

This is a draft. Please cite the published version, published by BRILL in *Grazer Philosophische Studien*, DOI:<https://doi.org/10.1163/18756735-000096>

Epistemic Infrastructure for a Scientific Metaphysics

Amanda Bryant

Abstract: A naturalistic impulse has taken speculative analytic metaphysics in its critical sights. Importantly, the claim that it is desirable or requisite to give metaphysics scientific moorings rests on underlying epistemological assumptions or principles. If the naturalistic impulse toward metaphysics is to be well-founded and its prescriptions to have normative force, those assumptions or principles should be spelled out and justified. In short, advocates of naturalized or scientific metaphysics require epistemic infrastructure. This paper begins to supply it. I first sketch my conception of suitably naturalized or scientific metaphysics. I then lay out a number of candidate epistemic principles centring around the notion of theoretical constraint. I offer several arguments for the principles, based on statistical likeliness, agreement, falsity avoidance, and methodological efficiency and inefficiency. Finally, I show how scientific metaphysics satisfies the epistemic principles and is therefore preferable to its traditional rivals.

Keywords: naturalized metaphysics; scientific metaphysics; metametaphysics; epistemology of metaphysics; epistemic infrastructure

A naturalistic impulse pervades much recent philosophy and has taken speculative analytic metaphysics in its critical sights.¹ The call to *naturalize* metaphysics is, roughly, the call to constrain it with science. Its central motivations include pessimism about the epistemic import of philosophical intuitions, respect for the epistemic distinction of science, and recognition of fruitful points of contact between science and metaphysics.² The subject of this issue, *inductive metaphysics*, integrates the “inductive forms of inference that have been employed successfully in the empirical sciences” as well as “accepted empirical sources of justification and knowledge”, including “the (admittedly provisional) results of the empirical sciences” (Scholz 2018, 206). In that respect, naturalized and inductive metaphysics are continuous. One key difference, however, is that it is not essential to naturalized metaphysics

¹ Bryant 2017, Hawley 2006; Ladyman and Ross 2007; Maclaurin and Dyke 2012, 2013; Maddy 2007; Waters 2017.

² I borrow the phrase ‘epistemic distinction’ from Haack, who comments that “distinction, unlike privilege, has to be earned; and the natural sciences have earned, not our uncritical deference, but our tempered respect” (2003, 23).

that it assign an explicit role to inductive inference (though it often does employ such inferences).³ However, both naturalized and inductive metaphysics reflect the naturalistic impulse and have decidedly scientific moorings. I'll refer to metaphysics with such scientific moorings as *scientific metaphysics*.

The naturalistic impulse is thoroughly normative in spirit. Its proponents claim that a scientific metaphysics is *preferable* to its traditional rivals and that if metaphysics is to be any good, then it *ought* to be the scientific kind. Here we face the immediate question of what kind of normative force these claims have and of what gives them that force. While there may be pragmatist or other ways of cashing out the force of such claims, I believe the kind of normativity at issue is often epistemic. The thought is that scientific metaphysics is preferable *epistemically* to its rivals, and if metaphysics is to be any good *epistemically*, then it ought to be the scientific kind. If so, then the claim that it is desirable — or even requisite in some sense — to give metaphysics scientific moorings must rest on underlying epistemological assumptions or principles. If the naturalistic impulse toward metaphysics is to be well-founded and its methodological prescriptions to have normative force, it is crucial that those assumptions or principles be spelled out and justified. In short, advocates of scientific metaphysics require epistemic infrastructure.

Prima facie, one option would be to rest our naturalistic prescriptions on empiricist epistemology. On such a view, we should tie metaphysics to science because the road to knowledge is empirical, and because science seems to be a paragon of empirical investigation. However, empiricism doesn't adequately motivate naturalization. For one thing, science doesn't have a monopoly on the empirical. On the contrary, empirical evidence is pervasive outside of science. We constantly run everyday inferences from empirical evidence — for instance, from the sounds I'm hearing now, I conclude that there is traffic on the street outside. If what did the epistemic heavy-lifting were just empirical evidence, then the prescription would be for metaphysics to integrate empirical evidence; nothing would follow about the need for *scientific* evidence in particular — that is, evidence gathered in the course of scientific practice.⁴ But what makes naturalistic prescriptions *naturalistic* and not merely empirical is their explicit mention of science. For another thing, science isn't always so thoroughly empirical. As Chakravartty points out, a view according to which

³ A *deductive* naturalized metaphysics might proceed in the style of Quine, reading ontology off of science. Alternatively it might take inspiration from Popper, proceeding by conjecture, deduction, and falsification or disconfirmation.

⁴ The demarcation problem looms. You might think we cannot hope to motivate the naturalistic prescriptions at issue without first answering it. This issue requires serious consideration and I don't have the space to discuss it here. I will just say, briefly, that I don't think we need a precise demarcation in order to talk about science and to assign it (again borrowing Haack's phrase) epistemic distinction. It is enough that we can gesture toward the institution, point to its successes, and identify a family of methods, practices, and attitudes on which those successes rest.

science is comprehensively empirical is “a caricature” (2013, 33). That is because “not all sciences actually make novel predictions (evolutionary biology), or employ experiments (string theory), or are successful in manipulating things (cosmology)” (2013, 34). Since empiricism assigns no special status to *scientific* evidence in particular, and since it arguably calls into question the epistemic credentials of less empirical pockets of science, it can’t provide the needed epistemic infrastructure.

Another *prima facie* option would be to endorse a kind of scientism according to which science is the best or only good way of knowing, and therefore any form of inquiry worth its salt should proceed by appealing to or modelling itself after science. If it were the case that science should guide all inquiry, then of course it would follow that it should guide metaphysics. However, a view so sweeping is both difficult to defend given the many aims and targets of inquiry, and also far stronger than necessary. Far stronger because, for instance, we might expect the philosopher to check science when investigating the underlying, descriptive nature of reality, without likewise expecting her to do so when investigating the just, the moral, or the beautiful. At least *prima facie*, there is a clear rationale for the one that doesn’t obviously apply to the others — namely, that science tells us about the underlying, descriptive nature of reality! In other words, one can deem the naturalization of metaphysics epistemically desirable or requisite without buying into a sweeping scientism of this form.

A less radical epistemic infrastructure can do the work of supporting the restricted naturalistic prescriptions at issue and establishing the preferability of scientific metaphysics over its traditional rivals. This paper will contribute to that infrastructure. It is worth noting that I don’t intend to supply an exhaustive infrastructure to the exclusion of further contributions; I mean only to supply some candidate principles centring around what I will argue is an epistemically important feature of theories — namely their level of constraint. These candidate epistemic principles may be entirely consistent with additional principles, for which there may be independent arguments.

Section 1 of the paper will briefly sketch my conception of scientific metaphysics. Section 2 will lay out a number of candidate epistemic principles centring around the notion of theoretical constraint. Section 3 will deliver a number of arguments for the candidate principles, based on considerations of statistical likeliness, consensus, falsity avoidance, and methodological efficiency and inefficiency. Section 4 will show how scientific metaphysics satisfies the epistemic principles on offer, and finally, Section 5 will conclude.

1 Scientific Metaphysics

There are various ways of marking the distinction between what I’ve here called ‘scientific metaphysics’ and its contrast class. I call the contrast class ‘free range metaphysics’, the hallmark of which is that it is only nominally constrained by science and instead

constrained primarily by demands for simplicity, consistency, intuitive plausibility, and explanatory power (Bryant 2017). I said above that scientific metaphysics, on the other hand, is characterized by its ‘scientific moorings’. This is meant to suggest that scientific metaphysics bears an important relation to science. This relation may be characterized in a number of ways: a scientific metaphysics is motivated by, responsive to, engaged with, answerable to science, and so on. Any of these notions will do, but each needs spelling out. I tend to prefer the notion of engagement. So I will say that a scientific metaphysics is one that engages appropriately with science. But what is it for metaphysics to *engage* science?

Different advocates of naturalization prescribe different modes of engagement. For instance, Ladyman and Ross (2007) say naturalized metaphysics unifies science. Maclaurin and Dyke (2012) say it premises metaphysical claims on scientific results. In my view, metaphysics can engage science in a number of fruitful ways, including: integrating scientific posits or theses into metaphysical theories in a non-ad-hoc manner, drawing conclusions of philosophical interest on the basis of scientific evidence, interpreting scientific data, correcting scientists’ interpretation of the data where appropriate, revising metaphysical claims in light of new data or theoretical understanding, and formulating new or revising old metaphysical questions in light of new scientific understanding. The list is not meant to be exhaustive. For instance, I also support Ereshefsky’s (unpublished) and Waters’ (2014) recent claims that metaphysics should attend to scientific practice.⁵ A metaphysics that does at least some of these things conscientiously and painstakingly is a scientific metaphysics. At any rate, while there are a number of ways of envisaging scientific metaphysics, what unites its advocates is their shared view that metaphysical speculation should be reined in and that the products and practices of science have some crucial role in so doing.

2 Epistemic Principles

Before articulating any epistemic principles, it is important to note that the principles I will develop will be appropriate only relative to certain purposes.⁶ It should already be clear that those purposes are epistemic, rather than pragmatic or heuristic. In fact, there may be a number of distinct epistemic purposes for which my principles are apt. That said, the principles may not be apt for just *any* epistemic purpose. To see this, consider the distinction (famously introduced by Reichenbach 1938) between contexts of justification and of discovery. Independently of what Reichenbach meant, the distinction can be characterized in a number of ways (see Hoyningen-Huene 2006). For my purposes here, I will take a context

⁵ However, since my epistemological principles and arguments will focus on theories and their features, they will not motivate the practice-based approach to scientific metaphysics. Independent arguments would be needed for that purpose. I thank Ken Waters for drawing this to my attention.

⁶ I thank Carl Gillett for pointing this out.

of justification to be one in which the primary aim of the relevant epistemic agents is to formulate justified beliefs, claims, or theories; I take a context of discovery to be one in which the primary aim is to uncover previously unknown facts. Different methods may be more or less advantageous given different aims — for instance, creativity, inspiration, innovation, and the like may be more important for the purposes of discovery than for the purposes of justification.⁷ For reasons I will articulate below, the principles I will develop might not be apt within a context of discovery, but would be so within a context of justification — understood as a context in which the aim is convergence toward a theory consilient with and strongly confirmed by the best available evidence at the time (i.e. convergence toward a justified theory).⁸ I will return to this below.

My central thesis is that, supposing we're in a context of justification so understood, one of the features of theories we should care about is the extent to which their content is constrained. By theoretical *content*, I mean just what the theory says. We can understand and quantify this in terms of the atomic propositions that make up the theory.⁹ A theoretical *constraint* is a limit on which theoretical contents we take to be allowable into a given theory as constitutive parts. When we adopt a constraint, it limits the theoretical contents we could rationally countenance (the range of theories we might accept) and the theoretical contents that we do in fact accept. For instance, for any theory, there is some body of data that we wish the theory to account for, which acts as a constraint on the theory. We expect our theory to be *adequate* to the data — that is, to be roughly consistent with it. As a rule, we don't accept claims that are inconsistent with the data unless there are overriding factors, like they are indispensable aspects or products of an otherwise empirically successful theory. We also expect our theory to *explain* the data, so our explanatory expectations also constrain the sorts of theoretical contents we entertain. All sorts of things can constrain our theories aside from the demands for empirical adequacy and explanatory power. We might require our theories to be, *inter alia*, internally consistent, unificatory, consilient with the broader theoretical landscape, or virtuous (simple, beautiful, elegant, convenient, familiar, fruitful). Sometimes we require them to be consilient with political or ideological goals. So for any given theory, the constraints we place on its content can depend on our local epistemic

⁷ Precisely what a *context* is is a perennially interesting and difficult question, but one on which nothing presently hangs. I would happily translate the word away if pressed to do so.

⁸ I believe the principles would also be apt within a context where the aim is convergence on truth. However, I want my view to be neutral with respect to scientific realism and humble with respect to the epistemic prospects of a scientific metaphysics. In other words, I want to allow that justification at time *t* may be the best we can do, both in science and in metaphysics.

⁹ The qualification that the propositions we're counting are 'atomic' is needed, because otherwise a theory represented by one giant conjunctive proposition would count as having little theoretical content — and stringing together atomic propositions shouldn't effect how contentful we think a theory is.

norms and pragmatic goals, as well as on features of our broader intellectual and socio-political landscape.

Theoretical constraints fall on a spectrum of strength. If the data is scant, our adequacy to the data constraint will not be particularly strong, since a good deal of theoretical content is consistent with relatively little data. The more data, the stronger the constraint, since relatively more data is consistent with relatively less theoretical content. The strength of our adequacy constraint will also depend on our auxiliary hypotheses. Without any auxiliaries, most theories have no empirical consequences. With the right auxiliaries, just about any theory can be made consistent with any data. So if we gerrymander ad-hoc auxiliaries, our adequacy constraint will be weak. If we limit the sorts of auxiliaries we countenance, we strengthen the constraint considerably.

Call a theoretical constraint *robust* when it falls on the high end of the spectrum — that is, when it allows relatively few atomic propositions into the theory as constitutive parts and prohibits relatively many.¹⁰ Call a constraint *permissive* when the opposite holds — that is, when it allows many atomic propositions and prohibits few. Note that these notions are relative, not categorical. The more robust (or less permissive) a constraint, the less content it allows into the theory; the more permissive (or less robust), the more it allows. A robust theoretical constraint is like a selective bouncer at a club, who lets in only the few people who meet his strict criteria for inclusion. It is good that he does so, because (at least according to some) a club that is too inclusive is not a very good club. I will argue that the same is true of theories.

It is crucial for the success of the arguments I will advance that I be especially clear about the particular way I am using comparative phrases in this paper. On my view, a *more constrained* theory is one with *less* theoretical content (i.e. one with fewer atomic propositions). This follows from my definition of a constraint as a limit on theoretical content, and from my understanding of theoretical content in terms of the constitutive atomic propositions of the theory. I am not claiming that a theory is relatively well-constrained in virtue of being consistent with relatively fewer atomic propositions. Relative constraint is not primarily a matter of how many propositions are consistent with the claims of the theory.¹¹

¹⁰ With the language of ‘allowing’ and ‘prohibiting’, I mean to suggest that once we adopt a (sufficiently precisified) constraint, a proposition usually either satisfies it or doesn’t. I say ‘usually’ to allow that there may be vague cases. Likewise from the epistemic agent’s point of view, some cases will be clear-cut and some won’t. For instance, if a constraint demands internal consistency, then if we’re working with a group of relatively simple and non-vague propositions, we can easily tell whether the constraint permits or rules out some additional (simple and non-vague) proposition. A less clear-cut case might involve substantive judgments regarding evidence, confirmation, explanatory power, epistemic risk, and so on.

¹¹ I thank an anonymous reviewer for pressing me to make these clarifications.

I will not argue for the blanket claim that the more constrained the theory, the better. On the contrary, there is such a thing as too much constraint. Theories that are too constrained cannot say much at all. Think, for instance, of the *Meditations*, in which Descartes' system of knowledge is constrained by so demanding a standard of justification that, had he held true to it, it should have permitted few if any theses at all. My point will be, rather, that there is an ideal range on the spectrum between anything goes and nothing goes — and that the range is closer to the latter extreme than to the former.

Constraints are too robust or too permissive relative to certain aims that I take to be integral to the activity of theorizing. Theories by their very nature are explanatory endeavours. A theory that is too robustly constrained can say hardly anything and as such doesn't stand a chance of meeting its explanatory aims. That is, it is hardly a theory at all. Likewise, a theory whose constraints are too permissive cannot meet its explanatory aims either, but for different reasons. A permissive theory simply entertains or encompasses too much to be explanatory. The thing we aim to understand and account for gets lost in a flood of irrelevant details. By contrast, an adequately explanatory theory zeroes-in on relevant details. So it is in virtue of the kind of thing theories are that some constraints are too robust or too permissive.

Being robust is not sufficient reason to implement a theoretical constraint. For instance, suppose we adopt a theoretical constraint that allows us to countenance only theses that demonstrate the existence of God. The constraint would be robust since little, if anything, would make it into the theory. Nevertheless, there is no epistemic reason to adopt such a constraint (though there might be non-epistemic ones). This goes to show that our selection of theoretical constraints should be well-motivated, and of course it is a deep and important question which sorts of motivations are good ones. I take such qualitative considerations to be complementary to the quantitative considerations at issue here. While I don't have space to address them at this juncture, I take them to require addressing if the epistemic infrastructure is to be complete. At any rate, I don't wish to argue that we should adopt constraints solely on the grounds that they are robust.

Likewise, being permissive is not sufficient reason to abandon a constraint. Whether we should do so depends on whether it has a good rationale. Take the internal consistency constraint. It rules out infinitely many inconsistent theories, but permits infinitely many consistent ones. So it is a permissive constraint. That *doesn't* mean we should give up on internal consistency! If the internal consistency constraint is a good one, it is so quite apart from how robust or permissive it is, but rather because contradictions are (ordinarily) bad or problematic in some way. So again, I don't wish to argue that we should abandon constraints solely on the grounds that they are permissive.

I wish to argue, rather, that (in a context of justification) we should as a rule robustly constrain theories. There may be good reasons to accept permissive constraints, such as the internal consistency constraint, but we ought to supplement them with other, more robust

ones. Suppose our club bouncer is singularly dedicated to preventing weapon-carrying persons from entering the club. Supposing our club is in, say, an affluent part of Manhattan, that criterion will be relatively permissive. But the criterion is nevertheless a good one! This just goes to show that we don't *only* care about the exclusivity of our club; among other things, we also care about how safe it is. But to the extent that we do care about exclusivity, we should implement additional criteria at the door. Similarly, in the case of theories, I don't wish to argue that we should care *only* about how constrained our theories are. I wish to argue simply that we should care about it *inter alia*. If so, then when a theory is constrained only by a small number of permissive constraints, we should hold the theory accountable to additional, more robust ones.

Call a theory *robustly constrained* when it has sufficiently many suitably robust constraints on its content (or when enough individually permissive constraints have the same joint result). Note that the notion is vague, since there is no precise answer to what counts as 'sufficiently many' or 'suitably robust'. But again, we can identify the extreme ends of the spectrum of constraint based on whether theories say too much or too little to stand a chance of meeting their explanatory aims. Where those constraints allow too many claims for the purposes of explanation, they are insufficiently strong; where they allow too few, they are overly so.

So far I've spelled out my notions of theoretical constraint and of *robust* constraint in particular. I will now articulate general reasons for thinking that robustly constraining theories is a good epistemic policy. In particular, I will distinguish three candidate epistemic principles. If any of the principles is true, then we have some motivation for thinking that as a rule, we should robustly constrain our theories. The principles are as follows:

Weak Constraint Principle:

Ceteris paribus, theories that are robustly constrained are epistemically preferable to those that are not.

Moderate Constraint Principle:

Theories that are robustly constrained are epistemically preferable to those that are not.

Strong Constraint Principle:

To be epistemically adequate, a theory must be robustly constrained.

The weak principle makes robustly constrained theories preferable *ceteris paribus*; the moderate principle makes them preferable full-stop; the strong principle makes robust constraint a necessary condition of what I call epistemic adequacy. For a theory to be *epistemically preferable* is for it to better fulfill our epistemic aims (and recall that in this case,

I've stipulated that the aim is convergence on a justified theory). For a theory to be *epistemically adequate* is for it to be rationally acceptable at t given the evidence available at t , where 'accepting' is to be understood in L. Jonathan Cohen's sense:

to accept that p is to have or adopt a policy of deeming, positing, or postulating that p — i.e. of including that proposition or rule among one's premisses for deciding what to do or think in a particular context, whether or not one feels it to be true that p . (1992, 4)

In other words, a theory is epistemically adequate when, given the evidence, a good reasoner might well adopt a policy of positing the theory or its component propositions.¹² When I say that a theory is *epistemically inadequate*, I mean that from an epistemic standpoint, it isn't rationally acceptable (in Cohen's sense of acceptance).

It should now be clear why these principles wouldn't be appropriate in a context of discovery. While constraint is (I will argue) important for the purposes of convergence on a justified theory, the aim of discovery might sometimes be furthered by being *less* constrained and allowing ourselves to think outside the box. That is, allowing ourselves greater leeway with respect to the range of possibilities we countenance may be beneficial in such a context.

As for contexts of justification, I have said my principles are apt for such contexts on a certain understanding of them, but the story I tell here requires finesse. That is because constraint and justification may work against one another. In particular, robustly constraining a theory may lead not to its justification, but to its disconfirmation or falsification.¹³ For instance, suppose we have a strong adequacy to the data requirement, together with some reasonable auxiliaries and a continually growing body of data. What we can expect to happen in such a situation depends on the details — what the theory is about, how well-established and mature it is, and so forth. Nevertheless, in a lot of cases it wouldn't be terribly surprising if the growing body of data ultimately ended up disconfirming the theory. After all, oftentimes falsity is easy to get and truth hard. If the disconfirming evidence presents a significant and persistent explanatory challenge, eventually we'll no longer think we're justified in accepting the theory. So if our end is justification, constraint might not be a surefire means.

I grant that robust constraint can ultimately lead to a failure of justification. However, in my view, the occasional failure of an epistemic policy to secure its aim isn't particularly surprising or troubling. Just as infallibility is nearly always too strong an

¹² Note that I've simply cashed out 'rationally' in terms of what a good reasoner would do. I don't wish to avow any substantive theory of rationality.

¹³ I thank Helen Beebe for pointing this out.

expectation of epistemic agents even when true belief is the end-goal, surefire justification is too strong an expectation of epistemic policies even when justification is the end-goal. If the worry is not simply that constraint and justification can pull apart, but that they are inherently at odds, the argument from statistical likeliness below may motivate us to think the opposite.

At any rate, I have cashed out ‘context of justification’ in terms of ‘convergence on a justified theory’ because I want us to have in mind a *process* of justification. I wish to allow that disconfirmation may be part of that process, in that disconfirming a thesis or theory motivates revision and may lead to the development of a successor thesis or theory that *is* justified by the best available evidence at the time. We learn from our mistakes, after all. It is relative to such a process that the constraint principles are apt.

At any rate, the main purpose of this section was to introduce three candidate epistemic principles revolving around the notion of robust constraint: one weak, one moderate, and one strong. The principles claim that robust constraint is either an epistemic plus or an epistemic requirement. For my purposes, the importance of the principles is that if any of them is true, then we have some motivation for thinking that as a rule, we should robustly constrain our theories. I will now turn to defending the principles.

3 Epistemic Principles Defended

At its heart, the epistemological project at hand is normative. The notions of preferability and adequacy at play in my candidate principles are normative notions. *Prima facie*, to get epistemic normativity off the ground, we need to find the “basic normative factor – the ultimate goal or value – in the epistemic domain” (Klausen 2009, 161). We face a choice at the very foundations of normative epistemology, regarding which epistemic good or goods we take to be fundamental. The received view is that the ultimate epistemic good is true belief (Klausen 2009, 161), but there are other candidates, such as *justified* true belief (DePaul 2001, Feldman 2002, Smithies 2012) and *knowledge* (Adler 2002, Bird 2007, Engel 2007, Huemer 2007, Littlejohn 2013, McHugh 2011, Peacocke 1999, Sutton 2007, Williamson 2000). Still more candidates include rational acceptability, agreement, fitness or pragmatic success, and the avoidance of falsehood (Klausen 2009, 161). However, to motivate my epistemic principles, I do not need to take a stand on the foundational question. Rather, I will attempt to show only that my epistemic principles promote a number of putative epistemic goods, while remaining neutral with regard to whether those goods are foundational. My conclusions will be conditional on our recognizing these putative epistemic goods as genuine, but it will be outside the scope of the paper to establish that they are. Nevertheless, I believe it will be a rhetorical boon to show that constraint conduces to a

range of putative epistemic goods, including relative statistical likeliness, agreement, avoidance of falsity, and methodological expediency.

My first argument claims that by limiting the theoretical contents we countenance or accept, constraints make our theories more likely to be true. This is a purely statistical matter, relating to *how much* the theory claims. The argument goes as follows.

Argument from statistical likeliness:

- P1) *Ceteris paribus*, relatively simple theories are statistically likelier to be true than their more complex (non-nested) rivals.
- P2) *Ceteris paribus*, statistically likelier theories are epistemically preferable.
- P3) Robustly constrained theories are relatively simple.
- P4) Robustly constrained theories are statistically likelier to be true than their more complex (non-nested) rivals.
- C) *Ceteris paribus*, robustly constrained theories are epistemically preferable.

Importantly, the kind of simplicity at issue in P1 is not standard ontological parsimony, nor is it relative logical weakness. Rather, it is a Bayesian form of simplicity that is a matter of the theory having relatively few adjustable parameters. In this context, adjustable parameters are aspects of a theory that, like variables or blank lines, are open to specification or filling-in. Adjustable parameters make theories flexible; they provide wiggle-room with respect to the specific, determinate content of the theory. The idea behind P1 is that *ceteris paribus*, the more adjustable parameters your theory has, the less statistically likely it is to be true.

Now, let me acknowledge that Popper famously argued the opposite — that simpler theories can be less statistically likely. Take the following example from Sober:

(LIN) There exist numbers a_0 and a_1 such that $y = a_0 + a_1x$.

(PAR) There exist numbers a_0 , a_1 , and a_2 such that $y = a_0 + a_1x + a_2x^2$. (2015, 89)

Insofar as LIN contains fewer adjustable parameters than PAR (it lacks a_2), LIN is simpler in the Bayesian sense than PAR. But, Popper's point is that LIN cannot be more probable than PAR. That is because LIN is equivalent to $(PAR) \wedge (a_2 = 0)$, and a conjunction is never more probable than its individual conjuncts. So we have a case where the simpler of two theories is less statistically likely.¹⁴ However, Sober points out that “when Popper talks about examples like LIN and PAR, he takes these models to be *mutually compatible*; in fact, one of them logically entails the other. Statisticians describe this relationship by saying that LIN is *nested* inside of PAR” – and anyone claiming that the simpler theory is statistically likelier

¹⁴ I thank Peter Godfrey-Smith for pressing me to consider this point.

assumes that the rival models are not nested in that manner (2015, 92). So that is why P1 above contains the parenthetical qualification that the rival theories must not be nested.

Why think that P1 is true? Harold Jeffreys (1939) famously claimed that we should assign relatively higher prior probabilities to relatively simpler claims. However, without some epistemically-motivated reason to assign priors in that fashion, this begs the question at issue (see Kelly 2004 and Sober 2015). A more compelling reason to think that P1 is true is that “models that contain fewer adjustable parameters postulate a narrower range of possibilities over which probabilities must be distributed” and thus models with fewer adjustable parameters are likelier (Sober 2015, 125). Moreover, a Bayesian form of Ockham’s Razor (according to which, for statistical reasons, we shouldn’t multiply adjustable parameters beyond necessity) emerges naturally in statistical and scientific reasoning (Nichols et al. 2016, 534 fn. 1).¹⁵ If statistical and scientific reasoning have epistemic distinction (independently of the arguments here), then the razor’s role therein may be suggestive. Both considerations require a much fuller treatment than I am able to give here, but I hope they make at least a *prima facie* case for P1.

As for P2, I’ve already stated that it is beyond the scope of the paper to establish the value of putative epistemic goods such as relative statistical likeliness. So I will make the argument conditional on the truth of that premise. That said, one reason to think that likelier theories are preferable — vis-a-vis the aim of converging on a justified theory — is that there is a natural connection between justification and statistical likeliness understood in terms of Bayesian conditionalization. So understood, the likelihood of a theory is its posterior probability — i.e. its probability conditional on the evidence. While the nature of justification is a matter of longstanding controversy, it is natural to think that if the evidence makes a theory likely, then *ceteris paribus* the evidence justifies the theory (or justifies its acceptance).¹⁶

The *ceteris paribus* clause is crucial, because there are a good deal of other features of theories that we care to preserve, which may trade off against statistical likeliness. For instance, we have certain descriptive and (as I have already said) explanatory aims in formulating theories. We typically want our theories to be informative, to contribute to understanding, to describe their subject-matter in a certain level of detail, and to be as complete as possible. On the whole, it is better to have a descriptive and explanatory theory that is only somewhat likely than to have a theory that is not particularly descriptive or explanatory, but a good deal likelier. As an extreme case, a tautology is perfectly likely, but not terribly informative. So we *don’t* want a principle according to which the less our theory

¹⁵ On Bayesian Ockham’s Razor, see also Berger (1985), Blanchard et al. (2018), Box and Tao (1973), Henderson et al. (2010), Jefferys and Berger (1992), MacKay (2003), and Sober (2015).

¹⁶ Of course there is much more to justification. Note that I have just said that a relatively high posterior probability is only *sufficient* for justification *ceteris paribus*; I have not said it’s necessary.

says, the better. That would have the silly consequence that we should prefer less informative theories to more informative ones. We want, rather, a principle that tells us to prefer the simpler of two theories that are equally satisfactory in other important respects.

As for P3 — the claim that robustly constrained theories are relatively simple — I believe this follows as a matter of definition. That is because robustly constraining a theory is a matter of restricting degrees of freedom with respect to theoretical content. Robust constraints make theories less flexible; they give them less wiggle-room with respect to their specific, determinate content. For that reason, robust constraint seems to entail the Bayesian kind of simplicity at issue.

It follows from the claim that simpler theories are statistically likelier (P1) and the claim that robustly constrained theories are relatively simple (P3) that robustly constrained theories are statistically likelier (P4). Finally, having established P1-P4, it follows that robustly constrained theories are *ceteris paribus* preferable (C). So we have a probabilistic argument for our weak constraint principle.

We may motivate a second argument for the weak constraint principle by noting that robust constraint is conducive to agreement, which we might also take to be an epistemic good.

Argument from agreement:

- P1) *Ceteris paribus*, relatively simple theories are more conducive to agreement.
- P2) *Ceteris paribus*, theories that are more conducive to agreement are epistemically preferable.
- P3) Robustly constrained theories are relatively simple.
- C) *Ceteris paribus*, robustly constrained theories are epistemically preferable.

Just as a simpler theory is more statistically likely, a simpler theory is also, *ceteris paribus*, more conducive to agreement (P1). That is because, the less content the theory has, the less there is to disagree about. Notwithstanding the presence of disagreement across most epistemic contexts, rampant disagreement can be a symptom of poor constraint. The less well-constrained a theory is, the greater and more pervasive disagreement will be, and the wider the range of theoretical alternatives we might consider live options. Religion illustrates the point well, since human beings have posited such a rich variety of religious theories throughout history and since our evidence so dramatically underdetermines them. The point is that robust constraint and disagreement inversely correlate. The *ceteris paribus* clause in P1 is meant to catch unusual cases in which a fuller, more descriptive theory commits only to truisms, while a less descriptive one makes controversial claims (for instance, a theory that commits to hundreds of simple mathematical equations versus a theory that says only that abortion is morally wrong).

P2 claims that theories that are more conducive to agreement (let's say among rational agents) better fulfill the aim of convergence on a justified theory. Again, I don't have space to argue for this. A possible connection between justification and agreement is that we sometimes take agreement to be justificatory — as in cases of corroboration. At any rate, the argument is again conditional on the truth of this premise. If agreement is indeed instrumental to our epistemic aim, then P2 follows. The *ceteris paribus* clause is meant to signal that, as in the case of statistical likeliness, agreement can be trumped by other factors, given our descriptive and explanatory aims. A fully descriptive and explanatory theory that generates substantial disagreement is preferable to a theory that we can agree on, but that contributes little to our understanding. I established the claim that robustly constrained theories are relatively simple (P3) in my explanation of the argument from statistical likeliness above, and the conclusion that robustly constrained theories are *ceteris paribus* epistemically preferable follows. So if we take agreement to be an epistemic good, then we have further reason to accept the weak constraint principle.

In sum, when theoretical constraint is robust and therefore permits a narrow range of theoretical content, there are two potentially valuable outcomes: greater statistical likeliness of our theory and greater consensus with respect to that theory. If those outcomes are instrumental to our epistemic aim, it follows that we should prefer robustly constrained theories, other things being equal.

We have already seen an argument built around the claim that simpler theories are statistically likelier to be true. A related claim is that the more theoretical contents we rule out and the fewer we permit, the likelier we are to avoid substantial falsehood.¹⁷ I say 'substantial' falsehood because strict falsity is cheap — most theories are idealizations and are therefore strictly false.¹⁸ The argument goes as follows.

Argument from falsehood avoidance:

- P1) The more content we exclude from a theory and the less we permit, the more likely we are to avoid substantial falsity.
- P2) *Ceteris paribus*, theories that are more likely to avoid substantial falsity are epistemically preferable.

¹⁷ We might think that these are the very same claim. However, acquiring true beliefs and avoiding false ones are arguably distinct epistemic goods (see Klausen 2009, Whiting 2013, and Zagzebski 2003). If so, then having one's theoretical claims be relatively likely to be true and having them be relatively unlikely to be false are distinct epistemic goods.

¹⁸ I owe this point to Stathis Psillos, who proclaimed "Falsity is cheap, guys!" at his talk "From the Evidence of History to the History of Evidence: Re-thinking the Pessimistic Induction" at *The History of Science and Contemporary Scientific Realism Conference*, Indianapolis, February 19, 2016.

P3) Robustly constrained theories exclude relatively many putative theoretical contents and permit relatively few.

C) *Ceteris paribus*, robustly constrained theories are epistemically preferable.

The idea is that the less you say, the less likely you are to say something false. That's because it is much harder to say something true than false. If I simply make up a claim, what are the odds that it will be true? Not great. If so, then for any arbitrary claim, if we don't know anything about it or its context, we should assign a higher prior probability that it's false than that it's true. Now, it's certainly *possible* for a constraint to rule out more truths than falsehoods, but keep in mind that what I've just given is another statistical argument. Excluding relatively more content from our theories makes us more likely to avoid substantial falsity (P1) because of how much cheaper falsehood is than truth.

P2 claims that theories that are more likely to avoid substantial falsity better fulfill the aim of convergence on a justified theory. One reason might be that minimizing the chance of *being* wrong likewise minimizes the chance of *discovering* that we are wrong and thereby of showing ourselves to be unjustified. So while minimizing the chance of falsity doesn't exactly move us toward justifiedness, it prevents us from moving in the opposite direction — and that is valuable for those who wish to be moving in the right direction. As in the previous arguments, the *ceteris paribus* clause does needed work. It signals that making substantial falsity unlikely is an important epistemic aim only *inter alia*. We would of course prefer a theory that tells us something about our subject-matter and risks some falsity over a theory that tells us next-to-nothing and does not risk falsity.

The claim that robustly constrained theories exclude relatively many putative theoretical contents (P3) is just built into the notion of robust constraint. By definition, part of what makes a constraint robust is that it excludes a broad range of putative theoretical contents from consideration. It follows that robustly constrained theories are *ceteris paribus* epistemically preferable (C).

We have seen in the discussion so far the importance of explanatory and descriptive aims. Every theory aims to adequately account for some subject-matter. That is, there are facts the theory aims to acquaint us with. Assuming that not every claim about some subject-matter is true (assuming for instance the falsity of trivialism), then the facts we aim to discover are to some degree limited. A poor method of discovering them would be to countenance propositions at random, hoping by sheer luck to stumble upon some facts. A more efficient method would better *target* the desired facts.

Argument from methodological expediency:

P1) Theories that better target the relevant facts are epistemically preferable.

P2) *Ceteris paribus*, robustly constrained theories better target the relevant facts.

C) *Ceteris paribus*, robustly constrained theories are epistemically preferable.

Theories that better target the relevant facts (or perhaps theories resulting from methods that do) are epistemically preferable (P1) assuming that such theories better fulfill the aim of converging on a justified theory. While my aims in this paper don't include establishing the epistemic import of effective targeting, we might think that the better we are at targeting desired facts in the course of constructing a theory, the more justified we are in accepting that theory. While I've remained neutral on the nature of justification, one explanation open to externalists here would be that relatively effective targeting is relatively reliable.

Robust constraint is conducive to effective targeting (P2) because part of the job of a constraint is to help the theory target the facts. When theoretical constraints screen off large amounts of putative theoretical content, they narrow the range of theoretical contents that we could countenance or accept. Theoretical constraints direct our attention away from large swaths of possibility space, toward others. The more robustly constrained our theories, the more narrowly they target putative facts. The *ceteris paribus* clause captures cases where constraints are robust yet lack epistemic rationale or purchase — and where they in fact undermine targeting.¹⁹ For instance, if we place on a climate-scientific theory a constraint requiring theoretical content to please the fossil fuels lobby, our theory may well fail to effectively target the facts.

Note, however, that the screening off function *does not suffice* for effective targeting; it merely conduces to it (*ceteris paribus*). If targeting the desired facts is like directing a flashlight toward a specific region of a dark room, we want the beam of light to be 1) suitably narrow and 2) pointing in the right direction. Screening off putative theoretical content narrows the beam of light, but doesn't necessarily point it in the right direction. The additional requirement that our constraints be well-motivated would hopefully get us the latter. At any rate, since robust constraint helps with targeting, it follows that robustly constrained theories are epistemically preferable *ceteris paribus* (C).²⁰

Let's take stock. I have advanced several arguments, each of which supported what I have called the weak constraint principle, according to which we should prefer robustly constrained theories *ceteris paribus*. These arguments were based on considerations of relative statistical likeliness, agreement, falsehood avoidance, and methodological efficiency. In fact, I

¹⁹ I thank an anonymous reviewer for correctly insisting on this particular *ceteris paribus* clause.

²⁰ Compare Kelly (2004) who argues, similarly, that simplicity helps us to get at the truth *more efficiently*, since opting for simpler models will mean fewer theory changes en route to the truth. Kelly explains:

disregarding Ockham's advice opens you to a needless, extra U-turn or reversal in opinion prior to all the reversals that even the best of methods would have to perform if the same answer were true. So you ought to heed Ockham's advice. Simplicity doesn't indicate the truth, but it minimizes reversals along the way. (2004, 492)

think a stronger version of the argument from methodological efficiency is defensible, which goes as follows.

Argument from methodological inefficiency:

- P1) Theories that fail to effectively target the relevant facts are epistemically inadequate.
- P2) Theories that fail to be robustly constrained fail to effectively target the relevant facts.
- C) Theories that fail to be robustly constrained are epistemically inadequate.

If a theory doesn't target the facts well (or the methods by which it is produced don't), then where our epistemic aim is to converge on a justified theory, it doesn't make sense to accept the theory. I don't mean to say that a theory that fails to be true can't possibly be justified; by 'failing to effectively target the facts' I don't mean 'failing to be true'. Rather, I mean something methodological: whatever methods produce the theory are not effective at targeting, at least in this instance and perhaps more broadly. As a result, the flashlight beam is not fixed precisely on the things we wish to see. If so, it's not clear how we could form justified beliefs about them. If a theory fails to target the facts well, then it fails to inform us about them, and as such fails to be rationally acceptable (P1).

To see why a failure to robustly constrain a theory is a failure to effectively target the relevant facts (P2), think about what a theoretical constraint does. It limits what our theory can say, thus narrowing the flashlight beam. The fewer and more permissive the constraints, the more the theory can say. If the content of one's theory is not constrained, then the theory can say anything whatsoever. If it isn't *robustly* constrained, then the range of claims that could go into the theory is enormous. Where that is the case, the theory is far less likely to be true or approximately true on the whole. If it does get at the truth, it does so by luck. So in order to target the facts *effectively*, robust constraint is a requirement. It follows that robust constraint is a necessary condition of epistemic adequacy (C). This argument motivates the strong constraint principle.

In sum, a number of arguments support the epistemic importance of robust theoretical constraint. Considerations of statistical likeliness, agreement, falsehood avoidance, and methodological efficiency supported the weak constraint principle. Considerations of methodological inefficiency supported the strong constraint principle. At this juncture, a couple things are worth noting. First, the fact that the arguments support two different principles is not particularly troubling, since the weak and strong constraints appear to be compatible. Second, the fact that no argument supports the moderate constraint principle is perhaps unsurprising, since the principle suggests, implausibly, that no considerations could outweigh the epistemic importance of robust constraint. But we have seen several examples in which other considerations did outweigh its importance. Notwithstanding the implausibility of the moderate principle, it was useful to articulate it in

order to clearly distinguish the available positions. At any rate, so long as one or more of the principles is true, then as a rule, we should robustly constrain our theories.

4 Scientific Metaphysics and Robust Constraint

Scientific metaphysics is relatively robustly constrained because it introduces a scientific constraint on its theoretical content. The scientific constraint is robust because scientific claims, together with the appropriate auxiliary hypotheses, frequently have observable consequences and are generally held accountable to a wealth of empirical evidence — for example when they are expected to explain, predict, or retrodict that evidence.²¹ Scientific claims are routinely subject to confirmation and disconfirmation in the course of scientific practice. When a claim or theory is disconfirmed by evidence posing a significant and persistent explanatory challenge to it, it is often rejected. That's the bouncer in action, throwing people out of the club. Science has disconfirmed a vast number of claims (claims committed to spontaneous generation, caloric fluid, luminiferous aether, and so on), and will disconfirm vastly many more. Routine disconfirmation shows us that experimental practices strongly constrain the content of many scientific theories, and thus science is robustly constrained. Hence, the more a scientific metaphysics holds itself accountable to science, the more robust constraint it inherits.

Now, one might argue that the scientific constraint does not rule out crank hypotheses about gods and evil demons and so isn't very constrained after all. Granted, the scientific constraint doesn't rule out crank hypotheses in the sense of demonstrating their falsity. However, it does rule them out in the sense that science simply doesn't countenance them. That is, to the extent that god and evil demon hypotheses are not empirically tractable, they are not taken to be scientifically interesting, or the proper subject-matter of science. This is not merely a sociological fact. Again, science often entertains claims that can be held up to empirical evidence — and to the extent that no empirical evidence could disconfirm the existence of god or an evil demon, such claims fall outside the domain of scientific interest.

In order to make the comparative claim that scientific metaphysics is preferable to its scientifically-unmoored rival, free range metaphysics, I would have to do more than show that scientific metaphysics is relatively robustly constrained. I would also have to show that free range metaphysics is relatively poorly constrained. I have made that case elsewhere (Bryant 2017).

²¹ I say 'frequently' and 'generally' because some scientific claims may be empirically intractable now and for the foreseeable future. String theory may furnish some examples (if you agree that its claims are in principle empirically intractable and yet nonetheless scientific), as may other highly-theoretical forms of physics and cosmology. These are not meant as fighting words, but merely as a nod to physical and practical limits on human beings and our measuring devices.

5 Conclusion

The claim that we ought to pursue scientific metaphysics requires epistemological infrastructure, and it has been the aim of this paper to begin to supply it. I articulated a number of candidate epistemic principles revolving around the notion of robust constraint. I argued that if greater statistical likeliness, agreement, falsity avoidance, and methodological efficiency are epistemic goods, then we should accept at least one of the candidate principles. If any of those principles is true, then as a rule we should robustly constrain our theories. Since scientific metaphysics is robustly constrained, we now have a foundation for thinking that we ought to pursue it.

Acknowledgments

This research was supported by the Fundação para a ciência e a tecnologia, award #PTDC/FER-HFC/30665/2017. For invaluable feedback, I heartily thank Michael Devitt, Peter Godfrey-Smith, Ian Miller, David Papineau, David Yates, commentators Bryson Brown and Nicholas Danne, as well as an anonymous reviewer. I also wish to thank audiences for their questions and comments on presentations of this work at: *6th Ernst Mach Workshop: Epistemology of Metaphysics*, WCPA Annual Meeting 2017, ARPA Conference 2017, *Traditional and Inductive Metaphysics*, CPA Annual Congress 2018, CSHPs Annual Meeting 2018, The University of Bristol, and SMS Annual Conference 2019.

References

- Adler, Jonathan 2002. *Belief's Own Ethics*. Cambridge, MA: MIT Press.
- Berger, James 1985. *Statistical Decision Theory and Bayesian Analysis*. New York: Springer-Verlag.
- Bird, Alexander 2007. "Justified Judging." *Philosophy and Phenomenological Research* 74(1), 81–110.
- Blanchard, Thomas , Lombrozo, Tania and Nichols, Shaun 2018. "Bayesian Occam's Razor Is a Razor of the People." *Cognitive Science* 42(4), 1345-1359.
- Box, George, and George C. Tiao 1973. *Bayesian Inference and Statistical Analysis*. New York: John Wiley & Sons.
- Bryant, Amanda 2017. "Keep the Chickens Cooped: The Epistemic Inadequacy of Free Range Metaphysics." *Synthese*, 1-21.

- Chakravartty, Anjan 2013. "On the Prospects of Naturalized Metaphysics." In: *Scientific Metaphysics*. Edited by Don Ross, James Ladyman, & Harold Kincaid, Oxford: Oxford UP, 27-50.
- Cohen, L. Jonathan 1992. *An Essay on Belief and Acceptance*. Oxford: Clarendon Press.
- Dawid, Richard 2013. "Theory Assessment and Final Theory Claim in String Theory." *Foundations of Physics* 43(81), 81–100.
- DePaul, Michael 2001. "Value Monism in Epistemology." In: *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*. Edited by Matthias Steup, Oxford: Oxford UP, 170–183.
- Engel, Pascal 2007. "Belief and Normativity." *Disputatio* 2(23), 179–203.
- Ereshefsky, Marc unpublished. "Science and Metaphysics: Lessons from Microbiology."
- Feldman, Richard 2002. "Epistemological Duties." In: *The Oxford Handbook of Epistemology*. Edited by Paul Moser, Oxford: Oxford UP, 362–384.
- Haack, Susan 2003. *Defending Science - within Reason: Between Scientism and Cynicism*. Amherst: Prometheus Books.
- Hawley, Katherine 2006. "Science as a Guide to Metaphysics?" *Synthese*, 149(3), 451–470.
- Henderson, Leah, Noah Goodman, Joshua Tenenbaum, and James Woodward 2010. "The Structure and Dynamics of Scientific Theories: A Hierarchical Bayesian Perspective." *Philosophy of Science* 77(2), 172–200.
- Hoyningen-Huene, Paul 2006. "Context of discovery versus context of justification and Thomas Kuhn." In: *Revisiting Discovery and Justification*. Edited by Jutta Schickore and Friedrich Steinle, Dordrecht: Springer, 119-131.
- Jefferys, William and Berger, James 1992. "Ockham's razor and bayesian analysis." *American Scientist*, 80, 64– 72.
- Jeffreys Harold 1939. *Theory of Probability*. Oxford: Oxford UP.
- Kelly, Kevin 2004. "Justification as Truth-Finding Efficiency: How Ockham's Razor Works." *Minds and Machines* 14, 485–505.
- Kitcher, Philip 1981. "Explanatory Unification." *Philosophy of Science*, 48(4), 507–531.
- Klausen, Søren Harnow 2009. "Two Notions of Epistemic Normativity." *Theoria* 75, 161–178.

- Ladyman, James and Don Ross, with David Spurrett, and John Collier 2007. *Every Thing Must Go: Metaphysics Naturalized*. Oxford: Oxford UP.
- Littlejohn, Clayton 2013. "The Russellian Retreat." *Proceedings of the Aristotelian Society* 113(3 pt 3), 293–320.
- MacKay, David 2003. *Information Theory, Inference and Learning Algorithms*. Cambridge: Cambridge UP.
- Maclaurin, James and Dyke, Heather 2012. "What is Analytic Metaphysics For?" *Australasian Journal of Philosophy* 90(2), 291–306.
- 2013. "What Shall We Do with Analytic Metaphysics? A Response to McLeod and Parsons." *Australasian Journal of Philosophy* 91(1), 179–182.
- Maddy, Penelope 2007. *Second Philosophy: A Naturalistic Method*. Oxford: Oxford UP.
- McHugh, Conor 2011. "What Do We Aim At When We Believe?" *Dialectica* 65 (3), 369–392.
- Nichols, Shaun, Shikhar Kumar, Theresa Lopez, Alisabeth Ayars, and Hoi-Yee Chan 2016. "Rational Learners and Moral Rules." *Mind and Language* 31(5), 530–554.
- Peacocke, Christopher 1999. *Being Known*. Oxford: Oxford UP.
- Perfors, Amy, Joshua Tenenbaum, Thomas Griffiths, and Fei Xu 2011. "A Tutorial Introduction to Bayesian Models of Cognitive Development." *Cognition* 120, 302–321.
- Reichenbach, Hans 1938. *Experience and Prediction*. Chicago: The University of Chicago Press.
- Scholz, Oliver 2018. "Inductive Social Metaphysics---A Defence of Inference to the Best Explanation in the Metaphysics of Social Reality: Comments on Katherine Hawley." *Journal for General Philosophy of Science* 49(2), 199–210.
- Smithies, Declan 2012. "The Normative Role of Knowledge." *Noûs* 46(2), 265–288.
- Sober, Elliott 2015. *Ockham's Razors: A User's Manual*. Cambridge: Cambridge UP.
- Sutton, Jonathan 2007. *Without Justification*. Cambridge, MA: MIT Press.
- Tenenbaum, Joshua, and Thomas Griffiths 2001. "Generalization, Similarity, and Bayesian Inference." *Behavioral and Brain Sciences* 24, 629–640.
- Waters, Kenneth 2017. No General Structure. In: *Metaphysics and the Philosophy of Science: New Essays*. Edited by Matthew Slater & Zanja Yudell, Oxford: Oxford UP, 81–108.

Whiting, Daniel 2013. "The Good and the True (or the Bad and the False)." *Philosophy* 8(2), 219–242.

Williamson, Timothy 2000. *Knowledge and Its Limits*. Oxford: Oxford UP.

Zagzebski, Linda 2003. "Intellectual Motivation and the Good of Truth." In: *Intellectual Virtue: Perspectives from Ethics and Epistemology*. Edited by Linda Zagzebski & Michael DePaul, New York: Oxford UP, 135–154.