Paper presented by Kieran Quill to Kingston Philosophy Café, on 29 June 2016

COGNITIVE BIASES

Terminology: "Cognitive" processes (in the sense of this paper): our **judgement** and **decision-making** mental processes under **uncertainty**.

Introduction: We are not the rational "computors" and logic machines that the Age of Reason and the Age of Enlightenment envisioned we are, or could become.

The way we envision the world around us is different to how the world actually is. We develop beliefs about the world that cohere with our worldview* and self-image. Information about the world that conflicts with our worldview and self-image creates the discomfiting feeling of *cognitive dissonance.*

Our *cognitive biases* prevent us from seeing new information objectively. They help us preserve our worldview and self-image. These tendencies towards bias appear to be "pre-loaded" in us from birth, a function of our vertebrate sensory and nervous system.

(*A worldview is a set of assumptions about physical and social reality that may have potent effects on cognition and behaviour).

"*Dissonance and consonance are relations among* ***cognitions****, that is, among opinions, beliefs, knowledge of the environment, and knowledge of one's own actions and feelings. Two opinions, or beliefs, or items of knowledge are dissonant with each other if they do not fit together; that is, if they are inconsistent, or if, considering only the particular two items, one does not follow from the other*" (Festinger 1956*)*

Heuristics: Cognitive biases often result from our using shortcuts in thinking called *heuristics*. Such shortcuts allow us to make quick decisions about otherwise complex and uncertain situations by following rules of thumb, some of which are close to instinctive. Educated guess, working backwards, common-sense, drawing a picture, sense that something is too silly, examining a concrete example of an abstract problem, etc, are examples. We can also call them informal reasoning, received wisdom and intuition. Heuristics can lead us astray, leading to systematic errors- rather like the optical illusions we are all automatically prone to.

When our heuristics fail to produce a correct judgment, a cognitive bias may result. This is the tendency to draw an incorrect conclusion in certain circumstances, basically having to do with how randomness is perceived by us. We have a very poor appreciation of randomness (chance). We do not recognise it when we meet it and we cannot produce it when we try to.

Cognitive biases are "*...prejudiced or partial viewpoints that affect someone's interpretation of a problem."* (USA National Institute of Health)

Cognitive biases distort our view of the world. We habitually misjudge the role of chance in what's going on and make decisions that are irrational-irrational in the sense of decisions that are misaligned with our own best interests. They are a central part of our psychology. If we do not recognize them in our thinking, we will be seriously prone to error in thinking about the physical world, and the physical world ramifies through the biological, psychological and social worlds.

Cognitive biases: Formally, cognitive biases are rules of thumb that specify how humans systematically deviate from what is normatively expected from their decisions. Cognitive biases are "norm violations". (*Normative*: relating to an ideal standard or model, or being based on what is considered to be the normal or correct way of doing something). Thereby, inferences about situations and people may be drawn in an illogical fashion.

Psychologists have identified many cognitive biases. These are at work in our everyday thinking. We are predictably irrational. Animals also show cognitive biases. We see the world through the lens of biologically-evolved adaptions, as do animals in their own fashion.

{Technical: A bias usually takes the following form: when confronted with evidence of type X, a judge will consistently chose alternative B instead of the expected alternative A. Because **we are very consistent with our biases**, rules that describe such biases have large predictive power. A popular text book lists 53 such biases (Baron, 2008; Table 2.1).
See
http://annenberg.usc.edu/Research/Student/WorkProduct/~/media/PDFs/HilbertPsychBulletin.ashx

In the 1970s Daniel Kahneman and Amos Tversky and conducted an extensive research programme – the *heuristics and biases program* – addressing the question of how people make decisions given their limited resources. The program was motivated by Herbert Simon's **principle of bounded rationality**. For example, humans have only limited time, information, and cognitive capacity to decide which partner to choose, food to eat, or property to buy, and so may have to rely on simple decision strategies or heuristics to make their decisions. They usually decide on the basis of "good enough", in such complicated situations, though they recognize that the decisions may not be optimal. The heuristics and biases program followed the bounded rationality principle by attempting to identify the specific limitations or biases associated with peoples' judgment and decision-making.

The heuristics and biases program was inspired by previous research on perceptual biases, and proposed that the human cognitive system, like the perceptual system, is structured to make inferences about the world based on imperfect signals that could lead to errors in some situations.

The objective of the programme was to provide explanations of norm violations due to reliance on a minimum set of cognitive principles, the most common judgment and decision processes proposed being **representativeness** (a judgment is based on how much the hypothesis resembles available data), **availability** (a judgment is based on how easily an example can be brought to mind), and **anchoring-and-adjustment** (a judgment is based on a specific value or anchor and then adjusted to account for other factors). Kahneman was awarded the **2002 Nobel Memorial Prize in Economic Sciences** (shared with Vernon L. Smith; Amos Tversky had died by then) *"for having integrated*

insights from psychological research into economic science, especially concerning human judgment and decision-making under uncertainty"}

Heuristics and consequent cognitive biases are an explanation of why we hold confident beliefs, strange or otherwise.

Physical and biological reality exists independently of individual human minds, but our understanding of it depends on the beliefs we hold at any given time.

We form our beliefs in the context of our biological structure and the social environments created by family, friends, associates and culture and society at large, that is, for numerous subjective psychological reasons. *After* forming our beliefs, we then defend, justify and rationalize them with intellectual reasons, persuasive arguments and rational explanations.

"Beliefs come first; explanations for beliefs follow… Once we form beliefs and make commitments to them, we maintain and reinforce them through a number of powerful cognitive biases that distort our percepts to fit belief concepts…We want to believe!" Michael Schermer, American historian of science, "*The Believing Brain"*, Constable and Robinson Ltd, 2012.

Heuristics and biases are not necessarily "bad": Anyone who is reflective is somewhat aware of their heuristics and biases (though they may not know this terminology!).

Our beliefs, even though subject to biases, become "us", we identify with them, integrate them into our self-esteem and sense of being in control of our lives. One could say that many of our assumptions about ourselves and the world are based on shared illusions. These illusions become self-fulfilling. "Perception is reality". Perceptions and standard explanations indeed *become* stable, workable reality, but we may not realise their basic frailty. Irish psychologist, Maureen Gaffney (in "*The Way We Live Now"*, Gill & MacMillan, 1996) points out that the "legions of the mildly depressed" are those with modest, warts-and-all, assessments of their own abilities and accomplishments, and of their prospects (in a sense, those who are "realistic" about themselves and their situation).

Those who are more robustly a little vain are less prone to depression, and, of course, "success breeds success" in life. See also the "*illusionary superiority*" bias, below. Our biases, feeding our need for a decent dose of narcissism, keep us sane in the face of the novelty, ambiguity and randomness of the world. We understand and relate to other people through them, and give more breaks to the confident and assertive. We "dress for success".

We are not normally very aware of these psychological phenomena in ourselves (and are more aware of others' biases than our own). However, we may well notice them in a personal crisis: for example, if we are trying to quit smoking or

another addiction. Then cognitive biases - all 50+ of them- will come at you like hornets!

However, the physical world is very counter-intuitive (and the complexity of the socio-economic-political world is intensifying): We are not naturally adapted to grasping how the physical world works. More precisely, we are not well-adapted to understanding **the world analysed and conceptualised as a machine** where a machine is "a combination of resistant bodies so arranged that by their means the mechanical *forces of nature* can be compelled to do work accompanied by certain determinate *motion*" (Franz Reuleaux, mechanical engineer, 1829 – 1905, the "father of kinematics"). This classical concept of a machine can be readily extended to modern devices such as digital computers and robots (programmable machines). The definition implies machines are deterministic, and so, in principle, predictable and retrodictable. Quantum phenomena (quantum machines) show inherent indeterminism and unpredictability (at least in the standard interpretations of the mathematical formalism), but only for individual "observations"; the governing equation, the Schrödinger equation, evolves deterministically into the future. (The awkwardness of this situation is a source of a great deal of philosophizing). That the world, or fragments of it, can be conceptualised as machinery is a form of (meaningful) *patternicity*. The patternicity is formalised by the first and second laws of thermodynamics, which were articulated in the 19th century from the analysis of machines (heat engines, e.g., steam engines).

Even so-called "elementary" mechanics foxes most of us, and we are very prone to systematic misconceptions when learning physics. See, as an example, the footnote about common misconceptions about electricity.

The story of natural science is practically the story of how to control cognitive biases. As well as the notorious difficulties of the concepts of force and motion which are very fundamental in all physics explanations, in practice all natural phenomena exhibit random fluctuations (many are chaotic) and probability and randomness are another notoriously slippery area.

Much of the terrestrial physical world *looks* quite orderly, and indeed is quite orderly, up to a point, otherwise we couldn't exist. But that is because we live on a cool, compacted planet, where geometric patterns and structures of aggregated quantum mechanical systems arise.

This geometric orderliness arises ultimately from the isotropy of space and time (or spacetime in relativity): this leads to the conservation of momentum and energy (Noether's theorems) and hence to periodic propagations of probability amplitudes (wavefunctions) hence allowing constructive and destructive interferences. These interferences give rise to geometric regularities, primarily the angular directionality of chemical bonding. Thus the characteristic "shapes" of molecules- the basis of most chemical explanations. But quantum phenomena themselves are inherently probabilistic effects (Heisenberg-type uncertainty principles); the apparent order of aggregations belying the underlying randomness.

And when we move out from the carefully prepared environment of the laboratory experiment (approximating well to a machine) to the actual world, we encounter a ferociously baffling mix of randomness, complexity and unpredictability- the weather, the climate, earthquakes, the EU, financial systems, heart disease, cancers, epidemics, politics. Technically such systems are nonlinear. Nonlinearity is the rule not the exception; the world is intrinsically nonlinear. Cause-and-effect relations are very difficult to elucidate, and unpredictability (due to complexity or chaos or both) is the order of the day. A sequence of "good" or "bad" events need not have any discernible cause at all.

In these realms (teaching mechanics, making decisions about our finances, managing diseases, managing ecosystems) we really do need to understand our natural heuristics and cognitive biases because they will certainly lead to gross errors in reasoning and decision making when we detect pattern and structure, cause-and-effect, where there is none.

Complexity, chaos and randomness are an intricate set of concepts, but can be organised through rigorous mathematical ideas. In particular, complex and chaotic systems (two different concepts) display nonlinearity (the superposition principle cannot be used). Complex systems are not necessarily chaotic (although they can be), and chaotic systems are not necessarily complex. Both types of dynamics (evolution through time) are actually not intrinsically random, they are in fact deterministic, but they can certainly appear to be so for all practical purposes (apparent randomness), and are highly unpredictable. An **indeterministic** system is one without a unique future trajectory- its future evolution is inherently random.

Anthony Hilton (*London Evening Standard*): *We are not designed for managing money, and that holds for professional wealth managers as much as for the rest of us.*

"Patternicity" and "agenticity" (the neologisms are Michael Schermer's):

Patternicity ensures that we will seek *and find* patterns in both meaningful (orderly) and meaningless (random) events and data. Our vertebrate sensory system is adapted for organising data and observations into "useful" patterns, generating a built-in bias towards seeing patterns where there is randomness. More technically, we have a built-in bias towards rejecting null hypotheses (we will look at this below). We detect patterns everywhere (and not just visual patterns). The psychological and cognitive science term of art *apophenia* is given to the human (and animal) tendency to perceive meaningful patterns within random data.

{Technical: **Randomness** comes in two qualitatively different forms. **Apparent randomness** can result from ignorance or lack of control of degrees of freedom in the system. In contrast, **intrinsic randomness** should not be ascribable to any such cause. While classical systems only possess the first kind of randomness, quantum systems are believed to show some intrinsic randomness. **In general, any observed random process includes both forms of randomness**.

Randomness is the lack of pattern or predictability in events. A **pattern** is a discernible **regularity** in the world. E.g., A causes B (causality). Or, If A then B.

The concept of randomness can be rigorously treated mathematically (measure theory): roughly stated, the information embodied in a random series of numbers cannot be ``compressed", a random sequence of bits (0s and 1s) of a given length cannot be generated by any computer program (itself a sequence of 0s and 1s) shorter than the random sequence itself. A sequence that is not random (though it may be very irregular) can be generated by a program shorter than the sequence. https://www.cs.auckland.ac.nz/~chaitin/sciamer.html

But we'll stick with our intuitive understanding).

All of us are very familiar with this: We "connect the dots" and create meaning- we see shapes in the clouds, we tell ourselves stories in our heads - even if the "dots" are actually random. "Patternicity" is powerful. Animals do the same (*associative learning*).

Perception and conception (humans' anyway) are not straight consequences of reality, but, rather, acts of imagination. The data we encounter in living our lives are never complete, and very often ambiguous.

This applies even at the autonomic (automatic) physiological level, e.g., the processing of the optical "data" that falls on our retina. The data on the retina is very "pixilated", unstable, most of it at very low resolution: but our visual systems process it automatically, "filling in gaps" to yield a sharp and clear visual scene before us.

(Telling ourselves a story about our own life is called **confabulation**- the tendency to conflate memories with imagination and other people's accounts as one's own. Confabulation is the framework from which we understand ourselves. I am a story I tell myself. But we create quite fictional narratives much more often than we realise, to explain who we are, what we have done and why- but without realizing it. Because this is so pervasive we cannot be at all sure of how much of what we consider to be the honest truth about our past is accurate. Again it is patternicity- filling in the gaps).

Agenticity is a strong tendency to imbue real or apparent patterns with intention and agency, and a belief that these invisible intentional agents control the world (e.g., spirits, gods, aliens, demons, government conspiricists, etc- one of the latest is that we are living in a computer simulation, our world is software (simulants) in the computer programs of some powerful civilisation- the *simulation hypothesis* of Bostrum and others: see footnote.

I don't know if agenticity is graced with a Greek term of art, but it's related to the concept of *teleology* in philosophy.

An entertaining example of *apophenia* and agenticity is the **gambler's fallacy** (the Monte Carlo fallacy or the fallacy of the maturity of chances): for example, tossing a fair coin. If you toss a fair coin seven times, say, and it comes up heads each time, then many people think that the next coin flip almost *must* come up tails. Of course the probability that the next toss comes up as heads again is ½. A fair coin has no memory whatsoever of what's already happened.

There is a strong tendency to think that "nature" "knows what's happening" and *must be trying* intentionally to correct the "imbalance of the probabilities".

A famous example happened in Monte Carlo in 1913, when the roulette ball fell in black 26 times in a row. As the session progressed, black occurring successively again and again, gamblers present went into a frenzy and lost millions of francs betting against black. They reasoned incorrectly that the streak of black was causing an imbalance in the randomness of the wheel and that it had to be followed by a streak of red.

It has been noted (by Bristol psychologist Bruce Hood) that *"many highly intelligent and educated individuals experience a powerful sense that there are patterns, forces, energies, and entities operating in the world…(though) such experiences are not substantiated by a body of reliable evidence, which is why they are supernatural and unscientific…"*

(We also have a strong tendency to belief that objects, people and animals contain an **essence**, something at the core of their being, and that this essence may be transmitted from objects or people to ourselves. This is likely rooted in natural fears of contagion and disease. For example, most people will not wear any clothes that had been worn by a murderer, or buy a house associated with some heinous event –but will cherish objects that belonged to a positive person such as a famous star or saint.)

So, *"You feel you are a rational, logical person, who sees the world as it really is. The truth is, you are as deluded as the rest of us, but that's OK, it keeps you sane"*. (Book: "*You are not so smart*", David McRaney, Oneworld Publications, 2012)

---------- ---------- ----------

Evolutionary explanation:

The physical world can be baffling, and the physical world also operates in our inner subjective mental world. The ever-changing spectacle before us includes a great amount of randomness.

Our ancestors lived in a dangerous world, and we still do, even if the dangers are usually less elemental and more controllable. Our Palaeolithic ancestors developed their cognitive and psychological powers to cope with one overriding concern- to survive in the wild. We retain their biology intact. But being able to maintain orderly relations between the organism and its physical environment, which we can certainly do, is not the same thing as the organism understanding how the physical world actually works.

The outcome of the evolutionary process is that we tend to make false positive errors (*Type I errors* in statistics jargon), that something is real when it isn't (we find non-existent patterns). This tendency is shared with all vertebrate animals.

A rustle in the vegetation- it could be due to the wind, or it could be due to a predator.

If you are a very relaxed creature, then you will assume it's just the wind (harmless).

If, in fact, it is a dangerous snake, then the relaxed creature has made a Type II error in cognition (false negative): believing something is not real when it is real: the creature has missed a real pattern.

Such relaxed creatures become dinner for the snake, their genes won't be around for long. Thus, we humans are not relaxed creatures, any more than other vertebrates are relaxed.

In general, the cost (in terms of survival or consumption of vital resources) of Type I errors is less than the cost of Type II errors in cognition (you are less likely to be someone's dinner). This is especially so in the wild, when survival to the next meal or mating opportunity is imperative. But it holds true, too, in our modern social lives.

Assessing the relative cost is itself resource-consuming (calories consumed, time it takes) so the default position is to assume all patterns are real and important.

Evolutionary survival imperative: All rustles in the grass are a sign of danger and not due to the wind.

That way, you survive to the next meal and mating opportunity, and can raise enough offspring to pass on one's genes to the next generation.

We do not naturally have an error detection mechanism to moderate the pattern-recognition "engine" of our brain system. Patternicity will occur whenever the cost of making a type I error is less than the cost of a type II error:

We can say: *Patternicity if* $COST_{Type\ I} < COST_{Type\ II}$

Another major factor (for people but presumably not animals) is the need to exercise personal control over relevant events, or at least to *feel* that we control events. However, to the extent that events are random, we are *not* in control; if we are in control then events are *not* random. Therefore, there is a fundamental conflict between our need to feel in control and our ability to identify randomness. This is one of the fundamental reasons for our inborn tendency to misinterpret random events. Experiments show that even if we intellectually know about the role of chance we still behave as if chance events are subject to

our control. This phenomenon is widely seen in business organisations- where in reality there are many uncontrollable elements (e.g., unpredictable market forces), yet often, business plans and strategies ignore such factors, with inevitable "crises" such as CEOs being reactively fired, layoffs, etc (example given by Leonard Mlodinow, *The Drunkard's Walk*, Penguin Books, 2008*)*.

See this: http://www.scientificamerican.com/article/the-believing-brain/

Also a critique: https://www.psychologytoday.com/blog/theory-knowledge/201207/why-i-am-not-my-brain

---------- ---------- ----------

A brief, incomplete, survey of cognitive biases: A large number of cognitive biases has been identified or proposed. (Wikipedia lists about 170 of them, including the Ikea Effect (!): *The tendency for people to place a disproportionately high value on objects that they partially assembled themselves, such as furniture from IKEA, regardless of the quality of the end result)*.

The Confirmation Bias, also called confirmatory bias or myside bias: the tendency to search for, interpret, favour, and recall information in a way that confirms one's pre-existing beliefs or hypotheses, while giving disproportionately less consideration to alternative possibilities. It is "the mother of all cognitive biases" in that it underlies most of the other biases.

People become emotionally attached to ideas, as mentioned already. We experience cognitive dissonance when faced with evidence that suggests that our ideas are wrong, which is an unpleasant experience, and we are very reluctant to let go of patterns we are "sure" are operating.

The confirmation bias is a systematic error in inductive reasoning, and a major impairment to our ability to escape from misinterpreting randomness. We are very efficient at pattern recognition, but (confirmation bias) we are focussed on finding and confirming patterns rather than minimising false conclusions.

The effect is stronger for emotionally charged issues and for deeply-entrenched beliefs. People also tend to interpret ambiguous evidence as supporting their existing position –and evidence is often ambiguous.

Confirmation biases contribute to overconfidence in personal beliefs and can maintain or strengthen beliefs in the face of contrary evidence.

People like to be told what they already know, and don't like new stuff!

A prevalent example of the workings of the confirmation bias is the occurrences of geographical **cancer clusters**, seemingly caused by local environmental conditions such as pollution from an industrial activity. This social phenomenon is also an example of the **Texas sharpshooter fallacy** (Wikipedia: The Texas sharpshooter fallacy is an informal fallacy which is committed when differences in data are ignored, but similarities are stressed… The fallacy is characterized by a lack of a specific hypothesis **prior** to the gathering of data, or the formulation of a hypothesis only after data have already been gathered and examined. Thus, it typically does not apply if one had an *ex ante*, or prior, expectation of the particular relationship in question before examining the data.).

If one searches for "clusters" one will find them and be susceptible to erroneous conclusions.

The development of a cancer usually requires successive mutations of somatic genes. This means the population is exposed to a high concentration of a carcinogen, or else a very long exposure to a carcinogen. Cancer clusters that epidemiologists are called on to investigate would require exposures that are usually believable only in people exposed to certain work conditions or chemotherapy, i.e., much higher concentrations than local people do receive in contaminated districts. However, understandably, given the nature of cancer, people are very resistant to any explanation that the clusters are random fluctuations. Every year in the USA and elsewhere 100s of comprehensive analyses of clusters are published-none of which credibly identify an underlying environmental explanation. They have been described by an epidemiologist as "an absolute, total, and complete waste of taxpayer dollars" (example taken from Mlodinow, ibid.).

Punditry is an industry built on confirmation bias. E.g., newspapers…*The Daily Mail*…or *The Guardian*. The media provide fuel for beliefs, they pre-filter the world to match existing worldviews. If their filter is like your filter you love them. If it isn't you hate them! You read not for information but for confirmation! (McRaney, ibid.)

We are all aware of the problems of asking loaded questions. *Seek and ye shall find* (Matthew 7:7).

In court trials, observers and juries invent a narrative in their mind about what happened (deciding for guilt or innocence), instead of listening to and evaluating the evidence first and then coming to a conclusion. They scan through the evidence and pick out what most closely fits the story in their mind.

Stereotyping is a common result of the confirmation bias- e.g., the "behaviour" or characteristics of minorities.

In the context of scientific research, confirmation biases can sustain theories or research programmes in the face of inadequate or even contradictory evidence;

the field of parapsychology has been particularly affected, it seems, but the problem is certainly not only for that subject.

Parapsychology is the investigation of paranormal and psychic phenomena which include telepathy, precognition, clairvoyance, psychokinesis, near-death experiences, reincarnation, apparitional experiences, and other paranormal claims.

"Scientific method" (see later) is the most systematic way to deal with this bias. But just to be aware that *randomness produces patterns* is a good start. And, also, to devote serious effort to looking for evidence *against* a belief (the basis of Karl Popper's philosophy of scientific method).

---------- ---------- ---------- ----------

Hindsight bias*: the tendency to reconstruct the past to fit with present knowledge (time reversed confirmation bias).*

Also known as the knew-it-all-along effect or creeping determinism, it is the inclination, after an event has occurred, to see the event as having been predictable, despite there having been little or no objective basis for predicting it.

The crystal-ball view of events is possible, unfortunately, only after things happen –Mlodinow, ibid.

Hindsight bias may cause memory distortion. It has been suggested that the effect can cause serious methodological problems while trying to analyse results in experimental studies. A basic example of the hindsight bias is when, after viewing the outcome of a potentially unforeseeable event, a person believes they "knew it all along". Such examples are presented by historians, physicians recalling clinical trials, and in judicial systems trying to attribute responsibility.

Hindsight bias is not only affected by whether or not the outcome is favourable or unfavourable, but also by the severity of the negative outcome.

> In medical malpractice suits, the more severe the negative outcome the more dramatic the juror's hindsight bias. In a perfectly objective case, the verdict would be based on the physician's standard of care instead of the outcome of the treatment; however, studies show that cases that end in severe negative outcomes such as death result in higher levels of hindsight bias.

Research suggests that people still exhibit the hindsight bias even when they are aware of it or possess the intention of eradicating it. It is not possible to completely eliminate hindsight bias, but only to adopt ways to reduce it.

Hindsight bias is prominent after major accidents/ disasters: highly improbable and unpredictable events become not only probable but *practically certain* after

they happen. Had certainty really existed before the fact, then, of course, different actions would have been taken.

> E.g., 9/11: there were 100s of intelligence memos tracking the various comings and goings of Al Qaeda, and its potential targets. They included references to hijacking planes, attacking airports (LAX, Washington), bombing the World Trade Center. It was not at all clear when, where, or how or if, such attacks might actually happen.
>
> E.g., adverse outcomes in the NHS and social services are a major source of news for our media. Accidents are prone to happen in any human enterprise, but accidents occurring within the healthcare system seem more significant and severe due to their serious effects on the lives of those involved, sometimes resulting in patient deaths.
>
> In a healthcare system, various methods are systematically used to review accidents and mistakes made by physicians and other staff. However the reviewers always **already know** the outcome of the case. Methods include morbidity and mortality analyses, autopsies, case analysis, medical malpractice claims analysis, staff interviews, and observations of patients. Hindsight bias has been shown to interfere with assessing errors in these cases. Many of these errors are considered preventable *after the fact*, clearly indicating hindsight bias. There are two sides to the debate on how case reviews should be approached: *error elimination strategy* and *safety management strategy*. The error elimination strategy aims to find the cause of errors, relying heavily on hindsight (therefore more subject to the hindsight bias).The safety management strategy relies less on hindsight and identifies possible limitations during the decision making process of that case. However, it is not invulnerable to error either.
>
> None of this means that we should not develop plans and scenarios for the future, but we would do better to develop our abilities to react effectively to unforeseeable or unlikely events, and qualities such as perseverance and flexibility.
>
> History: Richard Henry Tawney (historian) observed*: "Historians give the appearance of inevitability, by dragging into prominence the forces which have triumphed and pushing into the background those which have been swallowed up"* (quoted by Mlodinow, ibid).

Rosy retrospection: — the tendency for people to rate past events more positively than they had actually rated them when the event occurred.

Self-Justification bias: the tendency to rationalise decisions after the fact to convince ourselves that what we did was the best thing we could have done. (Related to the hindsight bias).

You don't evaluate yourself based on past successes and failures: you excuse your failures and see yourself as more successful more intelligent and more skilled than you are!

You could say: *You are incredibly egocentric- just like everyone else!* (McRaney, ibid.)

---------- ---------- ---------- ----------

A self-serving bias (and, also, illusionary superiority) is any cognitive or perceptual process that is distorted by the need to maintain and enhance self-esteem, or the tendency to perceive oneself in an exaggeratedly favourable way. Individuals tend to ascribe success to their own abilities and efforts, but ascribe failure to external factors.

— The tendency to claim more responsibility for successes than failures. It may also manifest itself as a tendency for people to evaluate ambiguous information in a way beneficial to their interests (there is also *the group-serving bias*).

> Clinically depressed patients tend to show less of a self-serving bias than people in the general population.
>
> *Self-esteem is mostly self-delusion* (David McRaney, ibid.), but if you were to stop and truly become aware of your failures and faults you would risk paralysis by doubt and fear. However, your personal "hype machine" can sputter to a stop: result depression and over-anxiety. It has been noted that some cultures, the USA especially, reinforce personal hype through a culture of exceptionalism (American exceptionalism).
>
> Much research has shown that just about everyone (in the US anyway) thinks they are more competent, ethical, intelligent, less prejudiced, better drivers, younger-looking than people of the same age, etc, etc. than the average. You do not think you are an average person, but you do believe that everyone else is!
>
> **{Depressive realism** is the hypothesis developed by Lauren Alloy and Lyn Yvonne Abramson that depressed individuals make more realistic inferences than do non-depressed individuals. Although depressed individuals are thought to have a negative cognitive bias that results in persistent, negative, automatic thoughts, maladaptive behaviours, and dysfunctional worldviews, depressive realism argues not only that this negativity may reflect a more accurate appraisal of the world but also that non-depressed individuals' appraisals are positively biased. This theory remains quite controversial, as it brings into question the mechanism of change that cognitive behavioural therapy, CBT, for depression purports to target.}

Unfortunately, we are biased against people who come out at the bottom (perhaps often we do this unconsciously): more than a few regard the

unsuccessful as deserving their lot. This attitude reflects a strong need to understand situations in terms of cause and effect (a form of patternicity).

---------- ---------- ---------- ----------

Expectation bias (or experimenters bias)— the tendency for experimenters to believe and publish data that agree with their expectations for the outcome of an experiment, and to disbelieve, reject, or demote the corresponding weightings for data that appear to conflict with those expectations.

This is linked with other biases (Confirmation Bias, Selection Bias, and Publication Bias)

> Wikipedia: The **observer-expectancy effect** (also called the **experimenter-expectancy effect**, **expectancy bias**, **observer effect**, or **experimenter effect**) is a form of reactivity in which a researcher's cognitive bias causes them to **subconsciously** influence the participants of an experiment. Confirmation bias can lead to the experimenter interpreting results incorrectly because of the tendency to look for information that conforms to their hypothesis, and overlook information that argues against it. It is a significant threat to a study's internal validity, and is therefore typically controlled using **a double-blind experimental design**.

---------- ---------- ---------- ----------

The normalcy bias: This is quite familiar to professionals involved in planning for and dealing with disasters, whether sudden in onset, or slow-moving like business or societal collapses. It is also called "*negative panic*". Our expectation is that our fight-or-flight instincts will be activated when disaster strikes, with panic and suchlike reactions, or superhuman acts of determination to escape.

However, in disasters, many people become abnormally calm and pretend that everything is normal *("Crisis, what crisis?"- The Sun* headline, in response to James Callaghan's nonchalance).

Normalcy bias is also strikingly contagious. If everyone else is milling around, waiting for information, YOU will too! (probably). The tendency is, first, to interpret the situation within the context of what is normal and familiar, and thus to greatly misjudge the severity of the danger. It happens no matter what the scale of the problem or how obvious it is. There are many famous examples, such as the responses of survivors after the collision of two fully-laden Boeing 747s on the runway at Tenerife in 1977. A more recent event, 9/11, displayed plenty of this behaviour too (including the US Airforce, whose pilots kept asking "is this real world?" or just another exercise?).

It seems that an overflow of ambiguous and novel information can be overwhelming, and many people do nothing at all, and perish when they could escape. It is estimated that as much as 75% of people find it impossible to reason and act appropriately during catastrophic events. Another 7% react with confused panic. Only 7% react with unimpaired, heightened awareness.

> The **normalcy bias**, or **normality bias**, is stalling during a crisis and pretending that everything will continue to be fine and predictable, just as it was. Prior to an impending crisis, e.g., a predicted hurricane, it also causes people to underestimate both the possibility of a disaster and its possible effects.
>
> The assumption that is made in the case of the normalcy bias is that since a disaster never has occurred, it never will occur. It can result in the inability of people to cope with a disaster when it does occur. People with a normalcy bias have difficulties reacting to something they have not experienced before. People also tend to interpret warnings in the most optimistic way possible, seizing on any ambiguities to infer a less serious situation.
>
> The opposite of normalcy bias is overreaction, or **"worst-case thinking" bias**, in which small deviations from normality are taken as signalling an impending catastrophe.
>
> Normalcy bias is not the same as *fear bradycardia* (to suddenly stop moving in hopes for the best, an automatic and involuntary response at the first signs of danger- *tonic immobility*). This is an ancient mammalian reaction evolved to trick a predator's motion-tracking abilities by blending into the background.

Normalcy bias seems to be based on the fact that much of our behaviour is an attempt to lower anxiety, in ourselves primarily (self-soothing), but also in others. You attempt to make everything OK by a sheer act of belief. We pretty much all have a supreme need to feel safe: when it becomes all too obvious that this is now impossible you may well drift into a daydream where you are, after all, safe and secure.

We are "wired" to see everyday life as plain and routine. If we were not so wired, we would never be able to handle the information overload: life would be all noise and no signal.

We are not only talking about obvious disasters here: global climate change might well be a disaster that we are actually beginning to go through. Other possible examples: the continuing aftermath of the crash of 2008; the euro crisis; growing world inequality in wealth; antibiotics-resistant bacteria. We are arguably failing to act adequately and won't until it becomes impossible to ignore; it is too difficult to imagine how abnormal life could become if the predictions are true.

Sunk-Cost bias: Familiar to finance directors, but it affects a great many beliefs beyond finance.

It is the tendency to believe in something because of the cost already sunk into that belief. We hold onto unprofitable investments, failing businesses and

relationships, fabricating rational reasons to justify those beliefs and behaviours in which we have made sizeable investments of money or energy. The fallacy is that past investment should influence future decisions. Rationally, we should instead compute the chances of succeeding from this point forwards, and then decide.

Example: George Bush: 2006: *"I'm not going to allow the sacrifice of 2,527 troops who have died in Iraq to be in vain by pulling out before the job is done".*

Anchoring bias: **Anchoring** or **focalism** is a cognitive bias that describes the common human tendency to rely too heavily on the first piece of information offered (the "anchor") when making decisions. During decision making, anchoring occurs when individuals use an initial piece of information to make subsequent judgments. Once an anchor is set, other judgments are made by adjusting away from that anchor, and there is a bias toward interpreting other information around the anchor. Sellers and buyers often use this anchoring effect in setting an anchor point for the price. For example, the initial price offered by a seller of a used car sets the standard for the rest of the negotiation, so that prices lower than the initial price seem more reasonable even if they are still higher than what the car is really worth.

Various studies have shown that anchoring is very difficult to avoid. Anchoring can occur even when a completely unrelated concept is presented prior to considering the problem.

> Anchoring bias influences medical decisions far more often than medical providers would like to admit. Consider a patient who is referred to a specialist due to suspicion of a certain disease ('Disease X'). The specialist may fail to consider other (perhaps more likely) diagnoses because the patient has already been labelled with 'Disease X'. We may become anchored to a diagnosis due to other biases such as the **availability heuristic** or the **recency effect.**

Selection bias is the selection of individuals, groups, or data for analysis in such a way that proper randomization is not achieved, thereby ensuring that the sample obtained is not representative of the population intended to be analysed. It is sometimes referred to as the **selection effect**. The phrase "selection bias" most often refers to the distortion of a statistical analysis, resulting from the

method of collecting samples. If the selection bias is not taken into account, then some conclusions of the study may not be accurate.

> Scientific researchers should employ **randomization** to reduce the selection bias. By randomly selecting study groups and control groups from the same pool of subjects, researchers are less likely to conduct biased studies. There are well-founded methods for randomisation.

---------- ---------- ----------

Attribution bias: -- Systematic errors made when people evaluate or try to find reasons for their own and others' behaviour.

It is a kind of personal spin-doctoring.

People often attribute their own beliefs (especially in social and political matters) to reasoned intellectual choices, but attribute other persons' opinions on the same subject to their emotional needs. The attribution bias of perceiving intellectual reasons for belief as superior to emotional reasons is a manifestation of a broader self-serving bias.

The flip-side of this is that people tend to blame their own failures on bad luck or circumstances, and others' failures on their ineptitude. "I failed because the teacher has it in for me! They failed because they were lazy."

---------- ---------- ----------

Inattentional blindness bias (perceptual blindness): "*there are none so blind as those who will not see".* Matthew Henry (1662-1714), English Presbyterian minister and writer.

Also known as perceptual blindness, it is a psychological lack of attention that is not associated with any ophthalmic problem. The individual fails to recognize an unexpected stimulus that is in plain sight. When it becomes very hard for one to attend to all the stimuli in a given situation, a temporary *blindness effect* can take place as a result; that is, individuals fail to see objects or stimuli that are unexpected and quite often prominent.

This is a very powerful shaper of beliefs. It is the tendency to miss something obvious and general while attending to something special and specific. A type of "elephant in the room" situation.

We might imagine that our eyes as like a video camera and that our mental processes neutrally capture what the camera detects. Not so. Both our

perceptual systems and central nervous systems are deeply influenced by beliefs already held.

---------- ---------- ----------

Bias Blind Spot: The tendency to recognise cognitive biases in other people but to be blind to their influence upon our own beliefs.

> In a sample of more than 600 residents of the United States, more than 85% believed they were less biased than the average American. Only one participant believed that he or she was more biased than the average American. People vary with regard to the extent they commit the bias blind spot. It appears to be a stable individual difference that is measurable (for a scale, see Scopelliti, Morewedge, McCormick, Min, LeBrecht, & Kassam, 2015).

People tend to attribute bias in an uneven way. When people reach different conclusions from each other, they each tend to label the other person as biased, and themselves as being accurate and unbiased.

People who are high in bias blind spot are more likely to ignore the advice of other people, and are less likely to benefit from training geared to reduce their commission of other biases (Wikipedia).

---------- ---------- ---------- ----------

The availability heuristic is a mental shortcut that relies on immediate examples that come to a given person's mind when evaluating a specific topic, concept, method or decision. The availability heuristic operates on the notion that if something can be recalled, it must be important, or at least more important than alternative solutions which are not as readily recalled. Subsequently, under the availability heuristic, people tend to heavily weigh their judgments toward more recent information, making new opinions biased toward that latest news.

> — estimating what is more likely by what is more available in memory, which is biased toward vivid, unusual, or emotionally charged examples. This leads to using hasty generalizations and anecdotal evidence in arguments.

> For instance, consider the 2014 Ebola outbreak in West Africa. This received widespread attention in the news. Ebola is a frightening infection with no effective treatment. It was a serious situation for under-resourced African nations. The news reported on the thousands of Africans who died of the disease. A few people were diagnosed in the United States after traveling from Africa. People, including doctors, heard about it constantly. As a result, when patients presented to their doctor at that time with a fever, the consideration of Ebola as a cause of the fever was perceived by some as being more likely than other illnesses (such as influenza) which were orders of magnitude more likely than Ebola. The availability of the idea of Ebola via the heavy news exposure caused many patients and health care providers to weigh Ebola's likelihood higher than it actually was.

The role of news and media outlets is very prominent here- terrorist attacks, plane crashes, etc. For example, drug use is a low ranking risk factor for serious illness and death, yet it receives as much attention as poor diet and lack of exercise, which rank second (after heart disease).

---------- ---------- ---------- ----------

How to control and minimise Cognitive Biases?

"Science", in a word. This is a rhetorical oversimplification (!) as literature, poetry, drama, philosophy and theology are rich in insights into the foibles and delusions of mankind. But the "scientific method" has particularly explicit procedures for addressing bias. These are based on an explicit analysis of the role of chance.

Much popular science and even much of the science taught in school and university is *science knowledge* (facts about the world, and descriptions of various theories –the big bang, the conservation of energy, acids and bases, dinosaurs, natural selection, etc).

In contrast, how science works is often neglected by our education systems.

"How science works" is, in practice, a complex matter (and the subject of philosophy, sociology and economics of science and technology). But it can be understood in good enough terms by any interested person and it is valuable knowledge.

The US National Science Foundation has defined understanding the scientific process for gaining knowledge as grasping:

- Probability
- The experimental method
- Hypothesis testing

Not all science is rocket science, and it is valuable for us to understand some of the big scientific issues around health, nutrition, genetics, molecular biology, climate change, agriculture (if not necessarily fundamental physics and technical chemistry and biology). I recommend Ben Goldacre for his very good books: "*Bad Science*", and "*Bad Pharma."*

Probability theory is implicit in all I've covered (our propensity to misjudge the effects of natural randomness in phenomena and in our measurements).

Probability theory is, however, a challenging subject. We have already seen an example, the gambler's fallacy (and some argue that it is not always a fallacy). But it gets much more involved than that. Stories of mathematicians getting it wrong are legion (e.g., the Monty Hall TV game-show strategy problem that wrong-footed even one of the 20th C's most brilliant mathematicians, Paul Erdős, much to his chagrin, as the correct strategy was published by an amateur columnist in popular *Parade* magazine, Marilyn vos Savant).
http://www.wired.com/2014/11/monty-hall-erdos-limited-minds/

Suffice it to say we are not wired up for probabilistic thinking- we are adverse to randomness (cognitive bias again).

Statistics (statistical inference) is the application of probability theory (a branch of pure mathematics) to empirical reality. There are controversies about its applications in various situations, but statistical inference (from samples to populations) is an extremely powerful set of methods, and the key to controlling and minimising cognitive biases.

Hypothesis testing and the Null hypothesis: In all published research, statistical methods are applied to decide whether a set of measurements or observations provides good support for a hypothesis, or whether, on the contrary, the apparent support is probably due to chance.

Scientific method begins with what is called the *null hypothesis*. This is basically the same logic as the "innocent until proven guilty" presumption of innocence in criminal justice. It is also related to Ockham's razor, the Principle of Parsimony (a heuristic, not a principle of logic): the preference for simple hypotheses over complicated ones.

A scientific hypothesis is not true, or "null", until proved otherwise. A null hypothesis states that X does *not* cause Y.

If you think X *does* cause Y, a positive claim, then the burden of proof is on YOU, *not on me,* to provide convincing experimental data to reject the null hypothesis.

This constraint, that the burden of proof is on you, not on me, is a very beautiful formal aesthetic feature of science, and liberating in my view (I have enough burdens already, thank you very much…). It also captures the formal beauty of the fundamental honesty and modesty of scientific method. Indeed science is inherently conservative, cautious in the face of uncertainty, cautious in the face of doubt.

Unfortunately the contemporary authority and political-economic prestige of "science" can lead to arrogance, and influential aberrations such as "scientism." Scientism can considered as science-as-ideology (social and political ideology). Scientistic stances, very apparent in much contemporary popular and even professional science (e.g., the TED lectures) can lead to unwarranted, misleading and, often, thoughtless doctrines such as that promulgated by Francis Crick: *""You", your joys and your sorrows, your memories and your ambitions, your sense of personal identity and free will, <u>are in fact no more than</u> the behaviour of a vast assembly of nerve cells and their associated molecules"* (Quoted in M R Bennett and P M S Hacker, "*Philosophical Foundations of Neuroscience*", Blackwell Publishing, 2003). The human body and even "you" *can* be viewed, as a meaningful pattern, as a molecular machine, for various purposes such as medical diagnosis and treatment.

But to say "you" are no more than a machine is not a scientific hypothesis (testable by observation and measurement) but a metaphysical statement. One is under no obligation whatsoever to accept it. As Hacker and Bennett point out (ibid): *"Such assertions as these –namely that human beings are machines, or that the behaviour of a human being is no more than the behaviour of their nerve cells, or that decisions are taken in and (apparently) by the brain –are not science but metaphysics… [and thus] they are not open to scientific confirmation or disconfirmation".* One can equally, and more precisely given our conceptual schemes, say that living creatures are not machines. Biology after all is not a branch of physics. Furthermore, Hacker and Bennett point out *"I am not the behaviour of the nerve and other cells of which I consist, since I am not the behaviour of anything -not even of myself. I am a human being, an animal of a certain kind with very distinctive kinds of capacities."* Another example: Richard Feynman (The Feynman Lectures on Physics): "*Everything that living things do can be understood in terms of the jiggling and wiggling of atoms".* Contemporary philosopher Daniel Dennett claims that we are "moist robots", quoting cartoonist Dilbert (Scott Adams).

Not a few scientists (at least those who write popular books) seem to revel in the conclusion, apparently foisted upon us if we are "scientific", that it's all, the universe and us, meaningless, just a "product" of elementary forces or powers.

(Short diversion on scientism: Philosopher Isaiah Berlin analysed the phenomenon perfectly in 1964 in his lectures on *The Roots of Romanticism*: [The view is] that there is a nature of things such that, if you know this nature, and know yourself in relation to this nature, and…understand the relationships between everything that composes the universe, then your goals as well as the facts about yourself must become clear to you…About all these things disagreement may occur, but that there is such knowledge-that is the foundation of the entire Western tradition…The view is that of a jigsaw puzzle of which we must fit in the fragments, of a secret treasure which we must seek.

The essence of this view is that there is a body of facts to which we must submit. Science is submission, science is being guided by the nature of things, scrupulous regard for what there is, non-deviation from the facts, understanding, knowledge, adaption. Quoted by Curtis White, *The Science Delusion*, Melville House, 2013)

There are many crucial nuances to the "scientific method": there is no such thing as *proof* in the empirical sciences (proof is only possible in the domain of the *a priori* formal sciences of mathematics and logic); at best a hypothesis or overall theory is *well-supported* by the evidence, and it may be considered "confirmed" for all practical purposes. "Confirmed" means that the hypothesis or theory is a very reliable account of specified phenomena, in a given domain, under "typical conditions". In chemistry, for example, a proposed reaction mechanism may be a very reliable predictor of outcomes when the pressure is "fairly low", say up to about 20 bar. But chemists know that different outcomes are possible at high pressures. There is also the need to clearly distinguish a scientific hypothesis from a statistical hypothesis; there can be "informed" null hypotheses (informed by "prior knowledge"); and there are different "scientific methods" such as (Bayesian) induction and the hypothetico-deductive method (Popper, etc). But it is not necessary for non-specialists to understand all this for most purposes.

In addition *"there is no single, uniform subject called "science", that is pursued by all scientists in their investigations any more than there is a single subject matter called "reality" that furnishes them with a subject matter"* (Hacker and Bennett, ibid.). "*There are numerous sciences, each with different concerns. Scientists, at their best, are dedicated to the pursuit of truth within their subject".* Ibid. Thus, it is excruciatingly awkward and often pointless to attempt to "reduce" chemical concepts such as acid or base, or even chemical bond, to quantum mechanics (even though QM is "supposed" to be the fundamental theory underlying chemistry).

We don't say, vaguely, that a hypothesis is "well-supported" or "confirmed". We can *measure* our confidence in a hypothesis-with a definite number!

The most used measure is the *p-value* (statistical significance). The *p*-value is, formally, a function of the observed sample results (a test statistic) relative to a statistical model (*Fisher's probability value,* after the great English mathematician and biologist Ronald Fisher, who carefully analysed the overall concept and its use in the 1920s).

So, you can judge *for yourself* whether a scientific claim is well-supported. If so, then the hypothesis is *likely true, but* it may, in fact, be false.

It is evident that too many professional scientists do not understand what a *p*-value actually measures! That's probability for you! The august American Statistical Association felt moved recently to issue a formal statement (uniquely in its history) on the use of *p*-values, such is the degree of misuse and confusion (see foot notes, below).

To give a flavour of the *p*-value and similar statistical concepts: *The absence of evidence is not evidence of absence; failure to reject a null hypothesis is not equivalent to accepting a null hypothesis (the nub of the confusions is that it often is thought that it is equivalent).*

Where do we draw the line between accepting and rejecting competing hypotheses? There is actually no "magic line" that says that this particular hypothesis is definitely the true one. Significance testing (*p*-values and the rest) is a formal process for calculating the probability of our having observed what we observed **if** the hypothesis we are examining (often the null hypothesis) is true. If the calculated probability is "low enough" (a practical decision: how much risk do we want to accept by our concluding that an effect or difference exists when there is no actual effect or difference) we *reject* the hypothesis. We *accept* it if the probability is likewise "high enough". Even with data significant at, say, the 3% level, our conclusion may, due to pure chance, be incorrect. Nor is mere statistical significance enough- that one's data are significant statistically, does not mean the effect is in any way important at all. Expert knowledge and judgement of the subject area is essential, as is proper experimental design. Unfortunately, there are many examples of technical expertise in statistics being misused to misleading or even fraudulent purposes (see Ben Goldacre: *"Bad Science"*).

Inferential sciences such as the historical sciences of palaeontology, cosmology etc, or aspects of the observational sciences such as ecology, epidemiology, zoology, geology, etc do not fit the model of manipulative experimental laboratory sciences (physics, chemistry, biochemistry). Still, rigorous hypothesis testing is applied, for example, the convergence of evidence from many lines of research, and the possibility of studying "natural experiments of history" or "natural experiments of society" (in anthropology, sociology and economics).

By these various methods we can avoid or reduce the cognitive biases that will otherwise certainly colour our interpretations of data and observations. But it is always *caveat emptor* for anyone wanting to use scientific findings.

The principle of positive evidence follows from the burden-of-proof requirement. A claimant must have positive evidence in favour of a theory, not just use negative evidence against a rival theory ("if science cannot explain alien UFO-type phenomena, then that shows that UFOs are real". It doesn't.)

Plenty of natural puzzles go unsolved, until another day, if ever.

Philosophy- Metaphysics: Metaphysical hypotheses are not scientific hypotheses. They are not empirically testable: no observation could cause a believer to reject a metaphysical explanation. (Empirically testable: verifiable or falsifiable).

Any metaphysical explanation must be taken on *faith* (though it may take the form of a "metaphorical faith": the metaphors creating a "spirit world", or a "supernatural world" so to speak, through symbols, a world that discovers its

own autonomy in respect to the physical and biological domain: see "*The Science Delusion*" by Curtis White). Kant said of his own work: "*I have ruled out knowledge to make room for faith"* (he was probably a Christian believer, certainly "officially", though he may have been agnostic in reality).

Nevertheless, it is defensible to use evidence from the natural world in making a metaphysical decision, so long as we are clear about the logical status of what we are doing. Where no direct knowledge is ever possible, even in principle, as in the case of direct knowledge of *the noumenal* (what lies permanently beyond the range of any possible human experience and thought) it is not irrational to decide to have faith in a metaphysical explanation. Natural theology is potentially an example of this.

Metaphysics depends on the use of *a priori* reasoning. This can lead to a kind of **philosophical cognitive bias:** rationalism. We are supposed to know, by pure reason, what kind of thing *must* cause some effect, and what could *not* cause it. Some kinds of causation seem especially intelligible, like pushing directly on a material object (mechanical causation). Others, like action at a distance (gravity, electromagnetic radiation like light), or the interactions of mind and body seem very mysterious. It was also held that there must be absolute space and absolute time. Late 19th C (Maxwell's equations with their non-Galilean "*c*" constant) and especially 20th C relativity and quantum mechanics have destroyed our naivety! The physical world just is not what we expect it "should" be.

But in the hands of its greatest practitioners, like Kant, *a priori* reasoning from very general considerations can lead to some tremendous insights. This is the case (I think) with the *transcendental idealism* of Kant and Schopenhauer.

The Ultimate Cognitive Bias- we are and *must be*, forever, in a "Kantian fog".

The below quote is from Bryan Magee, *Confessions of a Philosopher, A Journey through Western Philosophy*, Phoenix Paperback, 1997: *"Reality does not correspond to our conception of it.* (KQ: In opposition to this is philosophical realism: that there is an independently existing material world to which our perceptions and conceptions correspond). *There is no intelligible sense in which our system of the world can be said to be "like" the world as it is in itself because the former can exist only in terms of mind-dependent and sense-dependent categories and there are no other kinds of category in terms of which any comparison between those and the world could be made by us. And even leaving aside the impossibility of comparison, categories as such are applicable only to experience, they are the forms of experience: they categorise perceptions, conceptions, and whatever else may be available to us in consciousness or awareness: what they are is such that there is no way in which*

they could be features of things as those are in themselves, independent of consciousness or awareness.

Reality in itself cannot be what we all grow up taking it to be, but whatever it is independently of what we take it to be is something radically unconceptualisable by us.

It is an analytic truth that our whole conception of the world is in mind-dependent categories which could not possibly apply to anything independently of awareness".

KQ: Note that Kant and Schopenhauer are **not** denying the reality of our empirical experiences, nor denying that the sciences reach truths about the empirical world. They are **not** saying we are making up reality "in our heads".

OVERALL CONCLUSIONS:

Cognitive biases, arising from our biological structure, are inevitable, and not necessarily maladaptive. But they can and do lead to serious errors in physical and biological reasoning, and scientific and statistical methods are essential to minimise such errors. Physical reality (its "machine-like" and random and chaotic nature) ramifies through all aspects of our world, and cognitive biases are a source of error in decision-making in complex social and economic situations.

More general forms of thinking, such as philosophy, are also susceptible to cognitive biases.

From a personal perspective- keep going! But with your eyes open to the pervasive role of randomness (which after all makes the world more interesting, open, and adventurous).

Even a tossed coin weighed towards failure sometimes lands on success! (Mlodinow, ibid).

"If you want to succeed, double your failure rate!"- Thomas J Watson Sr, Chairman of IBM.

Foot note: "Scientific method", *p*-values, null hypotheses and hypotheses testing:

As most people know, science involves "doing experiments" and "making observations" so as to establish truths about nature and about the social and psychological domains. Scientists

strive to arrive at just one, unique, explanation or "model" (the truth) about the given things or processes to be found in nature and other domains

These matters ("scientific method") can be quite rigorously set out. This link gives an excellent account of the various aspects of scientific methodology: Chapter 4 (Framing and Testing Hypotheses), Gotelli and Ellison, *A Primer of Ecological Statistics,* 2004:

https://www.uvm.edu/~ngotelli/Bio%20264/Gotelli&EllisonChapter4disputed.pdf)

***P*-values**: An assessment of "statistical significance" by an index called the *p*-value underpins conclusions of a great many publications of original research in the sciences.

The *p*-value was first formally introduced by Karl Pearson, in 1914, but Laplace had much earlier, in the 1770s, used basically the same idea.

The use of the *p*-value in statistics was popularized by English mathematician and biologist Ronald Fisher. In his influential book *Statistical Methods for Research Workers* (1925), Fisher proposes the level $p = 0.05$, or a 1 in 20 chance of being exceeded by chance, as a limit for statistical significance, and applies this to a normal distribution (as a two-tailed test), thus yielding the rule of two standard deviations (on a normal distribution) for statistical significance (called the "68–95–99.7 rule"). Fisher also published "*The Design of Experiments"* in 1935, also highly influential, introducing, among other concepts, that of the null hypothesis (the famous "lady tasting tea" experiment). Fisher: The null hypothesis is "*never proved or established, but is possibly disproved, in the course of experimentation"*

But the concept of the *p*-value is widely misused and misinterpreted by writers and their readers. As has been noted: "*P-values often determine what studies get published and what projects get funding".*

There is high pressure on researchers to find low *p*-values. Obtaining a low *p*-value for a hypothesis test is considered critical professionally, because it can lead to funding, papers in leading journals, and kudos. Statistical significance is everything!

The American Statistical Association, in a statement on the use of *p*-values, upholds the usefulness of properly interpreted *p*-values, but cautions that *p*-values are "*commonly misused and misinterpreted"* and continue, *"The widespread use of "statistical significance" (generally interpreted as "$p \leq 0.05$") as a license for making a claim of a scientific finding (or implied truth) leads to considerable distortion of the scientific process".*

A *p*-value is the probability of obtaining an effect at least as extreme as the one in the data, **assuming** the truth of the null hypothesis. Being a probability, its value is in the range 0 to 1.

Most usually, the null hypothesis postulates the absence of an effect, such as no difference between two groups, or the absence of a relationship between a factor and an outcome. It is the devil's advocate stance, and is exceedingly important because it is entirely possible that there will be an apparent effect in the experiment due to random error –in fact it is highly unlikely that the experimental data will exactly equal the null distribution of outcomes.

P-values have been criticised because they are, as noted, widely misunderstood, and they don't tell scientists what they want to know (that "the hypothesis is confirmed"). Here are a few examples of misinterpretation of this index:

- $P = 0.05$ does not mean there is only a 5% chance that the null hypothesis is true.
- $P = 0.05$ does not mean there is a 5% chance of a Type I error (i.e. false positive).
- $P = 0.05$ does not mean there is a 95% chance that the results would replicate if the study were repeated.
- $P > 0.05$ does not mean there is no difference between groups.
- $P < 0.05$ does not mean the experimental hypothesis is confirmed.

A *p*-value means **only one thing!**: *The probability of getting the results you did (or more extreme results) given that the null hypothesis is true.*

Note that a low *p*-value indicates that your data are *unlikely* assuming a true null, but it cannot decide which of two competing situations is more likely:

- The null is true but the sample was unusual.
- The null is false.

The *p*-value is not the **error rate** (the probability of incorrectly rejecting a true null hypothesis.

Determining which case is more likely requires subject area knowledge and replicate studies. Replication is crucial.

This is a good example of how slippery probability theory is, but also, correctly understood, how precise probability thinking is, and how important it is -and, more generally, the difficulties of correct statistical inference from sample to population.

Footnote on difficulty of physics due to cognitive bias: As one example, all teachers of physics are aware of the "constant current" misconception in understanding electric currents. This is the misunderstanding that there is a constant current provided by the battery in the circuit, that "wants" to flow, and that circuits "try" to resist or split this current. This misconception causes endless confusion. At base, there is a **universal anthropomorphic bias** at work, that batteries want to produce a current, want to force it through any given circuit, etc. Actually, batteries supply a constant voltage (energy per electric charge), and this is due to the very atoms that constitute it. Such anthropomorphic (teleological) mechanisms are always out of place in mechanistic explanations and calculations. https://www.psychologytoday.com/blog/psyched/201309/all-paths-lead-magical-thinking -anthropomorphism

Footnote: Are we living in a computer simulation-example of our tendency to agenticity? http://www.scientificamerican.com/article/are-we-living-in-a-computer-simulation/

From Wikipedia entry: In 2003, the philosopher Nick Bostrom proposed a trilemma that he called "the simulation argument".

Actually Bostrom's "simulation argument" does not directly argue that we live in a simulation; instead, Bostrom's trilemma argues that one of three unlikely-seeming propositions must be true. The trilemma points out that a technologically mature "post-human" civilization would have enormous

computing power; if even a tiny percentage of them were to run "ancestor simulations" (that is, "high-fidelity" simulations of ancestral life that would be indistinguishable from reality to the simulated ancestor), the total number of simulated ancestors, or "Sims", in the universe (or multiverse, if it exists) would greatly exceed the total number of actual ancestors. Therefore, at least one of the following three propositions is almost certainly true:

1. "The fraction of human-level civilizations that reach a post human stage (that is, one capable of running high-fidelity ancestor simulations) is very close to zero", **or**
2. "The fraction of post human civilizations that are interested in running ancestor-simulations is very close to zero", **or**
3. "The fraction of all people with our kind of experiences that are living in a simulation is very close to one"

Bostrom goes on to use anthropic reasoning to claim that, *if* the third proposition is the one of those three that is true, and almost all people with our kind of experiences live in simulations, *then* we are almost certainly living in a simulation.

Bostrom claims his argument goes beyond the classical, ancient "sceptical hypothesis", claiming that "...we have interesting empirical reasons to believe that a certain disjunctive claim about the world is true", the third of the three disjunctive propositions being that we are almost certainly living in a simulation". Thus, Bostrom, and writers in agreement with Bostrom such as well-known philosopher David Chalmers, argue there might be empirical reasons for the "simulation hypothesis", and that therefore the simulation hypothesis is not a sceptical hypothesis but rather a "metaphysical hypothesis". Bostrom states he sees no strong argument for which of the three trilemma propositions is the true one: "If (1) is true, then we will almost certainly go extinct before reaching post humanity. If (2) is true, then there must be a strong convergence among the courses of advanced civilizations so that virtually none contains any relatively wealthy individuals who desire to run ancestor-simulations and are free to do so. If (3) is true, then we almost certainly live in a simulation. In the dark forest of our current ignorance, it seems sensible to apportion one's credence roughly evenly between (1), (2), and (3)... I note that people who hear about the simulation argument often react by saying, 'Yes, I accept the argument, and it is obvious that it is possibility #*n* that obtains.' But different people pick a different *n*. Some think it obvious that (1) is true, others that (2) is true, yet others that (3) is true."

As a corollary to the trilemma, Bostrom states that "Unless we are now living in a simulation, our descendants will almost certainly never run an ancestor-simulation".

END

Kieran Quill, 26 June 2016 ©

I thank Jeremy Smithers for his very helpful comments on this paper.

Extra: "20 dots exercise": next page:

Example of 20 dots randomly distributed

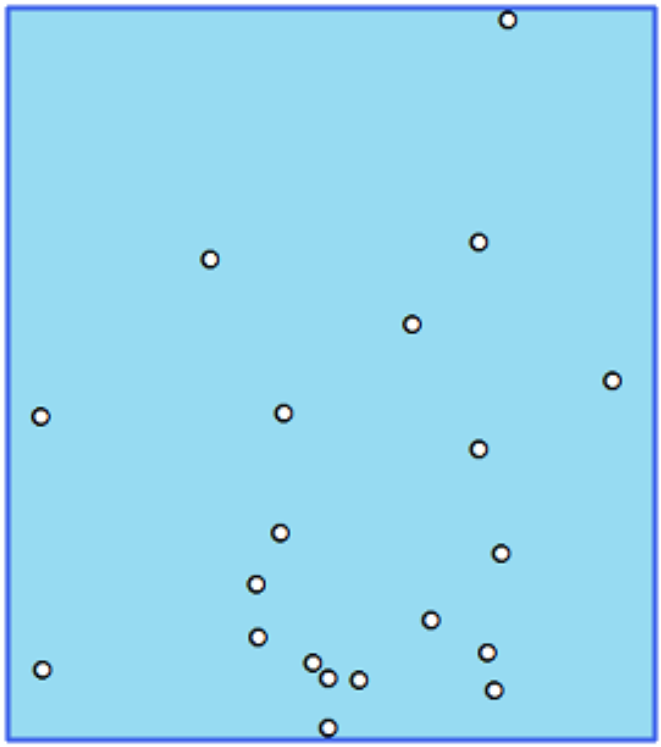

Most people, if asked to draw 20 dots randomly distributed (before being shown pictures like this) will more or less fill the page with fairly evenly spaced dots. Some will draw a fairly circular cluster of dots filling only a part of the page, but again, the dots will be rather regularly spaced.

People who have some prior knowledge of statistics may make a drawing that is somewhat like this example, deliberately introducing a degree of inhomogeneity, and maybe some outliers.

However, none of these attempts will produce a random array (that would pass statistical tests for randomness). By a fluke, someone might produce a very close approximation, but when asked to draw another picture will not be successful the second time. This is because we are inherently *biased* against randomness, we are much too *intentional* to be successful at this kind of exercise.