Chapter 5
Inferences and Human Inference Abilities

By Dr. Charles Wallis

Last Revision: 9/12/2020

Chapter Outline

5.1 Characterizing Inferences

Most critical thinking textbooks begin with a discussion of arguments. A central tenet structuring this text, however, is that one must understand and teach critical thinking tools, like the use and evaluation of arguments, against a background of our native patterns of thought. Appreciating the tools introduced in a critical thinking course requires an understanding of their relationship to our normal cognition. Thus, this text and lectures begins with a discussion of the native and artifactual means by which our brains gather information (information ecosystems). In this chapter the text and lectures turn to our native information processing resources and the information processing strategies. After a discussion of human inference abilities in this chapter and lecture the next chapter turns to arguments, human artifacts that model inferences.

5.1.a The Pervasiveness of Inferences

Humans constantly make inferences. Sometimes humans make inferences with a full conscious awareness of the information and the inference steps. For instance, when people balance their check book, they consciously follow a series of steps to consciously manipulate information in their working memory and on paper. People also make inferences during which they possess only a partial conscious awareness of what information their brains use and what inferential steps their brain makes—call these inferences semi-conscious inferences. When one drives, one makes many inferences, inferences about one's speed, distance from other cars, one's current position, relative position, etc.. Some of the information one utilizes makes its way into consciousness as do some of the inferential steps. However, not all of the information, nor all the inferential steps, go through conscious processing. One might notice one's distance from other cars if that distance suddenly or unexpectedly changes. A driver might consciously infer that they need to ease off the gas in order to keep a safe distance. But, often drivers do not maintain conscious awareness of all such information, nor do drivers always consciously infer that the situation requires a correction. Finally, one makes huge numbers of inferences during which the information and the inference steps never enter into consciousness. When one identifies an object using vision, one's brain makes a series of extremely complicated inferences using information that never enters consciousness. Light sensitive cells in the back of the eye called rods and cones gather information about the presence or absence of light reflected from objects in the environment. The information collected by rods and cones consists (to oversimplify) in a two-dimensional collection (array) of values. Thus, the brain starts with information much like the array of values collected by the light sensors at the back of a digital camera. The brain uses this information to determine the outlines of the scene. It infers how those outlines go together to form objects. It also reconstructs the relative positions of these objects in the third dimension—depth. Considering the inferential nature of even such seemingly ever-present tasks as vision, therefore, can help one to appreciate just how pervasive a role inferences play in one's everyday life. One can likewise see that while humans perform some inferences consciously, they also perform many inferences only semi-consciously (only partially utilizing consciousness), and they perform many, many inferences unconsciously (with neither the specific information, nor the inferential steps ever reaching consciousness).

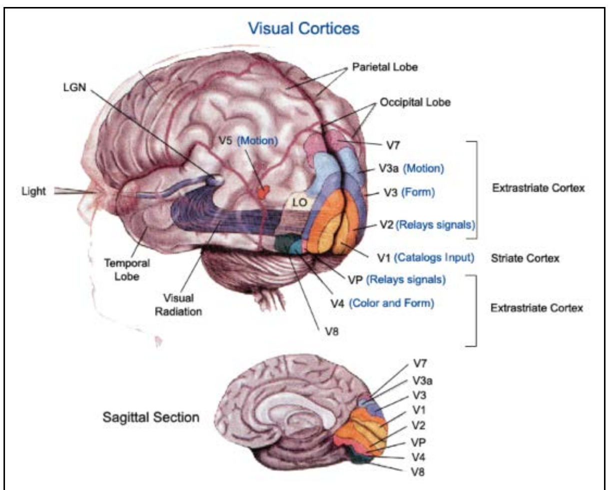

Diagram depicting the primary visual pathway responsible for the initial processing of visual information. Under normal conditions, humans can process the information about reflected and projected light captured at the back of the retina into a 3-D representation of the visual scene and recognize objects in that scene in about 300 milliseconds (about 1/3 of a second). Only the final results of this incredible series of inferences ever make it into working memory and consciousness. From: manumissio

People often fail to appreciate the pervasiveness of inference in their everyday lives.   Even while reading this text, you are making inferences at an unconscious level. Your brain processes the information about the light projected from your computer screen or reflected from your paper into information about letters and their relative positions. Your brain combines these letters into words and integrates these words into sentences. Finally, your brain determines the meanings of those sentences. The inferences your brain makes to extract the content from the sentences on this page seems effortless, but these inferences are actually quite complex. Huge parts of your brain continuously process visual stimuli in order to generate explicit and available representations of objects, properties, events, and relations in the environment.

Of course, not all inferences that you perform occur unconsciously. Some inferences, like when you multiply numbers using the Hindu-Arabic numeral positional method, involve consciously transforming information in a step-by-step fashion. Still other inferences occur with only partial conscious awareness. For example, when you walked to class today, you made a number of semi-conscious inferences: You parked your car, inferring that it was safe to leave it in the spot you chose. You inferred that class would be in the same place as always. You inferred that you could follow the same sidewalk you took last time. You saw a door, inferred that it pulls open. At any point, you might have had to alter these actions because your inferences proved incorrect. For instance, if you parked and then noticed that there was a sign saying "no student parking today" at the entrance to the lot, you would infer that you needed to move your car and it was not safe to leave it in the spot you chose.

5.1.b What are Inferences?

Despite the ubiquitous nature of inferences in human cognition, most people would have difficulty stating (1) the nature of inferences, (2) why people need them, or (3) what features or outcomes one might want to optimize in inferences. This section presents some answers to these three questions. So, what are inferences? In the most general sense, inferences are transformations of information available and explicit to a person or to some cognitive process. These transformations use available and explicit information to create new available and explicit information. More precisely, inferences take two general forms; some inferences create information previously unavailable to the person (induction), while other inferences make inexplicit information explicit and available for use (deduction).

This characterization might seem abstract to the point of vacuity to some readers, so let us examine the notion a little more carefully with a few examples to help make things more concrete. For starters, what do I mean by the notions of information being explicit and available? An analogy might help clarify: Suppose you go to a store looking for an item. A store may have the item, but the item comes in the wrong size, the wrong form, or it costs more money than you can afford. The store has the item, but not in a form or price that allows you to use it. Implicit information exists in a person's brain, but not in a readily usable form, just like the incorrectly sized or too expensive item in the store. Therefore, one's brain cannot use implicit information directly, without further modification. In contrast, a person explicitly possesses information if they have it stored somewhere in a fashion that allows them to use it for the task at hand. Put another way, one's brain encodes that information in a manner that allows one's brain to use it directly without further modification.

Consider two different ways of writing *pi*: π and 3.14159…. 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 allows one to explicitly refer to that constant, say when writing the equation for the area of a circle: Area = $\pi \cdot r^2$. 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.

What about the notions of available and unavailable information? Returning to the store analogy; the store may have the item just as needed—but if the employees have not put that item on the shelf it is not available for the moment. Consumers cannot buy the item unless someone makes it available by placing it on the shelf. Similarly, one might have information stored somewhere, but one cannot recall it in order to use it—that information is unavailable to the person at the moment. Of course, the store simply may not stock the item at all. This too will make the item unavailable. Likewise, one may not have the information stored in one's brain at all, in any form, making the information unavailable.

Thus, information counts as explicit only if one's brain encodes it in a manner that allows for the direct utilization of the information without further modification. Information counts as available only if the information is both encoded explicitly **and** accessible immediately for use. Explicit information, therefore, is available only if it is immediately accessible for use by some inference process. Implicit information always proves unavailable, just as information not encoded by one's brain proves unavailable.

Thus, one can think of information as falling into three classes; explicitly encoded, implicitly encoded, and not encoded. When one's brain has explicitly encoded information, that information exists in a form that facilitates its use in inferences. Such explicitly encoded information may prove either available for immediate use in inferences, or it may prove unavailable in which case it remains inaccessible for inferences at that time. When one's brain has implicitly encoded information that information might exist in the brain, but not in a form that facilitates its use in inferences. As a result, the information remains unavailable for inferences. Finally, when one's brain does not encode information, that information does not exist in the brain and as such proves unavailable for inferences. The table below depicts the different relationships between availability and encoding.

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

It is probably easiest for students to see the notions of explicit and available information in working memory. For instance, suppose your instructor came into your English class and said, "Guten morgen Klasse. Öffnen Sie bitte Ihre gelben Bücher." The instructor has given you some information. However, unless you speak German that information is neither explicit nor available. In contrast, suppose your instructor tells you, "Good morning class. Please open your yellow books." Now, the information is explicit and available for you in your working memory. Consider another case.

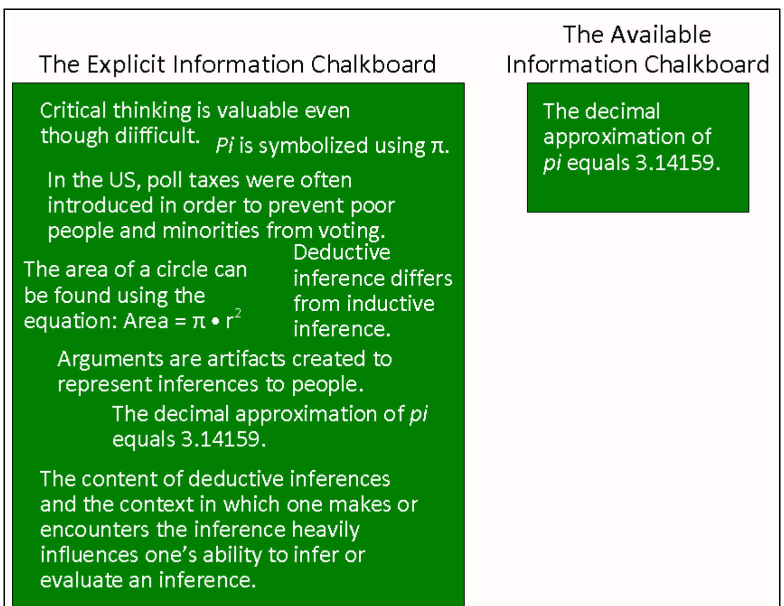

Diagram illustrating the notion of explicit information and available information.

Imagine trying to remember the name of the bias that occurs when people preferentially seek out (or interpret) information to confirm their existing attitudes or beliefs. You feel like the name is on the tip of your tongue, but you just cannot recall. In all likelihood the name is explicitly encoded in your long-term memory, but it remains unavailable because your brain temporarily cannot transfer that information from long-term to working memory.

One might think about the explicit and available information by imagining two chalkboards: On one huge chalkboard in a back room one records everything that one knows in a manner that makes the information useful for solving problems. The other, much smaller chalkboard has a different purpose, whenever one needs a piece of information, one goes to the back room, finds the information on the big chalkboard, writes that information down on the smaller chalkboard, and returns to the front room to work—the information on the smaller chalkboard then becomes available for use. One cannot copy information on the small chalkboard that does not already exist as explicit information on the large chalkboard. However, one can use explicit information from the big chalkboard to make new information explicit for further use or to transfer to the big board. The diagram (above) shows the various relationships we've just discussed.

5.1.c What are the Functions of Inferences

```
 13
125
x 6
---
750
```

The Hindu-Arabic positional method transforms explicit and available multiplicand information into explicit and available information about their product.

Needless to say, even the smartest and best informed person does not have every bit of information they need written on their big chalkboards. A person's big chalkboard may not integrate all its information particularly well. Likewise, one's information might not even prove consistent take as a whole. **Inferences function to help people to adapt to the world by transforming information, by generating new information, and even by allowing one to discover bad information and inconsistent information in order to correct or discard those inconsistencies and inaccuracies**. One can find a simple illustration of inference as a transformational process in the way one takes two numbers and uses the Hindu-Arabic positional technique to generate their product. One has explicit and available information for each multiplicand and one simply transforms that information to create an explicit and available representation of their product.

Thus, inferences occur at many levels; unconsciously as with visual recognition, semi-consciously as with inferences about routes to take to class, as well as explicitly and consciously as when solving a math problem. Similarly, information counts as available and explicit for these processes if the processes themselves can use the information in its current form. For example, the light reflectance information collected by photosensitive cells at the back of the eyes becomes available and explicit to one's unconscious visual processing system. In contrast, only the products of visual processing are available to one's conscious mind; the raw light reflectance information collected by photosensitive cells never becomes available to conscious processes. In short, making an inference and having available and explicit information do not necessarily mean conscious information or conscious inferences. Inferences are just information transformations and explicit, available information is just information encoded so that someone **or** some cognitive process can utilize it.

5.1.d 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. For instance, if you correctly believe that an assignment is due on Tuesday, you can adjust your schedule so that you complete the assignment by its due date. In contrast, if you incorrectly believe the assignment is due on Wednesday, you might well fail to complete the assignment by its due date. If one's inference strategies preserve truth so that from true initial information one generates more true information, then one can depend upon the products of those inference strategies.

However, an inference strategy might optimize one or more other features thereby making it better than other potential inference strategies. For instance, inferential power proves very desirable in inferences. You might recall Sherlock Holmes' amazing abilities to make remarkably unobvious inferences on very little information. Such powerful inferences prove both necessary and desirable in everyday life in that these inferences can greatly extend one's initial knowledge. Every time one utilizes information from one's past experiences to guide one's actions in the present one uses inference to extend one's knowledge of the past into knowledge in the future. Likewise, when one's brain generates representations of objects and their relative position in the three dimensional environment, the visual system makes a series of powerful inferences to transform information about two-dimensional light values into information about the objects reflecting or projecting that light and their relative positions in three-dimensional space.

In fact, the human brain can make these inferences and recognize objects in about 300ms (1/3 of a second). This incredible speed proves important since it allows one to quickly identify threats or needed objects. Speed, therefore, represents another feature of inferences that one might want to optimize. Speed, power, and truth…all have potential value in inference strategies. Unfortunately, a given inference strategy must usually trade-off strength in one or more features for strength 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 moves in 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 has the computational resources and speed to execute such a program. Thus, the "generate-all-possible games" strategy represents a non-viable solution to the computer science 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. In fact, an American computer scientist and cryptographer named Claude Shannon (1916-2001) has proven 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.[1] The Shannon number poses a problem for the computer science student; 10^{120} moves means that the number of all possible moves in every permutation of an average chess game exceeds the number of seconds since the big bang. The computer science student's program plays wonderful theoretical chess, but would prove impossibly slow for real use.[2-4]

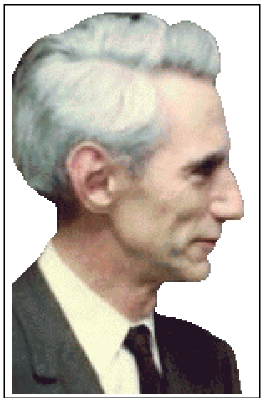
Claude Shannon (1916-2001)
From: netzspannung.org

Thus, the fourth important 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.

Indeed, of the various potentially desirable properties of inference strategies tractability might well prove the most basic. Your brain always tries to find a solution to problems—even if the solution isn't perfect. People exhibit stress responses when confronted with unsolvable problems—some researchers even suppose that subjects develop the classic stress response, learned helplessness, when confronted with unsolvable problems.[5-7] Successful problem solving and decision-making has also been shown to activate the brain's reward system, whereas failure triggers a differential response.[8-10] Such findings suggest that the brain has a solutions imperative, a strong desire to solve problems and/or avoid unsolvable problems.

5.2 Innate Reasoning Abilities: Origins and Elements

To understand human inference abilities one must first understand the origins of those abilities. Indeed, the origins of humans has greatly shaped two central elements of human inference abilities—the human brain and the native 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 native human inference abilities. I begin this section by discussing the origins of humans and proto-humans (called Hominini by scientists[11]). 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.

5.2.a Human Origins

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.[12-26] If one includes the fossil species *Sahelanthropus tchadensis*, then the period extends to over 7 million years.[22] As one follows the fossil record of Homininis one notes most species 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 the *Homo sapiens* starting around 200,000 years ago as surviving by subsistence foraging and hunting.[23, 24] The hunter-gatherer existence represents the exclusive mode of human existence until a mere 10,000 years ago when the Mesolithic era ended.[20] The Neolithic Revolution marks the end of the Mesolithic era and signals the slow spread of humans who domesticate animals, develop agriculture, and live in larger, relatively permanent groups.[25, 26]

Scientists currently hypothesize that human languages develop during the Paleolithic Era approximately 100,000 to 50,000 years ago.[27-33] 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 relatively unintuitive given the integral role that language—spoken and written—plays in contemporary life. Nevertheless, a substantial body of scientific research seems relatively homogeneous in concluding that language has only shaped human thinking for a relatively short period of time of the limited period during which.

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, the typical hunter-gatherer 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 to 7 million years of Hominini hunter-gather existence—ice ages, for example—most Hominini do not live long enough to experience much change. The mean hunter-gatherer lifespan is probably 21-37 years.[34] 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, a small, stable environment means that a hunter-gatherer likely 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. In such an environment, an individual Hominini's experiences probably represent pretty accurate samplings of the environment overall. Similarly, since the environment remains stable and homogenous during an individual hunter-gatherer's lifetime, their inferences are largely reactions to the specific problems at hand. That is, their inferences need to work for the specific content (problem) and context (situation).

Thus, researchers characterize the environment in which humans and proto-humans solve problems for something like 7 to 4.4 million years as relatively small, stable, and homogenous. In such an environment, an individual human's experiences provide them with fairly accurate samplings of the environment overall. Their experiences, in other words, are likely typical of the sorts of situations and problems that they will encounter. In similar fashion, typical Hominini

		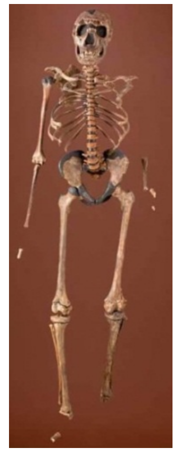	
Picture of the recovered skeleton of *Ardipithecus ramidus*. From: Wikipedia	Artist rendering of *Ardipithecus ramidus*. From: Wikipedia	*Homo erctus* skeleton. From: Smithsonianscience.org	*Homo erctus* statue at the Smithsonian museum. From: newsdesk.si.edu

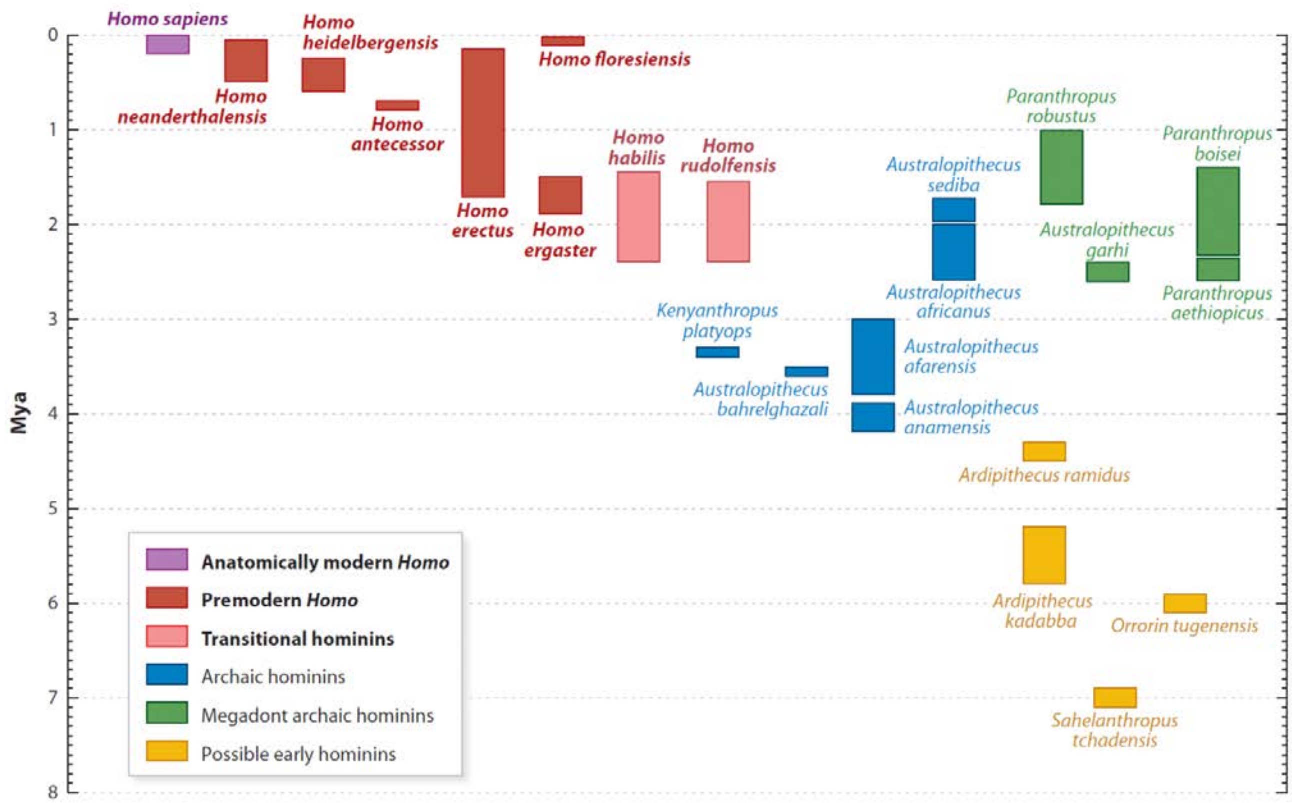

Diagram illustrating the various fossil specimens classified by their relationship to/within the Hominin taxonomy and the era during which scientists suppose that they flourished. From: Wood's and Baker's *Evolution in the Genus Homo*.[14]

problem-solving likely revolves around reactive and rather immediate responses to specific contents (problems) and in specific contexts (situations). Most problems probably involve objects and events in the immediate physical environment and in the current moment. Researchers have found very little evidence to suggest that most species of Hominini plan far into the future. Rather, they probably live in the here-and-now. Likewise, nearly all species of Hominini have extremely limited abilities to share information and to externalize information. In other words, their inferences must rely primarily upon their own sensory information in combination with long-term and working memory.

5.2.b Two Elements of Inference Ability: The Brain and The 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, how and when they collaborate, when they fail, and when they fail to interact with each other.

5.2.b.1 The Brain: Conscious vs Unconscious Inference

Each of the brain's native strategies has its strengths and weaknesses. Unconscious inference strategies are automatic, relatively fast, they can often handle larger and more complex bodies of information, and they tend to be robust. However, they can also prove relatively inflexible, especially when the problem violates one or more of the assumptions implicitly driving that approach to inferences. In contrast, conscious inferences tend to exhibit greater flexibility and adaptability. But conscious inference strategies 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 rely upon 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.

5.2.b.2 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 consciously 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.

Can you find Jason Statham's face in the coffee beans? From: The Huffington Post

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 can actually enter into working memory at any given moment. In 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 picture (left above) illustrates just how little of the visual scene can actually make it into consciousness at any one time. You may seem to see all of the coffee beans in the picture, but do you see the face of Jason Statham? Most people have difficulty finding Statham's face even when carefully searching the picture.

Video illustrating the dissociation between conscious perception of object position when wearing prismatic glasses and the unconscious adaptation of the brain despite no conscious adaptation. From: Youtube

Consider another example: Unconscious inference processes like vision tend to operate automatically and robustly. However, what happens when one or more of the assumptions implicit in vision is violated? The movie (left) gives an excellent example of how even vision can prove unreliable when a situation violates even one assumption implicit in its operation. The movie also illustrates the difficulties in modifying automatic and unconscious processes when some situation violates their underlying assumptions. Even though the basketball players can easily detect the problem resulting from the shift in their visual image, they cannot simply and immediately adapt their shooting. When practice allows them to adapt, removing the glasses again causes them to miss their shots. 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*:[35]

> 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.[36-45] Additionally, situational and idiosyncratic factors like familiarity during development (humans exhibit a genetic disposition towards incest avoidance), dissimilarities in major histocompatibility complex (humans appear to find potential suitors with different immune responses more attractive), hormone levels (fertility cycles in women appear to affect the features that drive attraction in both genders)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.[41, 42, 46-51]

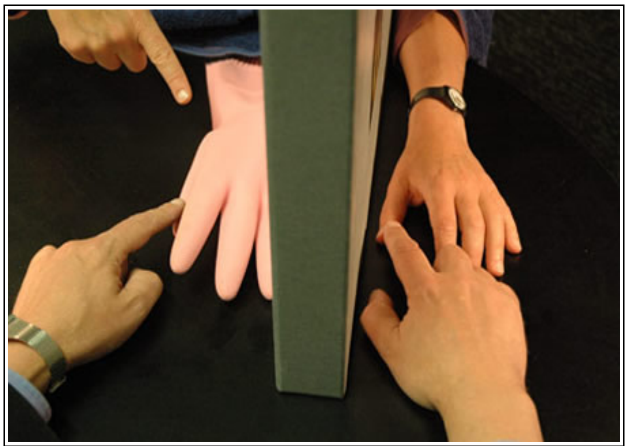

Video illustrating an illusion of bodily location. From: Youtube

Consider the unconscious processing involved in identifying one's body and its place in space relative to other objects. Everyone has had the experience of identifying some object to grab, looking away while grabbing it, and fumbling the pick-up. The video (left) shows just how dissociable conscious perceptions of our body are from our brain's inferences about the locations of our body parts.

5.2.b.3 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, then, 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 functions as a simple container in which the brain stores shorter-term memories. This supposition, in fact, does not reflect 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 date back to a 1974 paper by Allen Baddeley and Graham Hitch.[52] 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."[52] (p.35) By 2003 Baddeley refines his initial model into 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

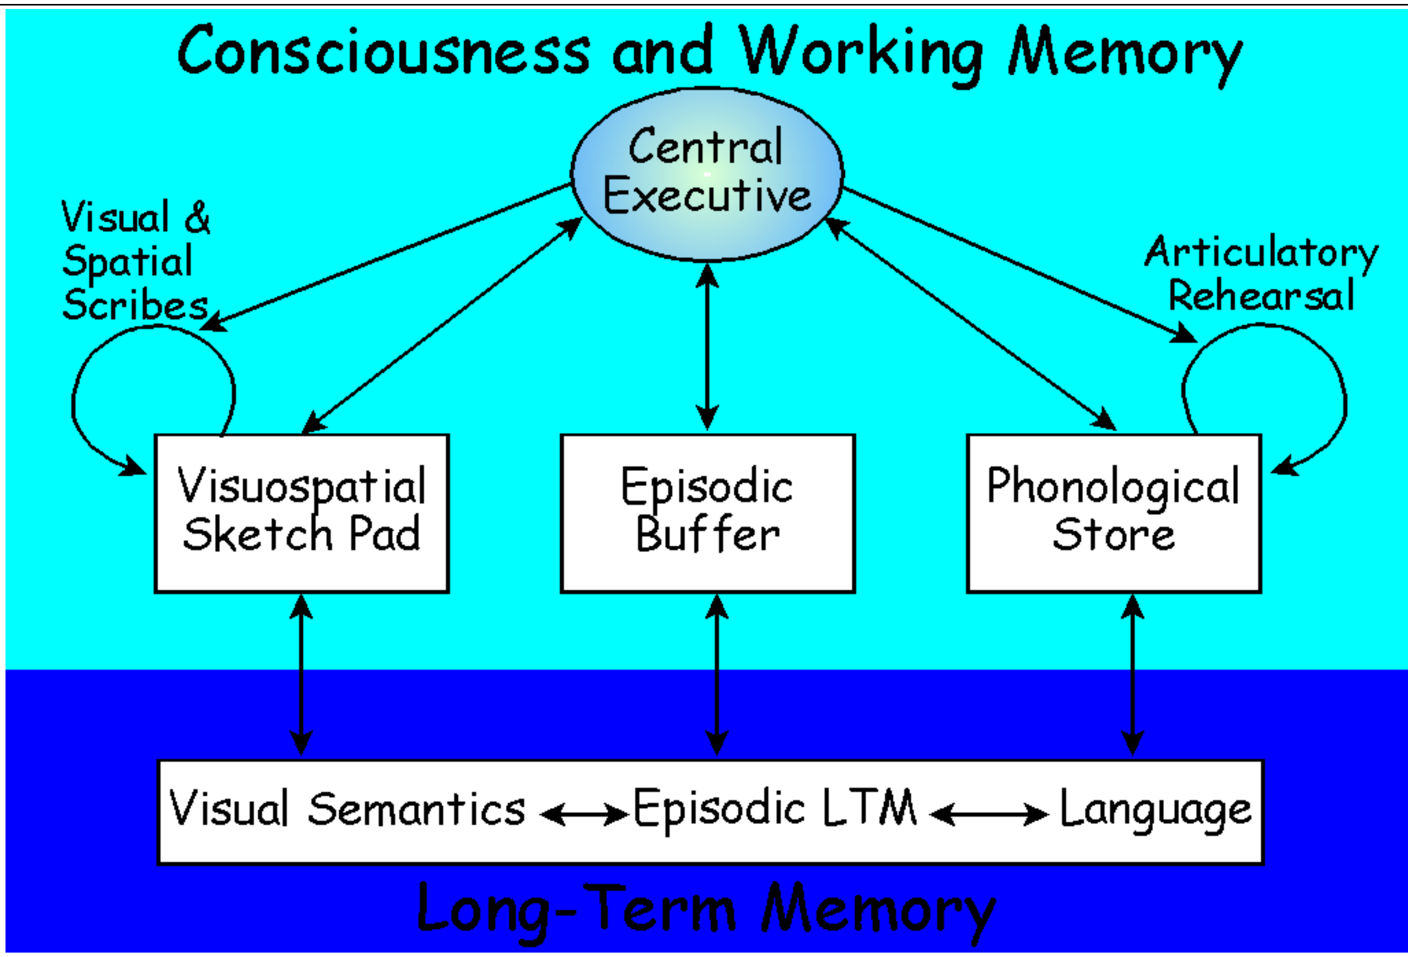

Diagram depicting the Baddeley multi-component model of working memory adapted from his 2003 diagram.[53]

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 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.[53, 54] 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 maintained by the process of articulatory rehearsal. For example, once you hear a series of numbers you must rehearse those numbers to maintain them in the phonological loop. The final working memory store, the episodic buffer, encodes 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. The episodic buffer likewise 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 long-term memory (LTM) and working memory.[53, 54]

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 and information, coordinates processing for a task, and coordinates between tasks when multi- tasking.[53, 54]

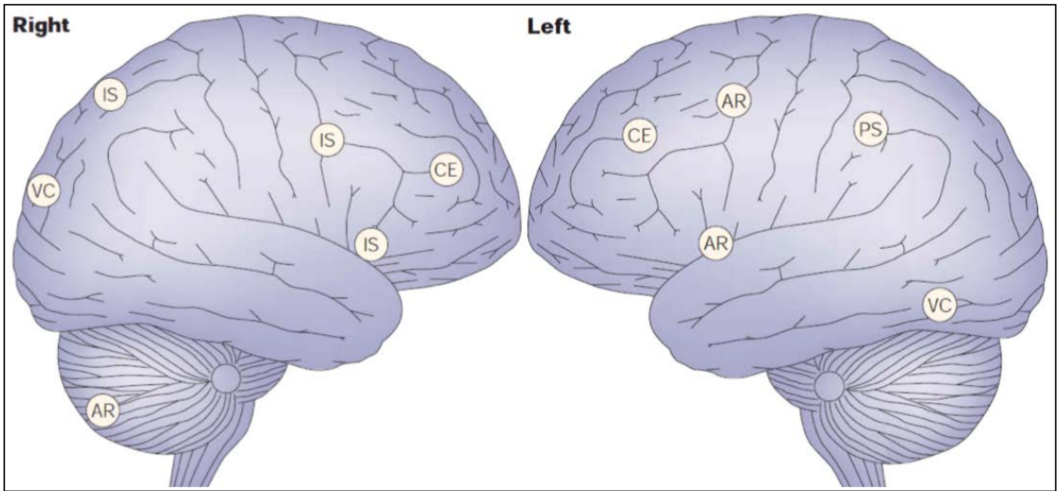

Diagram depicting the various areas of the brain associated with specific components of the multiple-components model of working memory. CE = Central Executive, AR = Articulatory Rehearsal, IS =Visual Scribe, PS= Phonological Store, VC = Visual Stetchpad From: Baddeley 2003.[53]

Students interested in the hypothesized anatomical embodiment of the various working memory modules in Baddeley's model of working memory can consult the diagram (left). Consistent with the multi-component model, sub-systems of working memory appear to have distinct anatomical centers. Also consistent with its integrative nature, working memory draws information from every cortical brain region (lobe) and both brain hemispheres as well as the cerebellum.[53, 54]

5.2.b.3.a Working Memory is Relatively Small

Earlier chapters note that 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 five 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.

5.2.b.3.b Limits on Amount and Complexity of Information in Working Memory

Measures of working memory that indicate a capacity ranging between five and nine items predate the concept of working memory itself.[55] Probably the most famous measure of working memory capacity appears in George Miller's 1956 paper, "The Magic Number Seven, Plus or Minus Two."[56, 57] Contemporary researchers tie capacity estimates for working memory 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 the time it takes to speak those words increases. 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. Cognizers can mitigate these limits somewhat by chunking information items together. For example, remembering sequences of three-digit chunks often allows one to remember more digits than remembering each digit individually.[54, 56, 58] 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 individual item complexity does not seem to affect capacity. However, visual working memory capacity appears to have a hard limit of three to four items.[54, 59-61]

Students who wonder if one might overcome the limits in information capacities just discussed 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.[62-65] 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 training stops. Moreover, like many cognitive functions, the capacity of working memory appears to decrease with age.[66, 67] Some evidence suggests 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.[68] 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 problem-solving.

5.3 Human Inference Strategies and their Typical Deployment

So far the discussion in this chapter and lecture characterize inferences and the properties that can distinguish good inferences from less useful inferences. It then distinguishes between two strategies employed by the human brain in making inferences—unconscious and conscious strategies. Important and interesting questions might occur to readers when contemplating these strategies. For instance, students might wonder, "do human inferences tend to have these properties?" Students might likewise ask "which strategy proves better?" To answer these questions, I find it useful 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. However, the term "system" proves 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.[69-71] Vinod Goel suggests the following characterization of the data from neuroscience:[70]

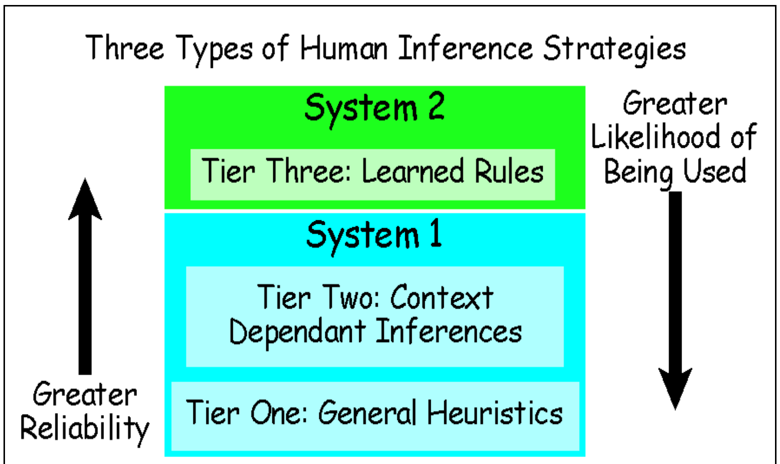

Diagram depicting the three kinds of inference strategies deployed by human beings, the relative likelihood of each being used in a given circumstance, and the relative general reliability of each kind of strategy. The two tiers (classes of inference) in blue collectively form what many psychologists 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. The inference processes in System 2 require conscious awareness to choose and conscious attention to execute. Click on diagram to display animated version.

> 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, though none has gained wide acceptance. Daniel Kahneman often uses the terms "fast" and "slow."[72] Jonathan Evans and Keith Stanovich adopt the categories "type1" and "type2."[73] Other researchers like Adam Darlow and Steven Sloman adopt the categories "intuitive" and "deliberative."[74]

5.3.a What are General Heuristics?

To understand general heuristics one must first understand 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."[75] Psychologists and computer scientists both call simple, efficient rules of thumb "heuristics" or "heuristic knowledge." One employs a heuristic when confronted with a complex problem or when one has incomplete or partially inaccurate information. In other words, a heuristic represents a strategy that trades a degree of truth-preservation in one's inference in order to gain the inferential power, speed, or tractability necessary to generate an answer. As we will see, heuristics tend to implicitly presuppose certain facts about the world and/or the problem in order to facilitate a solution. Ideally, these implicit presuppositions prove true most of the time, though such presuppositions often have significant exceptions. As a result, heuristics can work well under most circumstances, but in certain cases reliance of heuristics leads to systematic errors in reasoning.

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 features 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 famously for formulate the native judgment heuristics humans seem to employ for estimating probability and revising such estimates.[76-81] 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 of 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.

5.3.b System 1

The first two tiers or classes of inference strategies encompass strategies that represent part of the human native brain architecture and functioning. In other words, many of these inference processes are innate, developing without any explicit instruction. These inference processes also operate relatively automatically with little conscious oversight. For this reason, psychologists tend to group them together into a single system, often called "system 1" or "type 1."[73, 82, 83] Thus, 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, though in some cases learned strategies become consolidated and automated by the brain over time thereby reducing or eliminating the need for attention.[84] System 1 processes exhibit (2) autonomy in that they operate largely outside of working memory. As a result, people tend to exhibit limited conscious (2a) awareness, (2b) oversight, and (2c) insight into the operation of System 1 processes. In other words, (1) one employs a System 1 inference as a natural reaction to a situation and (2a) without having a conscious awareness of doing so. One has very little (2b) ability to affect the operation of a heuristic and (2c) very little insight into how one actually solves the problem. System 1 processes also tend to exhibit high levels of (3) contextualization and often (4) 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). One can better understand contextualization and associative processing by considering two different approaches to shooting projectiles like arrows. One often learns basic archery in a highly contextualized manner. To wit, one learns using a specific bow and set of arrows. While underlying structural features like the force transferred to the arrow from the bow, wind resistance, the arrow's mass, etc. determine the distance and accuracy of a given shot, these underlying structural features remain largely implicit in determining a given shot. Instead of explicitly representing these structural features and their relationships an archer learns to implicitly associate these features and their relationships through repetition and practice. The resulting ability of the archer becomes attuned to the specific context and content—their bow and arrows and their typical shooting conditions—and the associations they have developed through repetition and practice. As a result, the archer may well need to recalibrate if they get a different bow, different arrows, or if they are shooting in novel or unusual conditions—just as we say the basketball players do when they put on the prismatic goggles in the video in section 5.2.b.2 above.

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 probability that she will like an item by unconsciously comparing it (4) 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 without consciously inhibiting its use and explicitly employing a different strategy. You likewise exert little to no control over the information upon which the heuristic draws. Finally, since the representativeness heuristic relies heavily upon the content and context of an inferential situation (3), your shopping inference would prove quite different were you shopping for someone else, say, 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 probably will not consider possible gifts, for instance, that the site does not explicitly present 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 consist of native inferential and decision-making dispositions operating automatically in reaction to problems one encounters. These strategies exhibit very limited conscious awareness, oversight, and insight in their operations because they operate largely outside of working memory. These dispositions rely heavily upon information regarding the specific problem and the specific manner in which the problem is presented in experience. Finally, these dispositions often operate through the slow accretion of information about useful associations between the specific objects, properties, and relationships from similar contexts in the past. Thus, psychologists often characterize these processes as forming one strategy or approach to inference—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 strategies for quickly and efficiently processing 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.

5.3.c System 2

In contrast to System 1, psychologists differentiate a second class of human inference and decision-making processes that embody a different problem-solving strategy—System 2. System 2 encompasses the third and final tier or class of human inference strategies in the above diagram--consciously executed inference strategies. Unlike System 1 inference strategies, System 2 strategies tend to rely heavily on working memory and require conscious effort—both in deciding to use the strategy and in executing 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 a previous estimate of the probability of an event. Naturally, since one tends to deploy these inference strategies consciously one has much more ready access to their functioning when one uses them in problem solving. System 2 inference and decision processes tend to be learned and often leverage underlying structural features common to a class of inference or decision problems to generate a solution. Thus, psychologists categorize these processes as separate strategy for solving problems—System 2.

5.3.d The Relationship Between System 1 and System 2

The above diagram illustrates two important points about these two inferential and decision-making strategies and how processes implementing the 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. However, if one looks at the general reliability of these processes the reverse relationship holds—System 2 strategies (e.x. 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 employing 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. While reasoners can inhibit these System 1 processes and employ more appropriate System 2 strategies, fact that System 1 operates

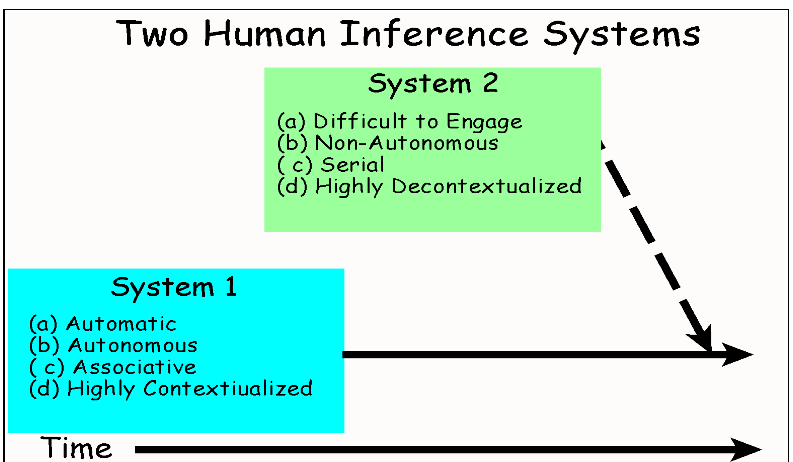

Diagram depicting the two human inference categories and their respective properties System 1 consists of both general heuristics and context dependent inference strategies. This system has evolved so that humans have innate dispositions that automatically engage when humans face a problem. System 1 processes tend to contextualize the problem 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 video depicting the potential roles of each category.

largely outside of conscious awareness coupled with the severe limitations of working memory dictate that System 2 interventions prove much less common and much more difficult. In general, humans can actively intervene only when (1) an appropriate System 2 process is readily available and either (2a) the context of the inference or decision suggests the appropriateness of the System 2 process or (2b) the failure or inadequacy of the System 1 solution becomes manifest. The relationship between System 1 and System 2 is much like that of the relationship between a train and the train's distractible engineer. Once engaged System 1, like the train, barrels down its predetermined track towards a solution. Like a train's engineer, System 2 monitors and modulates System 1 to avoid or at least minimize potential problems. But, like the distractible engineer, System 1 proves inadequate to regularly and reliably detect and deter any but the most obvious obstacles to optimal inferences and decisions. 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.

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

The last section suggests that one can distinguish System 1 inference and decision strategies from System 2 strategies by noticing that System 1 strategies represent genetically encoded dispositions to develop specific patterns of brain functioning and cognitive architecture solutions originating in evolutionary selection in response to a specific kind of environment and set of problems. In contrast, many of the most important and widespread System 2 inference strategies have their origins in the cultural heritage of the last approximately 10,000-12,000 years—with the greatest number of these strategies emerging within the last few hundred years. Reasoners must learn System 2 inference strategies—often from others--and reasoners must consciously choose to employ those strategies. 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 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.

5.4.a Deductive Inferences

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 transformational outcome. 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. Yet, in another sense, deductive systems do increase the reasoner's stock of truths. Specifically, deductive inferences transform implicit 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 in that these inferences do not increase (amplify) the number of potential truths (explicit and inexplicit information) that the reasoner possesses.

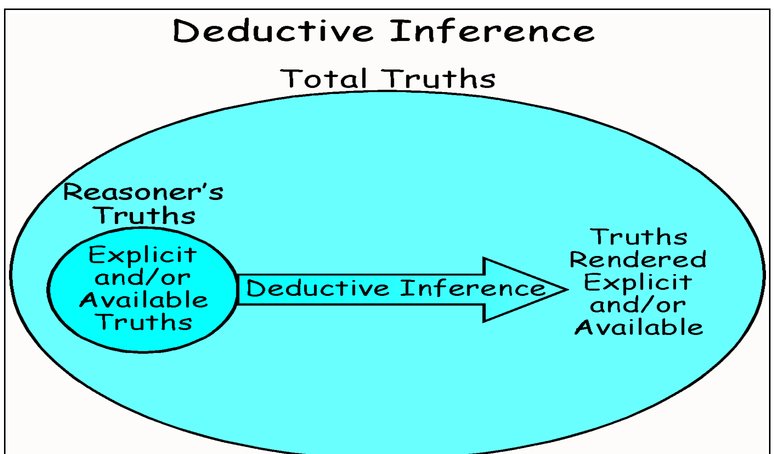

Diagram depicting deductive inferences in terms of the relationship between the initial information explicit and available to the reasoner and the information generated by the inferential transformation. Click on image to see animation.

Furthermore, deductive inferences also serve another important function: 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 at the same time. 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 also held by that individual. 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.

5.4.b Inductive Inferences

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. 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 the 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--they are generally the actual mechanisms of information transformation. For instance, inductive inferences suppose (at least) that new situations will resemble old situations in some respect and to some degree. An inductive generalization nicely illustrates this feature: Suppose that you notice that on those occasions when it rains you lose your internet connection. You might generalize your experience to the future by concluding that all times when it rains will be times you lose your internet connection. Your inference implicitly assumes that the correlation you have observed in the past between rain and lost connections will continue in the future.

So, inductive inferences extend one's stock of truths by implicitly assuming one or more structural or dynamic regularities. As a result of these implicit assumptions, inductive inference strategies introduce a degree of risk into one's inferences. Specifically, the implicit presupposition driving some inference strategy may prove false in a given inferential situation, thereby generating a false belief. Returning to the internet example, 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. In other words, even a very good inductive inference can result in a conclusion that proves false. The virtue of inductive inference, then, does not lie in the perfect preservation of truth from initial information to the conclusion. 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. In other words, the conclusion of a good inductive inference proves very likely true.

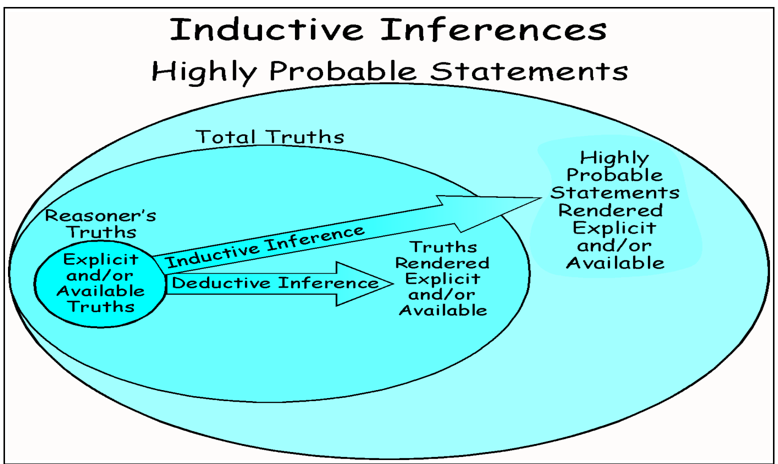

Diagram depicting inductive and deductive inferences in terms of their respective relationships between explicit and available initial information and the information generated by the inferential transformation. Click on image to see animation.

5.5 Innate Inductive Abilities, Inabilities & Biases: Inductive Inferences

Indeed, the fact that an individual hunter-gatherer's experiences typify their environment proves crucial for understanding general heuristics. For example, the representativeness heuristic acts so that a person judges an object, property, event, or relation more or less probable based upon how typical the object, property, event, or relation is in their own experiences. Specifically, the representativeness heuristic estimates the probability of the object, property, event, or relation based upon how typical the object, property, event, or relation appears to be given their concepts and schemas—the executive summaries of their experiences. In other words, the representativeness heuristic judges the likelihood of an event in the real world by judging the extent to which that event typifies the essential or salient features of one's own models and concepts.

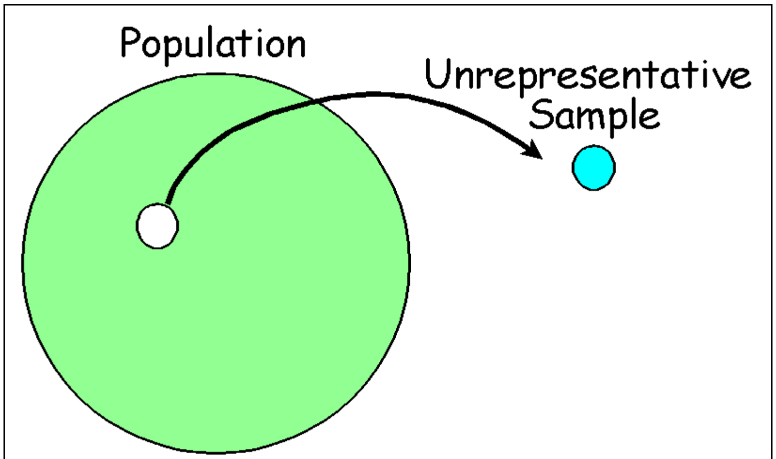

Diagram depicting (1) the relationship between a sample and a population, (2) the relationship between a representative sample and a population, and (3) the relationship between an unrepresentative sample and a population. Click diagram to see animation.

Two psychologists who have studied human inductive inference abilities, Amos Tversky and Daniel Kahneman, characterize representativeness as follows:[85] "*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) In other words, 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 partial information about objects, properties, events, or relations in some population—a sample--to information making claims about those objects, properties, events, or relations 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. 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.

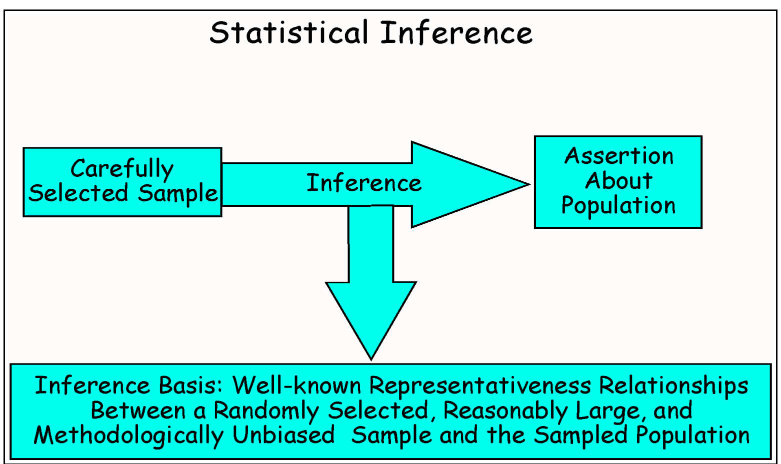

Schematic drawing depicting statistical inference and its underlying assumptions. The inference takes information about the sample and infers a similar range of values in the population based upon well-known representativeness relationships between randomly selected samples of certain sizes and the population.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

Later lectures illustrate the role of this inference strategy in statistics. 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 less than 150 years. These rules and methods act so that the dimensions and degrees of representativeness between the sample and the population remain relatively constant and high. That is, the sample consistently corresponds to the population with regard to some target object, property, event, or relation with relatively small variations. In short, the value in the sample provides an excellent basis for estimating the value in the population. 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.

5.5.a Example of Inductive Bias: The Representativeness Heuristic

The representativeness heuristic, in contrast, uses one's own concepts and schemas as samples of the population. 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 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 the object, property, event, or relation typifies the essential or salient features of one's own models and concepts. Thus, the representativeness heuristic embodies a contextualized inference strategy in that (1) the content and context (e.x. the presentation of the problem) partially determine the concepts one takes as samples, and (2) the samples one employs—one's own concepts and schemas—can prove idiosyncratic. 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.

So, 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 result, 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) tend to provide a much less representative sample, and often generate bad estimates. For example, 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 relative likelihood of dying in a terrorist attack compared to the likelihood of dying from accidental suffocation, people tend to rate terrorism more likely. However, according to the U.S. State Department 56 U.S. citizens died world-wide from terrorism in 2005,[86] while on average about 6,000 U.S. citizens die of accidental suffocation each year.[87]

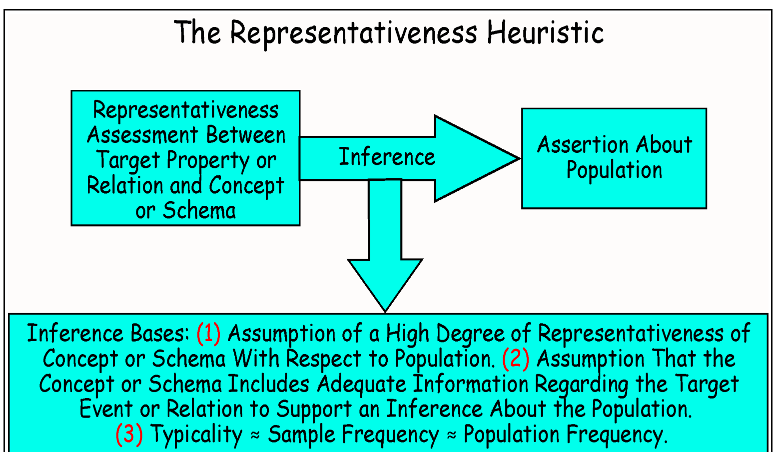

Schematic drawing depicting the representativeness heuristic inference and its underlying assumptions. The heuristic determines the probability of a property, object, event, or relation in the sample—in this case a concept or schema representation—it then infers a similar range of values in the population. The bases (underlying implicit assumptions) behind the representativeness heuristic are: (1) The assumption that the concept or schema is representative of the population. (2) The assumption of the representativeness of the concept or schema with respect to the target parameter in the population such that the value for the target parameter as given by the concept or schema is representative of the value for the target parameter in the population. Click on diagram to display video.

5.6 Innate Deductive Abilities, Inabilities & Biases: Deductive Inferences

When one turns to an examination of human innate deductive abilities, one finds two general trends. First, humans have a very limited ability to process large amounts of data or to processes complex, highly inter-related data when making inferences involving working memory. Thus, humans tend to perform poorly when formulating or evaluating complicated or long deductive inferences. Second, the content of individual arguments and the context of individual inferences (presentation, circumstance, etc.) drive human formulations of deductive inferences as well as human evaluations of deductive inferences.

5.6.a The Resources Difficulty of Deductive Reasoning

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 an argument taken from Charles Lutwidge Dodgson, better known as Lewis Carroll (1832-1898):[88]

> No interesting poems are unpopular among people of real taste.
> No modern poetry is free from affectation.
> All your poems are on the subject of soap-bubbles.
> No affected poetry is popular among people of real taste.
> No ancient poem is on the subject of soap-bubbles.
> --
> Therefore, your poetry is not interesting. (p.118)

Is the above argument a good deductive argument? Most people have almost no idea. It seems like rambling, rather 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).[89-97] 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 become too complex and involve too much information for native human reasoning abilities rather quickly.

5.6.b Context and Content Effects in Deductive Reasoning

So, 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 when formulating and evaluating arguments. Indeed, researchers have demonstrated 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. Consider the following table of arguments:

<table>
<tr><td>1
All elephants are big things.
All elephants are mammals.

Some mammals are big things.</td><td>2
No C are B
All A are B

No A are C</td></tr>
<tr><td>3
No trersnks are yrdogs.
All batgobs are trersnks.

All batgobs are yrdogs.</td><td>4
No U.S. presidents are women.
All women are people who can reproduce.

No people who can reproduce are U.S. presidents.</td></tr>
</table>

Of the four arguments, logicians would designate as invalid or deductively bad only the 3rd argument (bottom, left). Logicians would designate the other three arguments as valid or deductively good arguments. People tend to find that the difficulty in correctly evaluating these arguments increases as they move from box 1 to box 4. 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 (i.e., 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.[98-100] Arguments employing pseudo-content (meaningless word-like content) prove even more difficult for most people to intuitively 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). For instance, the example in the fourth box, though perfectly valid, seems like a bad argument to many people because both the premises and the conclusion are false. Judging the argument in the fourth box invalid illustrates a systematic bias in innate human deductive reasoning resulting from the tendency to contextualize (i.e., rely heavily on content and context) 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 disbelieved conclusions. Researchers call this tendency "Belief Bias".[101-105] 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 they try to formulate or evaluate the argument. Importantly, both the familiarity of the content, the type of content, and the relationship between content and underlying logical structure affect human performance on deductive reasoning tasks.

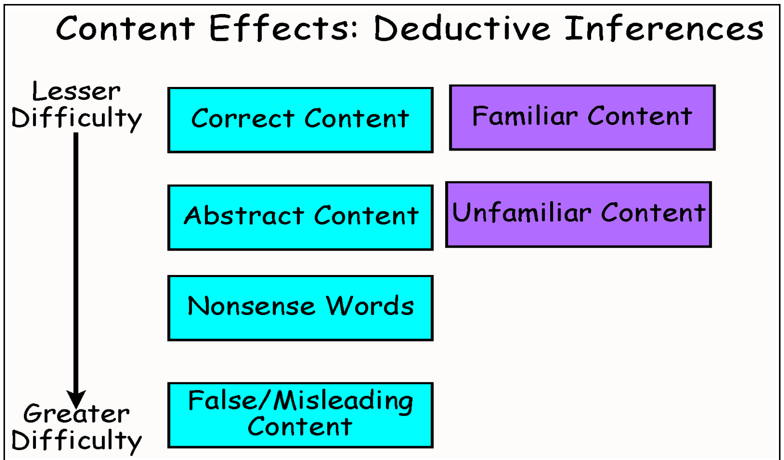

Diagram indicating the relative difficulty of making or evaluating deductive inferences with various kinds of content. The easiest types of inferences to make or evaluate involve familiar content (i.e., are about, familiar objects, properties, events, or relations) where the premises and conclusion are 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.

5.7 Context Dependent Inference Strategies

The last two sections discuss general heuristics and general tendencies of deductive inference. This section turns to the second tier of processes in System 1—context dependent strategies. One must carefully distinguish contextualization, a general feature of System 1 inference strategies, with context-dependent inference strategies. Context-dependent inference strategies form a class of inference strategies people use only in very specific contexts, and which they do not employ outside of those contexts. Belief contexts influence deductive inferences because the inference processes are influenced by the content of beliefs. Context can also affect deductive reasoning when the context triggers a context-specific inference strategy. Such strategies do not operate as general strategies. Rather, they operate in relatively specific contexts.

5.7.a Example: Conditional inferences

One of the more striking examples of a context-dependent inference strategy involves the innate human ability to reason using conditional statements. Conditional statements function to relate the truth of two component statements.

Specifically, the conditional relates the truth of the antecedent, the condition, to the truth of the consequent. For instance, the conditional sentence, "If you read this chapter, then you can better understand the lecture," claims that the truth of the antecedent--you read this chapter--insures the truth of the consequent--you can better understand the lecture. Conditional statements, despite their ubiquity and utility in general reasoning, prove difficult for humans to process. Consider the following problems first investigated by Peter Cathcart Wason, and subsequently entitled the Wason Selection Task.[106-109]

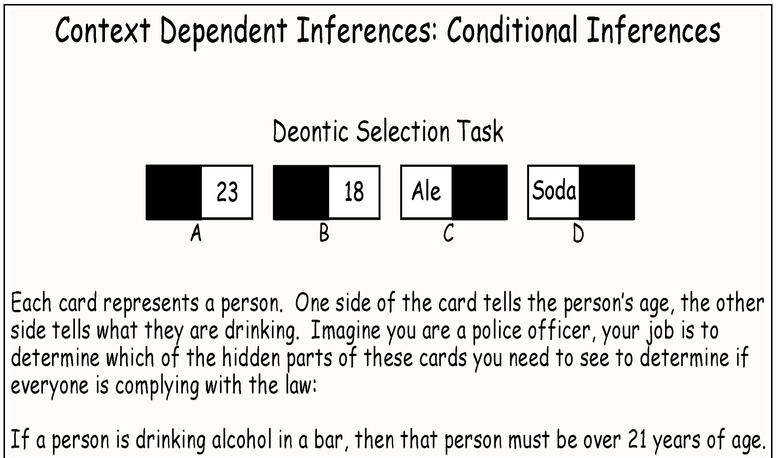

Two examples of a conditional reasoning task Wason explores in adults and that Cummins, Chao, and Cheng[110-112] in explore in development. The first example employs a deontic selection task involving a permission rule. The second task employs a standard, non-deontic conditional of the type explored by Wason. Wason's and subsequent studies support the hypothesis that people have difficulty making and evaluating conditional inferences in non-deontic cases as evinced by the strong tendency to incorrectly evaluate the truth of conditionals in non-deontic cases. In the second example, one ought to list A and D as cards that need to be checked. In contrast, people tend to have better, more robust performance on reasoning tasks involving conditionals within deontic (permission and obligation) tasks.

People tend to have better, more robust performance on reasoning tasks involving conditionals within deontic (permission and obligation) tasks, like the first example given in the video (left) and depicted in the cover image. This difference in typical human performance seems to emerge early in development and persist into adulthood. Indeed, this pattern of relative ease in conditional reasoning and evaluation tasks within deontic contexts has led researchers to suppose that performance in these contexts is either part of an innate context-specific mechanism for reasoning, or that humans possess an innate disposition to learn such rules in deontic contexts.[107, 110-113] In the second example the video presents a selection task using a standard non-deontic and relatively unfamiliar conditional. Researchers find that in non-deontic selection tasks like the second example given in the video, people demonstrate a significant inability to reason with or to evaluate conditional statements or related arguments. Moreover, difficulties that arise in these non-deontic cases appear in ordinary situations in which people perform or evaluate simple conditional inferences. In fact, people's conditional reasoning in non-deontic cases exhibits the same sorts of general content effects described earlier in reference to general deductive inference abilities.[107, 110-113]

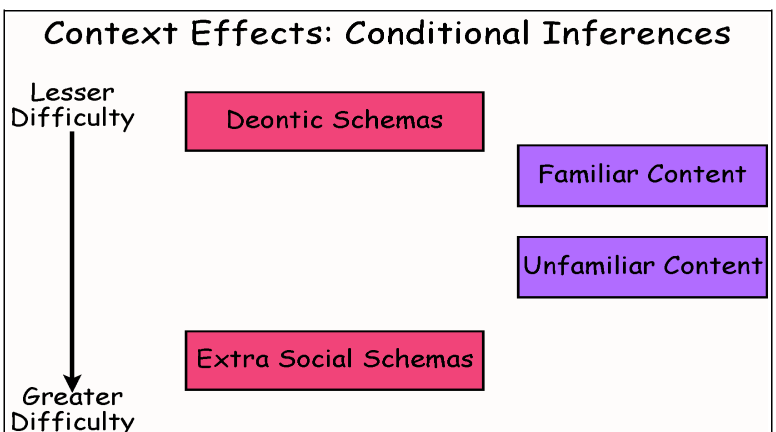

Diagram indicating the relative difficulty of making or evaluating the truth or falsity of conditionals and/or performing conditional inferences with various kinds of contexts. Click diagram to play animation.

The video (left) illustrates the relationships between a conditional argument's content and the difficulty it presents to typical humans when the try to formulate or evaluate inferences involving that conditional. The easiest types of inferences involve deontic contexts in which the conditionals and/or inferences concern cases of permission, duty, obligation, etc.. Outside of deontic contexts performance drops significantly following the pattern for other deductive inferences and evaluations. Outside deontic contexts the easiest types of conditional inferences and conditional evaluations involve familiar content (i.e., conditionals 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 conditional inferences and conditional 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 correct perform or evaluate.

5.7.b Example: Probability Assignments

A now famous problem originally formulated by Steve Selvin and often called the Monty Hall Problem provides yet another example of content-dependent inferences.[114-117] Consider the three boxes in the video and cover diagram below. Suppose that you and I will play a game many times. The goal of that game is to win as much money as possible from me. Here's how the game works. While you wait outside the room, I randomly choose one of the three boxes and put a five dollar bill inside the box. When you return, I ask you pick one of the three boxes. Suppose that I put the money in box number two and you pick box number two. I randomly choose one of the other two empty boxes, say number 3, and show you the contents, i.e., that it is empty. I discard box three, so now there are two boxes left. I push these two boxes forward and ask you if you would like to stick with your original box or switch to the other remaining box. Now, suppose that I put the money in box number two, but you choose box number one. I show you the contents (empty) of box number three and set it aside, leaving two boxes. I push these two boxes forward and ask you if you would like to stick with your original box or switch to the other remaining box. Your job is to decide which of the following three strategies will result in your winning the most money over the long-run: (1) Switch from your box to the other remaining box, (2) Stay with your original box, (3) It does not matter to your long term winnings whether you switch, so you can switch or not as it suits you.

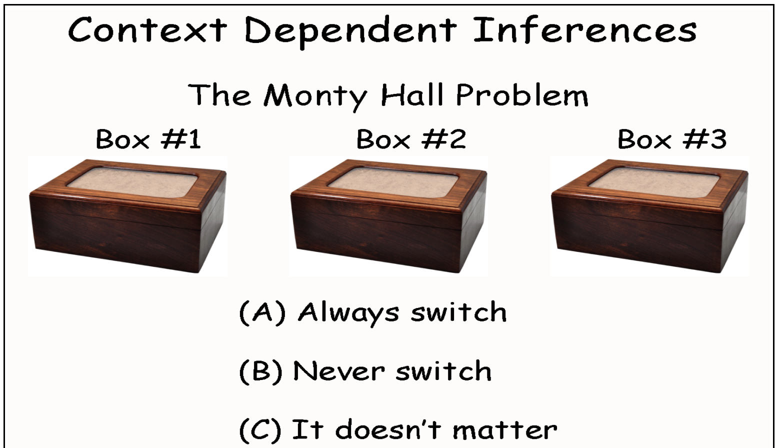

Diagram depicting the set-up and choices in the Monty Hall problem. Consider the three boxes in the diagram below. Suppose that you and I will play a game many times. The goal of that game is to win as much money as possible from me. Here's how the game works. While you wait outside the room, I randomly choose one of the three boxes and put a five dollar bill inside the box. When you return, I ask you pick one of the three boxes. In case 1 you pick number two and I put the money in number two. I randomly choose one of the other two empty boxes, 3, and show you the contents, i.e., that it is empty, and discard it. Now there are two boxes left. I push these two boxes forward and ask you if you would like to stick with your original box or switch to the other remaining box. In case 2 you choose number one, but I put the money in box number two. I show you the contents (empty) of box number three and set it aside. Now there are two boxes left. I push these two boxes forward and ask you if you would like to stick with your original box or switch to the other remaining box. Your job is to decide which of the following three strategies will result in your winning the most money over the long-run: (1) Switch from your box to the other remaining box, (2) Stay with your original box, (3) It doesn't matter whether you switch, so you can switch or not as it suits you. People are often very surprised to discover that one should always switch boxes to maximize willing in the long run. Click on diagram to play animation.

In general, people will estimate the probability of a random event by considering the number of current possibilities. Since people perceive two boxes in the context of the choice to switch or stick with a box, they estimate the probability as one out of two or 50-50; two boxes, so two chances and one five dollar bill, so one possible winner. However, the choice remains governed by the probability of the original choice, one out of three. The mechanics of the game merely disguise the fact that you choose between the contents of your one, original box and the contents of the other two boxes. How do I disguise this choice? I go through the show of revealing an empty box from the two boxes before asking you to choose. This bit of stagecraft allows you to discount that box in your calculations, despite the fact that you will always receive the contents of the two boxes you did not choose if you switch. People have a difficult time wrapping their head around this problem, so make sure you attend lecture to get Wallis' extra explanation.

5.8 Chapter Summary

This chapter characterizes inferences and discusses the various types of innate human inference strategies standardly divided by cognitive scientists into two general strategy categories—System 1 and System 2. The discussion emphasizes that these strategies evolved like the human brain itself during the hunter-gatherer phase of Hominini evolution. As such evolution has optimized these strategies for an environment that is relatively small, stable, and homogenous. In such an environment, an individual human's experiences are pretty accurate samplings of the environment overall.

Similarly, hunter-gather problem-solving was likely limited largely to reactive, relatively simple, and concrete problem-solving linked to specific contents (problems) and contexts (situations).

The combination of automaticity, limited conscious access, and contextualization in System 1 inference strategies represents an approach to inferences that typically results in relatively fast, concrete, resource sparing inferences best suited to reactive responses to environmental circumstances. Indeed, automatic inferences increase reaction time in that they respond without lengthy consciously mediated recognition or evaluation processes. Likewise, automaticity and limited conscious access minimize the need to employ the very limited resources of conscious attention in problem-solving. Finally, contextualization represents a strategy for quick and highly fluid problem-solving driven by one's current situation. Add to this a strong genetically determined disposition towards the development of such inference strategies, and the need for a long, resource intensive learning period disappears as well.

However, the advantages conferred by the system 1 problem-solving strategy depend upon (implicitly presume) a particular sort of environment that presents a particular sort of problem. When one employs automatic, contextualized inference strategies to which one has with little conscious access in environments or on problems that violate the presuppositions of the strategy, systematic errors will occur and these errors will often prove difficult to identify and correct. If these inference strategies also prove largely innate, then the reasoner will have very limited ability to alter this basic architecture. In his famous 1990 book, *Who is Rational?*, Keith Stanovich tells readers that,[83]

> Because this tendency toward the contextualization of information processing by System 1 is so pervasive, it is termed here the *fundamental computational bias* in human cognition. The fundamental computational bias is meant to be a global term that captures the pervasive bias toward the contextualization of all informational encounters. It conjoins the following processing tendencies: (a) the tendency to adhere to Gricean conversational principles even in situations that lack many conversational features (Adler, 1984; Hilton, 1995), (b) the tendency to contextualize a problem with as much prior knowledge as is easily accessible, even when the problem is formal and the only solution is a content-free rule (Evans, 1982, 1989; Evans et al., 1983), (c) the tendency to see design and pattern in situations that are either undesigned, unpatterned, or random (Levinson, 1995), (d) the tendency to reason enthymematically--to make assumptions not stated in a problem and then reason from those assumptions (Henle, 1962; Rescher, 1988), and (e) the tendency toward a narrative mode of thought (Bruner, 1986, 1990).
>
> All of these properties conjoined together represent a cognitive tendency toward radical contextualization. The bias is termed fundamental because it is thought to stem largely from System 1 and that system is assumed to be primary in that it permeates virtually all of our thinking (e.g., Evans & Over, 1996). If the properties of this system are not to be the dominant factors in our thinking, then they must be overridden by System 2 processes. (pp. 192-93)

The material in this course illustrates the historical development of and real value of the body knowledge and techniques designed to compensate for the weaknesses and biases inherent in our innate inference strategies. For the most part, this body of knowledge and techniques represent strongly decontextualized inference strategies and knowledge. However, though this cultural heritage complements and compensates for weaknesses in native human abilities, it is not a panacea for poor reasoning. The inability to utilize this body of knowledge and techniques in a consistent and pervasive fashion dramatically mitigates the potential of these techniques and knowledge. Nevertheless, one can easily observe the compounded positive (or negative) impact of individual human beliefs and decisions regarding diet, transportation, manufacturing and distributing goods and services, etc.. The consequences of individual beliefs and decisions manifest themselves in the current change occurring in the earth's climate, the dramatically increasing incidence of obesity in the United States and all its related health problems, etc.. More importantly, competent, literate, and effective thinkers and decision makers benefit from; (1) better, more highly evinced, and integrated belief systems, (2) better, more informed decisions yielding more highly-valued outcomes, and (3) a greater awareness of the world and its multifarious opportunities and possibilities together with their associated benefits and pitfalls. Likewise, societies—particularly industrialized democratic societies—rely upon informed, effective thinkers and decision makers to exist and function. With both challenges and benefits in mind, the next lectures will turn to arguments, their structure, and techniques for extracting and evaluating them from written text and spoken passages.

5.9 Key Terms

Ampliative vs Non-ampliative inferences: Ampliative inferences extend one's conclusion beyond what one's knowledge guarantees true. Ampliative inferences thereby broaden or extend our knowledge. However, in order go beyond known truths ampliative inferences must take on epistemic risk—risk that the conclusion can be false even when the premises are true. Inductive inferences, for instance, are ampliative inferences. Even when someone presents a strong inductive inference with true premises it remains possible that the conclusion, though highly likely to be true, is actually false. Non-ampliative inferences, in contrast, work to render otherwise implicit and unavailable information explicit and available. Since non-ampliative inferences seek to enlarge the body of explicit and available true information given a reasoner's current information, non-ampliative inferences act to optimize truth-preservation across the informational transformation. Thus, non-ampliative inferences do not increase epistemic risk through their operation.

Automatic Inference and Decision-Making Strategies: Automatic inference and decision making strategies engage in reaction to problems a reasoner encounters without the reasoner having to consciously evaluate the problem or choose the strategy. For example, general heuristics like the representativeness heuristic operate automatically in reaction the situations in which one must estimate likelihoods.

Autonomous Inference and Decision-Making Strategies: Autonomous inference and decision-making strategies operate without drawing significantly upon working memory resources. As a result, autonomous inference strategies tend to operate largely outside of conscious awareness. These inference processes also often sidestep the information capacity limitations of working memory and perform in a relatively uniform manner across different levels of fluid intelligence.

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 those processes and transform that initial information 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…er, hopefully.

Inferential Power: Inferential power refers to the property of an inference strategy to generate information that goes beyond the explicit and implicit information guaranteed to be true given a reasoner's initial information. Thus, powerful inference strategies are also ampliative inference strategies that broaden or extend a reasoner's knowledge beyond what was guaranteed to be true before the inference. As a result, powerful inferential strategies must take on a degree of epistemic risk—risk that the conclusion can be false even when the making an inference from true initial information. For example, when you infer that you can make it to school before your class starts you cannot guarantee that you will not get into an accident, develop car trouble, or run into unusually heavy traffic.

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 Representativeness Heuristic: The representativeness heuristic 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 collection of inference and decision-making processes sharing a common problem-solving strategy evolved so that humans develop native dispositions that automatically engage when encountering inference and decision problems. The strategy tends to contextualize these problems by relying upon information about the specific context and content of the current problem. System 1 inference processes function relatively independently of working memory and therefore 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. However, these processes also prove less susceptible to the limitations on the amount and complexity of information inherent in working memory and perform relatively uniformly across individual variations in fluid intelligence. These processes also often operate through implicit associations and in a relatively fast manner.

System 2: System 2 inference and decision-making strategies, in contrast, to System 1 consist primarily of learned knowledge and techniques. Strategies in System 2 do not automatically engage when a reasoner faces a problem. Indeed, they often prove difficult to engage. System 2 strategies require conscious awareness and oversight to operate and tax working memory resources significantly. 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—strategies explicitly driven by the underlying structural features of problems. System 2 inference strategies also provide humans with greater conscious insight and oversight into their inference and decision-making processes.

Tractability: Tractability refers to the property of an inference strategy to complete the inference in a reasonable amount of time (or even at all) utilizing only the available cognitive resources. For instance, inferring the product of two eight digit numbers within seconds using only working memory proves to be an intractable strategy for most people. However, using the Hindu-Arabic positional method to compute the produce using pen and paper proves tractable. Likewise, using a calculator also proves tractable.

Working Memory: Contemporary theories of working memory characterize working memory as a brain system that functions to hold and manipulate information during conscious problem solving and decision making. Working memory can incorporate information from different modalities. Two important properties of working memory are: (1) The contents of working memory are consciously available. (2) Working memory capacity is extremely small both in terms of the amount of information and in terms of the complexity of information.

5.10 Bibliography

1. Wikipedia. in Wikipedia (The Wikimedia Foundation, San Francisco, CA, 2013).
2. Wikipedia. in Wikipedia (The Wikimedia Foundation, San Francisco, CA, 2012).
3. Wikipedia. in Wikipedia (Wikimedia Foundation, San Francisco, 2012).
4. Wikipedia. in Wikipedia (Wikimedia Foundation, San Francisco, 2012).
5. Mikulincer, M. Coping and learned helplessness: effects of coping strategies on performance following unsolvable problems. *European Journal of Personality* **3**, 181-194 (1989).
6. Snyder, M.L. & Frankel, A. Egotism versus learned helplessness as an explanation for the unsolvable problem effect: Comment on Kofta and Sédek (1989). *Journal of Experimental Psychology: General* **118**, 409-412 (1989).
7. Dor-Shav, N.K. & Mikulincer, M. Learned helplessness, causal attribution, and response to frustration. *Journal of General Psychology* **117**, 47-58 (1990).
8. Tricomi, E., Delgado, M.R., McCandliss, B.D., McClelland, J.L. & Fiez, J.A. Performance Feedback Drives Caudate Activation in a Phonological Learning Task. *Journal of Cognitive Neuroscience* **18**, 1029-1043 (2006).

9. Li, J., Delgado, M.R., Phelps, E.A. & Smith, E.E. How instructed knowledge modulates the neural systems of reward learning. *Proceedings of the National Academy of Sciences of the United States of America* **108**, 55-60 (2011).
10. Nakai, T., Nakatani, H., Hosoda, C., Nonaka, Y. & Okanoya, K. Sense of Accomplishment Is Modulated by a Proper Level of Instruction and Represented in the Brain Reward System. *PLoS ONE* **12**, 1-20 (2017).
11. Wikipedia. in Wikipedia (Wikimedia Foundation, San Francisco, CA, 2013).
12. 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).
13. 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).
14. Wood, B. & Baker, J. Evolution in the Genus Homo. *Annual Review of Ecology, Evolution & Systematics* **42**, 47-69 (2011).
15. Wood, B. & Harrison, T. The Evolutionary Context of the First Hominins. *Nature* **470**, 347-352 (2011).
16. Hanson, B. Light on the Origin of Man. *Science* **326**, 60-61 (2009).
17. Lovejoy, C.O. Reexamining Human Origins in Light of Ardipithecus ramidus. *Science* **326**, 74, 74e1-74e8 (2009).
18. Suwa, G. et al. The Ardipithecus ramidus Skull and Its Implications for Hominid Origins. *Science* **326**, 68, 68e1-68e7 (2009).
19. White, T.D. et al. Ardipithecus ramidus and the Paleobiology of Early Hominids. *Science* **326**, 75-86 (2009).
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 (Wikimedia Foundation, San Francisco, CA, 2013).
23. Wikipedia. in Wikipedia (The Wikimedia Foundation, San Francisco, CA, 2016).
24. Wikipedia. in Wikipedia (The Wikimedia Foundation, San Francisco, CA, 2016).
25. Wikipedia. in Wikipedia (The Wikimedia Foundation, San Francisco, CA, 2016).
26. Wikipedia. in Wikipedia (The Wikimedia Foundation, San Francisco, CA, 2016).
27. Fitch, W.T. The Evolution of Language (Approaches to the Evolution of Language) (Cambridge University Press, Cambridge, England, 2010).
28. Wikipedia. in Wikipedia (The Wikimedia Foundation, San Francisco, CA, 2015).
29. Wikipedia. in Wikipedia (The Wikimedia Foundation, San Francisco, CA, 2015).
30. Wikipedia. in Wikipedia (The Wikimedia Foundation, San Francisco, CA, 2015).
31. Wikipedia. in Wikipedia (The Wikimedia Foundation, San Francisco, CA, 2015).
32. Wikipedia. in Wikipedia (The Wikimedia Foundation, San Francisco, CA, 2015).
33. Hillert, D. The Nature of Language: Evolution, Paradigms and Circuits (Springer Science + Business Media, New York, NY, US, 2014).
34. Gurven, M. & Kaplan, H. Longevity Among Hunter-Gatherers: A Cross-Cultural Examination. *Population and Development Review* **33**, 321-365 (2007).
35. Bargh, J.A. & Chartrand, T.L. The Unbearable Automaticity of Being. *American Psychologist* **54**, 462 (1999).
36. Bronstad, P.M., Langlois, J.H. & Russell, R. Computational Models of Facial Attractiveness Judgments. *Perception* **37**, 126-142 (2008).
37. Bull, R. & Shead, G. Pupil Dilation, Sex of Stimulus, and Age and Sex of Observer. *Perceptual and Motor Skills* **49**, 27-30 (1979).
38. Hess, E.H. Attitude and Pupil Size. *Scientific American* **212**, 46-54 (1965).
39. Hess, E.H. The Role of Pupil Size in Communication. *Scientific American* **233**, 110-119 (1975).
40. 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).
41. 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).
42. 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).
43. Anthony, C.L. Facial Attractiveness: Evolutionary Based Research. *Philosophical Transactions of the Royal Society B: Biological Sciences* **366**, 1638-1659 (2011).

44. Bzdok, D. et al. ALE Meta-analysis on Facial Judgments of Trustworthiness and Attractiveness. *Brain Structure & Function* **215**, 209-223 (2011).
45. 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).
46. 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).
47. 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).
48. Havlicek, J. & Roberts, S.C. MHC-Correlated Mate Choice in Humans: A Review. *Psychoneuroendocrinology* **34**, 497-512 (2009).
49. 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).
50. Vartanian, L.R. & Hopkinson, M.M. Social Connectedness, Conformity, and Internalization of Societal Standards of Attractiveness. *Body Image* **7**, 86-89 (2010).
51. Zaki, J., Schirmer, J. & Mitchell, J.P. Social Influence Modulates the Neural Computation of Value. *Psychological Science* **22**, 894-900 (2011).
52. AD, B. & GJ, H. in Recent Advances in Learning and Motivation (ed. Bower, G.) 47-90 (Academic Press, New York, NY, 1974).
53. Baddeley, A. Working memory: looking back and looking forward. *Nature Reviews Neuroscience* **4**, 829-839 (2003).
54. RepovŠ, G. & Baddeley, A. The multi-component model of working memory: Explorations in experimental cognitive psychology. *Neuroscience* **139**, 5-21 (2006).
55. Brener, R. An experimental investigation of memory span. *Journal of Experimental Psychology* **26**, 467-482 (1940).
56. Miller, G.A. The Magical Number Seven, Plus or Minus Two. *Psychological Review* **63**, 81-97 (1956).
57. Miller, G.A., Galanter, E. & Pribram, K.H. Plans and the Structure of Behavior (Henry Holt and Co, New York, NY US, 1960).
58. 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).
59. 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).
60. 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).
61. 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).
62. 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).
63. 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).
64. Ando, J., Ono, Y. & Wright, M.J. Genetic Structure of Spatial and Verbal Working Memory. *Behavior Genetics* **31**, 615-624 (2001).
65. Holmes, J., Woolgar, F., Hampshire, A. & Gathercole, S.E. Are working memory training effects paradigm-specific? *Frontiers in Psychology* **10** (2019).
66. 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).
67. Bo, J. Age-Related Declines in Visuospatial Working Memory Correlate With Deficits in Explicit Motor Sequence Learning. *Journal of Neurophysiology* **102**, 2744-2754 (2009).
68. 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).

69. De Neys, W. & Goel, V. in Neuroscience of decision making. (eds. Vartanian, O. & Mandel, D.R.) 125-141 (Psychology Press, New York, NY, 2011).
70. Goel, V. Anatomy of deductive reasoning. *Trends in Cognitive Sciences* **11**, 435-441 (2007).
71. 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).
72. Kahneman, D. Thinking, Fast and Slow (Thinking, Fast and Slow New York, NY, 2011).
73. Evans, J.S.B.T. & Stanovich, K.E. Dual-Process Theories of Higher Cognition. *Perspectives on Psychological Science* **8**, 223-241 (2013).
74. Darlow, A.L. & Sloman, S.A. Two systems of reasoning: Architecture and relation to emotion. *WIREs Cognitive Science* **1**, 382-392 (2010).
75. Wikipedia. in Wikipedia (The Wikimedia Foundation, San Francisco, CA, 2015).
76. Kahneman, D. & Tversky, A. Subjective Probability: A Judgment of Representativeness. *Cognitive Psychology* **3**, 430-454 (1972).
77. Kahneman, D. & Tversky, A. On the Psychology of Prediction. *Psychological Review* **80**, 237-251 (1973).
78. Tversky, A. & Kahneman, D. Availability: A Heuristic for Judging Frequency and Probability. *Cognitive Psychology* **5**, 207-232 (1973).
79. Tversky, A. & Kahneman, D. Judgment Under Uncertainty: Heuristics and Biases. *Science* **185**, 1124-1131 (1974).
80. Wikipedia. in Wikipedia (The Wikimedia Foundation, San Francisco, CA, 2015).
81. Wikipedia. in Wikipedia (The Wikimedia Foundation, San Francisco, CA, 2015).
82. Evans, J.S.B.T. In two minds: Dual-process Accounts of Reasoning. *TRENDS in Cognitive Sciences* **7**, 454-459 (2003).
83. Stanovich, K.E. Who is Rational?: Studies of Individual Differences in Reasoning (Lawrence Erlbaum Associates Publishers, Mahwah, NJ US, 1999).
84. Song, S. Consciousness and the consolidation of motor learning. *Behavioural brain research* **196**, 180-186 (2009).
85. 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).
86. Department, U.S.S. (ed. Department, U.S.S.) (U.S. State Department Washington, D.C., 2005).
87. Xu, J., Kochanek, K.D., Murphy, S.L. & Tejada-Vera, B. (ed. Statistics, N.C.f.H.) (U.S. Center for Desease Control, Atlanta, GA, 2010).
88. Carroll, L. (eds. Browne, T., Melwani, G., Weeks, G., Smith, L.L. & Team, T.O.D.P.) (Project Gutenburg, 1897 (2009)).
89. Dawes, R., Faust, D., & Meehl, P. Clinical Versus Actuarial Judgment. *Science* **243**, 1668-1674 (1989).
90. 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).
91. Faust, D. The Limits of Scientific Reasoning (University of Minnesota Press, Minneapolis, MN 1984).
92. Goldberg, L.R. Simple Models or Simple Processes? Some Research on Clinical Judgments. *American Psychologist* **23**, 483-496 (1968).
93. 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).
94. 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).
95. 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).
96. Sarbin, T.R. A Contribution to the Study of Actuarial and Individual Methods or Prediction. *American Journal of Sociology* **48**, 593-602 (1943).
97. Sawyer, J. Measurement and Prediction, Clinical and Statistical. *Psychological Bulletin* **66**, 178-200 (1966).
98. Evans, J.S. Logic and Human Reasoning: An Assessment of the Deduction Paradigm. *Psychological Bulletin* **128**, 978-996 (2001).
99. 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).

100. 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).
101. Klauer, K.C., Musch, J. & Naumer, B. On Belief Bias in Syllogistic Reasoning. *Psychological Review* **107**, 852-884 (2000).
102. Wikipedia. in Wikipedia (Wikimedia Foundation Inc. , San Francisco, CA, 2013).
103. 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).
104. 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).
105. Banks, A. The Influence of Activation Level on Belief Bias in Relational Reasoning. *Cognitive Science* **37**, 1-34 (2012).
106. Wason, P.C. The Processing of Positive and Negative Information. *The Quarterly Journal of Experimental Psychology* **11**, 92-107 (1959).
107. Wason, P.C. On the Failure to Eliminate Hypotheses in a Conceptual Task. *The Quarterly Journal of Experimental Psychology* **12**, 129-140 (1960).
108. Wikipedia. in Wikipedia (The Wikimedia Foundation, San Francisco, CA, 2016).
109. Wikipedia. in Wikipedia (The Wikimedia Foundation, San Francisco, CA, 2016).
110. 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).
111. Cummins, D.D. Evidence of Deontic Reasoning in 3- and 4-year-old Children. *Memory & Cognition* **24**, 823 (1996).
112. Cummins Dellarosa, D. Evidence for the Innateness of Deontic Reasoning. *Mind & Language* **11**, 160-190 (1996).
113. Wason, P.C. & Johnson-Laird, P.N. A Conflict Between Selecting and Evaluating Information in an Inferential Task. *British Journal of Psychology* **61**, 509-515 (1970).
114. Wikipedia. in Wikipedia (The Wikimedia Foundation, San Francisco, CA, 2016).
115. Selvin, S. 67 (1975).
116. Selvin, S. On the Monty Hall Problem (Letter to Editor). *The American Statistician* **29**, 134 (1975).
117. Wikipedia. in Wikipedia (The Wikimedia Foundation, San Francisco, Ca, 2016).