

Unification: Not Just a Thing of Beauty

Ioannis VOTSIS

Received: 27/08/14
Final Version: 06/12/14

BIBLID 0495-4548(2015)30:1p.97-114
DOI: 10.1387/theoria.12695

ABSTRACT: There is a strong tendency in science to opt for simpler and more unified hypotheses. A view that has often been voiced is that such qualities, though aesthetically pleasing or beautiful, are at best pragmatic considerations in matters of choosing between rival hypotheses. This essay offers a novel conception and an associated measure of unification, both of which are manifestly more than just pragmatic considerations. The discussion commences with a brief survey of some failed attempts to conceptualise unification. It then proceeds to an analysis of the notions of confirmational connectedness and disconnectedness, as these are essential ingredients in the proposed conception of unification and its associated measure. Roughly speaking, the notions attempt to capture the way support flows or fails to flow between the content parts of a hypothesis. Equally roughly, the more the content of a hypothesis is confirmationally connected, i.e. support flows between its content parts, the more that content is unified. Since the confirmational connectedness of two content parts is determined by purely objective matters of fact, the proposed notion and measure of unification are themselves strictly objective, i.e. not merely pragmatic. The essay concludes with a discussion of how the proposed measure handles several examples but also how it relates to the debate over measures of coherence.

Keywords: Unification, confirmation theory, objective probabilities, ad hoc, relevant deduction, holism and coherence.

RESUMEN: En la ciencia hay una marcada tendencia a preferir las hipótesis más simples y unificadas. Una opinión mantenida a menudo es que tales cualidades, aun siendo atractivas o estéticamente satisfactorias, constituyen consideraciones pragmáticas, a lo sumo, en el asunto de la elección entre teorías rivales. Este ensayo ofrece una concepción novedosa de unificación y una medida asociada a ella, ambas claramente algo más que meras consideraciones pragmáticas. La discusión comienza con un breve repaso de algunos intentos fallidos de conceptualizar la unificación. Después se analizan las nociones de conexión y desconexión confirmacional, componentes esenciales en la noción de unificación y la medida asociada que aquí se proponen. Dicho brevemente, esas nociones pretenden captar el modo en que el apoyo discurre o no entre las partes del contenido de las hipótesis. Simplificando, cuanto más conectado confirmacionalmente está el contenido de una hipótesis, más unificado está. Dado que la conectividad confirmacional de dos partes del contenido está determinada por cuestiones de hecho objetivas, la noción y la medida que propongo son también estrictamente objetivas, esto es, su valor no es meramente pragmático. El ensayo concluye con una discusión sobre cómo la medida propuesta afronta diversos ejemplos y sobre su relación con el debate sobre las diferentes medidas de coherencia.

Palabras clave: Unificación, teoría de la confirmación, probabilidades objetivas, ad hoc, deducción, holismo y coherencia

1. Introduction

We often hear that simplicity and unification, though aesthetically pleasing or beautiful qualities for a hypothesis to possess, are at best pragmatic considerations when deciding between rival hypotheses. This essay proposes a novel conception and an associated measure of unification. In short, the degree to which the content of a hypothesis is unified gets re-



flected in the extent to which, roughly speaking, its content parts are confirmationally connected. Since such connections are settled by purely objective matters of fact, the proposed conception and measure of unification turn out to be objective, and hence not merely pragmatic, considerations we can appeal to when faced with the prospect of selecting between competing hypotheses. The structure of the essay is as follows: Section 2 contains a brief survey of some prominent but failed attempts to conceptualise unification. Section 3 offers an informal introduction to the notions of confirmational connectedness and disconnectedness. To give a rough idea, the two notions attempt to capture the way support flows or fails to flow between the content parts of a hypothesis or, equivalently, between the contents of two or more hypotheses. Confirmational connectedness and disconnectedness are given formal expositions in Section 4. Section 5 then puts forth a measure for the so-called ‘monstrousness’ of a hypothesis. This sets up the proposed measure of unification, explicated in Section 6, which is roughly the inverse of a variant of the monstrousness measure. A discussion of how the proposed measure handles various cases ensues. Section 7 delves into a hefty difficulty facing the current proposal, namely practicability. Finally, the last section explores in passing the relations between the proposed measure of unification and measures of coherence.

2. *Unification: A Brief Overview*

Attempts to devise a satisfactory conception of unification abound. One of the earliest to have left an indelible mark in the literature is Friedman (1974) where it is argued that an intimate connection exists between explanation, understanding and unification. To be precise, Friedman argues that an adequate theory of explanation must show how explanation generates understanding. In his view, understanding is generated when we trim down the number of independently acceptable law-like assumptions that feature as *explanantia* in the derivation of an *explanandum*. The lower that number the more unified an explanation. Friedman’s account is in great part motivated by a desire to avoid a problem with the deductive-nomological (DN) account of explanation identified in Hempel and Oppenheim (1948). According to the problem of trivial explanations, deriving an *explanandum* from one or more laws, as the DN account dictates, is not sufficient to turn those premises into a genuine explanation since we can derive any statement, and hence any law, from itself. Friedman seeks to avoid this problem by limiting the derivations that yield genuine explanations to those that unify phenomena. To give one of his examples, the kinetic theory of gases unifies (and therefore explains) phenomena that previously fell under the distinct explanatory spheres of the Boyle-Charles law, Graham’s law and the assumption that gases have certain specific heat capacities because we now need one, instead of three, independently acceptable law-like assumption(s).

Though highly influential, Friedman’s account was beleaguered from the onset. In a well-known contribution, Kitcher (1976) argues that Friedman is only successful in ruling out trivial explanations at the expense of also ruling out some genuine ones. That is, there are genuine scientific explanations that do not satisfy all the conditions of Friedman’s account. One such condition is that only *K*-atomic statements explain. A statement *S* is *K*-atomic if it is not logically equivalent to the conjunction of $n \geq 2$ law-like statements that are acceptable independently of *S*. And a statement is independently acceptable to another statement when evidence adequate for accepting the former is inadequate for accept-

ing the latter. Kitcher cites the (genuine) explanation of the law of adiabatic expansion of ideal gases by the conjunction of the first law of thermodynamics and the Boyle-Charles law as a counter-example. Each conjunct can be accepted independently of the conjunction and thus the conjunction is not *K*-atomic. Another (related) objection is due to Salmon (1998). He points out that no fundamental law statements seem to qualify as *K*-atomic and utilises Newton's law of universal gravitation as an example. This law can be split up into three law-like statements: the first ranging over pairs of masses of astronomical dimensions, the second ranging over pairs of masses where one member is astronomical in its dimensions and the other non-astronomical (i.e. smaller) and the third ranging over pairs of masses where both members are of non-astronomical dimensions. Presumably, each such statement is independently acceptable. Moreover, as Salmon is quick to highlight, this treatment seems generalisable: “[i]t seems possible to partition virtually any universal statement into two or more independently acceptable generalizations” (p. 70).

Kitcher is not only a detractor of some of Friedman's ideas but also an exponent of others. This is no more clearly evident than in Kitcher (1989). For a start, Kitcher there agrees with Friedman that explanation and unification are intimately connected. Moreover, he sees eye to eye with him on the vital role derivation plays in providing successful explanations. What is a successful explanation for Kitcher? Roughly speaking, it is a derivation drawn from what he calls the 'explanatory store' $E(K)$, where K is a set of statements underwritten by the scientific community at a particular time. A little more precisely, $E(K)$ is “the set of derivations that best systematizes K ... the criterion for systematization [being] unification. $E(K)$, then, is the set of derivations that best unifies K ” (p. 431). How are we meant to understand unification here? To unify, Kitcher appears to be telling us, is to “minimiz[e] the number of patterns of derivation employed and maximiz[e] the number of conclusions generated” (p. 432).

Schurz and Lambert (1994) object that by conceiving of unification as a form of 'deductive systematization', Kitcher (but also Friedman) misses out on “interesting cases [in] science” (p. 73). What they have in mind are cases where the systematization is, broadly construed, inductive. Schurz and Lambert also indicate that Kitcher's solution to the problem of 'spurious explanations' is not restrictive enough. To put things in perspective, suppose that a spurious explanation is, simplistically speaking, one that *merely* postulates a deductive relation $H \rightarrow E$ between H a given hypothesis and E some empirically true statement. Kitcher attempts to rule out such explanations by barring argument patterns where what is being derived is substitutable by *any* statement whatsoever. With this approach, for example, he manages to eliminate argument patterns of the form “God wants P , and whatever God wants is the case, therefore: P ” (1981, pp. 527f). His solution is not, however, general enough as it does not apply to cases where the content of P is internally restricted by the spurious theory in question, e.g. when our theory about God specifies that P concerns *only* weather patterns (Schurz and Lambert 1994, 115, f22).

The final approach to unification I will be mentioning can be found in the Schurz and Lambert essay. Following in Friedman's footsteps, the two authors uphold the view that the concepts of understanding, unification and explanation are closely related. For simplicity, we here restrict our attention to their account of the first two concepts. Loosely speaking, to understand a phenomenon is “to know how [a statement of that phenomenon] fits into one's background knowledge” or, as Schurz and Lambert prefer to call it, “cognitive corpus C ” (p. 66). The cognitive corpus C of a subject S consists of S 's state of knowledge and belief. More precisely, it consists of a pair $\langle K, I \rangle$, where K is a set of declarative sen-

tences known/believed by *S* and *I* is a set of inferences mastered by *S*. The key notion of fitting deserves closer inspection. According to Schurz and Lambert, a subject fits statements of phenomena into their cognitive corpus by inferentially connecting them with parts of that corpus such that the latter's unification increases. The inferential connections at stake are conceived of as "argument[s] ibs (in the broad sense)" (p. 71). How broad this category of arguments is meant to be is not immediately clear from the discussion but the general idea is that the said arguments are not merely deductive, as in the case of Friedman's and Kitcher's projects, but must also include some sufficiently strong inductive ones.¹ What exactly does it mean for there to be an increase in the unification of a given cognitive corpus? There's no quick and easy way to answer this question but, to give the reader a sketch, the idea is that on the basis of certain weight assignments and rules one can calculate the costs and gains of shifting from a cognitive corpus *C* not fitted with a statement *P* to a modified version *C** that is so fitted. If the overall result of this calculation is a positive/negative number then *C** is more/less unified than *C*. Otherwise, the two corpuses are equally unified. As an example, consider one of the addition rules provided (1994, pp. 79-80). The addition of a hypothesis to (what they identify as the base of) a cognitive corpus has a certain intrinsic cost since, intuitively, adding theoretical information to our cognitive corpus complicates its content. Having said this, a good hypothesis may at the same time bring with it an extrinsic gain if it allows for the systematization of correct or accurate data. The implication being that in such cases the extrinsic gain outweighs the cost and hence the resulting corpus is more unified than the original one.

The main disadvantage of this approach, as I see it, is that it is only as good as (a) the exact characterisation of the class of ibs arguments, (b) the weight assignments we make and (c) the rules for calculating costs and gains. So much is up for grabs that calculating the costs and gains of a shift from one corpus to another remains a strongly subjective matter. Sure, Schurz and Lambert provide some guidance regarding (a) and (b), but nowhere near enough to secure their account against relatively obvious objections. Take the aforementioned rule. Although we are told that the addition of a hypothesis that systematizes many data (and indeed correctly predicts new data) brings with it a net gain, it is unclear how much data needs to be systematized in order to tip the balance toward such a gain.

3. *Disconnectedness: A First Glance*

While they ultimately fail, the above accounts do at least get one fundamental thing right. By emphasising the role of the *acceptability* of law-like assumptions in the case of Friedman, of the *endorsement* of a set of statements *K* by the scientific community at a given time in the case of Kitcher, and of the systematization of *correct* or *accurate* data in the case of

¹ As Schurz and Lambert note: "In probabilistic terms, this means that the probability (degree of belief) of *Con* [the conclusion of such an argument] is increased by *Prem* [its premises] to a 'sufficiently' great value. This probabilistic characterization figures only as an intuitively necessary condition for being a correct argument ibs... Since it is doubtful that a sufficient condition for correct arguments ibs can be given in general, we will specify a *list* of the kinds of correct arguments ibs, including *deductive*, *approximative-deductive* and *inductive arguments*, which extensionally defines the correctness of an argument ibs" (p. 71) [original emphasis].

Schurz and Lambert, these accounts place a premium on the link between unification and confirmation.² The account to be unveiled in the sections that follow is in tune with this appraisal and indeed elevates the link with confirmation to the single most important ingredient in our quest to understand unification. According to this account, unification is to be understood as a measure of confirmational connectedness. But what is confirmational connectedness and its opposite confirmational disconnectedness? To understand these notions we need to take a rather long detour into another very important topic, namely the thorny topic of ad hoc hypotheses.

What makes a hypothesis ad hoc has been the subject of great controversy over the years. The account that follows does not aim to do justice to the notion of ad hoc-ness as it is used in the scientific literature. No single account can do that for the simple reason that the notion has not been used univocally (see Holton 1969, p. 178). A more manageable undertaking involves the provision of an account of some features that are common to ad hoc hypotheses. Such an undertaking can be found in Votsis (forthcoming). There the aim is to articulate one undesirable feature often associated with ad hoc hypotheses, namely what can best be described as the artificial or contrived nature of the content of those hypotheses. In what follows, I rehearse, and expand on, the main ideas behind this account.

Let us call this undesirable characteristic the ‘monstrousness’ of a hypothesis. It indicates the extent to which a hypothesis is compiled out of confirmationally disconnected parts, much like the monster in Mary Shelley’s *Frankenstein* is compiled out of miscellaneous ill-fitting parts. Equivalently, it indicates the extent to which two or more hypotheses are confirmationally disconnected. More technical details can be found in the next section. For now let us try to flesh out the ideas of monstrousness and confirmational disconnectedness with two toy examples, their main virtue being that they are fairly straightforward to understand.

In some cases, confirmational disconnectedness and therefore monstrousness follows from the fact that even though different parts of a single hypothesis are confirmed, well-confirmed or even true, the said parts are nonetheless confirmationally unrelated. That is, the support each receives does not affect the other parts. Consider an example adapted from Goodman (1983). Take the true proposition ‘8497 is not a prime number and the other side of the moon is not flat and Elizabeth the First was crowned on a Sunday’.³ Confirming any one of these conjuncts presumably has nothing to do with confirming any of the others. Thus, having confirmed, well-confirmed, or even true parts is not sufficient for confirmational connectedness. An unqualifiedly true hypothesis like the above may be highly confirmationally disconnected and hence monstrous.

² Schurz and Lambert are more explicit about this link than Friedman and Kitcher. They insist that “scientific unification must yield empirical confirmation” (p. 73). The rationale for this demand is given a few sentences earlier: “The appeal to reality in the search for unification prevents artificial unification; science can’t produce more unification than really exists in the world” (p. 73). This is an extremely important point that I also subscribe to!

³ The original version of the proposition is offered by Goodman as an example of a contrived hypothesis: “8497 is a prime number and the other side of the moon is flat and Elizabeth the First was crowned on a Tuesday” (p. 69). All three conjuncts in this version are in fact false.

In other cases confirmational disconnectedness and monstrousness follows from the fact that some parts of a hypothesis cannot be supported altogether. As an example, consider the false proposition ‘The moon’s escape velocity is 2.38km/s and the other side of the moon is flat’. The first conjunct enjoys quite a bit of support and is validated, among other things, by the successful return of the crews from the six manned moon-landing missions. Assume, for the sake of the argument, that this is genuine support. Assume moreover, as it seems to be the case, that the second conjunct cannot be supported. Since, in my view, the question of whether two hypotheses or hypothesis-parts are confirmationally disconnected depends entirely on facts about the universe, if such facts leave one hypothesis or hypothesis-part entirely without support, then that hypothesis or hypothesis-part could not be confirmationally connected to the other. That is to say, the two hypotheses or hypothesis-parts must be confirmationally disconnected. Note that this point holds also in cases where both hypotheses or hypothesis-parts are entirely devoid of support. It cannot be stressed enough that the determination of confirmational (dis-/)connectedness is not an a priori matter in the view being proposed but rather an empirical one that depends entirely on facts about the universe, namely those same facts that we cite as supporting a given hypothesis.

4. (Dis-/)connectedness: A Formal Look

Let us use ‘ $x \vdash_r y$ ’ to denote that y is a *relevant deductive consequence* of x – more on this notion below. Confirmational disconnectedness can be articulated thus:⁴

Any two content parts of a non-self-contradictory proposition Γ expressed as propositions A, B are *confirmationally disconnected* if, and only if, for all pairs of internally non-superfluous propositions α, β where $A \vdash_r \alpha$ and $B \vdash_r \beta$: (i) there is no true or partly true proposition γ such that $\alpha \vdash_r \gamma$ and $\beta \vdash_r \gamma$ and (ii) where $0 < P(\alpha), P(\beta) < 1, P(\alpha/\beta) = P(\alpha)$ and (iii) there is no atomic proposition δ such that $\alpha \wedge \beta \vdash_r \delta, \alpha \not\vdash_r \delta$ and $\beta \not\vdash_r \delta$.

And so connectedness can be articulated thus:

Any two content parts of a non-self-contradictory proposition Γ expressed as propositions A, B are *confirmationally connected* if, and only if, for some pair of internally non-superfluous propositions α, β where $A \vdash_r \alpha$ and $B \vdash_r \beta$: either (1) there is a true or partly true proposition γ such that $\alpha \vdash_r \gamma$ and $\beta \vdash_r \gamma$, or (2) where $0 < P(\alpha), P(\beta) < 1, P(\alpha/\beta) \neq P(\alpha)$ or (3) there is at least one atomic proposition δ such that $\alpha \wedge \beta \vdash_r \delta, \alpha \not\vdash_r \delta$ and $\beta \not\vdash_r \delta$.

In what follows, we restrict our attention to the notion of disconnectedness, though, *mutatis mutandis*, the same analysis can be given for the converse notion of connectedness. Let us consider each of the notions appearing in the *analysans*. First off, what is a content part? To fully make sense of the answer I am about to give requires comprehension of the concept of relevant deductive consequence. That concept is explicated a few paragraphs below so I ask

⁴ I first proposed a version of these notions in Votsis (2014).

the reader to be patient. For now, suffice it to say that a proposition c is a (non-trivial) content part of a proposition Γ if, and only if, $\Gamma \vdash_r c$ and $c \not\vdash_r \Gamma$. In more colloquial terms, c is a (non-trivial) consequence of Γ but not vice-versa and hence c has strictly less content than Γ .

Next up, consider the non-superfluous-ness clause. As the name suggests, this is incorporated into the definition to remove superfluous or redundant content and thus to reduce the complexity of the evaluation. Formally, a proposition ν is non-superfluous just in case if its content is expressible as a conjunction, there is no conjunct of that conjunction the removal of which yields a proposition τ such that ν and τ are logically equivalent.⁵ In other words, ν is non-superfluous if, and only if, none of its content parts are duplicated. The simplest case of such duplication is through iterated conjuncts, e.g. $\alpha \wedge \alpha$. Take proposition κ : $\alpha \wedge \alpha$. Removal of one conjunct yields another proposition κ' : α . Since κ and κ' are logically equivalent that makes κ superfluous. A point that needs clarifying is that the non-superfluous-ness clause is meant to apply to the content of α and of β but not to the joint content of those propositions, i.e. to $\alpha \wedge \beta$. That is, the clause prohibits content duplication only within each proposition, i.e. within α and within β , but not across these propositions, i.e. between α and β . That's why we insist on the characterisation '*internally* non-superfluous'. Thus, the contents of α and β are allowed to overlap. We permit this kind of duplication across a pair α, β because it allows us to detect one kind of confirmational connection.

Now consider clause (i). To get a better grip on this clause consider first a version where the notion of relevant deductive consequence is replaced with the notion of (normal) deductive consequence. That is, suppose that the clause demands that there is no proposition γ such that $\alpha \vdash \gamma$ and $\beta \vdash \gamma$. It should be clear that if there were such a proposition γ , α and β would have some common content and hence any confirmation of that content would confirm at least a content part of α and at least a content part of β . So, if there were such a proposition γ , we would have to accept that α and β and hence A and B have some sort of confirmational connection. That's why the first clause (of confirmational disconnectedness) demands that there is no such proposition γ .

Why employ a version of the clause that utilises the notion of relevant deductive consequence instead of (normal) deductive consequence? Because there is a fatal flaw with the latter. Whatever the content inherent in two propositions α, β , one can always deduce a proposition that is common to both, e.g. $\alpha \vee \beta$ or a tautology. Such trivial consequences are a well-known feature of classical logic. So, were we to stick to the normal deductive entailment formulation of clause (i) we would *guarantee* the existence of a proposition γ that meets the expressed condition. But that would in effect mean that all pairs of propositions α, β are *connected* thereby rendering the concept of disconnectedness unsatisfiable. Banning such trivial deductive consequences is thus of paramount importance. That's why the notion of relevant deductive consequence is called upon. Its purpose is to home in on the non-trivial content of a proposition by ruling out its trivial content as irrelevant.

We say that a conclusion or consequence of "a valid deduction is relevant iff no subformula of the conclusion is replaceable on some of its occurrences by any other formula salva validitate of the deduction" (Schurz 1991, p. 391). Thus, a conclusion or consequence of a valid deduction is irrelevant if, and only if, it contains a subformula whose substitution

⁵ In other words, if a proposition ν is non-superfluous then the resulting τ is logically weaker than ν . On a different note, it is worth pondering whether a stronger form of equivalence is required in the definition just given.

on some of its occurrences by any other formula maintains the validity of the deduction.⁶ Otherwise put, a deductive consequence is relevant if, and only if, it is not irrelevant. A few examples are in order. Suppose we deduce $\alpha \vee \beta$ from α . Is $\alpha \vee \beta$ a relevant consequence of α ? No! For $\alpha \vee x$ where x stands for any other formula whatsoever still classically follows from α . The same holds with a bunch of other classically derivable consequences of propositions that are widely deemed to be undesirable in the context of confirmation including: (D1) $p \rightarrow q \vdash (p \wedge r) \rightarrow q$, (D2) $p \vdash q \rightarrow p$ and (D3) $(p \wedge q) \rightarrow r \vdash (p \rightarrow r) \vee (q \rightarrow r)$ where the replaceable formulae are in bold. Note that derivations like D1 are undesirable because they help underwrite cases that involve the tacking (by conjunction) paradox – in short, that tacking, i.e. affixing, any irrelevant conjunct to a hypothesis confirmed by some evidence absurdly leads to that conjunct’s confirmation. And derivations like D2 are undesirable because they help underwrite cases that involve spurious explanations. Derivations like D3 are undesirable for reasons that will become apparent in our discussion of clause (iii) below. On the positive side of the above definition, note that various classically derivable consequences that are deemed desirable in the context of confirmation qualify as relevant consequences. These include instances of crucial derivation rules such as: (R1) $p, p \rightarrow q \vdash q$, (R2) $\sim q, p \rightarrow q \vdash \sim p$ and (R3) $p \vee q, \sim p \vdash q$.

Now consider clause (ii). Let us begin with the intended interpretation of probabilities. The probabilities here are meant to be objective. That is, probability statements are meant to indicate true relative frequencies and/or true propensities of things happening like events, states-of-affairs or property instantiations. An objective interpretation captures the earlier intuition that the confirmational disconnectedness or connectedness of the content of a hypothesis is determined by facts about the world. It is not a subjective matter. At this stage, we put aside worries about our epistemic access to objective probabilities, though we return to this issue in Section 7.

The notion of probabilistic independence employed in clause (ii) is meant to help us scoop up additional ways through which two propositions fail to be confirmationally connected. The first kind of confirmational connection we encountered was in the form of a (relevant) deductive link. Sometimes such a connection may be absent, as required by clause (i), but another one may be present in the form of a probabilistic link. To ensure that there is no probabilistic link between any pair α, β we require that such pairs are probabilistically independent. That means that the probability of the one is unrelated to the probability of the other. And, if that is the case, then we cannot possibly argue that the two propositions are confirmationally connected.

What I just said is not quite right. To secure confirmational disconnectedness between a pair of propositions A, B at the probabilistic level we must inspect not only the total content of each proposition but also the content of their parts. After all, two propositions may have consequences that are probabilistically dependent even though the propositions themselves are probabilistically independent. That’s why we rummage through all the possible pairs of relevant deductive consequences of each in our evaluations.⁷

⁶ This is not an appeal to relevant logic which axiomatises the notion of relevance. Rather, it is a restriction of the kinds of inferences that are allowed within the framework of classical deductive logic. The special turnstile “ \vdash_r ” is no more than a short-hand way of referring to the latter.

⁷ The condition that $P(\alpha), P(\beta) > 0$ reflects the common practice of assigning a non-zero value to false-propositions to avoid undefinability. The condition that $P(\alpha), P(\beta) < 1$ is not strictly necessary. Take a

Finally consider clause (iii). This clause is meant to scoop up residual ways through which two propositions fail to be confirmationally connected. Two propositions A , B may satisfy clauses (i) and (ii) and yet still be confirmationally connected. What form do these connections take? As a first approximation, the kind of connection I have in mind is through a proposition δ that is a relevant deductive consequence of $\alpha \wedge \beta$ but is not a relevant deductive consequence of either α or β when each is considered on its own. The first thing to note is that this sort of situation would not even be possible if we did not filter out irrelevant deductive consequences. That's because classical deductive logic allows the derivation of $(p \rightarrow r) \vee (q \rightarrow r)$ from $(p \wedge q) \rightarrow r$. But it seems obvious that there are states of affairs where the presence of two conditions is sufficient to bring about an event even though the presence of each condition on its own is insufficient to do so. If we want our logic to capture such states of affairs, derivations like the above must not be allowed. Thankfully, the restriction to relevant deductive consequences rules them out as irrelevant - see D3.

A second thing to note is that if clause (iii) required only that there is no δ that any α , β jointly (but not individually) and relevantly entail then the concept of disconnectedness would once again be rendered unsatisfiable.⁸ In other words, there would always be such a δ regardless of the content α , β are allowed by the clause to possess. One such δ , for example, is $\alpha \wedge \beta$. To rule out such guaranteed joint consequences I first considered turning to the notion of *relevant deductive content elements*, originally proposed by Schurz (1991). I now believe that this notion, though instructive in its general spirit, is inadequate for the task at hand. Here's why. First off, let us consider what's packed into the notion. A proposition ψ is a relevant deductive content element of a proposition φ if, and only if, ψ is a relevant consequence of φ and ψ cannot be decomposed into a logically equivalent conjunction of propositions, each of which has less content than ψ . Intuitively speaking, a relevant deductive content element is a proposition whose content cannot be decomposed into smaller content parts. How does this help rule out guaranteed joint consequences like $\alpha \wedge \beta$? Well, a proposition δ that is a relevant deductive content element of $\alpha \wedge \beta$ could not be logically equivalent to $\alpha \wedge \beta$ since that would mean that the content of both α , β is smaller compared to that of δ . Note, moreover, that all conjunctive relevant consequences of $\alpha \wedge \beta$ that have less content than $\alpha \wedge \beta$ (and that are not relevant consequences of α or β on their own) are also ruled out for the same reasons, as they should be since they are also guaranteed to exist.

Despite appearances to the contrary, the notion of relevant deductive content elements is not fully fit for the job. This needs some explaining. A supposition sometimes made in the metaphysics literature (e.g. Armstrong 1989), one that I also share, is that the world has atomic states. The universe could not instantiate any part of such a state without instantiating the whole state for the simple reason that, provided this supposition holds, the state has no smaller parts. The notion of a content element might be thought of as the conceptual analogue of an atomic state. But immediately there is a problem. Conceivability is a radically unrestrained tool. Take a proposition ψ_i describing an atomic

case where $P(\alpha) = 1$ and/or $P(\beta) = 1$. Such an assignment guarantees that $P(\alpha/\beta) = P(\alpha)$ regardless of the specific content of either proposition. It would thus seem to offhandedly rule out any connection between them. It is unnecessary to demand that $P(\alpha)$, $P(\beta) < 1$ because any connection between such propositions, i.e. where $P(\alpha) = 1$ and/or $P(\beta) = 1$, would at any case be deductive and hence capturable by the other two clauses in the definition of confirmational connectedness.

⁸ Note that where $\alpha \wedge \beta \vdash_r \delta$, $\alpha \not\vdash_r \delta$ and $\beta \not\vdash_r \delta$, α and β cannot have the same content.

state. Nothing prevents us from conceiving a decomposition of ψ_i into a logically equivalent conjunctive proposition where each conjunct has less content than ψ_i . The claim is not that conceptually-driven decomposition is limitless. Rather the claim is that if there is such a limit it need not, and in fact it is not likely, to coincide with the limit of a decomposition which is solely determined by the amount of content required to represent an atomic state. That's why appeal is made to the notion of atomic propositions instead of to the notion of content elements. The general motivation behind this move is to let atomic states dictate the amount of content that atomic propositions possess. A proposition ψ is atomic if, and only if, ψ is non-superfluous and truthfully represents all and only the content of an atomic state. Such a proposition does not permit any conjunctively atomic non-superfluous decompositions, i.e. a decomposition of ψ into n propositions $\omega_1, \dots, \omega_n$ where $n \geq 2$ such that $\psi \vdash_r \omega_1 \wedge \dots \wedge \omega_n$ and each ω_i represents all and only the content of a distinct atomic state. Thus, the ability to rule out guaranteed joint consequences is preserved.

Allow me to bring this section to a close by pointing out that the foregoing discussion links up nicely with one of the burning issues in confirmation theory, namely how support gets transmitted. On the one hand, it might be argued that simply being a content part of a hypothesis or of a set of hypotheses means that evidence for the truth of one part is automatically transmitted to all the other parts. Such a globally holistic account surely can't be right as it would lead us straight into the clutches of the tacking paradox. On the other hand, it might be argued that only direct support for a content part is genuine. This is denigratingly described as 'content-cutting' by Ken Gemes (1998). Such a globally anti-holistic account counters the very idea of induction since the transmission of support is strictly prohibited. Take the general hypothesis 'Objects with mass bend space-time around them'. On such a globally anti-holistic view, any evidence of massive objects bending space-time around them would be strictly limited in its support to those parts of the general hypothesis that concern the given objects. This means, among other things, that no such evidence would be capable of supporting the hypothesis that 'The next observed object with mass will bend space-time around it'. Indeed, the contradictory hypothesis, i.e. 'It is *not* the case that the next observed object with mass will bend space-time around it', would according to this view be deemed to be confirmationally on the same footing. Surely such judgments are incorrect.

If the truth cannot be found at either of these two extremes, it must lie somewhere in between. That's where my account slots in. It asserts that evidence for one content part will not indiscriminately spread to another content part but that spreading sometimes takes place. Whether or not it does depends on the confirmational connections between content parts in the terms specified by the aforementioned clauses. And that, as I have already repeatedly stressed, cannot be decided in an a priori fashion but is a matter of whether or not corresponding connections exist in the world.⁹

5. *Monstrousness: A Measure*

A non-self-contradictory molecular proposition Δ is *monstrous* if, and only if, some of its content parts are confirmationally disconnected. What does molecular mean? It means

⁹ As things stand in my account, such connections need not be causal.

that its content can be decomposed into a conjunction whose conjuncts include at least two distinct atomic propositions. The reason for this requirement is that if Δ were an atomic or even a sub-atomic proposition it would not have enough content to allow us to make the sort of judgments about it we'd like to make, e.g. whether there is an atomic proposition δ such that $\alpha \wedge \beta \vdash_r \delta$, $\alpha \not\vdash_r \delta$ and $\beta \not\vdash_r \delta$ where, of course, α and β are content parts of Δ . Note also that although the *analysandum* makes reference to an individual proposition, e.g. a hypothesis, we can still use it to pass judgment on the monstrosity of two or more propositions, e.g. two or more hypotheses, so long as these are mutually consistent. To achieve this we just need to form a conjunction that takes each distinct proposition as a conjunct. That conjunction then constitutes our individual proposition Δ . Henceforth, and unless otherwise noted, whenever I speak of a proposition (or a hypothesis) Δ I will mean either a natively individual one or one that we constructed as an individual from a number of others.

Most hypotheses we encounter in science probably contain at least some disconnected parts and hence are monstrous to some extent. It is thus not enough to say that a hypothesis is monstrous. We must also find a way to gauge the extent of that monstrosity. That is, we need to put forth a monstrosity measure. Let us start with an informal characterisation: The degree of monstrosity of a proposition Δ is given by the ratio of the sum of disconnected pairs of its parts to the sum of the total number of pairs of its parts, i.e. connected and disconnected. In both cases, these sums are computed on the basis of all distinct ways of distributing the total content of Δ into parts. To give an example of two distinct distributions: Δ may be distributed into propositions A, B where $A: \zeta_1 \wedge \zeta_2 \wedge \zeta_3$ and $B: \xi_1$ and Δ may be distributed into propositions A', B' where $A': \zeta_2 \wedge \zeta_3$ and $B': \xi_1 \wedge \zeta_1$. Why compute monstrosity on the basis of all distinct ways of distributing the total content of Δ ? Because different ways of distributing that content can have an effect on its degree of monstrosity. To neutralise this problem we avoid arbitrarily choosing one such distribution over all others and instead take all of them into account. Thus, any whiff of arbitrariness disappears and we cannot even be accused of leaving any one distribution out - without a vote so-to-speak.

Formally, the extent of the monstrosity μ of a proposition Δ is given by the following function:

$$\mu(\Delta) = \frac{\sum_{i=1}^n d_i^{\alpha,\beta}}{\sum_{i=1}^n t_i^{\alpha,\beta}}$$

where $d_i^{\alpha,\beta}$ denotes the number of disconnected pairs α, β in a given content distribution i , $t_i^{\alpha,\beta}$ denotes the total number of connected plus disconnected pairs α, β in a given distribution i and n denotes the total number of content distributions. To determine the number of disconnected pairs in a given content distribution we count how many times a *different* pair of relevant deductive consequences α, β turns out to satisfy clauses (i), (ii) and (iii) simultaneously.¹⁰ Any pair that is not disconnected is counted as connected. The higher (/lower) the value of $\mu(\Delta)$ the more (/less) monstrous the content of Δ .

¹⁰ The emphasis on *different* pairs is meant to rule out counting a pair more than once within a given content distribution. Thus, if $\alpha_1: \varphi, \beta_1: \psi, \alpha_2: \psi, \beta_2: \varphi$, we evaluate either pair α_1, β_1 or pair α_2, β_2 but not both.

It is important to note that μ is only intended as a provisional measure of the contrived-ness of hypotheses. That is to say, I am fully aware that this proposal needs revisions.¹¹ Even so, I consider it a solid first step in the right direction. One major advantage is that the measure is purely objective since its output is entirely dependent on the notions of confirmational connectedness and disconnectedness, themselves wholly determined by facts about the world such as the true relative frequencies and/or true propensities of events, states-of-affairs, properties, etc. Another major advantage is its wide-ranging applicability. The account doesn't place any debilitating restrictions on propositions Δ . As a result, the said propositions can be drawn from a large pool of entries, which includes central hypotheses, auxiliaries, explanantia and explananda. This not only allows us to gauge the monstrosity of the most commonly touted relations, e.g. the relation between a central hypothesis and an auxiliary hypothesis or the relation between an *explanans* and an *explanandum*, but also of any other relation we can think of, e.g. the relation between one auxiliary hypothesis and another.

6. Unification as Inverted Monstrosity*

If the monstrosity measure does indeed capture the degree to which the content of a hypothesis is contrived, artificially put together or forcibly united, then it is not unreasonable to expect that the less monstrous a hypothesis the more unified its content. That is, unification may be understood as the inverse of monstrosity. Well, not exactly... We first need to tweak the notion of monstrosity and its associated measure. That's because there is one kind of confirmational connectedness that is neither here nor there in cases of unification, namely that conveyed by clause (1). This clause gets activated when two content parts share content. But sharing content is a trivial way of unifying. Suppose a given pair α_1, β_1 is such that $\alpha_1: \varphi$ and $\beta_1: \varphi$ where φ is a true or partly true proposition. This ensures the satisfaction of clause (1) as there is a true or partly true proposition γ such that $\alpha_1 \vdash_r \gamma$ and $\beta_1 \vdash_r \gamma$ and hence α_1 and β_1 are confirmationally connected. But surely unifying requires bringing together content, not repeating it. The first thing to do then is to drop clauses (1) and (i) from their respective notions. Moreover, it is important to explicitly forbid content duplication across a pair α, β , i.e. at the level of a pair's joint content $\alpha \wedge \beta$, by introducing an external non-superfluous-ness qualification. That's because such content duplication adversely affects our evaluations through the remaining clauses. Take, for example, clause (2). Without the aforesaid qualification, we would be allowed to compare α_1 and β_1 but these turn out to be confirmationally connected: $P(\alpha_1/\beta_1) \neq P(\alpha_1)$ since $P(\varphi/\varphi) = 1$ and $P(\varphi) < 1$. But, as already mentioned, we want to avoid such trivial ways of unifying. Hence the appeal to the external non-superfluous-ness qualification.

Let us call the revised notion 'confirmational connectedness*'. It can be expressed as follows:

¹¹ I can already foresee problems that demand the modification of some of the things I said. Alas, there is no time or space to develop them here.

Any two content parts of a non-self-contradictory proposition Γ expressed as propositions A, B are *confirmationally connected** if, and only if, for some pair of internally and externally non-superfluous propositions α, β where $A \vdash_r \alpha$ and $B \vdash_r \beta$: either (1') where $0 < P(\alpha), P(\beta) < 1, P(\alpha/\beta) \neq P(\alpha)$ or (2') there is at least one atomic proposition δ such that $\alpha \wedge \beta \vdash_r \delta, \alpha \nmid_r \delta$ and $\beta \nmid_r \delta$.

Mutatis mutandis, we can obtain the correlative notion of confirmational disconnectedness*. On the basis of these two notions we can assert that a non-self-contradictory molecular proposition Δ is *monstrous** if, and only if, some of its content parts are confirmationally disconnected*. Finally, we can revise our monstrousness measure in light of the above alterations. Thus, $\mu^*(\Delta)$ is almost exactly like $\mu(\Delta)$, the only difference being that the notions of confirmational connectedness and disconnectedness are replaced by the notions of confirmational connectedness* and disconnectedness*. We are now ready to express the unification u of a proposition Δ with the following function:

$$u(\Delta) = 1 - \mu^*(\Delta)$$

Of central importance here is the fact that the objectivity of the monstrousness* measure carries over to the unification measure since the latter is advanced as the inverse of the former. In other words, the objectivity of the unification measure is guaranteed by the fact that its output is entirely dependent on the notions of confirmational connectedness* and disconnectedness*. Whether these notions are satisfied for any given pair of propositions is, as we have already witnessed, a matter wholly determined by the way the world is like. Herein lies the greatest strength of this approach. The world itself decides what is physically connected to what, in what way and to what degree. In short, unification, under the proposed measure, is not in the eye of the beholder.

Let us motivate this relation between monstrousness* and unification with some examples. Take a true hypothesis that conjoins propositions expressing disparate facts:

H_1 : q_1 is a black raven on Earth \wedge q_2 is a decaying particle on Jupiter $\wedge \dots \wedge$ q_n is an exploding neutron star.

Such a hypothesis should come out as highly disunified.

Is this what our measure tells us? That sure seems so. There are presumably very few, if any, confirmational connections* between the different conjuncts. That's because there are presumably very few real world connections between the expressed facts. For example, whether or not q_1 is a black raven on Earth is presumably probabilistically independent of whether or not q_2 is a decaying particle on Jupiter. Thus, were we to go through our calculations, H_1 would get a high monstrousness* and hence a low unification score. And that's precisely what one would expect from such an obviously disunified hypothesis. Compare H_1 to the following true hypothesis that conjoins propositions expressing related facts:

H_2 : r_1 is a black raven \wedge r_2 is a black raven $\wedge \dots \wedge$ r_m is a black raven.

There are presumably very many confirmational connections* between its conjuncts. That's because there are presumably very many real world connections between the expressed facts. One systematic confirmational connection* that holds between all pairs of conjuncts is through the following proposition:

R_1 : There are at least two black ravens.

Note that R_1 is *not* a relevant deductive consequence of any given individual conjunct belonging to H_2 . But it is a relevant deductive consequence of any pairing of conjuncts, e.g. r_1 is a black raven $\wedge r_2$ is a black raven *or* r_1 is a black raven $\wedge r_3$ is a black raven. Now R_1 is clearly not an atomic proposition but, by that very fact, at least one of its content parts is an atomic proposition. Let us denote any such content part as follows: R_1^c . We have just established that between all such pairs of conjuncts there is always a proposition δ , namely R_1^c , exactly as demanded by condition (2') in the characterisation of confirmational connectedness*. That is to say, these conjuncts are systematically confirmationally connected*. Thus, were we to run through our calculations, H_2 would get a very low monstrosity* and hence a very high unification score.

So far, so good! Now consider a slight peculiarity. H_2 would not get a lower unification score when compared to a statement expressing the same content in a condensed form:

H_3 : Ravens r_1 to r_m are black.

Critics may find this consequence unacceptable. Such critics would be particularly aggravated in the case where we compare a universally quantified statement, e.g. all ravens are black, to a logically equivalent long conjunction, e.g. r_1 is a black raven $\wedge r_2$ is a black raven $\wedge \dots \wedge r_n$ is a black raven \wedge nothing else is a raven.¹² They might argue that a universally quantified statement brings out the unity of the content much more than a long conjunction and hence deserves to be awarded a higher unification score. I would now like to put this view to rest. It should be clear that one formulation could not be more confirmed than the other, for the simple reason that they both express the same content. That is to say, without a difference in content there cannot be a confirmational difference. But, if all that matters in judging the unity of a hypothesis are these confirmational minutiae, as I have been suggesting with my notions of confirmationally (dis-/)connectedness*, then it should also be clear that one formulation could not be more unified than the other. In other words, how the same content is sheathed should have no effect on its unification score. The measure of unification proposed above respects this point. It respects it by computing a unification score on the basis of *all* distinct ways of distributing some content into parts. Since the two formulations express the same content they have *all and only* the same content distributions. But that just means that they are equally connected* and hence equally unified.

I would like to draw this section to a close with an admittedly cursory discussion of a historical example. Fresnel's equations for the reflection and transmission of light are derivable from Maxwell's theory of electro-magnetism. The latter is famous for having unified the behaviour of all sorts of phenomena. Visible light, according to this theory, is but a

¹² An unresolved puzzle is how to deal with the conjunct 'nothing else is a raven'. One intuition is to exclude such conjuncts from monstrosity* and unification evaluations.

specific kind of electro-magnetic wave or radiation. The very fact that visible light satisfies the same well-confirmed equations as other kinds of electro-magnetic radiation, e.g. radio waves, means that Maxwell's theory makes substantial strides in unifying its content. The unification exhibited in this case is unification in its purest form: Various phenomena under study are treated as instances of the same regularity. What does our account say about this sort of case? Well, among other things, the fact that all kinds of electro-magnetic radiation satisfy the same equations means that there is a true or at least partly true proposition corresponding to each instance of this regularity, e.g. m_1 : ' E_1 is an electromagnetic wave satisfying Maxwell's equations', m_2 : ' E_2 is an electromagnetic wave satisfying Maxwell's equations', etc. But now note that, when conjoined, any two such propositions relevantly entail a third, namely 'There are at least two electromagnetic waves satisfying Maxwell's equations', that neither relevantly entails on its own. This last proposition is molecular. Therefore, at least one of its content parts is atomic. Call the atomic part R_2^c . All such pairings of conjuncts like $m_1 \wedge m_2$ relevantly entail an atomic proposition δ , namely R_2^c , exactly as demanded by condition (2'). In other words, these conjuncts are systematically confirmationally connected*. And that demonstrates one of the ways in which the content of Maxwell's theory is unified.

7. Practicability

I already pointed out that the current view's greatest strength is that it hitches a ride on the objective notions of confirmational connectedness* and disconnectedness*. That same ride also proves to be its greatest weakness. Even when we seek to compare two propositions whose content is relatively small, a great number of objective probabilities need to be taken into account. Since we have at best limited access to such probabilities the proposed measure is severely handicapped in its practicability.

Three points help mitigate the despair caused by this objection. The first concerns the proper aim of this essay. That aim was to propose an objective conception and measure of unification whose judgments accord well with clear cases of unified and disunified hypotheses. There was never an aspiration to provide a measure of unification that is fully practicable. Perhaps we cannot have a fully objective and fully practicable measure of unification. If the keys to unification are indeed the notions of confirmational connectedness* and disconnectedness* and these notions need to be spelt out in more or less the way they have been above then we might just have to live with a fully objective but less-than-fully practicable measure. If that were the case, people would then be free to choose between such a measure and less-than-fully objective ones that may be fully practicable.

The second point picks up from where the first one left off. There is no question of the measure being completely impracticable. We needn't have access to a complete score to get a good sense of which of two hypotheses is the more unified. As the examples above show, all sorts of information is accessible and can help make that judgement. This includes information about the relevant deductive relations between different content parts. For example, we know that each conjoined pair of H_2 's conjuncts entails R_1^c . That's sufficient to indicate a systematic confirmational connection* between the hypothesis' content parts. In other words, although these judgments may in some cases be more uncertain than others, the measure does not leave us completely in the dark.

The third and final point concerns the complexity of the measure. At present, getting the measure to yield a complete score involves a complex calculation requiring copious amounts of information, not all of which, as we have already seen, is readily available or even accessible. But what if there was a way to simplify the measure without sacrificing too much of what it tries to track. One idea potentially worth pursuing is to run computer simulations on all sorts of environments where the ‘real world’ connections are stipulated in advance. We can then check under what conditions more simplistic measures yield similar judgments to our proposed measure. It might turn out, for example, that for the kinds of conditions that are similar to those usually encountered in the real world it is more often than not sufficient to compute unification scores on the basis of very limited information. The measure proposed in this essay could thus act as a lighthouse, helping to guide the construction of another measure, one that is more, and perhaps even fully, practicable. Alas, though it may make for a good sequel essay, this idea is not one that I can pursue further here.

8. *A Word on Coherence*

When I first started digging into the unification literature I was almost instantly led to the related literature on coherence. The latter notion is traditionally, and very roughly, understood as a relation of ‘agreement’ or ‘mutual support’ between testimonies, beliefs, memories, propositions, etc. (Olsson 2005, Ch 2). There are clearly similarities between my unification project and various projects that seek to understand coherence, not least of which are the probabilistic relations of support that both attempt to lean on. The two projects may have much in common but they are not the selfsame. Accounts of coherence stress the importance of rewarding content duplication. That is, the more content gets duplicated the higher the coherence. As Olsson notes, “Now if [testimonial agreement and, in particular, testifying the same proposition] is not a case of coherence, then, I must confess, I have no idea of what that notion could possibly involve... Testimonial agreement is more than just coherent; it is very coherent” (p. 16). Indeed, he goes on to fortify this idea: “... the degree of coherence equals 1 if and only if $P(A \wedge B) = P(A \vee B)$, i.e. just in case A and B coincide” (p. 99). By contrast, accounts of unification like the one presented above ignore content duplication.¹³ That’s because, intuitively, to unify is to bring together things that are at least in some minimal respect distinct. To conclude, insofar as the projects of unification and coherence overlap I very much hope that any contributions made in this essay shed light not only on the coveted notion of unification but also on the equally coveted notion of coherence.

Acknowledgements

For useful feedback I would like to thank my colleagues Gerhard Schurz and Paul Thorn as well as two anonymous referees. I gratefully acknowledge the German Research Founda-

¹³ Schurz (1999) is keenly aware of this difference between the two notions: “*coherence minus circularity equals unification*” (p. 98) [original emphasis].

tion (Deutsche Forschungsgemeinschaft) for funding my research under project B4 of Collaborative Research Centre 991: The Structure of Representations in Language, Cognition, and Science. Part of this paper has been written while working on the project “Aspects and Prospects of Realism in the Philosophy of Science and Mathematics” (APRePoSMa) during a visiting fellowship at the University of Athens. The project and my visits are co-financed by the European Union (European Social Fund—ESF) and Greek national funds through the Operational Program “Education and Lifelong Learning” of the National Strategic Reference Framework (NSRF)—Research Funding Program: THALIS—UOA.

REFERENCES

- Armstrong, David Malet. 1989. *A Combinatorial Theory of Possibility*. Cambridge: Cambridge University Press.
- Friedman, Michael. 1974. Explanation and Scientific Understanding. *Journal of Philosophy* 71(1): 5-19.
- Gemes, Ken. 1998. Hypothetico-Deductivism: The Current State of Play; The Criterion of Empirical Significance: Endgame. *Erkenntnis* 49: 1-20.
- Goodman, Nelson. 1983. *Fact, Fiction and Forecast* (4th ed.). Cambridge, MA: Harvard University Press.
- Hempel, Carl Gustav and Paul Oppenheim. 1948. Studies in the Logic of Explanation. *Philosophy of Science* 15(2): 135-175.
- Holton, Gerald. 1969. Einstein, Michelson, and the ‘Crucial’ Experiment. *Isis* 60(2): 132-197.
- Kitcher, Philip. 1976. Explanation, Conjunction, and Unification. *Journal of Philosophy* 73(8): 207-212.
- . 1989. Explanatory Unification and the Causal Structure of the World. In *Scientific Explanation: Minnesota Studies in the Philosophy of Science*. Vol. 13, eds. Philip Kitcher and Wesley Salmon, 410–505, Minneapolis: University of Minnesota Press.
- Salmon, Wesley. 1998. *Causality and Explanation*. Oxford: Oxford University Press.
- Schurz, Gerhard. 1991. Relevant Deduction: From Solving Paradoxes towards a General Theory. *Erkenntnis* 35: 391-437.
- . 1999. Explanation as Unification. *Synthese* 120: 95-114.
- Schurz, Gerhard and Karel Lambert. 1994. Outline of a Theory of Scientific Understanding. *Synthese* 101: 65-120.
- Votsis, Ioannis. 2014. Objectivity in Confirmation: Post Hoc Monsters and Novel Predictions. *Studies in History and Philosophy of Science* 45: 70-78.
- , forthcoming. Ad Hoc Hypotheses and the Monsters within. In *Fundamental Issues of Artificial Intelligence Synthese Library*, ed. Vincent C. Müller, Berlin: Springer.

IOANNIS VOTSI is a Senior Lecturer at the New College of the Humanities (London) and Assistant Director of the Düsseldorf Center for Logic and Philosophy of Science at the University of Düsseldorf. His main research area is the philosophy of science but he also has active interests in the philosophy of language, philosophy of logic, philosophy of artificial intelligence and meta-philosophy. He has co-edited a number of special issues and volumes on themes like scientific realism, novel predictions, underdetermination and theory-ladenness.