

What kind of is-ought gap is there and what kind ought there be?

P.D. Magnus and Jon Mandle

forthcoming in *Journal of Moral Philosophy*

final authors' draft as of 17oct2016; originally posted on-line 26aug2014

<https://www.fecundity.com/job/paper.php?item=isought>

ABSTRACT: Some philosophers think that there is a gap between *is* and *ought* which necessarily makes normative enquiry a different kind of thing than empirical science. This position gains support from our ability to explicate our inferential practices in a way that makes it impermissible to move from descriptive premises to a normative conclusion. But we can also explicate them in a way that allows such moves. So there is no categorical answer as to whether there is or is not a gap. The question of an is-ought gap is a practical and strategic matter rather than a logical one, and it may properly be answered in different ways for different questions or at different times.

Introduction

It is common to think that there is a gap between *is* and *ought* — that is, that premises which are entirely descriptive or factual cannot legitimately yield a conclusion which is normative or ethical.

This is often called *Hume's principle* or *Hume's Law*, because of a passage in Hume's *Treatise* where he seems to invoke it (pp. 469-70). However, the question of Hume exegesis is vexed. Arguably, Hume was concerned more with the origin of our moral concepts than with how we might come to apply them in particular instances. Rachel Cohon writes that the real point of the much quoted passage is that “our basic recognition of morality — our grasp of what vice and virtue fundamentally are — is not acquired by inferring” (Cohon 2008, p. 94). This is separate from the issue of how the moral concepts, once acquired, might be used in arguments. Moreover, there are other theses about inference that have been credited to Hume. Gillian Russell (2010) is led to formulate both Hume's Law and Hume's Second Law, only one of which is about ethics. To avoid pitfalls of exegesis and historical reference, we will call it the *is-ought gap thesis* or just the *gap thesis*.

As we have said, some philosophers think that the gap thesis is obviously true. Others, convinced by counterexamples offered by A.N. Prior (1960) and John Searle (1964), take it to be simply false. Still others feel that the counterexamples can be avoided by some refinement of the gap thesis and that the refined thesis — although not transparently obvious — can be proven to be true.¹

¹ For example, Singer (2015) argues that the purported counter-examples of Prior and others have conclusions that “do not offer any substantive normative claims

This dialectic presumes that as a matter of logic or perhaps semantics, the thesis is either simply true or simply false. We argue that this dichotomy does not hold for the gap thesis, at least in the form that is relevant for considering the relation between empirical science and ethics. Instead, whether there are legitimate inferences to normative conclusions from purely factual premises depends on how we explicate our practices of ampliative inference. There *can* be an is-ought gap when appropriate in specific contexts, but there does not need to be. The is-ought gap is not something forced on us by logic or semantics. Rather we might choose to enforce it, meaning that we regiment our discourse so as to block inferences from purely descriptive premises to normative conclusions. This would be advantageous sometimes, so in some contexts we would have good reasons to reject such inferences, but not always. Whether the gap thesis holds in some context is thus a question of strategy — indeed, of values — rather than of logic or metaphysics.

Suppose that someone advocates a specific course of action, is challenged to justify themselves, and cites particular facts as reasons for the recommendation. To make it specific, suppose someone says we ought to go on a picnic and (when asked why) answers that the weather is gorgeous. How can we make sense of their inference from the fact that the weather is nice to the claim that we ought to go on a picnic? Of course, it is not deductively valid to infer from the single premise that the weather is nice to the conclusion that we ought to go on a picnic. So we could treat this as an enthymeme, an implicit deductive inference with a suppressed premise like "Nice weather is a *pro tanto* reason to have a picnic." Alternatively, we could treat the inference as ampliative, relying on the regular connection between nice weather and appropriate conditions for a picnic.² As an ampliative inference, it would require no further premise. The former explication respects the gap thesis by refusing a normative conclusion without there being a (perhaps unstated) normative premise, while the latter explication violates the gap thesis by allowing the inference to go through without such a premise. Our central claim is that the choice between the two explications is a practical matter of how to regiment the inference. The choice between explications, and therefore whether the gap thesis holds, is not dictated by logic or metaphysics. Whether the gap thesis holds in a particular context, depends on the practical reasons that we have to choose one explication rather than another.

In sec. 1, we look at Ronald Dworkin's use of the gap thesis to drive a wedge between scientific and normative discourse. The version of the gap thesis that interests us is precisely the one invoked by Dworkin. In sec. 2, we consider the alleged counterexamples offered by Prior and Searle. One lesson to draw from them is that the gap thesis parallels the problem of induction. So in sec. 3, we consider how ampliative inference works in science. Although scientific inference involves inductive risk, that risk can be carried either by explicit premises (which

to someone who accepts the premises." A more precise statement of the is-ought gap thesis that avoids such valid but vacuous arguments is the following: "There are no valid arguments from non-normative premises to a relevantly normative conclusion" (pp.199, 200).

² This could, but need not, be understood as a semantic relation between "nice weather" and "picnic".

assert a connection between observations and unobserved phenomena) or by rules of inference (which authorize moving from observations to unobserved conclusions without further premises). Which of these construals of scientific inference is to be preferred is a *practical* rather than a logical matter. In sec. 4, we turn back to the gap thesis. Like scientific inference, our normative inferences can be explicated in either of two ways: one in which the gap thesis holds (and the relevant inferences can only be authorized by explicitly normative premises) and another in which it does not (and the appropriate inferences do not require such premises). Importantly, this is not a choice we need to make once and finally. We might have a gap for some topics at some times but not for other topics or at other times.

1. Science, ethics, and the gap thesis

Ronald Dworkin (2011) argues that empirical science and normative enquiry are fundamentally separate activities. The former employs what he calls criterial or natural kind concepts, whereas the latter employs what he calls interpretive concepts. The justification of empirical scientific claims ultimately bottoms out in what the natural world is like, but moral judgments cannot be based on brute facts alone. At crucial moments in his argument, in order to establish this as a stark and fundamental divide, he invokes the principle that it is impossible to draw a normative conclusion from strictly descriptive or factual premises.

The gap thesis which Dworkin has in mind is that “no series of propositions about how the world is, as a matter of scientific or metaphysical fact, can provide a successful case on its own — without some value judgement hidden in the interstices — for any conclusion about what ought to be the case” (2011, p. 44). Dworkin writes that this seems obviously true to him, but he also goes on to give a brief argument for the thesis:

Consider this attempt to violate it: “Jack is in great pain and you could easily help him. Therefore, just for that reason, you have a moral duty to help him.” If this is a good argument, as it stands, then some principle about what makes an argument a good one must be at work. What is that principle? It cannot be any form of induction or generalization, because these would assume that you have had a moral duty in the past, which is a moral assumption. It cannot be a principle of deduction or semantic entailment. It needs something more, and that must be something — a hidden premise or an assumption about the nature of good moral reasoning — that is infused with moral force. (2011, pp. 44-5)

The strategy is to take a would-be counterexample to the gap thesis and show that it is not a counterexample after all. He presents an argument which purports to infer a moral conclusion from factual premises, asks what could possibly justify such an inference, and then argues that only a hidden moral (or normative) premise or assumption could do the work. Why does Dworkin think there must be a hidden moral premise? He considers and rejects three alternatives.

First, he claims that the inference cannot simply be a case of enumerative induction. For that, we would need to have seen moral duty accompany opportunities to help individuals in pain many times in the past. That would mean incorporating all of those past duties as premises for the inference to the moral conclusion, thus vindicating the gap thesis. This follows for 'induction' taken in a

narrow sense, to mean inference from past, observed instances to future, as-yet-unobserved instances of the same kind. However, this kind of projective induction does not exhaust scientific inference. Other non-deductive inferences allow scientists to posit unobservable entities, and the term 'induction' is sometimes used in a broader sense which also includes those ampliative inferences. Dworkin's point does not follow if we take *induction* in the broader sense.

Second, he claims that the argument is not deductively valid, something which is obvious if we reconstruct it formally.

Third, he claims that the inference is not one of semantic entailment. We take it that his point is that one can fully understand what all the terms in the premise mean and still fail to acknowledge that the conclusion follows. Yet, even if he is right about this case, his claim arguably does not hold in all cases. This allows the construction of informal counterexamples.

Dworkin's enumeration fails to address ampliative inferences beyond projective induction (a point we return to in sec. 3), and it fails to rule out more craftily constructed counterexamples (which we turn to in sec. 2).

2. Alleged counter-examples to the gap thesis

Counterexamples to the gap thesis can be divided broadly into two kinds: The first kind relies on the usual rules of deduction, so we call them formal. The second kind relies on semantic implication, so we call them informal.³

Numerous formal counterexamples are given by A.N. Prior (1960). One, which he attributes it to T.H. Mott, relies on the rules for disjunction. Call this ANP1:

Tea-drinking is common in England.

Therefore: Either tea-drinking is common in England or all New Zealanders ought to be shot.

The idea is that the premise is descriptive and the conclusion is normative. Yet the argument is valid, because the inference is allowed by the disjunction introduction rule. So it shows how to get a normative conclusion from merely descriptive premises.

This purported counterexample to the gap thesis takes for granted that the premise is properly classified as descriptive while the conclusion is normative, a presupposition that can be questioned. A gap-thesis enthusiast might reply by saying that the conclusion of ANP1 is descriptive rather than normative, because it does not actually license or prohibit any action.⁴ But this shifts the problem without solving it, as we can see if we consider ANP2:

Either tea-drinking is common in England or all New Zealanders ought to be shot.

Tea-drinking is not common in England.

Therefore: All New Zealanders ought to be shot.

³ Pigden (2016) marks this as the distinction between counterexamples to 'logical autonomy' and counterexamples to 'semantic autonomy'.

⁴ Ultimately, the burden is on a defender of the gap thesis to establish grounds for distinguishing factual from normative claims. There must be a meaningful distinction in order for there to be a meaningful gap thesis.

The reply to ANP1 was to say that ANP1's conclusion (which is the first premise of ANP2) is descriptive. The second premise of ANP2 is clearly descriptive, as well. The conclusion, validly derived, is normative. So if ANP1 is defused by saying that its conclusion is descriptive, then ANP2 is a counterexample to the gap thesis. One might attempt to reply to this by insisting that the conclusion of ANP2 is descriptive rather than normative, too, but it is of the canonically normative form 'All Fs ought to be G-ed' — if that is not normative, what is?

Moreover, any sentence could have been embedded as a disjunct in the conclusion of ANP1. For whatever one accepts as a normative sentence, an ANP-style counterexample can be constructed. Yet this very generality might suggest that the formal counterexamples do not show anything especially interesting. The normative part of the conclusion of ANP1 could as easily have been different. It could have been about different subjects, different actions, or different normative statuses. We could just as easily have concluded, 'Either tea-drinking is common in England or it is never permissible to shoot a New Zealander.' The normative content of the conclusion of ANP1 is *vacuous*.⁵ Various philosophers have attempted to give a precise characterization of this, but for our purposes it will suffice to say that the normative content of a conclusion is vacuous if it can be exchanged for any other content without changing the validity of the inference.

Once we have seen ANP1, it is easy to construct other formal counterexamples to the gap thesis which rely on other deductive rules: Given inconsistent descriptive premises, one can validly derive any normative conclusion one pleases; given any premises, one can derive logical truths which have normative subsentences; and so on. In all these cases, the normative sentence or formula in the conclusion is vacuous because the rules of deduction permit us to derive parallel conclusions with any arbitrary sentence or formula. So there is a somewhat refined version of the gap thesis which escapes such formal counterexamples: There is no normative conclusion that can be drawn from merely descriptive premises in which the normative content of the conclusion is non-vacuous.⁶

A famous informal counterexample is given by John Searle (1964). Call this JS:

Jones uttered the words "I hereby promise to pay you, Smith, five dollars."

Therefore: Jones ought to pay Smith five dollars.

Searle constructs the argument with several more steps, but the lynchpin of it is that promising is a speech act which Jones performs (as a matter of fact) and by which he becomes obligated to do what he said he would do (as an ethical matter). The inference is not deductive, and so not a formal counterexample. The conclusion follows only *ceteris paribus*. If Jones uttered the words while he and

⁵ Prior calls it contingently vacuous, because it would have non-vacuous normative consequences if the first disjunct were false. Given that tea-drinking actually is common in England, though, the disjunction tells us nothing about the bullet-worthiness of New Zealanders.

⁶ Working this out in the details involves numerous complications. See, for example, Pigden 1989, various contributors to Pigden 2010, and Singer 2015.

Smith were in character in a play he would not incur an obligation, Jones ought not pay the money to Smith (all things considered) if Smith would use it to commit atrocities, and so on. Searle takes it that the inference in JS is justified because of what the word ‘promise’ means.

Gillian Russell has recently discussed a variety of informal arguments and claimed that all of the purported counterexamples would, if successful, do too much. She thinks it is possible to prove the gap thesis, understood as just a special case of what she calls an *implication barrier thesis*. Other implication barriers are posed by the problem of induction: From claims just about particular things, it is impossible to derive general or universal conclusions. From claims just about the past, it is impossible to derive conclusions about the future.⁷

Russell suggests that the problem of induction and the gap thesis are structurally analogous to each other because each involves blocking certain inferences, and this underwrites what she calls “a good company argument” for the gap thesis (Russell 2010, p. 159). Nobody thinks that counterexamples like those she discusses solve the problem of induction, so there must be a way of formulating the particular-universal and past-future barrier theses which evades them — that is, there must be sophisticated formulations of the inductive barrier theses that stand in the face of the purported counter-examples. By analogy, the is-ought gap thesis might also have a formulation which avoids the counterexamples. Her idea is to suggest that, just as a precise formulation of the problem of induction will allow it to survive purported counter-examples, so too, a precise statement of the is-ought gap will allow it to survive its own purported counter-examples. In work with Greg Restall (Restall and Russell 2010), Russell provides a general formulation of barrier theses which she thinks does this.

Taking Russell’s good company argument seriously, we suggest a conclusion very different from the one that Russell draws. Despite the problem of induction, scientists nevertheless do generalize and draw inferences about the future from past, particular observations. And they are right to do so. Russell does not think that scientists should stop making generalizations and be content with simply cataloging discrete facts. Since the problem of induction does not establish an uncrossable inferential barrier between past and future, the “good company argument” suggests that there might not be an uncrossable barrier between is and ought, either. The ability of scientists to properly make inductive inferences does not disprove Restall and Russell’s barriers to deduction, but the barriers to deduction can be surmounted by ampliative inferences. Scientists are not (and should not be) limited to deductive inferences. There are legitimate ampliative inferences that scientists make which allow them to start from particular observations and arrive at general conclusions or conclusions about the future. One might hope that the gap between the descriptive and the normative could be bridged in a similar way.

To make sense of this, we need to look more closely at how ampliative inference works in science.

⁷ Russell calls these (Bertrand) Russell’s Law and Hume’s Second Law, respectively (2010, p. 155).

3. How ampliative inference is possible in science

To speak in sweeping terms, the old problem of induction teaches us that inductive generalization is not a matter of deductive inference. Instead, it is *ampliative*, meaning that the negation of the conclusion is logically compatible with the premises. As Hume argues, facts about the past do not logically necessitate facts about the future. Goodman's new riddle of induction shows further that inductive rules cannot be simply a matter of logical form. Any formal rule which licenses legitimate inductions can, by changing natural predicates to grue-ified ones, be made to license illegitimate inductions. The contents of the terms in the argument, not only the form, matter to the legitimacy of the inference. An exclusive focus on deductive arguments cannot capture this distinction.

The term 'induction' is sometimes meant narrowly just to mean generalization: inferring from the particular to the general and from the past to the future. Ampliative inference in science goes well beyond this. Scientists posit unobservable entities such as electrons. This is not inductively generalizing from observed cases to unobserved cases, but instead positing entities which are utterly unlike those which are directly observed on the basis of observation. The introduction of unobservable entities on the basis of observation is another implication barrier (in Russell's sense). So the problem in science is not just the narrow issues that concern Hume and Goodman. The lesson of their problems generalize, however. The legitimacy of an ampliative inference is never simply a matter of the formal (deductive) structure of the argument.⁸

John Norton (2003) articulates this lesson in what he calls *the material view of induction*. All ampliative inference is local. It relies on facts which hold only in specific domains. For example, scientists melt a sample of bismuth in the lab and infer the general melting point of bismuth, but they do not infer a general melting point for wax from a sample in the lab (Norton 2003, p. 649). Bismuth varies neither in composition nor in melting point, but wax varies too much. The typical operation of science, Norton claims, is to explicitly articulate the relevant, domain-dependent facts which license an inference. Scientists do not simply generalize about bismuth but not about wax. They explain this legitimate difference in inferential practice by saying that bismuth is an element while wax is an indeterminate mixture. The composition of a chemical element is invariant, which supports generalizations about its structure and properties. The composition of bits of wax — and so also their properties — are highly variable. So what starts as a two-line inference from a sample in the lab to a general claim can be supplemented with a *material postulate* which says explicitly that samples of bismuth are all the same element. If the material postulates are made sufficiently precise, then eventually the argument can be rewritten as a deductively valid one. Any ampliative inference can be made deductive this way, by explicitly spelling out the material postulates which underwrite it.

Norton argues that it is more advantageous to make material postulates explicit, because scientists are more reliable when assessing the truth of explicit

⁸ The term 'induction' is often stretched to label the broad class of legitimate but deductively invalid arguments that scientists make — that is, as a synonym for 'ampliative'. The lesson of Goodman's new riddle holds as much for induction *sensu lato* as induction *sensu stricto*.

material postulates than they are when assessing the reliability of a general rule. He writes, “As long as the inductive risk resides within the [inference] schema, we must assess it through a highly problematic judgment of the overall reliability of the relevant schema. We have little chance of coming to a clear judgment let alone determining how to reduce the risk. However once the risk is relocated in a material postulate in some local domain, our assessment of the inductive risk will depend in large measure on our confidence in the material postulate” (Norton 2003, p. 665).

Ingo Brigandt (2010) also argues that scientific reasoning is material—that it is legitimate because of its content rather than merely because of its form. However, he rejects Norton’s idea that domain-specific content is best represented in explicit material postulates. Rather, Brigandt argues that scientific inferences are justified by *material inference principles*, rules that are relied upon without being explicitly written down as part of the argument.⁹ On his account, “a material inference is justified based on the meaning of the various concepts contained in the premises and conclusion” (2010, p. 33). For example, having a reliable melting point is just part of what it means to be *bismuth*.

Brigandt argues that relying on inexplicit material inference principles allows for an interplay between inference, explanation, and discovery. Scientists are typically only able to write down enough premises to make their arguments deductive late in an enquiry. When charting unfamiliar phenomena, they nevertheless must make some inferences for the enquiry to progress. In science underway, it is crucial that they can rely on potentially risky inferences which they cannot (yet) fully articulate. He writes that, in practice, “it is doubtful whether the meaning of terms can always fully be reconstructed by explicit definitions — at least the ability to use language and draw valid inferences does not require this” (2010, p. 41).

So Norton and Brigandt agree on a material view of induction, but differ on how the domain-specific knowledge should be understood. Should it be written down as material postulates (Norton)? Or should it be left to the tacit knowledge of material inference principles (Brigandt)?

Scientific inference is typically ampliative rather than deductive, so there are implication barriers in science. Does this mean that it is illegitimate to draw a conclusion about the future from premises which are just about the past? *Yes* if we insist that all domain-specific details must be presented as explicit material postulates, but *no* if we allow material inference principles. To continue with our example: If we insist that the justification of any inference from a particular observation to a general conclusion about the melting point of bismuth in the future requires additional premises, then we will always need to add some explicit *premise* about the future to license such an inference. If we allow that the meaning of ‘bismuth’ involves it having the same melting point in the future as in the past, then the inference is legitimate without any explicit mention of the future in the premises. So, one might think, whether there is a past-future gap depends just on which of Norton or Brigandt is right.

⁹ Brigandt adopts the idea of material inference principles from Wilfred Sellars (1953) and Robert Brandom (1994). Brandom is himself influenced by Sellars.

However, this is a false dilemma. When observations and evidence are used as premises of an ampliative inference, there is risk involved. The success of the argument depends on domain-specific conditions. If those conditions fail to hold, then the conclusion may turn out to be false. The same risk is taken if additional premises are added which state explicitly that the conditions do obtain. In the former case, the risk is that the rule will lead us astray, while in the latter, the risk is that the premises will be false. As P.D. Magnus (2008) argues, the risk is same for a specified inference regardless of whether it relies on an explicit material postulate or on a material inference principle.

Imagine two scientists who agree on both the evidence and on what conclusions may be drawn from the evidence, except that one relies on material postulates and the other relies on material inference principles. The two stand to reap the same rewards and suffer the same accidents based on their inferences. In terms of immediate risk, what matters is just which inferences they accept, not how they choose to express those inferences. As Magnus writes, “It might... be sound advice to lift with the muscle of premises rather than the skeleton of inference, but the weight is the same in either case” (2008, p. 313).

Magnus argues by analogy with systems of deductive proof. Take first a natural deduction system which has lots of rules of inference but no axioms. Contrast that with an axiomatic system which has just one rule and a great many axioms. Such systems license the same proofs, so there is in a sense no formal difference between them. Any proof in one can be translated into a proof in the other. The choice between a natural deduction system and an axiomatic system cannot be made on the basis of formal power, but instead must be made on practical grounds. A natural deduction system may be deemed to better represent the meaning of the logical operators, for example, while an axiomatic system might be easier to automate. Neither system is the right one *tout court*, but each might be more useful than the other for some purposes. To take an example: There is a secure inference from $(P \& Q)$ to Q , so we can ask what licenses or secures the inference. In an axiomatic system, we might need to write down an axiom $(P \& Q) \rightarrow Q$ and perform modus ponens. In a natural deduction system, it is a direct step — licensed by the conjunction elimination rule, a rule that is in the metalanguage rather than in the object language of the inference. The rule is part of our account of proper inferences, rather an axiom in the proper inference itself. There is nothing about the inference itself which highlights one system as correctly capturing the move's security or legitimacy.¹⁰

To apply the analogy: Material *postulates* are like axioms, statements or claims which arguments explicitly invoke. Material *inference principles* are rules, steps that can be taken without writing them down explicitly. Just as equivalent deductive systems can be constructed which place different emphasis on axioms

¹⁰ Actual deductive arguments (for example, in mathematics) are rarely made with the precision and detail required by a formal proof system. When mathematicians make arguments, it is possible to explicate their arguments in a proof system which uses axioms or one which uses rules. Although there may be some cases in which a mathematician determinately has one in mind rather than the other, they usually do not— when they do not, there is no fact of the matter as to which system they are employing.

or rules, equivalent ampliative inferential practices might place different emphasis on postulates and inference principles. The difference is practical, and different emphases might be appropriate for different purposes.

The assumption that every ampliative inference relies on a hidden premise is not simply a view about scientific inference. As Robert Brandom points out, it is commonly believed that endorsing an inference from A to B always depends on belief in the conditional *If A, then B* (Brandom 2000, p.53). He gives the example of the inference from “It is raining” to “The streets will be wet” which the common view takes to be legitimate only if there is a suppressed premise in the form of a conditional (“If it rains, then the streets will be wet”). Brandom rejects this, holding instead that material inferences are fundamental. As he sees it, “formal logical vocabulary” such as the conditional plays an “explicating role”. Such vocabulary “permits the formulation, as explicit claims, of the inferential commitments that otherwise remain implicit and unexamined in the contents of material concepts” (Brandom 2000, pp.70-71). Brigandt and Brandom are committed to a throughgoing inferentialism, according to which grasping a concept “is mastering its *inferential* use: knowing ... what else one would be committing oneself to by applying the concept, what would entitle one to do so, and what would preclude such entitlement” (Brandom 2000, p.11). Yet even inferentialists must admit, at least for some concepts, that it is possible to make all of the relevant commitments explicit, and in those cases, deductive rules would capture all legitimate inferences.¹¹ So we need not decide for or against inferentialism in order to maintain that *some* arguments could be rendered to rely merely on deductive or quasi-deductive rules, as Norton recommends.

Because an argument presuming material postulates and a parallel argument employing material inference principles sanction the same conclusions, it might seem as though there are no grounds for choosing to express an argument in one way or the other. They are equivalent in terms of “the sort of rationality that... consists in making the appropriate inferential moves” (Brandom 1994, p. 106). Yet there may be context-dependent, practical reasons for choosing one or the other mode of expression. This suggests that we should not see the choice between postulates and inference principles as being timeless or binary. In the early days of an enquiry, we might exploit metaphors that we cannot make fully explicit. Part of the progress of enquiry involves spelling out precisely how the metaphor holds and how it does not. Ultimately, the metaphor may be exhausted and replaced by explicit statements about the target domain. In the early period of enquiry, Brigandt’s account may be most apt. Scientists rely on inferential principles without knowing precisely when they hold and when not, without being able to fully spell them out explicitly as postulates. In the later period, Norton’s might fit better. Still later, perhaps, there might be no further need to make the premises explicit as they recede into common knowledge. (Practicing scientists know that the melting point of an element is constant through time without having to make that an explicit premise of their reasoning.) The introduction of fruitful metaphors which are ultimately replaced by literal

¹¹ Brandom claims that it is an advantage of his approach that it shows how formal inference is possible, relative to a distinction between logical and non-logical vocabulary (Brandom 1994, pp. 104-5).

descriptions is, arguably, a general pattern in science (Boyd 1979). However, it suffices for our purposes here if it occurs even *sometimes*. Different explications work better for different moments of enquiry.

To take a different example: At the outset of enquiry, scientists might rely on regularities which they are only beginning to understand. They can state them only roughly, with crude *ceteris paribus* clauses to mark the indefinite limits of how far their statements can justify predictions. So the predictions they make are justified more by their practices of inference than by the explicit statements of regularity — that is, more by the material inference principles they enact than by material postulates which they state. As enquiry progresses, they learn more about when and where the regularity holds, and *ceteris paribus* clauses becomes less conspicuous and premises become more explicit. Philosophers of science disagree about whether true laws of nature could ever be written as categorical conditionals, rather than as ones that hold for the most part, but there are important moments in enquiry when it is premature to expect everything to be explicit and moments when it would be obscurantist to refuse to elaborate on details.

The upshot is that, for scientific inference, one cannot say simply that there is an inferential gap or that there is not. Whether it is sometimes permissible to generalize without further premises from samples about this kind of stuff to all instances of this stuff depends in part on how we express our system of inference and how we conduct our enquiry. A gap thesis in science is not a strictly logical matter.

The idea of a gap, as a matter of in-principle justification, drops away when we see that material postulates and material inference principles might license the same conclusions. Instead, gaps are a matter of how we regiment inferential practice. Maintaining or rejecting a gap is a practical choice which should be made based on costs and benefits rather than a factual determination of how the world is. Moreover, the choice need not be resolved once and for all. We will, at different times and in different contexts, maintain or collapse inferential gaps.

4. Lessons for the is-ought gap

What we saw in the previous section was that, for any inference system constructed using material postulates, an equivalent system could be constructed using ampliative rules instead. The considerations which underwrote that equivalence did not depend on the conclusions being general, being about the future, or being about unobservable entities. The same follows for inferences from facts to normative conclusions. We can understand these inferences as depending on postulates (explicitly normative premises which yield normative consequences when combined with other, factual premises) or principles (rules of inference that lead from factual premises directly to normative conclusions). Where Norton says that ampliative inference in science requires material postulates, the parallel gap thesis is that inference from facts to values requires explicitly normative postulates. This is what Dworkin and other defenders of the is-ought gap thesis hold. But if we follow Brigandt's suggestion and accept material inference principles, then there would be no such gap. We would authorize inferences from factual to normative claims, just as practicing scientists infer claims about the future from the past without explicit postulates. To repeat: if we *choose* to explicate such inference in terms that require postulates, then the

postulates needed to support inferences to normative conclusions will themselves be normative. Such an explication *introduces* rather than discovers the gap. If we insist that the normative work must be done by explicit premises, then the gap thesis holds. But we also have the option of allowing the normative work to be done by material inference principles that authorize descriptive-to-normative inferences, in which case the gap thesis does not hold.

We can understand our discourse and inferential practices as having an is-ought gap, or we can understand it as not having one. We can regiment our inferences by relying on more expansive inferential rules or by reducing our inferential rules (at the limit to deductive inferences alone) and insisting on making more normative postulates explicit. This choice has no immediate consequences for metaphysics, for whether facts and values are different kinds of things. Ampliative inferences in science can hold between things of the same kind (as when we infer from observed cases to as-yet unobserved ones) or between ontologically different things (as when we infer from observations to the existence of unobservable entities), so the general structure of the case is compatible with facts and values being the same or with them being different. Whether we make normative inferences by employing ampliative inference principles or by introducing logically-contingent postulates is a practical matter of how we explicate our inferential practices. The practical matter depends not on mere logic or ontology but instead on the interests at stake, on the risks and benefits of various possible outcomes, and on our cognitive resources and powers. It also depends on the social context of enquiry, the discourse to which the inference contributes.

Consider a homey example, akin to the one which Dworkin entertains (sec. 1, above): “Nate is about to unwittingly step into an open pit. Therefore, I ought to warn him.” Like Searle’s argument JS, the inference is defeasible. Perhaps I am in no position to warn him, perhaps I know that he enjoys falling into pits, and so on. Nevertheless, I might do a better job of saving my friends from injury if I do not stop short and consider what further premises might be supplied to the argument in order to justify the conclusion. It would be better both for me and for Nate if I make the inference immediately, by a rule which gives a positive but defeasible presumption in favor of warning people who are about to unwittingly step into pits. It would be better for me not to pause to consider the truth of the postulate (and its possible exceptions) that would make this inference into a formally valid deduction — a postulate something like: “If you see someone about to fall into a pit, one should warn them.”

Suppose I yell out to Nate, and a third party asks why I did so. I tell them that I had to do it (it was morally required of me) because I saw that Nate was about to unwittingly step into an open pit. In normal circumstances, that is a sufficient justification. But notice that on this rendering, the facts alone underwrite the normative claim (that I should warn him). In normal circumstances, this inference is both perfectly acceptable and generally known to be so. However, we can imagine circumstances in which it is not. (Philosophers are especially adept at playing this game.) Suppose the observer comes from a land where falling into pits is considered good luck or is not dangerous. Then, if challenged as to why the inference was a good one, I might represent it as a premise: “If you see someone about to fall into a pit, you should warn them (*ceteris paribus*).” By making the conditional explicit, we can then examine whether it is true. I can point out the value of helping others avoid danger, and

they can point out the loss of luck-enhancing opportunities. In the more ordinary case, in which the fact of Nate's imminent fall is sufficient to justify an obligation to call out to him, we can represent the shared background agreement either as acceptance of the conditional or as acceptance of the inference. Insisting that it *must* be represented as acceptance of a shared premise rather than acceptance of a shared inferential pattern seems to be the product of a philosopher's fixation on deductive inference. When he considers his analogous example, Dworkin acknowledges: "Yes, the fact that someone standing before you is evidently in great pain does seem by itself a reason why you ought to help him if you can. Nothing more needs to be said." But, contrary to appearances, Dworkin insists that this inference is only valid because "you unselfconsciously accept, as something that goes without saying, a general responsibility to help people in grave need when you easily can" (Dworkin 2011, p.45). Although we agree that the inference *can* be validated through a general postulate, Dworkin gives us no reason to think it must be or even that in ordinary circumstances there are advantages to doing so.

Despite the equivalence in terms of which actions are ultimately licensed, there is a difference in implementation between what we do when we make an explicit assumption and what we do when we adopt a rule. Scientists who offer quasi-deductive arguments with explicit material postulates are doing psychologically and sociologically different things than those who rely on inexplicit rules, even though they can come to license precisely the same conclusions. Similarly, the choice between presenting a quasi-deductive argument for a normative conclusion that relies on an explicitly normative premise and inferring a normative conclusion by relying on a normative inference principle depends not on ultimate metaphysical commitments or timeless logic but instead on context, the extent of shared background assumptions, the assessment of various risks, and so on.

The point of the previous section was that inference in science involves some risk of being wrong, and the same risk might be embodied in material postulates as in ampliative procedures. Similarly, the inference I make that I should warn Nate involves risk. In the comfort of a classroom, we can spell out reasons why the inferential tendency is a good one or fanciful circumstances in which it would not be, but this is a luxury rather than a requirement for accepting the argument. For any ampliative inference, we can ask whether the inference pattern we are relying on is normatively appropriate and replace the inference principle with a material postulate. This explication can in principle yield a deductive inference in the limit when all ampliative rules are replaced with explicit conditional postulates, but that limit is only rarely if ever achievable. Our lives are stitched together from beliefs about contingent matters gained by ampliative inference, and explication and interpretation must come to a stop somewhere.

We have argued that having an is-ought gap just means insisting that any inferences with normative conclusions must be spelled out with explicit normative postulates rather than relying on ampliative inferences to move from descriptive premises to a normative conclusion. Our homey example suggests that in uncontentious cases, ordinary normative justification can often do without an is-ought gap. What about contentious cases where sufficient background agreement is not in place? In such cases, it may be valuable to elaborate, even if just in crude terms, why I ought to warn Nate about the danger of the pit. The

inference in the homey example is not underwritten by a *sui generis* duty to warn Nate about pits, after all.¹²

As a general matter, sometimes an is-ought gap is appropriate and sometimes not. Enforcing the gap thesis insulates discussion of facts from questions of norms and values. It creates discourse which is at least superficially value-free. Although the notion that science is thoroughly value-free is false, the relative isolation established by a gap can serve a valuable purpose. Helen Longino writes, “The myth of scientific value neutrality... is... a functional myth. It clears the way, conceptually, for the elaboration of a particular approach to a set of phenomena once that approach has attracted the consensus of a significant portion of the relevant scientific community. This makes it possible to adjust a framework... to the brute facts” (1990, pp. 224-5). When scientists enforce an is-ought gap, they prohibit inferences to moral conclusions without explicitly introducing moral premises. And when they remain agnostic about such moral premises, they create a space for empirical investigation relatively free from normative disputes. By isolating an area of inquiry, scientists reap the significant rewards associated with specialization.

This is illustrated in the approach to risk analysis which divides science related to policymaking into two separate parts (see Douglas 2009, pp. 137-148). In *risk assessment*, scientists collect evidence and report on the nature and extent of a risk. In *risk management*, policymakers decide what to do about the risk. This is meant to insulate the science from the politics, and the insulation comes by way of an is-ought gap. Only descriptive considerations enter into risk assessment, and normative claims are introduced explicitly by policymakers in risk management. However, assessing the nature and extent of a risk involves judgments about which risks are *significant* and what is a *reasonable* margin of safety or level of confidence. Because significance and reasonability are normative, this threatens to undo the separation and introduce values into the moment of risk assessment. In order to maintain the separation, the model requires that there be explicit inference guidelines which specify how the judgments in risk assessment should be made. For example, an inference guideline might specify that male rats will be used as the model organism for estimating toxicity. Organizing risk analysis in this way has advantages. It is transparent and procedurally objective. Businesses know what to expect from the process, while the risk assessment is protected from business interests or political agendas in particular cases. However, this comes at the cost of rigidity. The details of particular cases may make the inference guideline inapplicable. Douglas gives the example of chemicals which cause cancer in male rats by a mechanism that does not obtain in female rats or in humans. She concludes, "An inference guideline that required the inference from rat cancer to human cancer would ignore important scientific evidence; an inference guideline that considered the additional evidence would require a case specific judgment... thus undermining the key advantages of the guideline

¹² Elaborating the general principles which justify warning him is, of course, inappropriate in the moment. If I take too long to puzzle about it, I will not act in time. But in one's leisure before travelling to a distant land, where norms for warning strangers of danger (and much else) may differ, it might make sense to review explicitly such differences.

(consistency, transparency, efficiency, and predictability)..." (2009, p. 145). Case specific judgments introduce values in assessments of what is significant, reasonable, and so on. Structuring risk analysis in line with an is-ought gap has advantages but also comes at a cost, and the decision of whether to enforce such a gap depends on the evaluating these costs and benefits.

Enforcing a sharp distinction between science and values maintains not only the is-ought gap but also what we might call the ought-is gap. It proscribes both inferring from descriptive claims to normative claims but also discussing normative claims when considering the justification for descriptive claims. This can make it difficult or impossible to recognize the value-sensitivity of the dominant framework and conclusions. Feminist philosophers of science have documented these effects at the intersection of science and gender. Sexist normative assumptions have often structured biomedical enquiry, for example, in a way that shapes the evidence scientists collect, the inferences which they draw, and the range of alternative hypotheses which they consider. Trying to settle these factual questions independently of normative ones hides important normative presumptions. Maintaining an is-ought gap makes it seem as if values are entirely another matter, separate from the factual questions which scientists were addressing. (Longino 1990, Shiebinger 1999) Although a monolithic myth that science is value-free is pernicious, it may nevertheless be sensible in some contexts to construct and maintain these gaps. Consider uncontentious science which has no direct consequences for policy, like astronomy. Whether 'Pluto' should count as a *planet* should be decided based on the structure of the solar system rather than based on our fondness for it or its association with beloved Disney characters.¹³

Our central claim does not depend on what model of risk analysis is ultimately preferable. Rather, what matters is that the outcomes turn on details about the costs and benefits of maintaining the gap. It is a practical difference, rather than a logical one. The question is not whether there is an is-ought gap tout court but instead when and for which enquiries we ought to make there be an is-ought gap. There will often be competing considerations, and particular enquiries might go either way.

Works cited

- Richard [N.] Boyd. Metaphor and theory change: What is "metaphor" a metaphor for? In Andrew Ortony, editor, *Metaphor and Thought*. Cambridge University Press, Cambridge, 1979.
- Robert B. Brandom. *Making It Explicit: Reasoning, Representing, and Discursive Commitment*. Harvard University Press, Cambridge, Massachusetts, 1994.
- Robert B. Brandom. *Articulating Reasons: An Introduction to Inferentialism*. Harvard University Press, Cambridge, Massachusetts, 2000.
- Ingo Brigandt. Scientific reasoning is material inference: Combining confirmation, discovery and explanation. *International Studies in the Philosophy of Science*, 24(1): 31-43, March 2010.
- Rachel Cohon. *Hume's Morality: Feeling and Fabrication*. Oxford University Press, Great Clarendon Street, Oxford, 2008.

¹³ Regarding the characterization of planets, see Magnus (2012, ch. 3, sec. A).

- Heather E. Douglas. *Science, Policy, and the Value-free Ideal*. University of Pittsburgh Press, 2009.
- Ronald Dworkin. *Justice for Hedgehogs*. Harvard University Press, Cambridge, Massachusetts, 2011.
- David Hume. *A Treatise of Human Nature*. Oxford University Press, 1978. Edited by L.A. Selby-Bigge and P.H. Nidditch.
- Helen Longino. *Science as Social Knowledge*. Princeton University Press, Princeton, New Jersey, 1990.
- P.D. Magnus. Demonstrative induction and the skeleton of inference. *International Studies in the Philosophy of Science*, 22(3): 303-315, October 2008.
- P.D. Magnus. *Scientific Enquiry and Natural Kinds: From Planets to Mallards*. Palgrave MacMillan, Basingstoke, Hampshire, 2012.
- John D. Norton. A material theory of induction. *Philosophy of Science*, 70(4): 647-670, October 2003.
- Charles R. Pigden, Logic and the autonomy of ethics. *Australasian Journal of Philosophy*, 67(2): 127-151, 1989.
- Charles R. Pigden, editor. *Hume on Is and Ought*. Palgrave MacMillan: Basingstoke, Hampshire, 2010.
- Charles R. Pigden. Hume on Is and Ought: Logic, Promises, and the Duke of Wellington. In Paul Russell, editor, *The Oxford Handbook of Hume*. Oxford University Press, 2016.
- A.N. Prior. The autonomy of ethics. *Australasian Journal of Philosophy*, 38(3): 199-206, 1960.
- Greg Restall and Gillian Russell. Barriers to implication. In Pigden 2010, pages 243-259.
- Gillian Russell. In defence of Hume's law. In Pigden 2010, pages 151-161.
- Londa Schiebinger. *Has Feminism Changed Science?* Harvard University Press, Cambridge, Massachusetts, 1999.
- John R. Searle. How to derive "ought" from "is". *The Philosophical Review*, 73(1): 43-58, January 1964.
- Wilfred Sellars. Inference and meaning. *Mind*, 62(247): 313-338, July 1953.
- Daniel J. Singer. Mind the Is-Ought Gap, *Journal of Philosophy*, 112 (4): 193-210, April, 2015.