The Theory-Ladenness of Observation and the Observation-Ladenness of Theory

Ioannis Votsis

1. Introduction

The topic of theory-ladenness has been a fixture of the philosophy of science since the middle of the twentieth century. Put simply, and expressed as a thesis, it holds that observation is laden with theory. As such, it presents a potential threat to the objectivity of science. After all, if what we observe is affected by theories, and those theories vary, then epistemic (e.g. confirmation) judgments based on those observations may not be as neutral and univocal as we would hope them to be. In this chapter, we explore some of the most pressing epistemic issues that surround this topic. We begin with a short section on terminology. We then proceed to a section that describes five prominent types of theory-ladenness. Following that is a section that articulates and scrutinises the threat theory-ladenness poses to the objectivity of science. The subsequent section introduces and explores the inverse notion of observation-ladenness. The penultimate section speculates about some impacts AI in science may have on theory-/observation-ladenness. We conclude with a summary of the chapter's main points.

2. Terminology

The terms 'theory' and 'observation' in the expression 'theory-ladenness of observation' have been employed in a variety of ways. There are narrow construals, where the two terms (or cognate terms) are taken at face value. But there are also broad construals, where the two terms (or cognates) are understood as umbrellas. The term 'theory' in such cases covers one or more of the following: hypotheses, conjectures, models, conceptual schemes, linguistic frameworks, prior beliefs, cognition, scientific instrument operation/design choices, neuro-physiological factors, etc. The term 'observation' in such cases also covers one or more of the following: perceptions, sensations, experiences, data, evidence, empirical results, measurements, records, detections, etc.

We find both the narrow and broad construals useful so long as one makes clear how each is employed. Unless otherwise noted, the expression 'theory-ladenness of observation' will be employed to make general points about all sorts of types of ladenness. Whenever specific points about certain types of ladenness need to be made, we will revert to a more fine-grained terminology, e.g. the theory-ladenness of perception, the theory-ladenness of measurement and the cognitive penetrability of perception.[1] What all types share is a 'functional' structure. Inputs (e.g. sources, samples, signals, stimuli, worldly systems) can be thought of as getting transformed into outputs (e.g.

[1] Raftopoulos (2001) offers a great introduction to the debate over the cognitive penetrability of perception and its connections to theory-ladenness.

perceptions, sensations, observations, data, evidence) by a ladening process (e.g. theories, hypotheses, models, conceptual schemes, prior beliefs).

Before we proceed to the next section, a couple of provisos are in order. First, some of the outputs may also play the role of inputs. This is arguably the case with raw data, which feed into and help create evidence. Second, there is admittedly something weird about calling instrument or neuro-physiological mechanisms 'theories'. We here disavow ourselves of the tedious task of engaging with the literature on the metaphysics of scientific theories. Instead, we ask whether such mechanisms, or at least our understanding of them, can be effectively described in terms of a set of theoretical claims or commitments. We answer in the affirmative. For example, weather radars are instruments whose purpose is to detect how fast particles of rain, snow, sleet or hail are moving towards or away from them. They do so by using the Doppler effect to calculate the change of a wave's frequency relative to the source of the wave (in this case the radar itself). The frequency increases as objects move closer; it decreases as they move away. The use of the term 'theories' is thus shorthand for the theoretical claims or commitments made by those mechanisms, or at least the way we understand them. Third, as we shall see in Section 5, talk of theory does not necessarily mean pure theory, i.e. theory ungrounded in any observation.

3. A Tentative Typology

In what follows, we briefly discuss some prominent types of theory-ladenness, namely attention, perception, data, and memory, largely following a typology championed by Brewer and Lambert (2001).[2] It is worth emphasising that this is neither an exhaustive identification of types nor one that sufficiently probes the various sub-types.

Let us begin with *attention*. According to this type of theory-ladenness, what we attend to with our sensory organs or instruments depends on our theoretical standpoint. Several experiments in cognitive psychology can be adduced in support of this type of theory-ladenness. This includes the 'invisible gorilla' experiment, where a person dressed as a gorilla goes unnoticed by participants in a study, who are instructed to simply count the number of basketball passes between players. The effect studied by this experiment is known as 'inattentional blindness' (Simons and Chabris 1999). In the history of science, there are well-known cases where attention, or, more properly, the lack thereof, prevents discoveries from happening earlier and raises questions whether some will happen at all. The case of Charon, Pluto's moon, is one such case. Pictures of Charon, appearing only

[2] Brewer and Lambert's own classification is in terms of attention, perception, data evaluation and interpretation, data production, memory and communication. A partly overlapping classification can be found in Frigg (2022), where five kinds of theory-ladenness are identified: perceptual, expectation, attention, operational and conceptual. Schindler (2011) distinguishes between two kinds: the theoretical irrelevance of data and theoretical assessment of the reliability of data.

as a smudge, go as far back as the mid-1960s, but the periodic nature of the smudge's movement over several pictures, and therefore the realisation that it may be an orbiting object, was only noticed in the late 1970s (Christy and Harrington 1980).

Next up let us consider *perception*. This is perhaps the most well-known and studied type of theory-ladenness. On this type, what we perceive also depends on our theories. Once again, several studies in cognitive psychology can be offered in support. They include the endlessly cited experiments of the old/young woman (Leeper 1935), the Muller-Lyer illusion (Segall et al. 1966), and the rat-man figure (Bugelski and Alampay 1961). A more recent and less prominent example involves different groups of participants being primed with different stories, a story about automobiles vs. a story about space exploration. Participants are then asked to describe some photographs, including a photograph of an ambiguous 'back of a road sign' vs. 'moon obstructed by a rod' object across the night sky (Balcetis and Dale 2007). Less than a tenth of those who read the space exploration story (vs. half of those who read the automobile story) reported seeing a sign. The suggested conclusion: prior beliefs (playing the role of our theory) affect what we perceive, or, at least how we interpret it.[3]

Consider next the theory-ladenness of *data*. What are data? Simply put, data are records of observations or measurements. According to this type of theory-ladenness, theoretical assumptions influence data at three different stages: acquisition, vetting and analysis. We look at each of these in turn.

In the acquisition stage, instruments are employed to generate raw data from various sources, signals or samples, hereafter just sources for simplicity. Theoretical assumptions are often made about both the instruments employed and the sources studied. An example of the latter comes from the alleged discovery of 'polywater' in the 1960s. Polywater was meant to be a polymerised form of water with a different boiling point, freezing point and viscosity. An assumption made at the time was that the samples being tested were pure. As it turned out, they contained various impurities (including chlorine and sodium), whose eventual detection led to the realisation that the different chemical signatures seen in the data were tracking the impurities, not a new form of water (McKinney 1991). Polywater was summarily debunked. An example of the latter, i.e. theoretical assumptions that concern instruments, involves the alleged detection of a heavy 17-keV neutrino in the mid-1980s. After several years of investigation, it was concluded that the detections were due to, among other reasons, instrument-produced artifacts. In one experiment, for instance, a ring was meant to catch escaping electrons that interfered with the measurements. The ring was in fact malfunctioning. That lead to

[3] As what gets affected by theory may be the interpretation that participants attach to their perceptions, instead of the perceptions themselves, the experiment was designed to guard against the former explanation. Whether it did so successfully is another matter, and not something we discuss here.

the realisation that the spikes in the data were at least partly due to this interference, and not a new type of heavy neutrino (Wietfeldt et al. 1995). The 17-keV neutrino was also summarily debunked.

In the vetting stage, raw data is cleaned or corrected from all sorts of undesirables. These include duplicates, irrelevances, glitches, known confounders, human errors, external interference, incorrect calibration, and incomplete records. Deciding what to remove or correct is, once again, a process rife with theoretical assumptions. In Astronomy, for example, foreground signals originating in our Galaxy, e.g. thermal dust emission, confound our measurements of Cosmic Microwave Background radiation (CMB) by blinding our instruments. To overcome this problem, theoretical assumptions are employed to help determine the contributions of the confounds in various regions of the map. Those regions are then blocked out via a so-called 'galactic mask' (Bennett et al. 2003). In Climate Science, one key type of measurement is gas trapped in ice, retrieved through ice core drilling. The gas gives us an indication of the chemical composition, e.g. CO_2 concentrations, of the air at several points in the past. Alas, the compactification of ice takes a variable amount of time to seal off the gas bubbles, during which time air continues to move between layers that have yet to be sealed off. This means that the gas age and the ice age surrounding it are different. Delta-age, as this difference is known, is theoretically modelled by, among other things, estimating the accumulation rate of snow, assumptions about the site the ice cores are taken from, and a diffusion model for how gas moves through different ice layers.

The final stage concerns the analysis of already vetted data. Theoretical assumptions infiltrate these decisions as well. They include choices about the statistical framework, baseline calculations, clustering methods, robustness measures, significant thresholds, correction methods and even software employed. Sometimes our hands are tied, i.e. there is no choice or hardly any choice to be made as the specific problem demands a specific solution. More often than not, however, multiple choices are available and they need not all lead to the same result. A case in point is a study by Botvinik-Nezer et al. (2020). Seventy independent teams were given the same fMRI dataset of participants completing a behavioural task relating to gambling. They were then asked to analyse the dataset to test nine *ex ante* hypotheses about which regions of the brain were active. Shockingly, no two teams opted for the same analysis pipeline and there was "sizeable variation in the results of hypothesis tests" (84). To be more precise, only one hypothesis was identified as having a high rate of significant results (84.7%).

Finally, let us consider the theory-ladenness of memory. There are two obvious cases of memory: biological and artificial. In both cases, how we store information, e.g. whether it is encoded in a neural net, a handwritten or printed note, or as bits on a solid state drive, and what happens to that information over time is affected by theory.

Take biological memory first. A long diatribe is not needed to convince the reader that human memory can be a fickle thing, especially in the long run. Still, it's important to mention that there are experiments in cognitive psychology experiments demonstrating how memory can be influenced by theory or belief. In one study, for example, participants were briefly shown an academic office and then asked to either describe or draw from memory what was in the room (Brewer and Treyens 1981). Many recalled objects that were not there, e.g. books, or failed to report objects that were there, e.g. a skull, simply because the respective objects fit or didn't fit the academic office 'schema' (read: stereotype). There are also well-known cases from the history of science. One such case is the commonly held belief that Galileo carried out falling body experiments from the Tower of Pisa. Galileo certainly described such experiments in the *Dialogues Concerning Two New Sciences*, though there is no evidence to suggest that these were anything more than mere thought experiments (Vicente and Brewer 1993). Collective memory was altered over time, perhaps because it makes for a simpler story.

Let us consider artificial memory next. All storage media, e.g. handwriting, printed text, digital (optical disks, hard disks, solid state drives, etc.), can, and do, obviously fail over time. Moreover, the format in which information is recorded may be difficult or next to impossible to understand at a later time. Some ancient writing systems, the Minoan civilisation's Linear A for example, are yet to be deciphered. Even when we do understand the format, there may be insufficient information to know what to make of the record. Anyone who conducted empirical research knows how difficult it is to follow a study unless it is accompanied with enough information about such things as the inclusion criteria, the protocol and even the units of measurement. As we saw earlier, this is especially pressing in modern science where numerous choices are involved in generating data. Indeed, one of the lessons urged by the authors of the aforesaid fMRI study is that more effort needs to go into recording the specifics of a study and ensuring that those specifics are easily accessible for posterity (Botvinik-Nezer et al. 2020).

To summarise, we have seen how attention, perception, data, and memory can be affected by various theoretical decisions. In some cases, it is clear that those decisions can shoehorn the output in ways that may shake our confidence in the science. But not all cases are like that. It is now time to examine exactly what the threat is and what can be said and done in reply.

4. Assessing the Threat

Recall that the theory-ladenness of observation presents a potential threat to the objectivity of science. How? By calling into question the ability of outputs like perceptions, observations, data to be neutral and objective adjudicators between rival beliefs, hypotheses, models, etc. After all, if what we attend to, perceive, produce/vet as

clean and relevant/analyse as data, or store in memory is affected by the specific theories endorsed, then our epistemic judgments based on those perceptions, data, etc., may also be variably affected. In this section, we articulate and assess this threat, focusing on certain prominent arguments/points and counters.

We begin with the proponents of the view that theory-ladenness undermines the objectivity of science. Kuhn (1970; 2022) and Feyerabend (1957; 1962), partly following Hanson (1958), both make much of the theory-ladenness results emerging out of the nascent science of experimental (cognitive) psychology as well as a handful of cases from the history of science.[4] In their view, the impact of theory on observation is unmistakable. Feyerabend, for example, asserts that "the interpretation of an observation-language is determined by the theories which we use to explain what we observe, and it changes as soon as those theories change" (1957: 163). Kuhn similarly asserts "[p]racticing in different worlds, the two groups of scientists see different things when they look from the same point in the same direction" (1962: 150). Both then conclude that theory-ladenness leads to observational incommensurability, literally lacking a common measure of observation, to add to other types of incommensurability (semantical and methodological) between rival theories or paradigms. Feyerabend, for example, asserts that "incommensurable theories may not possess any comparable consequences, observational or otherwise" (1962: 94). Similarly, Kuhn asserts that "[a]pplied to the relation between successive scientific theories, incommensurability meant no common lexicon" (2022: 51). Their claims and arguments go on to fuel extreme views about the nature and limits of science, including strong social constructivism (all scientific knowledge is constructed through social negotiation) and relativism (truth is relative to specific points of view).

Several counters to the view that theory-ladenness undermines the objectivity of science have emerged over the years. One type of counter targets incommensurability. Achinstein (1964), for example, indicates that rival theories or paradigms cannot be both semantically incommensurable and inconsistent, as Kuhn and Feyerabend insisted. Semantic incommensurability requires lack of shared meaning, but that is a surefire way to establish consistency between rival theories or paradigms. Others use case studies to point out that most communication between scientists of rival theories or paradigms does not break down. Votsis and Schurz (2012), for example, indicate that in the dispute between the caloric and the motion theories of heat, scientists understood each other quite well. They employed temperature and weight measurements to try to resolve their dispute in an empirical and rational manner, recognising the advantages and disadvantages of each theory in explaining phenomena. Here's a telling quotation from

[4] Hanson argued that observation is theory-laden, but not that it is irreversibly corrosive to science's objectivity.

two famous caloric theorists: “We will not decide at all between the two foregoing hypotheses [i.e. the caloric and the motion theories of heat]. Several phenomena seem favourable to the second [i.e. the motion theory], such as the heat produced by the friction of two solid bodies, for example; but there are others which are explained more simply by the other [i.e. the caloric theory]—perhaps they both hold at the same time. So, one must admit their common principle… [i.e. the conservation of heat]” (Laplace and Lavoisier 1783).

Another type of counter targets the informativeness of experimental results from cognitive psychology. For example, Brewer and Lambert (2001) question the plausibility of generalising from a few experimental results to *all* cases of perception. Stimuli in such experiments, they point out, are “either ambiguous, degraded, or requir[e] a difficult perceptual judgment” (2001, p. 179). Indeed, the ecological validity of both stimuli and experimental conditions is highly questionable (Votsis 2020). After all, neither ambiguous or degraded shapes like duck-rabbits nor contrived conditions such as short time-intervals between priming and perceptual judgments are frequently found outside of the cognitive psychology lab.[5] Taking this type of counter a step further, Frigg (2022) questions the very relevance of experimental results from cognitive psychology. His argument is that modern science increasingly relies on instrument outputs, not human perceptual judgments, and, as such, theory-laden perception becomes a “largely tangential” issue.

Those who raise doubts about the objectivity of science by insisting that theory infiltrates instrument outputs remain unshaken by this last point. Indeed, in its most vicious form, the threat from theory-ladenness takes on the form of the ‘experimenter’s regress’, which directly questions the neutrality of instruments (Collins 1985). Essentially, this is the accusation that all theory testing is circular because the instruments scientists utilise presuppose the very theories they want to test.[6] Collins cites the attempt to detect gravitational waves in the 1960s and 1970s as an example of such regress. His argument, reconstructed here for brevity, is roughly the following: As gravitational waves had never before been detected, it was not known how to validate such detections. All scientists could do, in constructing and using instruments like Weber bars, was to fall back on their own theories. Those theories included the general theory of relativity, which predicts gravitational waves. Thus, the theory of both the experiment and the instrument presupposed the theory being tested.

[5] When we are uncertain and we care about the outcome of our perceptions, we often take a closer look, and, if need be, reconsider our judgments. If the uncertainty remains, we can and do sometimes refrain from making such judgments.

[6] Collins appears to have toned down his rhetoric on the pervasiveness of the experimenter’s regress in subsequent publications.

Several counters have been given to this kind of argument. One reply is that the theory of the instrument or experiment need not be, and indeed frequently is not, the same as the theory under test (Franklin 1986). For example, to test the general theory of relativity, Eddington employed telescopes, photographic plates and micrometer devices, none of which relied on that theory but rather on optical and geometric theories (Cf. Schindler 2013). Another reply is that, even when the theories are (partly) shared, there is no guarantee of a positive outcome for the experiment. As Hacking (1983) puts it, experiments have a 'life of their own'. Versions of the Michelson Morley experiment, for example, were carried out numerous times by different teams. All of them failed to confirm the predicted aether drag, thereby disconfirming the received theory of the time, namely the luminiferous aether theory.

A key question raised by this discussion is how frequently results from experiments contradict or at least fail to support the tested hypotheses. Alas, this is not an easy question to answer. First of all, null results do not necessarily contradict the tested hypotheses. They just indicate that there is not enough evidence to support, as statistically significant, an effect predicted by those hypotheses. Second, not all null results get reported. As is well-known the authors of experiments either suppress such results because they haven't found anything they deem worth reporting or fight an uphill battle to publish them because editors and reviewers are less likely to be impressed by them (Brazil 2024). Third, positive results sometimes get undone by subsequent, larger and more statistically powerful studies (Ioannidis 2005). All three points make it difficult to measure the true rate of experiments that contradict or at least fail to support a tested hypothesis. Having said this, it should be obvious that experiments are far from guaranteed to support the tested hypotheses. As such, it cannot be the case that hypothesis testing is as self-affirming or self-indulgent as Collins suggested it might be.

A more careful examination of the claim that theory-ladenness poses a threat to the objectivity of science is now due. We have been assuming all along that the mere presence of theory-ladenness is sufficient cause for concern. But that assumption clearly needs warrant. The crucial question is whether theoretical influences distort the content of observations beyond their breaking point, where they no longer function as neutral and objective adjudicators, allowing us to discover the world. Some authors seem to suggest as much. Kuhn, for example, claims that "[w]hen paradigms change, the world itself changes with them" (1970: 111). Similarly, Burr claims that "… there can be no such thing as an objective fact… the very word 'discover' presupposes an existing, stable reality that can be revealed by observation and analysis, an idea quite opposed to social constructionism" (2015: 9–14).

Can we identify the distortion breaking point? Following Locke, we propose that point to be when structural information about the inputs is no longer mapped in the outputs. Let

us explain. Almost 350 years ago, Locke ([1689] 1996) pointed out that however different the qualitative experiences of distinct individuals, so long as the structure of the input is preserved in the output for each individual, they would still be able to talk about and point to the same/different objects in the world. Indeed, this would be so even if their qualitative experiences were inverted, now known as the 'inverted spectrum argument'. Here's a simple illustration: Whenever Tim's input is α, he forms experience e_1, and whenever it's β, he forms experience e_2. Liz's mapping is inverted: she forms experience e_1 on input β and experience e_2 on input α. That they have different, indeed inverted, experiences is immaterial. All that matters is that each individual consistently forms the same/different outputs to the same/different inputs. That allows them to identify the same/different inputs via their respective outputs. Moreover, since α, β, e_1 and e_2 are left unspecified, the point generalises to any pair of inputs and outputs. Further elaborations of this argument, including a loosening of the operative notions (from same/different to similar/dissimilar) as well as strictures on the range of conditions under which it functions can be found in Votsis (2015; 2020; 2025a).

Returning to the issue of distortion, what reasons are there to endorse the claim that theory distorts observation beyond its breaking point – henceforth, just 'distortion*' for brevity – and, moreover, does so frequently enough to bring the objectivity of science into disrepute? It's temping for supporters of this claim, let us call them 'distortionists*', to rely on the commonly held view in cognitive and neuro-scientific psychology that perception is a largely constructive process.[7] On this view, perception combines bottom-up (e.g. early sensory signals) and top-down (e.g. prior beliefs) states and processing to generate outputs. However, contrary to what distortionists* think, this view does not underwrite their position because construction ≠ distortion*, for much the same reason that influence ≠ distortion*. Indeed, the inverted spectrum case described earlier is a paradigmatic example involving construction but not distortion*. Although Tim's physiology constructs experience e_1 on α and e_2 on β, whereas Liz's constructs e_2 on α, and e_1 on β, the relevant structural information is preserved, namely two objects are consistently tracked and mapped as distinct. Nothing in the aforesaid cognitive and neuro-scientific psychology view demands that constructions are always, or even frequently, distortions.[8]

It is time to address the massive elephant in the distortionist* room. Suppose distortionists* insist on motivating their view with what they deem to be an abundance of evidence from cognitive psychology and the history of science. For this evidence to be

[7] We have received comments like this in some of our talks from members of the audience.

[8] Failed expectations in the perceptual system, i.e. where top-down and bottom-up processing clash, may be a case of distortion. Note, however, that despite initially winning out, top-down influences eventually get corrected by bottom-up signals. Just imagine what would happen if that were not frequently the case.

rationally compelling, it would have to be undistorted*. But if such evidence is undistorted*, then the distortionists* would be contracting themselves. After all, their view holds that no, or very little, evidence is undistorted*. Suppose, by contrast, that they drop the abundance of rationally compelling, and hence undistorted*, evidence requirement from their motivation. Then their view would have no leg to stand on, and it would beggar belief how they convince themselves, let alone others, that distortion* is ever-present or frequent.[9]

It gets worse for the distortionists*. The very supposition that distortionists* can successfully communicate their view is, from their own point of view, entirely lacking in plausibility. For if what they say is right, then observations of the sounds/text they generate when speaking/writing are hopelessly distorted*. A distorted* observation, recall, carries no structural information that can be shared across individuals. It then becomes a miracle how different distortionists* understand each other when they proclaim the virtues of, and support that exists for, their view, but also how they understand their opponents' (i.e. the non-distortionists*) view, and how their opponents understand them. More generally, it becomes miraculous how any communication takes place. Do you, the reader, genuinely understand the point of this paragraph? If so, then we are not talking past each other, and your observations of this text cannot be distorted*.

A final counterpoint to the distortionist* view is that theory-ladenness sometimes has a corrective (read: ameliorative) function. The example of a straw in a glass of water offers a clear illustration. Naked eye observations suggest the straw is bent but our theory of optics tells us that light just refracts to create the appearance of bentness. A more involved example is the case of adaptive optics in astronomy. Observations using Earth-based optical telescopes are obscured by atmospheric disturbance that shifts the apparent positions of celestial objects. To deal with this problem, adaptive optics relies on the projection of artificial stars at fixed points in the sky using powerful laser beams. The shift in the apparent position of the artificial stars can then be calculated to determine the profile of the atmospheric disturbance in real time. That can then be used to correct the observations of natural celestial objects by deforming a mirror that counteracts the atmospheric aberrations. In this case, both theory, namely diffraction theory and Kolmogorov turbulence theory among others, as well as observation, namely the observation of the apparent position of the artificial stars, are employed in the correction. Other examples of corrections that are either partly or wholly based on theory abound in science. They include GPS satellite measurements being corrected with the special and general theories of relativity, radiocarbon dating measurements being amended with cosmic ray flux theory, and sea surface temperature measurements being

[9] A cheeky reply that distortionists* may offer is that only or mostly the aforesaid cognitive psychology and history of science evidence is rationally compelling. Needless to say, this reply would be ad hoc, unless independent motivation could be given for the proposed restriction of rationally compelling evidence.

adjusted with a model of evaporative cooling. In sum, far from being always or even mostly distortive*, and thus detrimental to the objectivity of science, theory-ladenness can and does often exhibit corrective virtues.

Let us summarise this section's main points. Advocates of theory-ladenness utilise cases from the history of science and experimental results from cognitive psychology to motivate various forms of incommensurability. The threat incommensurability poses has been called into question by successive case studies demonstrating extensive communication and understanding between scientists endorsing rival theories or paradigms. It has also been called into question by considering the implausibility of generalising from a few experimental results in cognitive psychology to all perception. Beyond doubting the external validity of such results, the relevance of the theory-ladenness of perception has also been scrutinised based on the claim that modern science relies largely on scientific instruments, not perception. This counter, of course, leaves the supporters of the most vicious form of theory-ladenness, the experimenter's regress, unperturbed. That's because their accusations concern the outputs of instruments, which they claim to be circularly justified because the theories under test are allegedly the same ones behind the testing instruments. Replies to this accusation involve stressing that the said theories need not be, and indeed are frequently not, shared. Moreover, even when shared, it is far from guaranteed that the results will be positive. At this point in the discussion, a key question emerged: whether theory-ladenness distorts* observations, rendering them incapable of acting as neutral and objective adjudicators. Though it is temping to argue for the pervasiveness of distortive* effects on the commonly held view that perception is a largely constructive process, we saw that construction ≠ distortion*. Indeed, science itself does not give us reason to suppose that distorted* perceptions are frequent.

Those who advocate widespread distortion* face an unpalatable dilemma: On the one hand, they can endorse the abundance of rationally compelling, and therefore undistorted*, evidence from cognitive psychology and the history of science to motivate their view, but, in so doing, contradict themselves, as that would mean that distortion* is not widespread (at least in those areas). On the other hand, they can reject the abundance of rationally compelling, and therefore undistorted*, evidence, but, in so doing, leave themselves without a strong motivation for their view. The situation becomes even more dire for the distortionists*, as, by their own lights, they should be unable to successfully communicate their position/arguments to each other and to others, if there is a massive failure in the objectivity of observation. Finally, it is worth noting that far from always being a threat to objectivity, theory-ladenness sometimes plays corrective role.

All in all, it seems that the threat from theory-ladenness, while present and requiring vigilance and countermeasures, is not as debilitating to science as some make it out to

be. In the next section, we turn the tables on observation and ask to what extent theory is observation-laden and what this means for the objectivity of science. It is only by pursuing this inquiry that we will be able to form a more comprehensive assessment of the overall threat.

5. The Observation-Ladenness of Theory

In the previous section, we considered cases where theory affects observation, sometimes distorting* and sometimes correcting it.[10] A question that naturally arises at this point is whether observation affects theory, at times distorting* and at times correcting it. In what follows, we consider both distorting and correcting cases of what might be called the 'obervation-ladenness of theory'.[11] As with the discussion in the previous section, unless otherwise noted, the expression 'observation' is employed to cover all of the cognate notions, i.e. perceptions, data, etc.

Let us first examine the type of case where observation corrects theory. Examples of this type of case can be found in abundance. One example involves the construction of mature versions of the Copernican model of astronomy prompted by, *inter alia*, observations of the retrograde motion of planets, i.e. the apparent backward motion of planets against the night sky as seen from Earth. In Ptolemaic models, where the Earth is stationary, the size and speed of epicycles had to be continually revised to try to account for this motion. Mature versions of Copernican models, instead did away with epicycles, explaining retrograde motion with the relative motion of the Earth against the motion of other planets.[12] The latter models thus utilised observations to correct the mistakes of the former ones. Other examples include the use of observations to correct auxiliary theories such as the correction of Prout's mistaken hypothesis that a chemical element's atoms have the same mass – revised hypothesis: atomic mass depends on the isotopic version of elements – based on early 20th century mass spectrography observations and the correction of the auxiliary assumption that there were seven planets in the solar system – revised hypothesis: eight planets – based on mid-19th century telescopic observations of the perturbation of Uranus' orbit.

Next up, let us consider the case where observation distorts* theory. Given what we said in the previous section about theory correcting observations, it should come as no surprise that observation may lead some theories astray. That means observation may sometimes pose a threat to the objectivity of science. Observations may distort* theories

[10] In the previous section, we also saw a third type of case, where theory doesn't really affect observation, i.e. where observation is left (largely) intact. We here ignore this neural third type of case.

[11] I first proposed this notion in a BSPS talk twenty years ago.

[12] Exact predictions of this motion involve the shift to elliptical orbits (Kepler's first law), variable orbital speeds (Kepler's second law) and the harmony between orbital periods and semi-major axes (Kepler's third law).

in a variety of ways. We consider three key ways here. The observations may contain random error, they may contain systematic error, or they may give an incomplete picture of the target domain. Let us examine each of these in turn.

Take distortions* of theory due to random errors in observation first. These are frequent though, for reasons that will become clear below, hardly ever troubling in science nowadays. The story of the 750 GeV diphoton excess is one such recent case in particle physics that counsels against hasty conclusions. The relevant observations were made in 2015 at the Large Hadron Collider by two independent experiments (ALTAS and CMS). Measurements of the invariant mass of photons produced after proton-proton collisions registered an excess of events at around 750 GeV/c2 compared to what was expected as per the Standard Model. This suggested the existence of a new particle leading to a flurry of theoretical papers attempting to make sense of the anomaly. However, as the observations accumulated, the excess turned out to be a mere statistical fluctuation. It is nonetheless a clear example of a case where random error in observation distorted*, at least initially, our theorical understanding.

Consider, next, distortions* of theory due to systematic errors in observation. These are also rather common. One particular type of systematic error involves confounding variables, i.e. variables that entice scientists to spurious relationships. In experiments, for example, they are associated with, but (typically) not caused by the independent variable, and predict the dependent variable. A notable example of how observations may distort* theory due to confounding variables involves the deadly Pellagra disease. Late in the 19th and early in the 20th centuries, the reigning hypothesis was that Pellagra was caused by germs. The reason was simple. It was often observed in poor neighbourhoods, where overcrowding and inadequate sanitation created a perfect environment for infectious disease. However, the actual cause, or at least a more potent intervention factor in so-called 'primary Pellagra', turned out to be dietary, namely niacin deficiency. Since this and other deficiencies are also widespread in poor neighbourhoods, it generated a spurious relationship, namely that germs cause Pellagra.

Finally, let us consider theory distortions* that emerge from incomplete observations, i.e. observations that paint an incomplete picture of the target domain. These are even more widespread, especially in the initial stages of investigation. An example involves early theoretical models of the distribution of exoplanet types. Gas giants with short orbital periods around their parent star(s), also known as 'Hot Jupiters', were initially thought to be more plentiful. It turns out, however, that this was a consequence of the sampling/detection methods utilised. The method of radial velocity, for example, relies on the wobble of stars to infer the mass and orbital characteristics of surrounding exoplanets. The trouble with this method is that it is biased towards exoplanets that are

both massive and have short orbital periods, as these create a stronger, and hence easier to detect, wobble on their parent star(s). Hence, the overestimation of Hot Jupiters.

It's important at this point to assess the threat posed by the observational distortion* of theory. The threat is only as worrisome, as any mitigating measures allow it to be. Let us briefly consider some such mitigating measures.

First, let us examine the threat as it emerges in the case of random error. What is needed to overcome the distortive* effects of observation here is simply more observation. That's because, by its very nature, this type of error distributes its distortive* effects evenly, at least in the long run. Any distortion* thus washes out when we take averages of observations. How quickly the distortion* is likely to wash out (i.e. how many observations are required) depends on the target domain. In particle physics, for example, statistically significant thresholds are much more stringent (5σ vs. $2\text{-}3\sigma$ in the social sciences) for this precisely this reason. Tellingly, the 750 GeV diphoton excess case never made it to the 5σ significance threshold. It was dismissed because the excess disappeared as more observations were made over the space of a few months after the initial announcement of its 'discovery'. All in all, then, the distortion* of observation on theory via random error doesn't appear to be that threatening to science's objectivity.

Next up, let us consider the threat generated by systematic error in observation and, more specifically, the case of confounding variables. There is no silver bullet here, but randomised control trials (RCTs) are a powerful weapon in the scientific arsenal. They effectively attempt to turn systematic into random error by trying to equally distribute any confounders between the experimental and control groups. Known confounders are not so much of a problem, as one can check whether they are evenly distributed between the two groups after the randomisation. The distribution of unknown confounders, on the other hand, cannot be checked. The extent to which randomisation helps with their distribution depends, among other things, on the number of different types of unknown confounding variables. As that number increases, so does the likelihood that at least some types will not be equally distributed, thereby leading to distortions*.

Even so, progress in uncovering the real causal structure, or at least more potent intervention factors, can be, and often has been, made in science. The Pellagra example is a case in point. Joseph Goldberg, a physician assigned by the Surgeon General of the United States to study the disease, noticed that certain groups of people (e.g. mental patients) were more susceptible than other groups who were in close proximity to them (e.g. nurses). He justifiably wondered what other differences between these groups may make a difference. One such difference was dietary, e.g. mental patients had less access to meat (which we now know to be a source high in niacin). By changing this variable, Goldberg was able to substantially reduce morbidity rates. All things considered, the

distortion* of observation on theory via systematic error, specifically the case of confounding variables, proves more challenging a threat to the objectivity of science than random error. Even so, it by no means renders that objectivity hopeless as the Pellagra and numerous other examples suggest.

Finally, let us examine the threat of theory distortions* from incomplete observations. The story of science is largely a story of piecing together a complex puzzle, initially with few pieces at hand and no idea of what the puzzle depicts or how many pieces are left to complete it. In working the puzzle, scientists use any pieces they get their hands on (i.e. the observations already made) to guess at the picture. At first, there just aren't enough such pieces for their guesses to stand any chance of getting at even a part of the picture. Still, they keep on guessing. Their initial guesses are understandably wide off the mark. But as the pieces accumulate, so does their ability to reconstruct parts of the picture. Indeed, when different parts of the puzzle begin to coalesce, guesses at the whole picture become less desperate, or, to use our favoured epistemic notion, less distorted*. To get to this point requires the development of distinct methods of observing that not only help validate each other, but also enable the discovery of pieces that would have otherwise been inaccessible.

The multiplicity of observing methods is how sampling problems like the one described in the Hot Jupiters example are mitigated. Besides radial velocity, there are several other ways of observing exoplanets. These include transit photometry and gravitational microlensing. The former method relies on sudden dips in luminosity when a planet transits a star. Like radial velocity, it can detect Hot Jupiters, something that helps validate both methods. Unlike radial velocity, it can also detect small and lower density exoplanets. The latter method also relies on changes in luminosity but in a more indirect way. Objects with mass, like stars and planets, can act as lenses, magnifying certain background regions and demagnifying others, by gravitationally bending light from those regions. Exploiting this phenomenon to observe background objects is known as gravitational lensing. Gravitational microlensing involves measuring a planet's relatively small – hence micro – contributions to its parent star's lensing effect by examining how much the planet increases (during magnification) or decreases (during demagnification) the luminosity of a background star. Like radial velocity, it can detect Icy giants, something that helps validate both methods. Unlike radial velocity, it can also detect Icy exoplanets that are not giants, i.e. those with Earth to super-Earth masses, several astronomical units away from their parent stars. None of these observations are without problems, but together they diversify sampling and help us build a better picture of the abundance of each type of exoplanet. Needless to say, the whole picture may be so complex as to remain forever out of reach.

Our discussion in this and the previous section does not assume that in theory-ladenness it is only theory that affects observation and in observation-ladenness it is only observation that affects theory. The various distortive* and corrective examples presented in these sections often involve both theory and observation playing those roles. Take, for example, the case of adaptive optics. It's not just theory, e.g. diffraction theory, that helps correct the given observation, i.e. the apparent position of the celestial objects. It's also another observation, namely the observation of atmospheric disturbance. Similarly, take the example of the Pellagra disease. It's not just observation, i.e. that those with the disease lived in poor neighbourhoods, that distorted* the theory about the cause of the disease, i.e. the mistaken belief that it was germs. It's also another theory, namely that the overcrowding and inadequate sanitation conditions of poor neighbourhoods rule out non-germ causes of the disease.

Before we bring this section to a close, it is worth going over the main points. Theory appears to be just as observation-laden as observation is theory-laden. Moreover, observation may affect theory in a corrective or distortive* way, just like theory does with observation. Examples of both corrective and distortive* cases of observation on theory were presented and ultimately scrutinised. The point was also made that in both theory-ladenness and observation-ladenness discussions, it is not only theory that affects observation and observation that affects theory. The general upshot of those discussions is that while the threat to science's objectivity is serious, requiring vigilance and countermeasures, it is not as detrimental to the objectivity of science as some claim. Even so, we should be cautious in our approach, not jumping to the conclusion that one day we may be able to acquire the whole picture of the puzzle and nothing but.

6. Doing Science with Artificial Intelligence

In this section, we briefly speculate about two potential impacts the increasing utilisation of AI in science may have on theory-/observation-ladenness. First, we speculate how this utilisation may affect the distinction between theory and observation. Second, we speculate how the marginalisation of humans in science may affect theory and observation distortions* and corrections.

Let us begin with the potential impact the increasing utilisation of AI in science may have on the distinction between theory and observation. It should be clear from the above discussion that, given the multifarious ways in which theory affects observation and observation affects theory, including the point about observation assisting theory affect other observations and theory assisting observations affect other theories, the distinction between the two concepts is rather blurry.[13] Does the increasing utilisation of

[13] The point extends to cognate pairs of concepts we tentatively placed under the umbrella 'theory' and 'observation'. For example, the distinction between the concepts 'cognition' and 'perception' is similarly blurry.

AI in science help sharpen or further blur the distinction? Prima facie, at least, it appears that the latter is a more likely outcome. That's because the most popular kind of AI technology currently applied to science, artificial neural nets (ANNs), treats data and theories (here understood as 'theoretically vested' models) as closer in kind than how these have been traditionally rendered.[14]

Let us explain. Training a model with an ANN is a largely data-driven process. Data is first vectorialised, if not already in that form, and then fed into an ANN of an appropriate architecture. The learning algorithms chosen, particularly the activation, loss and optimisation functions, iteratively adjust the initially random weights of the ANN until the resulting model's performance reaches a desirable level, typically decided by testing the model on a proper subset of the original dataset not used in the training.[15] Here's the crucial bit. In principle, the training may be performed with raw data. Moreover, the model/theory acquired may be (virtually) as general and as accurate as can be. That's because (some) ANNs are universal function approximators (Augustine 2024). But, if raw data can be turned into theories, then the need for intermediary justification or discovery steps, e.g. raw data $\rightarrow$ processed data $\rightarrow$ data models $\rightarrow$ theories, that some philosophers (Van Fraassen 2008) have been proposing to support their claim that there is a strict division between data and theories is obviated, or, at least, called into question (Norelli, Votsis and Williamson 2025). Moreover, as Dennett (1991) noted long ago, there is a sense in which any model (including ANN-produced models) is just data in compressed form. Understood thus, models have a Janus face. On one side they are data-faced (data points encoded and predicted by the model) and on the other side they are theory-faced (generalisations of those data points). All in all, then, it seems that the increasing utilisation of AI, and particularly the utilisation of ANNs, in science may be weakening, rather than strengthening, the distinction between observation and theory.

Let us now turn our attention to the second potential impact the increasing utilisation of AI in science may have on theory-/observation-ladenness, namely how the marginalisation of human scientists may affect theory and observation distortions* and corrections. This discussion needs a little stage setting. Recall Frigg's reply to the threat posed by the theory-ladenness of perception. In his view, experimental results from cognitive psychology are irrelevant to modern science because modern science utilises instruments rather than human perceptual judgments. We believe that this sentiment is a little premature for a number of reasons. First, some sciences are less instrument-dependent than others, e.g. field work in anthropology and clinical judgment in medicine rely largely on human perceptual abilities. Second, some sciences, e.g. psychology, depend on the perceptual abilities of human participants in behavioural experiments to

[14] The notion of theoretically vested models can be found in Noreli, Votsis and Williamson (2025).
[15] For a more detailed, though introductory, explanation, see Votsis (2025b).

produce data. Third, even when instruments take centre stage, human scientists still draw on perception to read their outputs, whether those are presented on the actual machines or on computer screens. Fourth, however the data has been acquired, human perception is still utilised to inspect visualisations of the data in an attempt to make sense of it. As such, human perception is still very much in the loop. Otherwise put, reports of its death have been greatly exaggerated.

Having said all this, we agree with the general tenor of Frigg's argument that, as automation increases, the role of human perception becomes more marginal. Indeed, we think that at some point in the not-too-distant future, AI-driven science will operate at such scales and speeds that it will be imprudent, not to mention irrational, to have human perception in any part of the loop. That, of course, still leaves the matter of instrument-produced observations directed by AI scientists and their relationship to theories. That is, any threats to the objectivity of science that are instrument-based will still have to be dealt with. Can they be?

There is every reason to think that AI technologies will not only exploit the overall success of the methods outlined above in guarding against the distortive* effects, and enabling the corrective effects, of instrument-based theory-ladenness and observation-ladenness, but that it will also boost that success and those methods. We are already getting a glimpse of what's possible with AI assistance. DeepMind recently teamed up with scientists from LIGO (Laser Interferometer Gravitational-Wave Observatory) to work on the problem of reducing instrument noise (Buchli et al. 2025). LIGO detects black holes and neutron stars by measuring the resulting gravitation waves that emanate from them. At lower frequencies, 10-30 Hz, the sensitivity of the instrument drops so much it becomes useless due to persistent noise generated by the mirror stabilisation mechanism. Using reinforcement learning, an AI learning algorithm, the joint team managed to reduce noise about 30 to 100-fold, depending on the sub-band, an impressive feat by any measure. More generally, AI technologies can help improve, speed up or develop new instruments, methods of calibration, algorithms to remove or reduce artifacts and noise, ways of dealing with systematic error, techniques that increase sensitivity, and procedures that streamline reproduction, to name but a few examples.

To summarise the main, albeit speculative, points of this section, the increasing utilisation of AI in science will have at least two potential effects on science that are pertinent to our discussion. First, it will likely further erode the distinction between data and theories. Second, by removing the human element, including human perception, from the equation, it will augment science's ability to discover more pieces of the puzzle, faster and with more confidence.

7. Conclusion

Ladenness, whether in its theoretical or observational guise, is a multi-faceted phenomenon. The notions ‘theory’ and ‘observation’ in the expressions ‘the theory-ladenness of observation’ and ‘the observation-ladenness of theory’ were here treated mostly as umbrella notions, covering several cognate concepts. Four prominent types of theory-ladenness (attention, perception, data, and memory) were discussed. These were neither presented as exhaustive nor as internally uniform types. We nonetheless illustrated each type with some examples and showed how attention, perception, data, and memory can be affected by various theoretical decisions. That immediately raised the question of whether epistemically critical outputs like perceptions, observations, and data can play the role of neutral and objective adjudicators between rival beliefs, hypotheses, models, etc. Theory-ladenness, it seems, poses a grave threat to the objectivity of science. We went on to discuss this threat, considering the main arguments, e.g. incommensurability and the experimenter's regress, of those who deem science's objectivity to be impeached. In our assessment, we presented several mitigation strategies and counters to that view. These included case studies demonstrating extensive communication and understanding between scientists endorsing rival theories, the implausibility of generalising from a few results in cognitive psychology, the irrelevance of the theory-ladenness of perception to the objectivity of science, the independence of the theories behind instruments and the theories being tested, and the claim that experiments have a life of their own.

Most importantly of all, we examined whether theory-ladenness is not just present but distorts* observations to the extent that it renders them incapable of playing the epistemic role of neutral and objective adjudication. We found that, despite perception being a constructive process, there is no strong reason to regard the distortion* of perceptions as prevalent. Indeed, those who endorse the claim that distortions* are pervasive face an unsavoury dilemma: If they are pervasive, then they cannot appeal to rationally compelling evidence from cognitive psychology and the history of science to support their claim without contradicting themselves, as that evidence would also have to be distorted*. If they are not pervasive, then they lose the only way they have of motivating their view. Beyond the dilemma, we pointed out that distortionists* have seriously miscalculated what their viewpoint entails, as successfully communicating their position/arguments to each other and to others, would be impossible if there was massive failure in the objectivity of observation. We also presented cases where, far from having a distortive* effect on observation, theory has a corrective effect. This discussion immediately led the question whether observation may also affect theory in distortive* or corrective ways. We looked at a few cases of each, and concluded that theory is as observation-laden as observation is theory-laden.

The upshot of both the theory-ladenness and observation-ladenness discussion is that, while the threat to science’s objectivity is not to be ignored, it is not as crippling as has

previously been suggested. That, of course, doesn't mean that we should be hopelessly optimistic about discovering everything, everywhere, all at once. We concluded with some speculations about the potential impacts the increasing utilisation of AI in science may have on theory-/observation-ladenness. We argued that this increased utilisation is likely to further erode the distinction between theory and observation. Moreover, we argued that doing science with AI but without human perception (or indeed involvement) will likely strengthen our defence against the threats posed by theory-/observation-ladenness, ultimately augmenting science's ability to tell us more about the world we inhabit.

References:

Augustine, M. T. (2024). A survey on universal approximation theorems. *arXiv preprint* arXiv:2407.12895.

Balcetis, E., & Dale, R. (2007). Conceptual set as a top—down constraint on visual object identification. *Perception*, *36*(4), 581-595.

Bennett, C. L., Hill, R. S., Hinshaw, G., Nolta, M. R., Odegard, N., Page, L., ... & Wollack, E. (2003). First-year Wilkinson microwave anisotropy probe (WMAP) observations: Foreground emission. *The Astrophysical Journal Supplement Series*, 148(1), 97-117.

Botvinik-Nezer, R., Holzmeister, F., Camerer, C. F., Dreber, A., Huber, J., Johannesson, M., ... & Rieck, J. R. (2020). Variability in the analysis of a single neuroimaging dataset by many teams. *Nature*, 582(7810), 84-88.

Brazil, R. (2024). Illuminating 'the ugly side of science': fresh incentives for reporting negative results. *Nature*, https://doi.org/10.1038/d41586-024-01389-7.

Brewer, W. F., & Treyens, J. C. (1981). Role of schemata in memory for places. *Cognitive psychology*, 13(2), 207-230.

Brewer, W. F., & Lambert, B. L. (2001). The theory-ladenness of observation and the theory-ladenness of the rest of the scientific process. *Philosophy of Science*, *68*(S3), S176-S186.

Buchli, J., Tracey, B., Andric, T., Wipf, C., Chiu, Y. H. J., Lochbrunner, M., ... & Mera, F. (2025). Improving cosmological reach of a gravitational wave observatory using Deep Loop Shaping. *Science*, 389(6764), 1012-1015.

Christy, J. W., & Harrington, R. S. (1980). The discovery and orbit of Charon. *Icarus*, 44(1), 38-40.

Collins, H. (1994). A Strong Confirmation of the Experimenters' Regress. *Studies in History and Philosophy of Modern Physics*, 25(3): 493–503.

Dennett, D. C. (1991). Real patterns. *The Journal of Philosophy*, 88(1), 27-51.

Feyerabend, P. K. (1957). An attempt at a realistic interpretation of experience. In *Proceedings of the aristotelian society* (Vol. 58, pp. 143-170). Aristotelian Society, Wiley.

Feyerabend, P. K. (1962). Explanation, Reduction and Empiricism. in *Scientific Explanation, Space and Time, Minnesota Studies in the Philosophy of Science*, Vol. III, H. Feigl and G. Maxwell (eds.), University of Minnesota Press: Minneapolis, pp. 28–97.

Frigg, R. (2022). *Models and theories: A philosophical inquiry*. Taylor & Francis.

Hanson, N. R. (1958). *Patterns of discovery*. Cambridge: Cambridge University Press.

Ioannidis, J. P. (2005). Contradicted and initially stronger effects in highly cited clinical research. *Jama*, 294(2), 218-228.

Kuhn, T. S. (1970). *The structure of scientific revolutions* (2nd ed.). Chicago: University of Chicago Press.

Kuhn, T. S. (2022). *The Last Writings of Thomas Kuhn. Incommensurability in Science*, ed. by Bojana Mladenović, Chicago, The University of Chicago Press.

Norelli, M. F., Votsis, I., & Williamson, J. (2025). The interplay of data, models, and theories in machine learning. *Philosophy of Science*, 92(5), 1383-1393.

McKinney, W. J. (1991). Experimenting on and experimenting with: Polywater and experimental realism. *The British Journal for the Philosophy of Science*, 42(3), 295-307.

Raftopoulos, A. (2001). Is perception informationally encapsulated? The issue of the theory-ladenness of perception. *Cognitive Science*, *25*(3), 423-451.

Schindler, S. (2011). Bogen and Woodward's data-phenomena distinction, forms of theory-ladenness, and the reliability of data. *Synthese*, 182(1), 39-55.

Schindler, S. (2013). Theory-laden experimentation. *Studies in History and Philosophy of Science Part A*, 44(1), 89-101.

Simons, D. J., & Chabris, C. F. (1999). Gorillas in our midst: Sustained inattentional blindness for dynamic events. *Perception*, 28(9), 1059–1074.

Van Fraassen, B. C. (2008). Scientific representation: Paradoxes of perspective. Oxford: Oxford University Press.

Vicente, K. J., & Brewer, W. F. (1993). Reconstructive remembering of the scientific literature. Cognition, 46(2), 101-128.

Votsis, I., & Schurz, G. (2012). A frame-theoretic analysis of two rival conceptions of heat. *Studies in History and Philosophy of Science Part A*, 43(1), 105-114.

Votsis, I. (2015). Perception and observation unladened. *Philosophical Studies*, 172(3), 563-585.

Votsis, I. (2020). Theory-ladenness: testing the 'untestable'. *Synthese*, 197(4), 1447-1465.

Votsis, I. (2025a). Grounded empiricism. *European Journal for Philosophy of Science*, 15(2), 22.

Votsis, I. (2025b). Concept and Feature Change in Scientific and Deep Neural Net Representations. In *Proceedings of the Annual Meeting of the Cognitive Science Society* (Vol. 47).

Wietfeldt, F. E., Norman, E. B., Chan, Y. D., da Cruz, M. T. F., García, A., Haller, E. E., ... & Žlimen, I. (1995). Further studies on the evidence for a 17-keV neutrino in a doped 14 germanium detector. *Physical Review C*, 52(2), 1028.