A CASE FOR DEDUCTIVISM

Draft of January 27, 2025

Matheus Silva

1. Introduction

This paper argues that every inductive inference is covertly deductive and that validity might involve different kinds of necessity (logical, nomological, metaphysical). The main distinction is between *inductivism*, the view that some inferences are genuinely inductive, and *deductivism*, the notion that all inferences are deductive. Section 2 presents the problems of inductivism, highlighting its conceptual deficiencies and advancing deductivism as an alternative. In section 3, I argue that the epistemic commitments of inductivism and deductivism presuppose fallibilism and infallibilism, respectively. Section 4 attempts to clarify the metaphysical assumptions of both positions and suggest a modal interpretation of probability that illuminates some puzzles regarding chancy counterfactuals. Section 5 concludes.

2. The problems of inductivism

Consider the well-worn inductive example that "all swans we have seen are white, therefore, all swans are white." This inductive inference was supported by all the past observations of white swans until Willem de Vlamingh discovered black swans in Australia, in 1697. The Dutch explorer had to abandon his previous belief based on this new information. This is supposed to be a paradigmatic example of how inductions are defeasible, but it can be argued that this inference is actually valid. The induction was made under the assumption that all swans possess the same color.

If we add this assumption to the premises, it follows that all swans are white (Goarke, 2009, p. 134). It could be objected that the premise that all swans have the same color is false, but the falsity of a premise doesn't make the inference invalid. The inductivist might also object that we can never know these kind of statements, but even if the truth of this skeptical hypothesis were conceded, the argument would still be deductive and valid. Note that if we remove the hidden premise that makes the inference deductive, the inductive inference *qua* inductive inference cannot work. If we deny the hidden premise that all swans have the same color, the conclusion will be false. So the assumption of the premise is required for inferring the conclusion. As a matter of fact, we can only understand how inductive inferences work by reconstructing them as deductive.

The inductivist claims that reasoners openly endorse some invalid inferences as long as the premises provide some evidential support for the conclusion. Now, suppose one makes an inductive inference and it turns out that the premises are true and the conclusion is false. Surely she would abandon her inference. But this means that no one would willingly endorse an invalid inference even if the premises provided good grounds to accept the conclusion. The inductivist might reply that she would abandon the inference only because it turns out that the premises don't support the conclusion after all, but this reasoning is implausible. If the premises provided good grounds to accept the conclusion when the inference was assumed to be invalid, they would still have to be good grounds to accept the conclusion now that the conclusion turns out to be false. If the inductivist concedes that it's rational to endorse an invalid inference before the conclusion was known to be false, she will have to maintain the inference after the conclusion was revealed as false. If the invalidity of an inference shouldn't deter us from making an inference, it shouldn't compel us to abandon it. But the invalidity of an inference does compel us to abandon it. Therefore, the invalidity of an inference should deter us from making an inference.

Perhaps the inductivist could claim that in relation to empirical matters there are no assurances that an inference is valid. So you have to bet that an inference will support the conclusion while being aware that a defeater is an open possibility. In other words, there is a bet that it's not the case that the premise is true and the conclusion is false. However, this implies that when an inductive inference is made, there is a bet that the inference will be valid. In this case, induction is not an inference type that is supposed to work when it's invalid, but an inference type that aims for validity without guarantees that it will be successful. Now, when a mathematician attempts to prove a conjecture, her inference is surely deductive, but there are no guarantees that it will be successful either. So the lack of assurances that the conclusion of an inference will not turn out to be false is not enough to classify this inference as non-deductive.

There is a common, yet misguided perception that deductive inferences must have certain conclusions. This perception results from a confusion between the claim that in a valid deductive inference the premises *necessarily entail* its conclusion and the claim that the conclusion is necessarily true if the premises are true. Let's make an interpretation of a deductive inference where P and C be the premises and conclusion, the box ($\Box$) to represent necessity and the turnstile ($\vdash$) represents entailment. Thus, we have $\Box (P \vdash C)$. If the conclusion of any deductive inference was certain, we would have to interpret deductive inferences as "$P \vDash \Box C$", which means that the premises of a deductive inference imply a necessary conclusion. But $\Box (P \vdash C)$ implies only that $\Box P \vdash \Box C$, i.e., if the premises are certain, so is the conclusion. So a deductive inferences are certainty preserving and not certainty establishing (Groarke, 1999, p. 3).

One criticism that can be directed against deductivism is that it trivializes human reasoning. When an epistemic agent makes an inference she is also trying to support the conclusion based on the premises. If every inference is deductive, and the premises proposed by the reasoner are known to be true, any and every conclusion would follow seamlessly. But this concern gets some things wrong. The hardest task for any reasoner is to establish the truth of the premises. This is particularly evident when the premises are inherently complex and harder to access, as it's usually the case as far as serious research is concerned, or speculative due to their own nature (for instance, if they concern deep foundational questions that cannot be easily addressed by empirical testing and observation). On top of being very difficult to establish the truth of the premises, it's also harder to find premises that are more plausible than the conclusion and avoid circular reasoning. Finally, the very task of making an inference where the conclusion can be deduced from suitable premises is a formidable challenge, as any philosopher, theoretical physicist or mathematician can confirm. Deductions may seem trivial when we are limited to baby logic conventions and ordinary examples of *modus ponens* and *modus tollens*, but can easily increase in complexity and defy the best minds when is actually needed for theoretical purposes.

Inductivists might object that reconstructing inductive arguments as deductively valid will distort them (Simard, 2007, p. 124). For the sake of argumentation, let's concede that in inductive inferences the premises give good grounds to accept the conclusion, but doesn't entail them. In that case, the inductive reasoner accepts the following principle: "if the premises give good grounds to accept the conclusion, you are entitled to accept the conclusion". If we add this fallibilist assumption to an inductive inference, we will have this meta-inference in our hands:

If P_l-P_n give good grounds for C, you are entitled to accept C.

P_l-P_n give good grounds for C.

You are entitled to accept C.

This meta-inference about inductive inferences is valid, however it implies that the corresponding inductive inferences that involve P_l-P_n and C are valid. If there are no circumstances where the premises support the conclusion and you are not entitled to accept it, the inference is valid. If it turns out that the conclusion is false, you are also entitled to abandon the assumption that P_l-P_n give good grounds for C. The reason why inductive inferences may seem invalid is that relevant epistemic assumptions are always swept under the rug. If we accept that if Pl-Pn give good grounds for C, you are entitled to accept C., the relevant inference that involves P_1-P_n and C will be valid, but without it, no inductive inference can be made.

Now consider the following statements:

(1) It's likely that it's not the case that *p* is true and *q* is false.

(2) It's unlikely that it's not the case that *p* is true and *q* is false.

(3) It's certain that it's not the case that *q* is true and *q* is false.

These propositions express different perspectives according to an epistemic reading. (1) satisfies an inductive standard, (2) fails to satisfy an inductive standard, and (3) satisfies a deductive one. (1)-(3) are propositional attitudes that have a proposition as their object, namely, "it's not the case that *A* is true and *B* is false," which is made true by truthmakers that are independent of probabilistic assessments. Consequently, the different degrees of confidence in the same inference have no influence on how successful it will be and thus cannot be a measure that specifies an inference category. If an epistemic reading of an inference were indicative of an inference type, its satisfaction would be sufficient to determine whether an inference type is successful. It isn't.

The inductive standard only expresses how confident one can be in an inference, but this standard doesn't express any inference in its own right. In the best case scenario, it could be described as a conclusion of a previous inference. But this very conclusion, "It's likely that it's not the case that *A* is true and *B* is false," would have to be regarded as true and not as merely likely to be true. Otherwise, we would have a reiteration of probability assessments such as "It's *likely that it's likely that* it's not the case that *A* is true and *B* is false," which only expresses how confident one can be in one's own epistemic inclinations to accept an inference. In order to make sense of (1), we need an actual commitment to the truth values involved instead of a mere epistemic reading of it.

The textbook examples of induction were introduced as a cautionary tale about the fallibility of empirical generalizations. It doesn't matter how large is your sample of white swans observations, the generalization that all swans are white can still be defeated by the counterexample of a newly discovered black swan. Inductive inferences are supposed to be defeasible. They require epistemic humility. This view is in stark contrast with the supposed infallibility of deductive reasoning. But if the failure of numerous mathematical proofs are not indicative of the nature of successful deductive inferences, the numerous failures of inductive inferences are not indicative of the nature of successful inductive inferences.

In an induction, the fact that an inference leads to an uncertain conclusion does not imply that the inference itself is uncertain. The nature of the propositional content presented in the conclusion doesn't reflect the reliability of the inference that draws the conclusion. To think otherwise would be a category mistake. It's also clear that inductive reasoning can lead to conclusions that will never face any counterexamples. For example, the generalization that the sum of two even numbers is even is definitive because all even numbers share the same properties. This inference is both inductive and certain. If we deny this result, induction will be reduced to uncertainty, rather than being regarded as distinct forms of inference such as generalizations and previsions.

Inferences that involve probabilistic factors are not necessarily uncertain inferences. For example, if I infer that the toss of a fair coin has a 50% chance of being a tails event, this is a reliable and predictable inference with a definite outcome. One could argue that what is meant by induction then is something entirely different: given an inherently random process such as a coin toss there is more than one possible event that can result from it. According to this interpretation, inductivism it's an ontological thesis about the nature of certain events and not a hypothesis about a particular type of inference. Now, if the distinction between induction and deduction is not a distinction between types of inference, but actually a distinction between the nature of different types of events, we can say that an event that has a 99% chance of occurring is inductive. So the mere addition of 1% chance would turn this inductive event into a deductive one. Intuitively, it seems that an event with 99% chance is nearly deductive, and that a nearly deductive event can't be distinguished from a deductive event with "the naked eye", so to speak. It's also plausible to think that relying on a cut off point of a measly 1% is a conventional decision that has no bearings in nature. Why not choose 0,5% or even 0,1% as a cut off point? When did the event becomes deductive?

In a non-deductive inference the conclusion should not be required by the premises, meaning that the conclusion is independent from the premises. It's not simply that the conclusion is not a consequence of the premises. The conclusion *cannot be* a consequence of the premises, because if a conclusion could be a consequence of the premises, it would be a consequence of the premises (since validity in one possible world implies validity in all worlds if we assume the accessibility relations of S5). But now we have to conclude that the conclusion results from the premise in a contingent manner. If the conclusion is connected with the premises only contingently, how it can be inferred from the premises or be supported by them? The epistemological concern is that you cannot be entitled to make an inference if the conclusion doesn't follow from the premises. If the premises are sufficient evidence for the conclusion, they ensure the conclusion. It follows by contraposition that if the premises don't ensure the conclusion, as it is the case in non-deductive inferences, they are insufficient evidence for the conclusion. Finally, if the premises are fallible, there are some possible worlds where the premises are true and the conclusion is false, which means that there are worlds where the premises are misleading. But there is no way of knowing if the premises are not misleading in the actual world. If any non-deductive inference is potentially misleading, there is no way of determining in which non-deductive inference we should trust. Thus, the acceptance of non-deductive standards of inference will lead us to a skeptical scenario because it implies the permanent possibility of falsity in even the best-grounded hypothesis.

Suppose I inferred *q* from *p* only because *q* was likely to be true given the truth of *p* and there were no other factors besides probabilistic considerations that motivated this inference. In this case, I had no assurances that the conclusion would be actually true given the premise since the inference is about a random process with a propensity for the occurrence of events that are favorable to the truth of the conclusion. But I didn't actually know that the conclusion would turn out to be true in this particular circumstance. So this was not really an inference to begin with, but merely an educated bet on a particular outcome. I bet that *q* would be true given the truth of *p* because this is what happens in most cases. So I was hopeful that *q* would be true given the truth of *p*. A genuine inference is not a bet, but a thinking process where the conclusion is drawn from the premises because it's believed to be a consequence of the premises.

What does it mean to say that *p* supports *q*, in an inductive inference? It means "*p* is a reason to accept *q*, but it doesn't entail *q*. It could be wrong". This description is infelicitous. If *p* is indeed a reason to accept *q*, it can't be misleading. So what should be said is something stronger, namely, "*p* is a reason to accept *q*, if it's not misleading". Now, according to fallibilism, any reasonable epistemic agent will be inclined to accept the following statement "I will rely on *p* in order to believe in *q* even if I can't know that *p* is not misleading". But this amounts to "I believe that *p* is a reason to accept *q*, but I don't know if *p* is a reason to accept *q*". The last statement is not incoherent, but it's very pessimistic. It means we will never be able to tell whether a supposed reason to accept a belief is a genuine reason. It turns every evidentiary concerns into blind guessing.

If a deductive inference is possibly valid, it's valid (since its validity in one world implies validity in all accessible worlds if we accept S5)[1]. But for a deductive inference to be defeasible it should be possibly misleading; this requires that there should be at least one possible world where the inference is thought to be valid when it's not. Now, if there is a possible world where a

1 It's not a coincidence that examples of deductivist reconstruction provide valid deductions. This is not because the deductivist is being charitable in her interpretation, but mainly because every deduction is valid. So according to deductivism, every inference is not only deductive, but also valid. So there are no invalid inferences, only failed attempts to make inferences. Propositions are asserted and there is an unsuccessful attempt to deduce one proposition from the others. But since this proposition is independent from the premises, there is no inference.

deductive inference is invalid, it cannot be valid, since this would amount to truth preservation across all possible worlds. So only invalid deductions would be defeasible. An inductive inference that aims for nomic necessity operates with a similar mechanism. If it's possibly valid (in the sense of being truth preserving in all nomically accessible worlds from the perspective of one nomically accessible world), then will be valid in this more restricted scope. But this inference can only be defeasible if there is one nomically accessible world where the inference is wrongly assumed to be valid, thus contradicting our initial assumption that the inference is possibly valid. Thus, only invalid inductive inferences, in our special sense, can be defeasible. Successful inductive inferences that aimed for important empirical discoveries couldn't be defeasible in principle.

If an inductive inference is uncertain it's perceived to be likely, but likely to be what? An inference cannot be likely to be true because only inferential constituents can have truth values, but a proposition about the combination of truth values of inferential constituents can have truth values. So let's say that accepting an induction is accepting that is likely that is not the case that the premise is true and the conclusion is false. But this means that the inference is likely to be valid, which aims the satisfaction of a deductive standard. If an induction is likely valid, then there are possible worlds where the propositional content of the premise implies the conclusion. Let's assume that these worlds are nomically accessible ones, so we can say that the conclusion is logically necessitated by the premise in most nomically accessible worlds.

In a strong induction, the inference is presented as a relation of evidential support where the conclusion should be accepted because is well supported by the available evidence described in the premise. But endorsing this inferential standard assumes the following requirement: "it's not the case that principles of evidential support are true and the proposition 'we should accept the conclusion supported by the evidence' is false". This requirement is deductive in character.

One could object that there is another way of interpreting inductive standards in terms of objective probabilities. According to this interpretation, the likelihood of an event occurring is an objective fact that is independent of epistemic elements. So an inductive inference would have to accept the following standard:

(4) It's unlikely that the premise is true and the conclusion is false.[2]

This proposition can have a metaphysical interpretation. But notice that inductive inferences in this sense are covertly deductive. Let's suppose that an event is likely because it has a 70% chance of occurring given the evidence. This statement can be summarized as follows:

(5) It's not the case that the observational statement is true and the event doesn't have a 70% chance of occurring.

And (5) can be interpreted as having the following quasi-propositional form:

(6) It's not the case that p is true and q is false.

---

2 This statement can be interpreted as being equivalent to (1):

(1) The conclusion is likely given the premise.

It's reasonable to interpret (1) as asserting that the conditional probability of the conclusion given the premise is high. If conditionals are interpreted as inferences, the inductive standard implies the equation:

(2) $\Pr(p \rightarrow q) = \Pr(q|p)$

Since the probability of an inference that satisfies the inductive standard will be measured by (1). But it's known that this will lead to the following triviality result:

(3) $\Pr(p \rightarrow q) = \Pr(q)$

Now, this counterintuitive result can be interpreted as saying that *the probability of an inference is the probability of the conclusion*, which is absurd. The result can be reinterpreted in a positive light if we assume that $\Pr(p \rightarrow q) = \Pr(p \supset q|p)$, which is less or equal to $\Pr(q)$, and this is reasonable. This means that an inference $p \rightarrow q$ is acceptable when $p \supset q$ is robust with respect to p, i.e., when $\Pr(p \supset q)$ is high and would remain high after learning that p. Since $p \supset q$ is logically equivalent to $\neg p \vee q$, the probability of $p \supset$ q given p is equal to the probability of $\neg p \vee q$ given p, which amounts to the probability of q given p. The truth conditions of material implication represent the most basic deductive standard for inferences. *So the inductive standard without modal qualifications is only acceptable when it's equivalent to the satisfaction of a weak deductive standard given the acceptance of the premise*.

The only problematic aspect of this inference is that it relies on conditional probability and the equation, but, one might argue, both assumptions are only plausible if we assume a suppositional view of conditionals, which are interpreted as a conditional assertion of the consequent given the antecedent. If we assume that conditionals are propositions (either if they are interpreted as connectives or claims to inference), this connection must be mistaken because it's the whole conditional that is asserted, not the consequent given the antecedent. In this case, what the triviality result shows is that there is no correspondence between (and we shouldn't identify) the probability of a conditional, understood as a proposition, and conditional probability. I attempt to explain the connection between inferential acts and claims of implication, including their relation to theories of conditional in the Appendix II.

Where p is the premise and q is the conclusion stating that the event is likely. But notice that (6) represents a classical deductive standard, not an inductive one. Moreover, these inferences can be interpreted as meta-observational reports about chancy events instead of actual inferences about which of these events take place. There is a difference between asserting "p" and asserting "probably p". In the first case it's assumed that one knows or at the very least believes in p. In the second case one believes that there are good chances that p without any additional commitments to its actual truth value. By asserting "probably p" no one is asserting p. Our uncertain beliefs, with rare exceptions, are not about probabilities. Probability assessments depend on the total available evidence, whereas true beliefs depend on truthmakers (Olin, 2003, p. 64). This is important because inferences in the real world will require not only probabilistic assessments about the chances of events, but actual conclusions about whether or not these events occur. The inferential act occurs when a conclusion is drawn from a premise. If the occurrence of the conclusion itself were merely likely no inference would take place.

It wouldn't be an exaggeration to suggest that the probability calculus is the cornerstone of inductivism, providing normative guidance to correct any of our mistaken intuitions about inference. But it's arguable that while statistical probability can be relevant to betting and chance phenomena, it has no significant relevance to rational belief (Pollock, 1983, p. 66). In fact, epistemic probability and degrees of confirmation are not in agreement with the axioms of the classical calculus (Olin, 2003, p. 77).

There is a tension between the notion of degrees of belief and the ordinary notion of belief according to which we either believe or not in something simpliciter. If we accept the probability calculus, if I'm ignorant about p I should adopt a degree of belief of 0.5 in p. This simple prescription will face counterexamples. Suppose that p and q are logically incompatible and I don't know if either p, or q, or the disjunction ($p \vee q$) is true. Let's say that p states that *the 327th* *automobile to pass through a certain intersection this morning is blue*, and q states that it's green. Since p and q are inconsistent, but have nonzero probabilities, there is no consistent attribution of probabilities where each proposition, including the disjunction, have 0.5 because according to the probability calculus pr($p \vee q$) > pr(p) and pr($p \vee q$) > pr(q) (Pollock, 2006, p. 86). Does it even make sense to suggest that indecision about p is a degree of belief of any kind? This can't be right. The indecision about p is the absence of belief in p.

According to the probability calculus, conditional probability should be measured by the ratio formula:

(RATIO) $\Pr(q \mid p) = \Pr(q \,\&\, p)/\Pr(p) \quad (\Pr(p) > 0)$

But there are good reasons to think that this definition of conditional probability is incorrect. Consider the probability that the Democrats win, given that the Democrats win. The corresponding conditional probability is 1, but the relevant unconditional probabilities in the ratio formula are vague. The same problem happens when we consider the conditional probability that the Democrats do not win, given the Democrats win. It's probability is 0, but the related unconditional probabilities are vague. Or suppose that T and F necessary and impossible, so (T, given the Democrats win) = 1 and (F, given the Democrats win) = 0. The ratio formula fails again in both cases. Another counterexample is (this coin lands heads, given the Democrats win) = ½, and the ratio formula will fail again due to the same reason (Hájek, 2003, pp. 293-294).

Another counterexample involves conjunctions. Suppose we have a series of probabilistically independent beliefs $A_1, \ldots, A_n$. The probability attributed to each belief is high but less than 1. The probability of the conjunction of A_is decreases with each additional conjunct that is added to this series. If the set of beliefs is sufficiently large, the probability of the conjunction is low. Since the probability of a statement and its negation are complementary according to the probability calculus,

$\Pr(p) = 1 - \Pr(\sim p)$

the negation of the conjunction has high probability (Orlin, 2003, pp. 71-72). Now suppose we have a long list of therapies that do not cure AIDS such as drinking eight glasses of water a day, taking aspirin etc. This list of therapies can be represented as $x_1, ..., x_n$. Let's also conceive an artificial predicate, "aica", which is defined as follows: an item is aica if and only if it is tested before *t* and is a cure for AIDS, or is not tested before *t* and is a cure for cancer.

Let's name the claim that x_i is not aica, A_i. Now, suppose each x_i in our long list of therapies was tested before *t* and it turns out that they are not the cure of AIDS, so they are not aica. Let's admit that we are justified in believing each A_i and their probabilities are equal. If the list of A_is is large enough, we are entitled to accept the negation of the conjunction of the A_is. But the negation of this conjunction implies:

(A) At least one of $x_1, ..., x_n$ is aica.

And this implies:

(A*) At least one of $x_1, ..., x_n$ is either a cure for AIDS or a cure for cancer.

But this consequence is absurd. None of the therapies in the list is a cure for AIDS or cancer (Olin, 2003, pp. 74-75).

It's one thing to claim that non-deductive standards are acceptable even if they fall short of deductive standards due to fallibilist assumptions. But it can be argued that non-deductive standards undermined deductive goals that should be within our grasp. If degrees of warrant worked like probabilities, it would be impossible to warrant a conclusion given a deductive argument that contains numerous premises that also happen to be uncertain. Consider a deductive argument with 100 independent premises, where each individual premise has a degree of warrant of .99. The conjunction of the premisses will have a degree of warrant of only .37, so these premises conjointly will never warrant us to draw a conclusion. For instance, imagine one surveys 100 people to know about their preferences related to two products. Each person express her preference by saying "I prefer x to y". 79 people preferred *A* to *B*. This conclusion is a deductive consequence of the data. Let's suppose the pollster is highly confident on each piece of data, but not certain. Her degree of warrant for each answer in the survey is .99. But if degrees of warrant are probabilistic, the conclusion that 79 out of the 100 surveyed reported preferring *A* to *B* would have a degree of warrant of measly .37, and so it wouldn't be warranted (Pollock, 2006, p. 91). So the notion that probability calculus and the theory of subjective probability provide a logic of induction is misguided.

Let's go back to inductions. Another characterization of induction that revolves around metaphysical assumptions is that in a deduction it's logically impossible for the premises to be true and the conclusion to be false, but in an induction the truth of the premises is consistent with the falsity of the conclusion. Thus the distinction between induction and deduction is expected to reflect the nature of the different modalities involved in each inference type. Maybe an inductive inference requires a form of necessity that is weaker than logical necessity or maybe the inference it's only about contingent truths. But as we shall see, there is no reason to assume that validity should be characterized solely in terms of logical necessity.

In some textbooks, deduction is presented as a limit case of induction where the conclusion is *certain* rather than probable given the premises. In this view, deduction is described through an epistemic lens where inference is described in terms of degrees of belief and evidential support. If the available evidence guarantees the truth of the conclusion, it's deductive. This epistemic reading is problematic because it implies that any inference that lacks truth preservation and has a certain conclusion is deductive. The other problematic consequence of this view is that any inference that preserves truth but has an uncertain conclusion must be regarded as inductive. Instead, what should be said is that if it's known that there are no circumstances where the premises are true and the conclusion is false, the conclusion is certain given the premises. But since knowing is factive, the epistemic aspect of a deduction must be a consequence of a metaphysical aspect that requires a certain configuration of truth values. The same interpretation holds for inductive inferences where a conclusion is probable given the premises if it's known that in most circumstances where the premises are true, the conclusion is true.

Instead of adopting an epistemic reading of deductions, we should adopt a metaphysical view where a deductive inference is successful when the conclusion is necessitated by the premises across all possible worlds, or, to put in other words, in which there are no possible worlds where the premises are true and the conclusion is false. This approach is independent from an epistemic reading for a multitude of reasons, including the facts that epistemic judgments have no bearing on metaphysical issues and that logical consequence may not be a luminous event.

Regarding deductive inferences, it's also important to observe that they have no distinctive propositional form because the supposed paradigmatic examples of deductions, such as *modus ponens* or hypothetical syllogism, are not inferential forms but coherence requirements for inferences (Silva, 2023b). The actual inferential forms in such requirements are the conditionals that are misinterpreted as premises (Silva, 2023c). Thus, the only thing we can safely say about deductive inferences is that they are such that the conclusion is necessitated by the premises. This freedom from logical form is important because it enables us to apply a modal necessity framework *to all forms of inference*. For instance, a generalization in mathematics, which is traditionally considered an inductive inference, is actually deductive because the conclusion is necessitated by the premises.

Inferences aim truth preservation within a modal range instead of mere degrees of confidence and evidential support. The mention of modal range is important because the essence of deduction is that in an inference the conclusion is necessitated by the premises within a specific modal range and not necessarily across all metaphysically possible worlds. There is a case to be made for deductions with varying modal ranges where inferences take place along a modal spectrum of necessity. For instance, it's perfectly reasonable to say that some successful inferences have a conclusion necessitated by the premises only within the actual world because they aim a narrow modal range. There are different levels of necessity for conclusions, different degrees of logical consequence, and a gradation of deductive inferences that is measured by different modal scopes[3].

The most important inductive inferences involve nomic necessities, but knowing how deeply the conclusion is necessitated by the premises depends on whether nomic and metaphysical necessities align. If they are co-extensive, the premises of inferences that rely on natural laws will necessitate the conclusion *metaphysically*. If they are not co-extensive, they can be divided into two groups: those where the conclusion is necessitated by the premises in all metaphysically possible worlds, and those where the conclusion is necessitated by the premises only relative to specific nomically possible worlds.

This deductivism hypothesis is a consequence of a conceptual analysis of the speech act "argument". Groarke says:

> We can see that it is always possible to deductively reconstruct an argument which is not transparently deductive by noting that any arguer is committed to the statement 'If the premises of my argument are true then the conclusion is true'. This follows directly from the implications of the speech acts 'argument' and 'assertion', for any arguer who argues for some conclusion C on the basis of some set of premises purports to believe that C is true and the her premises justify this belief, (cf. van Eemeren and Grootendorst: 1992 pp. 30-31). In this sense, their argument declares that they believe that these premises imply the conclusion and that the conclusion is true if the premise are true. It is perhaps worth noting that they are committed to the latter conditional not merely in the sense of material implication, but in the stronger sense that they must believe there is a relationship between their premises and their conclusion which makes it reasonable to base a belief in the latter on a belief in the former (Groarke, 1999, pp. 6-7).

3 This idea is motivated by the following observation: material implication has the same properties of formal implication (deduction), but is restricted to one world. It follows that there are different degrees of implication across a modal range. They all share the same properties, but differ in scope. The properties of classical implication can be generalized for all inferences and their respective modalities (logical, metaphysical, nomic etc). For a detailed defense of this thesis see Silva (2023).

This argument can be improved with additional modal distinctions. Any inference requires the bare minimum assumption that it's not the case that the premises are true and the conclusion is false. This assumption is the truth conditions of a material implication. If the reasoner intends the premises to support the conclusion in a wider modal range, she would assume that it's not the case that the premises are true and the conclusion is false in her presumed modal range. If this modal range includes all logically possible worlds, we now have the truth conditions of formal implication. If this range includes only nomically possible worlds, we have an intermediary implication (in this case, the truth conditions of material implication in all nomically possible worlds), and so on and so forth.

The inductivist will have to concede that even if an inductive inference doesn't aim any kind of validity, in the strong sense that there are no logically acessible worlds where the premises are true and the conclusion is false, she will have to concede that, at the very least, an inductive inference aims material validity, in the sense that it's not the case that the premises are true and the conclusion is false in the actual world. After all, if the premises are true and the conclusion is false, the inference will be abandoned. But since a material implication will have to satisfy *modus ponens*, *modus tollens* and other classical coherence requirements, it will be a deductive inference with deductive features.

There is also something to be said about cases of inductive inferences that are about matters of probability distribution as such. For instance, I may infer that a tails event has a 50 percent chance of occurring given a coin toss. It's arguable that this conclusion is necessitated by the premise in all epistemic possible worlds since probability calculus is an epistemic necessity. The modalities are as varied as the patterns involved in each inference. The predominant view is that epistemic possibilities are not objective. For instance, Williamson contrasts objective possibilities with epistemic possibilities where the first it "is not a matter of what any actual or hypothetical agent knows, or believes, or has some other psychological attitude to" (2016, p. 454). He insists that any question that is still open involve epistemic possibilities that "are not objective possibilities of any kind". The example he uses is that it's epistemically possible for us that the number of other planets that are inhabited is both bigger than 2 and smaller than 2 (Williamson, 2016, p. 454). But this is incorrect. The proposition that a given number *n* is *simultaneously* bigger than 2 and smaller than 2 is epistemically impossible given our mathematical knowledge. In fact, it's necessarily false. No number of any kind can be simultaneously bigger than and smaller than 2. What is epistemically possible is an entirely different proposition, namely, a disjunction that either the number of inhabited planets is bigger than 2 *or* smaller than 2. This epistemic possibility is also a logical, physical and metaphysical possibility that is determined by how things stand in the world, including the fact that at least one planet is inhabited.

McFetridge (1990, p. 137) also endorses the view that epistemic possibilities are not relevant to alethic modalities. But epistemic modalities are grounded on knowledge, which is ultimately grounded on truths (logical, nomological, metaphysical). If I know *p*, and *q* is epistemically possible in relation to *p*, then *q* is simply a possibility of a alethic kind that can be specified by the relevant knowledge in question. In fact, it's a truism. For instance, any claims to logical possibility will be made in relation to a given body of logical knowledge. So all epistemic possibilities in relation to logical knowledge are logical possibilities. In fact, we simply have no better guide to objective possibilities than epistemic possibilities.

The intuition that epistemic possibilities can be entirely subjective is probably motivated by the realization that our knowledge is always limited and in constant change. So what seems to be epistemically possible in the present might turn out to be a parochial prejudice that results from an imperfect understanding of reality. For instance, according to newtonian mechanics, it's epistemically possible that a body travels faster than light and gravity itself was believed to be a force that acts instantaneously. Both assumptions turned out to be physical impossibilities. There is a problem with this reasoning though. It confuses epistemic with doxastic possibilities, and limited understanding with empirical allegations that overreach their factual basis. The assumptions of newtonian mechanics were not knowledge based, so they don't count as genuine epistemic possibilities. Moreover, truths in any given domain must be consistent, so the notion that some proposition might be epistemically possible given the current knowledge and later reveal itself as impossible due to some new development is certainly false.

It's important to anticipate some objections. McFetridge (1990, p. 138) insists that "the claim that an argument is deductively valid involves a notion of necessity then it involves the strongest notion of necessity." He offers two arguments, that he presents as conditions, to support this claim. The first is that "it is a distinctive and important feature of deductive validity (one in which it contrasts with inductive strength) that adding extra premisses to a valid argument cannot destroy its validity" (McFetridge, 1990, p. 138). That monotonicity is implicit in inductive inferences can be illustrated as follows: let's assume that my belief in q can be justified based on the evidence of p, so I accept a conditional, $p \rightarrow q$. But I was told that inferences are non-monotonic, so I accept that my initial belief in $p \rightarrow q$ is compatible with $(p\&r)\rightarrow q$ being false (since antecedent strengthening represents monotonicity). But in order for that to happen, I would have to accept that my initial inference, $p \rightarrow q$, can be maintained simultaneously with its defeat by the new finding, r The only way to avoid this incoherence is to defend that $p \rightarrow q$ implies $(p\&r) \rightarrow q$. If the latter inference turn out to be unsustainable when it is reinforced by a proposition about new findings, it is because the initial inference that implies it was incorrect all along. It is conceptually intrinsic to the notion of inference that it should be perceived as resilient until it isn't[4].

4The general point I'm trying to convey is that belief revision is a consequence of a closure principle. If S realizes that p implies q, and it turns out that q is unacceptable given S's total evidence, then p is no longer justified for S. But this ability to change our beliefs retroactively given new information is only possible if there is a transmission of evidential support and justification closed under implication. See the Appendix I for some additional criticisms of fallibilism and its rejection of closure.

The second argument is that "there is this connection between deducing *q* from *p* and asserting a conditional: that on the basis of a deduction of *q* from *p* one is entitled to assert the conditional, indicative or subjunctive, if *p* then *q*." (McFetridge, 1990, p. 138)[5]. This condition is easily satisfied by other necessities. Suppose I deduce *q* from *p* because *p* nomically entails *q*. It's obvious that I would be entitled to assert the corresponding conditional, if *p* then *q.* Intuitively, this will hold for any varieties of necessity.

The conventional wisdom exemplified by McFetridge's position implicitly assumes that an inference can only be deemed deductive if it satisfies the following conditions:

(7) The conclusion is logically necessitated by the premises.

(8) It is knowable *a priori* that the conclusion follows from the premises.

Now, consider the following *modus ponens*[6]:

P1. If Hesperus is Phosphorus, then necessarily Hesperus is Phosphorus.

P2. Hesperus is Phosphorus.

C. Necessarily, Hesperus is Phosphorus.

---

5 McFetridge (1990, p. 138) discusses the following counterexample to his second argument: If "*p*, so *q*" is valid, then "*p* and not-*q* so *q*" is valid as well (due to the first condition). Consequently, the conditional "if *p* and not-*q* so *q*" should be assertable (due to the second condition), but it isn't. His answer is "No. For I here regard the antecedent as impossible in the fullest sense. The assertibility of this odd counterfactual can thus be seen as a special case of the vacuous truth of counterfactuals with fully impossible antecedents." (McFetridge, 1990, p. 138). But this answer is unhelpful because vacuous validity is a known property of deductive validity, or at least of classical validity. McFetridge has three choices: to abandon his second condition, to adopt a notion of deductive validity that is non-classical, or to provide an explanation of assertion of vacuous conditionals. If conditionals are interpreted as classical inferences, say, a material implication, a vacuous conditional can be asserted even it's epistemically useless.

6 Again, I don't think *modus ponens* is really an inference but rather a coherence requirement for inferences. However, this will not affect my argumentation, since one could simply reformulate Kripke's example as the following conditional: "If Hesperus is Phosphorus, then necessarily Hesperus is Phosphorus. Hesperus is Phosphorus. Necessarily, Hesperus is Phosphorus." Notice that the predominant assumption that *modus ponens* is a paradigmatic example of deductive inference is partially responsible for the notion that deduction should be exclusively *a priori*. If we accept that *modus ponens* is an inferential form, not only it's knowable *a priori* that the conclusion follows from the premises, as it's obvious, to the point that it could be regarded as an analytic matter since the conclusion is explicitly contained in the premises. This meager diet of examples induces the belief that the claim that a deductive inference is valid is analytic in the sense that one can determine whether the inference is valid simply by grasping the meaning of the propositions involved in the inference. It's natural to think that analytic statements are known *a priori*, so it will be natural to think that deductions are decided solely by *a priori* considerations. However, if we reinterpret these examples as coherence requirements, deductive inferences become free from logical form in a way that makes them harder to access. Consequently, any claim that a deductive inference is valid would lose its analytic status.

(Kripke, 1980, pp. 108-109)

P1 is known *a priori* because of the principle of necessity of identity, which is a theorem of modal logic. P2 is an empirical truth. The conclusion that follows from this combination of premises is a necessary truth, but that can be only learned through empirical investigation. This inference satisfies conditions (7) and (8) because it's an instance of *modus ponens*, so it's deductive in the usual sense; on the other hand, this inference is relevant precisely because it's not an *a priori* affair. Kripke's notion of the necessary *a posteriori* shows that some important truths can be obtained by simply mixing known premises from different epistemic modalities. So it's arguable that we can also say about this inference that:

(7*) The conclusion is metaphysically necessitated by the premises.

(8*) It is knowable *a posteriori* that the conclusion follows from the premises.

So is this inference deductive or not? The traditional view can claim that Kripke's inference is indeed deductive, but only because it satisfies (7) and (8), and not (7*) and (8*), which are regarded as extraneous for taxonomic purposes. This answer is unsatisfactory, since the inference is significant because of (7*) and (8*). The other possible reply is that Kripke's inference is not deductive, which seems to be even more unreasonable and prejudiced. Notice that (8) is assumed as corollary of (7), but the converse is not true. In principle, there is nothing that prevent us from conceiving a metaphysical necessity that is also knowable *a priori.* So there is no obvious connection between the fact that a conclusion is necessitated by the premises and how this fact is known. The crux of the matter is that what is essential about deductive inferences is not the types of epistemic and alethic modalities involved, but the fact that the conclusion is necessitated by the premises. The manner by which this fact is knowable should be completely incidental to our understanding of deductive inferences[7].

7 The very significance of the distinction between *a priori* and *a posteriori* has been questioned. For instance, intuitively, a rational intuition is a mental state that should count as an experience, but then *a priori* knowledge cannot be independent of experience. Hawthorne (2007, p. 201) claims that the importance of the distinction "has been grossly

It's also possible to interpret modal properties in degrees with the notions of modal force and comparative possibility. This idea was first suggested by Lewis (1973), sect. 2.5. and developed by many others (Kment, 2006). The degree of possibility of a proposition *p* is measured by a metric of comparative closeness. The closest a *p*-world is to the actual one, more easily it could have been actual, which in turn determine its degree of possibility. *p*'s degree of impossibility increases in the same proportion of departure from actuality that is required for it to be true. The degree of necessity of *p* increases in the inverse proportion of how lower is the degree of possibility of ~*p*. The degree of necessity of a proposition consists in its modal force, meaning its robustness or inexorability in relation to the variation presented by other words. The more stable or sturdy is a proposition, the more necessary it is. The contrast is with contingent truths that could easily have been false. Different modal ranges of necessity (metaphysical, nomic, logical) require different grades of modal force. For instance, metaphysical necessity would require a higher degree of modal force than nomic necessity. If we abandon the limit assumption, according to which there is exactly one closest *p*-world, and instead assume that there are infinitely many closest *p*-worlds (Lewis, 1973, p. 424), then there are infinitely many degrees of possibility and necessity.

Lange (2005, p. 278) also proposes a unified framework of all varieties of necessity (logical, metaphysical, conceptual, physical …). What the different varieties of necessity have in common is that a necessary truth is a truth that would still have been true under a certain range of counterfactual perturbations. There are various strata or grades of necessity according to their range of invariance given the counterfactual perturbations.

3. Induction and fallibilism

overestimated" and Williamson (2013, p. 295) that it "does not cut at the epistemological joints". I will not address these criticisms here. The main point, for the purposes of this article, is that if the same cognitive mechanisms are usually employed in both *a priori* and *a posteriori* knowledge, there is no significant epistemological difference between inferences that rely on the two.

I argued that paradigmatic examples of deductive inferential forms such as *modus ponens* and *hypothetical syllogism* are not really inferential forms, but coherence requirements for inferences. This implies that any inference must be deductive in order to be coherent. One immediate objection is that these are coherence requirements for the type of inference used in these examples, namely, material implication. Since material implication can be considered a formal implication restricted to one world, the argument begs the question against the critics. One reply is that if these examples were not coherence requirements for inferences, the following claim would be coherent:

> (9) It's not the case that the premise is true and the conclusion is false, but the conclusion may still be false when the premise is true.

But (9) is incoherent. The critic might object that she subscribes to the following claim:

> (9*) It's unlikely that the premise is true and the conclusion is false, but the conclusion may still be false when the premise is true.

But since (9*) is perfectly consistent, there is no incoherence in her objection. However, (9*) is a cop-out since it's motivated by a different claim, namely:

> (10) The truthmaker of the premise is what makes the conclusion true, but the conclusion may still be false when the premise is true.

And this is obviously problematic. The point is that (9*) seems plausible because it depends on an epistemic reading of inferences that still leaves basic metaphysical commitments out of the equation. There is only a caveat in the formulation of (10). It's arguable that in most inductions, it's the truthmaker of the conclusion that is responsible for the truth of the premise, and not the other way around. When I conclude that every metal will expand when heated because this behavior was observed on a piece of copper, it's the natural law of the conclusion that is instantiated by the premise. So we can have a different formulation:

(11) The truthmaker of the conclusion is what makes the premise true, but the conclusion might still be false when the premise is true.

The incoherence of (11) is even more glaring. If we consider a simple prevision, both the premise and the conclusion will be made true by the same truthmaker, for example, "since this piece of metal expanded when heated, the next piece of metal will expand in the same conditions." So we have something along the lines of:

(12) Both the premise and the conclusion are made true by the same truthmakers, but the conclusion might still be false when the premise is true.

Again, this seems nonsensical. Perhaps the inductivist can claim something as follows:

(13) The truthmaker of the premise makes the conclusion true most of the time, and the premise can be true and the conclusion false.

This is consistent, but then she would have to add the following clause that justifies her inference:

(14) The truthmaker of the premise makes the conclusion true most of the time; if the premise is true, the conclusion is true this time, yet the conclusion might be false when the premise is true.

And once again, we have an incoherent statement. The challenge faced by any inductivist is to provide a metaphysical interpretation of induction that is not incoherent. If the proposition that is likely to be true is also assumed to be known, then any claims that it might still be false will be incoherent because knowledge is factive. This suggests that indeterministic standards assumes, and thus inheres, the problems of fallibilism. In fact, it can be argued that fallibilism posits that remote possibilities that defeat our epistemic aspirations are always open. So *fallibilism can be considered as a subtle form of skepticism* that is usually presented as a matter of epistemic humility. This is ironic because fallibilism was developed as an antidote to infallibility and its excessive demanding epistemic criteria such as certainty. But if an epistemic agent cannot make any claims about her own knowledge without simultaneously conceding epistemic possibilities that contradict it, knowledge will be unattainable.

Thus, there is a case to be made for an alternative infallibilist position. The infallibilist assumption means that we the need to rethink our understanding of inferences. The prevalent notion is that most inferences are non-monotonic, since many of the past inferences were abandoned by their proponents after new discoveries. But no one advances an inference assuming that it is going to be abandoned. Instead, one argues for an inference given the expectation that will be resilient upon new discoveries. It would be an incoherence otherwise. It is conceptually intrinsic to the notion of inference that it should be perceived as resilient until it isn't. Similarly, if I think I know *p*, I can't sincerely question my belief that *p* will be resilient to new information because this would result in incoherent statements such as 'I know *p*, but I might be wrong about *p*'. If the inference doesn't meet this requirement, it is going to be abandoned. The fallibilist mindset requires something akin to a pessimistic induction based on previous mistakes, but the fact that a proposition that was assumed to be known turned out to be false doesn't imply that knowledge shouldn't be indefeasible any more than the fact that a perceptual error demonstrates that perception shouldn't be trusted[8].

That there is a connection between infallibilism and deductivism can be demonstrated as follows:

P. In order to know that *p*, one has to rule out the possibility that *p* is false.

C. Therefore, S knows *p* on the basis of *e* only if S knows *e* and *e* logically entails *p*.

Dutant (2007, p. 68) calls P and C, epistemic infallibilism and evidential infallibilism, respectively. If we adopt P, but reject C, evidentiary concerns will have to be inductive and will not satisfy P. If S

8 For a defense of infallibilism see Bird (2007), Blome-Tillmann (2009), Chisholm (1982), David (2001, p. 163), DeRose (1991), Dodd (2011), Dretske (1978), Dutant (2007, 2016), Goldman (1978), Lewis (1996), Merricks (1995, 1997), Moon (2012), Neta (2011), Nozick (1981), Plantinga (1993), Schaffer (2004), Sturgeon (1993), Williamson (2000) and Zagzebski (1994).

knows *p* on the basis of *e*, but *e* doesn't entail *p*, S cannot rule out the possibility that *p* is false. not-*p* will be unlikely given *e*, but not conclusively ruled out as a possibility. The connection between deductivism and infallibilism can also be demonstrated as follows: deductivism can be interpreted as the claim that inferences that are not absolutely conclusive (e.g., invalid inferences) are irrational (Stove, 1970, p. 88), whereas infallibilism is the claim that S knows *p* on the basis of *e* only if S knows *e* and *e* is absolutely conclusive evidence for *p*. So deductivism and infallibilism are interdependent.

The notion of deductivism defended in this article includes other alethic modalities, including inferences where the conclusion is necessitated by the premises nomically or metaphysically. So we have to adopt a version of infallibilism that encompass these modalities as well. This version was proposed by Dutant (2007, p. 74) in what he calls "modal infallibilism", where S knows that *p* only if S's belief that *p* could not have been wrong[9], and the notion of possibility requires an alethic modality, not an epistemic one. The only problem of this approach is that follows the usual assumption that epistemic possibilities are not relevant to alethic modalities.

It's important to observe that there is a difference between inductive inferences as such (generalization and prevision) and indeterministic standards that describe when an inductive inference can be accepted (when the conclusion is likely given the premises). Indeterministic standards are completely silent about the truth values of both the premises and the conclusion, and agnostic about whether the reasoner assumes the truth of the conclusion at all. The statement that "*p* is likely, but I don't know if it's true" amounts to "*p* is possibly true, but I don't know if it's actually true". What happens is that a substantial inference requires a commitment to the truth value of the conclusion, but a mere probabilistic assessment only describes the probability of the conclusion

9 Notice that Dutant himself rejects epistemic and evidential infallibilism as inadequate expressions of the infallibilist hypothesis, and favors his modal infallibilism hypothesis instead. But it's clear that modal infallibilism would require something very similar to evidential infallibilism even if the implication process involved other modalities, for instance, that *p* nomically entails *q*.

based on the available evidence. It's not enough to assert that the conclusion is highly likely given the premises. In order to make an inference, the reasoner has to claim that the conclusion is true given the premises, which means that a commitment to the truth value of the conclusion in the actual world must be made, otherwise nothing is inferred. In other words, inductivism doesn't describe the inferential act that is supposed to express, because it simply reinstates modal indeterminism in a less sophisticated fashion. This explains why inductive inferences can be reinterpreted as deductions without inconsistency. Deductivism fills the metaphysical vacuum left by inductivism.

It can also be objected that indeterministic standards assume that *the probability space maps the possibilities without telling us which possibility is actual*. "*p* is likely to be true, but I don't know if *p* is true" amounts to "*p* is possible, but I don't know if *p* is actual". But since there is no direct relation between probabilities and the actual world, there are no epistemic aspirations to knowledge based on probability that goes beyond probability distributions. This generates a gap between probability assessments and the actual world whenever the likelihood of an event is not maximum, which is all the time in empirical matters. Suppose you decide to participate in a fair lottery with one winning chance among twenty million tickets. It's highly likely that you bought a losing ticket, but you can't claim to know that you have a losing ticket despite the overwhelming odds against you. What you can say is that it's highly likely that you have a losing ticket, but you can't disregard the remote possibility that you have a winning ticket. If you could claim that your ticket will lose, you would have to conclude that every ticket will lose because every ticket has the same chances of winning, but this would contradict the initial assumption that there is one winning ticket (Kyburg, 1963, p. 30). The knowledge of the probability distribution involved in the lottery is not enough to assert whether a particular ticket will lose or not. The lottery example reinforces infallibilism. The admission that there is a remote possibility in which you have a winning ticket undermines any pretensions that you know that you have a losing ticket (Lewis, 1996, p. 551).

Notice that one of the main objections against infallibilism is that it leads to skepticism. If one knows *p* there is no chance that not-*p*, but, continues the argument, there are always chances that skeptical scenarios are real. So one doesn't know that *p*, where *p* can be interpreted as a proposition that one is not a brain in a vat. The problem with this reasoning is that it's circular. If one truly knows *p*, any intuition that not-*p* is possible is a modal illusion. In other words, the claim that a skeptical scenario is a possibility begs the question against the infallibilist. Moreover, infallibilists will interpret fallibilism as a form of skepticism, so this objection will not break any intuition tie in favor of fallibilists even if it were successful.

4. The metaphysical assumptions of inductivists and deductivists

There is something to be said about the metaphysical assumptions of inductivists and deductivists. The notion that some inductive inferences are inherently probabilistic and irreducible to deductive standards is motivated by the acceptance of modal indeterminism. But modal indeterminism without the acceptance of possibilism is not enough to justify inductivism because in this scenario any inference will collapse to a material implication. This occurs because if possibilism is false, then everything there is, is actual; then all we can say about an inference is whether or not there is a combination with true premises and a false conclusion. This would be true even if some phenomena are random since this would be the only world where these chancy events take place. So it is possibilism, and not indeterminism, that is carrying the epistemic burden for inductivism, and possibilism requires a defense of possibilia. To be fair, it seems that any deductivist would be inclined to accept possibilism, so the debate about the underlying metaphysical assumptions would revolve around the acceptance of indeterminism. Most people would be prone to accept inductivism as a formality, almost as a corollary that follows from the indeterminism of quantum mechanics. But knowing if indeterminism is indeed the only available theoretical framework for quantum mechanics is a philosophical notion that is open to debate.

Now, this framing of the inductivists' assumptions puts deductivists in a difficult spot. If inductivism requires indeterminism, should we infer that deductivism presupposes determinism? If we accept determinism, what follows is actualism because if every event is predetermined by antecedent events and conditions together with natural laws, the actual world is the only world that exists. Any counterfactual scenario that had the same initial conditions and laws of nature would be identical to the current world. So the deductivist will have to abandon possibilism, and with it, the very notion that in a deduction the conclusion is necessitated by the premises in a modal range. Every inference becomes a material implication, which is the weakest form of deduction. The only alternative is we assume that the initial conditions could be different. In this case, the total configuration of events could be different even if every event is determined. But in a world were all events are predetermined, it's reasonable to assume that initial conditions would be "set in stone" as well.

It can also be argued that deductivism is compatible with indeterminism. Suppose an inference that deals with indeterministic phenomena is such that it's not the case that the premise is true and the conclusion is false 99% of the time. This inference can be interpreted in a deductive manner as follows: each chance represents a nomically possible world. In a way, this inference is truth-preserving in 99 out of 100 nomically possible worlds, which means that the conclusion is necessitated-ish by the premises in a specific modal range. In this case, we can say that an event has a 10% chance of occurring, or of being distributed in 10% of all nomically accessible worlds.

The connection between probability and modality is self-evident. Suppose an event has a 90% chance of happening, but fails to happen in the actual world. Does this mean that it would happen in 9 out of 10 nomically possible worlds? If the answer is "no, it can also fail in every one of them due to a streak of bad results across multiple worlds", we would start to question whether this event really had a 90% chance after all. It seems that for each random result in the actual world, the remaining possibilities would still need to be accounted for in some modal range. Imagine that proton decay exists in nature even though no observation of this rare phenomenon was ever recorded in a particles accelerator. Perhaps its probability is so low and the conditions for its occurrence are so extreme that this phenomenon will never be displayed in the actual world. The way in which we could make sense of this phenomenon is that protons will decay in some counterfactual situation that departs from actuality. Its low probability also suggests that proton decay would have to be very distant from the actual world, otherwise it would be easily produced. Conversely, the more probable is an event, the closer it is to the actual world. Let's call this view the modal view of probability. There is still one additional detail we need to take in consideration. It can be argued that there are infinitely many possible worlds. In this case, the relation between possibility and probability can be expressed as an average of the results due to the law of large numbers.

This modal view of probably allows us to solve some puzzles related to counterfactuals in a chancy world. Hawthorne presents the following example: "Suppose I drop a plate. The wave function that describes the plate will reckon there to be a tiny chance of the particles comprising that plate flying off sideways." So the counterfactual "If I had dropped the plate, it would have fallen to the floor." is false because if I had dropped the plate, it might have flown off sideways (Hawthorne, 2005, p. 396). Hájek (2007, p. 6) makes a similar point. Suppose that in relation to a coin that will never be tossed one asserts the following counterfactual: "If the coin were tossed, it would land heads". This counterfactual is false. He enthusiastically quotes Jeffrey (1977) and Stalnaker (1984, pp. 164-165) for endorsing the same notion that there is no fact of the matter of how the coin would land, since there is no fact of the matter in a chancy world.

It seems that one of the solutions for the problem of contingent future can be applied to this case with our modal interpretation of probability. We should make a distinction between what will be the case with what should be the case. What will happen is not inevitable. It's not something that need to happen (Torre, 2011, p. 365). The same could be said about what could have happened in a counterfactual scenario. Some would have happened even if it shouldn't have happened. The fact that a coin would have landed tails doesn't imply that it should have landed tails. What would be the case doesn't need to be the case. Some events are chancy in the sense that they could have being different from what they are, but this doesn't change the fact that something will have to be the case in other words. Only one among a set of possibilities will be actual. This doesn't require second-guessing of an indeterministic process since it's a consequence of our understanding of probability.

Hájek (2007, p. 26) also states that events with probability 0 can happen. His example is that a coin that is to be tossed repeatedly infinitely many times, it might land tails on every toss, even though the chance of this is 0. Our understanding of probability contradicts this statement. Suppose I claim that there is no inconsistency in assuming the occurrence of a quasi-miracle event, say, an event of a fair coin that will land only heads in a potentially infinite series of tosses. Let's call the proposition about the occurrence of this event "*p*". If we accept that each individual result represents a possibility, then *p* will have to be a necessary truth. But this means that we will have to describe an inherently indeterministic event as a determined one. This is an inconsistency. So we will have to abandon our initial assumption of the possibility of a fair coin that only land heads indefinitely. If chances are possibilities, our probabilistic assumptions must be revised according to our modal intuitions, and not the other way around.

Appendix I
Fallibilism and closure

I mentioned before that the lottery example (the lottery paradox) reinforces infallibilism, so is a problem for fallibilists. Another problem is the fallibility paradox. If I have found mistakes in my past justified beliefs originated by a method M, it's reasonable to assume that there are mistakes in my present justified beliefs that are originated by the same method. Consequently, I have my present justified beliefs p_1, …, p_n, but I'm also entitled to believe in ~(p_1 & ... &p_n). Thus I'm entitle to believe in inconsistent statements (Olin, 2005, p. 238). This inconsistency results from the underlying skepticism implicit in fallibilism. The fallibilist wants to maintain a principle of epistemic modesty that we ought not to believe that all our justified beliefs are true, but that implies that we ought not to believe in all our own justified beliefs, which is incoherent. Worse, it implies that some of our justified beliefs are wrong, but we can't ever tell which one. Thus the principle of epistemic modesty leads to a skeptical scenario.

At its core, the fallibility argument is an inductive argument, but this means that it's possible that its conclusion is false even if its premises are true. One is entitled to have a justified belief that some of her other present justified beliefs is false based on her past mistaken beliefs. But this very belief about the fallibility of our beliefs can be one of her present mistakes, as far as she knows. So fallibilism may be mistaken due to its own criteria. The only way to remove this cloud of suspicion from fallibilism is assuming that belief in fallibilism is certain, but any requirement of this sort will be incoherent.

Note that a popular solution adopted by fallibilists in order to prevent these paradoxes is to abandon the closure principle and its variants, but it can be argued that by doing so they are actually abandoning consistency. If S is justified in believing *p*, and knows that *p* implies *q*, then S is justified in believing *q* as *a matter of consistency*. It's not surprising then that fallibilists are relying on these paradoxes to openly endorse the rationality of inconsistent beliefs (e.g., Olin (1989; 2005, p. 238); Foley (1970); Klein (1985)).

Now consider Multiple Premise Closure (MPC), which is one of the variants of the closure principle:

> If S is justified in believing $p_1, ..., p_n$, and knows that $p_1, ..., p_n$ jointly imply q, then S is justified in believing q.

Olin (2005) attempted to undermine MPC with the following example: suppose that Bernard has a lot of debt and giving his current income he will not be able to pay off his debt in the next month. But Bernard bought a ticket in a Super-Lotto. If he wins the prize in the Super-Lotto he would be able to pay off his debt next month. We can conclude that Bernard will not win the lottery by *modus tollens*. The inference can be summarized as follows:

(1) Bernard will not be able to pay off his debt next month.

(2) If Bernard wins the lottery, then he will be able to pay off his debt next month.

(3) Bernard will not win the lottery.

The problem is that if MPC is correct, premises (1) and (2) constitute evidence for (3), but this is intuitively wrong. The evidence that lead us to accept (1) and (2) (information about his financial situation and the lottery prize) doesn't constitute evidence that Bernard will not win the lottery (Olin, 2005, p. 243).

It can be objected that given the available facts in the context, the only means by which Bernard will be to pay off his debt next month is by winning the lottery. So (2) is actually a biconditional. In this case, any evidence that justifies the belief that Bernard will win the lottery also justifies the belief that he will be able to pay off his debt next month. By contraposition, any evidence that justifies the belief that he will not be able to pay off his debt next month also justifies the belief that Bernard will not win the lottery. These propositions are not independent, but are contingent on the epistemic status of each other. In order to make this interpretation more realistic, suppose I found out that Bernard, against all expectations, managed to pay off his debt in one month. It seems that given the known facts the best explanation is that he won the lottery prize. The converse is also true. If Bernard is broadcast in the news as the winner of the prize, I can safely conclude that he is now able to pay his debt. So the counterexample is disarmed.

The closure principle also motivates these principles:

Transmission of Evidential Support: If *e* is evidence for *p*, and *p* implies *q*, then *e* is evidence for *q*.

Transmission of Justification: If S is justified in believing *p*, and knows that *p* implies *q*, then S is justified in believing *q* on the basis of *p*.

Olin (2005) presents the following counterexample against these principles: suppose I'm justified in believing

p: Cynthia will watch television the entire evening on Wednesday.

because she announced her intention to do this; *p* implies

q: There will not be a power failure on Wednesday evening.

This is supposed to be a counterexample to both principles. The evidence for *p* is Cynthia's announced intention, but this is not evidence for *q*. I'm also justified in believing *p*, but that doesn't mean that I'm justified in believing *q* (Olin, 2005, p. 237). But this example is only plausible if *p* is interpreted as synonymous with "Cynthia *intends* to watch television the entire evening on Wednesday." and this proposition clearly doesn't imply *q*. Note that this and other apparent counterexamples to closure principles follow the same recipe:

- rely on an evidential basis for the premise that only satisfies fallibilist requirements;

- deny that the evidential support is transmitted due to infallibilist requirements for the conclusion;

But if the evidential support is not transmitted due to infallibilist reservations, then the initial premise should be disregarded as lacking in evidence for the same reason. After all, if *p* could be justified according to stronger infallibilist scruples, then *q* would follow as well, since any defeater that would undermine *p* would have to be disregarded from the get-go. A genuine counterexample to closure principles would have to observe either fallibilist or infallibilist guidelines, but never both. If the inference is governed by fallibilist guidelines across the board, the premises will not imply the conclusion, since there is always a possibility that the conclusion is false when the premise is true. On the other hand, if there is an attempt to observe infallibilist reservations, the transmission of evidential support and justification will follow seamless. In other words, a genuine counterexample to a closure principle is an impossibility since it would require fallibilist guidelines in a relation of logical entailment that prevents defeaters.

Another reasoning that shows that counterexamples to the closure principle are bounded to fail is the following: any argument against a closure principle is deductive, due to its own conceptual nature; but this argument will only be compelling if it satisfies closure principles such as MCP and transmission of both evidential support and justification. For imagine that a philosopher concocts an elaborate counterexample against closure, but these principles are false. Then even if I accept the evidence for (or I'm justified in believing in) her premises, and know that the premises jointly imply the conclusion, I would still not have evidence for (or being justified in believing in) the conclusion of her argument that disproves closure principles. So the way to defeat closure principles requires closure principles and thus it's incoherent.

The only alternative that avoids inconsistency is to present non-deductive arguments against closure, but any argument of this sort will be unconvincing due to the simple fact that it has to avoid deductive inferences in order to criticize closure. There is also another problem with this approach:

if epistemic support (or knowledge) are not transmitted with implication, there is even less reason to assume that they will be transmitted with non-deductive arguments. The reason is crystal: if inferences in which a conclusion's truth are ensured by the premise's truth can still fail the transmission test, there is simply no hope for inferences in which the premise's truth is no preserved. Of course, this is not a coincidence: the transmission of epistemic support and knowledge should at the very least require the transmission of truth, but transmission of truth only occurs in deductive arguments.

Appendix II

Conditionals and inference

Suppose I claim that *q* follows from *p*. This is a claim that it's logically impossible for *p* to be true and *q* false. This claim cannot be valid or invalid, but only true or false. Validity (logical consequence or entailment) is a relational property between two or more propositions' truth values across all logically possible worlds. Let's say that *p* implies *q* broadly iff there are no possible worlds where *p* is true and *q* is false. This is broad implication. We can also say that *p* implies *q* locally if it's not the case that *p* is true and *q* is false in a given world. Broad implication means that there are no worlds without local implication. Now consider that I actually infer *q* from *p*. This is an inferential act. The act itself is not true or false. Ok. But can it be valid? Broad validity means that it's logically impossible for *p* to be true and *q* false, but I cannot infer *q* from *p* in all possible worlds. My individual inferential act is local, but it's supposed to be the basis of a broad implication that has an immense (potentially infinite) modal range. I could say this: "When I infer *q* from *p*, I know in one go that *q* should be inferable from *p* in all possible worlds". But I would only be willing to infer *q* from *p* in all possible worlds If I accept the claim that *q* follows from *p*. The inferential act[10] and the corresponding proposition that the relevant inferential act is justified are different, but related and ultimately in agreement with each other.

Suppose I claim that *q* follows materially from *p*. This means that *p* locally implies *q* in one world (say, the actual world). Thus I accept that it's not the case that *p* is true and *q* is false in the actual world. I would be willing to infer *q* from *p* in the actual world. Now suppose *p* is false in the actual world. Here we have at least three answers. One is that this result satisfies the initial claim of material implication vacuously. So the inference is justified for vacuous reasons, but it will be useless and harmless, since *p* is false. The other intuition is that this inference is no longer justified, since I'm only willing to infer *q* from *p* in a context where *p* is true. But this answer contradicts our earlier assumption that material implication holds when it's not the case that *p* is true and *q* is false in the actual world. So this view assumes a different notion of implication where *p* implies *q* in a word *w* only if both *p* and *q* are true in *w*. Should we say that if *p* is false in *w*, *p* doesn't imply *q* in *w*? Not really. If *p* doesn't imply *q* in *w*, *p* is true and *q* is false in *w*. In other words, this view has to abandon the traditional notion of invalidity as well. There are two options here. The first is to accept that the inference of *q* from *p* is invalid when *p* is false or *p* is true and *q* is false. This seems too strong. The other is to argue that no inference is made when *p* is false, so claims about invalidity are meaningless in this case. But this means that if *p* is false in *w*, we cannot make any meaningful claims about whether *q* follows from *p* in *w*. This lead us to the third intuition that *p* implies *q* in a word *w* iff in the closest-*p*-worlds, *q* is true. The closest *p*-worlds may involve *w* or not. Now notice that when I make this assumption I'm making drastic changes: first, it no longer matters whether *p*&*q* is true in the world where the implication is supposed to take place; the rationale from this deviation is that we are supposed to find out whether there is a connection between *q* and *p* in a similar world. So even if *p* is impossible relatively to the past of the word *w* (in the sense that the

10 The term "inference" and its correlates (reasoning, argument …) are misleading. They suggest that the conclusion follows from the premises in a temporal sense, that it comes later as a consequence, that it has some movement. Maybe it does from an epistemic perspective, but from an objective perspective there is no movement of any kind. You only need a configuration of propositions and their respective truth values. So making new discoveries amounts to find new and unexpected "neighbors" of the same worlds in which the premises are true.

available facts were set and *p* turned out to be false) one would still push to a closer, but different *p*-world. You consider whether *q* follows from *p* in w knowing that *p* is false in w. Therefore, if one claims that *q* follows from *p* in all worlds, but *p* is necessarily false, one could still push to impossible *p*-worlds to judge this claim. This feature is not a bug, but is intended by design. The hypothesis was crafted to conceive worlds where the premises are true, even if they happen to be false.

The first intuition above is the material account of conditionals. Its downside is that it implies the possibility of vacuous validity and vacuous truths. This intuition is a consequence of the classical notion of logical implication. The second option is the suppositional view of conditionals. It's incompatible with the classical notion of implication because claims of implication are meaningless when the premises are known to be false. Notice that despite what its proponents would want us to believe, the suppositional view is not really inconsistent with a propositional view of conditionals, since you still have both an inferential act and a corresponding implication claim about this inferential act. This also implies that Quine's criticism of Bertrand Russell's use of material implication as a use-mention fallacy is also misguided. The third choice is that the possible world theory of conditionals. Its peculiarity is that it can evaluate an implication claim in impossible worlds.

Each theory has its strengths and weaknesses. The material account is particularly counterintuitive in accepting the existence of relations of vacuous implication that happen to be local. The reason is that there is always a temptation to bypass vacuous implication looking for a closer world where the premises are true. Notice that claims of unrestricted vacuous implication are usually regarded as harmless. But this is an incoherence from the detractors of vacuous implication. If vacuous implication is acceptable when it covers all possible worlds, it should be acceptable as well when is restricted to fewer worlds, say, one world. The suppositional view states that an inference can only occur when the premises are true, or assumed to be true. But this would render *modus tollens* invalid. If this consequence is not enough to discard this hypothesis, I don't know what will. The possible world theory implies that I can determine the validity of a local implication claim about a word *w* where the premise is false looking for a very similar but different world. This reasoning is faulty. In order to determine whether *p* broadly implies *q*, I consider each individual *p*-world in order to determine whether *q* is true. These are the rules of the implication game. The possible world theorist makes a sleight of hand when she decides that local implication is invalid in a world_n because of a similar, but distinct world_{n1}.

To understand this issue more closely, suppose *p* broadly implies *q* because *p* is necessarily false. Let's say that I checked all the possible worlds and they have a fixed number *n*. The possible world theorist will not hesitate in conceive an additional world, in this case, an impossible world, in order to disprove my claim of vacuous implication. But this reasoning is fallacious: the original claim was intended to cover *n* worlds, not $n + 1$ worlds. Of course, a claim of implication that covers n + 1 worlds may be false, but this is not the implication claim that is being evaluated by the possible world theorist. The same could be said about the evaluation of a claim that *p* locally implies *q* in *w* that also happens to be vacuous. If this claim is evaluated in w_1 it may be locally invalid in w_1, but *w* and w_1 will still be fundamentally different worlds. The possible world theorist might insist she is intuitively correct in her assessment, but this should give us pause. The whole point of a vocabulary with possible worlds is that it allows us to express our modal intuitions clearly. Once we decide to include impossible worlds in our modal vocabulary (words that cannot be conceived, by definition), this rationale is broken. Moreover, if we accept her reasoning, it follows that any inference is invalid. The reason is obvious: I can simply evaluate the premises and the conclusion in different worlds to obtain the combination of truth values I want. Possible world theories of conditionals are prone to this contextual fallacy because the observance of contextual requirements are only needed when we consider the relation of truth values between more than one proposition. If we consider a single conditional by adopting a truth condition that aims to be an ontological analogous of the Ramsey's test, then any questions about how this single conditional will remain in the context set is a non-starter. The problem becomes clear when we consider the relation between one conditional and other conditionals, or between one conditional and other non-conditional propositions. If a conditional is judged to be true because of truth value attributions in a given world *w*, then additional propositions in the same inference (conditional or not) will have to be analyzed in the same world. If for some reason, two propositions cannot be evaluated in the same worlds, say, if the closest world of a conditional is different from the closest world of another conditional in the same inference, then there is simply no inference in the way we currently understand the term.

Let's say the possible world theorist decides to bite the bullet and maintain that every inference is invalid. One can always argue that a logical analysis of conditional sentences will reveal that their truth conditions are indexed to a modal range. Of course, most people are not conscious of this implicit assumption, but they are important nonetheless. To see why that's the case, notice that any regular proposition contains temporal indexicals that can be articulated in its complete form, but in relation to which a speaker might not be fully aware. A proposition in its complete form would have to include every element that is required to determine its truth conditions, including dates, time, location and a modal index. If there is no modal index, we would have no means to determine the truth value of a proposition effectively. The proposition could concern any of infinitely many worlds as far we are concerned and we would have no means of knowing which. So we have to add a modal index. But once this detail is added, the possible world theory lose much of its appeal. Suppose I claim that p implies q in w. This means that p_w implies q_w. If p_w is false, the claim that p_w implies q_w is vacuously valid, but any counterfactual scenario of a similar world w_1 will be irrelevant, since we will evaluate the claim that p_{w1} implies q_{w1}, which is entirely different. The only drawbacks of postulating modal indexes is that the claim that p broadly implies q will have to be reinterpreted as a claim that p_w, p_{w1}... p_{wn} broadly implies q_w, q_{w1}... q_{wn}. This requires some revision in the way we speak, because each premise and conclusion should be a modal neighbor.

References

Bird, A. (2007). "Underdetermination and Evidence," In Bradley Monton (ed.), *Images of Empiricism* (pp. 62–82). Oxford: Oxford University Press.

Blome-Tillmann, M. (2009). "Knowledge and Presuppositions," *Mind*, 118: 470, 241–94.

Dodd, D. (2011). "Against Fallibilism," *Australasian Journal of Philosophy*, 89: 4, 665-685.

Chisholm, R. (1982). "Knowledge as justified true belief," In *The foundations of knowing* (pp. 45–47). Minneapolis: University of Minnesota Press.

David, M. (2001). "Truth as the epistemic goal," In M. Steup (ed.), *Knowledge, truth, and duty* (pp. 151-169). New York: Oxford UP.

DeRose, K. (1991). "Epistemic Possibilities," *The Philosophical Review*, 100: 4, 581-605.

Dretske, F. (1978). "Conclusive reasons," In G. Pappas & M. Swain (eds.), *Essays on knowledge and justification* (pp. 55–62). Ithaca, NY: Cornell University Press.

Dutant, J. (2007). *The case for infallibilism*. Latin Meeting in Analytic Philosophy, Genova, 20-22 September 2007. CEUR-WS Proceedings, ISSN 1613-0073, 59–84.

Dutant, J. (2016). "How to be an infallibilist," *Philosophical Issues,* 26: 1, 148–171.

Fetter, J. H. (1981). *Scientific Knowledge*. Dordrecht: D. Reidel Publishing Co.

Foley, R. (1970). "Justified Inconsistent Beliefs," *American Philosophical Quarterly*, 16, 247-57.

Gerritsen, Susan. (1994). "A Defense of Deductivism in Reconstructing Unexpressed Premises," In F.H. van Eemeren and R. Grootendorst (eds.), *Studies in Pragma-Dialectics*, (pp. 41–7). Amsterdam: International Centre for the Study of Argumentation.

Goldman, A. I. (1978). "A causal theory of knowing," In G. Pappas & M. Swain (eds.), *Essays on knowledge and justification* (pp. 65–86). Ithaca, NY: Cornell University Press.

Groarke, L. A. (1992). "In Defence of Deductivism: Replying to Govier," In Frans van Eemeren, Rob Grootendorst, J. Anthony Blair, and Charles Willard (eds.), *Argumentation Illuminated* (pp. 113–21). Amsterdam: International Society for the Study of Argument.

Groarke, L. A. (1995). "What pragma-dialectics can learn from deductivism, and what deductivism can learn from pragma-dialectics." In F.H. van Eemeren, R. Grootendorst, J.A. Blair, and C.A.

Willard (eds.), *Proceedings of the Third ISSA Conference on Argumentation* (pp. 138-145), vol.2 Analysis and Evaluation, Amsterdam: Sic Sat.

Groarke, L. A. (1999). "Deductivism within pragma-dialectics," *Argumentation,* 13, 1-16.

Groarke, L. A. (1999b). "The fox and the hedgehog: on logic, argument, and argumentation," *Protosociology,* 13, 29-45.

Groarke, L. A. (2002). "Johnson on the metaphysics of argument," *Argumentation*, 16, 277-286.

Groarke, Louis. (2009). "A 'Deductive' Account of Induction," In *An Aristotellian Account of Induction: Creating Something from Nothing*. Montreal & Kingston: McGill-Queen's University Press.

Hájek, A. (2003). "What conditional probability could not be," *Synthese*, 137, 273–323.

Hájek, A. (2022). Most counterfactuals are false. ms.

Hållsten, Henrik. (1999). Deductive Chauvinism. *Synthese*, 120, 49–59.

Hawthorne, J. (2005). "Chance and Counterfactuals," *Philosophy and Phenomenological Research*, LXX: 2, 396-405.

Hawthorne, J. (2007). A Priority and Externalism, in Sanford C. Goldberg (ed.), *Internalism and Externalism in Semantics and Epistemology* (pp. 201–218). Oxford: Oxford University Press.

Jeffrey, R. (1977). "Mises Redux," in Robert E. Butts and Jaakko Hintikka (eds.). *Basic Problems in Methodology and Linguistics* (pp. 213–222). Part Three of the Proceedings of the Fifth International Congress of Logic, Methodology and Philosophy of Science. Dordrecht: Reidel.

Johnson-Laird, P. N., Byrne, M. J. (1991). *Deduction*. East Sussex, U.K: Lawrence Erlbaum Associates.

Kitcher, P. (1981). "Explanatory Unification," *Philosophy of Science*, 48, 507–531.

Kitcher, P. (1989). "Explanatory unification and the causal structure of the world," In Philip Kitcher and Wesley Salmon (eds.), Scientific Explanation (pp. 172-177). Minneapolis: University of Minnesota Press.

Klein, P. (1985). "The Virtues of Inconsistency," *The Monist*, 68, 105–135.

Kment, B. (2006). "Counterfactuals and the Analysis of Necessity," *Philosophical Perspectives*, 20: 237–302.

Kripke, S. (1980). *Naming and Necessity,* Cambridge, MA: Harvard University Press.

Kyburg, H. E. (1963). "Probability and randomness," *Theoria*, 29: 1, 27–55.

Lambert, K., Ulrich, W. (1980). *The Nature of Argument*. New York: MacMillan.

Lange, M. (2005). "A Counterfactual Analysis of the Concepts of Logical Truth and Necessity," *Philosophical Studies*, 125: 277–303.

Lewis, D. (1973). “Counterfactuals and Comparative Possibility,” *Journal of Philosophical Logic*, 2, 418–446.

Lewis, D. (1996). “Elusive knowledge,” *Australasian Journal of Philosophy*, 74, 549–567.

McFetridge, I. (1990). “Logical Necessity: Some Issues”, in J. Haldane and R. Scruton (eds.), *Logical necessity, and other essays* (pp. 135–54). London: Aristotelian Society.

Merricks, T. (1995). “Warrant entails truth,” *Philosophy and Phenomenological Research,* 55: 4, 841–855.

Merricks, T. (1997). “More on warrant’s entailing truth,” *Philosophy and Phenomenological Research,* 57: 3, 627–631.

Moon, A. (2012). “Warrant *does* entail truth,” *Synthese*, 184: 3, 287–297.

Musgrave, Alan. (2012). “Deductivism Surpassed: Or, Foxing in its Margins?” *Journal for General Philosophy of Science*, 43: 1, 125–132.

Neta, R. (2011). “A refutation of Cartesian fallibilism,” *Noûs,* 45: 4, 658–695.

Nosich, Gerald. (1982). *Reasons and Arguments*. Belmont, ca: Wadsworth.

Nozick, R. (1981). *Philosophical explanations*. Cambridge, MA: Harvard University Press.

Olin, D. (1989). “The Fallibility Argument for Inconsistency,” *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, 56: 1, 95–102.

Olin, D. (2003). *Paradox*. Chesham: Acumen Publishing Limited.

Olin, D. (2005). “A Case Against Closure,” *VERITAS*. Porto Alegre, 50: 4, 235–247.

Plantinga, A. (1993). *Warrant and proper function*. New York, NY: Oxford University Press.

Pollock, J. (1983). “Epistemology and Probability,” *Noûs*, 17: 1: 65-67.

Pollock, J. (2006). *Thinking About Acting*. New York: Oxford University Press.

Salmon, W. (1989a). “Deductivism,” In Philip Kitcher and Wesley Salmon (eds.), *Scientific Explanation* (pp. 172-177). Minneapolis: University of Minnesota Press.

Salmon, W. (1988/1998). “Deductivism, Visited and Revisited”, in A. Grünbaum and W. Salmon (eds.), *The Limitations of Deductivism* (pp. 95–127), University of California Press, Berkeley/Los Angeles/London. Reprinted in *Causality and Explanation* (pp. 142–165), Oxford University Press, Oxford.

Schaffer, J. (2004). “From Contextualism to Contrastivism,” *Philosophical Studies*, 119: 1–2, 73–103.

Silva, M. (2023). The Inextricable Link Between Conditionals and Logical Consequence. ms.

Silva, M. (2023b). Coherence of Inferences. ms.

Silva, M. (2023c). ‘If-then’ as a version of ‘implies’. ms.

Simard, P. (2007). “On deductivism: a critical survey of deductivism in informal logic”. Thesis, University of Windsor.

Stalnaker, R. (1984). *Inquiry*, MIT Press, Cambridge, MA.

Stove, D. (1970). “Deductivism,” *Australasian Journal of Philosophy*, 48: 1, 76–98.

Sturgeon, S. (1993). “The Gettier problem,” *Analysis,* 53: 3, 156–160.

Suppe, F. (1997). “Science Without Induction,” in J. Earman and J. D. Norton (eds.), *The Cosmos of Science: Essays of Exploration* (pp. 386–429). University of Pittsburgh Press, Pittsburgh.

Thomas, S. (1981). *Practical Reasoning in Natural Language*. Englewood Cliffs, nj: Prentice Hall.

Torre, S. (2011). “The Open Future,” *Philosophy Compass*, 6: 5, 360–373.

Williamson, T. (2000). *Knowledge and Its Limits*, Oxford: Oxford University Press.

Williamson, T. (2013). How Deep is the Distinction between A Priori and A Posteriori Knowledge? In Albert Casullo & Joshua C. Thurow (eds.), *The a Priori in Philosophy* (pp. 291-312). Oxford: Oxford University Press UK.

Williamson, T. (2016). “Modal science,” *Canadian Journal of Philosophy*, 46: 4-5, 453–492.

Yezzi, R. (1992). *Practical logic*. Ronald Yezzi. Mankato (ed.), Minn.: G. Bruno.

Zagzebski, L. (1994). “The inescapability of Gettier problems,” *Philosophical Quarterly*, 44: 174, 65–73.