

Forthcoming in *Naturalism and Analysis*, Robert Nola and David Braddon-Mitchell (eds), MIT Press, December 2008.

## **Jackson's Armchair: The Only Chair in Town?**

**by Justine Kingsbury and Jonathan McKeown-Green**

### **0     Nothing-over-and-above, conceptual analysis and the Jackson thesis**

Are all the facts about nations, cultures and economies really just facts about people's mental states and their interactions? Are all of the properties which determine whether or not a thing is a work of art really just physical properties of that thing? Is linguistics, the scientific investigation of language, best understood as a branch of psychology, the scientific investigation of the mind? Can psychology be reduced to biology? Can all biological phenomena be explained chemically? Is chemistry really just part of physics? Is there anything going on in the world which isn't a physical thing? Can there be freely-chosen, autonomous human action in a purely physical world? Frank Jackson has made a controversial claim about the way in which one should investigate questions like these. This paper is a qualified defence of that claim.

The questions are examples of *nothing-over-and-above*: each of them asks whether the subject-matter of one human classificatory practice or discipline is anything over and above the subject-matter of some other, putatively more inclusive, more fundamental, or better-understood human classificatory practice or discipline.

Lots of scientists and philosophers go in for nothing-over-and-above<sup>2</sup>ry. Because of where we are heading, it is worth noting that nothing-over-and-above<sup>2</sup>ry is a characteristic preoccupation of a certain breed of boffin, the *serious metaphysician*, as Frank Jackson calls her (Jackson 1998, 4-5). A serious metaphysical project is one that aims to produce, or at least explain why in principle it is possible to produce, a complete inventory of all the things or phenomena in some specified class by appealing to some restricted class of more primitive things or phenomena. For example, suppose your preferred ontological primitives are all physical. Then the job of the boffin who is interested in the mental and the social is to catalogue and theorise about everything mental or social in purely physical terms, or at least to argue persuasively that this could be accomplished. The boffin's challenge is to explain away the appearance that mental and social facts belong to a different ontological category: one that requires us to believe in minds and cultures *in addition* to atoms and forces. She is engaged in nothing-over-and-above<sup>2</sup>ry. Jackson suggests that this challenge of locating one subject-matter within another is what makes serious metaphysics both interesting and serious.

Nothing-over-and-above<sup>2</sup>ry is one of the two central strands in our discussion. The other is *conceptual analysis*. Conceptual analysis is what I do when I reflect on whether rules can have exceptions, whether robots can be moral agents or whether a fully-scored chamber work for cellos and alto saxophone gets to count as jazz if it is written by the African-American composer Henry Thredgill. It is what German mathematicians did in the late nineteenth century when they thought about infinity and noted that two infinite collections of objects, unlike two finite ones, can be the same size even when one is only a subcollection of the other. We indulge in conceptual analysis when we reflect fairly systematically on some notion - like rules, moral agency, jazz, or infinity - which forms part of some human practice. Systematic reflection means, roughly, deciding which

possible scenarios are scenarios in which we have something to which the notion applies,<sup>3</sup> which possible scenarios are ones in which the notion is not applicable, and which scenarios are ones in which it is indeterminate. Thus, we distinguish cases of jazz from cases of non-jazz and both from borderline cases. As Grice put it:

If I philosophize about the notion of cause, or about perception or about knowledge and belief, I expect to find myself considering, among other things, in what sort of situations we should, in our ordinary talk, be willing to speak (or again be unwilling to speak) of something as causing something else to happen; or again of someone as seeing a tree; or again of someone as knowing rather than merely believing that something is the case. (Grice, 1958, 172.)

In line with this approach, Jackson explicitly takes conceptual analysis to be *the method of possible cases* (Jackson 1998, 31-32.)

Jackson has written a lot about philosophical methodology in the last decade. His efforts culminate, but do not terminate, in *From Metaphysics to Ethics: A Defence of Conceptual Analysis* (1998). He argues for the following thesis.

**Jackson's Thesis: Conceptual analysis is part and parcel of any nothing-over-and-above. In other words, the method of eliciting intuitions about possible cases is standard issue whenever we try to redescribe the subject-matter associated with one kind of human classificatory practice in terms of the subject-matter associated with another, more inclusive, classificatory practice. A fortiori, conceptual analysis is part of what needs to be done by the serious metaphysician: it is part of what you need to do to give a complete account of**

**what sort of stuff there is in terms of an austere package of ontologically primitive descriptors.<sup>1</sup>**

4

This thesis has attracted a lot of attention. Some people have written about its role in arguments for particular philosophical claims. For example, David Chalmers, and in former times Jackson himself, have appealed Jacksonianly to intuitions when arguing that qualia cannot be explained physicalistically, so we often find criticisms of Jackson's thesis tangled up with criticisms of non-physicalism about the mental.<sup>2</sup> Commentators on Jackson have also seized upon the fine details of his approach to conceptual analysis and especially on his views about the nature of possibility and the a priori, which invoke the apparatus that Krister Segerberg (1973) called *two dimensional modal logic*.<sup>3</sup>

We will argue that Jackson's Thesis is correct, but that his own formulations of it misdescribe conceptual analysis. Jackson, we think, wrongly assimilates conceptual analysis to the semantic analysis of ordinary language. We will keep the discussion very general and very methodological. We will not go into detailed case studies about the mind, the colours, ethics or even jazz. We will also not engage with the details of Jackson's two-dimensionalism. We will argue for Jackson's thesis by presenting a series of objections coupled with replies on Jackson's behalf. The objections are distillations of criticisms which can be found in the literature, but we use them primarily as dialectical tools for elaborating Jackson's position and clearing up misunderstandings. Initially, our replies will be compilations or restatements of points that Jackson himself makes. Soon enough, though, our replies will exceed or depart from Jackson at least in emphasis. We cannot promise that Jackson would endorse our replies to the later objections. We will end with the promised worry of our own - not about Jackson's Thesis as such, but about his preferred formulation of it. As noted above, we believe that Jackson misconstrues the nature of

conceptual analysis. There will be no reply to this final objection since we think it succeeds.

Before we present the objections, a brief terminological digression for those and only those who are uncomfortable with our introduction of "nothing-over-and-abovey." Traditionally, 'reduction' and 'reducibility' have been used where we use 'nothing-over-and-abovey'. But which examples of nothing-over-and-abovey count as speculations about reducibility is controversial. Typically, *reducing* a theory T of the subject-matter of some practice to some other theory T+ is taken to require that two criteria be met:

1. All the laws, equations or other principles in T are derivable from or entailed by laws, equations or principles in T+, perhaps in conjunction with intertheoretic bridge principles.

2. Every item in the subject-matter of T is identical or nomically coextensive with an item, or is constituted by items, from the subject-matter of T+.

However, there are theories T and T+ for which at least one of these criteria arguably fails and yet the contention that the subject-matter of T is nothing-over-and-above the subject-matter of T+ remains a live option. Consider, for example, the relationship between the mental and the physical. Davidson's anomalous monism violates 1. Davidson (1970) says that there are no psychological laws and no laws linking psychological generalisations with laws of physics. Yet Davidson maintains that every mental event is a physical event and this makes him a nothing-over-and-aboveyist. Of course, it also means that he endorses 2, but his violation of 1 makes him a non-reductionist in the eyes of many. A variation on his view shows us how both the criteria for reduction can be violated at once consistent with nothing-over-and-abovey. Because it is an anomalist view, it violates 1. However, this variation allows that there is no single physical state, event or cluster of states or events which is identical or coextensive with or which constitutes my belief that *strong leadership*

*is wrong leadership*. To say that I have this belief is to report on certain consequences of <sup>6</sup> facts about my total mental state, namely actual or possible interactions between my brain, certain input stimuli and certain output behaviours. Hence this anomalist holist view also violates 2. (To make the violation even more extreme, make the holism highly externalist: whether my total mental-state/mental-processing is such that I believe that *strong leadership is wrong leadership* depends both on features of my present environment and on the details of how my brain came to be in its present configuration.) All of this is compatible with the claim that the mental is nothing-over-and-above the physical. Our holist (externalist) anomalous monist is a nothing-over-and-aboverist.<sup>4</sup>

Because the requirements on reduction are typically taken to be more stringent than those that concern us in what follows, we will stick to our own term, 'nothing-over-and-abovery.' If you dislike it, you are welcome to substitute 'ontic dependence', 'ontic determination', or even 'reducibility'. If you have a developed account according to which the idea can be captured in terms of supervenience theses, then make yourself at home.

Here the digression ends and the objections begin.

## **1 First objection : The record of forays into conceptual analysis is dismal**

It is well-known that attempts to capture the rules we implicitly follow when we attribute *knowledge* to somebody, *meaning* to an utterer or *causation* to a sequence of events have run out of steam because they have failed to produce satisfactory analyses. Famously, Edmund Gettier (1963) challenged the conceptual analysis of KNOWLEDGE as justified true belief. He exhibited counterexamples to it and these prompted new definitions which generated new counterexamples. Nobody has yet proposed an analysis which satisfies the philosophical community. Similarly, Paul Grice (1957) offered an account of what it is for

an agent  $S$  to mean that  $p$  on a particular occasion when uttering  $x$ . This reconstruction of Grice's original definition comes from Schiffer 1987, 243: <sup>7</sup>

$S$  means that  $p$  in uttering  $x$  iff, for some person  $A$  and feature  $\Phi$ ,  $S$  intends

(1)  $A$  to recognise that  $x$  has  $\Phi$ ;

(2)  $A$  to think, at least partly on the basis of thinking that  $x$  has  $\Phi$ , that  $S$  uttered  $x$  intending  $A$  to think that  $p$ ;

(3)  $A$  to think, at least partly on the basis of thinking that  $S$  uttered  $x$  intending  $A$  to think that  $p$ , that  $p$ .

Like the justified true belief analysis of KNOWLEDGE, Grice's already convoluted definition of speaker meaning (as it came to be called) fell prey to sundry counterexamples, many of them Grice's own, and spawned increasingly long-winded definitions until well into the 1980s. Likewise, David Lewis' (1973) counterfactual analysis of causation generated increasingly complicated definitions with more and more epicycles. <sup>5</sup>

The definitions that typified these projects appeared more and more unnatural as time passed and the counterexamples kept on a-comin'. We can argue by induction that there is no point in pursuing these projects any further and we can speculate that something about the way our minds operate or about the nature of the concepts means that definitions just are not available for them. The prospects for successful analysis seem equally grim in other areas of philosophy, so we should conclude (albeit tentatively) that this lesson extends to analyses of BELIEF, DESIRE, THE GOOD and the rest. We should give up on conceptual analyses of the kind that have filled the journals (especially *Analysis*) with definitions and counterexamples. Instead, we should channel our energies into working out what it is about the mind or the concepts that gives them this elusive character. If all of this is right, then Jackson is wrong to hold out any hope that conceptual analysis can pay

dividends as a tool for assimilating one subject matter to another, or as a tool for anything<sup>8</sup> at all.

## **Reply**

Note first that much successful conceptual analysis occurs quietly out of the limelight. If I tell you that you are doing alright because your home is worth \$700,000 you might well complain. On analysis, the claim that your house is worth \$700,000 means, roughly, that if you sold it now you could expect to get about \$700,000. But you don't want to sell it. Your rates have gone up. You're not doing alright at all.

Some successful conceptual analysis is not so work-a-day. An example of this is our earlier one from mathematics of successful (albeit partial) efforts to tame the concept of the infinite within a tractable mathematical theory.

True, the projects which have flooded the pages of the philosophy journals have not been completed successfully. But they do concern philosophically intractable notions and we shouldn't expect the job to be easy. It is precisely because these projects are long and hard that they take up so much space in the journals.

Complete analyses of even these hard philosophical chestnuts must be available in principle. After all, if you describe a possible scenario to your informant in enough detail, she can tell you whether the concept as she uses it applies or not or whether she is unsure. Tell her a story about something painted by a robot and hung in an art museum and she *will* have something to say, however noncommittal, about whether or not it's art - even if all she says is that she's not sure. Tell her a story about a planet where the clear, drinkable liquid that fills the oceans is not H<sub>2</sub>O but something else and she will have something to say about whether or not this liquid is water or why we can't decide. Indeed, the reason we

know that most of the post-Gettier definitions of knowledge have counterexamples is because our informants can identify cases where the definition is satisfied but where there is not knowledge, or vice versa. All of this means that there is in principle a story to be told about our informant's responses to all possible cases and this story is the right analysis of the concept. It will say, for every one of infinitely-many possible scenarios, the degree to which one takes it that the concept applies in that scenario. Not an elegant definition, maybe, but an analysis of the concept as our informant employs it nonetheless.

Although any complete analysis of a concept has to be a statement of necessary and sufficient conditions for its application - a definition, if you like - it might not be a definition of the sort one would naturally wish for. For a start, we should not be surprised if some of the rules we implicitly follow when classifying things as jazz or non-jazz, right or wrong, art or non-art, water or non-water, are horrendously complicated and gerrymandered. It would be marvellous if the quest for a conceptual analysis often terminated in a more or less pithy statement of necessary and sufficient conditions like Grice's candidate for a definition of awe: 'a mixture of fear and admiration' (Grice 1958, 176). Unfortunately, however, our informal, everyday classificatory practices evolve over time in response to changing circumstances and are not subject to careful review by the philosophy department or the Crown Law Office. The question of whether to include a particular musical event among the jazz may well depend on vagaries: what were the journalists working for *Downbeat Magazine* most interested in during the late '50s? What kinds of music did most of the jazz players from the thirties move on to when the big bands disbanded? Perhaps the best definition of the concept JAZZ will be a motley disjunction of conditions. And perhaps the best definition of the concept MEANS will be of the long-winded Gricean type illustrated above with exceptions, caveats and epicycles on-board.

We should also not be surprised if a complete analysis is noncommittal on certain<sup>10</sup> matters. There are plenty of ordinary users of WATER who genuinely have no view about whether XYZ counts as the water on Twin Earth, and maybe a correct analysis of WATER (as the concept is used by those speakers or by the whole community) should reflect this.

These points about what we can expect from a definition suffice to undermine the usual anti-analysis complaint that increasingly complex definitions spell increasing likelihood that the search for a complete account is faltering. Hedged, complicated analyses might be just what we should expect.

Nevertheless, our objector might rejoin, there's something wrong somewhere with the yarn we have been spinning. There must, she might insist, be some concepts for which complete analyses are in principle unavailable. Otherwise, an unthinkable infinite regress of definitions looms - unthinkable because of natural language's expressive limitations.<sup>6</sup>

To meet this worry, we concede that there might well be concepts which resist complete analysis in languages with finite expressive resources and our spade may indeed be turned by some of these if we keep defining complex concepts in terms of simpler concepts for long enough. But consider one more point about completed analyses. We should not be too surprised if some completed conceptual analyses resist formulation in finite languages. Suppose we are analysing mental states. We ask, what is it for something to count as the belief that *strong leadership is wrong leadership*? We ask, in other words, which of the infinitely-many possible states of the world are those about which an ordinary, folksy, wielder of the BELIEF concept would claim that I hold this belief? If we take seriously the idea that conceptual analysis is the method of reflecting on (the infinitely-many) possible cases, we may discover that only an infinitely-long list of the possibilities to which the concept applies captures the rule governing the concept's use. We might hope for better than this, but it may not always be obvious that optimism is warranted. This

recourse to infinitely long lists of conditions will surely apply too in the case of any concepts whose unpacking would take us beyond the expressive resources of finite languages.

Since completed analyses need not yield satisfying, quotable sound-bites, conceptual analysis must not be identified with what Stephen Stich calls the 'method of proposing definitions and hunting for intuitive counterexamples' (Stich 1993, 354). To be sure, a fun way to elicit intuitive judgements about whether a concept (JAZZ, for instance) applies to a possible case or not is to entertain a definition of it and search for counterexamples. However, conceptual analysis, as Jackson understands it, is the process of reflecting on possible cases and this need not, though it may, consist in auditioning, testing, refining, rejecting and replacing definitions.

Our original objector has one more worry. If complete analyses of many philosophically interesting concepts are so elusive, what's the point of the enterprise? Why go in for analysis if there's often little hope of finitary completion?

Our reply: We don't always need a complete analysis of a concept in order to learn something worthwhile about it from the business of analysing it. Failure to net the necessary and sufficient conditions for the application of some concept to a circumstance need not spell a failure to learn more about the application of the concept. It might still be of interest to uncover, as we have done in the case of knowledge, necessary but insufficient, or sufficient but unnecessary, conditions for its application, and we should not regard the unavailability of even provisional definitions of some concept as a sign that we are getting nowhere with the analysis of it. Philosophers probably know much more about knowledge as a result of the discarded analyses thrown up in Gettier's wake. Much the same can be said about our insights into meaning, causation and especially infinity which owe so much to incomplete or downright incorrect analyses.

In sum, conceptual analysis, regarded as the probing of our intuitions about the application of some notion in various possible scenarios, is alive and well: it makes unobtrusive contributions to scientific, philosophical and ordinary reasoning. It pays dividends even when it does not converge on an exceptionless criterion of application for a concept. So for all our objector says, it is available to play its allotted role in Jackson's programme.

## **2 Second objection: Armchair reflection is irrelevant to finding out what the world is like**

There is something odd about the idea that one can sit, secluded from life's laboratory, and make substantial progress towards discovering what the world is like or what sort of stuff it contains. This is what we learned when we finally realised that natural science relies primarily on a posteriori discoveries. To flesh this out a bit (and this is how Jackson himself puts the objection in Jackson 1994, 2) the appeal to intuitions about possible cases is *doubly* irrelevant to finding out what the world is like. First, it is an appeal to intuitions, not to facts. Second, it is an appeal to the possible, not the actual. We could augment this objection by recalling, third, that Kripke (1980) and Putnam (1975) made a big deal of the fact that much that is necessarily true - such as the fact that gold is an element whereas salt is a compound - is learnable only a posteriori. In general, access to the nature of things is not a priori access.

## **Reply**

Part of what it takes to discover whether there are Ks and what they are like is determining what it would take for us to regard something as a K. This latter is a job for the conceptual analyst. To explain: A serious metaphysician, or any scientist, if it comes to that, wants to know whether the world contains certain things - mesons, gods, minds, values, world souls, or whatever. She also wants to know what these things are like. Hers is the project of discovering whether there are any Ks, which things (if any) are the Ks and what (if anything) the Ks are like. However, the methodological starting-point for any such project is the question "Under which circumstances would we judge that there are Ks?" Only by shedding light on this question can we proceed, since only then can we go out into the world and discover whether any of the circumstances in which we would judge that there are Ks obtain. And only then can we say whether there are Ks, which things are the Ks and what the Ks are like. This does not mean that we must completely settle the question of what we would regard as a K before we write the surveys and perform the experiments. The two complementary projects - a priori conceptual analysis and a posteriori empirical examination - might interlock and co-occur. If conceptual analysis has any priority over empirical investigation, it is logical or methodological, not temporal.

Our objector will already be restless. It is not obvious that we need to know what we would regard as a K, as a god, for example, in order to look for one. Maybe we need to know what we would regard as God in order to go out and look for the sort of thing that we would regard as God. But a search for God is different from a search for what we regard as God. We are interested in whether there are real gods, real mesons and so on. This is not the same as being interested in whether there are items that we would regard as gods or as mesons.

Here's the response. We can't embark on a search for Ks unless we have some idea what we are looking for. As Jackson says:

When bounty hunters go searching, they are searching for a person and not a handbill. But they will not get very far if they fail to attend to the representational properties of the handbill on the wanted person. These properties give them their target, or, if you like, define the subject of their search. Likewise, metaphysicians will not get very far with questions like: Are there *Ks*? Are *Ks* nothing over and above *Js*? And, Is the *K* way the world is fully determined by the *J* way the world is? in the absence of some conception of what counts as a *K*, and what counts as a *J*. (Jackson 1998, 30-31)

Sometimes what we are looking for will be obvious, and we hardly need to glance at the handbill. If it's marbles, then we already know what marbles are, so we can just go out there and look for them. If it's mesons, well, we might have to consult our scientific theory, but often we will have done that long before we started looking for the mesons, so it's as good as done. However, if it's God we seek, we might have to consider carefully the sort of thing we are looking for, methodologically prior to the search. The only reason that it makes any sense to go out and look for a God is that we have a practice of talking and thinking about God and of fine-tuning our behaviour depending on whether it seems to us that there is a God or not. The search for God is, at least initially, a search for the thing that meets our God-profile. This is why the search for God confronts us early on with the question "What would a thing have to be like before we would regard it as God?" In order to get any sort of fix on our target, we must have a conception of Godiness - of the circumstances under which we would ascribe it - that informs our search, even if the conception is something we develop in the process of the search.

Now we can see how conceptual analysis comes in at the methodological starting-point. This question "What would a thing have to be like for us to regard it as God?" just is the question "Under what clearly-specified circumstances would we say that we have a God, which situations would clearly be Godless and when don't we have a clue?" And that is just a question about possibilities - one which we can ask and answer from the armchair so long as we have a rich enough description of the cases on which we are to rule. (We might ask, for instance, whether a circumstance in which the world was created by a mortal, but omniscient mouse would be one in which God created the world. Our answer would determine whether, should we discover that this was indeed the case, we would say that God in fact created the world.)

What has all this to do with nothing-over-and-above? Why is conceptual analysis needed there? Well, questions of the form <What does it take for something to be regarded by us as a K?> always crop up when we are considering whether and to what extent the subject-matter of one classificatory practice can be assimilated to that of another. Suppose you hold the view that economic facts are just constructs out of psychological facts: you believe that the economic is nothing over and above the psychological. If you are 'seriously' going to claim that posits like supply and demand are constructs out of the psychological, you had better have to hand a story about how to translate between theories within psychology and the economic theories in which supply and demand feature. And to do that you need to know something fairly detailed about the role played in economic theory by the posits. You need to know what it would take for us to regard something as a fall in demand or as a rise in supply. In particular, you need to know which scenarios, described in austere psychological terms, count as falls in demand, which as rises in supply and so on.

Or suppose you are wondering whether freely-chosen action is possible in a world where every action is causally determined by the physical facts up to the moment that the

action is performed. You are wondering whether, in such a world, free choice is anything<sup>16</sup> over and above something physical. To approach the matter, you must consult intuitions about which acts are free and which are not. In particular, you need to consider whether there are any physicalistically determined acts which would count, according to those intuitions, as free acts.

Now we can answer the original objection point by point. First, yes, conceptual analysis appeals to intuitions, not to facts, but intuitions about cases are highly relevant to the search for empirical facts. They are relevant because it is only when we have an intuitive idea of what sort of thing we seek that we can go out and seek it. As Grice has it: "It is a very old idea in philosophy that you cannot ask, in a philosophical way, what something is unless (in a sense) you already know what it is." (1958, 173.) Second, yes, conceptual analysis appeals to the possible, not the actual, but appeals to merely possible cases are relevant because we conduct our search in partial ignorance of the way things actually are; the point of the search is to see whether one of the possibilities in which there are Ks is the actual case. Third, yes, access to the nature of things is not in general a priori access, but the role we have specified for conceptual analysis need not have anything to do with the discovery of essential natures. The discovery that gold is by nature an element was indeed an empirical discovery. Jackson's point, however, is that we were only able to make it because we knew, by consulting our intuitions, what sort of actual shiny, valuable thing with samples of which we are acquainted would be regarded by us as gold. The intuitive knowledge that gold is shiny and valuable enables us to pick up the scent, and we need some way of picking up the scent; so conceptual analysis was needed for the research. Nobody need claim that gold's shininess and comparatively high value are essential properties of it. They are just properties we need to know about in order to strike gold.

### **3 Third objection : Jackson's description of scientific and philosophical discovery is inaccurate**

A common worry about Jackson's story is that assimilations of one scientific subject-matter to another do not in practice require sustained or systematic reflection about possible cases in the light of our intuitions. Instead, they are a combination of a posteriori discovery and ontological parsimony. Consider, for example, the discovery that the temperature of a gas is simply the mean kinetic energy of the molecules that constitute it. The story is sometimes told this way. The kinetic theory of gases was a theory about what gases were made of: it said they were made of small molecules. Within this theory, it was possible to explain the fact that a gas had a certain temperature in terms of the motion of the molecules. This was an empirical discovery. But then general methodological considerations about how science should be done kicked in. We want our overall picture to be uncluttered. It is far simpler to rule that the mean molecular kinetic energy of a gas *is* its temperature than to opt for a more complicated theory in which we distinguish mean molecular kinetic energy from temperature.

#### **Reply**

Ontological extravagance was never an option. Here's an alternative account of this example which invokes Jackson's thesis. The temperature of a gas is a posit in the thermodynamic theory of gases, but that theory does not pretend to be a fundamental theory. We can consider various possible cases in which gases are constituted in various ways and ask, for each such case, what would count as the temperature of a gas, if anything. Of course, we don't need to run through all these cases in advance of doing some

more fundamental science, but when we discovered that gases are composed of independently-moving particles, we knew enough about the commitments of the less fundamental thermodynamic theory to know that mean molecular kinetic energy does the work that temperature does in that latter theory.

The advantage of this Jacksonian account is that it doesn't require an appeal to simplicity. Once we knew the molecular dynamics of gases, we knew what temperature really was. It was not like deciding whether we really needed ungainly mechanisms to make our theory work. It was just a case of redescribing a macro-phenomenon in micro-phenomenal terms.<sup>7</sup>

#### **4 Fourth objection: Metaphysics should not be hostage to conceptual analysis**

Surely it's okay if, in the end, the things in the world that turn out to be the Ks lack many, perhaps even most, of the characteristics which our unsophisticated folk theories attributed to Ks. Bill Lycan is interested in showing that mental states like beliefs and desires are really just physical. He is a nothing-over-and-aboverist about the mental, but he doesn't like the Jackson thesis.

I am entirely willing to give up fairly large chunks of our commonsensical or platitudinous theory of belief or of desire... and decide we were just plain wrong about a lot of things, without drawing the inference that we are no longer talking about belief or desire. (Lycan 1988, 31)

If Lycan is right then, in the case of beliefs and desires anyway, it's not such a big deal which kinds of physical scenario would count for us as cases where we have a belief or a

desire. Lycan suggests that it doesn't really matter if we don't pin down our subject-matter by analysing our practices of ascribing mental states.

### **Reply**

Lycan can't really be right. If what you discover is fairly unlike what you started looking for, then either you have changed the subject or else your concept evolved while the search was underway (and perhaps as a result of the search). If it really is BELIEF or DESIRE that we are interested in assimilating to the physical, it seems to matter a great deal that we not stray far from what the people who talk about beliefs and desires have in mind. Otherwise, it seems that we are no longer talking about beliefs and desires.<sup>8</sup> Of course, what the folk regard as criterial for something's being a belief or a desire is almost certainly a moveable feast. Different folk will have different views about it; most folk are unreflective about it; and the impact of scientific discovery, social evolution, apparent inconsistencies among the criteria, or just new kinds of experience may prompt radical criterial revision. We cannot assume that our concepts are uncontested, determinate, or stable over time. And of course, if we haven't done the conceptual analysis yet, we won't know just which things the folk would be prepared to count as beliefs or desires. We also won't know what changes they would be prepared to make to their concepts of belief and desire in response to scientific investigation.

In this connection it is interesting to note what Lycan goes on to say:

I think that the ordinary word 'belief' (qua theoretical term of folk psychology) points dimly towards a natural kind that we have not fully grasped and that only mature psychology will reveal. (Lycan 1988, 32)

This might well be right. But if it is right, it is a piece of accurate conceptual analysis. It <sup>20</sup> is right because although the folk are fairly elastic about what they will count as a belief when a situation is described to them in physical terms, they will probably want beliefs to be fairly natural physical kinds. They will probably weight the intuition that a belief is a natural kind very highly and be prepared to say that they were wrong in lots of ways about what it meant to believe something so long as beliefs are fairly natural things, according to the best physical ontology. That's if Lycan's right. In other words, Lycan is wrong to think that it's okay for the theorist to give up on large, weighty chunks of folk theory about beliefs and still say we have beliefs, but it might turn out that the very things he wants to give up on are such that the folk will let him be their guest.

Having said all of this, there's surely something right in what Lycan says. Surely science need not be subject to the tyranny of our a priori conceptions of things. Here's where we start making claims which Jackson might not endorse while still, we think, defending his thesis that conceptual analysis is needed for nothing-over-and-above. Consider another objection.

## **5 Fifth objection: Philosophy rightly understood is continuous with the natural sciences, and therefore does not require conceptual analysis**

Kim Sterelny once wrote:

My approach [to the mind] is not just physicalist, it is naturalist. Naturalists are physicalists... [B]ut naturalists have methodological views about philosophy as well; we think philosophy is continuous with the natural sciences. On this view, philosophical theories are conjectures whose fate is ultimately determined by empirical investigation.... An alternative conception is to see philosophy as an

investigation into conceptual truths proceeding by thought experiments that probe the way we understand our own concepts. (Sterelny 1990, *ix*.)

## Reply

First, Sterelny's characterisation of the second, less-than-naturalistic alternative philosophical programme is not a useful characterisation of Jackson-style nothing-over-and-above. Jackson does think we must investigate conceptual truths in order to nothing-over-and-above, but, as we have seen, the role he assigns them is only that of providing a methodological starting-point for our empirical research. Second, it's not clear that there is any useful question to address about which work counts as philosophy and which counts as natural science. Yes, Jackson wants to say that philosophers should do conceptual analysis, but he also thinks that scientists should (and do) too, and he says nothing about whether there is anything else for a philosopher to do.

These preliminaries aside, though, we agree with Sterelny, at least to this extent. Surely you can put forward an empirical hypothesis about the sorts of phenomena dealt with in a high-level theory without being a hostage to that theory. There are projects which are naturalistic in Sterelny's sense. However, they are different from nothing-over-and-above, so Jackson's thesis does not apply to them. And not all major debates in philosophy and science are instances. (Sterelny 2003 offers a revised assessment of the methodological situation which is in line with our remarks here.)

To illustrate the difference between Jackson-style nothing-over-and-above on the one hand and Sterelny-style naturalism on the other, suppose I am interested in the mind: in what we are going on about when we talk of beliefs and desires, what consciousness is, etc. Suppose I think too that there is nothing in the world but physical stuff and I want to account for all this physicalistically. Well, one option is to do some nothing-over-and-

abovary: to work out what, if anything, the ordinary people are doing when they ascribe <sup>22</sup> and entertain beliefs and desires or manage to stay conscious. This is surely a worthwhile project: mental states are discussed, considered and apparently experienced by ordinary people who also invoke them to explain and predict. If you buy into this project, we have argued, you need to do conceptual analysis. But there's another way to go. You start by thinking about some interesting puzzle that arises within a physicalistic framework for humans and for other animals. How do these creatures manage to perform in a way which shows they are tracking the environment, planning ahead, responding to certain stimuli and ignoring others, grouping perceptual information together in particular ways and so on? You could begin to construct a physicalist theory of how all this is done. You know, of course, that people attribute mental states to each other and to animals as well and you might use the data about when they do this, and what they attribute, as data in the construction of your physicalistic (your *naturalistic*) theory. But your job is not to discover the extent to which these folk-psychological practices and their mentalistic ontology are grounded in the physical. You are not trying to redescribe or rehabilitate folk psychology. You are trying to solve some puzzles physicalistically and these happen to coincide more or less with the sorts of puzzles which we ordinarily explain in terms of our theory of the mental. You would therefore do well not to say that you have developed an empirical theory about the nature of *belief*— a theory which might tell us what belief truly is. Before you are entitled to do that, you must amalgamate your project with one which involves analysis of the relevant folk concepts. As Sterelny (2003: ix) puts it, you would need “both a well-developed account of... folk commitments about belief, and a theory of reference for folk psychological vocabulary telling us the extent to which folk psychology’s vocabulary depends on the accuracy of folk psychology’s picture of the mind.”

As Sterelny says, any theories you develop in the course of this investigation are <sup>23</sup> empirical conjectures. This means, of course, that they are subject to the assessment procedures that apply to any empirical claims in the natural sciences. If Lakatos (Lakatos and Zahar 1978) is right, for instance, these empirical conjectures had better produce novel predictions soon and those novel predictions had better be confirmed or the scientific community will want to chuck your theory on the scrap-heap even if you don't.

We think that projects of this naturalistic kind are fairly common. Take mainstream Chomskyan linguistics, where a *grammar* is an empirical conjecture about the native speaker's linguistic competence: a theory of the internal structures that correspond to the rules she has mastered. Such a theory is primarily a hypothesis about which on-board rules enable the speaker to understand or produce the infinitely many sentences she can understand or produce and it is part of a psychological explanation of how those rules could have been mastered. It is *not* an attempt to collate what a lay native speaker regards (or could be made to regard via intuition probing) as the rules of her language. Even so, gathering the raw data for such a grammar involves eliciting folk responses to questions about possible expression-types: "Is this one grammatical?" "Are these two synonymous?" and so on. The aim in pumping these intuitions is not to produce an analysis of the folk concepts of grammaticality, synonymy and the like. Still less is it to show that such concepts are nothing over and above psychological processes. Rather, it is to amass fallible information for use in the construction or testing of a theory about what rules the native speaker has internalised. Although such a theory is called 'a grammar', mainstream linguists do not take it for granted that the rules of their theory must be rules that the native speaker could be made to endorse. Intuitions are merely empirical evidence that certain rules are among those that have been mastered. We do not claim that grammatical theories developed within a broadly Chomskyan framework rely less for their plausibility on folk

intuitions than, say, theories about the nature of folk psychology do. Surely they don't. 24

Folk psychology, or folk ethics, or the folk account of free will can be theorised about empirically just as the internalised rules that explain our linguistic competence can. Our claim is rather that Chomsky's starting point, the puzzle of linguistic competence, identifies his project as distinct from that of a nothing-over-and-above-ist. Chomsky explores folk intuitions in order to explain a folk capacity physicalistically. The physicalistic nothing-over-and-above-ist explores folk intuitions in order to investigate the physical natures of things which the folk talk and think about.

Other examples of Sterelny-style naturalism are Ruth Millikan's analysis of mental content in terms of proper function (Millikan 1986) and Michael Devitt's semantic programme (Devitt 1996). They are plausibly not attempts to reduce the folk theory of mental content attributions or of meaning to something more ontologically pristine. However much conceptual work is involved in their construction, Jackson's thesis is not in play.

## **6 Sixth objection: Jackson should not equate conceptual analysis with investigations into semantic competence**

To background this last objection - the one we endorse - it will be useful to recapitulate, in Jackson's words, what his thesis commits him to.

Serious metaphysics requires us to address when matters described in one vocabulary are made true by matters described in another. But how could we possibly address this question in the absence of a consideration of when it is right to describe matters in the terms of the various vocabularies? And to do that is to

reflect on which possible cases fall under which descriptions. And that in turn is to do conceptual analysis. Only that way do we define our subject - or, rather, only that way do we define our subject as the subject we folk suppose is up for discussion. (Jackson 1998, 41-42)

As noted earlier, our objection is not to Jackson's thesis as such. It is that when he explains it his own way, he smuggles *language* into his characterisation of conceptual analysis. We have talked about the *subject-matters* associated with different *human classificatory practices*. Jackson talks about 'stories told in different vocabularies'. He adopts an unashamedly linguistic outlook on conceptual analysis, and he is in illustrious company. In the passage from Grice quoted earlier - the one where he championed conceptual analysis - he talked about our talk: our talk about cause, about perception, about knowledge. Linguistic analysis has been central to philosophising in many traditions throughout the twentieth century. Even so, Jackson would be better off without the linguistic turn, or so we will argue. Let's look at more of what he says and then at why it creates unnecessary problems for him.

Our subject is really the elucidation of the possible situations covered by the *words* we use to ask our questions—concerning free action, knowledge, and the relation between the physical and the psychological, or whatever. I use the word 'concept' partly in deference to the traditional terminology which talks of *conceptual* analysis, and partly to emphasize that though our subject is the elucidation of the various situations covered by bits of language according to one or another language user, or by the folk in general, it is divorced from considerations local to any particular language. When we ask English users in

English for their intuitive responses to whether certain cases are or are not cases of knowledge, we get information (fallible information...) about the cases they do and do not count as covered by the English word 'knowledge'. But our focus is on getting clear about the cases covered rather than on what does the covering, the word, *per se*. We mark this by talking of conceptual analysis rather than word or sentence analysis. (Jackson 1998, 33-34)

We extract three claims from this passage.

1. When Jackson talks about conceptual analysis, he is not really talking about the analysis of concepts. Instead, he is talking about how to understand the way that linguistic expressions function.
2. The sense in which he is interested in how linguistic expressions function can be explained along the following lines: Those of us who use these words to pick out things in the world - things with particular properties - are able to do so because of certain beliefs, knowledge, or dispositions which count as determining a rule for the use of the term. Maybe we can't articulate that rule, but there must be one, and it must be one that we have internalised somehow because we are able to say - given a full description of a possible case - whether the term would apply to that case or not. So when we say we are doing conceptual analysis, what we are really doing is looking for the knowledge, disposition, or whatever that determines, or constitutes, the rules governing our uses of terms - of the sorts of terms we find in metaphysics and science especially. We can think of this as the project of finding out what makes us *semantically competent* with these terms.<sup>9</sup>
3. That isn't quite the right characterisation of conceptual analysis. The English word 'water' and the French word 'eau' and all the other translations of these words into various languages pick out the same kind of stuff. Strictly, it's not the word 'water' (or 'cause', or 'person', or 'belief', or ...) whose use we are investigating. Rather, it is something which

each of these words has in common with its translations into other languages. This does <sup>27</sup> not really alter any of what was said in 1. and 2. because if two terms really are exact translations of each other, the same rule is required for semantic competence with either - each is usable as a description of the same contexts as the other.

We have two worries about all this. First, Jackson's idea about the nature of conceptual analysis is hostage to his own theory of semantic competence with words, his theory that I have access by mere reflection to the knowledge in virtue of which I understand a verbal usage. Now Chomsky denies that there is such a thing as semantic competence, and there are other broadly Chomskyan semanticists, like Peter Larsen and Gabriel Segal (1995), who think that we do have knowledge that makes us semantically competent, but unlike Jackson, they believe it is locked away in a *semantic module*. For Larsen and Segal, this is the knowledge of how to prove the T-theorems of a truth theory and it is not accessible to any kind of reflection. Jackson is welcome to his theory of semantic competence, but it seems as though conceptual analysis ought to make sense independently of who is right about semantic competence with words.

One gesture towards fixing this problem, while sticking to Jackson's linguistic take on conceptual analysis, involves remembering that it is how we would *use* words in various scenarios that matters to conceptual analysis. Would I call that object with the lovely eyes and the beautiful singing voice a person, even if I knew it was a computer, or a Martian? The answer might depend a lot on who I am talking to, what information we share, whether I need or ought to be precise, polite, and so on. I won't refrain from calling the computer a person if my refraining will confuse or horrify you. Specifying all this stuff is part of specifying the sort of case we are ruling on. This sounds as though it should be dealt with as part of pragmatics - of how we do things with words - rather than within the theory of what enables us to master and understand the meanings of words.

Second worry. When we are interested in what it takes for an action to be right, or what causes what, or who are the people, it really doesn't seem that we are interested primarily in what rules we follow in virtue of which we understand the words 'right', 'cause' and 'person'. For one thing, it may be that a being without language can distinguish people from non-people, can act in a way which requires an understanding of cause and effect and so on. It would be operating with the same concepts as the ones we have words for. For a second thing, there might well be philosophically interesting *concepts* for which we lack words. For a third thing, even among the linguistically endowed, verbal conduct is only one of the behaviours that manifest our conceptual commitments. Who are the persons is evidenced not only by when we are disposed to call them persons but also by the fact that we wonder whether they should marry our daughters. What counts as a vote for Bush is evidenced not only by what gets called a vote but also by which chads we don't leave hanging. Indeed, if we wanted to find the translation into Jungle of the English word 'person' we would look to see what gets punished, what gets married and so on - not merely because these are useful nonverbal signs of which word we are looking for but because the word, whatever it is, normally gets used within a comprehensive person-tracking practice. Now perhaps some of these tell-tale non-verbal activities and the kinds of thinking they require would not be available to a nonlinguistic being; perhaps the acquisition of many concepts requires linguistic competence. Still, it doesn't follow even from this that what we are interested in analysing are the words we use. Finally, suppose that despite the above misgivings, conceptual analysts steadfastly restrict their attention to linguistic data. Still, one need not employ or even possess a word that covers all and only the persons in order to betray linguistically that one has a concept of persons.<sup>10</sup> Remarks like "Nobody alive can remember the Crimean War", "If you object to abortion, you should object to euthanasia"

and “War crimes should be punished” speak volumes about person concepts, whether or <sup>29</sup> not our language has a word that translates ‘person’.

All in all, we should not regard the objects of analysis as linguistic terms. Analysing our talk may be the best way to analyse our classificatory practices, but surely it is only really a heuristic. The practices themselves - or the concepts or dispositions which generate them - are our real quarry. True, we have not really said what it is we are analysing when we do conceptual analysis. We have said only that it is not language. But successful conceptual analysis can and does proceed without a clear fix on what concepts are, just as successful arithmetic proceeds without a clear fix on what numbers are. We should not assume that the success of conceptual analysis and the fact that we owe much of this success to our facility with the meanings of words deliver the verdict that words are our analysandum. This verdict is corroborated by Jackson's admission that it is the semantic rules shared by different languages, not the words themselves, that really interest conceptual analysts. If we are interested in the common background machinery that is shared across languages, then, yes, we might invoke very similar semantic rules in order to be competent with terms in different languages, but surely this is only because these different words are all getting at the same underlying thing - the common practice or the common concept.

We end by noting two likely reasons for Jackson's assimilation of conceptual analysis to an investigation into natural language semantics. The first reason is that linguistic tokens are admirably public, inspectable and physical. Even so, they are not what we are really after and we should not pretend that they are, just because the job of finding what we really are after is hard. The second reason is more interesting. One thing that Jackson's linguistic focus enables him to do is talk about our classificatory practices as though they were collections of sentences - roughly, as though they were theories in the

logician's sense. He talks about 'belief' as a kind term in a theory called 'folk psychology',<sup>30</sup> just as one might say that 'temperature-in-a-gas' is a kind term in the thermodynamic theory of gases. When Jackson describes conceptual analysis in a piece of nothing-over-and-above - let's say in a quest to assimilate the mental to the physical - he likes to think of it as a process of stating, for every case described in the language of a physical theory, what if anything would be the right thing to say in the language of folk psychology. This enables him to think of the method of cases as ideally involving a translation between the language of one theory and the language of another, more pristine theory. For example, there is a long true statement in the language of physical theory which entails the statement that Fred believes that strong leadership is wrong leadership. This way of thinking of our practices as theories and of our theories as linguistic objects is very useful to Jackson. It immediately licenses him to discuss questions about entailment and consequence that arise when we consider the relationships between different theories - say physics and psychology - and it lets him use as much formal apparatus for alluding to the inferential relationships among sentences as he likes. He can (and does) model the reduction of one practice to another by appealing to Ramsey sentences.<sup>11</sup> He can elaborate his conception of possibility by invoking the two-dimensional modal formalism developed by Davies and Humberstone (1980).<sup>12</sup> These formal procedures have been rigorously presented over the years and have become well-understood, in just the way that folk practices have not. Hence, there is a definite pay-off for Jackson when he treats questions about the nature of belief as though they were questions about one's competence with the word 'belief'.

We won't push a view on whether it is useful to talk of folk psychology, or our views on persons, or our moral practice, as a collection of sentences in a language. However, suppose it is okay to think in this way. Jackson could retain his talk of theories

and his talk of them as logico-linguistic items but ditch the claim that the terms in these theories are words as we ordinarily use them.

The way to do it is to think of the languages of these theories not as languages spoken by the folk or the scientists, but as languages constructed by the theorist-anthropologist-conceptual analyst. Suppose I am interested in finding out what beliefs are. Well, according to Jackson, my first step is to find out what we mean by the word 'belief,' which might be a bigger job than that because we probably need to understand the whole of folk psychology and hence the network of sentences which give 'belief' its meaning. Anyway, we are supposed to get out of this some idea of what it would take for a thing to be a belief. Now we have claimed that what Jackson should be saying is different. He should be saying that my job is to analyse all our belief-related practices. This more inclusive task will involve discovering the conditions under which one can rightly ascribe beliefs to people and the situations in which the word 'belief' can be felicitously uttered, but it will also involve discovering when our actions are best explained by the hypothesis that so-and-so believes that  $p$  and not that  $q$  and situations where non-linguistic creatures betray a notion of belief. Having done all of this, I might be able to attribute a folk theory - folk psychology - to the folk (and maybe to the relevant nonlinguistic creatures, if there are any). I will want to be able to express this theory verbally (or at least symbolically) in order to say anything about what it would take to realise the folk psychological roles, so it makes sense for there to be a language in which I express it. Maybe I will use the word 'belief' in that artificial language as a term for the folk psychological posit which ordinary English-speaking folk are getting at when they too use the word 'belief'. But it is to be understood that there might be more packed into the semantics of 'belief' in the language of my theory (my reconstruction of folk psychology) than is packed into the semantics of 'belief' in English. I will have stipulated that my neologism ('belief') applies to the practices of non-

English speakers and maybe of certain non-humans to some extent. This stipulation is not to be abandoned if, as it turns out, English-speakers don't use 'belief' to cover such cases.

Armed with a special vocabulary for my theory, I, the analyst, may blissfully nothing-over-and-aboveise. Good luck to me.<sup>13</sup>

- 
1. For Jackson's formulation of the thesis, see Jackson 1998, Chapter 2.
  2. See Block and Stalnaker 1999, Byrne 1999 and Yablo 2000.
  3. See the papers listed in the previous note and also Stalnaker 2003.
  4. See Child 1993 for such a view and Haugeland 1982 for more cases where both criteria for reduction fail.
  5. See Menzies 2001.
  6. See Williamson 1995 for a good discussion of this worry. Williamson says, however, that the worry might not arise for 'sophisticated' versions of conceptual analysis like that defended by Michael Smith (1994). Williamson would almost certainly include Jackson's among the suitably sophisticated versions.
  7. This approach is inspired by David Lewis's 1966 "An Argument for the Identity Theory."
  8. See Jackson 1998, 37-41.
  9. The connection between conceptual analysis and competence is made explicit in Jackson and Pettit 2002. More indirect evidence comes from Jackson 2003 where the account of semantic competence with referring terms mirrors Jackson's 1998 story about conceptual analysis.
  10. We owe this point and the example types to Denis Robinson: see Robinson 2004.
  11. See Jackson 1998, Chapter 6. Here Jackson draws on work in Lewis 1970.
  12. See Jackson 1998, Chapters 2 and 3.
  13. Many thanks to Fred Kroon, David Lumsden, Alex Morgan, Tahua O'Leary and Denis Robinson for their comments on drafts of this paper.

## References

Block, Ned & Stalnaker, Robert. 1999. "Conceptual analysis, dualism, and the explanatory gap," *Philosophical Review* 108, 1-46.

Byrne, Alex. 1999. "Cosmic hermeneutics," *Philosophical Perspectives* 13, 347-383.

Chalmers, David J. 1996. *The Conscious Mind*. Oxford: Oxford University Press.

Child, William. 1993. "Anomalism, Un-codifiability and Psycho-physical Relations," *The Philosophical Review* 102, 215-243.

Chomsky, Noam. 1986. *Knowledge of Language*. New York: Praeger.

Davidson, Donald. 1970. "Mental Events," reprinted Davidson 1980, 207-225.

Davidson, Donald. 1980. *Essays on Actions and Events*. Oxford: Clarendon Press.

Davies, M., and L. Humberstone. 1980. "Two notions of necessity," *Philosophical Studies* 38, 1-30.

Devitt, Michael. 1996. *Coming to our Senses: a Naturalistic Program for Semantic Localism*. New York: Cambridge University Press.

- Gettier, Edmund. 1963. "Is Justified True Belief Knowledge?" *Analysis* 23, 121-123.
- Grice, Paul. 1957. "Meaning," in Grice 1989, 213-223.
- Grice, Paul. 1958. "Postwar Oxford Philosophy," in Grice 1989, 171-180.
- Grice, Paul. 1989. *Studies in the Way of Words*. Cambridge, Mass.: Harvard University Press.
- Haugeland, John. 1982. "Weak Supervenience," *American Philosophical Quarterly* 19, 93-104.
- Jackson, Frank. 1994. "Metaphysics by Possible Cases," *Monist* 77, 93-110.
- Jackson, Frank. 1998. *From Metaphysics to Ethics*. Oxford: Oxford University Press.
- Jackson, Frank. 2004. "Why We Need A-Intensions," *Philosophical Studies* 118, 279-298.
- Jackson, Frank and Philip Pettit. 2002. "Response Dependence Without Tears," *Nous* Supplement 12, 97-117.
- Kripke, Saul A. 1980. *Naming and Necessity*. Basil Blackwell.

---

Lakatos, Imre and E Zahar. 1978. "Why Copernicus' Programme Superseded Ptolemy's" in John Worrall & Greg Currie (eds): *The Methodology of Scientific Research Programmes* (Cambridge: Cambridge University Press), 168-192.

Larson, Richard & Gabriel Segal. 1995. *Knowledge of Meaning*, Cambridge, Mass.: MIT Press.

Lewis, David. 1966. "An argument for the identity theory," reprinted in Lewis 1983, 99-107.

Lewis, David. 1970. "How to define theoretical terms," reprinted in Lewis 1983, 78-95.

Lewis, David, 1973. "Causation," reprinted in Lewis 1986, 159-172.

Lewis, David. 1983. *Philosophical Papers* Volume I, Oxford: Oxford University Press.

Lewis, David. 1986. *Philosophical Papers* Volume II. Oxford: Oxford University Press.

Lycan, William. 1988. *Judgement and Justification*. New York: Cambridge University Press.

Menzies, Peter. 2001. "Counterfactual Theories of Causation," in *The Stanford Encyclopedia of Philosophy*.

---

Millikan, Ruth. 1986. "Thoughts Without Laws; Cognitive Science With Content."

*Philosophical Review* 95, 47-80.

Putnam, H. 1975. "The meaning of 'meaning'," reprinted in Putnam, *Mind,*

*Language and Reality*, Cambridge University Press, 215-271.

Robinson, Denis. 2004. "Failing to agree or failing to disagree: Personal identity quasi-relativism," *The Monist* 87, 512–536.

Schiffer, Stephen. 1987. *Remnants of Meaning*. Cambridge, Mass.: MIT Press.

Seegerberg, Krister. 1973. "Two-dimensional Modal Logic," *Journal of Philosophical Logic* 2, 77-96.

Smith, Michael. 1994. *The Moral Problem*. Oxford: Blackwell.

Stalnaker, R. 2003. "Conceptual Truth and Metaphysical Necessity," in Stalnaker, *Ways a World Might Be*, Oxford University Press, 2003, 201-215.

Sterelny, Kim. 1990. *The Representational Theory of Mind*. Oxford: Blackwell.

Sterelny, Kim, 2003. *Thought in a Hostile World*. Oxford: Blackwell.

---

Stich, Stephen, 1993, "What is a Theory of Mental Representation?" in Stich and Warfield (eds), *Mental Representation: A Reader* (Oxford: Blackwell), 347-364.

Williamson, Timothy. 1995. "Is Knowing a State of Mind?" *Mind* 104, 533-65.

Yablo, Stephen. 2000. "Textbook Kripkeanism and the open texture of concepts," *Pacific Philosophical Quarterly* 81, 98-122.

<sup>13</sup> Many thanks to Fred Kroon, David Lumsden, Alex Morgan, Tahua O'Leary and Denis Robinson for their comments on drafts of this paper.