A Non-cognitivist Theory of Necessity

(Work in progress. Most proofs and some other parts are, for the moment, unfortunately missing in action. The general line of thought should however be identifiable.)

Staffan Angere
University of Lund
staffan.angere@fil.lu.se

ABSTRACT

I present here an outline of a theory of necessity, inspired by thinking about the Carnap-Quine debate on analyticity. According to this theory, and contrary to what may be called the received view in contemporary metaphysics, there is no such thing as absolute necessity. Instead all necessity is relative to a framework, and what framework we use is primarily context-dependent. Claims of necessity are not strictly true or false, but are instead used to specify a framework. I furthermore apply the theory to questions of analyticity, causality and knowledge, in order to show how it may help us understand the connections between these problems.

1. Quine and Carnap on analyticity

The necessary supposedly comes in many different guises: we have the causal, logical, analytical, biological, psychological and practical kinds, for instance. Attempts to reduce one or several of these to the others have been common, as well as more general arguments that, even if a straightforward reduction cannot be made, one of them gives the sense of what is 'truly', fundamentally necessary. In contrast to these, I will argue here that there is no such thing as necessity *per se*, and no privileged interpretation of what 'necessary' means.

This might at first seem paradoxical. Isn't it *really* necessary that 2 + 2 = 4, or that everything is self-identical, for instance? According to the theory of necessity presented here, the answer is 'no'. Necessity is relative to a conceptual framework, and no framework can encompass all others.

This idea of framework-relative necessity goes back to Carnap (1937, 1956), who defines the necessary as that which follows from rules of language. There is, however, another tradition that passes through Quine, who identifies the necessary with those beliefs we are least likely to revise, and arrives at a concept of necessity that admits of degrees. We will gather inspiration from both of these approaches for our own theory, so first of all we will take a look at the debate between these philosophers.

Although W. v. Quine and Rudolf Carnap share many views on philosophy – most of all their staunch empiricism - their approaches to the question of necessity show important differences. As Quine puts it, ‘Though no one has affected my philosophical thought more than Carnap, an issue has persisted between us for years over questions of analyticity and ontology’ (1976, p. 203). The fundamental among these, Quine holds, is that of analyticity or logical truth; it is from his theory of analyticity that Carnap derives his theory of ontology, with its distinction between internal and external questions of existence, for instance (Carnap 1950).

Carnap’s theory of necessity begin with his characterization of language as calculus in *The Logical Syntax of Language* (1937, p. 4). An *analytic* sentence is one that is ‘L-valid’, i.e. valid in terms of logical consequence alone. *Logical* consequence, as opposed to *physical* consequence, is singled out by being stable under replacement of descriptive symbols. What characterizes the logical (non-descriptive) symbols is that the semantic rules of the language suffice to determine the truth-value of every sentence made up from just these symbols. (ibid. §50, §51).

What is logical and what is non-logical is thus determined by semantical rules (more specifically rules of consequence, or *c*-rules), and these are part of the definition of the language in question. During a series of papers (most notably *Empiricism, Semantics and Ontology* and *Meaning Postulates*), Carnap stresses our creative freedom in designing such languages, since the introduction of one is nothing but the proposal to adopt a certain system of conventions, much like one may make up the rules for a new game at will. These systems, rather than being true or false, may then be chosen according to how useful or perspicuous they are for a given purpose. The adoption of a system *as such* does not commit us to anything substantial; only holding certain nonanalytical sentences in the system as true or false does.

Quine’s views on analyticity are well known from his *Two Dogmas of Empiricism*, where his rejection of the analytic-synthetic distinction is coupled to his holistic theory of knowledge. According to this theory (the ‘web of belief’ theory), our beliefs may be ordered according to how close they are to experience, and the closest ones are the ones most likely to be given up in face of an anomaly. Analyticity is thus a matter of degree, rather than an all-or-nothing matter: the beliefs commonly thought to be ‘analytic’ are roughly those that are most entrenched. But this means that Carnap's distinction between internal questions (which are about 'facts'), and external questions (which are about what language to adopt), becomes blurred as well. In *On Carnap's Views on Ontology* (1951), he summarizes the difference and his own 'thorough-going pragmatism' as follows:

> Carnap maintains that ontological questions [i.e. external existence questions – *S. A.*], and likewise questions of logical or mathematical principle, are questions not of face but of choosing a convienient conceptual scheme or framework for science; and with this I agree only if the same be conceded for every scientific hypothesis. (1976, p. 211).

Looking closer we can see that the seeds of this collapse between framework and content were present already in Carnap's own work. Most of the time, Carnap seems to take what language we work in as fixed. But surely he would accept the idea that scientific change can involve changing the language? Suppose then that we switch the language of physics from one whose space-time is Euclidean to one where it is not. We do not then need to throw out all our old observations; we can reinterpret them in our new language instead. But this means, in a way, that the *same* fact (for instance, that the global curvature of space-time is zero) may be analytical in one language and synthetic in another. Since the adoption of either of these languages is a matter of convention, however, the question of whether the curvature of space-time is a matter of fact or convention becomes a matter of convention itself.[1]

Quine's blurring of the fact-convention distinction should not make us believe that nothing is fact and everything is convention, however. The point is rather this: there is no principled way to distinguish between a framework and a substantial theory. Yet, it seems Carnap is correct in pointing out the importance of such a distinction for everyday science. When we engage in inquiry, we do not regard 'everything' as up for revision; it is even doubtful what that would mean. At the very least, we have to have something to stand on in order to ask questions, since these always come with some kind of presuppositions.

Trying to steer a middle course between Quine and Carnap, we may hold the following: in a given scientific context, we generally work inside a conceptual framework, whose semantic rules constitute the logic of our inquiry. In another context, some of these rules may be regarded as factual rather than logical. But this means that what is factual and what is not a

[1] Carnap himself comes close to this standpoint in his later writings; in *Meaning Postulates*, for instance, he tells us that the question of whether 'all ravens are black' is to be taken as analytic or not is one of free decision.

matter of fact itself – indeed, the very question becomes pointless. This is the reason why we have allowed ourselves use a discussion of analyticity, or logical necessity, as basis for our upcoming theorizing about necessity in general. Giving up the logical-factual distinction simply means that any kind of necessity can be regarded as logical necessity in some system we may adopt.

2. Theories and frameworks

In view of our discussion of Quine and Carnap, we will present a theory of necessity that identifies what is *necessary* with that which we are holding to be part of our framework in the current context of inquiry. Following Quine, however, we will not make any essential difference between such a framework and a substantial factual theory. The present section contains an outline of a theory of theories that makes no such distinction.

Let a theory A be a closure operator C_A (A's *consequence operator*) on the powerset of some set Σ_A (A's *subject matter*). It does not matter what the elements of Σ_A are; we may take them to be any kind of entities that may be true or false, such as beliefs, sentence tokens, sentences, depictions etc. Generally, we refer to them as *possible claims* (or just claims), since they delimit what may be claimed to be the case using the theory in question. That C_A is a closure operation means that it is a function C_A: $\wp(\Sigma_A) \mapsto \wp(\Sigma_A)$ that fulfils the following conditions, for all $\Gamma, \Delta \subseteq \Sigma_A$:[2]

(*expansivity*)	$\Gamma \subseteq C_A(\Gamma)$
(*idempotence*)	$C_A(\Gamma) = C_A(C_A(\Gamma))$
(*monotonicity*)	$\Gamma \subseteq \Delta \Rightarrow C_A(\Gamma) \subseteq C_A(\Delta)$

A theory, as we interpret it here, is structurally indistinguishable from a *deductive system* in

[2] Alternatively, we could define a theory using a *consequence relation* $\vDash_A$ of the Gentzen type, defined as holding between sets of claims and single claims. It is easy to show that every consequence operator gives rise to a unique consequence relation, and *vice versa*.

Tarski's sense (cf. Tarski 1955, ch. XII), although we do not use this term for it in order to avoid its connotations of aprioricity or analyticity. In particular, a theory is meant to be able to fulfill the role of a conceptual framework or language. It is what we use to justify inferences and to express ourselves. In the limit, we call p *A-true* iff $p \in C_A(\varnothing)$. We refer to the set of all A-truths as $\top_A$.

Say that the sets of claims Γ and Δ, or single claims p and q, of the theory A are *A-equivalent* (symbolically $\Gamma \equiv_A \Delta$) iff $C_A(\Gamma) = C_A(\Delta)$ (or $C_A(\{p\}) = C_A(\{q\})$. Given any theory A, we can define its *reduct* $\rho(A)$ as the theory whose subject matter $\Sigma_{\rho(A)}$ consists of the equivalence classes of Σ_A under the relation $\equiv_A$ and whose consequence operator $C_{\rho(A)}$ is defined by $C_{\rho(A)}(\Gamma) = C_A(\cup\Gamma)$, for all $\Gamma \subseteq \Sigma_{\rho(A)}$. Since, as is easy to prove, $\equiv_A$ is a congruence with respect to the operator C_A, $\rho(A)$ will have essentially the same structure as A, but will identify claims that are A-equivalent.

Theories can be naturally ordered as stronger or weaker: A is stronger than B iff they have the same set of claims, and every inference allowed in B is allowed in A as well. Furthermore, say that A is a theory *in B* (in symbols $A \preccurlyeq B$) iff A and B share subject matter, and there is a $\Delta \subseteq \Sigma_A$ such that $C_A(\Gamma) = C_B(\Gamma \cup \Delta)$, for all $\Gamma \subseteq \Sigma_A$. It follows directly that if A is a theory in B, then A is stronger than B, although the converse does not hold. Instead, the theories in B correspond to those stronger theories obtained by "fixing" the set of claims Δ as true, while keeping the general structure.

The reason why $A \preccurlyeq B$ expresses an important intertheoretic relation is that if it holds, then the consequence operator of A itself can be expressed using nothing but the claims and consequence operator of B. In other words, all that the theory A says can be said in B as well. But this is not the only possible relationship between theories: a theory can also be *part* of another theory (in symbols $A \sqsubseteq B$), if A's set of claims is a subset of B's, and their consequence relations agree on this subset. It is easy to show that $\sqsubseteq$ is a partial order, and the following theorem proves the same for $\preccurlyeq$.

Theorem: $\preccurlyeq$ is a partial order.

Proof: A partial order is a reflexive, transitive and antisymmetric relation. The first two properties are trivial, and the third is obtained by letting $C_A(\Gamma) = C_B(\Gamma \cup \top_A)$ and $C_B(\Gamma) = C_A(\Gamma \cup \top_B)$, for all $\Gamma \subseteq \Sigma_A$, and showing that this entails that $C_A = C_B$. Let X be an arbitrary subset of Σ_A. Then, $C_A(X) = C_B(X \cup \top_A)$. Letting $Y = X \cup \top_A$, we obtain that $C_B(Y) = C_A(Y \cup \top_B) = C_A(X \cup \top_A \cup \top_B) = C_A(X \cup \top_B)$, so $C_A(X) = C_A(X \cup \top_B)$. In the special case $X = \varnothing$, expansivity then gives us that $C_A(\varnothing) = \top_A \supseteq \top_B$. A symmetric argument shows that $\top_B \supseteq \top_A$, so $\top_A = \top_B$. But replacing $\top_A$ by $\top_A$ in $C_A(\Gamma) = C_B(\Gamma \cup \top_A)$ gives us that $C_A(\Gamma) = C_B(\Gamma \cup \top_B)$, so $C_A(\Gamma) = C_B(\Gamma)$. *Q. E. D.*

These two relationships may be collected into one: we say that B *frames* A (written as $A \leq B$) iff $\Sigma_A \subseteq \Sigma_B$ and there is some $\Delta \subseteq \Sigma_B$ such that $C_A(\Gamma) = C_B(\Gamma \cup \Delta) \cap \Sigma_A$, for all $\Gamma \subseteq \Sigma_A$. We will also express this as the fact that A can be *expressed* in B. The following theorems show that $\leq$ can replace the relations $\preccurlyeq$ and $\sqsubseteq$, and that $\leq$ is a partial order as well.

Theorem: $A \leq B$ iff there is a theory C such that $A \sqsubseteq C$ and $C \preccurlyeq B$.

Proof: (to appear).

Theorem: $\leq$ is a partial order.

Proof: (to appear).

That B frames A means that B can be used as a conceptual framework for expressing the theory A in, in the sense that whenever we hold A as true, we could just as well have held a certain set of claims in *B* as true as well, since as far as *A*'s subject matter is concerned, the consequence relations that we are committed to are identical.

This characterization of the relation between framework and theory brings out the insights of Quine and others clearer than the traditional language/sentence dichotomy. A language can not be used by itself for making a claim, and a set of sentences does not constitute a language. In contrast, a theory as such is neither a claim nor a framework, but can be used for making a claim, or as a framework for expressing other theories in, as we choose. We *can* of course express Carnap's sentences and languages as theories. Let $\mathcal{L}$ be a language, or in Carnap's terms a 'sort of calculus, that is to say, a system of formation and transformation rules concerning *expressions*, i.e. finite, ordered series of elements of any kind, namely, what are called *symbols* (…)' (Carnap 1937, §46). Among the transformation rules are the rules of consequence, which lay down what sentences follow from what sets of sentences (§48). We can then define a theory *L* with a subject matter that consists of the set of well-formed sentences in $\mathcal{L}$, and a consequence operator C_L that takes every subset Γ of *L* to the set of sentences in $\mathcal{L}$ that are consequences of Γ.

Now consider a sentence *p* of $\mathcal{L}$. Holding *p* as true means that we are entitled to draw inferences from *p* as well as those that can otherwise be drawn in $\mathcal{L}$, which corresponds to accepting a theory *P* such that $C_P(\Gamma) = C_L(\Gamma \cup \{p\})$. Thus, every sentence in $\mathcal{L}$ corresponds to a theory as well.

When we are engaged in scientific inquiry, we are moving *inside* a framework. There are inferences we are allowed to make that do not require external justification. For instance, a chemist who uses the Pauli exclusion principle does not need to explain why it holds; physical laws are taken as fixed by chemists, unlike the properties of specific chemical substances, which are the topic of inquiry. Moving upwards in the hierarchy, the biologist does not question the chemical properties of substances, the neurologist takes cellular biology as fixed, etc.

Rather than a matter of convenience, this division of labor must be recognized as a prerequisite for all science. As Wittgenstein once put it, 'all explanation ends somewhere', and we may likewise say that all justification must begin somewhere, or nothing can be justified at all. We have to take something as fixed – at the very least a logical system, but often much more – in order to even begin giving justifications. Even the radical skeptic, who argues for our not having any justified grounds for belief, must use some standards for what constitutes a valid or invalid argument.

Justification does not always have to start at the same place, though, and is not even a matter of having to justify 'more' or 'less', as may have been suspected from my mention of a hierarchy. We can, for instance, fix the axioms of ZF set theory, and consider how the theory itself varies when we change the underlying logic from classical to intuitionistic, or we can fix the logic and consider various systems of set theory. We can hold microbiological theory fixed, and study how different chemical laws could give rise to the same behaviour, or *vice versa*. This highlights a central difference between this theory of frameworks and Quine's 'web of belief' theory: for Quine, the remoteness from experience is an objective fact about the individual belief. For us, what is more or less entrenched is fundamentally context-dependent, and varies with the context of inquiry.

3. Necessity, analyticity and logic

What framework to use is not uniquely determined given a theory, since many conceptual frameworks can be used to express the same theory. Adding a context may narrow down the class of theories that can be taken as frameworks: some inferences are generally allowed in some contexts. Still, in normal inquiry the framework we use will generally be quite underdetermined, which leads to the notions of valid inference and justification being vague.

Precification of the context, and limiting this vagueness, is one use of necessity. If I claim that *p* is *necessarily* true, that is an indication that I am working (or at least wish to be working) within a framework *A* such that *p* is *A*-true. One may even interpret it as a kind of prescription: 'let us hold the truth *p* fixed', and this is why this theory of necessity is called non-cognitivist. It's prescriptivist aspects somewhat resemble those Hare claim for ethical language in *The Language of Morals* (1952).

It does however contain a factual part as well: necessary truth is a species of truth, so by claiming *p* to be necessary, I also hold it to be true. This makes the claim that *p* is necessary possible to falsify by showing that *p* is false. If *p* indeed *is* true, however, strict truth or falsity are not applicable to the claim that *p* is necessary. Still, not all contexts may allow us to work in any kind of framework. So *given* a context, there are certainly frameworks that are acceptable and unacceptable. Just like we can talk about what a specific ethical theory allows or disallows even if we are ethical non-cognitivists, we can talk about what is necessary or contingent in a specific context.

This context-dependence can be interpreted as the source of all our different notions of necessity: since different contexts allow different inferences, we may classify contexts by the kinds of frameworks are acceptable in them. The physically necessary may, for instance, be classified as that which freely may be assumed in all contexts where physics is presupposed. But we can also use a claim such as '*p* is necessary', without qualification, to *talk about* the frameworks acceptable in the context of utterance, rather than to specify or change that context. Used this way, we can reclaim some of the force of truth and falsity of a necessity-claim as follows:

> '*p* is necessary', uttered in context *c*, is *true* iff *p* is an *X*-truth for every theory *X* acceptable as framework in *c*, and *p* is true.

> '*p* is necessary', uttered in context *c*, is *false* iff *p* is an *X*-truth for no theory *X* acceptable as framework in *c*, or *p* is false.

It is easy to see that if the context is not specified enough to allow just a single framework (such as when we do formal logic, for instance), '*p* is necessary' may fail to be either true or false. This is, as has been pointed out regularly by Dummett (most systematically in *The Logical Basis of Metaphysics*), a sign of antirealism. This should not surprise us. The distinction between necessary and contingent truth does not exist without us, and although it does not have to be consciously laid down as a convention by us, still it arises as a product of the social nature of scientific inquiry.

Claims of necessity thus have two functions: first, to specify the context, and second, to talk about it. I believe that which of these is applicable to a given utterance will have to be settled on a case by case basis. If we have specified what frameworks are acceptable beforehand, the second reading will make best sense, while otherwise the first may be more reasonable. It is, in any case, not my intention to give a linguistic analysis of necessity-claims here, but rather to theorize in general about what lies behind them.

On our theory, *p* is necessary iff *p* is a truth of every currently acceptable framework, and because what frameworks we accept in a given context at least partly are a matter of convention, necessity becomes dependent on convention as well. But this invites the objection that not all adoptions of a framework seem to be like this. Certainly, in some cases, we may lay down conventions at will that allow us to infer one truth from another. But in other cases, the necessity as such simply does not seem to be up to us.

Answering this line of criticism will take most of the rest of this essay, and involve considerations pertaining to different kinds of necessity (i.e. different contexts, on our theory). First of all, though, we should note that in many cases, it is based on intuition. It does not seem to us that we could alter the laws of logic, mathematics or physics simply by *fiat*, but this is of course just what it seems like. Scientific method may establish sentences as true or false, but does it really prove them to be necessary or impossible, in some substantial meaning of these terms? Our intuitions are notoriously poor guides, on their own, to what is true or false.

So, why does it appear impossible that 2 + 2 = 5? One reason may be the Kantian one: it is because of how we are as beings. We cannot picture such a situation, so we see it as impossible. The same goes for any kind of analytic (or even *a priori*) necessity. A judgment, according to Kant, is analytic when the predicate is contained in the subject. Since both predicate and subject are ideas, analyticity becomes a psychological phenomenon on this conception. But the fact that we cannot imagine something as true does not mean that it is false. Consider, for instance, spacetime in the general theory of relativity or string theory, or the wave-particle duality, or quantum field theory. All of these may be held to be impossible to picture accurately, but that does not make them false.

The fundamental problem with placing too much weight in appearances of necessity is the same as that which appears in placing too much weight on conventions: the link from necessity to truth is severed. That something is unimaginable does not entail that it is false, just as something's being conventionally decided does not have to make it true. In short, these philosophies of modality do not afford a justification of the T axiom, which is crucial for the theory to be one about *necessity*, rather than modality in general. In our own theory, we have gone around this by explicitly requiring that the necessary consists only of those *truths* that we hold fixed.

A better option for a definition of analyticity may be to abandon Kantian psychologism, and, like most philosophers during the 20th century, take a sentence to be analytic iff its truth follows from the meanings of its terms. So, assume that we have a theory L (our language) whose subject-matter consists of sentences. Let μ_L: $\Sigma_L \mapsto \Sigma_M$, where Σ_M is a set (informally a set of sentence-meanings, whatever they might be) be a *meaning function* for L. An *interpretation* of the theory L is an assignment of such a meaning function to L.

Given an interpretation μ_L, we say that q *follows analytically* from the set of sentences $\Gamma = \{p_1, \ldots, p_n\}$ iff a certain relation $\vDash_M$ holds between $\mu_L[\Gamma]$ and $\mu_L(q)$. What $\vDash_M$ is will depend on what we take meanings to be: if μ_L, for instance, takes every sentence in Σ_L to the set of worlds where it is true, then $\Gamma \vDash_M a$ can be interpreted as the relation that holds iff $\bigcap\Gamma \subseteq a$. In the limit, p is *analytic* iff $\mu_L[\varnothing] \vDash_M \mu_L(p)$, i.e. iff $\varnothing \vDash_M \mu_L(p)$. If the meanings are arbitrary sets of worlds, this means that p is analytic iff $\mu_L(p)$ is the set of all worlds. In general, since $\vDash_M$ must be able to give rise to a consequence operator in L, we take it to fulfill the usual Gentzen conditions on a consequence relation, for all subsets Γ, Δ of Σ_M and all claims $p, q \in \Sigma_M$:

(*reflexivity*)	If $p \in \Gamma$ then $\Gamma \vDash_M p$.
(*monotonicity*)	If $\Gamma \vDash_M p$ and $\Gamma \subseteq \Delta$, then $\Delta \vDash_M p$.
(*cut*)	If $\Gamma \cup \{p\} \vDash_M q$ and $\Gamma \vDash_M p$ then $\Gamma \vDash_M q$.

Analyticity, thus defined, is obviously relative to a meaning function μ_L. How do we determine it? The first point, of course, is that it does not determine itself. The context helps us a lot: if a sentence appears in the middle of an English text, we have reason to interpret it in English, for instance. But such reasons are always defeasible. Imagine that someone found an early writing by me, which made it clear that I constantly use the word 'bachelor' as a euphemism for ‘Bach-impersonator’. Then the truth of an occurrence of ‘Peter is a bachelor’ here would not guarantee the truth of ‘Peter is unmarried’, because the word ‘bachelor’ would not be used any of its usual senses.

This may seem like a trivial point, and many have seen some of Quine’s criticisms of analyticity in the same light: if all it means that ‘all bachelors are unmarried’ could be false is that ‘bachelor’ could have meant something else, why is that philosophically interesting? But this is the wrong way to view the matter: that ‘bachelor’ could have had a different meaning than it in fact does is not the same thing as the fact that a specific occurrence of ‘bachelor’ does not interpret itself. The first point is counterfactual, the second epistemic.

This all boils down to *interpretation* being a theoretical matter, and *language* as much subject to theoretical uncertainty as any other field of knowledge. To settle for an interpretation of a theory A is, at least partly, to adopt a theory B of what claims in the first theory are synonymous (i.e. have the same meaning), and of which claims it is that follow from which. As we have defined it, it is obvious that the meaning set Σ_M has to be the subject matter of a theory M, with consequence operator $C_M(\Gamma)$ defined from the relation $\vDash_M$, in order for meaning to be something that can carry the burden of analytic consequence.

We call two theories A and B *equivalent* iff there are functions (let us call them *translations*) $\varphi\colon \Sigma_A \mapsto \Sigma_B$ and $\psi\colon \Sigma_B \mapsto \Sigma_A$ such that the following conditions hold:

(*i*) $\varphi[C_A(\Gamma)] = C_B(\varphi[\Gamma])$ for all $\Gamma \subseteq \Sigma_A$, and $\psi[C_B(\Delta)] = C_A(\psi[\Delta])$ for all $\Delta \subseteq \Sigma_B$.

(*ii*) $\psi(\varphi(p)) \equiv_A p$ for all $p \in \Sigma_A$, and $\varphi(\psi(q)) \equiv_B q$ for all $q \in \Sigma_B$.

Condition (*i*) says that it does not matter if we apply φ or the consequence operator first when we translate from A to B (or, in mathematical terms, that φ commutes with consequence), and likewise for translating from B to A. (*ii*) guarantees that if we translate back and forth using φ and ψ, we will always end up with something equivalent to that which we started with. Together, they formalize the claim that the structures of A and B are essentially the same. Call A and B *isomorphic* iff there is a bijection $\iota\colon \Sigma_A \mapsto \Sigma_B$ such that $\iota(C_A(\Gamma)) = C_B(\iota[\Gamma])$, for all $\Gamma \subseteq \Sigma_A$ (equivalently, this condition may be written using consequence relations, as $\Gamma \vDash_A p$ *iff* $\iota[\Gamma] \vDash_B \iota(p)$, for all $p \in \Sigma_A$). We can then show that A and B are equivalent iff their reducts, $\rho(A)$ and $\rho(B)$, are isomorphic.

Theorem: For any theories A and B, A and B are equivalent iff $\rho(A)$ and $\rho(B)$ are isomorphic.

Proof: (to appear).

It follows as an easy corollary that a theory always is equivalent to its reduct. However, having the same structure may not be sufficient for claims in one theory to be interchangeable with claims in another, since they may be about different things. The claims are interchangeable only if the truth of one guarantees the truth of the other, i.e. if they are equivalent. But since A and B do not have to have the same subject matter sets, this equivalence must be something that holds in a theory that includes A and B. We say that a theory C *licenses* the equivalence between A and B iff $A \sqsubseteq C$, $B \sqsubseteq C$, A and B are equivalent under the translations φ and ψ, and $\varphi(p) \equiv_C p$ and $\psi(q) \equiv_C q$, for all $p \in \Sigma_A$ and $q \in \Sigma_B$.

That C licenses the equivalence of A and B can thus be interpreted as a claim that, as far as C is concerned, the theories A and B are fully intertranslatable. In that case, we also say that A and B are C-equivalent. Using this terminology, we can show how adopting an interpretation of a theory A amounts to adopting another theory. The new theory adopted as interpretation is not, however, a theory on the same level as A, but rather one that frames it, as the following theorem shows:

Theorem: if μ_A: $\Sigma_A \mapsto \Sigma_M$ is an interpretation of A, then there are theories B and C, such $A \preccurlyeq B$ and B and M are C-equivalent.

Proof: (to appear).

An intuitive reading of this theorem is as follows: if μ_A gives the interpretation of A in the meaning theory M, then there is some theory B such that A is a theory in B, B is equivalent to M and some theory C licenses this equivalence. Thus, M is equivalent to some framework for A, and holding μ_A to be a *correct* interpretation can be interpreted as adopting as a framework some theory that licenses this equivalence.

Why do we usually not notice the theoretical aspect of interpretation? It may be because we, as philosophers, usually *presuppose* a fixed disambiguated language, and a fixed meaning function. This means, essentially, that we keep the interpretation theory fixed, i.e. regard it as necessary. But it is still a *theoretical* matter. We have no duty to fix the interpretation, since we always could revise our theory of it. In many cases a radical revision may be extremely improbable – one might even say, so improbable that it is *practically impossible*, and therefore we are free to disregard it when commonly reading texts. But this is exactly the strategy of fixing part of the theory in a context of inquiry: in this case, the theory of interpretation. That we do not have to fix it thusly should be obvious to any student of literature, where the interpretation itself is the primary question, and the question of whether something actually is true or false secondary.

The standard way of avoiding problems like these for a philosopher involves taking consequence bearers to be propositions. They supposedly wear their interpretation 'on their sleeve'. But why is this? I would say that the considerations of varying interpretations are not applicable to propositions precisely because we have already applied them. Taking propositions to be consequence bearers is a signal that we do not allow interpretation or context to vary, i.e. that the only thing we need to add to obtain a unique reference is what world is actual. But is there anything else to it than this? The proposition as such seems to be nothing but a posit made to ensure this kind of fixity. A similar function may be played by prefixing *that* to a statement.

This is why it is not *absolutely* necessary that 2 + 2 = 4: the 'that' signals that we are to hold the interpretation of '2 + 2 = 4' fixed, and this makes '2 + 2 = 4' necessary (i.e. part of the framework), but nothing makes this necessity an absolute matter. By their appearance in this text, we probably should infer that the symbols in question should be interpreted so as to fulfill the Peano axioms, although we should not jump to the conclusion that they thereby have a *unique* interpretation: we have no ground for taking '2' to stand for $\{\emptyset,\{\emptyset\}\}$ rather than $\{\{\emptyset\}\}$ or the set of two-membered sets.

Indeed, we have a perfectly well-known form of necessity that classifies '2 + 2 = 4' as contingent: the logical kind. Since neither '+', '2' nor '4' are logical signs, the Tarskian theory of logical truth requires us to vary their interpretation. Another way of putting the matter is that for logical necessity, the notion of model includes interpretations of the non-logical constants. Bolzano's original version goes even further, and explicitly requires us to specify which words are taken as fixed and which are variable (cf. his *Wissenschaftslehre*). In a sense, his may be seen as a precursor to the current theory, although we have replaced his systems of propositions and ideas with our theories.

Not even outright stipulation gives us any such thing as absolute necessity. If I were to begin this paper with a section where I state that, with ‘bachelor’, I henceforth will mean ‘Bach-impersonator’, that by itself *creates* a context in which only frameworks where ‘bachelor’ and ‘Bach-impersonator’ are equivalent are acceptable. ‘A bachelor is a kind of impersonator’ would then be a necessary truth, but only relative to the context created. The ‘absolutely necessary’ never appears on the scene.

4. Metaphysical aspects of necessity

Although the notion of necessity described in the preceding section may be antirealist in spirit, nothing precludes necessity from having a real ontological basis. One example of this may be causality. Humean worries aside, we seem to have rather good reason to believe that there are causal relations in nature (at the very least, regularities may be quite real). Since at least Mill’s *System of Logic*, we have studied empirical methods of finding them.

First of all, causal necessity as such is mostly a philosophical term of art. Parts of it are captured in notions of natural laws, but collecting all of these to causal or natural necessity is something that primarily philosophers do. A scientist does not generally regard *all* interactions, of all sciences, as fixed, but only a subset of them. Likewise, in everyday situations, we commonly hold much more than causal Yet, the philosopher’s position is useful when we make more general claims.

For simplicity, assume that causal relations hold between facts. We write ‘the facts $f_1, \ldots, f_n$ cause fact f_0’, i.e. that $f_1, \ldots, f_n$ jointly cause f_0, as $\{ f_1, \ldots, f_n\} \vDash_c f_0$. If causality is an objective feature of the world, then it is an objective truth that these relations hold. Where convention enters is when we take causal relations as justificatory for making inferences, i.e. when we regard a causal link as sufficient for consequence. Let A be our current framework, and let Φ be a function from claims in A to facts, such that Φ takes elements of Σ_A to the facts these elements express (for now, let us assume that each claim in A expresses one and only one fact). The theory A then validates causal relations iff $C_A(\Gamma) \supseteq \{p \in \Sigma_A \mid \Phi[\Gamma] \vDash_c \Phi(p)\}$, i.e. if the consequences we may draw from a set Γ of claims contain those claims whose truth follow from the truth of Γ causally. If C_A is the *smallest* consequence operator for which this holds (for instance, if we can replace the set containment symbol in the expression above with an identity), we may call A a theory of *causal necessity*.

An advantage with this tying of causal necessitation to theoretical inference is that if we weaken our notion of a theory somewhat, we have a powerful tool to handle causation with *ceteris paribus* conditions. For instance, ingesting (enough) cyanide causes a lethal poisoning only under certain conditions, such as that the ingester does not take a suitable antidote. We may of course then say that it is not the ingestion of cyanide that causes the poisoning, but that the ingestion in combination with the absence of antidote ingestion does. The problems with this are of course familiar: for one thing, we may have severe metaphysical misgivings about an absence being an existing entity, and let alone being able to cause anything. On the other hand, even if we do allow that, there may be an unspecified amount of other things that may hinder the working of either the cyanide or its antidote, or any other relevant factors, so that an exhaustive analysis may be quite inaccessible.

What we need here is a rendering of 'p causes q' such that if p is true, then *in normal cases*, q is true as well. This is obtainable by widening the theory notion and replacing the monotonicity condition on a theory's consequence operator by some weakening thereof, such as Makinson's *cumulative monotonicity* (1988). The result would be a nonmonotonic theory, related to the usual notion of theory as nonmonotonic logics are to standard, monotonic ones. We will however not go into the details of this modification here, as it is peripheral to the main theme of this text.

Another type of necessity, often thought related to the causal one, is the *nomic* (or *natural* or *physical*) kind. The canonical form of law statements in the literature (cf. Armstrong 1983, etc.) is 'Everything that is F is G', where F and G are predicates. It is, however, supposed to differ in analysis from the straightforward rendering '$(\forall x)(Fx \rightarrow Gx)$' in that it carries some kind of modal force as well. As the details of what this force is based on varies with the metaphysicial theory we adopt, I will only sketch how it may be treated.

Assume, for simplicity, a realist theory of properties, such as Armstrong's (1978) and a theory of laws as being on higher-order relations between universals. Armstrong (1983) expresses the instantiation of such a relation by the universals F and G as '$N(F, G)$', and postulates that the existence of this relation entails the truth of $(\forall x)(Fx \rightarrow Gx)$, but is not entailed itself thereby. This entailment is questioned by van Fraassen (1987, 1989), who points out that it is far from clear *why* the instantiation of two universals by a certain relation should allow us to infer that a particular instantiates one of them if it instantiates the other. This is what he calls the 'inference problem' for a theory of natural laws.

Our theory allows us to give an answer to the inference problem: no relation, or state of affairs, *ever* allows us to infer Gx from Fx on its own. It is always we ourselves who allow or disallow that, by taking part in a context in which the existence of the state of affairs $N(F, G)$ is accepted as sufficient for us to draw the conclusion that all F's are G's. In contrast to the case of causality, however, we have much less developed means for how to identify the existence of such relations (i.e. mainly philosophical ones), so perhaps this type of necessity is less practically useful to incorporate in a framework.

A presumed third kind of necessity is grounded in an altogether different way of thinking about these matters. We have regarded necessity as largely conventional, and claims of necessity as not strictly true or false. But to the more metaphysically inclined, just about everything about this theory may seem wrong: isn't what is necessary something that is determined by what the world is like? Perhaps we *can* vary anything, or question everything, in our minds, but does it make sense to do so? Maybe we just land ourselves in falsity if we do not keep some things fixed. Necessity, as this brand of metaphysician sees it, is just not a matter of convention at all, but of how the world is (or even how *other worlds* are).

Now, there *is* something to this view, even in our theory: we have defined necessity in terms of consequence, and consequence is related to truth and falsity. Truth and falsity are determined by what the world is like. Thus, if all claims in the set Γ are true, and $p \in C(\Gamma)$, then p is true as well, and if p is false, then some claim in Γ is false too. But, this presupposes nothing more than a material implication, and so has no modal importance. It is simply the observation that if p is necessary, then p is true, or as it is more commonly known, the T axiom.

On the theory presented here, what 'necessarily, p' says, except just 'p is true', is that p should be held fixed. The metaphysical understanding, on the other hand, holds that p's necessity means that *p could not* have been false, in some absolute, fundamental understanding of 'could'. I think this sort of reading gained its current popularity with Kripke's *Naming and Necessity*, and to some part with Putnam's work on natural kinds.[3] Kripke considers the Russellian attempt to give the sense of 'Aristotle' as 'the man who taught Alexander the Great'. The problem with this is that it becomes analytic that Aristotle taught Alexander the Great, and impossible that he did not. Even the 'cluster-of-descriptions'-theory leads to the similar problems: according to Kripke, it does not have to be analytic that

[3] Putnam may however be less to blame for it, especially since he has taken exception to Kripkean ideas later on (cf. his paper 'Is Water Necessarily H_2O?').

'Aristotle' satisfies most of the properties commonly attributed to him, since his life *could* have taken a completely different course.

How would *we* interpret the claim that it is possible that Aristotle did not do any of the things he is commonly credited with? Simply as an indication that we do not intend the truth of these statements to be held fixed. There is no 'metaphysical' punch to this, not even a reference to some extra-theoretic entity such as Aristotle. Certainly, Kripke is entitled to make such a stipulation, and it may very well be one that captures some philosophers' intuitions about how they commonly use the word 'Aristotle'. But this does not show anything more than that there sometimes may be reason to fix the reference of 'Aristotle' in a Kripkean rather than a Russellian manner, not that it actually *is* fixed in that way.

I think that one of Kripke's fundamental mistakes here is the presupposition that 'Aristotle' has a univocal meaning. If by 'Aristotle' we mean 'the man who taught Alexander the Great', then there is nothing strange about that being analytic: we have simply fixed the meaning that way. Most of the time, however, the meanings of words we use are not completely fixed at all. Kripke's theory, which fixes words' meanings causally, may be one way of specifying what we mean when uttering a proper name, but that should not make us believe that it gives us *the* meaning of 'Aristotle'. The same holds for Putnam's famous claim that 'water = H_2O' is necessary. If we by 'water' mean 'H_2O' (or rather 'the chemical compound with structure H_2O'), then of course that is necessary. If we mean 'the fluid stuff found in our lakes and oceans', then it is not. Most of the time when we use the word, it does not matter how we interpret it, as long as the interpretation lies within some reasonable bounds.

I believe we can explain both why Kripke's claims of necessity (for instance, that a person's origin is essential to that person) and Putnam's (that necessarily water = H_2O) have seemed reasonable to many philosophers: they correspond to *acceptable* ways of specifying the use of certain words. The principle of charity then makes us assume that what Kripke means by 'Nixon' is 'the person with that very specific origin', since that specific interpretation makes his claim true.

The discussion in this section should also bring out a central difference between this so-called metaphysical necessity and the causal or mathematical necessities. All three are necessities only because of convention, or relative to a context. The latter, though, are grounded in theories that actually are used as frameworks in scientific inquiry, and therefore worth investigating and formalizing from the perspective of the theory of science. Furthermore, if causal relations are real, then the fact that something is taken to be causally necessary even reflects a physical commitment. Metaphysical necessity, on the other hand, is a product of intuitions: Kripke is quite clear on this. But, as we have already remarked, intuitions are notoriously subjective; they are formed by our upbringing, our education and other particularities of every individual's life. The study of metaphysical necessity thus becomes nothing more than the study of a certain class of intuitions compatible with those Kripke had. Although all necessity may be conventional, only certain kinds correspond to conventions worthy of study, and only a few may be grounded in objective features of the world.

5. Conclusion

We have sketched a theory of necessity, and applied it to problems in logic, metaphysics, and the philosophy of science. I do not claim it to be a completely new theory; it has the philosophies of Carnap and Quine to thank for too much for that. I do believe, however, that it is worth defending, especially since metaphysically thicker notions of necessity have become more and more accepted.

With Carnap, we have held that necessity always is relative to our framework (or in Carnap's words, our language), and the choice of framework is a matter of convention. As there is no fundamental framework, there is no such thing as absolute necessity. With Quine, we have repudiated the idea of an essential distinction between framework and theory. We still, however, hold that a theory may be *used* as a framework, and then be *held fixed*, and we have also argued that this is how things must be if we are to engage in inquiry at all. What is *necessary* becomes the same as that which we currently hold fixed.

So although we started out from Carnap and Quine, we have arrived at a theory quite different from either of theirs. Carnap did see claims of necessity as true or false, if only true or false because of convention. Quine's challenges to the idea that truth can be obtained by convention made him reject as meaningless the entire idea of a necessary/contingent distinction, except possibly as asserting the existence of a regularity (cf. Quine 1963). We, on the other hand, have found another use for the necessary, as expressing what is to be held constant and what is to be open to question.

This non-cognitivistic character of necessity resembles that which Blackburn (1984, ch. 6 and 1987) gives it. According to Blackburn, modalizing consists of laying down norms for what beliefs may be revised, as well as of emotivistic expressions of obviousness or inability to imagine certain states of affairs. However, neither of these functions explain why we *have to* modalize. Although the ‘cautious man’, who regards everything as contingent, may have a more complicated life than the scientist in general, still he is a possibility, and even worthy of some admiration. According to my theory, he is not. Or at least not if he is to avoid being a radical (and probably inconsistent) skeptic. Since every justification consists of giving some kind of argument, we can give no justifications at all if we are not allowed to hold some things fixed, at least for the duration of that argument. This applies to the skeptic trying to argue against our ability to have knowledge, as well as to his opponent.

Is our version of necessity therefore a species of epistemic necessity? In the narrow sense of the term, no. What is necessary according to us is not given by what anyone believes or knows. On the other hand, there *is* a connection to the epistemic concept of justification, since what is necessary is coextensive with what does not need to be justified. We can bring out this relationship more clearly by expressing a notion of knowledge directly in our theory.

Let us use *evidence* as a blanket term for anything that we can have that supports the truth of a belief, and let E_a be the set of all evidence person a has. This may include having been involved in a certain causal process, having a mathematical proof, having a ‘clear and distinct idea’, or any of several other things epistemologists have licensed as a basis for knowledge. Assume that C_c is the consequence relation of the theory used as framework in the present context. We can then frame a notion of knowledge as follows:

a knows that p $\stackrel{\text{def}}{=}$ (i) a believes that p, and

(ii) $p \in C_c(E_a)$.

The notion carrying all the theoretical weight here is of course the consequence operation C_c, specific to the context. In many cases, it will express what follows by *practical* necessity: the probability that it does not hold in the current context is so low that we may disregard it. In contexts of epistemological inquiry, as that invoked by a skeptic, we probably have to use purely logical necessity, perhaps with the addition of some basic foundational layer such as ‘direct experience’ if we are believers in such things.

For practical necessity, as possibly also for causal necessity, the axioms given for the consequence operation above will have to be weakened: again monotonicity is unlikely to hold in unmodified form. In common sense reasoning, for instance, adding a premise may invalidate inference that is valid without that premise, since we usually have a large amounts of contextually determined presuppositions in a conversation, that may be falsified by new assumptions.

The definition of knowledge is a version of Lewis's in *Elusive Knowledge*, but framed in terms of consequence instead of excluded possibilities. It shares that account's context-sensitivity, which lets it explain why everyday ascriptions of knowledge may be perfectly true, but skepticism still inevitable when we challenge the skeptic on his own playing field. Our presentation, however, shows connections between justification, knowledge and necessity. What is necessary is that which we do not need to justify, and what we do not need to justify is that which we do not require any further evidence to infer. The present version of the knowledge definition also lets us save one more pretheoretic intuition about knowledge: if *a* knows that *p*, then *p* cannot possibly be false. Thus, we are offered a way to avoid both the *Scylla* of fallibilism and the *Charybdis* of skepticism.

REFERENCES

(to be filled out).