Intuitions and Conceptual Analysis

This is essentially a course in philosophical methodology. What we are interested in, above everything else, is finding out which methods should be used to find the intension of a particular concept, that is, to find which actual and possible things satisfy that concept. Since a large part of what we do in, say, ethics, epistemology, metaphysics, philosophy of science and philosophy of language is concerned with this question, what we discover should at least potentially be relevant to large chunks of our philosophical practice.

As illustrations: at least part of what we do in ethics is deciding which (possible) actions or persons are good and which are not; at least part of what we do in epistemology is deciding which (possible) agents know which propositions and which do not; at least part of what we do in metaphysics is deciding which (possible) events (or whatever) cause other events (or whatever); at least part of what we do in (Bayesian) philosophy of science is deciding which (possible) decisions are prudent, and at least part of what we do in philosophy of language is deciding the meaning of particular terms.

In other words, we are interested to some extent in a conceptual analysis of goodness, knowledge, causation, prudential decision and meaning. Or if we are not interested in conceptual analysis, we are at least interested in the kinds of questions which would have been traditionally dealt with by conceptual analysis.

One of the funny things about this game is that, despite the fact that we only play the game because we don't know the answers, in playing we often assume that we do know the answers, or at least large parts of them. We assume that we know whether specific agents are good or evil, whether other agents do or don't know certain propositions, and so on. Often all the evidence we have for this will be an intuition that it is the case. A possible case which appears to cause problems for an analysis because our intuitive classification of it is inconsistent with the analysis will be called an intuitive counterexample, or just a counterexample, for short.

If we take counterexamples seriously, then several interesting questions immediately arise. First, can we always trust these intuitions? If not, can we say when we can and can't trust them? Secondly, if we can't, where does this leave the project of uncovering the nature of these concepts?

So there are some worries to begin with about the data that will be used as inputs for the program of conceptual analysis. Many of you will already be familiar with a much more pressing concern for this program, the apparent impossibility of ever achieving a true conceptual analysis. Even bracketing the concerns about intuition, we might wonder (a) whether conceptual analysis is possible, (b) whether it is worthwhile and (c) if so, how it should be done.

This course will be interested in large part with answering these questions, or at least providing materials for you to come to your own opinions about what the answers may turn out to be. To spoil at least some of the surprise, I'll be arguing that (a) and (b) should be answered affirmatively.

Examples of Counterexamples

Exhibit One: Gettier

Ayer, Chisholm and possibly Plato held that *X* knows that *p* is true iff *X* has a justified true belief that *p*. Gettier invited us to consider possible cases like the following.

Smith has a justified, but false, belief that Jones, one of his co-workers, owns a Ford. (Smith has been presented with misleading evidence to this effect.) He infers that someone in his office owns a Ford, and presumably this belief is also justified. Now as it turns out, someone else does own a Ford, call her Brown, but Smith is completely unaware that she does. So Smith has a justified true belief that someone in his office owns a Ford, but intuitively he does not know this. So knowledge is not justified true belief.

Exhibit Two: The Hanging Judge

Analytic Utilitarianism, as I will use the term, is the doctrine that by definition it is right to do what will produce the best consequences. There are tricky questions about how to measure consequences, whether expected or actual consequences matter, and so on, but let's ignore them.

You are a judge in a small southern post-bellum town. A white woman was raped by a black man, setting off riotous demonstrations by the white town-folk. A black man has been brought into your court charged with the offence, and though he is plainly innocent, the town-folk are convinced of his guilt. Indeed they are so convinced, that should you do anything other than sentence him to death, they will surely riot, lynching this man and probably causing other deaths in the process. The best consequences will be produced by hanging the poor defendent, who is bound to die one way or the other, and avoiding the riot. But intuitvely this is the morally wrong thing to do. So the right action is not always that which produces the best consequences.

Exhibit Three: Blockhead

Some early defenders of functionalism thought that we could define a mental state by its functional role, meaning by that, its typical causes and effects. It is an advance on behaviourism, because it allows that a mental state may count as a belief or desire or whatever, by virtue of the dispositions it induces, even if these dispositions are never realised. So by the functionalist lights, a creature is intelligent iff it is disposed to react to a suitably wide variety of environmental inputs in the right kind of way. Cashing out the right-hand side is non-trivial, but irrelevant to this example.

Blockhead looks and feels and like an intelligent human, but he has a quite different internal composition. Instead of a neural network in his head, he has a giant look-up table, stating the appropriate response to *any* possible environmental stimulus. The table is so big, that he has all the dispositions to react to his environment that any person could have, yet intuitively Blockhead is not intelligent. So (this) functionalist analysis of intelligence is refuted.

Books

The primary text for the course is:

Frank Jackson, *From Metaphysics to Ethics*, Clarendon Press, 1998. (**Jackson)**

Other books which would be very well worth having (in order of priority) are:

Michael dePaul and William Ramsey (eds), *Rethinking Intuition*, Rowman & Littlefield, 1998. **(Intuition)**

David Lewis, *Papers in Metaphysics and Epistemology*, Cambridge, 1999. **(Lewis)**

(and perhaps) Robert Stalnaker, *Context and Content*, Oxford, 1999. **(Stalnaker)**

The last three are all collections of essays, many of which are not relevant to this course. Indeed, the Stalnaker is really only relevant for the 'two-dimensional' account of possibility and *a prioricity* which he outlines, and we'll discuss. Several other articles will also be distributed as photocopies during the course.

On a more administrative note, **Jackson**, **Intuition** and **Lewis** are available from the bookshop in Marshall Square. Also, all of these books are available from Barnes&Noble.com, and if it turns out to be necessary I'll put in a bulk order for them to cut down on shipping costs. **Jackson** and **Intuition** cost about $30 each (inc. tax and shipping), and **Lewis** and **Stalnaker** cost $25 each (again inc. tax and shipping). Any savings on those prices we make by making a collective order I will dutifully pass on to anyone participating in the scheme.

Assessment

The assessment for the course will consist in:

Three short assignments ('weeklies') on the readings in the early part of the course;

A mid-sized paper on counterexamples due in mid-November, and

A term paper due on the last day of classes (December 10)

Because there is other work involved, I don't expect the term paper to be as long as it otherwise might have been. The questions for the weeklies and the paper on counterexamples are attached to this handout. The questions for the term paper are not yet written, but should be available soon. In any case, there is quite a bit of flexibility on topics for the term paper. The three pieces of assessment (taking the three weeklies as a single piece of work for this purpose) will count for roughly equal weight in determining final grades.

As you will see in the handouts, the paper on counterexamples can be replaced by a presentation on a philosophically interesting counterexample. It might be fun by the time we get to week seven or eight to have a change of host for a while, so it would be good if some students took up this option.

Outline of course

The course structure is not nearly as rigid as this outline makes it appear, but I thought it might be helpful to have some idea of where the course is headed. In particular the case studies which I have planned for the last few weeks are quite modular, and can easily be replaced if there are other topics that people want covered. That disclaimer aside, here's the plan.

Other than the chapters from **Jackson**, I will be making available copies of the papers listed here a week or two before the class. The readings with a star beside them indicate that there will be a short assignment on that paper, due the Friday before the class at which the paper will be discussed. The readings are listed in rough order of relevance to that week's discussion. The first one or two papers will be discussed in class, the rest are background.

Week One: Introduction

The classical model of analysis

The inductive argument against the possibility of analysis (Harman)

Harman, "Doubts about Conceptual Analysis" in Hawthorne and Michael (eds) *Philosophy in Mind* Kluwer, 1994, pp 43-48.

Week Two: Psycho-analysis

Alternative models of concepts

Empirical psychological data bearing on the plausibility of each model

McNamara and Sternberg "Mental Models of Word Meaning" *Journal of Verbal Behavior and Learning* 22 (1985): 449-474.

*Rosch and Mervis, "Family Resemblances: Studies in the Internal Structure of Categories," **Intuition**, 17-44

Murphy and Medin "The Role of Theories in Conceptual Coherence" *Psychological Review* 92 (1983): 289-316.

Week Three: Eliminating Analysis

Questioning the worth of analysis: Stich and Tye argue that since analysis can only tell us about our language, it cannot reveal the deep nature of things; at best it can reveal our implicit assumptions about the deep nature of things.

*Stich, "Reflective Equilibrium, Analytic Epistemology, and the Problem of Cognitive Diversity," **Intuition**, 95-112.

Stich, "What is a Theory of Mental Representation", *Mind* 101 (1992): 193-232.

Tye, "Naturalism and the Mental", *Mind* 101 (1992): 421-441.

Cummins "Reflections on Reflective Equilibrium" **Intuition**, 113-127.

DePaul "Why Bother with Reflective Equilibrium" **Intuition**, 293-309

Week Four: Physicalism & Serious Metaphysics

What does physicalism commit us to? Is there a clear statement of physicalism which makes it (plausibly) non-vacuously true?

The Location Problem and the Entry by Entailment Solution.

Jackson, ch. 1

Crane and Mellor, "There is No Question of Physicalism" *Mind* 99 (1990): 185-206.

Lewis, "New Work for a Theory of Universals", in **Lewis**, pp 8-55, particularly pp 33-39.

Kirk, "From Physical Explicability to Full-Blooded Materialism" *Philosophical Quarterly* 29 (1979): 229-237.

Week Five: What Analysis Can Do For You.

Why does analysis have a role in metaphysics?

How is analysis related to intuitions about possible cases?

How is the program of analysis consistent with what we learned from Kripke about reference?

***Jackson**, ch. 2, pp 28-42.

Slote, "The Theory of Important Criteria" *Journal of Philosophy* 63 (1966): 211-224.

Geach, "Some Problems About Time" in *Logic Matters*, Blackwell, 1972, pp 302-318.

Kroon, "Causal Descriptivism" *Australasian Journal of Philosophy* 65 (1987): 1-17.

Week Six: When are Counterexamples Relevant?

Why do we react so differently to problem cases in different parts of philosophy?

How can we tell whether different parties to a dispute have opposing opinions about a single concept and when they fail to share a concept?

Smart, "The Methods of Ethics and the Methods of Science" *Journal of Philosophy* 62 (1965): 344-349.

Shope, *The Analysis of Knowing: A Decade of Research*, Princeton, 1983, chs 1 & 2

Unger, *Living High and Letting Die*, Oxford, 1996, ch 4

my "What Good Are Counterexamples?", manuscript

"Humean Supervenience Debugged", **Lewis**, 224-247, esp sect. 10

"New Work for a Theory of Universals", **Lewis**, 8-55, esp pp 45-55

"Putnam's Paradox", **Lewis**, 56-77

Week Seven: How to Beat a Counterexample.

First strategy: Guilt by Association. If the intuition behind the counterexample clearly leads to other fallacious beliefs, the intuition is mistaken. This is the thought behind Horowitz's paper.

Second strategy: Explaining Away. It hurts to say that intuition is wrong. It hurts less if you can explain why it is wrong. One popular way of doing this is based around H. P. Grice's influential theory of conversation.

Horowitz, "Philosophical Intuitions and Psychological Theory", **Intuition**, 143-160.

Grice, *Studies in the Way of Words*, Harvard, 1989, chs 1 & 2

Lewis, "Causation as Influence", manuscript (http://www.nd.edu/~mmmm/lewis.pdf)

Jackson, "Assertion and Indicative Conditionals" *Philosophical Review*, 88 (1979) 565-589.

Week Eight: Presentations on Counterexamples

See 'Alternative Assessment Task' on the handout for assignment four.

Readings will be set by students doing presentations.

If there are few presentations, we will just move forward the rest of the course.

Week Nine: Conceptual and Metaphysical Necessity

Preliminaries: Replies to some objections to conceptual analysis (**Jackson**, 56-67)

The distinction between being necessary and being knowable *a priori* (Kripke, Tichy)

Formal representations of this distinction (Stalnaker, Davies and Humberstone)

Using the distinction to show how conceptual analysis delivers *a priori* results (**Jackson**, 46-52)

A reply to the objection that Kripke shows we cannot perform *a priori* analysis (**Jackson**, 67-83)

Jackson, pp 46-86

Stalnaker, "Assertion", *Syntax and Semantics* vol. 9. Also in **Stalnaker**, 78-95.

Davies and Humberston, "Two Notions of Necessity" *Philosophical Studies* 38 (1980): 1-30.

Kripke, *Naming and Necessity*, Blackwell, 1980.

Tichy, "Kripke on Necessity A Posteriori" *Philosophical Studies* 43 (1983): 225-241.

Week Ten: Replies to Jackson's Use of the Two Kinds of Necessity

There are more ways for modal intuitions to go wrong than Jackson allows (Yablo)

Jackson has misrepresented the way the two types of necessity interact (Block and Stalnaker)

Yablo, "Textbook Kripkeanism & the Open Texture of Concepts," *Pacific Philosophical Quarterly*, forthcoming (preprint at http://www.mit.edu/~yablo/tk.html)

Block and Stalnaker, "Conceptual Analysis, Dualism and the Explanatory Gap" *Philosophical Review*, forthcoming (preprint at http://www.nyu.edu/gsas/dept/philo/faculty/block/papers/ExplanatoryGap.html)

Chalmers, *The Conscious Mind*, Oxford, 1996, Chs 2 & 4

Week Eleven: Moral Realism and Ramsey Sentences

Moral functionalism (**Jackson**, Ch. 5)

Ramsey sentences (Lewis)

Analytic Descriptivism as a brand of modal realism (**Jackson**, Ch. 6)

Alternatives to analytic descriptivism (Smith, Darwall *et al*, Boyd)

Jackson, Ch. 5 and Ch. 6

Lewis, "How to Define Theoretical Terms" *Journal of Philosophy* 67 (1970): 427-446

"Psychophysical and Theoretical Identifications" **Lewis**, 248-261

Smith, *The Moral Problem*, Blackwell, 1994, esp ch 2

Darwall, Gibbard and Railton, "Towards *Fin de Siècle* Ethics" *Philosophical Review* 101 (1992): 115-189.

Boyd, "How to be a Moral Realist" in Sayre-McCord (ed) *Essays on Moral Realism*, Cornell, 1988, pp 181-228.

Week Twelve: Colour and Eliminativism

The Dispositional Theory of Colour (Johnston)

The objection from the causal role of colour (Jackson)

Comparisons with other arguments about causal roles (Kim, Yablo, Merricks)

Jackson, Ch. 4

Johnston, "How to Speak of the Colors" *Philosophical Studies* 68 (1992): 221-263

Kim, *Supervenience and Mind*, Cambridge, 1993, essay 14

Yablo, "Cause and Essence" *Synthese* 93 (1992): 403-449

Merricks, "Epiphenomenalism and Eliminativism" manuscript, preprint at http://www.nd.edu/~mmmm/trenton.pdf

Week Thirteen: More Case Studies to Finish With

Blockhead as a counterexample to mental functionalism (Block, Braddon-Mitchell and Jackson)

Newcomb's Problem as a counterexample to evidential decision theory (Gibbard and Harper, Lewis)

Block, "Psychologism and Behaviourism" *Philosophical Review* 90 (1981): 5-43

Braddon-Mitchell and Jackson, *Philosophy of Mind and Cognition*, Blackwell, 1996, especially pp 111-121.

Gibbard and Harper "Counterfactuals and Two Kinds of Expected Utility" in Hooker *et al* (eds) *Foundations and Applications of Decision Theory*, Reidel, 1978.

Lewis, "Causal Decision Theory" *Australasian Journal of Philosophy* 59 (1981): 5-30.

Week One: Classical Analysis

The traditional, or classical, theory of concepts says that for any concept *F*, there are conditions G_1, G_2, ..., G_n such that for any possible *a*, *a* is an *F* iff it is a G_1 and it is a G_2 and ... and it is a G_n, for some finite (and hopefully small *n*). In other words, that there are necessary and sufficient conditions for the application of any particular concept. For starters, we'll look at some of the commitments we incur if we buy the classical theory, and a relatively unimportant argument against its plausibility. Let's start with the commitments.

We **are not** committed to the view that *F* is not vague. Wittgenstein seems to assume something like this in the *Investigations* in his discussion of games. At one stage he suggests that we cannot provide an analysis of 'game' (I presume this means a short list of necessary and sufficient conditions) because there is no sharp boundary between games and not-games (See §§68-71). But the traditional theory is compatible with *F* being vague, provided at least some of the G_i are also vague. As Lewis puts it in "Counterfactuals and Comparative Possibility", the terms on either side of the analysis might be equally vague, they might 'sway together'. I suppose the traditional theory is incompatible with the existence of precisely one vague concept, for then there would be no vague terms for the analysans. I think traditionalists need not fear this possibility.

We **are not** committed to the view that *F* is an 'on/off' concept. That is, the traditional theory is compatible with *F* coming in degrees. Murphy and Medin (1985: 311) seem to think that it is devastating for the traditional theory that speakers regard terms like 'vehicle' as admitting of degree. And the first paragraph of the Rosch and Mervis paper in the readings suggests the existence of degree concepts means we have to amend the traditional theory, though they might be stipulating that what they call the 'classical view' rules out degree terms. As in the previous paragraph, it is hard to see what the problem is. If G_1 is a degree concept, and *x* is an *F* iff it is a G_1, a G_2 and a G_3, then presumably *F* will be a degree concept. The classical view seems to be inconsistent with the position that there is precisely one degree concept, but again this possibility seems unlikely. There are tricky questions as to what to say when more than one of the defining attributes of a concept are degree concepts, but (a) there are several well-known answers to this and (b) this problem isn't unique to holders of the traditional view of concepts, it is just the problem of how to understand connectives when we admit anything like degrees of truth.

We **are not** committed to the view that the defining attributes of *F* will tell us anything about what it is to be a prototypical *F*. Assume that *x* is an *F* iff it is a G_1, G_2 and G_3, and that most of the things which are G_1, G_2 and G_3 around here also have properties H_4, H_5, ..., H_{10}. Then presumably what it will be to be a prototypical *F* will be to have all or most of these ten properties. This is all consistent with saying that the *H* are no part of the meaning of *F*. (Note we should distinguish this from the case where *F* is defined as the stuff around here which has most of the properties H_4, ..., H_{10}, and this happens to be the stuff which is G_1, G_2 and G_3. In that case the *H*'s are in an interesting sense part of the meaning of *F*.)

In general, we should be wary of claims about how much the classical theory is meant to show. I call unwarranted assumptions about the intended scope of the theory instances of the 'Red Sox fallacy', for the following reasons. I would like to know more than I now do about concepts, or at least about conceptual terms. I would also like to know whether the Red Sox will ever win another World Series. The classical theory tells me nothing about the second question; should I conclude it is useless as an answer to the first question? Of course not, and I suppose everyone realises that. But replace the question about the Red Sox with one somewhat less absurd, and similar reasoning is all too often endorsed.

For example, the Murphy and Medin paper we'll look at next week is explicitly dealing with the question, why do some concepts seem more 'coherent' than others? They fault the classical view for not distinguishing between the coherent-seeming and non-coherent-seeming concepts: both have necessary and sufficient conditions of applicability. There is a clear implication that they think this is bad news for the classical theory as an account of when concepts apply at all. But why should the classical theory need to answer Murphy and Medin's question about coherence any more than it needs to answer my question about the Red Sox? (As an aside, Murphy and Medin seem to ignore the most exciting possible answer to their question: that some concepts seems more coherent than others *because they really are* and, presumably, because we are not incompetent at detecting such coherence. If that is true, if there is some objective coherence to concepts which we are detecting by our judgements of coherence, then it is a fascinating area for philosophical exploration.)

We **are not** (I think) committed to saying that competent users of the term '*F*' will know that something is an *F* iff it has the defining attributes. I think, and this is very controversial, that linguistic competence requires little or no knowledge of truth conditions for terms in the language whatsoever. Rather, all competence requires is that speakers know when terms can be properly used. Now there might be no short list of necessary and sufficient conditions for that, even if the classical view is right. We'll return to this much later in the course, when we discuss the importance of Grice's theory of communication for conceptual analysis. And even if we are committed to competent speakers having knowledge of analyses (or at least assertibility conditions) we are not committed to this knowledge being explicit, or being capable of verbal expression.

We **are** (if the classical theory is to be non-trivially true) committed to the defining attributes being not too gruesome, I think. If this really is a commitment, I don't think the classical theory is true for all concepts, though I do think it is true for most. If the defining attributes are allowed to be as gruesome as you like, it seems to become trivial. Say that someone tries to advance a particular non-classical view, say that *x* is *F* iff it bears a suitable resemblance relationship to an exemplar, then the classical theorist can simply colonise this position. There is a defining attribute for being *F*, it is bearing a suitable resemblance relationship to an exemplar. I think what is wrong with this attempt at colonisation is that it concedes that any defining attributes for *F*-ness are particularly gruesome. (See **Jackson**, pg 61 for an example of this move.)

We **are** (if the classical theory is to be non-trivially true) committed to some distinction among the properties so we can say the analysis is non-circular. We don't want our analysis to be *x* is an *F* iff it is an *F*. So we can't use *F* in the analysis. We probably can't mention it either, though Jackson denies this at one point. That is, we don't want our analysis to be *x* is an *F* iff it falls within the intension of '*F*', or would be called '*F*' by a competent speaker. Now in practice drawing such boundaries is a surprisingly difficult task. Despite years of trying, there is no clear consensus on where to draw the boundary between the normative and the non-normative, the mental and the non-mental, the observational and the non-observational, the intrinsic and the extrinisic, and so on. If our analysis is to be non-circular in any interesting way, I think we need to be able to specify which properties cannot be used in a non-circular list of defining attributes of *F*-ness. Obviously *F* will be on the list. Presumably for most *G*s, the attribute of being a *F* or *G* will also be on the list. Else we will be able to define an *F* as something which is an (*F* or *G*) **and** an (*F* or not-*G*). In practice, the property of being an *F* or *G* will be pretty gruesome, so maybe these two constraints will overlap to some extent. We will come back to the issue of circularity when looking at Jackson's discussion of Ramsey sentences.

The Inductive Argument Against Conceptual Analysis

Philosophers have been trying for 2500 years to find workable analyses of philosophically interesting terms into necessary and sufficient conditions. You would think that if there were any they would have found them by now. But, famously, all such attempts have come up short. Indeed, they have all failed so spectacularly that the very distinction between facts about meaning and facts about the world which this account presupposes has been called into question. By induction on the history of philosophy, we conclude that there are no such analyses to be found.

This argument is surprisingly popular. Here's four recent endorsements of it. All references are to papers on the reading list.

"Perhaps the safest bet is that whatever the mental mechanism underlying intentional categorization may be, it will not utilize "classical" concepts—the sort that can be defined with a set of necessary and sufficient conditions. The argument here is straightforwardly inductive: *No* commonsense concept that has been studied has turned out to be analyzable into a set of necessary and sufficient conditions. Indeed, given currently available evidence, it looks like there are no classical concepts" (Stich 1992: 249-50).

"[Analytic naturalism] does not sit well with the long and miserable failure of philosophical reflection to have produced any unobvious, satisfactory conceptual analyses. Surely the lesson to be learnt from the past here is that concepts generally don't have a priori discoverable necessary and sufficient conditions. This is also strongly suggested by recent work in cognitive psychology..." (Tye 1992: 424-5).

"When Quine, Putnam, Winograd and a host of others raised objections to the analytic-synthetic distinction, they did not mention controversial philosophical analyses. When problems were raised about particular conceptual claims, they were problems about the examples that had been offered as seemingly clear cases of *a priori* truth—the principles of Euclidean geometry, the law of excluded middle, 'cats are animals', 'unmarried adult male humans are bachelors', 'women are female' and 'red is a color.' … Speakers do not consider the Pope a bachelor. People will not apply the term 'bachelor' to a man who lives with the same woman over a long enough period of time even if they are not married. Society pages in newspapers will identify as eligible 'bachelors' men who are in the process of being divorced but are still married." (Harman 1994: 45)

"The failure of analytic philosophy to produce an uncontroversial, completely satisfactory analysis of the vast majority of abstract concepts should by itself suggest something is amiss." (Ramsey 1998: 174)

This argument does not propose simply to show that the classical theory is wrong for some words, but that it is wrong for all words! Five comments on this argument.

First, the evidence does provide better inductive support for a denial of the classical theory than it does for a denial of the analytic/synthetic distinction. Just showing that sentences like, "All bachelors are unmarried", or "All cats are animals" does not show there are no analytic truths to be found, even excluding logical truths. Even if we can't find an analytic positive universal, we can find all sorts of analytic-sounding negative universals, like "No bachelor is a grapefruit", or "No cat is a proper class."[1] If these are analytic it is bad news for the Quinean project of denying analyticity. But since these are analytic on both classical and non-classical theories of concepts, these are no help to the classical theorist. In short, necessary negative universals are good news for the project of saving analyticity, but no help for the project of saving the classical theory.[2]

Secondly, Tye's insistence that conceptual analysis provide us with unobvious truths seems a bit stringent. Why would it be a problem with the reliability of conceptual analysis if all of its output was obvious? It isn't good if a theory is platitudinous; it isn't good if a theory is wrong; but these flaws should not be confused. Being banal is not evidence of being wrong, it is evidence of being right! Tye's complaint seems to be like the kind of response Lewis is justly criticising here: "We are capable of all sorts of behaviour that would seem bizarre and unintelligible, and this is exactly the behaviour that folk psychology predicts, rightly, will seldom occur. But we take a special interest in questions that like beyond the

[1] In *Parts of Classes* Lewis suggests that at least some cat might be a set: the null set. But the null set isn't a class, so this doesn't look like much evidence that a cat may be a proper class!

[2] Harman makes a logical mistake on this point. It is perfectly consistent with the claim he says he is discussing, "Unmarried adult males are bachelors" that separated, but still married, adult males are *also* bachelors. This is inconsistent with the *analysis* of 'bachelor' which I suppose Harman thought he was discussing; but if so he probably should have said this.

predictive power of folk psychology; wherefore ingrates may fairly complain of a lack of *interesting* predictions!" (The quote is from "Reduction of Mind")

Thirdly, as Jackson says (somewhere) the impression that all analyses are subject to counterexample may be caused by a biased sample. We are especially interested in live philosophical questions. Hence we only look at analyses which are 'up for grabs', and ignore those which are settled as being right. So the sample of analyses currently under philosophical discussion will, in general, have a higher proportion which are vulnerable to counterexamples than the set of analyses in general. The quote from Harman above is intended as part of a response to this. He argues that the kind of counterexamples Quine, Putnam and others came up with in the 50s were counterexamples to what had been offered as paradigm cases of analysis. If even these failed, then we really did have reason to believe that no analysis could succeed. And this seems fair enough; the reason these analyses were being used was not because they were intrinsically interesting, as an analysis of 'knowledge' would be, but because either (a) they were being used as examples of analytic truths for undergraduates or (b) they were being attacked by anti-analytics.

Fourthly, it might be that all this shows is that we were looking in the wrong places for paradigm analyses. A quick browse through the OED (my favourite source of analyses) suggests that the following analyses are all correct:

dime = coin + worth tenth of a dollar
hexagon = polygon + six sides
harp = musical instrument + roughly triangular frame + consists of series of strings + played by plucking strings
flag = piece of material + attachable to staff or haylard + used as standard, ensign, decoration or display
brook = small + stream
receipt = written object + acknowledges something has been received
shop = building + where goods are made or prepared for sale and sold
newspaper = printed publication + consisting of folded sheets + containing news
subway = tunnel + for use of pedestrians or vehicles
noun = word + used as name or designation

If any of these are right, they show that the classical theory is right some of the time. But if the theory is ever right, it is hard to see what the force of the inductive argument is.

Finally, the inductive argument relies (implausibly to my mind) on our ability to detect the applicability of conceptual terms. Let's go back to Harman's example of the separated but not yet divorced man. The classical analysis says that he is not a bachelor, many folk say that he is. Why is this a problem for the classical analysis, as opposed to being a problem for the folk? Why not say that the folk are just wrong about this? I'll have quite a bit to say about mistaken

intuitions in later weeks, so I don't want to stress that point now. I just want to note that the inductive argument makes a very strong assumption about the infallibility of these folk intuitions, and without that assumption the classical analyses it questions may well be right.

To see this, compare the folk's dispositions to apply the term 'good person', in the sense of morally good. To adopt one of Harman's examples, the folk, or at least a distressing large percentage of them, will not apply the term 'good person' to a man who lives with another man over a long enough period of time (even if they are not married). This doesn't mean that people living in gay relationships are not good people; most of them are, and those members of the folk who disagree are just wrong. I think that just as the folk routinely make mistakes about who is good and who isn't, they routinely make mistakes about who is a bachelor and who isn't. Why should we ever have thought otherwise?

There is a possible rejoinder to this move which I've sometimes heard made turning on the fact that these are competent speakers making these ethical blunders. We'll discuss this more in later weeks, so I just want to flag it now. The rejoinder goes vaguely along the lines of: 'good' is defined by the way competent speakers use it, these people are not making verbal errors, they are making ethical ones, so we should take their use as constitutive of the meaning of 'good'. For now, I'll leave the issue of how to respond to that argument as an exercise for the reader.

1. For Next Week

The readings for week two (3 psychology papers) are set out on the other handout. There is a small assignment on the Rosch and Mervis paper which is also detailed on that handout.

Some people have expressed an interest in getting some background reading in before the rest of the course. The following readings will be most useful.

Any part of the **Jackson** book will be good to read, but especially chapters 1 and 2. We will not look at chapters 4 to 6 until the end of semester, and chapter 3 is quite tough unless you are well acquainted with the relevant literature.

Any of the papers in **Intuition** are valuable to read, but the main ones I would suggest looking at to get a feel for the issues we'll be discussing are those by Gopnik and Schwitzgabel; Cummins; Ramsey; Pust and Goldman; and dePaul.

Finally, I have put copies of two of the methodology papers Jackson wrote before *From Metaphysics to Ethics* in the filing cabinet. These are "Metaphysics by Possible Cases" and "Armchair Metaphysics". In each case the copy is from his collection *Mind, Method and Conditionals*, Routledge, 1998. Both papers are relevant to this course, and the first paper contains some discussions which are intrinsically interesting concerning temporal parts, and laws of nature.