential world falls short of the best of all possible inferential worlds, but its flaws remain largely manageable. Indeed, good ampliative inference processes generate highly probable information from true the initial information. Very good inductive inferences transform true information to generate one's conclusion in such a way that the conclusion's being false proves very unlikely. On other words, the conclusion of a good inductive inference proves very likely true. Moreover, the best inductive inference processes manage epistemic risk by keeping track of the risks associated with a given inference and with techniques to mediating that risk. Statistics, in fact, shines as an example of just such an inference method.

12.9.e Statistics

Two psychologists who have studied human inductive inference abilities, Amos Tversky and Daniel Kahneman, characterize representativeness as follows:[114]

> *Representativeness* is an assessment of the degree of correspondence between a sample and a population, an instance and a category, an act and an actor or, more generally, between and outcome and a model. (p. 22)

The representativeness relation holds between a population and some bit of knowledge had by the reasoner—a sample of the population. This relationship between the real world population and a sample—a small subset of instances taken from the population--provides the key to understanding most ampliative inferences. Ampliative inferences move from a sample--partial information about objects, properties, events, or relations in some population to information making claims about those features in the entire population—a generalized conclusion. The sample, the partial information, serves as the data or evidence taken from the population, and the ampliative inference extrapolates from that sample—that data or evidence--to make explicit claims about the entire population or novel members of that population.

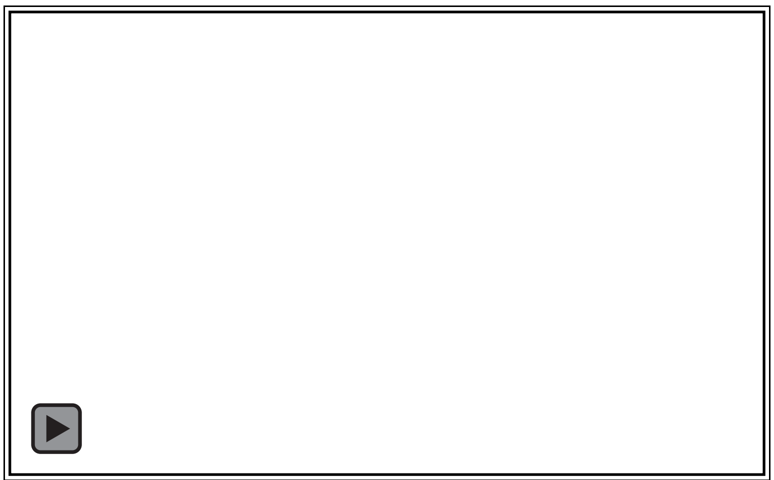

Diagram depicting (1) the relationship between a sample and a population, (2) the relationship between a unrepresentative sample and a population, and (3) the relationship between an representative sample and a population. Click diagram to see animation.

Thus, for Tversky and Kahneman representativeness provides the basis for statistical inference in that it uses the incidence of objects, properties, events, and/or relations within a sample (subset of the population) to infer the incidence of those objects, properties, events, and/or relations within a population.

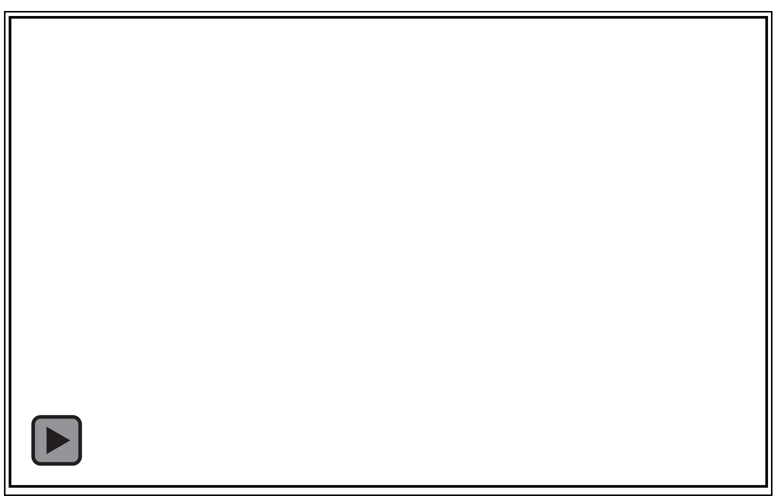

Video outlining the mathematical exploration of the relationship between samples and populations underlying statistics. Click on the diagram to watch.

The material on statistics discusses how thinkers like Abraham De Moivre and Pierre Laplace developed human understanding of this relationship. Specifically, mathematicians inferred the probabilities of drawing samples having specific frequencies when one knows the population frequency. The most probable sample has a frequency equal to the frequency in the population. Disappointingly, none of the samples prove particularly probable. However, the distribution of those frequencies based on their probability created a symmetrical bell-shaped curve around the population frequency. Moreover, the larger the samples the more closely the distribution clustered around the sample frequency matching the population frequency. Mathematicians could then define intervals on that curve such that a reasoner could manipulate the probability of drawing a sample with a frequency that fell within the interval by manipulating its size. Mathematicians defined a standard deviation as the interval within which a sample frequency would fall 67% of the time. Two standard deviations defined and interval within which a sample frequency would fall 95% of the time. Finally, three standard deviations specify an interval such that the sample frequency falls within it 99% of the time. While these moves were insightful, they only licensed inferences from known population frequencies to expected samples frequencies. However, with the advent of the central limit theorem statistics became possible. Roughly speaking the central limit theorem demonstrates that the relationship between populations and likely sample frequencies also holds under certain conditions between sample frequencies and the likely populations that could give rise to that sample. One could then define an interval, the margin of error, around the sample frequency that would capture the range of population frequencies such that 95% of the time the population frequency would fall within that range.

Statistical inference proves reliable because it operates by collecting and analyzing samples in accordance with a set of methods and rules that intelligent and insightful individuals have been developing for almost 300 years. Specifically, one tries to gather a sample from the population about which one wants to make an inference. The reasoner randomly selects a reasonably large sample. The reasoner then measures the features of each member of the sample in a manner that gives unbiased measurements. These rules and methods act so that the dimensions and degrees of representativeness between the sample and the population remain relatively constant. That is, the sample consistently corresponds to the population with regard to some target feature with relatively small variations. In short, the value in the sample provides an excellent basis for estimating the value in the population. The reasoner then formulates a hypothesis about range of possible population values in a manner accounting for sample error by specifying a range of values around the sample value. The history of statistics has largely been a history of refining and expanding upon this basic inference strategy to make increasingly powerful and varied inferences.

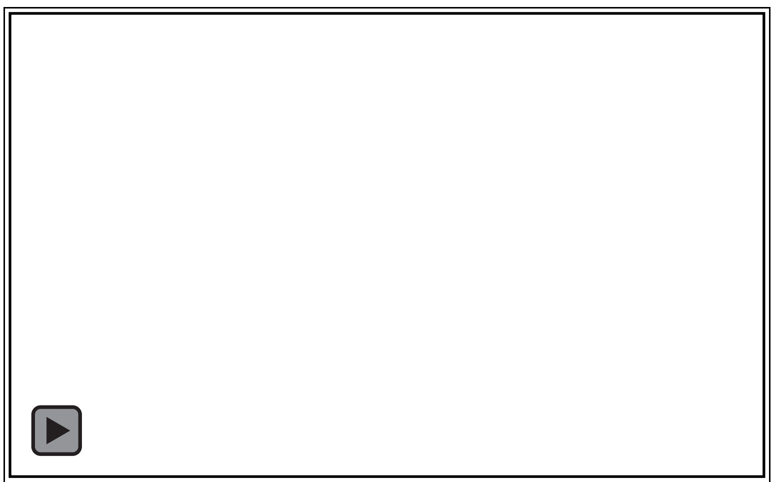

Schematic drawing depicting statistical inference and its underlying assumptions. The inference takes information about the sample and infers a similar range of values in the population based upon well-known representativeness relationships between randomly selected samples of certain sizes and the population. Click on diagram to display animated version.

12.10 Innate Inductive Inference Strategies: General Heuristics

How do inductive inferences relate to System 1 and System 2? General heuristics, the first tier of System 1, transform information inductively. Specifically, these heuristics rely upon assumptions about the world to make ampliative inferences. When these assumptions hold true for the context of an inference, then general heuristics prove good inductive inference processes. Indeed, the fact that an individual hunter-gatherer's experiences typify their environment proves crucial for understanding many general heuristics. Similarly, many heuristics operate to modify our estimates of likelihood in light of new evidence as Bayes Theorem did in probability theory. However, as we will see our native inductive inferences evolved to solve problems in relatively small, relatively stable, and relatively homogenous environments in which an individual's experiences likely prove representative samples. Systematic error can result when the inference context violates these inferential presuppositions. Worse still, since System 1 inference processes tend towards autonomy (minimal reliance on working memory) reasoners often fail to notice when problems or potential problems arise with System 1 inference processes.

12.10.a The Representativeness Heuristic

For example, consider a System 1 heuristic called the representativeness heuristic.[78, 79, 81, 115-118] The representativeness heuristic, in contrast to contemporary statistics, uses one's own concepts and schemas as samples of the population. Humans and proto-humans did not have the time or resources to select random, relatively large and methodologically unbiased samples from the population for each inference. Instead, for most of the 4.4 to 7 million years of human and proto-human existence hunter-gatherers used their own experiences as a general purpose sample. Specifically, the representativeness heuristic works to infer that the probability of an object, property, event, or relation in the world corresponds to how typical the object, property, event, or relation seems given one's concepts and schemas-- the executive summaries of one's experiences. Using one's concept as a sample, the representativeness heuristic estimates real-world probability based upon how typical the object, property, event, or relation appears to be given one's concept. In other words, the representativeness heuristic judges the likelihood of an object, property, event, or relation in the real world by judging the extent to which the object, property, event, or relation typifies the essential or salient features of one's own models and concepts. As a result, the representativeness heuristic relies on three assumptions: (1) It assumes a high degree of representativeness of between one's concepts and schemas and the parent population. (2) It assumes that one's concepts also contain adequate information about the specific relationships targeted in the current inference. (3) The representativeness heuristic assumes that it can measure the frequency of a feature in the sample/concept by how typical it seems given the concept and then infer that the sample frequency and the frequency in the parent population are roughly equal. In short, typical features are very common and have high frequency while atypical features prove uncommon and have low frequency in the sample. High frequency in the sample (measured by typicality) equals (roughly) high frequency in the population low sample frequencies equal low population frequencies. Typicality = Sample Frequency = Population Frequency.

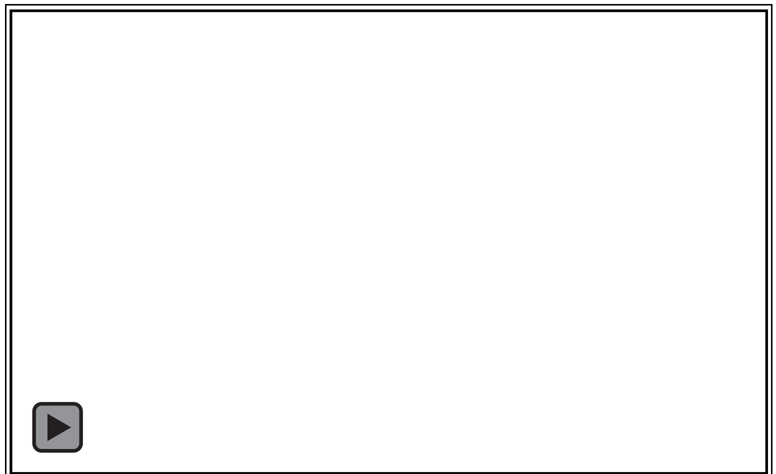

Schematic drawing depicting the representativeness heuristic inference and its underlying assumptions.

For example, suppose that I ask you to estimate the respective probabilities that the fruit in my lunch is an apple, a watermelon, or an olive. You will likely base the estimates you give me for the probabilities of each kind of fruit based upon typicality, .i.e., how typical each kind of fruit--apple, watermelon, and olive—is of a fruit given your fruit concept, i.e., how representative it is of your fruit concept. Since people in North America tend to find apples very typical examples of fruits given their fruit concept, you will likely rate an apple as most likely. Since olives no not have high typicality ratings, you will likely rate olives as the least probable fruit in my lunch.

The representativeness heuristic, as an instance of inductive inference, relies upon the truth of its presuppositions in order to extend one's knowledge beyond one's experiences. As a consequence, the representativeness heuristic generates good probability estimates for objects, properties, events, and/or relations in the real world whenever those presuppositions apply. Conversely, the representativeness heuristic systematically generates poor estimates whenever its presuppositions fail to apply. Specifically, when one deals with a relatively small, stable, and homogenous population one's experiences (concepts and schemas) are much more likely to provide a representative sample, and generate good estimates. When one deals with larger, dynamic, and heterogeneous populations one's experiences (concepts and schemas) prove less likely to provide a representative sample, and tend to generate bad estimates. For instance, people expect chance events to look random. When asked to rate the relative likelihood of the following two sequences of rolls of a fair die, people tend to rate the later sequence as far less likely: 1,3,5,2 or 3,3,3,3. In fact, probability theory dictates that the two sequences are equiprobable (1/1296). Similarly, when asked to rate the likelihood of dying on the job when working as a soldier, a police officer, or a refuse and recyclable materials collectors, people incorrectly rate police officer and soldier as more hazardous. Refuse and recyclable materials collectors have a 33/1,000 fatal injury rate compared to 10.8/1,000 for police officers and 4.45/1,000 for U.S. soldiers serving in Iraq in 2006.[119, 120]

12.10.b The Availability Heuristic

The availability heuristic, also a general heuristic, generates estimates of real-world probability using one's experiences as a sample.[80, 81, 118, 121-123] Instead of relying upon one's concepts and schemas as a ready-made general purpose sample, the availability heuristic implicitly assumes that one's experiences in long-term memory are an equally accessible, representative sample of one's environment. It then "counts" or estimates the number of instances in your experience of a given property, object, event, or relation using ease of recall and/or imagination. The easier one can recall instances of some event, the more common the event in one's experience according to the availability heuristic. The more common an event in one's experience, the more probable the availability heuristic infers the event is in the real world.

Biases of availability arise from the ease or difficulty associated with retrievability or imaginability of instances. If one finds it easier to pull examples of one thing than another thing from memory; that first thing will appear more numerous, even if there are equal numbers of both. Thus, availability can fail to support an inference in two ways: (1) If one's experiences fail to constitute a representative sample of one's environment, then measuring probability by measuring ease of recall or ease of imagination will lead to systematic errors--a failure of representativeness. (2) Since the availability heuristics measures frequency in the sample by ease of recall and/or ease of imagination, it often introduces methodological biases. Unfortunately, factors such as salience, one's life experiences, even the structure of memory can affect retrievability and imaginability. Finally, (3) if recall is not a good measure of sample frequency and/or sample error drives the sample frequency away from the population frequency, then equating the sample frequency with the population frequency will lead to error.

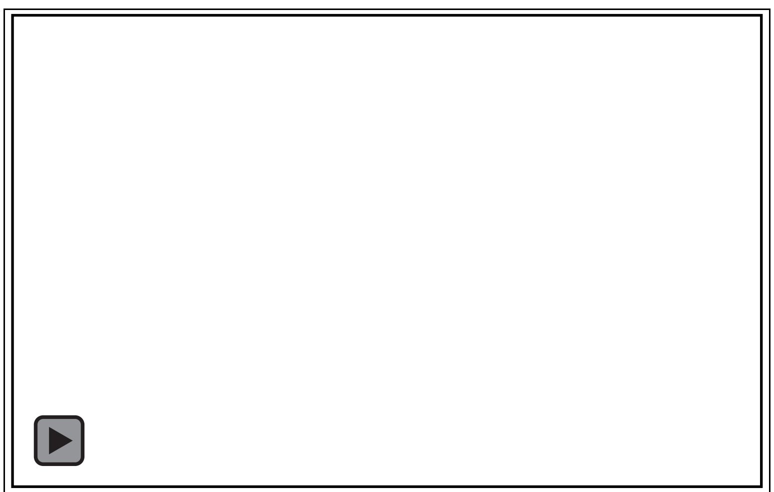

Video discussing the operation of the availability heuristic together with its implicit assumptions. Click on diagram to view video.

How might one's experiences fail to achieve an adequate level of representativeness? When one deals with a relatively small, stable, and homogenous population one's experiences are much more likely to provide a representative sample and generate good estimates. When one deals with larger, dynamic, and heterogeneous populations one's experiences do not provide a representative sample and generate bad estimates. For example, people misjudge the likelihood of two events co-occurring given the strength or weakness of their associations between the two. Insofar as the strength or weakness of association aids or undermines ease of recall or imagination, these errors result from the availability heuristic. To wit, many people believe that red cars get more speeding tickets. This oft repeated urban legend has made it to the level official debunking by Snopes.com. Though very little systematic study addressing such red-bias exists—none of it supports a difference between red and other-colored cars when it comes to speeding tickets. Similarly, researchers investigated the subjective estimates of the likelihood of various lethal events like murder.[124] Researchers asked students to estimate the relative likelihood of two lethal events. For example, subjects compared the relative likelihood of dying of a stroke vs all forms of accidental death. At that time, fatal strokes were 85% more likely than all forms of accidental death combined. However, only 20% of the students judged strokes more likely than all forms of accidental death. The researchers conclude that:

> The judgments exhibited a highly consistent but systematically biased subjective scale of frequency. Two kinds of bias were identified: (a) a tendency to overestimate small frequencies and underestimate larger ones; and (b) a tendency to exaggerate the frequency of some specific causes and to underestimate the frequency of others, at any given level of objective frequency. These biases were traced to a number of possible sources, including disproportionate exposure, memorability, or imaginability of various events. S[ubject]s were unable to correct for these sources of bias when specifically instructed to avoid them. (p. 551)

In other words, the subjects of the study seemed to estimate the frequency of each event in their experiences based upon the ease of recall or imaginability. They then seem to infer a similar rate in the population. That is, they seem to use the availability heuristic.

How might the availability inference itself fail? If one's ease of recall and/or imagination does not accurately reflect the relative frequency in one's experiences, then the inference process itself—and not one's experiences—fails. A classic example of such a failure occurs when one asks people to estimate likelihood of a randomly selected word starting with the letter "k" vs the likelihood of the word having "k" as the third letter.[80] People overwhelmingly judge the likelihood of the word starting with a "k" more probable despite the fact that words with "k" as the third letter vastly outnumber words starting with "k". Ease of recall fails to accurately estimate the ratio in such cases because the human lexical (word) memory easily recalls words by first and last letter place, but cannot recall words easily based upon other letter places. This indexing of events for ease of recall under some cues is not at all uncommon in human cognition.[122] Such strategies again trade speed and power for truth.

12.10.c The Affect Heuristic

Previous chapters discuss the role of positive and negative affect upon inferences and decisions. Researchers in the heuristics and biases literature in cognitive science sometimes refer to the influence and potential biases of affective framing as the affect heuristic.[76, 125-130] The affect heuristic refers to the impact of affective framing, i.e., mood, on belief formation and likelihood estimation. The affect heuristic does not infer beliefs or make likelihood estimates. Rather, it impacts other processes that make such inferences. Slovic and colleagues describe the operation of the affect heuristic as follows:[128]

> …people judge a risk not only by what they think about it but also by how they feel about it. If their feelings toward an activity are favorable, they tend to judge the risks as low and the benefits as high; if their feelings toward the activity are unfavorable, they tend to make the opposite judgment—high risk and low benefit (i.e., the affect heuristic; Finucane, Alhakami, Slovic, & Johnson, 2000). (p.323)

Previous chapter discussed how affective framing can influence decisions in the concepts, meaning, and definition chapter. Indeed, persuasive definitions operate to create affective framing. Similarly, the informal fallacies of appeal to force and *ad hominem* or argument against the person employed affect framing to influence reasoning. A study by Wright and Bower links mood to subjective probability assessments in a similar fashion:[130]

> Subjective probabilities were reported by subjects in a happy, neutral, or sad mood for personal and nonpersonal events. … Large, consistent mood effects are indicated. Relative to control subjects, happy people are "optimistic; " i.e., they report higher probabilities for positive events and lower probabilities for negative events. Conversely, sad people are "pessimistic," providing lower (higher) probabilities for positive (negative) events. Mood-state-dependent retrieval of information is indicated. (p.276)

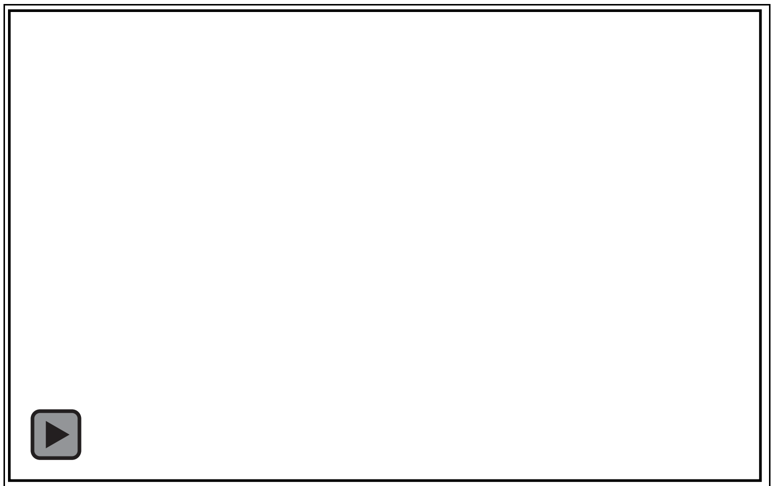

Video discussing the operation of the affect heuristic together with its implicit assumptions. Click on diagram to view video.

Why might affective state (mood) have such a strong influence upon estimates of likelihood and belief formation? Affective states provide reasoners with a rough index of their current state of well-being and sense of the overall threat level. Positive affective states signal higher levels of well-being/lower threat levels. Negative affective states signal lower levels of well-being/higher threat levels. As a result, when one enters into a negative affective state the brain begins to adopt more conservative belief formation stance as well as more conservative approaches to likelihood estimation. In effect, it follows the rule, "when in danger be careful." In contrast, when one enters into a positive affective state the brain begins to adopt less conservative belief formation stance as well as less conservative approaches to likelihood estimation. In effect, living by the creed, "don't worry, be happy."

12.3.d Anchoring and Adjustment

The next general heuristic used to infer likelihoods from evidence goes under the name of, "anchoring and adjustment."[81, 118, 131, 132] Anchoring and adjustment, or just anchoring or cognitive anchoring, refers to a process by which one tends to anchor one's estimates of the likelihood of a property, object, event, or relation to information currently available—regardless of its evidential relevance. Daniel Kahneman and Emos Tversky study anchoring and adjustment as a general heuristic estimating likelihood of a property, object, event, or relation based upon information currently available regardless of its evidential merit. Their rather startling experiment exposes the potential pitfalls of anchoring and adjustment.[133] Kahneman and Tversky give subjects a random anchor point, i.e., some irrelevant data, by spinning a numbered wheel (0-100) rigged to land on either 10 or 65. Kahneman and Tversky then ask subjects to estimate whether the percentage of African countries in the United Nations falls above or below the wheel's number (10 or 65). Finally, the researchers ask subjects to estimate the percentage of African countries in the United Nations. Those subjects anchored with 10 average estimates of 25%, while those anchored with 65 average estimates of 45%. In other words, the anchor values, though seemingly random and definitely irrelevant nevertheless affect subject estimates.[81]

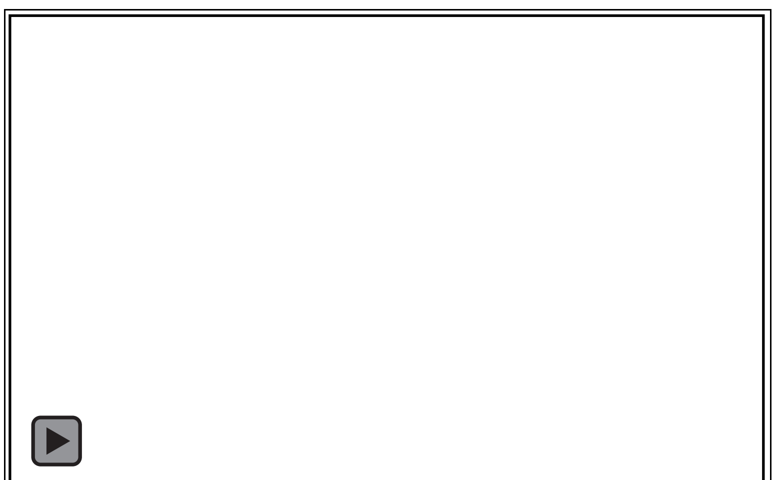

Video discussing the operation of anchoring and adjustment as well as the implicit assumptions that form its evidential basis. Click on diagram to watch video.

Similarly Chen and Kemp report the anchoring effects of self-assessments in the evaluation of retention and tenure cases in universities in their brilliantly named paper, "Lie hard: The effect of self-assessments on academic promotion decisions."[134] Self-assessments are notoriously unreliable. In fact, research indicates that self-assessments prove poorly calibrated and often inversely proportional to actual competence (i.e., very high self-appraisals tend to be given by less competent people).[135-139] Chen and Kemp summarize their findings in a rather understated fashion:[134]

> Practically, the results confirm the common wisdom that it is unwise to use self-assessments for organizational administrative purposes. Raters, supervisors, and decision-makers in promotion processes that use self-assessments are all likely to be biased by them, even when objective evidence is also available, and there does

> not appear to be a simple way to guard against the influence. For applicants the message is that modesty in self-assessment is unlikely to serve them well. (p.587)

Anchoring and adjustment operates by implicitly assuming that the contextual information from the immediate environment is likely to be both adequate and representative for making inferences and decisions within that context. As a result, more sophisticated techniques for evaluating evidential merit need not be employed. Such an assumption looks more plausible given that for most of the 4.4 to 7 million years of hunter-gatherer existence the vast majority of inferences and decisions consisted of reactive, relatively simple, and concrete problem-solving linked to specific contents (problems) and contexts (situations). Systematic error occurs when problems require information that is not present or salient in the problem context or when salient information proves irrelevant or unrepresentative.

12.3.e Confirmation Bias

So far the chapter has focused on processes whereby people make inferences about likelihoods based upon current or past experience. Confirmation bias (also called "confirmatory bias" and "myside bias") speaks to how people seek out information, remember information, and utilize that information.[140-145] Confirmation bias works to shape how humans gather, remember, and utilize information. Specifically, human beings exhibit confirmation bias when they preferentially seek out or interpret information confirming their existing or potential attitudes, beliefs or likelihood assessments. Though the term bias suggests prejudice, the use of bias here only indicates a stronger disposition towards searching for, remembering, and utilizing confirming evidence over disconfirming evidence. Confirmation bias works to shape how humans gather, remember, and utilize information, but it is not an inference process that forms beliefs directly or makes estimates of likelihood. Rather, it shapes how people seek-out and recall information when forming beliefs and/or estimating likelihoods.

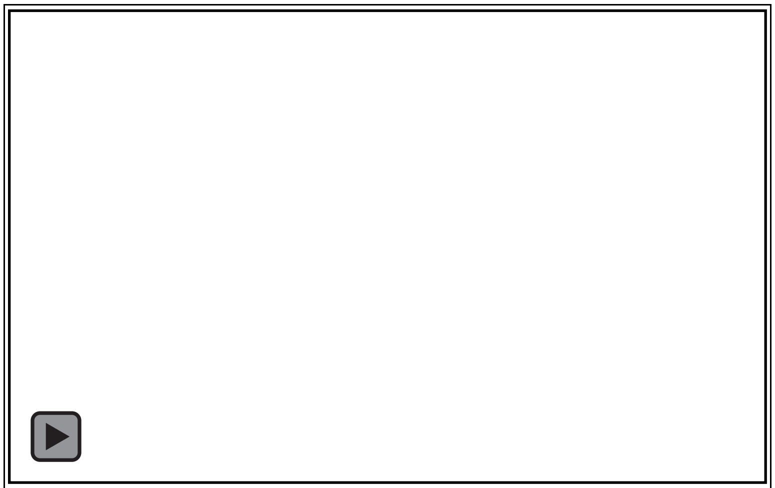

Video discussing the operation of confirmation bias together with its implicit assumptions. Click on diagram to view video.

One can think of confirmation bias as serving a useful purpose insofar as it leads someone to look for information that will provide additional evidence for their beliefs. Nevertheless, confirmation bias also acts to reinforce one's beliefs--even in the face of strong disconfirming evidence. For instance, Robert Carroll gives a nice example of confirmation bias in the skeptic's dictionary. Carroll notes that many psychic researchers have adopted a practice of letting the psychic or the experimenter decide when the trials for an experiment start and stop. As Carroll notes,[146]

> In many tests of psychic powers the subject is allowed to start or stop whenever he or she feels like it. For example, the subject may go through some warm-ups trying to psychically receive numbers or Zener card icons being psychically transmitted by another person. The responses of the warm-ups are recorded, however, and if they look good (i.e., seem to be better than would be expected by chance) then the responses are counted in the experimental data. If not, then the data is discarded. Likewise, if the psychic has had a good run at guessing numbers of card suits and starts to have a bad run, he can call it quits.

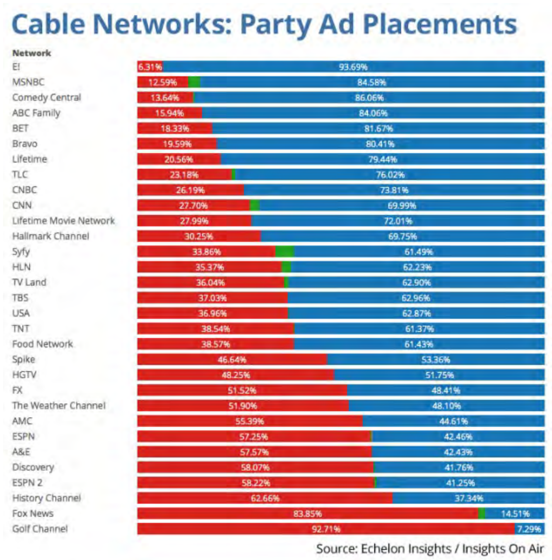

Poltical ads reflect how one's choice of news channels reflects one's political orientation. From: The Daily Signal

In other words, the optional start/stop allows a researcher, even a conscientious researcher, to bias the results of an experiment either for or against psychic powers. How could this happen? It's easy. Optional start/stop allows the experimenter or subject, in effect, to choose the data that counts as evidence for or against psychic powers. Indeed, concerns about the bias of the experimenter or subject affecting data prompted the introduction of "blind studies." In blind studies the experimenter, the subjects, or both do not know, for instance, whether an individual subject is getting the medication or an inert substance. Indeed, experimenter and subject bias are two important sources of error in scientific experiments across all areas of science.

But confirmation bias is not limited to scientific research and scientists. One can find recognition of confirmation bias by thinkers of all stripes and in their thinking. Leo Tolstoy, the Russian fiction writer, opines in his book, "What Is Art?":[147]

> I know that most men—not only those considered clever, but even those who are very clever, and capable of understanding most difficult scientific, mathematical, or philosophic problems—can very seldom discern even the simplest and most obvious truth if it be such as to oblige them to admit the falsity of conclusions they have formed, perhaps with much difficulty—conclusions of which they are proud, which they have taught to others, and on which they have built their lives. (p.198)

Likewise, mundane examples of confirmation bias abound. For example, if you suspect your significant other of cheating on you, you might well look for more evidence to support your suspicion rather than evidence of their faithfulness. The third information ecosystems chapter encouraged people to fight confirmation bias—and for good reason. As the chart (above, left) indicates people tend to watch news programs that reflect their political orientation—reflecting the confirmation bias of their viewership. MSNBC ad spending is dominated by democratic ads at 84.58%, while Fox News represents the inverse with 83.85% of spending going for republican ads.

12.3.f Belief Perseverance

Belief perseverance differs from the previous judgment heuristics in that it functions, not to form new beliefs or likelihood estimates, but to modify those beliefs and estimates in light of new evidence. Belief perseverance an be illustrated by a joke one of my undergraduate math professors told me when I announced I planned to major in philosophy and not math:

> The chairs of the philosophy department, the math department, and the physics department are sitting in the dean's office waiting to meet him. They sit in silence until at last the mathematics chair turns to the physics chair and proclaims: "You physicists make me sick. You need all these gadgets and experiments to discover the nature of reality. All a mathematician needs is some paper, a pencil, and a wastepaper basket." At which point the philosopher looks up, clearly puzzled, and says, "Wastepaper basket?!?!?"

Belief perseverance tends to anchor to one's current estimate of likelihood or one's degree of belief and fail to adjust appropriately when new, relevant evidence presents itself. Lee Ross and his colleagues demonstrate belief perseverance—in a telling experiment:[148] Ross and colleagues ask high school juniors and seniors to examine 25 cards, each containing one real and one fictitious suicide note and determine which of the pair was actually written by a suicide victim. Subjects are given a number for the average correct responses in the task. After each card, the researchers give subjects feedback as to the accuracy of their judgments. However, the feedback each subject receives does not reflect actual performance. The type of feedback depends upon which experimental group the subject is assigned. Some subjects receive greater positive (correct) feedback than expected average, others receive less positive feedback than the expect average, and others receive the expected average positive feedback. Once subjects finish their cards and take a break, the experimenters inform the subjects of the inaccurate nature of their experimentally manipulated feedback and ask the subjects to acknowledge the inaccuracy and its purpose. Following the debunking of their performance on the artificial task, subjects estimate their likely performance on a second, equally difficult set of 25 cards as well as their actual performance relative to the average performance on the original set of 25 cards. Ross and his colleagues summarize the results of their experiments as follows:[148]

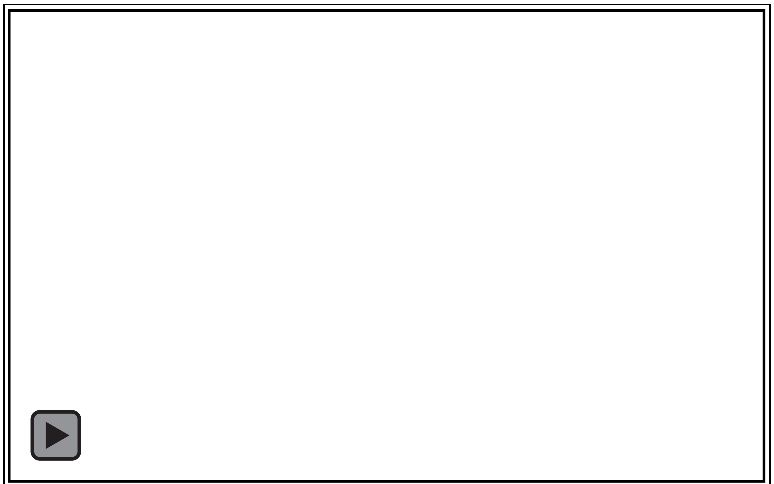

Video depicting the operation and implicit assumptions of belief perseverance. Click of diagram to watch video.

> These assessments of performance and ability showed a clear perseverance effect. Despite thorough debriefing, the greater the subject's apparent initial success, the higher were the scores she estimated for past and future performances. … Only 3 of the 20 success condition subjects thought that their actual scores were worse than average, while only *3* failure condition subjects estimated that their scores as better than average (p.884)

Anchoring and adjustment facilitates inferences from initial experiences to beliefs about or likelihood estimates for properties, objects, events, or relations. Perseverance facilitates the revision of confidence levels for beliefs or likelihood estimates in light of new evidence.[148, 149] As with the previous general judgment heuristics, these processes implicitly assume a relationship between a reasoner's information (sample) and the real world values (populations). Both heuristics assume a homogenous, stable, and small environment. Anchoring and Adjustment takes the contextual information as part of its sample assuming that the information that is available in a reasoner's experience of the immediate environment and that the reasoner associates with a property, object, event, or relation for the purpose of making an inference will prove highly representative of that property, object, event, or relation in the population. Perseverance relies upon two implicit assumptions. **(1)** One ought to assign a relatively low default evidential weight to new evidence absent other indications. **(2)** Concrete beliefs and estimates of likelihood based upon relatively few, generally reliable sources of information prove veridical on the whole and should be revised conservatively.

As with the previous judgment heuristics, systematic error results from (1) a failure of representativeness in one's sample, i.e., experiences, or (2) a failure of the inferential process itself due to a violation of its implicit assumptions. Thus, one ought to ask, "How might one's experiences fail to achieve an adequate level of representativeness?" In the case of anchoring and adjustment, failures of representativeness result when irrelevant or inaccurate information becomes associated with the property, object, event, or relation in one's immediate experiences. Thus, by asking subjects to estimate the percentage of African countries in the U.N. relative to the value of the wheel, Kahneman and Tversky associate irrelevant information with the property of U.N. membership—thereby anchoring the student's estimates. With perseverance, inaccurate likelihood estimates or beliefs prove unrepresentative. Perseverance acts to shield those unrepresentative beliefs and estimates from future evidence.

How might the perseverance inference itself fail? One might argue that perseverance always fails as process of revision for beliefs and likelihood estimates. Recall that the probability lecture introduces Bayes Theorem. Bayes Theorem, as a theorem of probability theory, always correctly modifies likelihood estimates in light of new evidence and the conditional probability of the event given the evidence. Perseverance, so the argument goes, always acts to discount the evidential value of new information—violating Bayes Theorem. However, the case proves more complicated than the just-rehearsed argument suggests. As with all ampliative inference rules, perseverance implicitly assumes that one already has true beliefs and/or accurate likelihood estimates. When one already has highly evinced beliefs or very accurate likelihood estimates, new evidence will only result in small revisions—even using Bayes theorem. Additionally, Bayes Theorem requires that one know the specific strength of the evidence. But one rarely knows the exact strength of the evidence one considers. Thus, one can see that perseverance acts to provide reasoners with a conservative guess about the evidential value of new evidence. Finally, as we learned Bayes Theorem requires one to carry out somewhat complex computations. Perseverance trades truth (or accuracy) for speed and power. It allows reasoners to update their beliefs and likelihood estimates without complex computations that heavily tax finite cognitive resources. It likewise operates in cases absent concrete numbers for evidential strength. The cost for these inferential benefits comes with the increased risk of false beliefs or inaccurate confidence levels/likelihood estimates that these assumptions bring.

12.10.f The Backfire Effect

The third ecosystems chapter recommends exposing oneself to sources with high levels of accuracy and fecundity but of different slants. The idea is that by exposing oneself to other perspectives and opinions one can potentially temper the impact of confirmation bias and belief perseverance. However, this recommendation comes with a caveat (warning). Some research suggests that when people are exposed to evidence that challenges emotionally or cognitively salient beliefs—like the beliefs at the core of one's worldview—these people often report increased confidence in their belief.[150-154] Researchers label this tendency to respond in a contrary fashion to disconfirming evidence "the backfire effect."

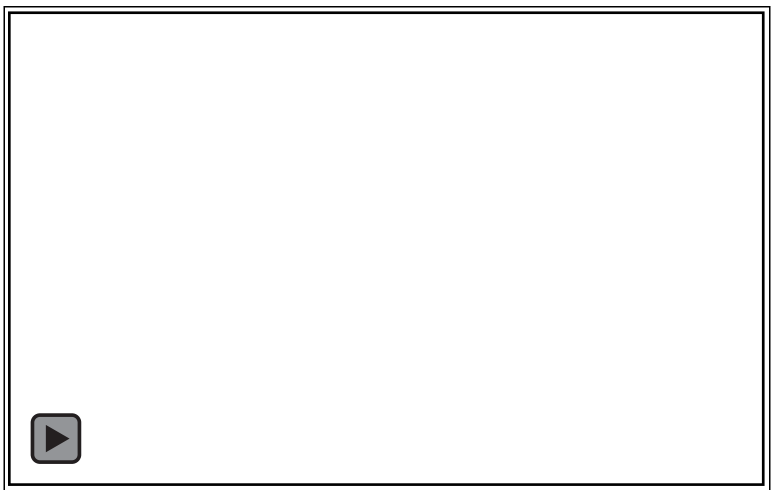

Video discussing the operation of the backfire effect together with its implicit assumptions. Click on diagram to view video.

If we speculate why the backfire effect operates in the manner that it does, it might have to do with the importance of a stable core set of beliefs in one's worldview. As noted with respect to belief perseverance, beliefs based upon relatively few, generally reliable sources of information gathered in a relatively homogenous, stable, and small environment tend to prove trustworthy. Additionally, when one lives a subsistence living like most proto-human and human hunter-gatherers likely did a relatively conservative attitude towards the core elements of one's worldview might well prove beneficial. Thus, cognitively and/or emotionally salient or central components of one's worldview should be strongly conserved in such circumstances as abandonment likely poses a greater risk.

Does the backfire effect mean that people can do nothing to mitigate confirmation bias and belief perseverance? I would say no. However, effective critical thinkers must cultivate a desire to form only true or highly evinced beliefs and accurate likelihood estimates. One must maintain an open mind—even when the emotional and/or ideological stakes are high. Likewise, research does suggest that repeatedly exposing oneself to potentially disconfirming evidence does mitigate belief perseverance and the backfire effect.

12.11 The Nature of Cognitive Illusions

This chapter has sought to rehearse the basics of human inference abilities, limitations and biases. Serious challenges exist for reasoners as they make inferences in a complex and variable world. These challenges are often exacerbated by the fact that human cognitive architecture as well as inference strategies have evolved to solve significantly different problems than the problems faced by contemporary humans. Furthermore, the environment in which humans solve problems has transformed from a relatively small, stable, and homogenous environment to a relatively large, dynamic, and heterogeneous environment. While most inference and decision problems faced by hunter-gatherers consisted of reactive, relatively simple, and concrete problem-solving linked to specific contents (problems) and contexts (situations), a more variable and complex environment demands greater sensitivity to the underlying structural features of problems holding across contexts. Finally, limitations in working memory capacity render conscious oversight of inferences and decisions limited and somewhat unreliable. Thus, cognitive illusions appear in the intellectual lives of human reasoners in a manner similar to visual illusions. Consider, for instance, the St. Louis arch; the arch appears to have a much greater height than width. In fact,it the arch's height and width are an identical 630 feet. The arch appears asymmetrical because the visual system processes the image utilizing false assumptions about the relative size of the base and the apex. While one can come to conscious awareness of this illusion, the mechanisms that give rise to it remain in operation. Conscious awareness can mitigate such illusions, but it cannot alter the architecture and strategies hardwired into the visual brain. The same holds for cognitive illusions. System 1 and System 2 have evolved as architecture and solution strategies over millions of years. Their basic properties cannot be altered significantly.

Picture of the St. Louis arch—a large, human constructed visual illusion. From: Smithsonian Magazine

12.11.a The Moral

Given the fixed points of the human cognitive condition, one can only hope to manage and mitigate the costs of such illusions by cultivating habits and knowledge to affect change in belief and decision outcomes when it matters most. The techniques and knowledge covered in this class represent the insights and bug fixes for the human cognitive system diagnosed and developed by the greatest minds of the last nearly three millennia. Students who utilize these methods and insights can significantly improve their lives and the lives of everyone they affect. It is my fervent hope that this text and course can provide the seeds from which students can cultivate a more productive, sustainable, and satisfying life.

Key Terms

Belief Perseverance: Belief perseverance operates automatically and relatively autonomously in a manner consistent with strategy 1. Belief perseverance acts as a mechanism for reassessing one's beliefs and/or confidence in (probability of) one's beliefs in light of new evidence. It, therefore, operates in a fashion analogous to Bayes Theorem in probability theory. Once a person comes to believe something or makes an assessment of its likelihood, they tend to resist revising that assessment or abandoning that belief in light of new evidence. People will systematically assign lower evidential weight to new, negative evidence even when that evidence is overwhelming and even when the altering confidence or abandoning the belief has little risk or cost to them.

Content-dependent Inference Strategies: Context-dependent inference strategies automatically guide inferences, but do so only in specific kinds of situations. For instance, human conditional reasoning and the evaluation of conditional statements proves much better in deontic (below) situations. Like general heuristics (below), context-dependent inference strategies exhibit (a) innateness, (b) automaticity (they work automatically without having to think about or choose them) (c) contextualization (i.e., System 1 inference strategies operate by bringing contextual and content-relevant information to bear on the problem), as well as exhibiting limited conscious (d) awareness, (e) oversight, and (f) insight.

Contextualized (Contextualization): A term used to describe how human reasoning and assessment of one's own reasoning and the reasoning of others is strongly shaped by the content of one's inferences or argument as well as the context of those inferences or arguments. For example, people tend to judge arguments as better when they agree with the conclusion of the argument and worse when they disagree with the conclusion. This particular content effect is called the belief bias.

Deontic: Deontic is an adjective indicating that the noun is related somehow to permission, duty, obligation, or similar normative concepts. For example, deontic contexts specify a set of contexts in which permission, duty, or obligation issues arise. "Should I run this stop light?" is a deontic question in that it concerns one's actions in relation to norms. As an aside, never run stoplights.

Deductive inference: Deductive inferences work to preserve the truth of the initial information across the inferential transformation. Thus, good deductive inferences (valid deductive inferences) operate such that if the initial information is true, the inferential transformation generates necessarily true information. Deductive inferences, as a result, can only reveal what must be true given the truth of one's initial information.

General Heuristics: General heuristics consist of innate, automatic, inference strategies one utilizes in general problem solving (that's the general part) and which involve the implicit presupposition of various facts about the problem or the world in order to generate solutions in a timely fashion given the information available (that's the heuristic part).

Inductive Inference: Inductive inference extends one's stock of truths by implicitly or explicitly assuming the truth of one or more assumptions regarding the structure of the world or assumptions regarding one or more regularities in the way the world changes. Inductive inferences, by making such assumptions, introduce a degree of risk into one's inferences. Specifically, the implicit presupposition may prove false, thereby generating a false belief. Your internet provider might have a specific problem that it identifies and fixes before the next rain. In such a case, the inductive inference that your internet service will fail with the next rain generates a false belief. The imperfect relationship between the truth of one's initial information and the truth of the resulting inferentially generated information means that inductive inferences trade truth for inferential power. The truth of one's initial information **does not** guarantee the truth of the conclusion, but good inductive inferences generate highly probable information from true the initial information.

Inference: Inferences are psychological processes that take the explicit information available to them and transform it into new explicit information that is now available for some other process, to store in memory, or for guiding action. For example, when one uses the manufacturer's instructions to assemble some furniture, one takes explicit information about the steps involved in assembly gathered through vision to infer sequences of motor actions that will bring out the complete, assembled piece of furniture—hopefully.

Population: Statisticians refer to the larger real world collection of individuals from which one takes a sample as the population or as the target population. For instance, the U.S. Census took a sample from the target population of humans living in the U.S..

The Backfire Effect: When people are exposed to evidence that challenges emotionally or cognitively salient beliefs—like the beliefs at the core of one's worldview—these people often report increasing their confidence in their belief. Researchers label this tendency to respond in a contrary fashion to disconfirming evidence "the backfire effect."

The Representativeness Heuristic: The representativeness heuristic operates using strategy 1. It automatically and autonomously infers that the probability of an object, property, event, or relation in the world corresponds to how typical the object, property, event, or relation seems in one's own experiences. Specifically, the representativeness heuristic estimates real-world probability based upon how typical the object, property, event, or relation appears to be given one's concepts and schemas—the executive summaries of one's experiences. In other words, the representativeness heuristic judges the likelihood of an object, property, event, or relation in the real world by judging the extent to which it typifies the essential or salient features of one's own models and concepts. For example, people often judge a series of rolls of a die that yields 3,3,3 less probable than a series that yields 4,2,6 because the latter seems more representative of a series that would result from a random processes like rolling dice.

Sample: In statistics researchers refer to a sample as a comparatively small group of individuals or objects from a larger, real-world population (target population). The researchers collect information from the sample in order to make statistical inferences about the individuals in the real-world target population. For example, news organizations regularly interview a sample of "likely voters" from the U.S. population. Based upon the information from these likely voters regarding likely choice in an election, news organizations make inferences about who voters in the U.S. population overall are likely to choose in an election.

System 1: System 1 consists of both general heuristics and context-dependent inference strategies. This system evolved so that humans have innate dispositions that automatically engage when facing a problem. The strategies tend to contextualize the problem by relying upon the specific context and content of the problem. System 1 inference strategies require little conscious awareness and oversight to operate. As a result, these strategies allow for very little conscious access into their functioning and very little conscious oversight of their operations. System 1 strategies share the properties of (a) innateness, (b) automaticity (they work automatically without having to think about or choose them) (c) contextualization (i.e., System 1 inference strategies operate by bringing contextual and content-relevant information to bear on the problem), as well as exhibiting limited conscious (d) awareness, (e) oversight, and (f) insight.

System 2: System 2 inference strategies, in contrast, consist of learned knowledge and techniques. Strategies in System 2 do not automatically engage when a reasoner faces a problem. Indeed, they prove difficult to engage. System 2 strategies require conscious awareness and oversight to operate. However, they tend to compensate for the sorts of weaknesses inherent in System 1 strategies and prove more generally reliable because they tend embody more decontextualized solution strategies. System 2 inference strategies also provide humans with greater conscious insight and oversight into their inferences.

12.6 Bibliography

1. Wong, J.C. in The Guardian (The Guardian, New York, NY, 2020).
2. Cohen, E. in CNN Health (Cable News Network, Atlanta, GA, 2020).
3. Wikipedia. in Wikipedia (The Wikimedia Foundation, San Francisco, CA, 2013).
4. Wikipedia. in Wikipedia (The Wikimedia Foundation, San Francisco, CA, 2013).
5. Wikipedia. in Wikipedia (The Wikimedia Foundation, San Francisco, CA, 2012).
6. Wikipedia. in Wikipedia (Wikimedia Foundation, San Francisco, 2012).
7. Wikipedia. in Wikipedia (Wikimedia Foundation, San Francisco, 2012).
8. Wikipedia. in Wikipedia (Wikimedia Foundation, San Francisco, CA, 2013).
9. Lovejoy, C.O., Suwa, G., Simpson, S.W., Matternes, J.H. & Whites, T.D. The Great Divides: Ardipithecus ramidus Reveals the Postcrania of Our Last Common Ancestors with African Apes. *Science* **326**, 100-106 (2009).

10. Wood, B. Reconstructing Human Evolution: Achievements, Challenges, and Opportunities. *Proceedings of the National Academy of Sciences of the United States of America* **107**, 8902-8909 (2010).
11. Wood, B. & Baker, J. Evolution in the Genus Homo. *Annual Review of Ecology, Evolution & Systematics* **42**, 47-69 (2011).
12. Wood, B. & Harrison, T. The Evolutionary Context of the First Hominins. *Nature* **470**, 347-352 (2011).
13. Hanson, B. Light on the Origin of Man. *Science* **326**, 60-61 (2009).
14. Lovejoy, C.O. Reexamining Human Origins in Light of Ardipithecus ramidus. *Science* **326**, 74, 74e1-74e8 (2009).
15. Suwa, G. et al. The Ardipithecus ramidus Skull and Its Implications for Hominid Origins. *Science* **326**, 68, 68e1-68e7 (2009).
16. White, T.D. et al. Ardipithecus ramidus and the Paleobiology of Early Hominids. *Science* **326**, 75-86 (2009).
17. Wikipedia. in Wikipedia (The Wikimedia Foundation, San Francisco, CA, 2016).
18. Wikipedia. in Wikipedia (The Wikimedia Foundation, San Francisco, CA, 2016).
19. Wikipedia. in Wikipedia (Wikimedia Foundation, San Francisco, CA, 2013).
20. Wikipedia. in Wikipedia (The Wikimedia Foundation, San Francisco, CA, 2016).
21. Wikipedia. in Wikipedia (The Wikimedia Foundation, San Francisco, CA, 2016).
22. Wikipedia. in Wikipedia (The Wikimedia Foundation, San Francisco, CA, 2016).
23. Wikipedia. in Wikipedia (The Wikimedia Foundation, San Francisco, CA, 2016).
24. Fitch, W.T. The Evolution of Language (Approaches to the Evolution of Language) (Cambridge University Press, Cambridge, England, 2010).
25. Wikipedia. in Wikipedia (The Wikimedia Foundation, San Francisco, CA, 2015).
26. Wikipedia. in Wikipedia (The Wikimedia Foundation, San Francisco, CA, 2015).
27. Wikipedia. in Wikipedia (The Wikimedia Foundation, San Francisco, CA, 2015).
28. Wikipedia. in Wikipedia (The Wikimedia Foundation, San Francisco, CA, 2015).
29. Wikipedia. in Wikipedia (The Wikimedia Foundation, San Francisco, CA, 2015).
30. Hillert, D. The Nature of Language: Evolution, Paradigms and Circuits (Springer Science + Business Media, New York, NY, US, 2014).
31. Gurven, M. & Kaplan, H. Longevity Among Hunter-Gatherers: A Cross-Cultural Examination. *Population and Development Review* **33**, 321-365 (2007).
32. Bargh, J.A. & Chartrand, T.L. The Unbearable Automaticity of Being. *American Psychologist* **54**, 462 (1999).
33. Bronstad, P.M., Langlois, J.H. & Russell, R. Computational Models of Facial Attractiveness Judgments. *Perception* **37**, 126-142 (2008).
34. Bull, R. & Shead, G. Pupil Dilation, Sex of Stimulus, and Age and Sex of Observer. *Perceptual and Motor Skills* **49**, 27-30 (1979).
35. Hess, E.H. Attitude and Pupil Size. *Scientific American* **212**, 46-54 (1965).
36. Hess, E.H. The Role of Pupil Size in Communication. *Scientific American* **233**, 110-119 (1975).
37. Korichi, R., Pelle-de-Queral, D., Gazano, G. & Aubert, A. Relation Between Facial Morphology, Personality and the Functions of Facial Make-up in Women. *International Journal of Cosmetic Science* **33**, 338-345 (2011).
38. Little, A.C., Jones, B.C., Burt, D.M. & Perrett, D.I. Preferences for Symmetry in Faces Change Across the Menstrual Cycle. *Biological Psychology* **76**, 209-216 (2007).
39. Little, A.C., Jones, B.C. & DeBruine, L.M. Preferences for Variation in Masculinity in Real Male Faces Change Across the Menstrual Cycle: Women Prefer More Masculine Faces When They are More Fertile. *Personality and Individual Differences* **45**, 478-482 (2008).
40. Anthony, C.L. Facial Attractiveness: Evolutionary Based Research. *Philosophical Transactions of the Royal Society B: Biological Sciences* **366**, 1638-1659 (2011).
41. Bzdok, D. et al. ALE Meta-analysis on Facial Judgments of Trustworthiness and Attractiveness. *Brain Structure & Function* **215**, 209-223 (2011).
42. Fink, B., Neave, N., Manning, J.T. & Grammer, K. Facial Symmetry and Judgements of attractiveness, Health and Personality. *Personality & Individual Differences* **41**, 491-499 (2006).
43. Bereczkei, T., Gyuris, P., Koves, P. & Bernath, L. Homogamy, genetic similarity, and imprinting; parental influence on mate choice preferences. *Personality and Individual Differences* **33**, 677-690 (2002).
44. Gyuris, P., Járai, R. & Bereczkei, T. The Effect of Childhood Experiences on Mate Choice in Personality Traits: Homogamy and Sexual Imprinting. *Personality and Individual Differences* **49**, 467-472 (2010).

45. Havlicek, J. & Roberts, S.C. MHC-Correlated Mate Choice in Humans: A Review. *Psychoneuroendocrinology* **34**, 497-512 (2009).
46. Kniffin, K.M. & Wilson, D.S. The Effect of Nonphysical traits on the Perception of Physical Attractiveness: Three Naturalistic Studies. *Evolution and Human Behavior* **25**, 88-101 (2004).
47. Vartanian, L.R. & Hopkinson, M.M. Social Connectedness, Conformity, and Internalization of Societal Standards of Attractiveness. *Body Image* **7**, 86-89 (2010).
48. Zaki, J., Schirmer, J. & Mitchell, J.P. Social Influence Modulates the Neural Computation of Value. *Psychological Science* **22**, 894-900 (2011).
49. AD, B. & GJ, H. in Recent Advances in Learning and Motivation (ed. Bower, G.) 47-90 (Academic Press, New York, NY, 1974).
50. Baddeley, A. Working memory: looking back and looking forward. *Nature Reviews Neuroscience* **4**, 829-839 (2003).
51. RepovŠ, G. & Baddeley, A. The multi-component model of working memory: Explorations in experimental cognitive psychology. *Neuroscience* **139**, 5-21 (2006).
52. Brener, R. An experimental investigation of memory span. *Journal of Experimental Psychology* **26**, 467-482 (1940).
53. Miller, G.A. The Magical Number Seven, Plus or Minus Two. *Psychological Review* **63**, 81-97 (1956).
54. Miller, G.A., Galanter, E. & Pribram, K.H. Plans and the Structure of Behavior (Henry Holt and Co, New York, NY US, 1960).
55. Mueller, S.T., Seymour, T.L., Kieras, D.E. & Meyer, D.E. Theoretical Implications of Articulatory Duration, Phonological Similarity, and Phonological Complexity in Verbal Working Memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition* **29**, 1353-1380 (2003).
56. Joo-seok, H., Woodman, G.F., Vogel, E.K., Hollingworth, A. & Luck, S.J. The Comparison of Visual Working Memory Representations With Perceptual Inputs. *Journal of Experimental Psychology. Human Perception & Performance* **35**, 1140-1160 (2009).
57. Miranda, S., Edward, K.V. & Edward, A. Perceptual expertise enhances the resolution but not the number of representations in working memory. *Psychonomic Bulletin & Review* **15**, 215-222 (2008).
58. Awh, E., Barton, B. & Vogel, E.K. Visual working memory represents a fixed number of items regardless of complexity. *Psychological Science* **18**, 622-628 (2007).
59. Friedman, N.P., Miyake, A., Robinson, J.L. & Hewitt, J.K. Developmental Trajectories in Toddlers' Self-restraint Predict Individual Differences in Executive Functions 14 Years Later: A Behavioral Genetic Analysis. *Developmental Psychology* **47**, 1410-1430 (2011).
60. Friedman, N.P. et al. Individual Differences in Executive Functions are Almost Entirely Genetic in Origin. *Journal of Experimental Psychology: General* **137**, 201-225 (2008).
61. Ando, J., Ono, Y. & Wright, M.J. Genetic Structure of Spatial and Verbal Working Memory. *Behavior Genetics* **31**, 615-624 (2001).
62. Gilchrist, A.L., Cowan, N. & Naveh-Benjamin, M. Working memory capacity for spoken sentences decreases with adult ageing: Recall of fewer but not smaller chunks in older adults. *Memory* **16**, 773-787 (2008).
63. Bo, J. Age-Related Declines in Visuospatial Working Memory Correlate With Deficits in Explicit Motor Sequence Learning. *Journal of Neurophysiology* **102**, 2744-2754 (2009).
64. Chuderski, A. & Necka, E. The Contribution of Working Memory to Fluid Reasoning: Capacity, Control, or Both? *Journal of Experimental Psychology: Learning, Memory, and Cognition* **38**, 1689-1710 (2012).
65. Evans, J.S.B.T. In two minds: Dual-process Accounts of Reasoning. *TRENDS in Cognitive Sciences* **7**, 454-459 (2003).
66. Evans, J.S.B.T. & Over, D.E. Rationality and Reasoning (Psychology/Erlbaum (Uk) Taylor & Fr, Oxford England, 1996).
67. Sloman, S.A. The Empirical Case for Two Systems of reasoning. *Psychological Bulletin* **119**, 3-22 (1996).
68. Stanovich, K.E. Who is Rational?: Studies of Individual Differences in Reasoning (Lawrence Erlbaum Associates Publishers, Mahwah, NJ US, 1999).
69. Stanovich, K.E., Davidson, J.E. & Sternberg, R.J. in The psychology of problem solving. 291-342 (Cambridge University Press, New York, NY US, 2003).
70. Stanovich, K.E. & West, R.F. On the Relative Independence of Thinking Biases and Cognitive Ability. *Journal of Personality and Social Psychology* **94**, 672-695 (2008).

71. De Neys, W. & Goel, V. in Neuroscience of decision making. (eds. Vartanian, O. & Mandel, D.R.) 125-141 (Psychology Press, New York, NY, 2011).
72. Goel, V. Anatomy of deductive reasoning. *Trends in Cognitive Sciences* **11**, 435-441 (2007).
73. Goel, V. in Handbook of neuroscience for the behavioral sciences, Vol 1. (eds. Berntson, G.G. & Cacioppo, J.T.) 417-430 (John Wiley & Sons Inc, Hoboken, NJ, 2009).
74. Kahneman, D. Thinking, Fast and Slow (Thinking, Fast and Slow New York, NY, 2011).
75. Evans, J.S.B.T. & Stanovich, K.E. Dual-Process Theories of Higher Cognition. *Perspectives on Psychological Science* **8**, 223-241 (2013).
76. Darlow, A.L. & Sloman, S.A. Two systems of reasoning: Architecture and relation to emotion. *WIREs Cognitive Science* **1**, 382-392 (2010).
77. Wikipedia. in Wikipedia (The Wikimedia Foundation, San Francisco, CA, 2015).
78. Kahneman, D. & Tversky, A. Subjective probability: A judgment of representativeness. *Cognitive Psychology* **3**, 430-454 (1972).
79. Kahneman, D. & Tversky, A. On the Psychology of Prediction. *Psychological Review* **80**, 237-251 (1973).
80. Tversky, A. & Kahneman, D. Availability: A Heuristic for Judging Frequency and Probability. *Cognitive Psychology* **5**, 207-232 (1973).
81. Tversky, A. & Kahneman, D. Judgment Under Uncertainty: Heuristics and Biases. *Science* **185**, 1124-1131 (1974).
82. Wikipedia. in Wikipedia (The Wikimedia Foundation, San Francisco, CA, 2015).
83. Wikipedia. in Wikipedia (The Wikimedia Foundation, San Francisco, CA, 2015).
84. Song, S. Consciousness and the consolidation of motor learning. *Behavioural brain research* **196**, 180-186 (2009).
85. Carroll, L. (eds. Browne, T., Melwani, G., Weeks, G., Smith, L.L. & Team, T.O.D.P.) (Project Gutenburg, 1897 (2009)).
86. Dawes, R., Faust, D., & Meehl, P. Clinical Versus Actuarial Judgment. *Science* **243**, 1668-1674 (1989).
87. Dawes, R.M., Faust, D. & Meehl, P.E. in A handbook for data analysis in the behavioral sciences: Methodological issues. (eds. Keren, G. & Lewis, C.) 351-367 (Lawrence Erlbaum Associates, Inc, Hillsdale, NJ England, 1993).
88. Faust, D. The Limits of Scientific Reasoning (University of Minnesota Press, Minneapolis, MN, 1984).
89. Goldberg, L.R. Simple Models or Simple Processes? Some Research on Clinical Judgments. *American Psychologist* **23**, 483-496 (1968).
90. Goldberg, L.R. Man Versus Model of Man: A Rationale, Plus Some Evidence, for a Method of Improving on Clinical Inferences. *Psychological Bulletin* **73**, 422-432 (1970).
91. Goldberg, L.R. Five Models of Clinical Judgment: An Empirical Comparison between Linear and Nonlinear Representations of the Human Inference Process. *Organizational Behavior & Human Performance* **6**, 458-479 (1971).
92. Grove, W.M. & Meehl, P.E. Comparative Efficiency of Informal (subjective, impressionistic) and Formal (mechanical, algorithmic) Prediction Procedures: The Clinical–Statistical Controversy. *Psychology, Public Policy, and Law* **2**, 293-323 (1996).
93. Sarbin, T.R. A Contribution to the Study of Actuarial and Individual Methods or Prediction. *American Journal of Sociology* **48**, 593-602 (1943).
94. Sawyer, J. Measurement and Prediction, Clinical and Statistical. *Psychological Bulletin* **66**, 178-200 (1966).
95. Evans, J.S. Logic and Human Reasoning: An Assessment of the Deduction Paradigm. *Psychological Bulletin* **128**, 978-996 (2001).
96. Evans, J.S.B.T., Handley, S.J. & Harper, C.N.J. Necessity, Possibility and Belief: A Study of Syllogistic Reasoning. *Quarterly Journal of Experimental Psychology: Section A* **54**, 935-958 (2001).
97. Evans, J.S.B.T., Handley, S.J., Harper, C.N.J. & Johnson-Laird, P.N. Reasoning about Necessity and Possibility: A Test of the Mental Model Theory of Deduction. *Journal of Experimental Psychology: Learning, Memory, and Cognition* **25**, 1495-1513 (1999).
98. Evans, J.S.B.T., Clibbens, J. & Rood, B. Bias in Conditional Inference: Implications for Mental Models and Mental Logic. *The Quarterly Journal of Experimental Psychology A: Human Experimental Psychology* **48A**, 644-670 (1995).
99. Evans, J.S.B. Reasoning with negatives. *British Journal of Psychology* **63**, 213-219 (1972).
100. Evans, J.S.B. Linguistic determinants of bias in conditional reasoning. *The Quarterly Journal of Experimental Psychology A: Human Experimental Psychology* **35A**, 635-644 (1983).

101. Chao, S.-J. & Cheng, P.W. The Emergence of Inferential Rules: The use of Pragmatic Reasoning Schemas by Preschoolers. *Cognitive Development* **15**, 39-62 (2000).
102. Cummins, D.D. Evidence of Deontic Reasoning in 3- and 4-year-old Children. *Memory & Cognition* **24**, 823 (1996).
103. Cummins Dellarosa, D. Evidence for the Innateness of Deontic Reasoning. *Mind & Language* **11**, 160-190 (1996).
104. Wason, P.C. On the Failure to Eliminate Hypotheses in a Conceptual Task. *The Quarterly Journal of Experimental Psychology* **12**, 129-140 (1960).
105. Wason, P.C. in New Horizons in Psychology (ed. Foss, B.) (Penguin, Harmondsworth, UK, 1966).
106. Wikipedia. in Wikipedia (The Wikimedia Foundation, San Francisco, CA, 2016).
107. Wikipedia. in Wikipedia (The Wikimedia Foundation, San Francisco, CA, 2016).
108. Klauer, K.C., Musch, J. & Naumer, B. On Belief Bias in Syllogistic Reasoning. *Psychological Review* **107**, 852-884 (2000).
109. Wikipedia. in Wikipedia (Wikimedia Foundation Inc. , San Francisco, CA, 2013).
110. Dube, C., Rotello, C.M. & Heit, E. The Belief Bias Effect is Aptly Named: A Reply to Klauer and Kellen (2011). *Psychological Review* **118**, 155-163 (2011).
111. Evans, J.S.B., Barston, J.L. & Pollard, P. On the Conflict Between Logic and Belief in Syllogistic Reasoning. *Memory & Cognition* **11**, 295-306 (1983).
112. Banks, A. The Influence of Activation Level on Belief Bias in Relational Reasoning. *Cognitive Science* **37**, 1-34 (2012).
113. Hume, D. A Treatise of Human Nature (ed. Levine, M.P.) (Barnes & Noble, New York, NY, 2005).
114. Tversky, A. & Kahneman, D. in Heuristics and biases: The psychology of intuitive judgment. (eds. Gilovich, T., Griffin, D. & Kahneman, D.) 19-48 (Cambridge University Press, New York, NY US, 2002).
115. Wells, G.L. The conjunction error and the representativeness heuristic. *Social Cognition* **3**, 266-279 (1985).
116. Metzger, M.A. & Krass, K. Replication of judgmental heuristics with simplified items and familiar content. *Psychological Reports* **63**, 671-677 (1988).
117. Macdonald, R.R. & Gilhooly, K.J. More about Linda: or Conjunctions in context. *European Journal of Cognitive Psychology* **2**, 57-70 (1990).
118. Plous, S. The psychology of judgment and decision making (New York : McGraw-Hill, New York, 1993).
119. Statistics, B.o.L. in Economic News Releases (U.S. Bureau of Labor Statistics, Washington, DC, 2018).
120. Preston, S.H. & Buzzell, E. in University of Pennsylvania Scholarly Commons (University of Pennsylvania, Philadephia, PA, 2006).
121. Carroll, J.S. The Effect of Imagining an Event on Expectations for the Event: An Interpretation in Terms of the Availability Heuristic. *Journal of Experimental Social Psychology* **14**, 88-96 (1978).
122. Taylor, H.A. & Tversky, B. Indexing Events in Memory: Evidence for Index Dominance. *Memory* **5**, 509-542 (1997).
123. Wikipedia. in Wikipedia (The Wikimedia Foundation, San Francisco, CA, 2020).
124. Sarah, L., Paul, S., Baruch, F., Mark, L. & Barbara, C. Judged Frequency of Lethal Events. *Journal of Experimental Psychology: Human Learning and Memory* **4**, 551-578 (1978).
125. Västfjäll, D., Peters, E. & Slovic, P. The affect heuristic, mortality salience, and risk: Domain-specific effects of a natural disaster on risk-benefit perception. *Scandinavian Journal of Psychology* **55**, 527-532 (2014).
126. Lin, W.U., Sheng, Z. & Yue, W.U. AFFECT HEURISTIC AND FORMAT EFFECT IN RISK PERCEPTION. *Social Behavior & Personality: an international journal* **46**, 1331-1344 (2018).
127. Pachur, T., Hertwig, R. & Steinmann, F. How do people judge risks: Availability heuristic, affect heuristic, or both? *Journal of Experimental Psychology: Applied* **18**, 314-330 (2012).
128. Slovic, P. & Peters, E. Risk Perception and Affect. *Current Directions in Psychological Science* **15**, 322-325 (2006).
129. Finucane, M.L., Alhakami, A., Slovic, P. & Johnson, S.M. The affect heuristic in judgments of risks and benefits. *Journal of Behavioral Decision Making* **13**, 1-17 (2000).
130. Wright, W.F. & Bower, G.H. Mood effects on subjective probability assessment. *Organizational Behavior and Human Decision Processes* **52**, 276-291 (1992).
131. Sherif, M., Taub, D. & Hovland, C.I. Assimilation and contrast effects of anchoring stimuli on judgments. *Journal of Experimental Psychology* **55**, 150-155 (1958).
132. Epley, N. & Gilovich, T. The Anchoring-and-Adjustment Heuristic. *Psychological Science (0956-7976)* **17**, 311-318 (2006).

133. Tversky, A. & Kahneman, D. The Framing of Decisions and the Psychology of Choice. *Science* **211**, 453-458 (1981).
134. Chen, Z. & Kemp, S. Lie hard: The effect of self-assessments on academic promotion decisions. *Journal of Economic Psychology* **33**, 578-589 (2012).
135. Ehrlinger, J., Johnson, K., Banner, M., Dunning, D. & Kruger, J. Why the unskilled are unaware: Further explorations of (absent) self-insight among the incompetent. *Organizational Behavior and Human Decision Processes* **105**, 98-121 (2008).
136. Kruger, J. & Dunning, D. Unskilled and unaware of it: How difficulties in recognizing one's own incompetence lead to inflated self-assessments. *Journal of Personality and Social Psychology* **77**, 1121-1134 (1999).
137. Moore, D.A. & Small, D.A. Error and bias in comparative judgment: On being both better and worse than we think we are. *Journal of Personality and Social Psychology* **92**, 972-989 (2007).
138. Dunning, D., Johnson, K., Ehrlinger, J. & Kruger, J. Why people fail to recognize their own incompetence. *Current Directions in Psychological Science* **12**, 83-87 (2003).
139. Williams, E.F. & Gilovich, T. Do People Really Believe they are above Average? *Journal of Experimental Social Psychology* **44**, 1121-1128 (2008).
140. Mynatt, C.R., Doherty, M.E. & Tweney, R.D. Confirmation bias in a simulated research environment: An experimental study of scientific inference. *The Quarterly Journal of Experimental Psychology* **29**, 85-95 (1977).
141. Fischhoff, B. & Beyth-Marom, R. Hypothesis evaluation from a Bayesian perspective. *Psychological Review* **90**, 239-260 (1983).
142. Miller, D.T. & Turnbull, W. Expectancies and interpersonal processes. *Annual Review of Psychology* **37**, 233-256 (1986).
143. Snyder, M. in Advances in Experimental Social Psychology (ed. Berkowitz, L.) 247-305 (Academic Press, 1984).
144. Darley, J.M. & Fazio, R.H. Expectancy confirmation processes arising in the social interaction sequence. *American Psychologist* **35**, 867-881 (1980).
145. Wikipedia. in Wikipedia (Wikimedia Foundation, San Francisco, 2012).
146. Carroll, R. in The Skeptic's Dictionary (The Skeptic's Dictionary, 2010).
147. Tolstoy, L. What is art?: and essays on art (ed. Maude, A.) (London, New York, Oxford university press, H. Milford, London, New York

London, 1932).
148. Ross, L., Lepper, M.R. & Hubbard, M. Perseverance in Self-perception and Social Perception: Biased Attributional Processes in the Debriefing Paradigm. *Journal of Personality and Social Psychology* **32**, 880-892 (1975).
149. Nestler, S. Belief perseverance: The role of accessible content and accessibility experiences. *Social Psychology* **41**, 35-41 (2010).
150. Nyhan, B. & Reifler, J. When Corrections Fail: The Persistence of Political Misperceptions. *Political Behavior* **32**, 303-330 (2010).
151. Lewandowsky, S., Ecker, U.K.H., Seifert, C.M., Schwarz, N. & Cook, J. Misinformation and its correction: Continued influence and successful debiasing. *Psychological Science in the Public Interest* **13**, 106-131 (2012).
152. Trevors, G.J., Muis, K.R., Pekrun, R., Sinatra, G.M. & Winne, P.H. Identity and epistemic emotions during knowledge revision: A potential account for the backfire effect. *Discourse Processes* **53**, 339-370 (2016).
153. Schwarz, N., Sanna, L.J., Skornik, I. & Yoon, C. Metacognitive experiences and the intricacies of setting people straight: Implications for debiasing and public information campaigns. *Advances in Experimental Social Psychology* **39**, 127-161 (2007).
154. Johnson, H.M. & Seifert, C.M. Sources of the Continued Influence Effect: When Misinformation in Memory Affects Later Inferences. *Journal of Experimental Psychology: Learning, Memory, and Cognition* **20**, 1420-36 (1994).