THE SEMANTICS OF MORAL COMMUNICATION

by

RICHARD BROWN

A dissertation submitted to the Graduate Faculty in
Philosophy in partial fulfillment of the requirements for
the degree of Doctor of Philosophy, The City University of
New York

2008

This manuscript has been read and accepted for the
Graduate Faculty in Philosophy in satisfaction of the
dissertation requirement for the degree of Doctor of
Philosophy.

09/03/2008	Michael Devitt
Date	Chair of Examining Committee
09/03/2008	Iakovos Vasiliou
Date	Executive Officer

Stefan Baumrin

Steven Ross

Robert Fiengo
Supervision Committee

THE CITY UNIVERSITY OF NEW YORK

Abstract

THE SEMANTICS OF MORAL COMMUNICATION

by

Richard Brown

Adviser: Professor Stefan Baumrin

In the first chapter I introduce the distinction between metaethics and normative ethics and argue that metaethics, properly conceived, is a part of cognitive science. For example, the debate between rationalism and sentimentalism can be informed by recent empirical work in psychology and the neurosciences. In the second chapter I argue that the traditional view that one's theory of semantics determines what one's theory of justification must be is mistaken. Though it has been the case that there are 'typical' combinations of semantic and justificational theories this is by no means forced on one. In the third chapter I examine two current kinds of expressivism; that of Blackburn and Copp. Each gives us an example of different combinations of semantic and justificatory theory. In Blackburn's case he combines a use theory of meaning with a classic emotivist theory of justification. In Copp's case it is a neo-Gricean philosophy of language with a realist's

theory of justification. I find both of these wanting as they ultimately collapse into relativistic subjectivism. In the fourth and fifth chapter I turn to developing Emotive Realism. The basic idea behind this view is that when someone says 'x is wrong/right' that person has (a) expressed a moral emotion about x and at the same time (b) expressed the belief that the emotion in (a) is the correct one to have. The belief expressed in (b) will be true or false depending on one's theory of justification. In the fourth chapter I argue that we need separate semantic theories for thoughts and sentences. In the fifth chapter I fill in the details of Emotive Realism as I see it.

Acknowledgments

Being the first in my family to attend college, let alone complete a dissertation in philosophy at a top university, makes this a very special moment in my life. But it being me, in particular, to be the one to do this makes it a momentous occasion in my life. I was the last person in my family anyone ever expected to go to college. I was a bit of a hoodlum and had an unusual childhood, to put it mildly. I will spare you the details but at the age of 13 I stole a motorcycle and attempted to run away from home. I was arrested and for the next 5 years was bounced back and forth between juvenile hall and various group-homes in California. Those were hard times, and there is a lot I could say about what happened during that time (and before it), but I made it through and was released at the age of 18.

I was on a very bad track with my life. Being locked up for five years had only made me more passive aggressive, and actually aggressive, than I was before. I had also learned to become a better criminal, and I am ashamed to admit my participation in multiple petty crimes. I bounced from job to job and city to city having a hard time adjusting to life outside a correctional institution. I ended up becoming a night manager at a Burger King in Morro

Bay California and playing drums in a Death Metal band. Like I said; a very bad path. In 1994 I found out from my mother that since I had been in a group home I was eligible to join a program that would help me go back to school. I figured I would give it a shot and with the help of the Private Industry Council, I enrolled in Cuesta Community College in San Luis Obispo at the age of 24.

My first semester there I took an introduction to philosophy course and immediately declared philosophy as my major. I remember being awe struck by the discovery that the questions that had so vexed me my whole life had a long history of being discussed. I wanted to be a part of that discussion. I can honestly say that that decision changed my life. Studying philosophy has changed me in more ways than I can count and I am extremely grateful to have had the opportunities that I have; though it has by no means been a cake walk! I owe a great many thanks to a great many people for helping me with the ideas in this dissertation and I want to take a moment to thank some of them now.

The basic idea behind this dissertation came to me very early in my philosophical career. It was in my junior year at San Francisco State University way back in the spring of 1997. I had just transferred from Cuesta and was taking my first upper division courses in philosophy (I was

also homeless and living in the 24 hour study section of the SFSU library, but that is another story). I had a 'philosophical analysis' course which focused on teaching us the basics of philosophical research and paper writing. We read several short pieces by prominent contemporary philosophers and had to write several short papers just laying out the positions and arguments. One of the pieces we read was Stevenson's "The Emotive Meaning of Ethical Terms". I had never heard of emotivism before and was immediately struck by it. It seemed to me to capture an obvious fact about moral discourse and I became an avid convert.

At the same time I was taking a philosophy of language course with Kent Bach which focused on meaning, speech acts and communication. This course made a big impact on me, one that has lasted to this day. I immediately saw a connection between the Austin/Strawson/Grice inspired philosophy of language and the views that Stevenson had. My paper was entitled 'Stevenson's Impact on Contemporary Philosophy of Language' and was my first attempt to spell out the ideas in chapter 2 of this dissertation.

But at the time I became much more interested in the philosophy of mind. In particular I was interested in issues about materialism and consciousness. In fact my

original plan was to complete my Masters in Philosophy at San Francisco State University under the supervision of Kent Bach. My proposed thesis was on the prospects for Representationist accounts of sensory qualities. I was at the time very incensed at accounts of consciousness that explained it away and very much a fan of the views of John Searle and Ned Block. I owe an immense debt to Kent and I want to thank him very much for his support and guidance. I have thought several times that if it weren't for him I would never have made it into either CUNY or UConn. He helped me edit my writing sample, my letter of intent, and my CV. He also wrote me a letter of reference, which though I have never seen, at least did not stop me from getting into the Grad center, so I assume it couldn't have been all bad.

Also from the SFSU philosophy department I would like to thank John J. Glanville and Helen Heise both of whom nurtured my philosophical development and worked very hard to help me come to a deeper philosophical understanding of myself and the world I live in. To this day I model my courses on aspects of their teaching style and strive to emulate them as researchers and teachers. In short studying philosophy at San Francisco State University changed my life forever and I am eternally grateful for the

opportunity they afforded me. Though I chose it purely because I wanted to live in San Francisco, I couldn't have made a better choice of places to study philosophy.

At the same time I was working in the psychology department with Mark Geisler doing EEG and ERP studies using scalp recordings. I completed all of the requirements for the MA in psychology but could not get both degrees simultaneously so my degree was to officially be in philosophy. Mark was very much a mentor to me and I owe him a huge debt of gratitude for taking me into his lab and being my friend. From him I learned how to be a professional academic and an engaging, approachable teacher.

In 2002 I transferred to the University of Connecticut, where I intended to get my PhD working with Austen Clark on naturalistic accounts of consciousness. While in Connecticut I met and studied with Ruth Millikan who has had a large impact on me. Austen and Ruth are two of the kindest philosophers I have ever met. I have never felt so supported and invested in as I did when I was in Connecticut. It is indeed, as Ruth often remarks, a wonderful place to do philosophy.

My first semester there I had a seminar on Moral Realism that I was not particularly looking forward to. But

as it turns out I enjoyed it immensely. It was there that I was introduced to Blackburn's kind of expressivism. We read *Ruling Passions* and I was quite taken by it, though I did not like the deflationism about truth or the anti-Gricean philosophy of language. It was there that I was also first introduced to Phillipa Foot's work, which has made a lasting impact on my thinking. Coincidently I was also at that time zeroxing some of Kent Bach's papers so that they could be scanned and turned into pdf's. I had access to a pdf maker and volunteered to turn his papers into pdf.

It was at that time that I first realized that if one had a certain kind of philosophy of language in mind one could formulate a kind of expressivism that could be realist. I called it a 'real emotivism' and spent the latter half of the Fall 2002 semester writing it up for the seminar; though I did not end up turning it in, for reasons that I won't go in to. But I want to acknowledge that this time in Connecticut was very fertile for me and without it I probably would never have come up with the specific idea for this dissertation.

As I said, I originally intended to get my Masters from SFSU and my PhD from UConn, but by the end of my first semester at UConn it became clear to me that I could not live there. Being from California I was not used to the

winters and being a city boy at heart I was not used to the total seclusion. In short, I freaked out and decided to go to New York and take CUNY up on its offer. For that reason I decided to take my Masters degree from Connecticut.

I went to the Graduate Center fully intending on working with David Rosenthal. Both Austen and Kent fully endorsed the move and the proposed working relation. Austen confided in me that Rosenthal was 'a real mench'; I found out later that that is a good thing. My first semester at the Graduate Center was extremely stimulating. I took a seminar with Blackburn and Horwhich which was very interesting. In one class I floated the idea I was then calling Real Emotivism and Blackburn said 'something like that might work'. Encouraged, I wrote up a new and improved version of the earlier paper and submitted it for the course.

I never received any comments from either Blackburn or Horwhich, but they must not have liked the paper because I received a B- in the course, the lowest grade I have ever received in graduate school, and which I take to be basically a failing grade. When I contacted Blackburn to inquire he said he did not remember why the paper had received that grade and that I should ask Horwhich. When I asked Horwhich he said he didn't grade the paper and that I

should ask Blackburn. At any rate, I do owe a lot of the content of chapter 3 to that seminar and discussion with Fritz McDonald that resulted from it so I am grateful for that.

But still, at this time, I completely anticipated working with David Rosenthal and I was taking all of his courses and participating in the Cognitive Science group that meets on Fridays. And though it is perhaps not evident from the work in this dissertation David's influence on me has been profound. He is a truly excellent philosopher and he has taught me more about how to actually do philosophy than anyone else that I can think of. I also owe to him the coinage of the word 'frigidity'. It was in a paper I wrote for David's philosophy of mind course. In it I talked about my interpretation of rigidity as being our ability to freeze the referent of some description for purposes of evaluating a counter-factual about it. David made a comment to the effect that 'if I wanted to use 'rigidity' that term was already taken by Kripke'. I remember thinking to myself right then 'fine, so I'll use "frigidity"'.

During my time at the Graduate Center I have written versions of the arguments in this dissertation for all of the members of my committee and I want to thank them all for their help and guidance in formulating the final

version I have here now.[1] I especially want to thank Steve Ross who has spent much time reading (and re-reading) my work and discussing these ideas with me. All of his comments and our discussions of them have been challenging and thought provoking. Incidentally, it is due to him that I changed the name to Emotive Realism.[2] Our interactions have made me fully understand what it is that I have been trying to say and it was in a seminar with him that I first entertained the idea that emotive realism could be a dissertation topic.

What really cinched it for me, though, was a combination of two separate events. The first was my lack of funding at the Graduate Center and my subsequent need to teach from my first semester in New York. When I moved from Connecticut I immediately applied to the local CUNY undergraduate schools and had interviews at John J. College and Brooklyn College. I ended up getting offers to adjunct at both colleges and I took the offer from Brooklyn College. It turned out that I made the right decision because I thoroughly enjoyed working at Brooklyn College.

---

[1] Also on rereading this I see that I have not properly thanked Bob Fiengo for his patience and support. He has gone above and beyond the call of duty both in the classroom and on my committee.

[2] Thinking back on this now I realize that at this time I was thinking that the primary speech act we performed was an expression of moral emotion and that we indirectly expressed the belief that that was the correct way to feel about. It was at Steve's urging that I came to see that it is better to simply say that both are expressed and abandon the direct and indirect business.

In particular I want to thank Emily Michael, the chairperson of the department of philosophy, for providing such a warm and supportive environment to learn how to teach philosophy. Emily offered me countless pieces of excellent advice and the option of teaching a great variety of courses. In particular it was at Brooklyn College that I first taught ethics. I found, to my amazement, that I quite enjoyed it.

The other event was the ethics qualifying exam at the Graduate Center. This was your typical four hour essay exam. I had taken one similar at SFSU on Plato, Descartes, Hume, and Kant so I was nervous but knew that if I worked hard I could do it. I studied furiously for that exam with the help of Ben Young and I quite enjoyed it, contrary to Ben who hated it. I took the exam and found out that I had received an Honors Pass, which was quite shocking. But I really started to notice that I was drawn to these metaethical issues.

I then audited a course with Stefan Baumrin on the British Moralists and found myself enthralled. My knowledge of ethics was increasing, and I found I was working more and more on my emotive realism voluntarily and at the expense of work due in other classes. I started to seriously consider doing a PhD in metaethics and not with

Rosenthal. It was the hardest decision I have made in graduate school, because I really do love those issues and plan to continue working in that area, and I really did admire and respect Rosenthal as a philosophical mentor. But I felt more inspired by the issues in metaethics, and I felt that I could actually complete this project, since I had been working on it for several years now. I approached Baumrin with the prospect and the rest is history.

Needless to say I want to thank Professor Baumrin for his help and guidance. Without his support I doubt that this project would ever have been completed. He is a first rate philosopher, and a first rate person who genuinely cares for his students. I for one have benefitted enormously from his tutelage. He has carefully read and thought about every word I have written and his criticisms and comments have improved this work, and the arguments therein, immensely.

I would also like to thank all of my cohorts in graduate school without whose countless hours of drunken conversation none of this stuff would have ever been worked out. In particular I would like to thank Josh Weisberg and David Pereplotchik for being such kind, generous, and patient compatriots as well as worthy philosophical adversaries.

Finally I would like to thank my girlfriend Jennifer Snider. Without her love and support these last five years I simply couldn't have done any of this. She is my best friend, lover, and companion and I can't imagine my life without her.

So, as you can see, this dissertation is the product of a life-long struggle with these ideas and I am deeply grateful to all of those people who helped and guided me along the way. My life will never be the same for having met you. It is with deep gratitude, humility, and respect that I submit this dissertation for partial completion of the requirements for the PhD in philosophy at the Graduate Center.

Richard Willie Hayes Brown
June 19th, 2008
Brooklyn, New York

Table of Contents

Chapter 1. Introduction

1.1. Philosophical Ethics

Of all the areas of philosophy ethics is easily the one that has the most relevance to our day to day lives as members of society. Indeed, the questions of how we ought to act or of what kind of person we ought to be are not unique to philosophy. These are some of the earliest questions that human beings asked; right up there with questions about the origins and nature of the world we find ourselves in. This is evidenced by the various codes we have from the people who lived in ancient Sumaria. Perhaps the best known of these early writings is the so-called Code of Hammurabi, believed to have been compiled some 1700 years before the birth of Christ.[3] What is perhaps most surprising about these very old rules is that the issues that are dealt with are the very practical issues of people who regularly interact with each other. They are the day to day issues of people who have to deal

---

[3] Scholars used to date his life to the 21st century before Christ, but there is a lot of evidence that has surfaced that dates his life to the 17th century. Here is a link to an article by the historian Immanual Velikovsky that details the change http://www.varchive.org/ce/hammurabi.html

with other people who don't keep their word, or are in other ways dishonest.

One might say that this pre-philosophic approach culminates in the ten commandments of the Mosaic tradition. In this pre-philosophic way of thinking the answer to the question 'why should I do this?' is always based on the authority of some personage, whether God or someone chosen by God. What is distinctive of the philosophical approach to ethics, then, is not (necessarily) the questions that are addressed, as these are perennial. Rather it is the approach that is adopted in answering them that distinguishes philosophical ethics from other forms of inquiry.

Socrates typically gets the credit for initiating the philosophical branch of ethics, 400 years before the birth of Christ, by advocating that human beings could determine the truth about morality via the use of reason and argumentation. Though he famously became disillusioned with the metaphysical investigations of his predecessors, (that is, if we are to believe the account given in the Phaedo), he never lost his passion for reason as a genuine source of knowledge about the right way to live, the kind of person we should strive to be, and the nature of moral properties and justification. We can figure out the answer

to questions like 'what is the right thing to do?', 'what is good about justice and courage?', and most importantly, 'why should I be moral?' independently of the commands of the gods or the wisdom of the poets. So, when Hammurabi or Moses claim to have the Laws handed down from God we, as rational creatures, can discover the reasons that God would so command us *via* the use of reason and argument. The point of the <u>Euthyphro</u> question is that if God commands us not to murder then it must be because there is something about murder which is wrong, and we can then endeavor to discover what it is about murder that would be objectionable to a supremely good being (should there turn out to be one).

1.1.1. Metaethics & Normative Ethics

This way of construing the task of philosophical ethics naturally leads to two different kinds of questions. On the one hand we have the question of what is the difference between good and bad actions or good and bad persons. By what mark do we classify actions or persons as right or wrong, good or bad, just or unjust, moral or immoral, etc? Answering this question is traditionally the province of what is called 'normative ethics'. It is in this branch of ethics that we find the traditional kinds of ethical theories like utilitarianism, deontology,

contractualism, constructivism, and virtue ethics. Each of these theoretical perspectives attempts to give a certain kind of answer to the normative questions. Though each type of theory does give a different kind of answer to the normative questions, they are all united by the assumption that there really are right actions and good people, as well as bad people and wrong actions. That is, they all assume that there is a correct way to sort actions or persons into good and bad or right and wrong. The main thing in dispute is how we do that.

This is of course to beg the question against the relativist. What if it is the case that there are no right actions because there is no such thing as an action's being truly right or wrong, good or bad? Maybe the truth of the matter is that standards of right and wrong are more like standards in fashion. Is there really a right answer to whether or not one ought to wear black dress socks with shorts? Or whether one should always tuck in a button-up shirt when wearing a tie? If there is a right answer to these questions it is probably only relative to what happens to be fashionable at the time. If something like this were true about morality then trying to give a normative ethical theory would then be a waste of time. This makes the issue about relativism in some sense prior

to the normative questions. The relativists' challenge must be met before we can even begin to theorize about the nature of right actions or good people.

This prior question falls under what has come to be called 'metaethics' because it is a question about the possibility of (normative) ethics itself. Socrates deals with this question in the Gorgias but Plato is perhaps the first to see that in order to answer this question one needs a theory that tells us how there could be a right answer. He is the first western philosopher to try to construct a theory that would explain how moral judgments could be true and so how relativism could be false. I suppose we should attribute the discovery of metaethics to him.

So in The Republic he begins by having Thrashymachus set out the challenge to morality in Book One and then he develops his theory of the forms as the foundation for the answer. Famously, the forms were non-physical, eternal, and unchanging objects and included in these were moral properties. Justice for Plato was a mind-independent property that was perfect and eternal. All just actions and people were just due to their 'participation' in this absolute form of Justice. To say that an action is just is to attribute to it the property of participating in the

Form of Justice. He then goes on to theorize that moral words stand for these moral properties just like ordinary physical predicates stand for ordinary physical properties.

Metaethics as so defined is primarily concerned with questions about the *meaning*, or status, of ethical terms like 'good,' 'evil,' 'ought,' 'obligation,' and 'right,' and the possibility of the *justification* of normative moral judgments like 'suicide bombing is morally wrong,' 'Uday Hussein was an evil man,' or 'Humans ought not to eat meat'. Now, though we are concerned with the possibility of the justification of these normative judgments, we are not concerned with giving a theory that would tell us, or purport to, whether these judgments are actually correct. Metaethical inquiry is concerned only with the nature of the kind of answer that can be given, not the actual answers that are given.

So, for instance, Plato's answer that there are eternal, perfect, and unchanging Forms of Justice and Courage tells us how a normative judgment could be true. It does not tell us which ones are. His account amounts to the claim the moral judgments are beliefs that are true in so far as they capture reality as it is in the Eternal Realm of the Forms; just as normal predicates work on his view. Telling

us what objects, or states of affairs, do participate in these Forms is the job of normative ethics. In Plato's case the normative theory takes the form of a virtue ethics based on his analogy between the parts of the soul and the parts of a city. While Plato's normative theory has fallen out of favor, his metaethical theory remains quite popular, but I shall not dwell on this here. My point is that the proper task of metaethics lies in giving a general theory about the nature of moral judgments, the semantics of moral words, and the status of moral properties that would explain how realism could be true, or is false, or whatever.

Metaethics has traditionally been seen as comprising a distinctively philosophical inquiry that is separate from any kind of empirical or scientific inquiry. As William Frankena puts it in Ethics (Frankena 1963),

> [metaethics] does not consist of empirical or historical inquiries, nor does it involve making or defending any normative or value judgments. It does not try to answer either particular or general questions about what is good, right, or obligatory. It asks and tries to answer logical, epistemological, or semantical questions. (p 4)

These logical, epistemological, and semantical questions occupied center stage in Anglophone philosophy through-out most of the Twentieth Century. So much so, in fact, that

they were widely taken to be the only interesting questions that philosophers could (or should) address.

This led to a rather strange situation in analytic philosophy. There they were, in the middle of some of the most gruesome and horrific carnages mankind has ever born witness to, by which I mean the two World Wars and the Holocaust, and philosophers studiously avoided any substantive normative debate, as any cursory glance at the history of this period will bear out (See, for instance, Soames 2006; Soames forthcoming). As the story goes, this stale environment was cleared in America by the fresh breeze of John Rawls and his work and in Europe by the work and influence of the existentialists. The situation is much improved now, though it is curious that Saul Kripke, who is widely regarded as perhaps the greatest living analytic philosopher, has no interest in normative ethics or metaethics (though I do think, as we will see, that his work is relevant to metaethics).

Another thing that has cleared, at least for the most part, is the view of philosophy that Frankena expresses in the passage quoted above. I agree that metaethics asks and tries to answer logical, epistemological, and semantical questions. These are certainly the questions that I will be addressing in this dissertation. I also agree that

metaethics does not make or defend any particular normative judgments. Nor should it, nor for that matter, shall I be concerned with any general or particular question about what is right, good or obligatory. That is the province of normative ethical theory and is a separate theoretical endeavor; in fact in the next chapter I will argue that no normative theory implies any metaethical theory or vice versa. My interest is only in the metaethical questions and I shall have virtually nothing to say about normative ethics other than arguing for its independence from metaethical theory (though I cannot fully contain my self and shall have something to say about these issues in the final chapter, but we will come back to that).

I do not, however, agree with Frankena's pronouncement that metaethics does not engage in any kind of empirical investigation. Frankena says this because at the time that he wrote philosophy was dominated by a kind of view which, among other claims, held that there was a strict split between empirical questions and philosophical ones. But since that time there has been something of a revolution in philosophy which we might call the 'cognitive turn'. This revolution was initiated by Quine's view that philosophy was continuous with the natural sciences and

has blossomed into the incorporation of philosophy into cognitive science. In the next section I will say something briefly about both of these claims as they will define the project I undertake in this dissertation.

1.2. Empiricism as the Default View

Let us define some terms. Following David Armstrong I will use *Naturalism* to name a metaphysical thesis which claims that everything that exists does so entirely in one single space-time system. Physicalism, or materialism, is the more restrictive (metaphysical) thesis that everything that exists in space-time is completely composed of matter, or what a completed micro-physics says. So, both idealists and materialists/physicalists can be naturalists in this sense as long as the non-material mind exists completely in one single space-time. The non-naturalist is then someone, like Plato, Descartes, and even Kant on some readings, who posits entities that exist outside of space and time. A non-*natural* property is one that does not even supervene on the physical. A non-*physical* property, by contrast, is simply one that is not composed of the postulates of a completed micro-physics, though it may still supervene on the physical.

These claims are to be distinguished from *Empiricism* which is the epistemological thesis that claims that the only way to acquire knowledge is the empirical *a posterori* way used by science. What is this method of acquiring justification? As Michael Devitt says,

> An answer starts from the metaphysical assumption that the worldly fact that p would make the belief that p true. The empirical justification of the belief is then to be found in its relationship to experiences that the worldly fact would cause. Justified beliefs are produced and/or sustained by experiences in a way that is appropriately sensitive to the way the world is. This is very brief…Still it is hard to say much more.

As against Empiricism we have Rationalism which is the view that some truths are known independently of experience or *a priori*. Usually it is claimed that we know necessary and universal facts about reality that are justified completely by reason and in no way by experience. So we can see that the tenants of empiricism and natuarlism go together while rationalism and non-natualism have usually gone together.

Armstrong's main argument for naturalism relies on what he calls the Eleatic principle (Armstrong 1978). This principle says that we ought not to posit the existence of entities that have no causal powers. What evidence could we possibly have that such a thing existed? So, take our proposed non-natural property. The question then is 'does

it have causal powers in the natural world or not?' If the answer is 'no' then the Eleatic principle councils us to reject the posited entity while if the answer is 'yes' we now have posited a kind of causation that is fundamentally different from the kind that we find in the natural world. How would such causation work? Besides which, we have good evidence, from science itself, that the natural world is causally-closed and so if the empirical way of knowing sketched above is indeed a way to get justification for our beliefs, whether it is the whole story or not, then we have reason to believe that the natural world is causally closed and so for naturalism.

The Eleatic principle is also wielded as an argument from naturalism to materialism. If a non-material entity has no causal powers then according to empiricism we cannot know about it and so would never have any reason to posit its existence. If it does have causal efficacy then it is of a type that is completely mysterious and unlike anything that we have hitherto encountered. Materialism is itself an empirical hypothesis and so far it has been very well borne out. The reason that we are able to posit the causal closure of the physical world just is that we have never had to appeal to any non-physical properties in a successful explanation.

So, if empiricism is correct then the Eleatic principle offers strong support for naturalism and from naturalism to materialism/physicalism. But why should we adopt empiricism? Aren't there many clear-cut examples of *a priori* knowledge? Isn't the fact that nothing can be both round and square at the same time or that the number 2 is even and prime obvious examples of *a priori* knowledge? Before answering this, we should point out that rationalism has suffered several notorious and embarrassing setbacks in the last few hundred years. Perhaps the first of which is Gallileo's showing empirically that Aristotle's physics was wrong in assuming that heavier bodies fall faster than lighter bodies. For centuries it had been held to be self-evident that a heavier body would fall faster than the lighter one. This was just too obvious to test.[4]

More dramatically, perhaps, is the discovery of non-Euclidean geometry which showed that various of the fundamental postulates of geometry were not necessary (like the parallel postulate). It was thought to be a necessary fact about reality that a triangle couldn't have more or less than 180 degrees of interior angles or

---

[4] Though I am aware that some Islamic scientists were testing this. I meant my claim to be limited to those in the Western tradition.

that parallel lines couldn't meet. But this was quite wrong. We should perhaps add to this the discovery of Russell's paradox and the various self-referential semantic/syntactic problems (i.e. Godel, the liar, etc) which have led to the development of alternative logics.

What this shows is that intellectual seemings are fallible. It cannot therefore be argued that something *seeming* to need no empirical support and *seeming* to be justified purely by reason is enough to establish that the fact in question is *really* justified independently of experience or not. So, take 1+1=2. It certainly seems that this is true, in fact it is hard for me to imagine how it could be otherwise. There is a strong subjective sense of certainty that I experience when I think about it. All of this is no doubt true. But we have as of yet no reason to think that it is REALLY necessary, or that its justification is independent of experience. This is because we can not tell *a priori* whether the intellectual seeming is indeed correct. This shows that rationalism is in serious trouble. There is no other reason to take rationalism seriously other than the strong pull these rational insights have on us.

One might want to reply by saying that this is overly skeptical. We shouldn't abandon *a priori* knowledge just

because **we** have mistakenly identified some things as necessary which weren't. So too, the objector continues, just because we hallucinate doesn't mean that we don't normally see objects. Fair enough. But then what we need is an actual account to back this up. What is the difference between the cases? We can give that in the empirical case. We can describe ways in which we could find out whether the person was hallucinating or not based on our ability to monitor the brain of the person in question and our visual impressions of the experimental set up. We can give a sketch, if not every detail, of a story which describes how the brain (causally) interacts with the environment it finds itself in and generates representations of that environment. But can you do the same for rationalism? To date no one has. What is an eternal, necessary, non-physical/non-natural object like a number or *modus ponens* really like? How do we interact with it? No one knows. How could they?

This is to echo Devitt's abductive argument against rationalism (Devitt forthcoming). He argues that when we have two competing theoretical explanations we should opt for the one that is better understood. In this case we have the thesis of empiricism and the attendant argument for materialism, on the other hand we have rationalism and

the attendant thesis of non-naturalism. The empirical thesis is worked out in some amount of detail, though by no means complete, whereas the rationalist thesis is not even worked out to anything like the level of detail that the empirical thesis is. In fact most accounts simply rely on analogies with the empirical way of knowing.[5]

Now this would be a pressing concern if it were impossible for us to fully understand the world we live in except for the truth of rationalism. But this certainly isn't the case. We have good candidates for materialistic accounts of every disputed area. For instance, with respect to the nature of the mind we have the mind-brain identity theory, the compositional theory of mind, and the higher-order theory of consciousness. I do not mean to say that we know that they are true, but only that they are viable candidates. For all we know right now they could be true. They have not been refuted by any a priori arguments, nor have they shown themselves to be inconsistent with the findings of science, quite the converse actually.

---

[5] For instance, Descartes compares an idea's being clear and distinct to 'taking something in hand in clear light and examining it'.

As for math and logic we have either constructivism or Mill's view that they are empirical generalizations,[6] or a modern version of deflationism or fictionalism. Again, none of these has been demonstrated to be correct. The point, rather, is that we should prefer naturalistic/materialistic accounts over their non-natural/non-physical competitors. They automatically become more plausible because of their reliance on the more plausible empiricist/scientific account of knowing.

The same is true of our metaethical theory. We should prefer a metaethical theory that is entirely materialistic and empirical in nature. We shall see in the next section that this puts us in a bind as it would seem to commit us to some form of relativism. The primary purpose of this dissertation is to show that we can have an acceptable naturalistic metaethical theory that does not commit us to relativism. I will get to that shortly but first let us turn to looking at the impact of a commitment to empiricism on one's metaethical view.

---

[6] I interpret this to mean that they are an attempt to model the way that the physical world works and to grow into Quine's indispensibility argument that the justification for mathematic is empirical

1.2.1. Metaethics Naturalized: Metaethics meets Cognitive Science

'Cognitive science' denotes a diverse field of inquiry which aims at understanding the human mind. It is comprised of researchers in such disjoint fields as neuroscience, psychology, sociology, anthropology, philosophy, computer science, and linguistics. Each of these disciplines has in its own way been involved in some aspect of the study of the nature of the human mind and its place in the natural world. The development of cognitive science as an independent area of enquiry signifies the growing recognition by researchers in all of the fields mentioned above that real progress can only be made by a sharing of information and results across disciplines. With this recognition comes a rethinking of the place of philosophy in relation to the sciences. It is no longer feasible for a philosopher to lock himself into a cabin for a month or two and set about deducing the *a priori* truths about reality. Experimental results are vital for philosophical theories and philosophers ignore them at their peril.

The cognitive turn in philosophy begins in the mid 1940's and '50's as a reaction to the behaviorism of the day and has had far reaching effects in philosophy. In the

philosophy of mind, for instance, the cognitive turn leads first to the mind-brain identity theory of Place, Smart and Feigl (and later Armstrong) (Feigl 1967; Armstrong 1968; Smart 1991; Place 2004). These philosophers saw that theoretical work in the brain sciences could be used to solve philosophical problems (Brown 2006). Then, when there seemed to be conceptual problems with the identity theory, and inspired by the developing field of computer science and artificial intelligence, we saw the rise of the now dominant view in the philosophy of mind known as functionalism (Putnam 1967/1991). Though even now naturalistically minded philosophers are arguing the merits of the identity theory over functionalism (Polger 2004).

Another area of significant interaction and success is the story of theories of concepts. Philosophers have long been interested in concepts; indeed it was a commonplace of the first-half of the last century that the role of philosophy lay in conceptual analysis. Sadly no conceptual analysis was ever completed and so, with the urging of Quine, philosophers reluctantly looked at what was going on in psychological work on concepts and became integral to developing proto-type theories of concepts as an alternative to the standard 'concepts are definitions'

line that had been abandoned in philosophy and linguistics. This area is still thriving with debate about Jerry Fodor's conceptual atomism (Fodor 1998). Fodor is critical of proto-type theories for failing to meet the 'compositionality' criterion.

Now I do not want to dwell on these subjects excessively but I am trying to point out that ethics sticks out as the one area of philosophical enquiry that has resisted influence from the cognitive turn. This is perhaps due in large part to the popularity of the fact/value distinction. Science is in the business of determining facts, ethics is in the business of evaluating values; never the twain shall meet. Virginia Held has voiced this concern in a particularly clear way (Held 1996). As she says,

> I…argue that cognitive science has rather little to offer ethics, and that what it has should be subordinate to rather than determinative of the agenda of moral philosophy. Moral philosophers often make clear at the outset that moral philosophy should not see the scientific or other explanation of behavior and moral belief, or the prediction and control that science has aimed at, as our primary concerns. Our primary concern is not explanation but recommendation. I start from this position: ethics is normative rather than descriptive. (p.69)

Held is not here denying that cognitive science has any application or relevance to ethics; she is denying that it can provide any real insight into the normative questions.

This may seem to be obviously true for normative ethics. It is hard for me to see how neuroscience or psychology could show us which theory of justification was in fact the correct one. Yet, of course, there are philosophers who think that cognitive science is vital for normative ethics. One such philosopher is Mark Johnston. Responding to the kind of criticism that Held advances he (Johnston 1998) says,

> Every moral tradition and every moral theory necessarily presupposes some specific view of how the mind works and what a person is. The cognitive sciences constitute our principle source of knowledge about human cognition and psychology. Consequently, the cognitive sciences are absolutely crucial to moral philosophy (p 691)

This is certainly true of normative theories. Think of the kind of utilitarianism that Mill advocates. Whether or not psychological hedonism is true is an empirical claim about human beings. It is the claim that, as a matter of fact, we are so built so as only to desire pleasure and the absence of pain. This is exactly the kind of question that the cognitive sciences are in a position to answer and so I find myself in agreement with Johnston here.

This is especially true if one thinks, as Mill did, that the truth of psychological hedonism establishes the truth of what is called ethical hedonism, which is the claim that only pleasure and the absence of pain are

intrinsically good or desirable for their own sake. If this is the case then the viability of the Greatest Happiness principle depends on the empirical question of whether psychological hedonism is true or not (in so far, that is, as the Greatest Happiness principle depends on ethical hedonism being true). However, I will leave the question about the relevance of cognitive science to normative theories to one side and instead turn to talking about the relation of cognitive science to metaethics.

Whatever the relevance of cognitive science is to normative ethical theory it seems obviously relevant to, and vital for, metaethical theories. Consider the traditional questions that metaethics addresses. These, as I have already said (section 1.1.1.), are questions about the meaning of moral terms and the nature and possibility of the justification of normative judgments. The question about the meaning of ethical words is one that has been addressed at length by philosophers of language and traditional metaethicists, but as it is a claim about semantics it should be informed by work in formal semantics, logic, linguistics, and psychology.

It is even more obvious that questions about the nature of moral judgments belong to cognitive science. Metaethical theories that emphasize the role that the

emotions play make all kinds of claims about the nature of the mental states in question and their relation to and influence of cognitive states like beliefs. There has been a lot of investigation into the nature of the emotions done in psychology and neuroscience and we want a theory that is informed by these advances in cognitive science. If it were to turn out that some empirical discovery in, say, neuroscience, showed that the emotions do not play these roles in our moral thinking this would be evidence that the theory in question is false. As a case in point let us briefly examine one of the major controversies in metaethics.

Looking at the history of moral philosophy shows two overarching camps in metaethics known as Rationalists and Sentimentalists which roughly correspond to the rationalist/non-natural and empiricist/natural split discussed in the previous section (1.2.). On the one hand we have the rationalist tradition that views moral truths as self-evident truths known via reason. This tradition is in the Platonic tradition and seems to be inspired by mathematical truths. On the other we have the sentimentalist tradition that views moral truths as either reactions to natural phenomena, like secondary-properties, or as emotional reactions like our attitudes of taste. The

guiding metaphor in this sentimentalist tradition is art (Gill 2007). Each of these traditions has important insights to offer and each tradition has a set of powerful motivations as well as a powerful set of objections.

The sentimentalists have always recognized the important role that the emotions play in our moral life, and this has begun to be borne out in recent neuroscientific and psychological studies. There are numerous studies that indicate that the emotions play more of a role in moral decision making than the rationalists would like to admit (for a nice survey of current research see (Greene and Haidt 2002)). For someone like me, who sees science as natural philosophy this means that sentimentalism captures some aspect of the empirical metaethical truth. The emotivists of the last century developed powerful tools in the philosophy of language to account for how our moral language is related to these emotional states and these insights should be learned from.

The rationalist has traditionally held that the emotions mislead us. Only reason, cool, calm, impartial reason, free from the fetters of passion, could deliver moral truth. On their side they have the very powerful pull of moral realism (what that is, I will put off until the next chapter); and I must say that I feel this pull. If the

sentimentalists are right then it seems to many that morality is just a matter of opinion; but how could it be a matter of opinion that one should keep their promises? Or that causing unnecessary suffering is evil? Whether Hitler, Stalin, Polpot, or Saddam Hussein were really evil would then also be merely a matter of opinion. But this offends against common sense! We could not say that one way of life was any better or any worse than any other. Killing innocent civilians for pleasure seems to be objectively wrong and the sentimentalist is traditionally seen as being unable to account for that.

The traditional problem with rationalism, especially in the Platonic tradition, has been its conflict with naturalism. The rationalists are seen as committed to an ontology that is not friendly to the naturalist. The idea that there are moral properties 'out there' that are not natural properties is very hard for a scientifically minded philosopher in the 21st century to take seriously. So we seem to be between a rock and a hard place. On the one hand we have our commitment to naturalism and so to the denial that there are any non-natural aspects to reality. Sentimentalism then seems to be the right theory and has some experimental support. A naturalist has got to

recognize that the emotions play an important role in making moral judgments.

On the other hand we have moral realism (which we will discuss in the next chapter) and the claim that there are some moral truths which conflicts with sentimentalism and so naturalism; but how can we give up the claim that the rape rooms that Saddam Hussein allowed under his dictatorship are anything but pure evil? How can we say that whether it is good to greet someone and wrong to slowly torture them to death depends simply on the way you happen to be feeling that day!?!?!?

One then seems forced to decide between naturalism with the bitter after taste of relativism or rationalism with its depressingly overpopulated ontology and incompatibility with science. And, as we have seen in the previous section, we have presumptive evidence for naturalism and, so it would seem, evidence that relativism is true. The incompatibility between naturalism and moral realism, it seems to me, is a far greater and more pressing problem than the incompatibility of relativity physics with quantum theory that physicists are dealing with.

Given these remarks as background we are now in a position to say something about the aims of this dissertation.

1.3. The Project at Hand

This dissertation is a study of the semantics of moral communication. I am interested in what kind of metaethical theory we can have given that we want (a) to be naturalists/empiricists and (b) a realist theory of justification. As such I am not concerned with what specific naturalistic realist theory of justification we give (see 1.1.1.). There are several promising candidates, ranging from the simple utilitarian ones to the more complex and subtle views of the Cornell realists and, more recently, constructivism of a Kantian kind to name some of the more popular. I shall have some brief remarks to make about these theories in chapter five but the main project of this dissertation is, I argue, completely independent of these claims about normative theory. This is because one of the findings of this dissertation is a kind of metaethical theory that clearly shows us the severability of these kind of normative considerations from our purely metaethical considerations. The split between semantics

and justification is mirrored by a split between normative theory and metaethics.

As we found in the previous sections, to construct such a metaethical theory we need

(i) an account of what moral judgments consist in that is consistent with the findings discussed in the previous section about the role of the emotions

(ii) an account of how it is possible that those judgments could be justified in a morally realist way and

(iii) an account of the meaning of the words and sentences that we use to express those moral judgments.

So the project is simple. I want to explore the options for accommodating i-iii.

The prospects for accommodating i-iii have traditionally been thought to be very slim. So, to begin this study we start by looking back at history of expressive metaethics and their traditional problems and objections from the moral realists. Doing this allows me to develop some basic ideas in the philosophy of language that we will use for the rest of the dissertation. It also allows me to discuss just what moral realism is supposed to be. What emerges from this study of the history of expressivism is that we have two distinct ways of carrying on the expressivist tradition depending on what kind of semantic theory we opt for.

On the one hand we have the kind of view descended from Ayer and the later Wittgenstein according to which the meaning of a word is simply its use. On the other hand we have the kind of view which is descended from Stevenson and Gricean theories of meaning which distinguish meaning from use. Though each of these semantic theories is in principle compatible with a realist theory of justification some combinations are more satisfactory than others.

To show this we next turn, in chapter three, to examining two such contemporary theories. The first is Simon Blackburn's quasi-realism. Blackburn is famous for arguing that his metaethical view can capture everything that the moral realist wants without any commitment to non-natural properties and at the same time having a fully sentimentalist account of moral judgment. If this were true then quasi-realism would be a very attractive candidate for satisfying i-iii. To deliver on his promise Blackburn invokes a use theory of meaning and a deflationary account of truth. After careful examination of this strategy I find that it is unable to satisfy condition ii.

That is, Blackburn's account, as given, is not really a morally realist kind of view in the sense to be defined

in the second chapter. He argues that his appeal to the use theory of meaning and deflationary account of truth obviates the need for any further account of the justification of moral judgments. We can say that our judgments are true. Blackburn asks; what more do you want? What else is there? But each part of the strategy fails. We have independent reasons not accept the use theory of meaning and the appeal to the deflationary account of truth does not excuse us from providing a real account of moral justification.

The second candidate for satisfying i-iii is David Copp's recent attempt to formulate a view he calls 'Realist-Expressivism'. He attempts to combine realism about justification with a Gricean philosophy of language that distinguishes between meaning and use and makes use of the pragmatic notion of implicature. This is an interesting and important kind of view but, as I will argue, it also fails to be able to be fully realist. The problem lies in his reliance on implicature. Without implausible assumptions his view turns out to be a radical kind of subjective relativism. The lesson from chapter three is that though one can in principle combine these theories (i.e. of meaning and justification) in any way one wants, one's choices are limited when one adds in the

constraint of a naturalistically acceptable realist theory of justification and no current metaethical theory meets these requirements.

This clears the way for us to consider other options for a semantic theory. In particular one that incorporates the lessons we have learned from our study so far. I think we can satisfy i-iii in a way that can actually be stated quite simply. In a nutshell the claim is that when someone sincerely and literally **says**, for instance, "eating meat is morally wrong", this person has

(a) **expressed** a **moral emotion** about eating meat

and, at the same time, with the same utterance

(b) expressed the **belief** that the moral emotion in (a) is the correct one to have

To **intentionally** express these two attitudes is to perform the speech act that I call **moral condemnation**. I think that, by obvious means, we could get a similar speech act for **moral approbation**. The belief expressed in (b) has a truth value that is determined by the correct naturalistic theory of justification, whichever one that happens to be. Our theory about how we use moral language to express moral judgments and our theory about the nature of the justification of those moral judgments are orthogonal issues.

I then turn to the task of constructing the alternative metaethical theory outlined above in the fourth and fifth chapters. The strategy I employ is as follows. What we need to do is first identify the most plausible candidate for a naturalistic semantic theory. This is the task of chapter four. I will argue that starting from this position we arrive at the view that we need separate semantic theories for thought and language. We need one theory for what sentence types in a language mean and another theory for what thoughts mean. I call these two kinds of theories P-semantics, for those that aim at explaining the psychological notion of how thoughts come to have meaning, and L-semantics, for those that aim at explaining how linguistic items, namely sentence types, come to have meaning. Seeing that we need distinct theories allows me to show how it is that one's semantic theory is independent of one's theory of justification; it is just the claim that L-semantics is independent of P-semantics.

To illustrate, as I said above I argue that a moral judgment consists in a combination of two mental states. One is a moral emotion (like condemnation or approbation) and the other is the belief that the emotion is the correct one to have. These mental states and their

contents can be explained in something like the causal theory of reference/language of thought hypothesis as defended by Michael Devitt. In brief, my condemnation of Saddam Hussein is about Saddam because of the causal relation my thought token stands to him. So too, my belief about the emotion is about that emotion due to its causal relations to the emotion. This kind of view is a molecular view. It holds that some thought contents get their content via causal links to the world and other thought contents get their content via inferential relations between it and other contents. This is a perfectly good P-semantic theory, one that I whole heartedly endorse, but this does not mean that we need to adopt this as an L-semantic view as well.

That is, we do not need to think that the English sentence type 'Saddam Hussein was a wicked man' means that the speaker has a certain moral emotion about a certain specific individual and believes that the emotion is the correct one to have towards the specific individual. I argue that one plausible candidate for word type meaning is simply whatever the speaker would need to know to count as using the word type correctly, which turns out to be roughly something like the dictionary meaning of the word.

We use these word types to build sentence types that we then use to express our mental states. An interesting consequence of this is that some moral sentence types will come out as analytically true. But we shouldn't let this disturb our naturalism as we can adopt the Quinian inspired account of analyticity as Devitt has argued.

What we end up with is a theory that is compatible with naturalism and realism. It allows moral sentence types to be straightforwardly true in a way that is not deflationary in nature but is nonetheless free of commitment to non-natural or non-physical properties. To put it very briefly, the separation of meaning and justification lets us say that the sentence type in which '…is wrong' occurs will be interpreted as having its dictionary meaning. So, the word 'wrong' will mean something like 'is morally blameworthy' and so a sentence like 'suicide bombing is wrong' will mean that suicide bombing is morally blameworthy, and the sentence will be true if it is and false if it is not.

Whether it is or is not morally blameworthy depends on what the correct account of justification turns out to be. If we have a naturalistic account of justification then the sentence will be true if our naturalistic account of justification determines that suicide bombing is a morally

blameworthy act. One nice thing about this kind of L-semantic theory is that it shows us how a key set of basic moral sentence types come out as analytically true (in an appropriately weakened naturalistically acceptable account of analyticity that severs the link between analyticity and *a priori* knowledge).

This allows us to see what moral debate is usually about. It is usually over how to classify some action. So, for example, suppose, as I argue, that 'murder is wrong' is an analytic truth. Then any action that is a murder is wrong. But a murder is an unjustified killing so to show that some action is a murder I must show that it is unjustified. This will immediately involve me in a normative debate that will require me to appeal to some theory of justification. I must show that the action doesn't promote the greatest happiness, or violates some duty. We do not argue over whether or not *murder* is wrong! We argue over whether what Smith did was murder.

None of this threatens our independent P-semantic theory that holds that a moral judgment is really a pair of mental states consisting of a moral emotion and a belief that the emotion is the correct one to have. Here too our theory of justification is separate from the specific account of semantics we have given. This is because we

assign truth-values to the belief aspect of the moral judgment that the sentence type is being used to express based on the best theory of justification. So at both semantic levels we see an independence from normative theory.

All of this brings up the question of whether it makes any sense to say that we can justify our emotional reactions to the world in any serious kind of way. In the final section of chapter five, where I discuss issues relating to justification, I argue that we can make sense of this claim and briefly suggest a constructivist theory of justification.

With this plan of attack in mind, let's begin.

Chapter 2. Language and Metaethics

In this chapter, to remind you, I want to do three related things. First I want to rehearse the history of emotivism in the 20[th] century. As we will see, emotivism starts as a theory that combines a distinctive thesis in the philosophy of language, viz. that moral words have no meaning and only serve to express our feelings, with a claim about the possibility of the justification of moral judgments, viz. that there isn't any such possibility. What will emerge from this rehearsal is a picture of a theory that gradually progresses towards separating the issues of the meaning of moral terms from the issues of the justification of moral judgments. I will argue in the final section of this chapter that separating these two issues from each other is absolutely vital to getting clear on the basic issues in metaethics and for developing a satisfactory theory of moral discourse that allows for the possibility of moral realism while at the same time incorporating the insights of the expressivist tradition.

The second thing that I want to do is to rehearse some basic distinctions in the philosophy of language that will be crucial for understanding the two views we will look at

in the next chapter. What will emerge from this twice-told tale is that there are two distinct ways to categorize metaethical theories. One such way is according to the views in the philosophy of language that are endorsed. The other is according to the theory of justification that is endorsed. The debate between moral realism and irrealism is a debate about the possibility of the justification of moral judgments and is independent of any theory about the meaning of moral words. This will suggest that we can mix and match these kinds of theories in any way that we want.

2.1. The Evolution of Emotivism

It had been long recognized that language can be used to do more than to merely describe the world. This is explicit in Berkeley, especially in Section 20 of the Introduction to his *Principles* (Berkeley 1710/1998). He there says,

> Besides, the communicating of ideas marked by words is not the chief and only end of language, as is commonly supposed. There are other ends, as the raising of some passion, the exciting to, or deterring from an action, the putting the mind in some particular disposition…I entreat the reader to reflect with himself, and see if it doth not often happen either in hearing or reading a discourse, that the passions of fear, love, hatred, admiration, disdain, and the like arise, immediately in his mind upon the perception of certain words, without any ideas coming between (p 99)

He even suggests that 'good' and 'danger' are examples of words that do not stand for ideas but rather serve to excite passions or exhort to action. This is mentioned in Warnock's <u>Ethics Since 1900</u> (Warnock 1960, p 64) but what she does not point out is that Berkeley is much more radical than this. He goes on in Section 20 to argue that even proper names "do not seem always spoken, with a design to bring into our view the Ideas of those individuals that are supposed to be marked by them." Sometimes they are used "to dispose me to embrace his opinion," as when I say that Aristotle held some view simply as a way of getting you to accept it. So, it had been a long standing view in the empiricist tradition that language could be used in ways that went contrary to their meanings and for more subtle purposes than to describe the world.

Perhaps even more surprising is that we can see this kind of idea being formulated as far back as <u>The Port Royal Logic</u> in 1662 (Arnauld and Nicole 1662/1861). In Part I chapter XIV the authors distinguish another kind of definition that involves what they call 'accessory ideas'. As they say,

> Now it often happens that a word, besides the principal idea, which we regard as the proper signification of that word, excites many other ideas,

which may be termed *accessory*, to which we pay but little attention, though the mind receives the impression of them.

For example, if one says to another, *You lied there*, and we regard only the principle signification of that expression, it is the same thing as if we had said to him, *You know the contrary of what you say.* But, besides this principle signification, these words convey an idea of contempt and outrage; and they inspire the belief, that he who uttered them would not hesitate to do us harm, which renders them offensive and injurious. (p 88)

So here we see the beginnings of the ideas which we will make much of in the coming pages. Though I should be quick to add that the authors of the Port Royal Logic were not themselves empiricists seeing as how there is a whole chapter on knowledge that cannot be derived from the senses.

In fact we can see an anticipation of the doctrine of emotive meaning in the next paragraphs that follow these. The authors go on to point out that often these accessory ideas are not annexed to the words by custom but are instead dependant on the manner in which the words are delivered; that is to say the expression of the accessory ideas depends on the tone of voice, or some gesture. But, they continue, "sometimes these *accessory ideas* are attached to the *words themselves* since they are excited commonly by all who pronounce them" (p 89).

But though there is this recognition that we do more with language than describe the world in 17th century writing on logic and language this is not marshaled to the aid of an empiricist theory of moral discourse. To see this idea blossom into the fruit of emotivism we need to go back to Moore's open question argument. Whether this argument is successful or not (I will discuss this more fully in section 2.3.1.) it was widely thought at the time, and for some time afterwards, to be fatal to the most plausible naturalist candidates for the meaning of ethical terms (at least utilitarianism and evolutionary accounts). It looked like one must then either be an intuitionist and think that goodness is a simple, non-natural property or, for those who wanted an empiricist ethical theory, sentimentalism seemed to be the only choice.

But when the sentimentalism of Hume was construed as a theory about the meaning of ethical terms it faced serious problems. Sidgwick had illustrated one of these problems nicely (Sidgwick 1874/1952, p 139-140). If the meaning of 'suicide bombing is morally wrong' is just that I morally disapprove of it, then when you say that suicide bombing is acceptable that just means that you morally *approve* of it. These two facts are consistent (that I approve and you

do not), but yet the sentences 'suicide bombing is morally wrong,' and 'suicide bombing is morally right' seem to contradict each other. Obviously the other side sees this and so to be charitable, Sidgwick supposes, we must take the sentimentalists to be talking about the justification for thinking that suicide bombing is wrong. The justification of our judgment is our feeling of moral disapproval. But if this is so it then becomes entirely mysterious what the sentences that we utter means.

Urmson sums up the situation nicely in the following passage

> …the emotive theory of ethics has its origins in epistemological despair. There is no account of the meaning of ethical utterances hitherto proposed which is epistemologically acceptable, since naturalism is unfortunately false and non-naturalism is abhorrent. (Urmson 1968 p. 19)

With this in mind let us turn to looking at this product of despair.

2.1.1. Classical Emotivism

The emotivists (Ayer 1936/1946; Stevenson 1937/1963; Stevenson 1944) responded by agreeing with Moore that 'good' cannot be defined in the sense that Moore had in mind, i.e. it cannot be equated with a set of natural properties which are its essence; 'good' had no meaning in that sense because it was not supposed to stand for a

property at all, it was supposed to merely *express* our feelings in much the same way that ejaculations and interjections do.

Ayer puts this most starkly and succinctly when he famously says,

> If I…say, "stealing money is wrong," I produce a sentence which has no factual meaning— that is, expresses no proposition which can be either true or false. It is as if I have written "Stealing money!!"— where the shape and thickness of the exclamation marks shows, by a suitable convention, that a special sort of moral disapproval is the feeling which is being expressed (LTL p 107)

I do not describe myself as having this feeling, rather I directly express the moral feeling in much the same way that 'ouch!' does not report that I am in pain but serves to express my pain. And just as 'Ouch!' can't be true or false and doesn't mean 'I am in pain', neither is "Saddam Hussein is evil" true or false nor does it mean that Saddam Hussein has the property of being evil. Moral sentiments can in no way be objective and are in the end simply brute. According to Ayer and Stevenson, though we may have some moral sentiments because of the factual beliefs that we have, in the end it is always possible that I approve of murder and you do not even though we agree on all the facts. Thus, according to Ayer, there is no such thing as real disagreement in ethics.

Ayer goes on to argue that expressing emotion is not the sole function of moral utterances. They are also used to exhort to action or to commend. He says,

> It is worth mentioning that ethical terms do not serve only to express feeling. They are calculated to arouse feeling, and so to stimulate action. Indeed, some of them are used in such a way as to give the sentences in which they occur the effect of commands" (LTL, p 108)

He is here recognizing that we do more than merely express our emotions when we engage in moral discourse; we also have what are called perlocutionary goals. We will talk more about these shortly but the basic idea is simple; we want to arouse feeling and stimulate action. This certainly is worth mentioning, I dare say it was good of him to do so!

He continues,

> We may define the meaning of various ethical words in terms both of the different feelings they are ordinarily taken to express, and also the different responses which they are calculated to provoke. (p108)

Taken together it looks as though what he has in mind is the idea that the meanings of ethical terms are given by the way that people use them and people use these words to express their moral emotions in virtue of existing conventions that pair the word to its use in expressing the emotion.

Stevenson's view is a bit different from Ayer's. The clearest statement of it is in his 1937 paper *The Emotive Meaning of Ethical Terms* (Stevenson 1937/1963), so let us take a closer look at the argument of that paper. He begins by pointing out that

> broadly speaking, there are two different *purposes* which lead us to use language. On the one hand we use words (as in science) to record, clarify, and communicate *beliefs*. On the other hand we use words to give vent to our feelings (interjections), or to create moods (poetry), or to incite people to actions or attitudes (oratory). The first use of words I shall call "descriptive," the second "dynamic." Note that the distinction depends solely upon the *purpose* of the *speaker*. (p 19)

For Stevenson dynamic language use is not restricted to moral discourse. It is a general feature of language use. Less broadly speaking he has in mind such specific purposes like encouraging someone, arousing sympathy, dropping a hint, satisfying a want, leading to believe, making a prediction, and communicating a belief. Some of these purposes result in descriptive uses of language while others result in dynamic uses.

Stevenson goes on to argue that the two different uses of language are not mutually exclusive. An utterance of 'hydrogen is the lightest known gas,' *may* be a descriptive use, which would mean that the purpose of the speaker was "to lead the hearer to believe this, or believe that the

speaker believes it," (ibid); but it may not. That is, depending on what the purpose of the speaker is, they could be *either* expressing a belief with the aim of communicating information about hydrogen, *or* they could be expressing their undying fascination with that most sublime of gases, hydrogen with the aim of making you appreciate it as well. Which one it is that we happen to be doing depends on our purpose, not on the semantic meaning of the sentence that we use.

The meaning of the sentence 'hydrogen is the lightest gas' remains the same whether the sentence is used dynamically or descriptively. In general Stevenson says that dynamic usage is independent from semantic meaning and that we cannot

> determine whether words are used dynamically or not merely by reading the dictionary—even assuming that everyone is faithful to dictionary meanings. Indeed, to know whether a person is using a word dynamically we must note his tone of voice, his gestures, the general circumstances under which he is speaking, and so on (ibid 19-20).

We must not only know what the person *said* but also what they intend to be expressing *in* saying what they did as well as what they intend to accomplish *by* saying what they did. So, any given sentence, on Stevenson's view, has a meaning as determined by the dictionary meaning of its parts and the grammar of English. But on any given

occasion of its use a person may mean exactly what it means, or what it means and more, or something else entirely. Which is the case will depend on what the person is trying to do.

This is quite different from Ayer's characterization of moral language. For Ayer, it seems, the meaning of the ethical terms is given simply by the conventions of the language that arise from the use that the language is put to. For Stevenson, however, we have that kind of meaning as well as the purpose of the speaker. We will see in the next section that this split corresponds to two different ways of taking the theory. We will come back to that. There is another difference that is worth mentioning. For Ayer ethical language is used to express an emotion, for Stevenson it is used to express what he calls an attitude. He defines an attitude as being for or against something and offers as examples purposes, aspirations, desires, and preferences (p 3, <u>Ethics and Language</u>). This is a much broader conception of what we are expressing with moral language than that offered by Ayer.

Stevenson then argues that moral discourse is a kind of dynamic usage. Moral utterances are 'instruments' used in trying to adjust the feelings and interests of others. In saying that X is wrong I do not mean to merely express my

attitude about X, I also want you to have the same attitude as well. When we say 'X is wrong' we are expressing something like 'I morally condemn X; do so as well' which Stevenson says is really to express the conjunction of attitudes 'I morally condemn X' and 'Do so as well,' at the same time with the same utterance. What he means to capture by 'do so as well' is what he calls a 'quasi-imperative'. We cannot command someone to have the same likes and dislike as we do rather what we do is to use 'suggestion'; we 'strongly recommend' that you have our attitude as well. More accurately then his analysis is that when we say "Saddam Hussein was evil" we express something like "I morally condemn Saddam Hussein; I strongly recommend that you do so as well".

This becomes his explanation of moral disagreement. Were you to counter that Saddam Hussein was not evil we would not have disagreement in our beliefs but rather disagreement in our interests. Our recommendations are different and we are each trying to change the interests of the other. In this way both our interests can't be satisfied and so we do have conflict in ethics. He suggests that most moral disagreements are rooted in disagreement in belief and so admit of factual justification. That is, a lot of moral disagreements hinge

on some factual claim and so if one can convince the other about the truth of the factual claim one will change their affective attitudes. So, for instance, if the morality of abortion ultimately depends on whether or not the fetus is a person then disagreement about this is disagreement in belief. What we have to determine is whether or not the fetus is, in fact, a person and convincing our opponent of this fact will resolve the disagreement.

This takes care of a lot of cases but real ethical disagreement for Stevenson is purely disagreement in interest (p 26). He asks us to consider the following scenario.

> A is of a sympathetic nature and B is not. They are arguing about whether a public dole would be good. Suppose that they discovered all the consequences of the dole. Is it not possible, even so, that A will say that it is good, and B that it is bad? The disagreement in interest may arise not from limited factual knowledge but simply from A's sympathy and B's coldness (p 29).

Ayer has a passage in the introduction to the Second edition of *Language, Truth, and Logic* where he echo's this sentiment. It would seem then that at bottom our attitudes are just brute and that ultimately disagreement in interests cannot be rationally settled. However all hope is not lost, according to Stevenson. This person's attitudes stem from their character and so in order to change their mind you will have to change their coldness

into sympathy; you must get your "opponent to see life through different eyes," (ibid) and that, according to Stevenson, is a persuasive task, not a rational one. This account of moral disagreement leaves most feeling unsatisfied.

The theory as stated so far met with fierce resistance. On one hand we have a problem for the semantic claim that arises from embedding ethical sentences into larger contexts. If what we have is really a claim about the *meaning* of ethical terms then it looks like the 'moral modus ponens' is no longer valid. Consider the following argument.

> P1- Suicide bombing civilians is immoral
> P2- If suicide bombing civilians is immoral, then recruiting young people to suicide bomb civilians is immoral
> So, recruiting young people to suicide bomb civilians is immoral

The emotivist claims that p1 really expresses a non-cognitive attitude about suicide bombing civilians, something like 'boo suicide bombing!!', but what are we to make of p2? It is not clear that p2 expresses any attitude. Thus, if the emotive theory of meaning is correct the argument turns out to actually look like this.

> P1- Boo suicide bombing civilians!!!
> P2- If suicide bombing civilians is immoral, then recruiting young people to suicide bomb civilians is immoral

> So, recruiting young people to suicide
>   bomb civilians is immoral

So it appears that the argument equivocates and so is invalid. This is the so-called Frege-Geach (Geach 1965) problem. Aside from the equivocation it seems as though the very idea of validity in moral arguments is lost. If 'murder is evil' can be neither true nor false then no argument that it occurs in can ever be valid in the formal sense. We will discuss this problem, and the solutions, that have been offered in more detail later.

On the other hand we have an objection to the claim about justification. In answer to the challenge from Sidgwick mentioned in the introduction to this chapter the emotivists claimed, as we saw, that there could be no real contradiction in ethics. But this is extremely implausible. It certainly seems like people are contradicting me when they approve of suicide bombing of civilians and I do not. In fact it may seem, as it did to Sidwick, that without the possibility of real disagreement we are not even talking about moral judgments at all. This is because, as Sidwick says,

> The peculiar emotion of moral approbation is, in my experience, inseparably bound up with the conviction, implicit or explicit, that the conduct approved is 'really' right—i.e. that it cannot, without error, be disapproved by any other mind. (Sidgwick 1874/1952, p 140)

This is contrasted with our utterances about our tastes in food. In these kinds of judgments, judgments about the taste of broccoli and the like, we will affirm, at least after reflection, that someone could disagree and not be in error. This reflects the realism inherent in our ordinary moral ways of thinking. Since we both can't be right, and since I think that I am right, that means that you must be in error. This is markedly different from our acceptance of irrealism in the realm of taste. We will come back to the issue of realism and irrealism in the final section of this chapter.

For now it is enough to note that these two challenges must be met by any kind of theory that purports to carry on the expressive thesis. We need to be able to square our theory of the meaning of ethical terms with the obvious fact that moral terms act just like descriptive terms when they are embedded and when they are used in the context of making an argument and that they are truth-evaluable. A theory that was unable to account for this obvious semantic fact would not likely turn out to be true. We can accept some 'collateral' damage to our pre-theoretic notions if they are outweighed by the acceptance of a theory that has advantages over another theory. This sort

of thing is inevitable. However, if your theory of physical objects denies that there are physical objects and can't explain why it seems to us as though there are, you know something has gone wrong. The claim that, semantically, moral predicates and non-moral predicates (like 'is red', 'is a cat', etc.) act exactly alike in embedding contexts and can be combined with noun phrases to generate truth-evaluable sentences which can be used to generate valid inferences is non-negotiable.

At the same time we also need to be able to account for the difference between the ways we talk about taste in food and the ways we talk about morality. It seems to me, as it did to Sidgwick, that the very idea of morality is tied up in the notion that two people who disagree can't both be correct, or that one must be in error. This is an important challenge. If a theory is not able to capture the difference between our talk about our tastes in food and our moral judgments then we should reject that theory.

We will come back to these issues in more detail in later chapters. For now let it suffice to say that these two objections were and still are considered devastating to classical emotivism. It seemed clear to everyone that some modifications needed to made to the theory. This

results in the next wave of British and American expressivism in the early 1950's.

2.1.2. **After the War**

The post-war ethical writers, Urmson (Urmson 1950), Hare (Hare 1952), Nowell-Smith (Nowell-Smith 1954), and Edwards (Edwards 1955, written in 1950) all deny that they are emotivists, but they are clearly carrying on the work began by Ayer and Stevenson. They accept the expressive thesis, though they disagree amongst themselves as to what it is that gets expressed. They all see moral discourse as an activity where we use language to do something.

So Urmson sees 'good' and 'bad' as grading terms. We use these terms to perform the action of putting people and actions into categories. Hare sees moral language as a subspecies of prescriptive language so when we say that stealing is wrong we express something akin to a universalizable prescription on stealing: 'don't steal'. Edwards sees moral language as something that is like a command or a request but that is distinctly moral and not really a command at all. The most sophisticated of these writers is Nowell-Smith. He (P. 72-73) argues that there are different kinds of adjectives. There are aptness words which suggest that a certain emotion is appropriate (like

'horrible' and 'sublime') and there are gerundive words which classify objects as praise-worthy or contemptible ('like 'liberal' and 'weed'). He then argues that one and the same word will often play all of these roles depending on the purposes of the speaker. He calls this the 'Janus Principle'.

These philosophers saw that once one took the claim about *meaning* really to be the expressive thesis, that is a thesis about language *use* rather than meaning, there was a sense in which moral judgments could be justified or objective.[7] This was in response to criticisms of Ayer and Stevenson by intuitionists like Ross and even fellow emotivists like Toulmin (Toulmin 1950/1960). Ayer seemed to hold that all moral disputes were really factual disputes. In the Introduction to the Second Edition of LTL he emphasizes his agreement with Stevenson by saying that even though most of our moral disagreements are really factual disagreements, it is always possible that two persons could agree on all the facts and yet hold differing moral sentiments. If this happened there would be no way to say that either of the two persons had the wrong moral sentiment.

---

[7] I am ignoring, for ease of exposition, that these ideas were the result of the influence of Grice and Austin. I will get to this in the next section.

He also emphasized the claim that a lot of the time when we say that something is wrong what we mean is that that thing falls under some moral principle that we accept. So if I have completely 'internalized' the utilitarian theory, when I say that free speech is good I may simply mean that free speech maximizes happiness, and then of course my statement is objective and has a determinate truth-value and, it would seem, the word 'good' would then have some ordinary descriptive content. But none the less he insisted that there was a use of the word that admits of no truth-value and so in a truly ethical use of the word it is still meaningless.

 Ross (Ross 1939) argued that Ayer's claim that most (if not all) moral disputes are really disputes about some matter of fact undermined his emotivism. The problem according to Ross is that when we are arguing over these factual beliefs,

> What [we] are attempting to do…is to convince each other that the like, or dislike, is justified, in other words that the act has a character that deserves to be liked or disliked, is good or bad (p 404)

Ross is here objecting to the emotivist's theory of justification in much the same spirit as Sidgwick's objection to sentimentalism. He is arguing that moral argument must be made sense of in a proper moral theory

and once you do that you find that emotivism is false. When we argue it looks like what we are doing is using the facts to justify our moral judgment.

But with the expressive thesis more explicitly formulated an answer to this objection emerged. There was some descriptive content to ethical judgments and so we could give reasons for our attitudes but that did not stand in the way of our using those terms to express our moral sentiments. So, Urmson for instance says,

> to say that something is good is to commend, approve of, recommend, grade highly (and so on according to context) that thing in the light of certain recognized criteria either for the kind to which it belongs or from the point of view which we have adopted (Urmson 1968, p 130).

So our moral judgments will be in that sense objective. But at the same time what we do when we say x is wrong is very different from what we do when we say that x is red or square.

Also, Paul Edwards takes a similar line. He argues that just as when we say that food at a certain restaurant is nice there is a certain range of properties that are the 'nice-making' properties so too when we say of something that it is good or bad. The moral words have descriptive content though they are used to express something that is akin to a command or a request, but is neither (see for

instance p139 of <u>The Logic of Moral Discourse</u>). Hare too follows suit (see p 94 of <u>The Language of Morals</u>). But even so this objectivity is not realism because two people may agree on all the facts and yet still have differing attitudes. Even Hare who, as we will see, makes great strides in showing how reason can play a role in non-cognitive theories, is forced to admit that we could have a perfectly consistent Nazi (the fanatic who is willing to accept that his maxim entails that he would be treated a certain way were he different). It is still the case that, ultimately at least, two people who disagree can both be right. Each may simply be using a set of standards that are different from the others' and there is no way to adjudicate between rival standards.

Even so this was a major advance for expressive theories. For what the post-war expressivists showed is that we can give an account of how our emotive reactions might be justified. This kind of view culminates in the work of the lesser known emotivist William T. Blackstone in his paper 'Objective Emotivism' (Blackstone 1958). Blackstone argued that by using Nowell-Smith's idea of contextual implication, and the distinction from Hutcheson between exciting reasons and justifying reasons, we get an emotive theory where to say that X is right is to express

our moral sentiment and also contextually imply that we have reasons which justify our attitudes.

Nowell-Smith formulated his notion of contextual implication as follows,

> A statement p contextually implies a statement q if anyone who knew the normal conventions of the language would be entitled to infer q from p in the context in which they occur. Logical implications are a sub-class of contextual implications, since if p logically implies q, we are entitled to infer q from p in any context whatever. (Ethics p 81)

So the view that Blackstone argued for holds that when we say 'X is wrong' we express our non-cognitive moral emotion and contextually imply that we have some reason which justifies our attitude. This reason may be a utilitarian one (as Blackstone argues Hutcheson himself held), a deontological one, or whatever. But this justifying reason, according to Blackstone, is not what motivates us to action. The motivating reason is the non-cognitive moral emotion. What this amounts to is the claim that any normal speaker of English would be entitled to infer that I did in fact have some justifying reasons for my moral sentiment.

This neatly capitalizes on the very weakness that Sidgwick, Toulmin, and Ross exploited in attacking sentimentalism and classical emotivism, respectively. Since moral judgments are 'inseparably bound' to the

feeling that the sentiment is 'really' right it is reasonable to infer the existence of a convention by which our interlocutors could deduce our belief that our moral sentiments are justified. This view of Blackstone's differs from the views offered in the early '50's by being fully realist. He allows for full-blown justification of moral judgments even though he also sides with the emotivists in holding that we express our moral emotions in moral discourse and that reason does not motivate us to action. This view offered by Blackstone is the first historical example of what I have been calling an emotive realism. This is why I say that emotive realism is a *mostly* neglected class of theories. But, as far as I can tell, Blackstone's views had little impact on the course of subsequent debate in metaethics despite his foreshadowing the way the debate would go.

A corollary of the criticism we have been talking about is the charge that emotivism renders ethics essentially irrational. If all there is to moral judgment is the way that I, or you, feel about something then moral reasoning appears mysterious. This is an analog of the Frege-Geach problem discussed in the last section. Hare takes up this challenge. Hare, recall, argued that moral judgments were a kind of prescription and that 'x is wrong' was used to

express a universal prescription on doing x of the form 'don't x'.

Hare went on to distinguish the *phrastic* from *neustic* in the following way. Consider the two sentences below.

(a) Don't eat any animals
(b) You don't eat any animals

They both share a common phrastic element which we express as (c),

(c) Your not eating any animals

They differ in terms of their neustic component. In order to get from (c) to (a) or (b) we will need to complete them is some way. Hare suggests the following as capturing the differing neustic elements:

(c1) Your not eating any animals, please
(c2) Your not eating any animals, yes

He then points out that 'any' and the other logical words go into the phrastic element of the sentence and so are a common element between imperatives and commands.[8]

Since the logical connectives appear in the phrastic elements they must have the same meaning in both. Thus, if 'any' is to have any meaning when it is completed with an indicative neustic as when it is completed with an imperative neustic. This means that logical relations that

---

[8] This is also clearly, I think, an anticipation of the content/force distinction we will see in the next section.

'any' brings with it must also be brought to imperatives. To illustrate take the following case. Suppose that you ask for something to drink and I say "Help yourself, but don't drink any of my Brooklyn Lagers". Now suppose you go to the refrigerator and look inside and find several Brooklyn Lagers, clearly labeled, as well as water, juice, and a couple of other beers. What should you do in this situation?

   Hare's claim is that in these kinds of situation we see a kind of inference being made. Something like,

    1. Don't take any of the Brooklyn Lagers
    2. Those are Brooklyn Lagers
    So, don't take those

Hare claims that even though 1 can't be true there is a sense in which we either affirm it or deny it. What he means by this is that we 'assent' to the command or we don't. In his terms, 'either take a Brooklyn Lager or don't take one' is analytic.

   Given that imperatives obey the law of the excluded middle in this sense we can not assent to 1, believe 2 to be true and yet fail to assent to 3 without some kind of inconsistency. Hare took this to show that there is a logic of imperatives. The basic inference pattern for this imperative logic was a modified version of *modus ponens* where the major premise was a universal prescription, the

minor premise was a particular statement of fact and the conclusion a particular prescription just as in the case above. Though Hare doesn't do this one could define prescriptive analogs of the truth-functional definitions of the logical operators and develop a complete and consistent imperative logic.[9] Some people are resistant to the idea of a logic of imperatives, but as Hare himself points out, the reasons for this are mostly historical (LoM p. 26 fn).

We can also see the beginnings of an answer to the Frege-Geach problem here, though Hare never, to my knowledge, explicitly addresses the argument. We can extrapolate what he might have said from his discussion of hypothetical imperatives in The Language of Morals (pps. 34-35). He begins by discussing a standard inference of the following kind; "if all mules are barren, then this animal is barren." This sentence is

> entailed by the sentence 'this animal is a mule' we only have to know the meanings of 'all' and the other words used in order to make the inference.

---

[9] For instance, we could introduce the notion of a command's being *fulfilled* as an analog of truth. An imperative is fulfilled if it is obeyed (or 'assented to' as Hare puts it). Given this 'sit down and shut up!' would be fulfilled if both of its conjuncts were fulfilled, otherwise not. 'Sit down or shut up' would be fulfilled if one (or both) of its disjuncts were fulfilled. 'Don't sit down' would be fulfilled if 'sit down' wasn't fulfilled, etc. The rest is just working out the details.

Of course this is needlessly complex. The simpler way of putting this is as follows.

    1. All mules are barren
    2. This animal is a mule
    So, this animal is barren

We get to the more complex inference by "taking away the major premise from its proper place, and adding it to the conclusion inside an hypothetical clause".

He then asks us to compare the following analogous simple kind of imperative inference.

    1. Go to the largest grocer in Oxford
    2. Grimbly Hughes is the largest grocer in Oxford
    So, go to Grimbly Hughes

We can take this inference and, by applying the method of transformation pointed out above, get this,

    1. Grimbly Hughes is the largest grocer in Oxford
    Therefore, if go to the largest grocer in Oxford, go
    to Grimbly Hughes

Now clearly 'if go to the largest grocer in Oxford, go to Grimbly Hughes' is not a well-formed English expression.

Hare is aware of this and goes on to say,

    In English, we write the conclusion in this form:
        If you want to go to the largest grocer in
        Oxford, go to Grimbly Hughes
    We have only to know the meanings of 'want' and the
    other words used in the conclusion (including the
    imperative verb form) in order to make the inference.

Now, Hare is here more concerned to be showing that 'want' does not always denote a mental state, as in these cases

where it is a hidden imperative in a subordinate clause like the conclusion of the Grimbly Hughes argument, but it is easy to see how this might be extended to answer the Frege-Geach problem.

This, you will recall was the problem that arose when the emotivist tries to account for arguments like the following.

    p1- Suicide bombing civilians is immoral
    P2- If suicide bombing civilians is immoral, then
        recruiting young people to suicide bomb civilians
        is immoral
    So, recruiting young people to suicide bomb civilians
        is immoral

This is because in 1 the phrase serves to express an emotion but in 2 it does not, as below.

    1- Boo suicide bombing civilians!!!
    2- If suicide bombing civilians is immoral, then
        recruiting young people to suicide bomb civilians
        is wrong
    So, recruiting young people to suicide bomb civilians
        is immoral

But applying Hare's analysis of hypothetical imperatives we would get something like,

    1. Don't suicide bomb civilians
    2. If don't suicide bomb civilians then don't recruit
        young people to suicide bomb civilians
    So, don't recruit young people to suicide bomb
        civilians

Now 2. Looks like it is no good but Hare might respond that in English we write 2. As 2-.

However, this approach seems to have problems. Look at 2- more closely and we can see that it is itself a complex inference of the kind discussed by Hare. It is valid because of the more basic inference.

    1. Don't recruit people to do what is wrong
    2. Suicide bombing civilians is wrong
    So, don't recruit people to suicide bomb civilians

But here 2 looks like an indicative sentence. Now I suppose that Hare might respond that the argument really looks like this.

    1. Don't recruit people to do what you are commanded
       not to do
    2. Don't suicide bomb civilians
    So, don't recruit people to suicide bomb civilians

All of Hare's examples involve arguments with a universal prescription a singular statement and a particular prescription. He doesn't give us any indication of how he would deal with these kinds of arguments. But there do seem to be some promising solutions waiting to be developed.

We will talk more about this problem when we get to chapter three and we see how contemporary thinkers have extended Hare's insight. In particular we will look at Blackburn's attempt to translate this kind of approach into a general logic of attitudes (as opposed to logic of imperatives). But first I want to turn to developing some

of the distinctions in the philosophy of language that we will need. This will allow us to come back to the issues discussed here with tools capable of finer precision than the ones we have been working with.

2.2. Meaning, Speech Acts, and Communication

In the late 1950's the expressive thesis was taken only as a thesis about the meaning of words. Both Grice and Austin start with it in that capacity and then go on to develop the idea in a more general and rigorous way.[10] In this section I want to trace the development of a particular kind of philosophy of language that combines Austin's theory of speech acts with Grice's theory of meaning into a general theory of human communication.

2.2.1. Grice on Meaning

Grice, in his well known paper on meaning (Grice 1957), targets Stevenson's theory of meaning. On Stevenson's view the meaning of a word was its disposition to cause attitudes in hearers and to be caused by attitudes in speakers. A word could have either emotive meaning, if it caused or was caused by some emotive state

---

[10] I know that Austin was developing speech-act theory since at least 1939, and as he says, it is partly inspired by the work that was going on in metaethics, emotivism in particular, but my point here is that by the '50's they were solely interested in the expressive thesis as a linguistic claim and there was no talk about the justification thesis at all.

of the speaker/hearer, or it could have cognitive meaning if it were caused by a cognitive state like a belief. Grice argues that there are two related problems with this view about meaning.

The first is that an umbrella, or indeed any kind of object, can come to have the kind of disposition that Stevenson calls the meaning of a word, but surely we do not want to say that the umbrella then comes to have meaning. So, for instance, it is easy to imagine that someone's having seen me take an umbrella with me every time it is raining might cause them to have the belief that I think that it is raining when they see me take it this time, this being due to a complex kind of conditioning, but surely the umbrella doesn't mean anything. The second related criticism is that many sentences will produce certain attitudes in hearers that it would be extremely odd to think of as the sentence's meaning. Grice's example is 'Jones is an athlete'. This sentence is very likely to cause someone to have the belief that Jones is tall, but it is positively absurd to say that 'Jones is an athlete' means 'Jones is tall'.

In response to this Grice gives his intentional theory of meaning. It is not that the term has a dispositional property but rather that people have

generally intended to bring about certain psychological effects by uttering the term. On this kind of view the word 'cat' means what it does in virtue of the fact that people have intended to be talking about a certain kind of animal when they use it. With this comes the distinction between speaker meaning and sentence meaning. I can on a given occasion mean what the sentence means (what people generally intend to convey) or rather some further thing (what I specifically intend). Grice's strategy is to explain linguistic meaning in terms of speaker meaning. This in fact turns out to be another criticism of Stevenson's view. Even if his causal theory could overcome the objections already mentioned it does not seem to be able to account for the difference between speaker meaning and sentence meaning.

When using the word on a particular occasion a person may mean by it what the word usually means or something else all together. This leads him to his famous theory of implicature (Grice 1975). There are two basic kinds of implicature. One is a *conversational* act carried off via an intentional violation of some conversational maxims which the hearer is intended to recognize. I shall not here rehearse the various maxims and the way that Grice thinks that they can be exploited or 'flouted' to pull off

a conversational implicature. It will suffice for our purposes to note that conversational implicatures require an inference on the part of the hearer to the effect that the implicated meaning is required in order to make sense of what the person said, or the way in which they said it. For example, if I say 'metaethics is so stimulating' with a certain intonation and a rolling of the eyes it would be reasonable for you to infer that I actually mean the opposite of what I say. These kinds of implicatures Grice call *particularized* conversational implicatures because they depend on the particular context of utterance.

Grice contrasts these with what he calls *generalized* conversational implicatures. Grice's example of these kind of implicatures is the utterance "X is meeting a woman this evening". Someone who said this would, "normally indicate that the person to be met was someone other than X's wife, mother, sister, or perhaps even close Platonic friend," (LC p 180). This isn't part of the meaning of 'a woman', or so Grice urges invoking what has become known as 'Grice's razor'. Why should we postulate a variety of meanings for 'an X' when we have a simpler more elegant solution available? The simpler theory is that a person who says this would normally, or usually, or barring

special circumstance, be implicating that X was not a wife, etc.

There are also *conventional* implicatures, which are not carried off by an intentional violation or flouting of the rules but are actually a product of the conventional meaning of the words. The standard example of one of these is the utterance of 'she is poor but honest'. This is typically taken to implicate that there is some kind of contrast between being poor and being honest. But, it is claimed, this is not part of the linguistic meaning of 'but' since it has just the same meaning as 'and' (the truth-functional meaning). It is an implicature that has become conventionalized. These kinds of implicatures are hard to distinguish from generalized conversational implicatures, but there are a couple of distinguishing features that all conversational implicatures share, and which conventional ones lack.

The first of these is that it is always possible to "explicitly cancel" a conversational implicature. So, suppose I say 'john is meeting a woman and that woman is his wife". No doubt this is an awkward locution, but if I were to say it I would have canceled the implicature that John was meeting a woman that wasn't his wife. Generalized conversational implicatures display what Grice calls 'non-

detachability' which he characterizes as the inability of finding "another way of saying the same thing, which simply lacks the implicature in question" (p 181). This is supposed to distinguish the generalized conversational implicatures from the conventional implicatures. Since the conventional implicatures are carried off due to a convention it will always be easy to find an alternate way of saying the same thing that 'simply' lacks the implicature.

So, to take our earlier example --'she is poor but honest'-- I can easily say 'she is poor *and* honest' which is equivalent in meaning but lacks the implicature that there is a contrast between the two. This is, of course, assuming that the substituted phrase does not carry some conventional or generalized implicature of its own. We will come back to these distinctions in more detail in the third chapter when we discuss the view of David Copp.

2.2.2. Austin on Speech Acts

At around the same time Austin was developing his speech act theory. Its canonical form is presented in How to Do Things with Words (Austin 1962). His main goal was to show that there are some utterances that look like normal subject-predicate sentences but are actually in a

completely different business than describing the world. In the opening lecture he makes some remarks about a revolution that is going on in philosophy. He even jokes that we might call it "the greatest and most salutary in [philosophical] history," (p. 3). He has in mind the 'piecemeal' exposure of what he would call 'the constative fallacy' which is the fallacy of assuming that all utterances are capable of being either true or false.

Part of this exposure came from logical positivism and verificationism, which claimed to have shown that some sentences were nonsensical, but it was also partly started by thinking in ethics. He says,

> It has come to be commonly held that many utterances which look like statements are either not intended at all, or only intended in part, to record or impart straightforward information about the facts: for example, 'ethical propositions' are perhaps intended, solely or partly, to evince emotion or to prescribe conduct or to influence it in special ways. (ibid)

It is clear from this that Austin is referring to the views of Ayer, Stevenson, and Hare. Ethical sentences are not nonsense, they masquerade as statements. So even though they look like they can be true or false they cannot. He then tells us that "the type of utterance we are to consider is not, of course, in general a type of nonsense…rather, it is one of our second class— the masqueraders," (p. 4). The kinds of sentences he is

interested in are part of the class in which ethical sentences are found. They are masqueraders which look like they report or describe but that are not intended to do so. He thus sees himself as participating in the bigger movement against the constative fallacy begun by emotivism and that his contribution lies in extending the early emotivists' point to sentences with "humdrum" verbs, thereby giving a general account of the various things we can do with language.

After the introduction Austin presents the distinction between constative and performative utterances. Constative utterances report, describe, or constate some state of affairs. They are the truth-apt utterances and include much of our garden variety utterances, like 'New York is on the Hudson,' and 'Brooklyn College is in Flatbush,' etc. Performative utterances are not capable of being true or false. They are used to perform some action and so are literally speech acts. Promising, christening, ordering, betting, etc. are prime examples of this kind of utterance.

For instance saying (sincerely) "I promise to pay you back" is to make a promise, saying "sit down!" is to issue an order, etc. Austin sees performative utterances as little rituals that are governed by conventions and so

they have 'felicity conditions' but no truth conditions. So, if I have no authority to order you, or I don't have the power to christen this boat, I have not really performed those actions. Of course towards the end of How to do Things with Words Austin rejects the distinction between constative and performative as untenable. The problem being that stating is a kind of speech act itself. An utterance of "I state that it is three o'clock", or "I report that John is ready" looks just like "I promise…" and "I christen…" and hence there is no reason to think that factual utterances are not performative as well.

    In its place Austin introduces the three-way distinction between locution, illocution and perlocution. Each of these is a distinct kind of act that one may perform with one and the same utterance. The locutionary act is just the act of uttering certain sounds, as words, with a certain definite sense and reference; that is to say the locutionary act is just saying a certain sentence that has a certain meaning. In the strict sense, then, what is said is simply the locutionary act. The illocutionary act we perform *in* uttering those words is what we know as the speech act proper. It is the promising, stating, asserting, pronouncing, etc. The perlocutionary act is that which we (hope to) accomplish

by saying what we said. Thus speech acts are distinguished by the illocutionary act performed and the perlocutionary goal that we have. As Austin says, we "distinguish the locutionary act 'he said that…' from the illocutionary act 'he argued that…' and the perlocutionary act 'he convinced me that…'" (p 102); so I might have *said* "hybrid cars are better for the environment," (locution) and in saying that have been *arguing* that hybrid cars are better than standard cars (illocution) in order to *convince* you that you should get one (perlocution).

By the end of the lectures he admits that we can perform speech acts without having the type of speech act explicitly in the sentence we utter. I can apologize for stepping on your cell phone without saying "I apologize". I can say 'I didn't mean to do it,' or even 'I didn't see it there!' if said correctly, can actually be apologizing. Or again consider my yelling 'watch out!' (a locutionary act). In saying that I am (most likely) performing the illocutionary act of warning you and I may also accomplish the perlocutionary act of getting you to stop what you were doing and avoid the imminent danger yet I did not say "I warn you…." It follows as a consequence of this that you cannot determine what speech act someone is performing just by looking at the locutionary act that they perform

or in other words that the meaning of the sentence used does not (completely) predict what someone may use it for.

How is it that we can do this? For Austin speech acts were all conventional. To perform a speech act was to conform to a certain convention, and whether or not that act was successful depended on the existence of that convention. And there certainly is a class of speech acts that act in just this way. These involve institutional speech acts like a judge saying 'case dismissed' or an umpire saying 'you're out!' So one way to explain how we are able to do this is in terms of usage giving rise to conventions. So, the story might go, people in the past used 'watch out' to warn people, it was successful, and so people started regularly using it that way. Over time a convention arose in English whereby just the utterance of 'watch out' counts as warning someone. The problem with this view is that it fails to explain how 'watch out' could ever have been used successfully to warn in the first place.

2.2.3. Strawson and Neo-Gricean Speech Act Theory

There is another way that we could perform speech acts in the absence of a linguistic convention. Strawson (Strawson 1964/1994) argued that, while some speech acts

were conventional there was another class of speech acts that were performed via a Gricean intention. So to express a belief (that is to perform the illocutionary act of asserting) is to utter a sentence with the intention that my interlocutor take my utterance as a reason to think that I have some belief. This intention is successful partially due to the fact that I intend that you recognize my intention. So I intend you to recognize my intention that you take my utterance as reason to think that I have a certain belief. This is a reflexive intention, not an iteration of separate intentions.

This allows us to merge Grice's distinction between speaker meaning and linguistic meaning with Austin's theory of speech acts. One way that this can be done, and the way that I am most influenced by, is that of Kent Bach and Robert Harnish (Bach and Harnish 1979); I will talk more about their view and its application later, but the basic view is that what speech act I perform is not dependent on the meaning of the sentence that I utter or the conventions that govern it.

Rather it depends on what they call 'R-intentions' which are intentions whose "fulfillment consists in their recognition," (p 15). They then use this notion of R-

intention to define what it means to express an attitude as follows;

> For S to express an attitude is for S to R-intend the hearer to take S's utterance as a reason to think that S has that attitude (ibid)

They go on to give us a simple example.

> In the case of statements the speaker expresses two attitudes: belief in a certain proposition and the intention that the hearer believe it as well. That is to say, for S's utterance of *e* to be a statement that P, S must R-intend H to take the utterance as reason to think (a) that S believes P and)b) that S intends H to believe P. Correlatively, for H to understand that S is stating that P in uttering *e*, H must take S's utterance of *e* as R-intended to be reason to think (a) and (b). (p 16)

So, depending on my intention, any given sentence may be used to perform a number of illocutionary acts. So, I can say 'I promise to x' not as a promise, but as a warning, or a threat. Which it is that I am actually doing depends on what it is that I R-intend to be doing and not on the meaning of the sentence that I use. This fits very well with Stevenson's claim that the purpose of the speaker is what matters in determining whether the speaker is using a given sentence descriptively or dynamically.

It is important to distinguish the illocutionary act from its perlocutionary success. So I am successful in stating that P if my hearer believes (a) and (b) above, but whether or not they come to believe P as well is a

matter of my perlocutionary success. I can perfectly well make a statement without you coming to believe the stated proposition.

With these distinctions in the philosophy language in hand and the issues and views of the early emotivist tradition we are now in a position to put together the pieces and say something about what moral realism is supposed to be.

2.3. Meaning, Justification, & Moral Realism

It is often argued that a philosopher's theory of the justification of moral judgments determines his theory of the semantics of moral words and that this semantic theory is the only way to tell the difference between someone who is a 'real-realist' and someone who is 'quasi'-realist like Simon Blackburn. For instance, we see David Copp (Copp 2001), in an article we will come back to in the third chapter, saying that the distinctive doctrine of moral realism is that the moral realist thinks that the moral predicates refer to *robust moral properties*. So to say that suicide bombing is morally wrong, according to this way of construing realism, is to assert (i.e. express a belief) that suicide bombing has the robust moral property of being wrong. I have even heard philosophers

say such things as 'All realists agree that normative judgments express beliefs with truth-conditional contents, and that normative predicates ascribe properties' (Hussain and Shah 2006).

To say that moral properties are 'robust', according to Copp, is to say that "[they] have the same basic metaphysical status as ordinary non-moral properties," (Copp 2001, p. 4). It is of course a matter of some controversy just what the status of ordinary non-moral properties is; but let us waive that for the moment. In moral contexts there are, broadly speaking, two candidates and the majority view is that the ethical realist is committed to one of them. On the one hand we might think that there are non-natural properties and that 'good' and other moral words pick some of those properties out. This is, of course, the Platonic route. On the other hand, we might think that non-natural properties don't exist and insist that moral properties must be natural properties and 'good' and other moral words pick some of those out. This is usually the view of the utilitarian, though it is available to any number of theories. The point is that it is almost universally accepted that moral realism entails a commitment to the metaphysical claim that there are moral properties and the semantic claim that moral

predicates refer to these moral properties. I plan to argue that both of these claims about realism are false.

We will get to that in due time but for now let's get back to the traditional story. As against moral realism we usually pit anti-realism or 'irrealism' as I'll call it. The irrealist, on this traditional story, is one who denies that the moral predicates refer to robust moral properties of either variety. The origins of irrealism are in sentimentalism, which, as we have seen, blooms into the fruit of emotivism in the early twentieth century. The most popular form of irrealism is now called 'expressivism,' which is an umbrella term that is supposed to cover everything from the emotivism of Ayer and Stevenson to the prescriptivism of Hare to the contemporary views of Blackburn and Gibbard.

The title 'expressivism' seems to come into the literature via the work of Gibbard. It makes its first appearance (in a specifically moral context) in his 1986 paper "*An Expressivistic Theory of Normative Discourse*" (Gibbard 1986). The standard title by that time had become 'non-cognitivism', and Gibbard mostly uses this term in his earlier work, and indeed in this paper as well. In that paper he characterizes a view like his as an 'expressivistic' one and in his 1990 book Wise Choices,

Apt Feelings (Gibbard 1990) he calls his view 'norm-expressivism'. By the mid '90's the view defended by Blackburn and Gibbard came to be known simply as expressivism.

Blackburn had used the term 'expressive theories' in his 1984 book Spreading the Word but doesn't seem to make an 'ism' out of it. He preferred to talk of 'projectivism', in order to bypass the connection to emotivism and draw instead on Hume, or his own coinage 'quasi-realism'. In his 1998 book, Ruling Passions when he is discussing what to call a view like his he says,

> What should a theory of this kind be called? I have called it 'projectivism', but that can sound misleading. It can make it sound as if projecting attitudes involves making some kind of mistake, like projecting our emotions onto the weather, or projecting our wishes onto the world by believing things we want to believe. This is emphatically not what is intended. Gibbard calls the view 'expressivism', and I now think that is better. (p 77)

From there the term became a general term for any of theories that I have just mentioned. And of course the reason for this is obvious.

All of these views have something in common. They all agree that when we say that something is morally wrong or right we are doing something that is fundamentally different from describing the world; we are *expressing*

some "emotive or conative motivational or non-cognitive state," (Copp, ibid). It is, in fact, a huge mistake according to the irrealists to think that moral predicates act like non-moral predicates and refer to, or denote, or whatever, some kind of property. Moral predicates are in a different kind of business all-together and only look as though they stand for properties. The differentiating claim between realist and irrealist on this traditional picture is one about the *meaning* of ethical terms. Realism entails one kind of semantics, irrealism another. The realist says that the meaning of moral predicates is given by the robust moral property that they stand for or name or denote or whatever, the irrealist denies this. So much, then, for the standard metaethical story.

But this is wrong; the debate between the realist and the irrealist is not a debate about semantics. It is really a debate about metaphysics (Devitt 2002). In particular these views about their semantics are really supposed to be diagnostic of their views about the justification of moral judgments. The robust moral properties that moral predicates refer to are supposed to be the truthmakers for moral judgments in exactly the same way that non-moral properties are supposed to be the truthmakers for non-moral judgments. The irrealist is then

someone who denies that there are such properties and instead claims that moral judgments are justified by the emotional, conative, or motivational states of people. But as Michael Devitt also pointed out (ibid) to construe the debate this way is to make it impossible for a nominalist to be a moral realist. If one is disposed to think that there are no properties at all then it will trivially be the case that there are no moral properties, robust or otherwise.

Because the metaethical theorists of the early twentieth century took the relation between meaning and justification for granted their strategy was to attack the theory of meaning as a means of attacking the theory of justification. For instance, it was utilitarianism taken as a theory of the meaning of ethical terms that Moore's famed open question argument was aimed at. It was also the target of Ayer's similar attack in Language Truth and Logic. Both took themselves to be showing that 'good' does not mean 'maximizes happiness' and so therefore maximizing happiness was not the correct account of moral justification. They took this conclusion one step further and proclaimed that any attempt to define moral concepts in terms of non-moral ones was doomed to fail.

However, one thing that has emerged in the various responses to Moore is that we need to separate the question of the meaning of ethical terms from the question about the justification of ethical judgments. So, even if it is the case that 'good' does not mean 'maximizes happiness' it may still be the case that maximizing happiness is what the good is metaphysically, and so is the correct theory of the justification of ethical judgments. This had been previously obscured primarily due to confusion in what people take a definition to be. Quine diagnosed this problem in another context (Quine 1966/1976). As he puts it, "words have definitions, but so also do things" (p 51). In the Aristotelian sense a definition is the essence of the thing defined (Robinson 1950). So to define 'good' in this sense would be to give the essential characteristic(s) of goodness.

But as we have seen there are other senses of meaning. There is a modern sense of meaning in which to give the meaning of a word is not to give the essence of the thing which the word stands for but to specify the conventions that govern the use of that word in the language. This is a purely linguistic conception of meaning, which is to say that we are talking about a 'defining phrase' rather than a 'defining property'. It is

an open question of our own at this point whether the defining phrase will have to make reference to a defining property. It seems that there is the possibility of finding a defining phrase for 'good', 'right', 'evil', etc., that would not make reference to some defining property of goodness (the thing), rightness (the thing) or evil (again, the thing).

So, one's theory of the meaning of words is independent of one's theory of the justification of moral judgments. For any moral theory, like utilitarianism say, we need to know if it is supposed to be an account of the meaning of ethical terms or whether it is to be an account of the justification of moral judgments or both. Moore's open question argument is thus merely an argument against defining the *word* 'good' in terms of maximizing happiness or anything else that isn't itself a moral term, but as against maximizing happiness as an account of the justification of moral judgments it is impotent. This in turn means that the main lesson of the open question argument is not that we cannot define goodness (in the Aristotelian sense) as maximizing happiness, but rather that we cannot define moral words in terms of that which is not moral, which certainly makes more sense (Baumrin 1968). We will come back to this very important point.

What this means is that no metaethical view entails any normative theory and no normative theory entails any metaethical view. These kinds of theories can be mixed and matched in almost any variety that one wants. This claim is sometimes disputed. One may think that the existence of the sentence 'eating meat is wrong' entails that there is such a thing as wrongness. For example, Hussain and Shah say,

> Similarly certain metaethical theories, reductive realism for example, will entail particular normative claims. One cannot claim that 'right' just means "maximizes utility" without its following that if an action maximizes utility, then it is right.

But one can claim this! It may turn out to be the case that the best theory of the meaning of 'right' is 'maximizing utility' (say because one thought that Grice was generally right, and that this is what people have generally been intending to express when they utter the word 'right'). One could also be an error theorist in the sense that, though this is the meaning of the word, it is not the essential characteristic of rightness. Of course, the reductive realist doesn't believe this but this is because of some other assumptions, not because of what it means to be a reductive realist.

Once could also hold the view that the moral words have no meaning and are used to express our moral

emotions, like Ayer did, but also think that there are non-natural properties that we react to. This could come about because one thought that a roughly neo-Wittgensteinian view about meaning was correct and so that a words meaning was given by the function that the word played in the various 'language games' that people use it in. In the moral case one may think that the function of the words is to express the speaker's non-cognitive emotional states. On could also think that the correct account of the nature of moral properties was the Platonic one. Goodness might be a non-natural, non-physical property that we apprehend via the use of reason. It still might not be the case that the function of the word in English was to track that property and assign it to objects in the world.

This all depends on what kind of semantic and metaphysical theory one is attracted to. I don't mean to say that this kind of theory is one that we should adopt. But it is one that we *could* adopt if we thought there was evidence for it. It is not contradictory or absurd in any way. There are philosophers who think that there are non-natural moral properties (Shafer-Landou is one modern

representative)[11] and some of these people may like the Neo-Wittgensteinian philosophy of language. If we find that people do not adopt this combination of views it will be because of independent philosophical convictions not because of some internal inconsistency or incompatibility. So our metaethical theory is completely separate from our normative theory. To be sure there have been what we might call 'characteristic combinations' of these two kinds of theories. So, it is definitely the case that typically, or usually, an irrealist has a certain view about the semantics of moral terms and the realist another but there is nothing in either of these kinds of theories that force us to go one way or the other.

So we cannot define realism in terms of a particular semantic theory. This prompts Devitt to suggest the following way of characterizing moral realism.

> [Moral realism is the view that] there are people and acts that are objectively morally good, bad, honest, deceitful, kind, unkind etc (virtues and vices), acts that one objectively ought and ought not to perform (duties); people who are objectively morally entitled to privacy, to a say in their lives, etc. (rights). That this is so is open to explanation and plays a role in causal explanations (p 6)

The upshot, then, is that moral realism is the claim that there are rights, duties, and virtues that are objective,

---

[11] Though what he calls non-naturalism I would call merely non-materialism or non-physicalism

knowable and causally efficacious in determining our behavior but this way of characterizing the distinction is also flawed. This is because, as we will see in the next chapter, an irrealist like Blackburn could agree with almost everything said above as long as they hold a deflationary view of the claims made. The only thing that might be resisted is the claim about the causal efficacy of these properties. But even then it is not clear that Blackburn can contain his deflationary views.

Another attempt to characterize moral realism comes from Russ Shafer-Landau's recent defense of non-naturalism (Shafer-Landau 2003). He there characterizes

> …the realist position…by reference to its endorsement of the *stance-independence* of moral reality. Realists believe that there are moral truths that obtain independently of any preferred perspective, in the sense that *the moral standards that fix the moral facts are not made true by virtue of their ratification from within any given actual or hypothetical perspective.* (p 15)

He characterizes moral realism in this way because he wants to argue that constructivism is not a form of moral realism. I plan to address the issue of constructivism and its relation to realism in chapter five, so I do not want to get too deep into this issue here; but there is a *prima facie* problem with this way of characterizing moral realism first pointed out by Geoffrey Sayre-McCord

(Sayre-McCord 1986). The issue is that this definition may seem to rule out realism about mental states and so may be too strict. Why should mental states be unreal simply because they are mental, and so dependant on minds?

But this objection is too quick. Shafer-Landa goes on to say that moral standards "are not *made true*, and in particular, are not correct in virtue of being vindicated by some process of (inter)personal election or approbation," (ibid). This allows him to avoid the problem raised by Sayre-McCord. Consider the case of a belief. If I believe that p there is a straight-forward sense in which that fact depends on something mental; namely my having the belief. But there is another sense in which my having the belief that p is mind-independent. This is the sense in which my having the belief does not depend on anyone thinking or judging that I have the belief. The claim that I have the belief is not *made* true by anyone attributing the belief to me.[12] I could believe that p even if no one thought that I did and in fact even if everyone explicitly believed that I did not believe that p.

So too moral claims may be mind-dependent in this attenuated sense (something which Shafer-Landau

---

[12] Though there are some, like Daniel Dennett, who do think that what beliefs a person has is a matter of how we interpret that person.

acknowledges, (p 16)); they would be mind-dependent in the sense that if there were no minds there would be no moral claims. But they would be mind-independent in the fuller sense. Their truth need not depend on anyone thinking or judging that they are true.

But it seems to me unnatural and unfair to define constructivism out of being a realist view. It may, of course, turn out that there is a fundamental distinction to be made as between these two kinds of moral theories but that would require an argument to show it. That is to say, that constructivism is opposed to moral realism needs to be shown by argument rather than stipulation. What we need is an independent characterization of moral realism that would allow us to assess various views to see whether they were realist in this respect or not. This is especially true when we see that most constructivists hold all of the same semantic and metaphysical views as the realist. The constructivist (typically) agrees that there are moral properties, with the stipulation that those properties are made by us. They are made by us in the sense that we construct moral theories and the associated moral principles. These principles then tell us what moral properties are.

In what follows I suggest to use the following criterion as the distinction between realism and irrealism. The difference is to be found in answering the following question: If two people disagree over some fundamental moral claim, like whether unjustified killing is morally permissible, can, in some sense, both be correct? A realist will claim that only one of them can be correct, whereas an irrealist will claim that they can both be correct. This captures everything that we want but does not make the issue one of semantics.

We can then see that there is a range of answers here between a hard realism that says that in every case of a moral disagreement one person, and only one person, can be correct and a more modest form of realism that would claim that only some basic moral disagreements will act this way. Both kinds of realist will think that, at least in some cases, one of the parties to the disagreement has a false belief just like a person who thinks that the Earth is flat, or that 2+2=5, has a false belief. One way in which we could account for error is in terms of moral properties and a semantics where moral terms refer to or stand for the moral properties (whatever they turn out to be). But the way that I have characterized the realism/irrealism contrasts leaves open the possibility

that there are other accounts that do not appeal to moral properties, which I plan to explore in the last chapter, as well as the possibility that constructivism is a realist view.

This way of characterizing the realism/irrealism debate closely resembles Crispin Wright's (Wright 1992) claim that realist discourse involves what he calls 'cognitive command'. A discourse that displays this feature is one where

> It is a priori that differences of opinion formulated within the discourse, unless excusable as a result of vagueness in a disputed statement, or in the standards of acceptability, or a variation in personal evidence thresholds, so to speak, will involve something which may properly regarded as a cognitive shortcoming. (p 144)

One problem that I have with this formulation involves its appeal to the notion of the *a priori*. I will address the issue of the *a priori*, and its relation to analyticity, in chapter four. But aside from that Wright's formulation is on to something important.

I have tried to capitalize on this in the way that I have formulated the realism/irrealism debate while leaving aside the commitment to the *a priori* nature of the commitment. Again, the way that I have formulated it takes notice of the lessons from the debate between the emotivists and intuitionists. From that debate we learned

that taking account of the existence of moral disagreement is an absolutely essential part of any adequate moral theory and the existence of moral disagreement is assumed in my characterization of the debate.

We also incorporate the lesson from cognitive command: when two people disagree over some fundamental moral claim only one of them can literally be correct. Moral discourse respects bi-valence of this sort and when I endeavor to determine whether a certain moral theory is realist or not I will have this criterion in mind. It seems to me that this is an essential component of any moral theory that is worth the name. We saw Sidwick making the same point in section 2.1.1. when he said,

> The peculiar emotion of moral approbation is, in my experience, inseparably bound up with the conviction, implicit or explicit, that the conduct approved is 'really' right—i.e. that it cannot, without error, be disapproved by any other mind.

So this seems to be the best way to characterize the realism/irrealism distinction. It is flexible enough to allow people with diverse views to agree on it and it captures the essential characteristic of realist intuitions. It also allows us to formulate hard-nosed and modest versions of moral realism.

2.4. Lessons Learned

Given what we have seen we can separate metaethical theories in two different ways. We can first separate theories based on their view of meaning. We have one tradition stemming from the work of the later Wittgenstein and espoused by Ayer, and Urmson, that the meaning of a word is given by the use of that word in a language community. On the other hand we have theories that separate meaning from use in the way that Gricean theories do. This is the route that Stevenson, Edwards and Nowell-Smith took. What Stevenson called a word's emotive meaning, by which he meant the psychological affect that a word was disposed to arouse, turns out to be a species of what Grice would call conventional implicature. But it should also be clear that one could have a view that used particularized or generalized implicatures. We see representatives of both traditions in the current literature on expressivism.

We have Blackburn's *(Blackburn 1984; Blackburn 1998)* quasi-realism which is an expressive irrealism as well as Copp's (2001) realist-expressivism. Each of these theorists agrees that when we engage in moral discourse we express a non-cognitive, motivational/emotional state like anger, disgust, or 'norm-acceptance', but they disagree

over the existence of moral properties. On Copp's view when we say that x is wrong we express some non-cognitive state, but we also assert that the action/person in question has some robust moral property. We can thus characterize his view as an assertive-expressivism as opposed to Blackburn's purely emotive-expressivism. According to Blackburn there is no sense in which we assert that some object has some robust moral property when we say that it was wrong.

So each of these theorists accepts a part of the traditional sentimentalist picture, but Copp is a realist because he thinks the primary function of moral language is to stand for robust moral properties, which of course, Blackburn denies. It would seem that there are many ways to combine the semantic claims that expressivism makes with a theory of justification. In the next chapter I examine each of these views and find each wanting. The primary problem that I find is that they are incapable of supporting moral realism. Blackburn's view, as I will argue, for all of its appeal to subtle and sophisticated philosophy of language, and the deflationary view of truth, is simply a modern version of the relativistic subjectivist kind of emotivism that Ayer held. I will then turn to looking at Copp's version of a hybrid

realist/expressivist view before turning to developing my alternative. The primary problem I have with Copp's view is that it too seems to lead to relativism mostly because of his reliance on Grice's notion of implicaturre.

The upshot of this discussion will be that if we want to have a theory that is both emotive (that is, holds that the primary purpose of moral discourse is to express moral emotions) as well as realist (that is holds that when two people disagree about some fundamental moral claim both cannot be correct) we would do best to update Stevenson's view in light of the Neo-Gricean philosophy of language proposed by Bach and Harnish. I will begin to do this in chapters four and five. But first I want to look at Blackburn and Copp.

Chapter 3. Blackburn & Copp

3.1. Simon Blackburn

By far the most well known expressive theory currently on the market is that of Simon Blackburn (Blackburn 1984; Blackburn 1998). It is indeed as well known as it is controversial and for good reason. We are told that we can have the best of both worlds; we can be thoroughgoing naturalists by embracing the sentimentalist tradition, which though not the only theory of ethics compatible with empiricism, is surely one of the most attractive, while at the same time capturing everything that the realist wants to say. We can say both that the judgment that Uday Hussein was an evil man is an expression of our moral sentiment, meaning that there are no moral properties out there in the world that we are describing, and that it is *really true* that he was evil.

The resulting view he calls 'quasi-realism' because it allows us to capture the 'surface features' of realism without actually being realists. We can talk like the realist, saying that of course suicide bombing civilians is wrong, ***really REALLY*** wrong (stamp foot; pound table), while avoiding their commitment to odd non-natural

properties or the mistake of defining moral words in terms of things which themselves are not moral.

Now, as I argued in the second chapter, the moral realist is committed to the claim that two people who disagree can't both be right (at least about some basic moral judgments). This leaves the door open for various realist accounts of justification that would make one right and the other wrong, and commitment to some kind of non-natural moral property would be one way of doing this (though one that I do not think needs to be taken seriously).

But waiving this point for now, Blackburn certainly does talk like a realist as I construe them. For instance in Ruling Passions he says,

> We should think in terms of a staircase of practical and emotional ascent. At the bottom are simple preferences, likes, and dislikes. More insistent is a basic hostility to some kind of action or character or situation: a primitive aversion to it, or a disposition to be disgusted by it…we can then ascend to reactions to such emotions…going up another step, the sentiment may even become *compulsory* in my eyes, meaning that I have become prepared to express hostility to those who do not themselves share it. Going up another level, I may also think that this hostility is compulsory, and be prepared to come into conflict with those who, while themselves concerned at what was done, tolerate those who do not care about it. I shall regard dissent as beyond the pale, unthinkable. (p9)

This final step up the emotional staircase is the point at which we think that anyone who does not share our

attitudes is simply wrong. There is no wiggle room at this point; "I think it is wrong to hurt children for fun, and there is nothing left of *de gustibus non disputandum,"* *(ibid).* This is for the obvious reason that "…ethics does not only concern actions: we…think that in some circumstances people ought to feel various ways," (p 12). Who could disagree with this? This is certainly the way that the realist talks, and as I have claimed in the second chapter, is a necessary part of any ethical theory that takes moral language as seriously as we, its users, do. So far, then, I am in agreement with Blackburn.

On the other hand, he also really does talk like someone who takes the expressive thesis seriously. For instance a little later in Ruling Passions he says that,

> Amongst the activities involved in ethics are these: valuing, grading, forbidding, permitting, forming resolves, backing off, communicating emotion such as anger or resentment, embarrassment or shame, voicing attitudes such as admiration, or disdain, or contempt, or even disgust, querying conduct, pressing attack, warding it off. (p 51)

He is evidently following someone like Nowell-Smith in claiming that there is not one distinctive activity that we do when we engage in moral discourse. I think that this is an important point, and I agree with it. The one thing that we do not do is to describe the world. There is little or no descriptive content to most ethical terms.

Although they look like they are in the business of attributing some property to some object, they are not. They are in an all together different kind of business. Realizing that this was the case for adjectives was as important for the semantics of moral terms as Russell's discovery that definite descriptions were really disguised quantifier phrases was for the semantics of definite descriptions. So far then I am still in complete agreement with Blackburn.

The problems begin when we start to look more closely at the claims that he makes about justification. He talks the talk of realism, of bi-valence for (at least some) basic moral attitudes, but he does not walk the walk. What I mean is that for all that he says there is no theory of justification that backs it up. Or, more precisely, the theory of justification that emerges when we get down to brass tacks is just the same old one that Sidgwick accused the sentimentalists of having in the first place: the justification for suicide bombing being morally wrong is that I disapprove of it, or maybe if we are feeling sophisticated, that *we* (whoever that turns out to be) disapprove of it.

But this cannot be a satisfactory metaethical account, as Blackburn himself often points out. The

problem with this account of justification is the familiar one from Kant that he aimed at the sentimentalists of his time, and adapted by the intuitionists like Ross aimed at the emotivists of his time. The fact that you *happen* to feel a certain way about suicide bombing is at best a fortunate accident, there is nothing really morally praiseworthy about that unless you feel that way for the right reasons, or in other words unless there is an account of why it is that feeling that way is justified. An account, that is, of why it is that we *ought* to feel one way rather than another about suicide bombing of civilians.

No one can accuse Blackburn of not trying to meet this challenge, and his response has become more sophisticated over the years. In his earlier writing he focused on a neo-Wittgensteinian distinction between external and internal questions. In the next section, section 3.1.1. 'Questions, Internal and External' I will examine this early response and show that it is inadequate as it simply begs the question. Blackburn has given up talking in terms of internal and external questions and his current strategy is to revert to the classic emotivist move of invoking the deflationary theory of truth (see section 3.1.2. 'Truth, Justification, and the Expressive Way'). In

fact, Blackburn really makes two claims involving truth that are often not distinguished. This is facilitated by the fact that he simply refers to 'the deflationary theory of truth' as the claim that to say that p is true is to say no more than p or: 'p' is true if and only if p (the famous Tarski-schema, or T-schema for short). But we need to be careful to distinguish the *redundancy* theory of truth as a theory about the meaning of '…is true' and the *deflationary* theory of truth which says that there is nothing more to the property of truth than the T-schema.[13]

Though Blackburn never explicitly puts his cards on the table, he seems to think that the deflationary theory of truth follows from the redundancy theory. He implicitly assumes that meaning is use and therefore that all there is to truth is its use. In section 3.1.3., 'The Meaning and Use of "is True"' I show that the redundancy theory of truth is inadequate as it neglects the important distinction between meaning and use. This is another version of Searle's famous speech act fallacy, which is the fallacy of assuming that since a certain word or phrase is used to perform some speech act it therefore has the speech act as its meaning.

---

[13] I am using these terms in a somewhat non-standard way, but don't let that distract you. The basic point is the familiar one from chapter two that we need to sharply distinguish semantics questions from metaphysical questions.

If this is right, then the implicit argument for the deflationary account of truth that relies on redundancy fails, which in turn means that we need to examine the deflationary account of truth on its own terms. Blackburn argues that adopting the deflationary account of truth excuses him from having to give any further account of justification. In section 3.1.4., 'Truth and the Serious Nature of Moral Discourse' I show that it is in fact the case that the deflationary theory of truth cannot save Blackburn from the charge that quasi-realism collapses into mere autobiography.

3.1.1. Questions, Internal and External

Let us begin by rehashing the internal/external distinction as Carnap (Carnap 1956) made it; after all he was a non-cognitivist in that he held that moral utterances were akin to imperatives (Carnap 1935) and he was avowedly influenced by Wittgenstein so even though Blackburn does not explicitly put it in these terms, it is fitting for us to begin here. What will emerge in this section is that the early Blackburn relied on this distinction to defuse objections to his quasi-realism from those who think that quasi-realism neglects the question of what the justification of our moral sentiments *really*

consists in. This is not to say that he is a logical positivist like Carnap, as he does modify the distinction in an important way, but there is so much in common between the two that it is interesting to point it out. Ultimately this strategy begs the question against the realist, but, as I will argue in this section, we can fix this after which the terminology developed in this section will prove to be useful.

So, Carnap famously introduced what he called a linguistic *framework*. To construct a new framework we needed to do two things. First we must introduce a predicate that lets us

> say of any particular entity that it belongs to this kind (e.g., "Red is a *property*," "Five is a *number*,"). Second, [we must introduce] variables of the new type. The new entities are the values of these variables. (p 17)

After we have constructed a linguistic framework we can ask various questions that are internal to that framework. So, in the framework that Carnap calls 'the system of numbers' we can ask questions like 'is five a number?' and 'is there a largest number?' or 'is every even number larger than the number two the result of the product of two numbers that are themselves prime?' (That the answer to this question is yes is the so-called 'Goldbach's conjecture') and so on. Some of these questions, on

Carnap's view, will follow from the definitions that are given in setting up the framework, in which case the resulting questions have answers that are analytic (like for instance the answer to the question 'is five a number?'), while others will have answers that are 'empirical', or in other words do not follow from the definitions (like the answer to Goldbach's conjecture). [14]

However, there is another kind of question that is strictly speaking nonsensical. These are questions that one tries to ask outside of the framework, or in Carnap's words, "philosophical questions concerning the existence or reality of the total system of the new entities," (p. 18). These kinds of question, on Carnap's view, boil down to practical questions of whether the framework is useful or not. Most of the debate in ontology and metaphysics, according to his diagnosis, results from the participants confusing these two kinds of questions. As Carnap says,

> …strong objections have been raised…against abstract entities as designata, e.g., against semantical statements of the following kind
>
>> (1)   "The word 'red' designates a property of things"
>> (2)   "The word 'five' designates a number"
>
> Those who criticize these statements do not…reject the use of the expressions in questions, like "red" or

---

[14] I am aware that Quine is thought by many to have shown that Carnap's distinction, relying as it does on the analytic/synthetic distinction, is untenable, I will take this up in the final chapter.

"five"; nor would they deny that these expressions are meaningful. But to say that they are meaningful, they would say, is not the same thing as having a meaning in the sense of an entity designated. They reject the belief, which they regard as implicitly supposed by those semantical statements, that to each expression of the types in question there is a particular entity to which the expression stands in the relation of designation. This belief is rejected as incompatible with the basic principles of empiricism, or of scientific thinking. Derogatory labels like "Platonic realism,"…are attached to it.

This dispute is resolved according to Carnap because, from the internal perspective, (1) and (2) are trivially, indeed analytically true, while from the external perspective the only question that makes any sense to be asking is the practical question of whether or not we should adopt the framework and of course we should adopt the system of numbers because it is extremely useful to do so. There are practical reasons like commerce and theoretical reasons like the indispensability of numbers to science that speak in favor of our adopting this framework.

But we cannot ask the question 'are there *really* numbers?' in the external sense, by which we mean something like 'do numbers exist apart from the framework?' This is because as soon as we start to talk about numbers and their existence we have entered the framework and once inside it, the answer is 'of course

they do!' Someone who denied that numbers exist would show that they do not understand the definitions used in its construction. What the person who is asking this question must mean, if they are to make any sense at all, is whether or not we should adopt this framework or not. And there is overwhelming practical considerations that speak in favor of doing so. So we can accept the framework and thereby accept sentences like (1) and (2) and deny the Platonic realism associated with these sentences as a misguided attempt to ask an internal question in an external tone of voice.

Now, if we take this sketch of Carnap's views and translate them into a meta*ethical* doctrine instead of a meta*physical* one what we end up with is pretty much identical to the approach that Blackburn takes in his earlier work, especially as presented in his 1988 paper "*How to Be an Ethical Antirealist*" (Blackburn 1988). Thus he says,

> The projectivist can say this vital thing: that it is not because of our responses, scrutinized and collective, or otherwise, that cruelty is wrong. The explanation flows from the way that quasi-realism has us deal with oblique contexts. It issues an "internal" reading of the statement of dependence according to which it amounts to an offensive ethical view, about (of course) what it is that makes cruelty wrong. Critics of this explanation allow the internal reading but complain that the quasi-realist is being willfully deaf to an intended "external" reading, according to which the

dependency is a philosophical thesis, and one to which the projectivist must assent. (p. 173)

This passage could be from Carnap if one just switched the terms 'good' and 'bad' with 'red' and 'number'! Saying that the wrongness of cruelty depends on us is as absurd as saying that 2 being even depends on us. Nobody in their right mind would assert that (in the internal sense). But what critics of quasi-realism are trying to ask is the external question about the existence of the totality of the system.

Blackburn goes on, in true Carnapian style, to deny that there is any way to ask the question about what makes cruelty wrong without asking "a moral question, with an answer in which no mention of our responses properly figures," (ibid). Of course he cheerfully admits that the external question would make sense if realism were true, but by realism he understands the view that there must be some fact or state of affairs out there ("the wrongness of cruelty") that our moral language tracks, which is of course absurd!

Does Blackburn follow Carnap in thinking that the acceptance of the moral framework is a matter of what is practical? Well, he says that his strategy is to "confine external questions of dependency to domains where real

states of affairs, with their causal relations, are in question," (p. 173). So, the view that he seems to be advocating is this. The external questions, that is, questions about what properties there are, is answered by our commitment to naturalism and so the answer is that there are no non-natural properties for moral terms to describe. Blackburn also finds Moore's open question argument convincing (see for instance p. 86 of <u>Ruling Passions</u>) and so he thinks that there are no natural moral properties either.[15] Given that he thinks this is the case, he then argues that the best account of morality from the external perspective is that of the sentimentalist tradition. This does not look like he is saying that it is merely a practical decision. We have real reasons to believe empiricism is true and so real reasons to accept a sentimentalist account of morality.

However, he continues, the internal questions, that is, questions about what is right or wrong, good or bad, have to be answered from within the moral framework and from within that framework how could cruelty NOT be wrong? And how could its being wrong depend on our feelings about it? To say otherwise is to be immoral, to express a repugnant

---

[15] We have seen (section 2.3.) that this argument fails, but the point is that Blackburn reasons this way.

moral attitude. As he says, "in saying these things I am, of course, voicing some elements of my own ethical stances, but as promised, it is only by doing this that ethical truth is found," (p. 176). This is because the 'ethical truth' can only be found within the moral framework and from within that framework cruelty is wrong, to say otherwise is to betray an immoral character. And if we try to ask 'yes, but why is cruelty wrong?' and mean by that the external question, we show that we are confused or some kind of anti-naturalist.

In response to the Frege-Geach problem Blackburn develops a general logic of attitudes. Hare, as we saw (in section 2.1.2), developed the rudiments of imperative logic but he did not extend this notion to a general logic of non-cognitive attitudes. Blackburn begins by asking the Frege-Geach question about beliefs. The basic Frege-Geach point was that it seems implausible to think that when 'suicide bombing civilians is wrong' occurs in the antecedent of a conditional it is there being used to express a feeling. But if the argument is to be valid it has to have the same meaning in both contexts. What happens when we ask this about beliefs? As Blackburn says,

> Suppose I say that the sentence 'Bears hibernate' expresses a belief. Well, it only does so when the sentence is put forward in an assertoric context. So

what happens when it is put forward in an indirect context, such as, 'if bears hibernate, they wake up hungry'?…The standard answer is to introduce a proposition or thought, regarded as a constant factor in both the assertoric and indirect contexts (RP p. 71)

In one case the thought or proposition is asserted in the other it is 'offered up conditionally or as a supposition'.

He then suggests that the expressivist should adopt this very same strategy. The expressivist introduces the 'attitude' which is postulated to be constant between the expressive use and the indirect use. In one case we express it in the other we offer it up conditionally or as a supposition. He then introduces his notion of 'being tied to a tree' which is the attitude logic equivalent of validity. As Blackburn says,

> Suppose I hold that either John did the deed or he is to blame. Then I am in a state in which *if* one side is closed off to me, I am to switch to the other —or withdraw the commitment. And this is what I express by saying 'Either John is to blame, or he didn't do the deed', or equally 'if John did the deed, he is to blame'. By advancing disjunctions and conditional we avow these more complex dispositional states (Blackburn 1999 p. 71)

He then introduces new operators for this attitude logic. He introduces 'H!' and 'B!' respectively so that, syntactically 'H!(p)' and 'B!(p) are well formed formula and semantically are interpreted as 'hooray p!' and 'Boo p!'.

In conditional contexts he introduces '|H!(p)|;|H!(q)|' as a way of saying that having a 'hooray q' attitude is dependent on having a 'hooray p' attitude (Blackburn 1984, pps 193-194). The ';' is supposed to function like an attitude analog of material implication and the '|' symbols around the antecedent and consequent of the attitude conditional are supposed to indicate that the attitude is being put forth in a conditional way.

He can then interpret the Frege-Geach argument as follows.

        1- B! (suicide bombing civilians)
        2- If |B! (Suicide bombing civilians)|; then
           |B! (Recruiting young people to suicide bomb
           Civilians)|
        So, B! (Recruiting young people to suicide bomb
           Civilians)

Someone who assented to 1 and 2 but did not have the resulting attitude would be taking away in one breath what he had just given in the previous breath and "we can make no sense of them" (p. 72). This is not exactly the same thing as logical validity but we can get all of the truth-functional analogs. This seems to me to be an important step forward for emotive-expressive theories of moral reasoning. It is an improvement on Hare's first attempt and seems a valuable asset to any expressive theory.

But what about Blackburn's attempt to collapse the internal/external distinction? Evidently we cannot be satisfied with this answer; for there is a way of meeting Blackburn's ontological challenge (p. 172). This is the Rawlsian inspired constructivist option as defended, for instance, by Steven Ross (Ross 1991; Ross 1998). This kind of view can accept pretty much everything that Blackburn has said so far. Where the disagreement arises is over what counts as real. So, the constructivist will gladly admit to naturalism and agree that this means that there are no mysterious moral properties. Moral properties are the result of moral principles which are constructed by us; nothing mysterious about it. On this view we do not find a world with the moral properties in them, but rather we construct those moral properties ourselves by constructing moral theories.

There are options out there for people who want to ask this external question. Threatening us with the title of Platonists or anti-naturalists is not enough to get us to swallow Blackburn's program. We can sensibly ask the external question that Blackburn thinks is impermissible without thereby committing ourselves to the weird claims that Blackburn is so understandably eager to avoid.

The upshot of this discussion is that the early Blackburn simply begs the question against the realist. If there is any way of being a naturalist/empiricist and a moral realist then we need an argument to the effect that the external question is one that cannot (or should not) be asked. This he has not given. He has in fact admitted this, and then simply tried to bully us with the threat of being labeled an anti-naturalist. He is of course right that cruelty is not wrong because of our reaction to it, but telling us that he, Blackburn, thinks that someone who denies this is immoral does not put my mind at ease. WHY is it immoral? What is it, other than our emotional reactions to cruelty, which grounds this judgment? This is the important question and this question is not answered, it is in fact, dismissed as a misguided attempt to ask an internal question externally.

So far, then Blackburn has said no more than Ayer did. Suicide bombing is wrong because I feel a certain way about it. Period; and we have already seen that this simply is not an acceptable account of justification. What needs to be explained is why it is proper or right to feel that way about suicide bombing. And this is the external question, not the internal one which simply and mind-numbingly intones that to say otherwise is to be immoral;

that surely is right, but WHY? The question, then, is 'how does the framework connect to reality? If it does not, as Blackburn seems to suggest, then who cares? There could then be many moral frameworks, and from the external perspective we would see that relativism was true. After all, isn't this just what it truly means to say that the empirical facts do not determine the moral facts? So adopting this strategy seems to lead quasi-realism to relativism.

But Blackburn is insistent that quasi-realism does not lead to relativism. People who are keen to press this kind of attack against him, he points out, do so by invoking truth. The question, the objector asks, is which framework *corresponds* to the way things are? Putting things this way makes the claim that the frameworks do not connect to reality just admitting that relativism is true. All there is to morality is people and their emotional responses to the world, and lots of people have lots of different emotional reactions. His current strategy, then, is to appeal to deflationism about truth as a way of fending off these kinds of charges. There is no way even to ask this question about how the frameworks connect to reality because there is no such thing a robust property of truth that would make that question intelligible. I

will now turn to discussing this move and show that it too is inadequate.

3.1.2. **Truth, Justification, and the Expressive Way**

At the common sense level, truth and justification are related to each other in the following way. Some theories of justification tell us that there are true moral judgments and other theories of justification deny this.[16] And as we saw in the second chapter, the primary distinction between the realist and the irrealist is that the realist thinks that there are at least some cases where two people disagree over some basic moral judgment and only one of them is actually right; or in other words one of them has got to be wrong.[17] So, to say that it is true that cruelty is wrong is to say that someone who thought that it was false would be mistaken in just the same way as someone who thought that 2+2=4, or that

---

[16] I suppose I ought to say that for constructivist the issue is not necessarily one of the truth of moral principles but will rather be one of the *reasonablness* of the principles that we construct (i.e. those that we can agree on behind the veil of ignorance); at any rate none of this will really affect my main point. For, even if one prefers this way of talking, it will still be the case that various moral judgments will be true in virtue of the constructed 'reasonable' principles and so there will be cases where two people disagree and only one them can be right. This is especially evident when we consider that even Rawls (Rawls 1980) thinks that the two principles of justice are maximally reasonable and that the possibility of their turning out to be unreasonable at some time in the future is only a logical possibility, more akin to the possibility of a talking pizza than to flying pigs (a nomological possibility).

[17] Though of course it may be the case that neither knows who is actually right.

hydrogen is the lightest gas was false would just be plain wrong. This is in fact the reason that so many moral realists are led to talking about moral properties, for how else could a moral judgment be true if it did not track some moral state of affairs?

The irrealist denies that there are any such cases, mostly because of a view about what it means for something to be true in this sense. The typical irrealist is someone who shares the realist's conception of what it means for something to be true, that is, they agree that if moral judgments were true or false they would have to track some external fact or property or state of affairs, and go on to deny that there are any such things as those. Stevenson is the classic example of an irrealist on my view. Ultimately two people who disagree (and yet agree on all the facts) simply feel different ways about the matter in just the same way that even were you and I to agree on all the facts about the world it could still be the case that I liked chocolate ice cream and you disliked it. There is no interesting sense in which you are wrong because you do not like it, or in which I am because I do. This truly is a case of *de gustibus non disputandum*.

Blackburn tries to collapse this distinction by invoking minimalism about truth. He makes it clear that he thinks

that once the deflationary theory is draped over his expressivism the problem of justification is automatically solved. The whole quasi-realist strategy is to argue that once truth is deflated we can say whatever we want, it all comes for free. This is because for something to be true, for the deflationist, is simply for it to conform to the T-schema. In fact it seems that for the most part the 'quasi' in quasi-realism *just is* minimalism about truth. It is his minimalism about truth that he thinks allows him to talk like the realists and avoid their commitment to moral properties.

Let's look at how this is supposed to work. The T-schema tells us that the proposition that p is true if and only if p is the case and the deflationist about truth holds that this is all there is to a philosophical theory of truth. As Blackburn says,

> …the detailed differences in the interpretation of the truth-schema are not important here. But according to the deflationist [the T-schema] encapsulates all we need to know about truth. (Ruling Passions p 75)

It is not simply that the deflationist accepts the T-schema, as that is supposed to be neutral as between philosophical theories of truth, even the correspondence theorist thinks that the T-schema captures something important about truth. The deflationist makes the further

claim that there is *nothing more* to truth than what the T-schema tells us. There is no interesting philosophical problem that demands a philosophical answer like truthmakers, or correspondence.

Applying this to a theory of ethical discourse gives us the following view. When we engage in moral discourse and say that things are right or wrong, or should be done or should not be done, what we do, according to Blackburn, is to express our moral sentiments. This is no more than what the emotivist irrealists have been saying all along. Blackburn then points out that according to the deflationary accounts of truth, to say that these sentiments are true or false is simply to say that we accept or do not accept them. It is simply to say that we approve or disapprove of them. To say 'it is true that the Taliban's practice of subjugating women is wrong' is to say no more than 'the Taliban's subjugating of women is wrong,' which will be true just in case the Taliban's subjugating of women is wrong as per the T-schema (Blackburn 1999). There is nothing more to the truth of this sentence than that.

Blackburn triumphantly trumpets this result saying "we can add flowers without end": that is, not only can we say

that 'the Taliban's subjugating of women is wrong' is true but, if we want, we can say that it is

> really true, or really factually true, or really in accord with the eternal harmonies and verities that govern the universe… (p 79)

We can even throw in some good old-fashioned table pounding and fist pumping if we feel like it. Doing so indicates our ascent up the emotional ladder. But since truth is deflated we do not have to worry about those pesky Platonic properties.

This, in turn, provides a new response to the Frege-Geach problem. Since we can say that these evaluative sentences are true (in the deflationary sense) there is no problem about moral inferences. So if one was worried that the 'tied to a tree' strategy discussed in the previous section was not enough because it wasn't formal validity we simply invoke deflationism about truth and get the logical relations for free. There is no problem in saying that moral arguments are truth-preserving because we are not thereby committing ourselves to moral properties or a correspondence claim. The argument we have been considering,

> 1- Suicide bombing civilians is immoral
> 2- If suicide bombing civilians is immoral, then recruiting young people to suicide bomb civilians is wrong
> So, recruiting young people to suicide bomb civilians

is immoral

comes out valid, in the real sense. Since if 'suicide bombing civilians' is true then the conclusion straight-forwardly follows from 1 and 2.

This strategy in emotive theorizing has a history as long as emotivism itself. Ayer and the logical positivists all accepted it, though only Stevenson explicitly acknowledges that adopting this theory of truth allows us to say that moral utterances are really true. He says this in <u>Ethics and Language</u> (p. 169-171), but he is exceptionally clear about it in his retrospective essay in <u>Facts and Values</u> (Stevenson 1963) where he says that it is absurd to say that ethical judgments are neither true nor false, as he says,

> Such a view would represent not an effort to preserve our normal habits of speech but rather an effort to reform them. And…this particular reform shows every sign of being so inconvenient that its advantages (if any) would fail to justify it, (p. 216).

He goes on to make the very important point that the fact that we can say of 'p' that it is true no matter what 'p' happens to be is a *syntactical* fact about language which shows "nothing whatsoever about whether [p] expresses a belief or an attitude or both or neither".[18]

---

[18] Interestingly he does point out that this syntactic fact should warn us about treating ethical judgments as imperatives, though they are akin to them

This of course means that it must be possible to give an account of what it means to say that it is true that cruelty is wrong. Now, he continues, this would be worrisome if that meant that we were committed to realism, but "as Ramsey has shown", to say that 'p' is true is simply to say that p in a louder tone of voice (see especially p 214-220 of <u>Facts and Values</u>). So, Blackburn's strategy is nothing new; given the redundancy theory of truth there is no problem of moral justification. All we can mean when we say that it is true that cruelty is wrong is that I really feel that it is wrong! So, Stevenson is the first quasi-realist.

Now so far we have been talking about the *meaning* of '…is true'. In Blackburn's eyes Wittgenstein and Ramsey showed us that '…is true' is used to express agreement and so that is its meaning. This is because he thinks that the function that a word serves in a language is the word's meaning. As I will put it in the next chapter, he thinks that the L-semantics of English just is the P-semantics of thought. It is for this reason that Blackburn, and Stevenson for that matter, seem to think that deflationism about truth (a theory about the metaphysical nature of truth, the property) follows from the redundancy theory (a theory about the meaning of 'true').

But when we look closer at this strategy it falls apart like gossamer in our hands. The reason turns out to be partly due to the virtue that attracted Blackburn in the first place. Since deflationism adds nothing, it can't add a theory of justification. So whatever theory was there in the first place will still be there after the deflationsism is added. To see why, as I have already indicated, I think that it is vital that we distinguish the redundancy theory of truth from the deflationary account of truth. The former is a claim about language use, the latter a claim about the nature of the property. These are two importantly different ways of cashing out the slogan that the T-schema captures all that we need to know about truth. Conflating them is another instance if what we saw in the last chapter. Semantical and metaphysical issues are separate.

However, as I will show in the next two sections, this appeal to minimalist accounts of truth is flawed in both cases. It is not an adequate theory of the *meaning* of the predicate '…is true,' though it may be an account of the way that we commonly *use* the predicate in conversation, and it in no way supplants the need for an adequate theory of the justification of moral judgments. Whatever theory of justification that is there in the first place will

continue to be there even after the deflationary account of truth is in place. In Blackburn's case this means that his view really is relativistic.

3.1.3. The Meaning and Use of 'is True' in English

As is characteristic of neo-Wittgensteinians, Blackburn collapses the distinction between meaning and use. On his view the meaning of words is given by the function that they perform in language. So in the case of moral words that will be the stuff listed in the quotation above and so accordingly their meaning is their function of grading, expressing sentiment, etc. So, since the function that 'is true' plays in English is one of expressing agreement, the T-schema gives us all there is to the meaning of the word. To say that p is true is just to say that p. So, is there any reason to think that the redundancy theory of truth captures all there is about the meaning of '…is true'?

No. While I do think that people often use the word 'true' as a way of communicating that they agree with either what they themselves, or someone else, has said this communicative *use* of the predicate '…is true' depends on its having the correspondence *meaning*. 'True' means something like 'being in accordance with the actual state of affairs' and so it is easy to see how I could use it to

express agreement with what has been said. To say that something is true is to say that it is really the way things are. So in conversation I am able to exploit that meaning of the word in order to indicate that I agree with something that has been said. For example when you say that Osama bin Laden is a moral criminal and I say 'that's true' you know that I am expressing my agreement because I am in effect saying 'yes, that is in accordance with the facts'. It would be very hard to explain how it is that the predicate '…is true' came to have the function that Blackburn and Horwich (Horwhich 2005) think that it does if it did not have the correspondence meaning. If I did not know that the meaning of '…is true' was something like 'in accordance with the facts' what would ever make me think that you would know that I was agreeing with what you said when I said that it was true?

We exploit the meanings of words in this way quite often. Searle (Searle 1969/2001, p. 142) pointed out a similar phenomenon with 'promise.' Suppose a parent says to their lazy child "clean your room or I promise I will take away your cell phone!" It is very odd to think the parent is actually promising to do anything here since the thing promised is not something that the child wants the parent to do. In fact this kind of utterance is most

likely a threat or a warning. Or again consider a professor confronting a student suspected of plagiarism. The professor says "this passage is taken from wikipedia" and the student says "I didn't plagiarize! I promise I didn't!" This doesn't look like a promise either, how can you promise that you did not do something? This is rather an emphatic denial of the professor's accusation.

How is this possible? As Searle points out,

> I think we use it here because "I promise" and "I hereby promise" are among the strongest illocutionary force indicating devices for commitment provided by the English language. For that reason we often use these expressions in the performance of speech acts which are not strictly speaking promises, but in which we wish to emphasize the degree of our commitment. (p 142)

That 'promise' is an illocutionary device for indicating commitment is part of the meaning of the word, which is something like 'a declaration to do something in the future'. Again, if this weren't so how could we possibly expect anyone to figure out what we were doing in the kinds of utterances where we use 'promise' to do something other than promising?

Since Blackburn takes meaning to be use this is lost to him. Since he thinks that the function that the predicate '…is true' has of indicating agreement is all that there is to the meaning of the word he thinks that there is no

property that 'is true' stands for. The use theory of meaning makes the correspondence theory of truth look like a mistake since it claims that '…is true' has some other meaning than this function of indicating agreement. It means something like 'corresponds to reality', but that is not the way people use it! It is in this way, I think, that people like Stevenson and Blackburn implicitly arrive at the conclusion that the redundancy theory, a theory about the meaning of a word in English, leads to the deflationary theory, which is a metaphysical theory about the nature of a property. Once we distinguish the meaning of a word from the use of that same word the deflationary theory of truth is no longer the only option. We can agree that we often do use 'true' as a way of *communicating* that we agree or accept some moral standard but since this is dependant on the inflationary or non-redundant meaning of the word it may turn out to be the case that we can give a truthmaker for moral judgments, or in other words we could have a realist theory of justification and make the charge of relativism stick against Blackburn.

So, the next question is, 'does the deflationary theory of truth help Blackburn with his account of justification?' In the next section I will argue that it does not.

3.1.4. Truth & the Serious Nature of Moral Discourse

The basic problem is that deflationary truth is just too cheap. This is especially evident when we consider the fact that whatever account of moral contradiction that the quasi-realist can give will also be the correct account of contradiction in matters of taste. So 'broccoli is disgusting' will be true if and only if broccoli is disgusting and someone who said that it was not would really be contradicting me. From within the 'taste framework' broccoli is disgusting and I can just see that the Broccoli-ban and their feelings about the taste of broccoli are just objectively wrong. Of course all that any of this means is that I accept or agree with the sentiment that I expressed when I said that broccoli was disgusting. The story we tell here exactly parallels the story that is told in the case of moral judgments about cruelty, the Taliban, or whatever.

But clearly there could not be more of a difference between these two kinds of judgments. In particular, it seems obvious that this story about broccoli is just wrong. Common sense tells us that our feelings about broccoli may depend on two things. One, we may think that broccoli has a certain specific kind of taste and some people like *that* taste and others dislike it, which one it

is may depend on what the person can taste, or it may depend on how they were raised, or just simply that they are disposed to like it or not and all of these vary from person to person. So there is nothing wrong with a person who thinks that broccoli tastes good, they simply have different tastes than ours and which you have doesn't really matter. On the other hand we might say that broccoli has no determinate taste, it all depends on the person who does the tasting and the way that their taste buds are constituted. Taste is a secondary property whose reality is totally mind dependant. So whether it is disgusting or not is relative to a person's make up. Either of these common sense explanations of what is going on in the broccoli case differs dramatically from the common sense view of moral discourse. Only a madman would claim that our feelings about Saddam Hussein, the slaughter of children, truth telling, or promise keeping depended on us in either of the two ways mentioned above. Even Blackburn is not that reckless! He explicitly denies that anything like this is the right way to characterize moral disagreement. But the problem is that there is no way to distinguish these kinds of claims from the theoretical stand point of quasi-realism.

Since the theory is unable to distinguish these obviously distinguishable kinds of judgments, there must be something seriously wrong with the strategy of appealing to deflationism about truth to supplant the need for a theory of justification. In fact, it seems obvious what is wrong with it. It very obviously and flagrantly turns moral matters into matters of personal taste. It does this by claiming that all there is to truth is its function in natural language of voicing agreement. To say that something is true is simply to repeat what we have said. Whether we happen to have said something about rape or the taste of broccoli makes no difference. Once we take the deflationary account of truth seriously we are no longer able to take moral discourse seriously.

Blackburn cannot respond that we can distinguish talk about broccoli and talk about genocide by the level of emotional commitment that we have to claims in one area as opposed to claims in the other because it is not inconsistent, on his view, that there be people who take broccoli as seriously as we take suffering. Thus the Broccoli-ban are every bit as serious about people who disagree with their feelings about the taste of broccoli, even to the point of putting dissenters to death. It may be the case that Simon Blackburn does not take talk about

broccoli that seriously, but so what? If this is to be anything more than a mere autobiographical report what we need is a way to say that someone who did take talk about broccoli as serious as the Broccoli-ban was *mistaken* and further that their being mistaken is not simply an opinion of mine. Something, in short, that allows us to distinguish our talk about what depends solely on us and what does not. Quasi-realism fares very badly here.

Not only does quasi-realism have no way to distinguish between the Taliban and the Broccoli-ban that is not mere autobiography we can see that the very same problem arises for other moral claims. Suppose someone from the Taliban were to respond to Blackburn that their views on women were the correct ones to have and that Blackburn was wrong when he says that they (the Taliban) are objectively wrong. Let us suppose that they laugh at the idea that women are equal to men in any serious way. Then, according to the analysis that is on offer we are to conclude that what they have said is true just in case they really hold the attitudes that they say they do. Blackburn then points out that they are 'blind to the nature of women and the possibilities open to them' and so on, but the important question of WHY it is that the Taliban have to agree with him on this point is left begging to be addressed. Of

course by this I do not merely mean that the Taliban may irrationally refuse to admit that the evidence against them is compelling but rather the stronger claim that in some deep sense there is no way to really say which is right here. Each is saying something true when they express their moral sentiments about women. This is, of course, nothing more than relativism.

Here is another way to make the argument. Take a sentence like 'suicide bombing civilians is wrong' this sentence will be true, *ex hypothesi*, if and only if suicide bombing civilians is wrong. But what are we to make of the right hand side of this equivalence? We cannot take it to represent some moral state of affairs or fact. We, presumably, must take it to have its usual meaning and so it seems we must give it Blackburn's usual analysis. Doing so yields the B-schema;

> B-Schema: 'Suicide bombing civilians is wrong' is true if, and only if, Boo suicide bombing civilians!!

But now it is clear that he has done no more than to dress up the standard irrealist theory of justification.

This, then, is his new way of defending the external/internal strategy. Since he thinks that deflationsim is the right account of the nature of truth he does not see any way for us to ask the external

question. But as we have seen there is no reason to take this account of the nature of truth seriously. The redundancy theory of truth is inadequate as an account of the truth predicate in language and the deflationary account does nothing to quell the charge of relativism that Blackburn is so anxious to avoid. In all of these cases, then, what we have seen is that Blackburn simply begs the question against his opponents. In the first instance he begs the question against the realist by assuming that the only realist account must be the Platonic one that appeals to other-wordly moral properties by neglecting, for instance, the constructivist approach, and in the second instance he begs the question against the realist by assuming that all there is to the meaning of '…is true' is its use to express agreement and so neglects the question 'what connects the moral frameworks to reality?'

Interestingly he betrays this in the appendix to Ruling Passions where he addresses some common questions, which I think is worth quoting in its entirety. In response to question 18, which was 'aren't you really trying to defend our right to talk 'as if' there were moral truths, although in your view *there aren't any really*?' He responds,

No, no, no. I don't say that we can talk as-if
kicking dogs were wrong, when 'really' it isn't wrong. I
say that it is wrong (so it is true that it is wrong, so
it is really true that it is wrong, so this is an
example of a moral truth, so there are moral truths).
This misinterpretation is curiously common. Anyone
advancing it must believe themselves to have some more
robust, metaphysically heavyweight conception of what it
would be for there to be moral truths REALLY, and
compared with this genuine article, I only have us
talking as-if there are moral truths REALLY. I deny that
there is any such coherent conception.
Even if we sorted truth into TRUTH [a correspondence
notion] and truth [the deflationary notion] and decided
that there was no moral TRUTH, this would mean only that
you don't walk into rights and duties… (p. 319)

He just completely overlooks the possibility that there

might be some notion of moral TRUTH that is not committed

to the truthmakers of moral judgments being actual

entities. If this were so, then his diagnosis of himself

in the second paragraph would be completely spot on.

Compared to this genuine article he does have us only

talking as-if.

This concludes my critique of quasi-realism. I turn in

the next section to discussing the view of David Copp as

expressed in his 2001 paper *"Realist-Expressivism: a

Neglected Option for Moral Realism"*. This view has more in

common with the view I will develop but, as I will argue,

is unsatisfactory.

3.2. David Copp

On Copp's view a moral utterance of 'suicide bombing civilians is wrong' expresses the belief that suicide bombing civilians has the robust moral property of being wrong. So his view is a realist view by the semantic standards discussed in the second chapter. He treats moral properties in just the same way that he treats other properties. However, he argues that when we make this assertion we also implicate that we subscribe to the norm from which this judgment follows. The state of mind that we express by so implicating is a non-cognitive motivational state akin to what Gibbard calls 'norm acceptance'.

So, claims Copp, his realist-expressivism accepts key theses of both realism and expressivism. It is realist in that it claims that there are moral properties and that when we say that things are wrong we assert that various things have these properties. It is expressivist because it holds that we also express a non-cognitive motivational state. He then argues that any theory of justification is compatible with the realism in realist-expressivism. That is, whatever one thinks is the correct view about justification will determine which state of norm-

acceptance one is in and this will determine the truth-conditions for the utterance.

I very much agree with what Copp calls the 'two proposition model' which is the claim that when we say that something is right or good we express more than one mental state. How should we construe the relation between these two things expressed? Is it a Grician implicature in the way that Copp suggests? Since Grice distinguished different kinds of implicatures there will be different versions of realist-expressivism, which I will discuss in the next section (section 3.2.1). In section 3.2.2., I argue that there seems to be a problem with modeling the relation as a conventional implicature (as Copp wants to) because doing so collapses the view into a crude form of relativist subjectivism. On the other hand to construe it as a conversational implicature seems wrong as well. It is enough to say that we express both states when we say that suicide bombing is immoral.

There is another problem with Copp's account. His kind of expressivism is not an *emotive*-expressivism because he thinks that the 'primary' speech act is the expressing of a belief. He defines 'assertion' as the expression of a belief and claims that in making a moral utterance we perform the speech act of asserting. For this reason I

call   Copp's   kind   of   expressivism   an   *assertive*-expressivism.

3.2.1. Different Versions of Assertive-Expressivism

The basic claim that Copp wants to make is that when we say that X is wrong we may express multiple propositions. So, if we can express multiple propositions then there are several different possibilities for what we take the relation between these propositions to be. This gives rise to the different versions of assertive-expressivism. There are at least five of them as Copp counts.[19] We have,

> (1) the no [connotive/motivational] proposition view, (2) the entailment view, (3a) the conversational implicature view, (3b) the conventional implicature view, and (3c) the neither of the above view. (p 25)

It is obvious that (1) is just an ordinary version of realism that denies that we express any non-cognitive state when we say that x is wrong. Our sole speech act would be to assert that some person or action has the moral property in question. The problem for this simple version of realism, according to Copp, is the immoralist. Someone who says that suicide bombing is wicked but then goes on to say that they do not approve or disapprove of it or that they are not motivated to avoid it and condemn those that practice it, etc, is possible but none the less

---

[19] I think that for each of these types of assertive-expressivism we have a corresponding kind of emotive-expressivism

strange. The strangeness of a person like this is *prima facie* evidence that we do typically express two attitudes when we moralize. This leaves us with (2), (3a), (3b), and (3c) as candidates.

The entailment view (2) is a form of subjectivism. What it claims is that when I say suicide bombing is wicked I assert that suicide bombing has the robust moral property of wrongness and this entails that I have the non-cognitive attitude that I express. So what I say is true when I do in fact have that attitude. This in turn is to say that moral judgments are true when I feel the way that I say I do. To say that my utterance entails that I have the attitude means that I cannot make the assertion that a certain person/action has a moral property without having the corresponding non-cognitive attitude.

Copp thinks that the immoralist poses a problem for this view as well because they are an example of someone who lacks the emotive state and yet it does not seem natural to think that what they say, namely that suicide bombing is wicked, is therefore false. It may still be true even if they lack the emotive/motivation state. So an utterance of 'suicide bombing is wicked' does not entail that I have the corresponding motivational state (p. 25).

That leaves the last group of three. (3a) and (3b) invoke the Grician apparatus of implicature. To repeat, a *conventional* implicature is one that is somehow carried off by the meaning of the utterance but that is not part of the truth conditions of what was said. So, to take the famous example, an utterance of 'she is poor but honest' is supposed to conventionally implicate that there is some kind of contrast between being poor and honest, but the truth conditions of the utterance are just the usual ones for 'and'. In short, it is not rendered false if the relevant contrast (the implicated proposition) does not hold. A *conversational* implicature, on the other hand, is one that is carried by a violation of some maxim or mutually held belief and does not depend on the meaning of what is said. Rather it depends on the hearer working out via inference what the person means by saying what they did. Copp argues for the conventional implicature view because he thinks that it captures better what he takes to be going on in moral discourse.

Copp's account of how this works is given by his account of attitude expression that he calls 'Frege-expressing'. The central idea that he has is that terms have 'coloring' in roughly the sense that Frege talked about. This is a property of the meaning of the terms. One

of his examples is the term 'Yankee.' In some circles this word is colored in such a way that someone who said it would be taken to be expressing contempt for the individual that they said it of; in particular it is a way of slamming Americans. He gives us the following four 'rough tests' for coloring.

> if a speaker's assertion of a sentence S implies that p, and if this implication is due to the coloring of a term T contained in S, then (1) the belief expressed by the person in asserting S might be true even if p is false, and the implication that p should be (2) detachable and (3) cancelable. Furthermore, (4) it would be a misuse of the term T for a speaker to assert S when she knows or believes that p is not he case. (p. 20)

So consider the sentence 'Alice is a Yankee'. This he claims would normally (we assume) be used to express contempt for Alice. But the truth of the sentence depends only on whether or not Alice is American and not on whether the speaker has contempt for Alice. Imagine, for instance, that this is the only word for Americans that someone knows but they think that Alice is one of the 'good ones' and so refer to her by it but do not mean to express contempt for her thereby. So the belief expressed may be true even if the implicated thing is false.

By *detachable* he means that we could express the same belief by substituting 'American' for 'Yankee'. Someone who said that Alice was an American would express the very

same belief as the person who says that she is a Yankee but without also expressing the contempt. The expression of contempt is therefore a detachable part of the speech act performed. A speaker could *cancel* the implication by saying that Alice was a Yankee, but that they did not have contempt for her. We can imagine that the person from the above paragraph who liked Alice but does not know any other word for American other than 'Yankee' might say something like this to cancel the implicature.

But even so, Copp claims that it would be a misuse of the term 'Yankee' to do so. It would need some kind of explanation as to why the person called Alice a 'Yankee' but did not feel contempt for her. He calls this an aspect of the meaning of the word because it is governed by linguistic conventions. This is closely related to Stevenson's notion of emotive meaning. Whereas for Stevenson the term had an emotive meaning because of its disposition to cause certain feelings or attitudes, for Copp a term's coloring is a matter of linguistic convention (p. 14). So 'coloring' is Copp's Grician version of emotive meaning. In the next section I will argue that this model of how moral utterances work ultimately collapses into the entailment view and so leads, like Blackburn's view, to a subjective relativism.

3.2.2. The Collapse into Subjectivism

There are two kinds of argument, both of which Copp is aware, which threaten to collapse his view into a subjectivist view. One is from Frank Jackson and Philip Petit in their paper "A Problem for Expressivism" (Jackson and Petit 1998) the other is from Kent Bach's paper "The Myth of Conventional Implicature" (Bach 1999). As I have said, Copp is aware of both of these attacks and he thinks they amount to same thing. The basic idea behind the attacks is that we cannot draw a distinction between conventional implication and entailment. But this lumping the two arguments together obscures the force of Bach's criticism. I want to look first at the Petit & Jackson challenge and then argue that Bach's is more serious and Copp does not really address it by addressing the Petit & Jackson argument.

I think Copp has the argument from Jackson and Petit right. It is a challenge to show how

> It could be that a person uttering a sentence communicates a proposition, in virtue of linguistic conventions governing the literal use of the terms in the sentence, without thereby asserting the proposition such that the truth of what the person says or asserts depends on the truth of the proposition. (Copp p. 26)

If the sentence's truth depends on whether the speaker has the state in question or not then it collapses into (2)

and is simply a version of relativism. And if there is a convention in English that connects the sentence 'x is wrong' to the norm that justifies it then we have the connection that leads to relativism. Copp responds to this attack by arguing that there are pragmatic conventions as well as semantic conventions.

He compares this to promising. There is, he says, a convention in English that "it is appropriate for us to promise…only when (we believe) the relevant sincerity condition is fulfilled". What this means is that there is a rule governing promising which says that we should only promise if we really intend to fulfill the promise. But it is "…clearly not the case…that a person making a promise reports that he intends to follow through…" (p. 26). If this was what happened and he did not have the intention to follow through then he would not really count as making a promise and so he would not be doing anything wrong when he said 'I promise to x' without the intention to x, which is a very odd result indeed!

The point Copp is making seems to be that in the promising case there is a pragmatic convention that connects promising to the intention to fulfill the promise but not in a way that makes the truth of 'I promise…' depend on having this intention. Applying this to the

Petit and Jackson argument Copp can say that though there is a convention of English which connects the assertion that x is wrong to the norm which justifies it, it is of the kind that governs promising. It is a pragmatic norm governing when it is appropriate to assert 'x is wrong' not a semantic convention governing the truth of the utterance. I think that this is a plausible way of answering this challenge.

Let us turn to discussing Bach's argument and see if Copp's view stands up to it. Bach's argument against conventional implicatures is actually very simple. He calls it the IQ test. It stems from the fact that we can report what is said in indirect quotation. So when I say 'so and so said blah blah blah' the 'blah blah blah' reports what the person actually said. Now when we say 'Alice said that Jane is poor but honest' the relevant contrast seems to be reported as well. I will have reported Alice as having said less than she did if I say that she said that the person in question is poor *and* honest. So that there is a relevant contrast is part of what is said and so is part of the truth conditions for the utterance (Bach 1999). This shows that this other proposition is not detachable in the way that is required to be a conventional implicature. That is, I cannot

express the same belief as I do when I say that Jane is poor but honest when I say that she is poor and honest.

What this means is that 'but' as used in 'she is poor but honest' does not mean the same thing as 'and' as used in 'she is poor and honest'. They should be translated into first-order logic differently.[20] Now consider this argument as applied to Copp's version of realist-expressivism. On his view moral words have coloring as discussed above and so when I utter the sentence 'suicide bombing civilians is wrong' my non-cognitive state of mind is implicated by the meaning of 'wrong' but it is not supposed to be part of the truth-conditions of the utterance that I do in fact have this state of mind. This utterance is supposed to pass the four tests for coloring that we discussed above. The most important of these is detachability.

To detach the implicature, you will recall, we need to replace the colored word with a non-colored word and still

---

[20] One gets the usual (Ex)(Px & Hx) where the other gets translated as ((Ex) (Px & Hx) & (My) (Py➔-Hy)) where 'M' has to be a 'most' operator (perhaps as a so-called 'fuzzy' operator, or as expressing a probability of greater than .5). So (2) is to be read as 'there is an x that is both poor and honest while most things that are poor are not honest'. If one does not like the introduction of a 'most' operator because it sounds like we have now expanded first-order logic to second-order logic, we can paraphrase 'most things that are poor are not honest' as 'only some poor things are honest'. This allows us to render it as (Ex) (Hx & Px) & (Ey)(-Hy & Py) & More(y, x). So when you report that I said someone was poor but honest you report something with this as a truth condition

be able to express the same proposition. So in Copp's example that we talked about previously, 'Alice is a Yankee' we could replace 'Yankee' with 'American' and still express the same proposition; where that means that the truth conditions are exactly the same. Copp argues that in the moral case we detach the color by placing the term in inverted commas. So, 'suicide bombing is "wicked"' is the de-colored version of 'suicide bombing is wicked' and so we can detach the implication in moral utterances.

But consider someone saying (1) 'Richard said that suicide bombing is wicked' versus someone saying (2) 'Richard said that suicide bombing is "wicked"'. Am I reported as saying the same thing in both cases? It seems that I am reported as saying more in (1) than I am in (2). This suggests that it is just part of the meaning of 'wrong' that the action or person in question violates some justifying reason. This means that Copp cannot claim that the non-cognitive state is detachable and so cannot claim that conventional implicature is the right model. In fact it looks like his view does collapse into the entailment view.[21]

---

[21] In chapter 2 I argued that no semantic theory entails any theory of the justification of moral judgments but here it looks like I am arguing that Copp's semantic theory entails subjectivism. But this is not quite right. What I argued was that no semantic theory need entail any theory of justification, so I think that Copp could fix his theory

On the other hand to construe it as a conversational implicature seems wrong as well. Usually what happens in the case of a conversational implicature is that we do not mean what our sentence means, and instead mean something else. So, if I say 'can I use the restroom?' I am not asking if I physically am able to use it, I am requesting permission to go and use it. But when this model is applied to ethical discourse we get the strange result that when I say that suicide bombing is morally wrong I do not mean what the sentence means (that it is morally wrong) but rather mean something else instead, like I disapprove of it. Now one may have a view like this, but how could a hearer ever figure out that I morally disapprove of suicide bombing if not by knowing that I said that it was wrong?

The upshot of all this is that it seems that if we want to endorse a 'many propositions' model then the best way to do it is to simply say that we express both propositions at the same time with one utterance. As Bach

---

so as to avoid this problem. He could, perhaps, insist that his account of conventional implicature was correct and so accept that the utterance will be true if I do have the mental state in question but be an error theorist in an opposite way than usual. That is, instead of thinking that the moral predicates (mistakenly) stand for properties that do not exist Copp could claim that they (mistakenly) do not stand for properties that do exist. One would have to argue for this, or something like it, or accept a relativistic account of justification. So, some semantic theories will exert a kind of pressure towards a certain theory of justification.

has pointed out, most of the resistance to this kind of view stems from the old Grammarians dictum that a sentence expresses a complete thought. That may be true for the sentence but there is no reason to think that when I utter that sentence I must be only expressing one thought.

Chapter 4. The Varieties of Semantics

4.1. Taking Stock

So where are we in the argument? We have established that issues of meaning and justification are separate. This follows from the more general truth that semantical issues are separate from metaphysical issues. No theory of the meaning of words alone commits us to any theory of the justification of moral judgments. But though we can mix and match these kinds of theories in any way we want, our choices are limited when we add in the constraint of a naturalistic moral realism. And though a semantic theory needn't commit one to any theory of justification when combined with other assumptions it can put pressure on you one way or another. Thus we saw that two contemporary kinds of expressivism currently on the market threaten to lead us into a radical kind of subjectivist account of justification.

It is not on account of Blackburn's theory of meaning that he is a relativist. It is because he holds a certain theory of the justification of moral judgments that allows for it to be the case that two people who disagree over some fundamental moral claim can both be, in some sense, right. Similarly it is not the semantic claim that commits

Copp to subjectivism. Rather it is that without further amendments which are themselves very unintuitive the view is a subjectivist view.

These views are unacceptable to me since they place unreasonable constraints on a naturalistic moral realism of the kind defined in the second chapter. To remind you, a theory of the justification of moral judgments is a morally realist view, I argued, when it predicts or dictates that when two people disagree about some fundamental moral claim, like for instance, the claim that killing newborn humans for sport is morally wrong, one of these people must be wrong. They cannot both be right about the way that they feel about this claim.

So, to take the most straight-forward example, the utilitarian theory of justification that says that an act is right in so far as it promotes the greatest good for the greatest number is clearly a realist theory of justification in my sense. Killing babies for sport cannot, without substantial back story, maximize happiness. If utilitarianism is the correct theory of justification and someone were to disagree with me about killing babies for sport only one of us could be right.

The kind of view that Blackstone introduced is also clearly a version of moral realism in this sense, or at

least it could be if that turned out to be the best theory of justification. Blackstone argued, recall, that we express our moral emotions, as Hutcheson and Hume had thought, and contextually implicate that we have a justifying reason for having the emotion. Blackstone went on to argue that this was very much like Hutcheson's own view with a utilitarian account of the justifying reasons.

As I have said, I want to re-introduce a theory much like the one that Blackstone had, but updated with some modern views about the philosophy of language. These views in the philosophy of language are not shared by everyone, but this doesn't matter. As I argued in chapter two, any theory of semantics is in principle compatible with any theory of the justification of moral judgments so in a sense it will not matter what one's views are in the philosophy of language. But we have seen some problems when we take views like Blackburn's and Copp's and use them to construct a metaethical theory.

This gives us some guidance in our search for a semantic theory of moral communication. In particular it has shown us that the meaning/use distinction is important to an adequate metaethical theory. A natural thing to think is that the meaning/use distinction captures the semantic/pragmatic distinction. On this reading semantic

theories are in the business of telling us what the meaning of words are and a pragmatic theory is in the business of telling us how people use those words to do things.

This way of dividing semantics from pragmatics would allow us to look back at the classical period of emotivism discussed in the second chapter and see that they were trying to articulate a theory of the use of moral language and not a theory of the meaning of moral words and so were involved in developing the pragmatics of moral communication. They all pretty much agreed that these words had no meaning in the technical sense of meaning endorsed at the time. As we have seen Stevenson's view is very amendable to interpreting as relying on a use/meaning distinction.

In this spirit the emotive realist agrees that on any given occasion, a person may use a sentence in an innumerable many ways. One of those many uses, a distinctively moral use, of moral language, which typically include words like 'good' 'right' and 'ought' but can be virtually anything in the right circumstances, is to express our moral sentiments and at the same time the belief that the sentiment is the correct one to have.

But the emotive realist is not committed to the claim that the sentence 'suicide bombing is morally wrong' *means* 'I morally condemn suicide bombing; that is the correct way to feel about suicide bombing' in just the same way that 'sure, I'll be there' doesn't mean that one has promised. This is what *someone who says the sentence* is expressing if they are being sincere, etc, and so is what the *person* means, not the *sentence*. A person may use that sentence in order to make a promise, but in that case it would not be because of the meaning of the sentence that he was successful. It depends on what the person is trying to do when they make the utterance. So at this point the next step would seem to be to spend some time talking about the pragmatic/semantic distinction and then formulating emotive realism in those terms.

But while the meaning/use and semantic/pragmatic distinctions seemingly mark real divides, determining precisely where to draw the line between semantics and pragmatics remains a hotly contested question. In what follows I will suggest a way of categorizing various theories that will let us avoid this vexed question while at the same time preserving the common sense insight we began with. This will take us deep into contemporary semantical theories but when we are finished we will

finally have the proper semantical tools to formulate an emotive realism, which we will do in chapter five.

4.2. The Problem with Semantics

One reason that I have been hesitant to get involved in semantics is that the issues tend to run on the technical side and can obscure the broader point that one is trying to make. This is because semantics is a field that, from the outside, seems to have lost its way. It appears to be overrun with various camps that have varying intuitions. So, for instance we have those with very strong intuitions about whether or not there is water on Twin Earth. Each side seems to think that their intuition is just 'obvious' and that the other side must surely be missing something important. This might even lead one to think that semantics is just some kind of 'sport philosophy' that isn't connected to anything that is important.

This is just the problem that the naturalistic/ empirical approach to philosophy is supposed to rectify. We are not to rely (solely) on our intuitions as a guide to constructing an empirical/naturalistic theory of moral communication. Every semantical theory that we have so far considered has in one way or another relied on intuitions about normal speakers and how they use language. So, what

we ought to do is to step back and re-examine the aims of semantics. What are the goals of a naturalistic science of semantics?

4.3. The Semantic Task: L-semantics and P-semantics

As we have seen, the traditional motivation for an emotive metaethical theory is that of a commitment to naturalism. Non-expressive ethical theories are thought to be committed to metaphysically questionable entities out there in the world. The expressivist denies this. Given that that project we are now involved in is that of finding a semantic theory of moral discourse capable of supporting a naturalistic moral realism we ought, following Devitt (Devitt 1996; Devitt 1997), to take this kind of 'metaphysics first' approach and apply it to the field of semantics.

If we are naturalists then we will surely want a semantic theory that is naturalistic. In this respect semantics is no different than physics. The way we should proceed in semantics, then, is as if from scratch. We need to identify what the semantic task is. Devitt begins this starting over by pointing out that

> the 'basic' semantic task is to say what meanings *are*, to explain their *natures*. It is thus analogous to such tasks as trying to figure out what genes, atoms, acids, echidnas, or pains *are*… (Devitt 1996 p. 54)

The problem is that we don't know what meanings are. We thus find ourselves "in the unusual position of having to specify a subject matter," (ibid).

What kinds of things do we ascribe meanings to? The answer is, again, obvious;

> We ascribe meanings to thoughts and utterances. So thoughts and utterances are the immediate phenomena of semantics (p. 57)

In particular we ascribe meanings, according to Devitt, to the *that*-clauses in attitude ascriptions. So were someone to say 'Richard believes that eating meat is immoral' it seems natural to take *eating meat is immoral* to ascribe a mental content, with a meaning, to me. In fact more is done. I am being described as holding a distinctive mental attitude towards this content; that of believing it.[22]

These ascriptions attribute things with meaning to people and these meanings play an explanatory role in our daily lives. Meanings then seem to be properties of thoughts that play a causal role in the explanation of

---

[22] The going hypothesis is that the mental attitudes can be functionally defined. So, typically what you get is people claiming that a belief that p is just having the content 'p' playing a certain functional role, to want p is to have the same content playing a different causal role. I have argued that this can't be right and that we should think of the mental attitudes as distinctively qualitative ways of feeling about p. To believe p is to feel that p is certain, etc. I mention this briefly in the next chapter but it is largely irrelevant to the argument I am making here. As long as you agree that a belief consists in a mental attitude held toward some content this argument will be fine.

behavior and which can be used a 'guide to reality'. For instance, here is how he characterizes the semantic task in the *précis* to Coming to our Senses: A Naturalistic Program for Semantic Localism.

> In Coming I seek a solution to this problem [i.e. identifying the semantic task] by focusing on the purposes for which we ascribe meanings (or contents) using `that' clauses ("t-clauses") in attitude ascriptions: in particular, the purposes of explaining intentional behavior and of using thoughts and utterances as guides to reality. I call these purposes "semantic."

The natural starting place in constructing a theory of meaning is in identifying the work that meanings are supposed to do for us. We identify a particular causal role. After we do this we can look to see what property, if any, fills that role. It may turn out that nothing plays that role and so that there were no meanings. But it may also turn out that we can find a naturalistically acceptable candidate for playing the role.

We can formulate this more precisely in terms of a semantic role. As Devitt says,

> …a property plays a "semantic" role if and only if it is a property of the sort specified by t-clauses, and, if it were the case that a token thought had the property, it would be in virtue of this fact that the token can explain the behavior of the thinker or be used as a guide to reality. We are then in the position to add the following explication to the statement of the basic task: A property is a meaning if and only if it plays a semantic role in that

sense. And the basic task is to explain the nature of meanings in that sense

We can see from this that Devitt takes meaning to be a property of thoughts and the semantic task is to explain what property they have which allows them to play the role in behavior that they do.

This is certainly an important task and just for the reason given by Devitt. We absolutely need a theory of what it is about thoughts that allows them to have meaning. Thoughts seem clearly to be the kinds of things that we ascribe meaning to and the meaning that we ascribe to them certainly seems to play an important role in explanations and predictions of behavior. This is true in general but, more to the point here, it is true for moral communication. An adequate metaethical theory has got to give an account of how our moral judgments come to have meaning. This is also, it would seem, something that we should legitimately call semantics.

But this is not the first thing that may come to mind when one thinks about the 'basic' semantic task. It may seem intuitively obvious that it is *sentences* which we ascribe meanings to, not *thoughts*. This is the approach that Kent Bach has taken in his work and he sees this conception of the semantic task as clearing up a lot of

confusions in the philosophy of language. He has, from this alternate conception of the semantic task, launched an attack on the Devittian kind of semantic task we have been talking about. He has also attacked Kripke's view, which shares the spirit of Devitt's semantic theory. In the next section we will look back at what Kripke has said about semantics with an eye to showing the similarities between his conception of the semantic task and Devitt's. We will then look at Bach's criticisms of this kind of view. What will emerge is that both are on to something important which deserves the name 'semantic task'.

4.3.1. What Kripke Really Thinks

The 'intuitive test' of rigidity that Kripke (Kripke 1972) gives is a metaphysical one. He says "although someone other than the U.S. President in 1970 could have been the U.S. president in 1970 (e.g. Humphrey could have), no one other than Nixon might have been Nixon," (p 48). This may at first seem puzzling, I mean, it seems true, but why does Kripke think that this is a semantic issue? Semantics, according to him is in the business of giving truth conditions but truth conditions for what? How do we get them? The answer turns out to be *utterances that express singular thoughts*.

Thus our intuitions about the truth of sentences like 5 and 6 are really intuitions about the thoughts that these sentences are taken to be used to express.

> 5. The first great analytic philosopher might not have been the first great analytic philosopher

> 6. Bertrand Russell might not have been Bertrand Russell

That is, our intuitions are about particular *uses* of sentences and Kripke is talking about token utterances. Since we are taking these sentences to be expressing certain thoughts we assume that the speaker must have someone in mind. It is because Kripke takes 'Russell might not have been Russell' to be expressing the thought that Russell, the actual guy (pointing at Russell), might not have been Russell (again pointing at the same guy) that he thinks the truthmake for the sentence is just one person.

This becomes much clearer when we look at Kripke's later article, "Speaker Reference and Semantic Reference," ((Kripke 1977/1991) henceforth SR) where he is much more explicit about what the jobs of semantics and pragmatics are. He begins, following Grice, by distinguishing between speaker and semantic *meaning*.

> The notion of what words can mean in a language is semantical: it is given by the conventions of our language. What they mean on a given occasion is determined, on a given occasion, by these conventions, together with the intentions of the

> speaker and various contextual features. Finally what the speaker meant, on a given occasion, in saying certain words, derives from various further special intentions, taken together with various general principles applicable to all human languages regardless of their special conventions (Cf. Grice's maxims) (SR p 84)

So, in effect we have three kinds of meaning. We have the general meaning of the words as given by the conventions of the language (what Grice called *timeless* meaning). Whatever we say about these conventions we can say that the dictionary is the canonical listing of them. Next we have the meaning of the words on a given occasion which is determined by the conventions of the language and the intentions of the speaker and 'various contextual features'. This allows us to resolve any ambiguities, figure out what any indexical or demonstrative elements are being used for and in general determine what proposition the sentence expresses.[23] Finally we have what the speaker meant on a given occasion in saying what they did, which may diverge from what the sentence that they said means on that occasion.

So far this all looks like the standard Gricean picture. Kripke then goes on to distinguish the semantic referent of a term from the speaker's referent of the

---

[23] If any (Bach forthcoming)

term, which he says are special cases of semantic and speaker meaning respectively.

> …the semantic referent of a designator is given by a general intention of the speaker to refer to a certain object whenever the designator is used. The speaker's referent is given by a specific intention, on a given occasion, to refer to a certain object (p 84)

The semantic referent is given by a general intention which means that I have a general intention to refer to Kripke by using 'Kripke'. This means that I have something like a standing intention; I generally use tokens of 'Kripke' as a way of referring to Kripke. As he says,

> …he uses 'Jones' as a name of Jones-elaborate this according to your favorite theory of proper names- and, on this occasion, simply wishes to use 'Jones' to refer to Jones," (ibid).

The reason that I am able to do this is elaborated by his favorite theory of names, which is a causal theory. So, who it is that I have in mind is determined by the causal link that the name has to an individual. This determines who the thought is about.

The picture he is developing seems to be like this. When things go right, the speaker has someone in mind, that is, has a singular thought about a particular person, has a word that he uses as a name for that person (this person is the semantic referent of the term) and on this

occasion he intends to refer to that person by using the name (this is the speaker reference). But what is a general intention? Consider an analogy. I generally use hammers to hammer nails. In fact I generally intend to use hammers that way, that is why I have a hammer in the house. On a given occasion that I need to use a hammer, I have an intention to hammer *this* nail, and on this occasion I have a specific intention to use the hammer that is an instance of my general intention to use hammers in this way. In the case of names what this means is that I generally intend to conform to the convention of using tokens of 'Saul Kripke' to name its bearer: the philosopher Saul Kripke. That is why I have the name 'Saul Kripke' as part of my vocabulary, in just the same way that my keeping a hammer around the house is because I generally intend to use it to hammer nails.[24]

Now let us return to the question at hand. When we evaluate sentences like 5 and 6 what is it that we are evaluating? I have suggested that Kripke thinks we are evaluating sentences that are taken to express singular thoughts. That is, thoughts about a particular individual. It is in this way that we can understand why Kripke thinks

---

[24] This of course opens the door to the possibility that these two intentions may come apart, we shall return to this point shortly.

that the test for rigidity is a metaphysical one. This emerges clearly when Kripke says,

> In practice it is usual to suppose that what is meant in a particular use of a sentence is understood from the context. In the present instance, that context made it clear that it was the conventional use of 'Aristotle' for the great philosopher that was in question. Then, *given* this fixed understanding of [Aristotle was fond of dogs], the question of rigidity is this: Is the correctness of [Aristotle was fond of dogs], *thus understood,* determined with respect to each counter-factual situation by whether a certain single person would have liked dogs (had that situation obtained)? …this question is entirely unaffected by the presence or absence in the language of other readings [of the sentence]. For each such particular reading separately, we can ask whether what is expressed would be true of a counter-factual situation if and only if some fixed individual has the appropriate property. (NN P.9)

Kripke means to be talking about sentences as used on a particular occasion. What he is concerned with is particular uses of sentences, and their meaning and reference are determined by who the speaker has in mind and his intentions.

So the big picture is that we express a singular thought about something and we want to know the truth conditions for it (the thought). The way we determine the truth conditions of the thought is by looking at the token utterance that we take as expressing it and as we have seen this is done partly by the speaker's intentions (since 'Aristotle' names more than one object), and partly

due to the context of utterance. This is called semantics because we are talking about truth-conditions, but it is important to see that on this view it is token sentences that get evaluated and that these token sentences are taken as expressing singular thoughts.

Now as I have said, I think that Devitt's approach here is correct as far as it goes. It does identify something that we can legitimately call a semantic task and it provides a very nice account of a naturalistically acceptable property to fulfill that task, but Devitt's method also identifies another sort of thing that we might naturally identify as a semantic task. For, as we have seen, on the Kripke/Devitt kind of view it is thoughts that are primarily the bearers of meaning. What are we to say about sentence types? Don't they have meaning? When we think about semantics we are apt to think of sentences and *their* meanings. One might then take the semantic task to be that of giving the meaning of sentences independently of their being used to express any thought.

This way of thinking about semantics has it as simply a part of grammar. To illustrate, if I say 'Saul Kripke likes tea' talking about my dog and you say it talking about Saul Kripke we both use the same sentence, though we refer to different objects. We do so in the sense that we

use something with the same physical structure but we also use something with a certain syntactic structure, something that has a noun phrase and a verb phrase as part of its structure like (7)

(7) [S [NP [proper noun, Saul Kripke]], [VP [verb, likes], [np, tea]]]

This is roughly Kent Bach's position. According to Bach the job of semantics is to provide an interpretation of (7) that explains how it can be used by English speakers to do the things that people do with it.

Here are a couple of quotes from his 1999 paper "The Semantic Pragmatic Distinction: What it is and Why it Matters" and his 2002 paper "Semantic, Pragmatic"

> I take the semantics of a sentence to be a projection of its syntax. That is, semantic structure is interpreted syntactic structure. Contents of sentences are determined compositionally; they are a function of the contents of the sentence's constituents and their syntactic relations. (Bach 2002)

> Semantic information about sentences is part of sentence grammar, and it includes information about expressions whose meanings are relevant to use rather than to truth conditions. Linguistically encoded information can pertain to how the present utterance relates to the previous, to the topic of the present utterance, or to what the speaker is doing. That there are these sorts of linguistically encoded information shows that the business of sentence semantics cannot be confined to giving the proposition it expresses. Sentences can do more than express propositions. Also, as we have seen, there are sentences which do less than express

> propositions, because they are semantically
> incomplete. (Bach 1999)

Semantic information about sentences, for Bach, is part of sentence grammar.

So here we have another way of characterizing the semantic task, and this one seems legitimate as well. There has been a lot of debate between Bach and Devitt over which is the best way to formulate the semantic task. But this debate seems like a waste of time to me. A lot of energy has gone into arguing for various ways to draw the semantics/pragmatics distinction with the end result being much spilled ink and no consensus. It seems to me that we have two different and legitimate conceptions of what the semantic task is.

Each of the goals identified above seem to me to be legitimate candidates for serious theorizing. Bach's conception of semantic information is 'linguistically encoded information' whereas Devitt's is 'properties of thoughts that explain behavior'. To avoid this debate altogether I will use 'P-semantics' for semantics in the psychological sense that we want to give a theory of the meaning of thoughts and 'L-semantics' for semantics in the linguistic sense that we want to give a theory of the meaning of sentences considered apart from their being

used to express any given thought. What is the relation between these two kinds of theories? There are three possibilities that present themselves.

The first is that L-semantics just is P-semantics which is to say that the semantics of English just is the semantics of thought. Broadly speaking this is the conception of semantics that P. F. Strawson (Strawson 1950/1985) had, i.e. the meaning of a word was given by instructions on how to use it. It is still popular, as for instance Jerry Fodor (Fodor 1998) when he says,

> …English has no semantics. Learning English…[is] learning how to associate its sentences with the corresponding thoughts" (p. 9)

Or one might think that P-semantics just is L-semantics, which is to say that the semantics of thought just is the semantics of English. I suppose that people who think this would be people like Sellers (Sellars 1956) who think that we start with sentences and then work back to thoughts, which are theoretical posits to explain verbal behavior.[25]

Finally, one might want to give separate accounts for each. It might turn out that we need distinct semantic theories for thought and language. To take an example which involves some foreshadowing, if our P-semantic

---

[25] Devitt does cite Sellers' method as an inspiration in Coming to Our Senses

theory needs to appeal to the syntactic category 'name' in a theory of thought content/intentionality it may turn out that the semantic interpretation we give that syntactic item in thought will differ from the semantic interpretation of that syntactic category in an English sentence type.

This distinction between P-semantics and L-semantics arises from trying to give a theoretical account of the nature of speaker meaning and linguistic meaning. And it has the advantage of letting us avoid all of the debates about pragmatics and semantics. Given this distinction it is clear that both kinds of semantics will be interested in sentences and truth-conditions because when we want to know the truth-conditions of a sentence there is a P-semantic/L-semantic ambiguity that often gets overlooked.

We could be taking this sentence to represent an utterance, an actual saying of it or a writing of it, and so as an expression of thought. The truth conditions for the sentence taken this way are really truth conditions for the thought it is being used to express. On the other hand we could take the sentence as a linguistic type and be trying to evaluate its truth conditions independently of any thought it may be used to express. What would a speaker of English need to know in order to use the

sentence correctly? Accordingly the distinction between meaning and use that I have relied on, in an intuitive way, throughout this dissertation turns out to be the distinction between L-semantics and P-semantics. What I have been calling meaning is the subject of our L-semantic theory, and what I have been calling use is the subject of our P-semantic theory. That is, in so far as by use we mean the sentence as it is used on a particular occasion what we are really evaluating is the thought that is being expressed.

4.4. The Independence of L-semantics & P-semantics

As we have just seen, a large part of the argument of this chapter relies on my claim that we need to give distinct L-semantic and P-semantic theories. I want to claim that the causal theory of reference of the kind that Michael Devitt defends (Devitt 1974; Devitt 1981a; Devitt and Sterelny 1999) is a very good theory of the content of mental states but that it can't be a L-semantic theory. In this section I will present an argument which shows that we need to give separate theories of these things.

Let us start with Kripke's notion of rigid designation. As I argued in the last section a name in thought has its reference determined by a causal/historical link to the

person/object it names.[26] Now if L-semantics were reducible to P-semantics then we should be able to model names as linguistic singular terms. The way that this claim is typically expressed is by modeling names in natural languages as individual constants in logical notation. There are many notorious difficulties for this model of names. These are the familiar ones about informative identities and the opacity of belief attributions, to name just a couple of the more familiar. Various moves have been made to get around these standard objections. I do not want to rehearse these moves here. Rather what I want to do is to add a new problem to the list. One which I think forces us to realize that L-semantics and P-semantics are best seen as separate.

The problem is that when we model names in natural language as rigid designators, like we did in our P-semantic theory, we are able to prove that every object must necessarily exist in a modal logic as strong as S5. What this shows is that our L-semantic theory cannot contain individual constants. This shows us that we do not want an L-semantic theory that treats moral predicates as

---

[26] A 'name in thought' is simply whatever it is that allows me to think about individuals; that is, to have singular thoughts about them. It may, as Fodor thinks, be syntactically akin to a name in language or it may, as Devitt thinks, simply 'respect' the syntax of natural language.

standing for some property, for that is to treat them as individual constants. It also shows us that we need to separate these two theories. Let us look at the problem in some more detail.

It is well known that we can prove that any object must necessarily exist in a modal logic as strong as S5 (Prior 1956). I have a standard proof of this as (8).

> (8) Proof that every object necessarily exists in S5
>   (from (Menzel 2005))
>
>   1 x=x
>   -axiom of identity
>
>   2 (y) (y≠x) → (x≠x)
>   -instance of quantifier axiom
>
>   3 (x=x) → ~(y) (y≠x)
>   -from 2 by contraposition
>
>   4 (x=x) → Ey (y=x)
>   -from 3 quantifier exchange
>
>   5 Ey (y=x)
>   -from 1 & 4 by Modus Ponens
>
>   6 □Ey (y=x)
>   -from 5 by rule of necessitation
>
>   7 (x)□Ey (y=x)
>   -from 6 by universal generalization

As you can see it follows directly from the axiom of identity and an instance of the quantifier axiom which is a statement of universal instantiation. So, briefly in words, premise 2 says that if no object is identical to x then x (being an object) isn't identical to x. This is

equivalent to saying that if x is self-identical then x exists, and it follows from this and the axiom of identity that x does exist. Since x's existence is a theorem of S5 we can say that it's necessary and universal.

Since S5 is the modal logic that most people agree is strong enough to actually be of use as a logic of necessity and possibility this may seem quite jarring. Necessary existence is usually reserved for such lofty beings as numbers and God; who would've thought that my computer and I kept such company?! So what are we to do? There are those who recommend that we have an open mind about this and entertain that everything might really necessarily exist (Williamson 2002), but I do not find myself able to have such an open mind. Rather, the existence of this kind of proof seems to me evidence that something has gone wrong with the formulation of S5.[27]

Luckily for those who feel like I do there is Kripke's well known solution to this problem (Kripke 1963).[28] Following Quine, he invokes the 'generality interpretation' of variables and requires that no free

---

[27] Williamson defends an argument that is similar to the proof I have given though it is couched in terms of models instead of proofs. I will not discuss his defense of it here but I do think it can be answered.

[28] Kripke is concerned with the Barcan and Converse Barcan formulae and not explicitly with the problem of necessary existence, but his strategy works for necessary existence just as well.

variables be allowed in our instances of axioms in a proof

(Quine 1940). So, (8) comes out invalid and premises 1 and

2 need to be reformulated as 1' and 2',

    1' (x) (x=x)

    2' (x) ((y) (y≠x) → (x≠x))

which are said to be 'closures' of the standard axioms,

and we can't get a proof of necessary existence from them.

All we can prove is that all objects are necessarily self-

identical, which is harmless. Here is the proof

    (8') Proof that every object is necessarily self
        identical in S5 (again from (Menzel 2005))

    1' (x) (x=x)

    2' (x) ((y) -(y=x) -> -(x=x)

    3' (x) ((x=x) -> Ey (y=x)
        -From 2' contraposition and quantifier exchange

    4' (x) (x=x) -> (x)Ey (y=x)
        -From 3' by quantifier distribution rule

    5' (x)Ey (y=x)
        -From 1' & 4' by modus ponens

    6' □(x)Ey (y=x)
        -From 5' by rule of necessitation

However, all is not yet well, for we can still construct

a proof of necessary existence for any given object that

we want, including myself, my computer, or unicorns, by

using singular terms, or 'non-logical constants' instead

of variables. As for example, in (9), which shows that we

can derive a contradiction from the assumption that it is
possible that Kripke doesn't exist, and so by reductio,
that he must.

    (9) Proof by reductio that Saul Kripke necessarily
        exists: □Ex (x=SK) (adaptation of a first-order
        proof from David Rosenthal)[29]

      1. ◊ ~Ex (x=SK)
         -assumption

      2. ◊(x) (x≠SK)
         -equivalent to 1.

      3. (x)□ (x=x)
         -modal axiom of identity

      4. □ (SK=SK)
         -Universal Instantiation (UI) of 3.

      5. ◊ (SK≠SK)
         -UI of 2.

      6. ~□ (SK=SK)
         -equivalent to 5.

      7. □ (SK=SK) & ~□ (SK=SK)
         -from 4. & 6. Conjunction Introduction

      8. □Ex (x=SK)
         -from 1. - 7. via Reductio

This proof does not appeal to any axioms that have
occurrences of free variables and so respects the
generality interpretation. Now one may think that the move
from 3. to 4. is suspect because 1. denies that SK exists.
But this objection is misguided. UI just says that I can

---

[29] Rosenthal's first-order version was offered, in a course on Quine
and Sellers I took in Spring 2006, in support of Quine's views about
regimentation. I have adapted it into S5 simply for shock value.

replace a universally bound variable with any constant and SK is a constant; it appears in the first premise!

However, even if one was convinced that this objection was correct we could reformulate (9) as a version of (8) with constants instead of variables. As in (9′).

(9′) Adaptation of (1) with constants

1′ (x)□ (x=x)
    -instance of closed axiom

2′ □ (SK=SK)
    -UI of 1′

3′ □ ((y) (y≠SK) → (SK≠SK))
    -Closed quantifier axiom

4′ □ ((SK=SK) → ~(y) (y≠SK)
    -contraposition

5′ □ (SK=SK) → □Ey (y=SK)
    -quantifier exchange, & distribution

6′ □Ey (y=SK)
    -from 2′ & 5′ by Modus Ponens

So if we are to avoid the implication that I or my computer necessarily exist we need more than just the generality interpretation. The obvious source of the problem is that SK is a singular term so Kripke's next move is to require that there be no singular terms in our formal language. His quantified modal logic includes only variables.

So how then do we say that it is possible that Kripke might not have existed? One option is to adopt Quine's suggestion that we use Russell's theory of descriptions so that when we analyze sentences like 'Saul Kripke exists' we get a logical statement free of singular terms. He, of course, recommended that we invent a description like 'the thing that Kripkisizes', or 'the Kripkisizer' so that we would render 'Kripke exists' as (Ex) K(x) where 'K' stands for the invented description. This was meant to be a purely technical device to solve the technical problems associated with existence statements about non-existent things. In particular one gets the feeling that it is only to be used when one knows that the thing in question doesn't exist but in principle this kind of device could be used to replace all names in a language. So, this strategy presents itself as an obvious way to avoid the embarrassment of arguments like (9) and (9'). But, as Quine makes clear, we need only revert to this strategy if we are unable to find a suitable description to translate the name. The –isizes device is available as a last resort, but it would be nice to have a more principled description.

One view that has got what I consider to be a bad rap is the view that Kent Bach calls Nominal Description

Theory (NDT) (Bach 1987; Bach 2002). NDT says that a name N is semantically equivalent to the description that mentions N, something like 'is called "N"' or, as Bach prefers, 'is the bearer of "N"'. So for instance on this view the sentence 'Saul Kripke likes tea,' would be rendered as (10),

(10) The English sentence 'Saul Kripke likes tea' is true if and only if

(Ex) (is called "Saul Kripke"(x) & Likes tea(x))

Its truth condition is that there is an object which is called 'Saul Kripke' and which likes tea. This view is often dismissed out of hand because many people think that it is addressed by Kripke's remarks in <u>Naming and Necessity</u> about circular theories of reference. But this is not the case because NDT is a theory about the meaning of names not about their reference. It is supposed to tell us what a competent speaker needs to know in order to use the word correctly, it is decidedly not a theory about the reference of names, or how their reference is determined. So NDT presents itself as a way of following through with Quine and Kripke's urge to free our formal language from singular terms and so from worries about necessary existence.

4.4.1. Introducing Frigidity

But how do we reconcile this with the thesis of rigid designation? Indeed, if rigidity is supposed to be a semantic property then one may think that adopting NDT as a way to get around (9) and (9′) is really admitting that there is no such semantic property. Who the description in (10) picks out can vary from counter-factual situation to counter-factual situation. The reference of the name does not depend on anything semantic. Rather it depends on pragmatic facts about how the name is *used* to express a singular thought.

This is something like the view that Bach arrives at. He argues that intuitions about the rigidity of names are merely a 'pragmatic illusion'. He says,

> If someone said, 'Ronald Regan might have been the older brother of the fortieth US president,' he could be referring either to Ronald Regan with respect to some counter-factual situation or to whoever bears that name in some counter-factual situation. He could mean that there is a possible situation in which the (actual) Ronald Regan is the older brother of the president or that there is a possible situation in which the older brother of the president bears the name 'Ronald Regan'. (TR p 153)

I have coined the term 'frigidity' and 'frigid designator' as a way to contrast a view like Bach's with the standard semantic conception of rigidity. Bach would characterize the contrast between rigidity and frigidity as a contrast between a semantic conception of reference and a pragmatic

conception of reference. But as already said we are looking for a neutral way to formulate these issues. I do not want to get embroiled in disputes about the boundaries between semantics and pragmatics.

Making the dispute about the 'right' way to define the semantic task has the effect of hiding the fact that the parties agree more than they disagree. So, take Bach's earlier objection. He takes himself to be objecting to Kripke, but all he is saying is that there are different readings of this sentence in which we take the speaker to be expressing different thoughts and that we cannot tell from the sentence type alone which one it is. But this is not something that Devitt would object to! Devitt and Kripke are in agreement that this is perfectly acceptable but the point in question is that given one such thought is expressed it's truth depends on just one person and of course Bach agrees with this.

This comes out again when Bach says that a sentence like 'Betrand Russell might not have been Betrand Russell' can be used to say that Russell might not have been called 'Russell' (Bach 2002). In this case the second occurrence of 'Russell' is as a predicate. Kripke is aware of this and agrees that in colloquial English 'Russell might not have been Russell' can be used to mean that 'Russell might

not have been called 'Russell'' (NN p. 62 fn). But Kripke objects that this is not what he is after, and that it is a mistake to think that this is a response to his point. Because what he means is that the person, Russell could not have been some other object. Not merely that this object, Russell, could have borne a different name.

This is because, as we have seen, Kripke takes it as given that we mean to be talking about Russell and so the thought we are evaluating is one to the effect that Russell might not have been self-identical. Bach can't believe that this is what Kripke and Devitt are arguing for. He says,

> With rigidity so defined, [the semantic thesis that names are rigid designators (RDT)] no longer even appears significant semantically. It is true all right, indeed trivially true, only because [they are] individuating uses of a name by each thing the name is used to refer to. By this method, a name has as many uses as it has bearers. No wonder that a name, as used in a particular way, rigidly designates! …it renders RDT so trivial as to apply to descriptions. For example 'the first man on the moon' would rigidly designate Neil Armstrong, since if Armsrtong had not been the first man on the moon and Buzz Aldrin had been instead, the use of 'the first man on the moon' to refer to Aldrin would count as a different use. Of course no one would dream of individuating uses of descriptions in this way... (TR 154)

But this is exactly what they have been arguing![30] Bach thinks that this is crazy because he assumes that the task of semantics is the L-semantic task, while Devitt and Kripke seem clearly to think that the task is the P-semantic task.

I can now neutrally formulate the distinction between rigidity and frigidity by saying that frigidity is the claim that there is no such L-semantic property of rigidity. There is no grammatical or syntactic category comprising rigid designators at the level of linguistic meaning. There are no singular terms in English qua English. When we construct a *linguistic theory* of the semantics of natural languages (as opposed to a psychological theory of thoughts) we should do it so that it is free of singular terms. So what I think (8) and (8′) show is that our L-semantic theory cannot contain rigid designators. When we treat names as singular terms our best logic goes off the rails.

But then how do we evaluate these singular thoughts if there are no singular terms? This is easy to answer if we take the causal theory of reference as a P-semantic theory. It says that we can have *singular thoughts*, given

---

[30] And Devitt does dare to dream that definite descriptions are ambiguous as between a referential and an attributive meaning and so that some definite descriptions would be rigid designators.

that the right kinds of causal/historical connections hold between certain thought contents and the world. But we express those thoughts using a language that itself does not have singular terms as per NDT. We can then formulate a description that singles out Kripke without any singular terms. We can say that there exists an object that is called 'Kripke' and which I am thinking about now. We then explain 'thinking about' in terms of a thought (or a proper part of the thought if one likes) having the right kind of causal relation to x. This is symbolized as (11),

(11) (E!x) (K(x) & (Ey) (is an occurent Thought(y) & is Causally related to(x, y)))

where '(E!x)', pronounced 'E-shriek x,' is shorthand for 'there is a unique object x such that'. So (11) says 'there is a unique object that is both called "Kripke" and is causally/historically related to my occurent thought in the right way'.[31] Thus, to say that it is possible that Saul Kripke does not exist is to say that it is possible that the object picked out by (11) in the actual world is absent in some possible world(s).

So then we can see that the thought that Saul Kripke likes tea and the English sentence 'Saul Kripke likes tea' will have different truth conditions. The truth conditions

---

[31] See (Bach 1987) especially pps 17-25 for a nice account of the various relations that will work for C

for the English sentence are just the conjunctive ones from (10), that there is an object called 'Saul Kripke' which likes tea. So 'Saul Kripke', the English proper noun, is not an L-semantic rigid designator. It is a description that can be used to refer to many different things. However, the thought that I have about Saul Kripke, the actual guy, will have the truth conditions specified by (11) but with 'likes tea' included as in (12),

> (12) My thought that Saul Kripke likes tea is true if and only if (E!x) ((K(x) & (Ey) (T(y) & C(x, y))) & L(x))

(12) is a *de dicto* rigid designator. It picks out Saul Kripke in this world since I am causally related to Kripke in the right kind of way required by C.[32] We then 'freeze' him as the object of interest (via stipulation). So linguistic names are frigid designators, their reference is determined by the thought that they are used to express. They themselves do not refer to anything. But because of this feature, namely that their meaning is simply that there is an object that bears the name in question, they can be used to express our singular thoughts. I am successful in communicating my thought if I

---

[32] If anyone cares 'E!x' is really just short for Ex (P(x) & (y) (P(y) → (y=x)))

get you to have a thought that also has a causal relation to the object picked out by (12).[33] But you understand the sentence if you understand (10).[34]

Of course, all of this is nothing more than an inconvenience unless we encounter circumstances where it becomes important. In normal circumstances we can just stipulate that SK is short hand for (12) where T is my thought about Kripke right now. In this respect the idea that there are singular terms in logic is akin to Newtonian Mechanics. Newton's equations work well enough for us to ignore the fact that they are technically incorrect. So unless we are dealing with things that are traveling near the speed of light or trying to describe the interactions of the very small they are all that we will ever need. So there is no reason to change the way that we teach first-order logic. We can continue to represent Saul Kripke with SK and sentences like 'Saul Kripke likes tea' and 'Saul Kripke exists' as L(SK) and (Ex) (x=SK). When the student progresses to the appropriate level of sophistication it then becomes necessary to rid the language of singular terms, and L(SK)

---

[33] Sometimes it matters if it has the same causal relation, sometimes it doesn't. Which is to say sometimes it matters if you think of the object in the same way as I do and sometimes it doesn't

[34] It may be that context and mutually shared beliefs can get (5) without having a singular thought.

becomes (12) and (Ex) (x=SK) becomes (11), both of which are horribly more complicated, but that's life. One might think of it along the lines of Field's (Field 1980) strategy in <u>Science without Numbers</u>. Kripke shows that we can axiomatize modal logic without constants but then we go on using constants anyway because they are useful. I do not mean to endorse Field's program, I merely offer this as a way to see why the claim made here is merely an inconvenience.

Now that we have these two kinds of semantic theories and have seen the need to distinguish them and give separate theories of each, let us now apply this to metaethics.

Chapter 5. (Re)Introducing an Emotive Moral Realism

I can now turn to staking out the metaethical theory I call emotive realism. I call this kind of theory an 'emotive realism' for two reasons. The first is to emphasize the claim that when we use moral language we are emoting. Thus I distinguish these kinds of views form those, like David Copp's Realist-Expressivism, which hold that we are asserting that some object in the world has some property and implicating, *a la* Grice, that we have a non-cognitive mental state that he calls, following Gibbard, *norm-acceptance*. Since he holds that the primary speech act that we perform is asserting, with the non-cognitive mental state being implicated, I call his kind of view an assertive-expressivism, and mine an emotive-expressivism.

Secondly, I call these kinds of theories 'emotive *realism*' (as opposed to Blackstone's 'objective emotivism') to emphasize that this view is, or can be, entirely, and whole-heartedly moral realist *in whatever sense of that word one favors*. So let me state the view. In a nutshell the claim is that when someone sincerely and

literally **says,** for instance, "eating meat is morally wrong," this person has

(a) **expressed** a **moral emotion** about eating meat and, at the same time, with the same utterance

(b) expressed the *belief* that the moral emotion in (a) is the correct one to have.

To **intentionally** express these two attitudes is to perform the speech act of **moral condemnation** or **moral approbation.**[35] The belief expressed in (b) has a truth value that is determined by the correct theory of justification, whichever one that happens to be. It might be the case that some moral emotion is the correct way to feel because of some utilitarian reason, a Kantian, a constructivist one, or whatever. So this emotive-expressivism can work with any particular version of moral realism that one likes.

I have been using this phrase to name the specific version of the theory that I hold because as far as I know I am the only adherent of the view. But as I have said, I think that there really are a bunch of different possible theories that could be called an emotive realism. All of them will accept that someone who says that something is wrong has expressed (a) and (b) but there will be a

---

[35] I do not think that these are the only moral speech acts that there are. These are just two common ones.

distinct kind of theory corresponding to each way of filling in the details. In what follows I will develop one particular version of emotive realism. This version is one that I think best captures what we do when we discuss ethics and at the same time avoids the problems that beset Blackburn and Copp.

5.1. The Mental Attitudes and Their Expression

A P-semantic theory will tell us two things. It will tell us how it is that our thoughts have meaning by specifying the properties that thoughts have in virtue of which they are able to play the role that they do in our mental lives. It will also tell us how we express those thoughts in communication. The mental attitudes, sometimes called propositional attitudes, consist in our taking some mental attitude towards some mental sentence. These are the things that we express. So what we need to do in the rest of this section is to give an account of the nature of the propositional attitudes, including both the emotions and beliefs, and then to give an account of how we express those mental states in speech.

5.1.1. The Moral Sentiments as Propositional Attitudes

As mentioned in the last chapter, I take the propositional attitudes to consist of mental attitudes

held towards represented propositions (Fodor 1991). Each mental state is therefore individuated by both its mental attitude and the represented proposition that the attitude is held towards. So the belief that George will arrive on time is differentiated from the fear that he will arrive on time in virtue of the differing mental attitudes held towards 'George will arrive on time' however the belief that George will arrive on time is distinguished from the belief that George will be late in virtue of the differing content that these mental attitudes are held towards. Included in this category will be things like belief and desire, but also emotions like anger and fear.

If the above account is correct then we can individuate the emotions in terms of their experienced quality, that is, in terms of their feeling. This is surely the common sense view of what they are, and the way they have usually been construed in the sentimentalist tradition. They differ among themselves in terms of experienced quality as well as what that experienced quality is directed at or held towards. So being angry about some violation of a moral code is a way of feeling about that violation. This way of feeling about the violation differs from the way that it feels to have some other mental attitude, like disgust, towards the

violation. The fact that the emotions occur unconsciously is no threat to this way of dealing with the emotions in just the same way that the fact that sometimes no is looking at a rock does not effect the existence of the rock.

Goldman (Goldman 1993) offers us a nice intuition pump. Imagine a Mary-like thought experiment with a super scientist we'll call Gary. Instead of never seeing color Gary has never had a certain propositional attitude, say disappointment.[36] Now imagine that he suddenly is disappointed that he will not be let out of his room. Won't he have learned something new? Namely won't he now know what it is like for him to be disappointed? It seems to me that this suggests that there is a qualitative aspect to these kinds of mental attitudes. I think we can extend this common sense idea to beliefs and desires. I have elsewhere argued that each propositional attitude consists in a distinctive qualitative mental attitude held towards some represented intentional content (Brown 2007).

The point I want to make here is that the moral emotions work the same way. They are mental attitudes that we can have towards various things. So, 'I ought to go

---

[36] There are perhaps better examples, like perhaps love, but this is Goldman's

pick up Fred from the airport' differs from 'I want to go pick up Fred from the airport' in the way that I feel about picking up Fred from the airport in just the same way that 'I am afraid of going to the airport to pick up Fred' differs from 'I am angry about going to the airport to pick up Fred'. On the other hand we have the attitudes of moral condemnation and moral approbation. These are ways of being for or against things in a particularly moral way. So to morally condemn is to have a strong feeling of disapproval towards something. The moral sentiments are particular (moral) ways of feeling about something and include moral condemnation, approbation, and obligation.

Moral approbation and condemnation are familiar from the sentimentalist tradition and the feeling of obligation is familiar from the rationalist tradition. This is the feeling that inspired Kant to compare it to the starry heavens above. The claim I want to make is that this feeling is a propositional attitude, like belief or anger, which consists in a distinct mental attitude held towards some represented content. There is an interesting argument that 'ought' is a mental state verb. Zeno Vendler (Vendler 1972) shows that there is a grammatical criterion that allows us to distinguish performative verbs from other

propositional verbs. The mark has to do with what he calls their 'time-schema'. A sentence with a performative verb in it, like 'promise' or 'warn', indicates that the speech act is occurring at that very moment; this is why the performative is the first-person singular present tense of the verb. The past tense ('I promised…') is not performative. So when I say 'I promise to pay you back' the sentence indicates that the promising is happening right at that moment.

The mental verbs that correspond to the propositional attitudes, 'believe', 'hope', 'fear', 'expect', etc, however, indicate a state that is had over time, not something that is happening at the moment. So, I can say 'I still believe that p' but I cannot say 'I still promise that p' (without some kind of explanation, that is. I can, of course, say this thing, but it is not standard in the way that the other construction is). The same is true of 'ought'. I can say 'I still ought to pick Fred up from the airport'. This makes 'ought' look like 'believe' in picking out a state as opposed to an act. Just as believing is a state that that a person can be in, corresponding to having a certain mental attitude held towards some representational content, so feeling

obligated is also a mental state held towards some representational content.

So it looks like we should postulate a feeling of obligation. That is, we should include obligation among the moral emotions of approbation and condemnation. This actually makes sense if one looks at the entomology of 'ought'. One of its many past uses was as the past tense of 'owe' as in 'I ought him twenty dollars, but I paid him off last week'.

As I have said the mental attitudes consist of two parts; a mental attitude held towards some representational content. The causal theory of reference as developed by Devitt gives us a very nice account of that representational content. I follow Devitt in taking the P-semantic theory to be a molecular theory. That is to say that some of our mental words, concepts, have the meanings because of the inferential relation that they have to other mental words. So the concept BACHELOR means UNMARRIED MALE because of the inferential relations between these mental items. So too in the moral case our concepts acquire their meaning from the inferential relations between them. So, the English sentence 'stealing is wrong' is a (weakly) analytic truth in virtue of the linguistic conventions that govern English (as I will

argue in the next section). The mental sentence STEALING IS WRONG is weakly analytic because of the inferential relations that hold between mental tokens of these concepts. It is also likely, as Hare argued, that the logical words and corresponding thought parts are similarly defined in terms of inferential relations. This allows us to take all of the lessons about the logic of imperatives and attitudes in general that we have learned and apply them here. But it does not follow that all mental tokens get their meaning in this way. Some of our concepts have to actually connect up to the world in order for out thoughts to be of any use at all in getting around. In particular mental names and indexicals have to depend on causal relations to the world for their reference.

5.1.2. Attitude Expression

Bach and Harnish (Bach and Harnish 1979) define expressing as follows. "For S to express an attitude is for S to [reflexively]-intend the hearer to take S's utterance as reason to think that S has that attitude". The reflexive intention that they speak of here is what Grice called a communicative intention but called 'reflexive' to distinguish it from the iterative

intentions that Grice spoke of in his later writings. A reflexive intention is an intention that has the special property that its fulfillment consists in its recognition. My intention that you take my utterance as a reason to think that I have an attitude is fulfilled if you so recognize that I intend, partly on the basis of what I said partly based on how I said it, etc, you to take me as having that attitude. Thus on this view to express the belief that it is raining is to intend for you to take my utterance as a reason to think that I do believe that it is raining. So, on the intentional view the form of the utterance does not have to match up with the mental attitude that I have in the way that the causal theory requires.

It strikes me that the difference between consciously expressing an attitude and doing so unconsciously just is the difference between my wanting you to know what attitude I have and my not wanting you to know. So it seems to me that the causal theory of expressing explains what it is to express an attitude non-consciously. This would explain why we sometimes say things that give our audience a reason to think we have a particular attitude but that we did not intend for them to know that we had that attitude, e.g. so called Freudian Slips. The

utterance is probably caused by some state that we were in, consciously or not, that we did not want our audience to know we were in. So even in communicative contexts the casual theory has some explanatory work to do. And non-verbal behavior can be used to consciously express our attitudes as well. So the distinction between the causal theory and the intentional theory of expressing tracks the conscious vs. non-conscious expressing of attitudes.

Given this theory of what it means to consciously express an attitude we can now define the various kinds of speech acts. So to assert something is to express a belief and a desire that the hearer form 'a like' belief (p 41). To promise is to express "the speaker's intention and belief that his utterance obligates him to do something". To apologize for something, that is to perform the speech act of apologizing is to express regret and the intention that the hearer believe that I regret the act.

I want to introduce a distinctively moral kind of speech act modeled on the way that Bach and Harnish define speech acts. That of morally condemning (or approving). To perform this speech act is to do the following

> *Moral Condemnation*:
> In uttering expression *e*, S morally condemns action *a* if S expresses:
>
> i. The moral feeling of condemnation of the action a

ii. The belief that condemnation is the correct moral feeling to have
iii. The intention that the Hearer H believe ii

By expressing the intention that the hearer also believe that condemnation is the correct attitude to have we are forcing a response from the hearer. This accounts for the demandingness of moral utterances. We can by similar method define the speech act of *moral approbation*.

Let me quickly say that when I say that an utterance of 'Suicide bombing is wicked' expresses an emotive state and the belief that the emotive state is the correct one to have, I mean that this is usually what a person intends to be doing when they say something like that or at least that there is such a basic moral speech act as so defined and people do sometimes perform it. I by no means mean to say that this is always what they do. There are plenty of times when people say those words and mean to be doing something else. I whole heartedly endorse Nowell-Smith's Janus principle. Our purposes when we speak are many and varied.

So, we can imagine someone who, for some unknown reason, genuinely approves of what Hitler did but at the same time knows that this is the wrong way to feel about Hitler. This person might say "Hitler was Evil, and I feel the wrong way about Hitler". It may seem that on my view

'Hitler was Evil' is used to express moral condemnation of Hitler, but this person does not morally condemn Hitler so it looks like he is contradicting himself, even though the sentence itself is not contradictory. However, in this case the person most likely means that they understand that moral condemnation is the appropriate attitude to take towards Hitler and is saying that the attitude that they actually have towards Hitler is the wrong attitude to have towards Hitler. This is not a problem for my view because I only claim that we *typically* use these sentences to express our moral sentiments, not that we do so in every case. This person is using the sentence in a non-standard way, but we often use sentences in non-standard ways. The only claim I want to make is, as already said, that the speech acts that I am pointing out are done by us as well. And that these speech acts capture what we might call distinctively moral speech acts.

But why think that we ever express an emotional state as well as a belief that the state is the correct one to have? That is, why should we think that *moral condemnation* as I have defined it is something that we actually do? This is part of what I think Copp got right, as we saw in chapter three (section 3.2.1.). He there argued that there is reason to think that we do routinely express two mental

states when we communicate about morality. The evidence for this is the queerness of the amoralist. Someone who says that x is wrong but goes on to say that they do not approve or disapprove of x is clearly someone that is abusing moral language. We can illustrate this point by considering the following dialogue.

Jones: Eating meat is immoral, Smith!

Smith: What? Why shouldn't I eat meat? As long as the animal lives a better life than it would have in the wild and it is killed in a humane manner how am I harming it?

Jones: Oh! I didn't mean that you should feel that way about eating meat. I was merely expressing my feelings about it.

In such a dialogue Jones is clearly acting in an unusual way. Which seems to suggest to me that it is part of our ordinary practice to take the utterance of 'x is wrong' to express the belief that the way the speaker feels is the correct way to feel about the subject. So the speech acts I have identified do seem to be the kinds of things that people expect us to be performing which is some evidence that we can perform those speech acts in the Gricean way. As speakers we can reasonably expect our audience to recognize our intention to be expressing those attitudes. So if we think that we typically express our moral emotions in moral communication, as I do, it is natural to

think that we also express our beliefs about those moral emotions in the same utterances.

But notice that at no time in this discussion have we been talking about the linguistic types that people use when they perform these various speech acts. All I have been talking about is the nature of the thoughts and emotions that we have when we make moral judgments. I claim that when we engage in moral discourse part of what we are doing is expressing our moral emotions and the belief that the moral emotion is the correct one to have. The mental states consist of a distinctive mental attitude held towards some representational content. The representational content is given by the causal theory of reference. We express these mental states by formulating special reflexive intentions. The things that we say, and what they mean, are part of what we expect a hearer to use in order to recognize our reflexive-intentions, and thereby fulfill those intentions. What we need to do now is to give an L-semantic theory. That is, a theory of the linguistic types that we use when we express our thoughts.

5.2. The L-Semantics of Emotive Realism

Recall that in section 4.3. I argued that we cannot treat names, that is, *linguistic items*, as singular terms,

or rigid designators. Rather they are what I call '*frigid designators*'. Frigid designators are basically definite descriptions, which following Russell I model as quantifier phrases. To meet this demand I suggested that we adopt Kent Bach's version of what he calls Nominal Description theory, being the view that a name N is semantically equivalent to the nominal description that mentions the name, as he puts it, 'the bearer of "N"'. So the conclusion of that section was that English has no singular terms. It is a purely descriptional device. Our theory of English, as a language, a tool that we English speaking Humans use to express our thoughts, consists purely in predicates, quantifiers, variables, and the logical constants and connectives.

The meaning of these linguistic items is either given by some conventional truth-functional or model-theoretic account as in the case of the quantifiers all '(x)' and some 'E(x)' and logical constants like '&' and '~' and '→' or in terms of satisfying a description in the case of predicates like '…is red', 'is a dog', and definite and indefinite descriptions like 'the first actor to play Hamlet in a Hollywood movie' or 'a dog bit me yesterday,' or as directions for the use of the item to do something (like refer, as in the case of 'it', 'he', etc. or make a

promise or to express pain (as in the case of 'ouch', etc.).[37] The semantical task at this level is that of giving an interpretation of these linguistic types reflected in our categories.

But as we have also seen all of this is an entirely separate matter from determining whether I perform the act of asserting or commending when I utter use these linguistic types to produce tokens. That is a question of what mental states or attitudes or emotions I am expressing (which I addressed in the previous sections). The linguistic meaning of the English words (and indeed, the syntax and grammar of English) will constrain how they can be used, but it will become clear that we can use one and the same linguistic phrase to express any number of mental states.

The linguistic meaning of the sentence that we use will be determined by the dictionary meaning of the words that occur in the sentence plus the grammar for English.[38] I then argue that linguistic meaning in this sense can be

---

[37] I do not mean to be saying that we have a completed L-semantic theory. Debate still rages about the status of deontic logics, tense logics, adverbial logics, etc. The point here is that we have good reason to think that some simple extension of first-order predicate logic with identity is sufficient to model natural language types.

[38] Many people speak 'off dictionary' but this isn't a threat to the claim that I am making. One may have a vague grasp of the L-semantic meaning of the word type and so manage to express a distinct thought using the word whose meaning you are ignorant of. To some degree this happens all the time. We are all more or less ignorant of English.

explained in the Gricean way that appeals to speaker intentions.

Putting all this together I can now argue as follows. 'Saddam Hussein was evil,' that is the English sentence type, has a determinate meaning in English as a matter of its semantics and grammar. So, to get the L-semantics right, and hence to get its (the sentence's) truth condition, you would have to settle the question about what property, IF ANY, 'evil' was expressing and whether or not Saddam Hussein had it. As we have seen candidates for this property can be either natural or non-natural. And if it turns out that there is no such property in the world (as I suspect) then the sentence in question is false. If, on the other hand it turns out that there is a property, 'evil, and Saddam had it then the sentence will be true. All too true. But none of this seems to have anything to do with how I *use* that sentence. The emotivist says that we use that sentence to express some moral sentiment of ours and to try to adjust the attitudes and interests of others; if the sentence we use to do that is semantically false, so what? We often use sentences that are literally false to say true things.

For instance if I said "he never shuts up," or "he never means what he says" the sentence that I utter is literally

false. 'Never' means 'not ever' or 'at no time' and it can't be true that the person in question at no time stops talking or that they do not ever mean what they say. But this does not stop me from communicating something true about this person by uttering these sentences. This kind of thing is absolutely commonplace amongst language users (Bach 2001). So if it turned out that I say something false in the course of expressing my moral condemnation of Saddam Hussein and my belief that that is the correct way to feel about Saddam would have no bearing on whether or not my belief was true, and so no bearing on whether or not I am a moral realist or not. This is nice since I argued before that ones theory of semantics and ones theory of justification are completely independent, and now we have seen that that truly is the case.

But as we have seen, the only reason to be married to this kind of semantic claim is because of a desire to distinguish one's realism from other 'cloaked' irrealisms. But since emotive realism is compatible with any theory of justification independently of the semantic story we tell we may want to be neutral over semantic theories. So, though we could take a L-semantic view like Copp's or Plato's, and hold that adjectives stand for some property and that their semantic function is to predicate that

property to some thing, this is not the only option for the semantics of the sentence type in question. As I have argued one could in principle adopt any semantic theory that one wanted to, but for my part I adopt the Gricean strategy of explaining word meaning in terms of what speakers have generally intended to mean in uttering those words. A reliable guide to what speakers have generally intended to mean when they use certain words is the definitions tracked by the dictionary.

So, if one were to look up the adjective 'evil' in the dictionary one would see that it means 'morally reprehensible'. Thus the sentence say, 'Uday Hussein was an evil man' by which I mean to be talking about the English sentence type that you and I could both use to express our thoughts and feelings about Uday Hussein, will mean something like (1) below,

    (1) the individual called "Uday Hussein" deserves moral condemnation

and this sentence will be true just in case Uday did deserve moral condemnation and false if he did not. This will depend on what the correct account of justification is and so is neutral as between all theories of justification just as promised.

Similarly, something like 'keeping one's promises is good/right' will mean something like (2),

    (2) doing what one is bound to do is commendable/in accordance with what is just, good, or proper

'bound' (the word type) is itself defined in terms of 'obligation'. We thus heed the real lesson of the open question argument. The moral words are all inter-defined and completely independent of any theory of justification. None of the definitions of the moral words mention, or need to, any essential property or attribute of the things that are good, just, or proper. That is the job of a normative theory, which I shall address in the final section. So, just as Moore pointed out, we cannot define 'Good' in terms of some other property, it is not reducible to any other kind of thing. The open question argument, taken solely as an attack on defining the linguistic meaning of a moral word, is successful. The meaning of the word must be given within the framework of morality. But that does not mean that there are non-natural or non-reducible moral properties. There are only regular non-moral properties that from within the moral framework matter or do not matter.

Either way you prefer to define 'evil' this immediately solves the Frege-Geach objection we discussed

in section 2.1.1. Sentences have all the usual logical properties and so behave completely normally: Logic is a syntactic theory not a pragmatic one. Here is an interesting passage from Davidson (Davidson 1967) where he recommends that we keep the distinction between "questions of logical form or grammar, and the analysis of individual concepts,"

> If we suppose questions of logical grammar settled, sentences like 'Bardot is good' raise no special problems for a truth definition. The deep differences between descriptive and evaluative (emotive, expressive, etc) terms do not show here. Even if we hold there is some important sense in which moral or evaluative sentences do not have any truth value (for example because they cannot be 'verified'), we ought not to boggle at "'Bardot is good' is true if and only if Bardot is good"; in a theory of truth, this consequence should follow with the rest, keeping track, as must be done, of the semantic location of such sentences in the language as a whole—of their relation to generalizations, their role in such compound sentences as "Bardot is good and Bardot is foolish," and so on. What is special to evaluative words is simply not touched…(p 105)

In this instance I agree with Davidson. It is simply a matter of the syntax and grammar of English that the logical connections hold. And since ethical sentences will have truth values just like any ordinary sentence the syntactical transformations will preserve truth in just the way that it always does. The Frege-Geach problem has never been a problem for expressivism since it can either

adopt deflationsism about truth or the L-semantic/P-semantic split I have been arguing for.

On this view to determine whether the sentence is true or false we have to imagine that it is being used to express a singular thought about a particular individual.[39] I have been intending to refer to the Saddam Hussein who was in fact the leader of Iraq and was ousted by the actual U.S. Government and who was actually subsequently executed through-out this paper, and I think you have so taken me as referring to Saddam. This means that the sentence 'Saddam Hussein was evil' will be true just in case Saddam deserves moral condemnation.[40] In other words it will be true just in case the moral sentiment typically expressed by people who use the sentence is the right sentiment to have.

5.2.1. Some Moral Truths are Analytic

An interesting consequence of this kind of view is that some moral truths will turn out to be analytic. Now, before any hackles are aroused, I agree with Devitt's

---

[39] Sentence types aren't really true or false. Sentence types have truth *conditions* determined by their syntax, grammar, and the L-semantic meaning of its word types. Only a token of the type, produced with the right kind of intentions, is truth-evaluable and then it is the thought that is true or false (see section 4.3.1.).
[40] Just to see the contrast, on my view the sentence type means 'the bearer of "Saddam Hussein" deserves moral condemnation'. I claim that this thing has no truth value. It is only the thoughts that it can be used to express which can be true or false.

Quinian characterization of analyticity in <u>Coming to Our Senses</u>. This kind of account rejects the traditional link between analytic truths and a priori knowledge. To say that a sentence is analytic is just to say that it can be reduced by definition to a logical truth. So, according to the L-semantic theory that I have argued for, the sentences 'murder is wrong' and 'you ought to keep your promises' will mean something like (3) and (4) below,

   (3) unjustified killing is unjustified

   (4) you have an obligation to do what you obligate
       yourself to do

which are both logical truths. The analytic truths serve as the major premises in moral arguments and are completely neutral as between theories of justification. Now, again, all of this has bee solely about sentences construed as linguistic types.

This kind of position has had a long and venerable history in Western philosophy going back at least to Aristotle (Aristotle 350 bce/1984) when he says in Chapter six of Book two of the <u>Nicomachean Ethics</u> that there are actions like adultery, murder, and theft, which, "..imply by their names that they are themselves bad," (1107^a 10). This is very plausibly interpreted as Aristotle claiming that 'theft is bad' comes out true by definition.

Locke also very clearly has this kind of view in mind in the <u>Essay</u> (Locke 1690/1975). Consider this passage from Book IV

> *Where there is no property, there is no injustice*, is a Proposition as certain as any Demonstration in Euclid: For the Idea of Property, being a right to any thing; and the idea to which the Name Injustice is given, being the Invasion or Violation of that right; it is evident, that because these Ideas being thus established, and these Names annexed to them, I can as certainly know this Proposition to be true, as that a Triangle has three Angles equal to two right ones. Again, No Government allows absolute Liberty: The Idea of Government being the establishment of Society upon certain Rules or Laws, which require Conformity to them; and the Idea of Absolute Liberty being for any one to do whatever he pleases; I am capable of being certain of the Truth of this Proposition, as of any in Mathematicks.

The position he is marking out is just the one I have been advocating. We start with the definitions of the concepts and deduce the moral propositions in just the same way that mathematicians start with definitions and deduce theorems. But for all of this I have not said anything about justification at all. These sentences just are logical truths in disguise and so, as we would expect, they are true under any interpretation. So whether utilitarianism, constructivism, or intuitionism is true we simply substitute that in for 'justified'.

Now, I also follow Devitt in thinking that the logical truths themselves depend (partly) on how the world is. As he says,

> In virtue of what is 'All unmarried's are unmarried' true?...it is true partly in virtue of what it means and partly in virtue of the way the world is, in virtue of all unmarried *being* unmarried. And because the truth of 'All bachelors are unmarried' depends on the logical truth, its truth also depends partly on how the world is.

Transferring this account into our theory of moral semantics we get the claim that though 'murder is wrong' is analytic in the sense that it can be reduced to the logical truth 'all unjustified killings are unjustified' it is not knowable *a priori* since 'unjustified killing is unjustified' depends partly on the empirical fact that all unjustified killings are in fact unjustified.

So my claiming that 'murder is bad' is analytic does not mean that I am committed to the claim that the logical truths are necessary or that they are knowable a priori, which are both views that have been shown to be incompatible with naturalism.[41] So it is an empirical truth that all unmarried things actually are unmarried and it is an empirical truth that all unjustified killings are unjustified. It is conceivable that we have evidence at

---

[41] Though it does commit me to the view that it is at least epistemically possible that an unjustified killing could also be justified.

some time in the future that would force us to change our views about the logical truths. So, to take an excellent example of David Rosenthal's, consider the following case.[42]

Suppose that I am married (I am not) and that my wife and I decide to get a divorce. We get lawyers and hold meetings and hammer out an agreement. The lawyers draw up the papers and my (soon to be) ex-wife signs the papers on Friday afternoon. I can't come in on Friday, and so I arrange with the lawyers to come in first thing Monday morning and sign the papers. Now let us suppose that I go out on Sunday night and meet another woman and let's further suppose that we end up sleeping together. Now it seems that I cannot be accused of committing adultery and since only married people can commit adultery there is a sense in which I am not married (and so a bachelor). But technically I am still married (I have not finalized the divorce). So it looks like on Sunday night I am a married bachelor; or at least it is conceivable that we could describe the situation this way.[43]

---

[42] This example was used by Rosenthal in the same seminar on Quine and Sellars I mentioned in chapter four. As far as I know it has never been published or discussed in the literature.

[43] In fact, as pointed out by Baumrin in personal communication, this would be the case even if I had signed because it needs judicial approval and so would not be official until Monday anyway.

Now, people usually resist the conclusion of this scenario, but let us suppose that it is accurate. Is it a counter-example to the alleged analytic truth that bachelors are unmarried males? I don't think so. Rather what is at issue here is what counts as being married. Once we settle whether I am really married or not then we will settle whether or not I am a bachelor. What seems to be the issue here is whether or not all marrieds are married. That is, these kinds of examples show that even the logical principles are subject to empirical verification.

Similarly, consider an alleged counter-example to the anniversary business. Suppose that I get married on Leap year. Then my anniversary will not occur one year later, right? Again, I do not think that this is a counter-example. What is at issue here is 'what counts as one year?' This is actually nice in that it allows us to see the true nature of moral debate. When people are debating about whether murder is allowed or whether Robin Hood really robbed from the rich they are debating, not about whether murder is wrong or right –no one disputes that--, or whether it is morally permissible to rob --everyone knows that it isn't--, they are debating whether or not

the particular action in question should be described as a murder or should count as a robbery.

So, in the moral case we will say that the analytic truths are true partly due to their meaning, and partly due to the empirical truth that the wrong things are wrong. I will address this when we turn to talking about normative issue in the last section. But what of Quine's arguments against analyticity, even in this weak form? The Philosophical Lexicon defines the verb 'to quine' as follows,

> **quine**, v. (1) To deny resolutely the existence o[r][44] importance of something real or significant. "Some philosophers have quined classes, and some have even quined physical objects." Occasionally used intr., e.g., "You think I quine, sir. I assure you I do not!"

The joke, of course, is that Quine denied the existence of the analytic/synthetic distinction which is seemingly both real and important (Grice and Strawson 1957). There are a lot of philosophers who think that Quine is right, so we need to briefly see if there any reason to think that there is something wrong with this distinction.

I think that Devitt has the right answer to this question. He notes (p. 23 Coming to our Senses) that there are two sorts of arguments in Quine's famous Two Dogmas

---

[44] The actual entry says "…the existence of the importance of something…" I interpret the first 'of' as 'or'.

paper. The first sort point out that there is no way to break out of a small circle of terms like synonymy, semantic rule, and definition. Given any one of these we could easily get analyticity, but they each assume analyticity and so it is hopeless. But Devitt correctly points out that this cannot prejudge all attempts to break out of the circle. Each semantic theory has to be judged on its own merit here. As Devitt says,

> [it] must be an open empirical question whether a scientifically respectable account of meaning can be given, and whether a theory that gives this will be one…with weak analyticity of the sort just noted. So I shall say no more of these arguments

He goes on to talk about the second sort of argument, which is aimed at establishing a certain epistemological claim. This epistemological claim is meaning wholism.

But if we are only thinking of our L-semantic theory, wholism is no threat at all. In fact it is what we would expect given such a complex human tool like language. What we have is a vast, complicated construction all of the parts of which are inter-defined. The linguistic meaning of a word is determined by the conventions of the language. The conventions of the language arise because of Gricean intentions that people have had. So, 'cat' means cat (I.e. is semantically equivalent to the phrase 'domesticated carnivorous mammal of the species *felis*

catus') because people have generally intended to be expressing thoughts about those kinds of mammals when uttering that word. So when we talk about what a speaker of English would have to know in order to use a word correctly this is what we mean. We are talking about knowing the L-semantic theory for English, or in other words, the meaning of the word types in English.

This also defuses Hare's argument against this kind of strategy. He says,

> …we might assert the self evidence (because analytic) of the principle that we ought to do our duty; and then we might argue that we could ascertain what our duty was by some fact-finding process…(ibid p. 44)

Given this we would be able to deduce what our specific duties were. As we will see in the next section this is exactly what I do think is the case. So what is the problem here?

Hare goes on to complain about the meaning of 'duty' in the inference we make. So consider the following example.

> 1. I ought to do my duty
> 2. I have a duty to keep my promises
> I therefore ought to keep my promises

He says,

> But this is an equivocation. If 'duty' is a value word, then we cannot decide what our duty is merely by consulting word usage or by seeing whether we have a certain psychological reaction, but only by making

> a moral decision. On the other hand, if 'duty' were not treated as a value word, but regarded as meaning either 'that which I have a certain recognizable psychological reaction' or 'that to which the name "duty" is commonly applied in my society', then the principle 'One ought always to do one's duty' would not be self-evident (p. 44)

On my view 'duty' is a value word and it means, by definition, something like 'a moral or legal obligation'. I know this because I know English, and if I were unsure I could verify it by checking any standard definition of the word given in a dictionary. But I can usually tell by reflection that this is the meaning of the word and so I so take (1) above to be in a sense self-evident. But I agree with Hare that it is not self-evident in the sense that his intuitionist opponents meant. It is true partially because all duties are duties and everything obligatory is obligatory, so it is true partially in virtue of how the world is, just like all 'weakly' analytic truths. But these regularities are so far so regular that that these kinds of inferences are on as stable ground as any in mathematics or logic.

What sorts of moral truths are analytic? Murder, theft, and adultery are the easy cases. What are we supposed to say about 'lying is wrong' or 'torture is wrong'? Are these analytic? It certainly does not seem, at least at first glance, that 'lying is wrong' is equivalent

to 'all unjustified falsehood telling is unjustified'. According to common sense a lie is a falsehood, or at least something believed to be a falsehood, told with intent to deceive. This is the L-semantic meaning of 'to lie' in the moral sense of the word.

One could if one wanted maintain that this was the meaning of the word and then argue that 'deception' itself involves justice as part of its definition. If this were the case then 'lying is wrong' would look like (5),

(5) Making an untrue statement with intent to deceive is an act marked by injustice or deception

This way of doing the L-semantics of 'lying' does have it be the case that 'lying is wrong' is analytic. The same can be said for 'torture'. 'Torture,' the English word type, is defined as 'inflicting intense pain in order to coerce, punish, or for pleasure'. Any glace at a dictionary can verify this. According to this definition is 'torture is wrong' analytic? It is. Part of the meaning of 'wrong' in English is that an act is *injurious*. That torture is injurious seems indubitable and beyond dispute. So we do end up with an analytic truth; namely that injurious actions are injurious.

Interestingly, Hare gives a similar argument. He says,

> …we might say that to tell a story about someone, which everyone knows is *ben trovato*, is not *lying*. We can say this, because 'lying' does not mean simply telling falsehoods, but telling falsehoods that are *reprehensible*. Thus we might, and sometimes do, make a distinction between lies proper and white lies; lies proper are reprehensible; a white lie on the other hand, is In the words of the *Oxford English Dictionary*, 'a consciously untrue statement which is not considered criminal: a falsehood rendered venial or praiseworthy by its motive' (LoM p. 53)

Thus 'white lies' are not actual lies. The fact that we call them white *lies* doesn't matter. Whatever we call this distinction we clearly recognize that some falsehoods are justifiable. So even if it is a bit revisionary of English usage we have reason to define 'lying' as 'unjustified falsehood' or 'reprehensible falsehood'.

When we call something a lie what we are doing is claiming that the thing in question was unjustified. That is why we have those kinds of words in our language. Those who disagree with us are disagreeing that the action in question is accurately described as a lie. This is why George Bush can say, honestly, 'the U.S. doesn't torture'. He of course agrees that torture is wrong. What he and I disagree about is what counts as a justifiable amount of pain and suffering which can be inflicted in the course of an interrogation. We both agree that there are some easy cases which obviously count as torture. In a recent memo released George Gonzales can be seen as arguing that all

amounts of pain are allowed except those pains equal in intensity to death or organ failure. I disagree that those levels of pain are justified and so think that some of the things we do are torture. So I say that the U.S. does, in some cases, torture in the course of interrogation. But yet we both agree that torture is wrong.

The only way to resolve this dispute —because according to the realist as I see him we cannot both be right about this; It cannot be the case that it is both true that the U.S. tortures and that it does not— is to immediately 'descend' to the normative level and begin to have a normative debate. This will involve invoking theories of justification and arguing their relative merits. So, the Bush argument is a roughly utilitarian one. If the information can save many lives then extreme measures are justified. The arguments against this stark utilitarian position usually appeals to something like Kant's second version of the categorical imperative, or to some natural rights that humans have. There is also a utilitarian argument to the effect that when we use excessive measures we send a message to the enemy that these methods are OK and so we encourage the use of these methods on our own troops.

So to sum up a bit. On this model of moral discourse a lot of moral arguments are about how we ought to categorize various actions or persons. We have an analytic moral truth, a particular statement asserting that an action or person satisfies the description (that is, is unjustified). A reason to think this is true is that it helps explain how the kinds of disagreement between the President and I could happen. We both accept that tautology 'unjustified infliction of suffering in interrogation is unjustified'. To foreshadow a bit, only someone who was morally and rationally blind could deny this. But what we may rationally disagree about is whether some action (like water boarding) is justified or not. I think that it is not. The reasons for thinking so come from a normative theory.

This, in my view, takes care of a lot of cases but there are cases where this model won't work. To take one example consider, 'Same sex marriage is moral and should be allowed and recognized by the U.S.' This doesn't look analytic. To deal with these kinds of questions, the interesting ones, we, again, need to immediately descend to the level of normative theories. This is, of course, the way in which Blackburn's quasi-realism appears to be

doing something legitimate. Normative questions need normative answers.

But there is a sensible further question; which normative answers are best? The quasi-realist tries to hear that question as itself a normative issue, which it is, but that doesn't mean there isn't a real answer to it. So, before we conclude let us turn to a quick discussion of some normative issues. I will discuss the issue just mentioned in the section after next (5.3.2.), but first I want to address the question of whether it makes sense to justify our feelings.

5.3. Justification and Emotive Realism

Up till now everything that I have been doing has been entirely at the metaethical level. I have explicitly tried to avoid saying anything about particular normative judgments. I have only been talking about the form a theory of language, thought, and communication would have to take in order for it to support a moral realism that is compatible with one key claim of the sentimentalist tradition (and thus naturalism) and at the same time compatible with moral realism. The view I have been defending, and elaborating, is that when we say that something is morally wrong we typically perform a speech

act I have called *moral condemnation* which involves simultaneously expressing a moral emotion and the belief that the emotion is the correct one to have.

This kind of view, as I have argued, is in principle compatible with any theory of justification. This amounts to saying that we have various theories about why the moral sentiment that we have is the correct one to have. Whichever theory of justification turns out to be the correct one is compatible with the view as so far stated. So one could in principle be an irrealist about justification and still adopt the view of moral condemnation that I have been arguing for.

Exploiting Hutcheson's (Frankena 1955) distinction between exciting and justifying reasons we can say that my moral emotion is my exciting reason, the justifying reason can be whatever the correct normative moral theory happens to be. One such theory could be that there is no such story that we can tell. This would be an emotive subjectivism. We think our attitudes are the right ones to have simply because they are our attitudes. This would be a form of subjectivism or relativism. This means that the kind of view I have defended *could* be a form of classical emotivism. One could think that we perform the speech act that I have pointed out but also think that we are simply

mistaken in believing that the emotion is the correct one to have, since there is no correct way to feel about the subject. I do not think this, but for all I have said so far this could be true. So the kind of view I am defending is even compatible with the Blackburnian kind of theory of justification examined in the third chapter. One could have an emotive irrealism that was based on the neo-Gricean philosophy of language that I have defended as an alternative to Blackburn's deflationary view of truth and bulldozing of the meaning/use (L-semantic/P-semantic) distinction.

However, as I have been keen to indicate, I am interested in the prospects for a naturalized moral realist account of justification and it is open for such a person to say that the attitude was the correct one to have because it maximized happiness, or conformed to the Categorical Imperative, or would be agreed upon by rational agents behind the veil of ignorance, or was conducive to the good of society/the species, or whatever. We could agree with Aristotle and Hume when they say that the correct attitudes to have are those that a rationally virtuous person would have or we could even hold some form of intuitionism. One could even be a Platonist and think that there were mind-independent moral properties and that

moral predicates referred to these properties and that when we say that something is wrong we perform the speech act of moral condemnation as I have defined it. These would all be versions of emotive realism. So in one sense it does not matter which theory of justification that one is attracted to. It will still be possible to adopt the view I have defended here.

But what every version of emotive realism will have in common is the claim that our moral emotions can be justified. Versions of the theory will be more or less realist in so far as they are more or less committed every moral disagreement (i.e. every case of conflicting beliefs about the correct attitude to have) is such that two people who disagree cannot both be right. We will have modest versions of moral realism, like Steve Ross' which only make this claim about a select group of moral disputes over basic moral judgments that all realist normative theories of justification agree on, and we will have more brash versions according to which most (maybe even all) moral disagreements are like this.

Because of this commonality it is necessary to spend some time examining the very possibility of having justifying reasons for our emotional responses to the world. In the next section I will defend the claim that we

can have justifying reasons for our moral emotions. Afterwards I conclude this chapter with a discussion of the relation between constructivism and realism.

5.3.1. Justifying our Feelings

Does the claim that there can be correct and incorrect emotional responses to the world make sense? It does. At the very least, as Foot pointed out in her article *Moral Beliefs* (Foot 1958), there are logical relations that hold between concepts like 'good,' and 'right' and the things of which they can be predicated. She begins by looking at the concept 'dangerous' and the related fear reaction. There are things that it is quite literally incorrect to fear and which we cannot properly call dangerous. Monty Python's *The Holy Grail* springs to mind. When Arthur and the Knights of the Round Table first saw the little bunny that Tim was so freighted of their first response was laughter because bunnies are not the proper kinds of things for humans to fear.

Of course, they find out that Tim was right. They should have feared the bunny because it leaps very far and it has big nasty teeth, just as Tim had warned them. Whether or not it is appropriate to fear something is entirely a matter of facts about it taken together with

facts about us. So, if the bunny had not had those properties (of leaping and having great big teeth) or if the facts about human beings were different (their necks were not vulnerable to gashing, etc) then it would have been Tim's affective attitude that was incorrect (the fear of the rabbit), which is to say that Arthur's belief that Tim's affective attitude was not the correct one to have would have been true. So emotions can be said to be correct or incorrect responses to the world based on relevant human characteristics.

Foot then points out that the case is the same with concepts like 'good,' and 'evil' and our moral sentiments. There are things, like a person's preferring chocolate ice cream to vanilla, that are quite wrong to call evil.[45] Someone who did so, and who was really expressing moral condemnation of a person who did not like chocolate ice cream, and who sincerely seemed to despise chocolate ice cream and believed that this was the right kind of attitude to have about this kind of ice cream, that is to say someone who was not engaging in hyperbole, or metaphor, but was really and sincerely saying that this preference was evil, would really be making an incorrect

---

[45] Assuming, of course, that there are no special circumstances. One can of course imagine cases where choice of ice cream flavor has moral implications. But in the normal everyday circumstances we find ourselves in this is not the.

judgment about the preference for this flavor of ice cream.[46] Moral condemnation is not appropriately taken with respect to flavor of ice cream (assuming no special circumstances) and my belief to that effect is true. Someone who disagreed with me about this would really be contradicting me. Our beliefs could not both be true. So here we have a case where someone's moral emotion can be evaluated as correct or incorrect.

Let us see if we can extend Foot's insight a little further. Imagine that there were a person who went with some friends to see *Schindler's List*. Now imagine further that this person found the movie to be genuinely *funny*. Suppose him to be laughing and smiling, pointing at the screen while holding his sides, looking at those around him to see that they got the joke, and etc. Surely my belief that his attitude towards the movie is not the correct one to have is true! Whatever the correct response to that movie is, it does not involve laughing like one was watching *There's Something about Mary*.

The interesting question now becomes 'what reasons do we have for thinking that the movie does not fall under the extension of 'funny'?' and this is itself a question

---

[46] The judgment in this case is that the preference for chocolate ice cream belongs to the class of things that 'habitually arouse' the moral attitude expressed by the use of 'Evil' (cf Ayer's remark on p21 of LTL)

about what normative theory we should adopt. It seems to me that we have reasons of two kinds. First, we feel that the suffering of the victims, as portrayed in the movie, is disregarded or regarded in the wrong way when someone laughs at it. Hitler created a lot of human suffering and suffering is bad. Secondly, we feel that Hitler's actions were unjust. It is not only that he increased suffering but also that he treated a certain group of people radically different from the way he treated alleged pure Aryans and that this difference in treatment was unfounded.[47]

Notice that these two kinds of reasons, the injustice of his actions and the unjustified suffering as the result, are the same reasons we would give if asked why someone who morally praised Hitler's Final Solution (as opposed to thought it was a laugh riot) is despicable. In fact these seem to be the very reasons that we give for saying that what Hitler did was evil. So the idea that we can say that there are correct emotional responses to the world does make sense. We can evaluate the emotions. So it is not simply because theft offends the emotions of mankind that it is wrong. The claim is that theft ought to

---

[47] Finally we might point out that he broke his word by violating a formal agreement he had with Russia, but this seems to be a petty crime compared to the others mentioned.

offend the emotions because some suitable normative theory tells us that this is the correct way to feel about it. The emotions need not play any justificatory role. Their purpose is to motivate us to action in combination with our moral principles.

At this point I could rest content. The possibility of this kind of theory, an emotive realism, allows the very best of both moral traditions to harmoniously co-exist. Whichever theory of justification, and whether it involved moral properties or not is a different project. This is still no small accomplishment. We have arrived at a view that is Aristotlean in inspiration. It says that there are some moral truths that are just logical truths. The logical truths combined with the empirical truths, in at least some instances, is enough for us to deduce some particular moral truths. This is the upshot of Locke's passage and Searle's example and the driving intuition of all intuitionists. It emphasizes the role that reasoning plays in ethics, it is a justificatory role, while recognizing as well the role that the sentiments play, a motivating role.

5.3.2. Moral Arguments and Justification

What then are we to say about moral properties? Can't we say that actions which are murders have the property of

being wrong? This is a harmless way to talk if what we mean is that a particular normative theory tells us that that action is a murder. An action can be said to have the property of being a murder in the sense that some normative theory classifies it as a murder and it will therefore have the property of being wrong, since the theory has as one of its axioms 'murder is wrong'. But this doesn't add any strange properties to our ontology. We can see normative theories as functions from natural properties to categorization as right or wrong. What we have are natural properties. These are the properties of human beings and the world they find themselves in. We then construct theories that make sense of these properties from a certain point of view.

This kind of view immediately gives rise to a problem pointed out by Steve Ross. Consider the following argument,

1. Murder is wrong
2. Action A is a murder
Therefore, action A is wrong.

(1) is, on my view, an analytic truth, which in my terms means that it can be reduced by definition to a logical truth, in this case to 'unjustified killing is unjustified'. This explains why the argument is formally valid and why emotive realism of this kind is immune to

the Frege-Geach problem (as discussed in section 5.2.). It really says the following.

> 1' unjustified killing is unjustified
> 2' Action A is an unjustified killing
> Therefore, action A is unjustified

This captures one of the insights behind arguments against relativism (cf. section 1.2.1.).

The relativist often starts from the claim that it is a fact that anyone can observe that cultures vary in their moral beliefs. This kind of variance is to be explained, says the relativist, by the absence of truth in that area. The typical realist response is to point out that there is less disagreement than there seems to be. The Eskimo mother who puts her child out on the ice does not think that she is committing murder. She thinks that she is performing a justified, though regrettable, act. When we disagree with her we are therefore arguing about the second premise. We are arguing about what counts as an unjustified killing. In general moral debates are about how to classify some action; Is it a murder or isn't it? This amounts to asking whether the action has justification and so is to immediately engage in a normative ethical dispute.

The problem is that since there are many theories of justification there will some theories according to which

action A is justified and some on which it is unjustified and so whether or not action A is a murder or not cannot be settled. Now if one is a constructivist in the sense that I have identified this may seem to be the end of the line. Since the project that we are engaged in can go either way on this there is no way to answer this question. Hitler's being evil is a consequence of every moral theory. As Steven Ross has argued (Ross 1998) this will allow us to definitively say of some actions that they are murders, or that they are thefts. But there are many cases where we do not have agreement and in those cases they may be no real answer to which of the theories has got it right. And this may be as far as we can get.

I have always been attracted to a 'mixed' theory of justification much like that argued for by William Frankena in his book *Ethics* (Frankena 1963). We need some combination of consequentialism and deontology to fully capture our intuitions about morality.[48] These two kinds of considerations strike me as *prima facie justifiers*.[49] So, on some occasions my belief that my emotion is the correct one to have (or that yours is incorrect) will be true because the actions/people that arouse my passions have

---

[48] Frankena argues for a principle of beneficence and a principle of justice
[49] Much like W.D. Ross's *prima facie* duties

negative or positive effects on the amount of suffering in the world. But on other occasions my belief will be true because we have duties and obligations.

Each kind of consideration is in itself *prima facie* always correct, but what is actually the right reaction will depend on working out what takes precedence in any given situation. It seems to me that what it means to be a moral agent is to develop the practical skill that allows one to know when to use which kind of reasoning. Having this skill is having what Aristotle called *phronesis* it is the skill of correctly determining what kind of reasoning to bring to bear. But what is it that grounds these two kinds of considerations?

Frankena, borrowing Bair's useful phrase, argues that these are the kind of things that withstand scrutiny from the *moral point of view*. As he says,

> My own position…is that one is taking the moral point of view if and only if (a) one is making normative judgments…(b) one is willing to universalize one's judgments; (c) one's reasons for those judgments consists of facts about what the things judged do to the lives of sentient beings in terms of promoting or distributing non-moral good or evil…(p 113)

On my view (a) consists in having the mental states described in the speech act of moral condemnation. But why should we accept what (b) and (c)?

The reason that universalizibility is important to the moral point of view is because it is the manifestation of the general ability of the human mind to recognize logical contradiction as an indication that something is amiss. 'If the law applies to everyone then the law applies to me' is a logical truth. Kant gives a similar argument in the <u>Groundwork Metaphysics of Morals</u> (Kant 1785/1997) when he says,

> If we now attend to ourselves in any transgression of a duty, we find that we do not really will that our maxim should become a universal law, as that is impossible for us, but that the opposite of our maxim should instead remain a universal law, only we take the liberty of making an exception to it for ourselves (or just this once) to the advantage of our inclination. Consequently, if we weighed all cases from one and the same point of view, namely that of reason, we would find a contradiction in our own will, namely that a certain principle be objectively necessary as universal law and yet subjectively not hold universally but allow exceptions.(p 33-34)

The reason this works, as pointed out by Baumrin (Baumrin 2004), is because there are certain interests which are universally shared by all sentient beings. Thus are we able to secure our negative duties. That is, our duties not to cause unnecessary suffering, or not to break our promises, or not to commit adultery or theft.

We are able to work out the logical implications of our actions and from the moral point of view this reveals that we have duties and obligations. So for instance from

the fact that I say 'I promise to pay you back' combined with the definition that promising just means that you have placed yourself under an obligation it follows that you ought to keep the promise. This just is Searle's famous derivation of an 'ought' from an 'is' (Searle 1964). This suggests that, in some cases at least, given the facts of the situation, the meaning of our words, and a theory of speech acts, we can derive what our duties and obligations are. It is in this sense that the intuitionists had it right, though they wouldn't have cashed it out in the way we have here. Our duties are deduced.

Some are tempted to criticize Searle by charging that he assumes a deontic premise when he says that making a promise just is placing ones self under an obligation. However, as we have seen accepting this as a tautology is merely accepting something that is true by definition (exactly like 'bachelors are unmarried males') and whose acceptance just is the acceptance of the moral point of view. A person who does not accept the tautology that making a promise is undertaking an obligation is not only immoral but also irrational.

They are, in Baumrin's useful phrase, 'morally blind' (Baumrin 1986). Being a rational animal just is being the

kind of animal that understands that promising is in itself a reason to keep the promise or that some thing's being the right thing to do is a reason to do it or that causing harm to some creature is a reason not to perform the action that results in the harm or that if something is true of all things it is true of some particular thing, etc. That is because at bottom most of these things are logical truths, or they are reducible to logical truths. Someone who denies this either fails to see the logical connection between the concept of promising and the moral point of view (or suffering and the moral point of view), and so lacks the capability of rational thought (like a dog) or they do see this connection but disregard it and so lack respect for the dictates of reason. Hence morally praising Hitler is both irrational and immoral. It is immoral for the reasons cited above (having to do with suffering and justice) and irrational because one who does so irrationally ignores the connection between these concepts and Hitler's actions.

So, given the way that we use our words and the way this world is, you cannot reasonably disagree with me when I say that murder is wrong. You can disagree with me about whether some particular action is a murder or not, and so whether the action really deserves moral condemnation, but

'murder is wrong', the sentence and the thought, are both of them (weakly) analytically true. To disagree is just to fail to understand the words that are used. But of course whether some particular action is a murder or not is not a matter of what our words mean. That is a matter of what our theory of justification is and how it classifies actions.

As said, thus are our negative moral duties grounded. But as also pointed out by Baumrin in the same essay (ibid.) this uiversalization procedure does not work so well for our positive duties (i.e. duties like giving to charity). This is because individual interests vary too much to make universalization relevant. These kinds of duties may really turn out to be relative in my preffered sense. That is, it may well turn out that two people who disagreed about what percentage of a paycheck one was morally obligated to donate to charity and both of them could in some sense be correct. Everyone agrees that we ought not to be indifferent to the needy, but to what degree we are obligated to them may vary.

This is the normative interpretation of the claim that some moral truths are (weakly) analytic. That is, it tells us how these analytic truths relate to practical reason. The answer looks to be a broadly constructivist

one. Since these things are constructed with our interests in mind we have good reason to do what morality dictates. It also offers a systematization of our moral judgments. This is completely separate from the metaethical theory that one holds. In my case I have been arguing that it justifies our emotional reactions to the world. Because we are rational creatures with a certain kind of goal we come together and construct a system of norms, via some procedure. These norms then give us guidance as to which emotional reactions to cultivate. Being a moral agent, on my view, just is having this kind of interplay between the emotions and reason. But we haven't yet really answered the burning question, for we haven't yet said what ultimately grounds the project that we are engaged in.

5.3.3. Meta-Metaethics: Constructivism and Realism

So what, then, is a metaethical theory? This question has become obscure. The issue begins by asking if there are moral truths. When we ask what that means we get a variety of metaethical views. The most famous is Plato's that sees no difference between our ordinary kinds of discourse and moral discourse. So moral wrongness, on this view is thought to exist as a property, just like tables and chairs exist as objects. There are moral truths

because there are mind independent properties that our thoughts/sentences correspond to. This is one sense of realism. The things which are supposed to be real in this case are mind-independent non-physical moral properties. Then we also have reductive realism, which we also know by the name of naturalism. The reductive realist agrees that there are moral truths and that moral predicates refer to properties. The naturalist then says that these moral properties just are natural properties. Of course against all of these views there is the relativist who claims that there really are no moral truths. There are no moral properties and we are all either making a drastic mistake when we say that certain things are good or right (this is the error theorists view) or we are simply expressing our emotions. One of the points that I have been trying to make is that we should be careful to note that all of this is neutral over our semantic theory. We could be irrealists and think that moral predicates are supposed to stand for moral properties. This would be the emotivist who was also an error theorist. We could also think that moral predicates do not refer or stand for moral properties and yet think that we are expressing a moral sentiment that can be justified. This is my view and earlier Blackstone's view.

But then we have the deflationist about truth jump in and try to collapse our distinction between normative ethics and metaethics. The deflationist tries to mimic the realist by capturing all of the surface features of our moral discourse without the ontological price that comes with realism. But as we saw earlier this attempt to collapse the metaethical into the normative fails. We can ask how the various moral frameworks hook onto the world. The answer may be that they do not, in which case we have a version of metaethical relativism or it may be that they track some natural or non-natural property. Now, again, I do not think that this kind of argument is decisive against Blackburn or people who like his kind of expressivist view. My point is that those who do like that kind of view need to be upfront about the actual theory of justification that is on offer, which is just to say that Blackburn's view is just a more sophisticated version of Ayer's emotivism..

At this point in the debate the constructivist shows up with a third alternative. They agree with the Blackburnian side of the debate that the emotions are an intimate part of our moral lives and the strong pull of internalism about reasons that comes with it. But they disagree that morality is simply voicing our brute

emotions. This is manifestly not how it is. The constructivist then goes on to argue that there are moral properties and there are moral truths. It is just that the moral properties are constructed by us in the sense that they follow from moral theories we construct. They are mind-dependent properties in the sense that they come into the world with human minds and do not exist apart from human activity. In this sense they are like facts about music, or literature. Our moral predicates then refer to these properties and the beliefs that we express are true or false in so far as they capture these mind-dependant truths.

This view is offered as an alternative to both realism and irrealism. The moral truths are constructed by us out of some procedure (the categorical imperative, the veil of ignorance, or whatever). But is this a metaethical view? Hussain and Shah (Hussain and Shah 2006) have recently argued that constructivism is not really a metaethical view at all, at least not one that is a genuine alternative to realism. Their argument involves looking in detail at one contemporary formulation of constructivism by Christine Korsgaard. But the main point of their argument can be put rather simply. They argue as follows. Constructivism is a normative theory that tells

us how the demands of morality relate to practical reason and it is also a normative theory that allows us to generate principles that we can use to guide our behavior. These principles are the *prima facie* justifiers I spoke of in the last section. But suppose that we were asked why it was that the procedure was the correct procedure? Why should we think that the categorical imperative is the right procedure to follow? Why the veil of ignorance? The only plausible answer, it seems to me, is that we are engaged in trying to construct a certain kind of system. What we want is a system of norms that will help us to live together. We want to induce cooperation.

But why is this the correct project that we should be engaged in? And even if it is, why is it humans as autonomous agents that we should be concerned with? At this point all of the metaethical options seem to be on the table. On might think that there is no way to answer this question. Given that this is our project, then of course unnecessary torture is going to be ruled out, that won't induce cooperation! But if we have a different project then it might very well turn out that unnecessary torture is fine, in that project. Of course, we may balk at calling that project a moral system, given the way that we use the term; but that is no matter. There is nothing

in principle that distinguishes the system we call the moral one from this one that endorses unnecessary suffering. If we had been interested in some other project then what we call morality would be very different.

This is where the analogy with art and music comes out very strongly. Had humans been differently constructed we would find very different properties artistically pleasing. It is only after we have become interested in a certain kind of expression that jazz music will come to have any value. But if this is ones view then this is just a version of relativism. Granted it is not the kind of relativism that has been advanced in the past given that it will endorse a certain kind of modest realism given that we are interested in morality as a project concerning rule governed cooperation.

But one could also think that the reason that the project was the correct one was because of some non-natural property or even that it was so because it maximized a certain natural property (say, happiness). So, constructivism is not an alternative metaethical theory. For it needs to bolstered by one of the standard metaethical theories. Unless of course, one could have some procedure or project that was self justifying. It is conceivable that the project says, as one of its

instructions, that it should be completed and so it may seem that the constructivist can take this route. But the problem is that if one does take this route then any number of procedures will work. There are an infinite set of procedures and projects that will meet that criterion. Norm-governed cooperation is just one out of that bunch. Why is that the one we should peruse? So either constructivism is not a metaethical theory or it is a version of relativism. Either way it is not a distinctive metaethical view that is a new alternative to metaethical relativism.

I see it as normative theory about justifying reasons. We construct principles which serve to guide us in the development of the appropriate feelings. It is the apprehension of contradiction which seems to me to do the trick. Given the laws of logic and the description of the world, including the description of us as rational free agents, the rest follows. These are all natural facts about human beings and so the rational is the right. With reason comes the ability to generate universal laws, physics and logic are two examples of this; And morality is another. So ultimately I view myself as a metaethical rationalist with a constructivist normative theory.

So, what are we to say about the problem of relativity? At the metaethical level relativism is ruled out by the arguments I have appealed to in various places; we have assumed from the start that we could give some naturalistically acceptable morally realist account of justification. The basic point I have been making so far is that given certain interests that all sentient beings share and the contingent fact that there are more than one of these kinds of creatures, the moral framework automatically and immediately shows up in the sentient creatures that are also rational.

Here is an analogy. Suppose that one had never heard jazz before and that one was very interested in Country music. One might not then like jazz. It might sound jarring and strange. What would change one's mind? One thing that would work is trying to play jazz oneself. That would be a way to see what was distinctive of it and why it was worthwhile. You might then begin to develop an appreciation for it. This point seems to generalize. It is when we try to do something, or even imagine doing it, that we begin to be interested in the differ ways it can be done. Now in the case of jazz, I have the choice to do it or not and so I may never come to appreciate jazz. But when it comes to actions we have no choice. We must act

and as soon as we do we adopt the moral framework. This does not mean that we immediately have access to the correct theory of justification. Rather what it means is that with rationality comes universalization and with the possibility of action come the question of what I ought to do.

What are we to make of normative relativism? That is, how are we to adjudicate between the competing normative theories out there? As already argued, there will be many, many cases where there is agreement among all serious theories of justification and so the problem will not arise. Given the way we are and the way the world is it just follows that murder is wrong. But this is not to recognize some mysterious property. It is just to recognize a logical tautology just like when I recognize that a triangle can't have four sides.

The pressing problem lies in cases like legalized prostitution, abortion, same sex marriage, etc, etc. It may seem that there is no way to resolve the debate about these issues. Reasonable people may just, at some point, and even agreeing on all the facts, fundamentally disagree about whether, say, same sex marriage is a rights issue. Thus the worry is that it may be indeterminate whether

same sex marriage really is right or wrong, and so indeterminate what the correct attitude towards it is.

But this seems to me to be no more than an argument for skepticism. It may be the case that we do not know what the right answer is to these questions. But, given that there are well-defined end points to the continuum between moral and immoral it seems to me reasonable to expect that there is a well-defined answer for each point in between. I find it hard to believe that there isn't an answer to the question of whether or not abortion, vegetarianism, or same sex marriage are really justified. I find myself thinking that all normative theories worth the name tell us that same sex marriage is justified, that we ought not to eat animals, and that abortion is permissible in some cases, and so that I ought to approve of these things. In fact I am as confident of these things as I am of my belief that the inverse square law correctly describes gravitational force. But that is an argument for another time.

Chapter 6. Conclusion

The way has been long and torturous and it is easy to lose the forest for the trees. We should therefore take stock of where we have been and where we have arrived.

One of the main findings of this dissertation is that a major confounding factor in metaethical arguments is the yoking together of semantic and justificational theories. This has resulted in the overlooking of a very promising candidate for a naturalistically acceptable metaethical theory. In particular it has obscured the possibility of an emotive realism, though not totally as the case of Blackstone shows. This has led such prominent thinkers as Simon Blackburn to try to soften the blow that the sentimentalist theory of justification has upon us. In particular he does it by invoking the deflationary view of truth in support of the view that we need not be threatened by the relativism seemingly inherent in the sentimentalist position. But the deflationism about truth is not enough to shield Blackburn's position from the traditional problem associated with the sentimentalists' theory of justification. Besides this we have seen that the use theory of meaning is questionable in its own right, as I will remind you of shortly.

Copp's view is more promising in that he invokes the distinction between meaning and use and wants to give a metaethical theory that is capable of supporting a naturalistic moral realism. Unfortunately Copp fails to distinguish issues of meaning and justification and so is led to the view that moral predicates stand for or refer to robust moral properties as a way to distinguish his view from an irrealist's view like Blackburn's. But as we have seen, this semantic issue is orthogonal to the real issue between the realist and irrealist which is their theory of justification. So we need not hold that a person performs the speech act of asserting that some object has a moral property as part of our analysis of moral language. The real issue, or the way to get at the real issue, is as I have argued, whether or not two people who disagree over some fundamental moral truth, like that murder is morally wrong, can in some sense both be right. The realist claims that they cannot while the irrealist claims that they can. Given this we are free to construct a theory that holds that the primary purpose of moral utterances is to express a moral emotion without having to include the assertion that some person or act has a moral property.

But we do learn from Copp that there is a plausible argument that in moral discourse we express two mental states. He argues that we express the belief that the act in question has a robust moral property and implicate via conventional Gricean apparatus the non-cognitive mental state that he calls norm-acceptance. I interpret this be something akin to having a feeling of approval towards the norm. But if we take Copp's account seriously there is the danger of being committed to the view that when I say that suicide bombing is an evil act I express my belief that suicide bombing has the robust moral property of being wrong and that entails that I approve of the norm from which this judgment follows. If this is the case it will turn out that whether 'suicide bombing is evil' is true or false depends on whether or not I really do accept the norm from which the judgment follows. We could still keep our realist theory of justification if we are willing to accept a rather strange kind of error theory but barring that the lesson is that we should not rely on implicature.

So, after taking stock of recent prominent metaethical theories from the standpoint of desiring a naturalistically minded morally realist perspective we see that there is room to improve. I then turned to working out the emotive realist alternative. To obtain a complete

metaethical theory we need at least three things. In the first instance we need an account of the justification of moral judgments. At the distinctly metaethical level this will involve the debate between the realist and the irrealist, between the relativist and the objevtivist. Once we opt for moral realism several specific theories of justification present themselves. That is as far as the strictly metaethical enterprise goes. Deciding which realist account of justification is the best is itself a normative question. The second thing that we require is an account of the nature of moral judgments and how, if at all, they are capable of being justified. Once we have settled on a theory of justification and the nature of moral judgments we then need a theory of moral language in which we express those moral judgments. There are, of course, other requirements for a satisfactory metaethical theory, like an account of how we come to know the moral truths, if there are any.

I have argued that a moral judgment consists of two mental states. One is a moral emotion, like approbation or condemnation; the other is the belief that the emotion is the correct one to have. To fill in the details we need a P-semantic theory. That is, a theory of how these mental states come to have the meaning that they do. I have

argued that one promising naturalistic account of P-semantics is the causal theory of reference as advocated by Kripke and refined by Devitt. Devitt's account is a molecular account in the sense that it holds that some mental tokens get their meaning from direct causal relations to the environment whereas other mental tokens get their meaning from inferential relations between tokens.

Accordingly, I argued that the moral emotions are a species of propositional attitudes and consist in a mental attitude held towards some representational content. These contents are determined via the causal theory of reference for names and natural kind terms and via inferential relations for others. In particular some moral terms are defined via inferential relations between other moral terms. We express these mental states via reflexive intentions. The mental attitudes themselves, as I have elsewhere argued, can be modeled as distinctive qualitative attitudes held towards the representational content.

This account of the nature of moral judgment easily allows us to offer a realist theory of justification. The belief, recall, is the belief that the moral emotion is the correct one to have. This belief will be true if the

emotion is in fact the correct one to have towards the act/person in question. And whether it is in fact the correct emotion to have with respect to the act/person in question depends on our specific theory of justification. So if one is a utilitarian one will think that the feeling of approbation ought to be held towards those actions which maximize utility. In this way the account I have offered is neutral with respect to theories of justification.

Which particular theory of justification is true is a normative issue and is strictly speaking a separate endeavor from the one that I have been engaged in. I have been primarily concerned to construct a semantic theory capable of supporting a naturalistic realist theory of justification. Having done so I consider the project a success. Even so, I cannot fully restrain myself and I briefly suggested a constructivist account of justification on which we construct moral principles via some appropriate methodology. What we get from this is a naturalistic grounding for negative moral judgments. Given basic features about sentient beings we are able to see that some actions are absolutely and universally wrong. The positive side of things is more complicated. Here things depend on interests which are rarely shared, as

opposed to the others which are grounded in interests which are constitutive of rational sentient creatures. It follows from this that some of the positive moral duties will be relative to a particular culture and their interests. Others will be true of all rational creatures.

Now it could be the case that this story we just told about the P-semantics of moral judgment turned out to be the correct L-semantic account of what the sentence types we utter mean. If that were the case then the L-semantic meaning of 'Saddam Hussein was evil' would be something like 'I morally condemn Saddamm Hussein; that is the correct way to feel about Saddam Hussein'; where the first part is meant to express the moral emotion of moral condemnation. This kind of theory would hold that the meaning of the semantic type is given by the mental state that the semantic type is used to express.

I have given two kinds of argument against doing this. The first, given in chapter 3 and aimed at Blackburn's use theory of meaning, shows that in order to explain the way that English is used we need to posit a separate meaning independent of the use. This meaning needs to be something that the speaker knows, and we have independent reason to think that this is not simply its reference. What the speaker needs to know will vary

depending on the word. In some cases the person will need to know how the word is used, others what description it is associated with.

The second argument comes from problems in formal semantics. When we try to model names, for instance, in natural language on the model that works for thoughts we run into several well known problems. These range from the usual ones about belief contexts, informative identity statements, and the problem of necessary existence, which is perhaps the most troubling from the naturalistic perspective. Since logic offers us a test of various semantic theories we can see that the problem results from the claim that the L-semantics of names is the same as their P-semantics. This shows that these two theories cannot be reduced to each other, in either direction.

In the context of a metaethical theory this amounts to the claim that we cannot endorse the above claim about the meaning of moral sentences. We should not endorse the use theory of meaning. But if the meaning of 'Saddam Hussein is evil' is not given by 'I morally condemn Saddam Hussein; that is the correct way to feel about him' then what is its meaning?

The natural candidate for L-semantic meaning has already shown itself. It is what the speaker needs to know

in order to count as using the word correctly. In the case of names it turns out that this is simply the nominal description that mentions the name. In the large majority of other words it ends up being roughly the dictionary meaning of the word.

I submit that the kind of semantical theory argued for here in this dissertation is the best candidate for a naturalistic moral realism. It allows us to keep the insights of most metaethical theorizing and jettison the baggage.

Bibliography [50]

Aristotle (350 bce/1984). The Nicomachean Ethics. <u>The Complete Works of Aristotle:</u>
<u>The revised Oxford Translation</u>
<u>Volume Two</u>. J. Barnes, Princeton University Press.

Armstrong, D. (1978). "Naturalism, Materialism, and First Philosophy." *Philosophia* **8**(2-3): 261-276.

Armstrong, D. M. (1968). <u>A Materialist Theory of Mind</u>. London, Routledge.

Arnauld, A. and P. Nicole (1662/1861). <u>Logic: The Art of Thinking</u>. London, Hamilton, Adams, and Company.

Austin, J. (1962). <u>How to Do Things with Words</u>. Cambridge, Harvard University Press.

Ayer, A. J. (1936/1946). <u>Language, Truth, & Logic</u>, Dover Publications Inc.

Bach, K. (1987). <u>Thought and Reference</u>, Oxford University Press.

Bach, K. (1999). "The Myth of Conventional Implicature." *Linguistics and Philosophy* **22**(4).

Bach, K. (2001). "Speaking Loosely: Sentence Nonliterality." *Midwest Studies in Philosophy* **25: Figurative Language**: 249-263.

Bach, K. (2002). "Giorgione was so called because of his Name." *Philosophical Perspectives* **16**: 73-103.

Bach, K. (forthcoming). "The Excluded Middle: Semantic Minimalism without Minimal Propositions." *Philosophy and Phenomenological Research*.

Bach, K. and R. Harnish (1979). <u>Linguistic Communication and Speech Acts</u>, The MIT Press.

---

[50] All dates are in the Common Era unless otherwise noted as Before the Common Era by 'bce'.

Baumrin, B. H. (1968). "Is There a Naturalistic Fallacy?" *American Philosophical Quarterly* **5**.

Baumrin, B. H. (1986). "Moral Blindness." *Metaphilosophy* **17**(4): 205-213.

Baumrin, S. B. (2004). "The Shoes of the Other." *The Philosophical Forum* **XXXV**(4): 397-410.

Berkeley, G. (1710/1998). A Treatise Concerning the Principles of Human Knowledge, Oxford University Press.

Blackburn, S. (1984). Spreading the Word, Oxford University Press.

Blackburn, S. (1988). "How to Be an Ethical Antirealist." *Midwest Studies in Philosophy* **12**: 361-375.

Blackburn, S. (1998). Ruling Passions, Oxford University Press.

Blackburn, S. (1999). "Is Objective Moral Justification Possible on a Quasi-realist Foundation?" *Inquiry* **42**: 213-228.

Blackstone, W. T. (1958). "Objective Emotivism." *Journal of Philosophy* **55**(24): 1054-1062.

Brown, R. (2006). "What is a Brain State?" *Philosophical Psychology* **19**(6): 729-742.

Brown, R. (2007). "The Mark of the Mental." *Southwest Philosophy Review* **23**(1): 117-124.

Carnap, R. (1935). Philosophy and Logical Syntax, AMS Press. Inc.

Carnap, R. (1956). Meaning and Necessity. Chicago, University of Chicago Press.

Copp, D. (2001). "Realist-Expressivism: A Neglected Option for Moral Realism." *Social Philosophy and Policy* **18**: 1-43.

Davidson, D. (1967). "Truth and Meaning." *Synthese* **17**.

Devitt, M. (1974). "Singular Terms." *Journal of Philosophy* **71**: 183-205.

Devitt, M. (1981a). Designation. New York, Columbia University Press.

Devitt, M. (1996). Coming to our Senses: a Naturalistic Program for Semantic Localism, Cambridge University Press.

Devitt, M. (1997). Precis of Coming to Our Senses: A Naturalistic Program for Semantic Localism. Philosophical Issues. E. Villanueva. Atascadero, Ridgeview Publishing Company. 8: 325-49.

Devitt, M. (2002). "Moral Realism: A Naturalistic Perspective." *Croation Journal of Philosophy* 2: 1-15.

Devitt, M. (forthcoming). No Place for the A Priori. New Views of the A Priori in Physical Theory. M. J. Shaffer and M. Weber, Rodooi Press.

Devitt, M. and K. Sterelny (1999). Language and Reality, The MIT Press.

Edwards, P. (1955). The Logic of Moral Discourse. Glencoe, Illinois, The Free Press.

Feigl, H. (1967). The "Mental" and the "Physical": The Essay and a Postscript. Minneapolis, University of Minnesota Press.

Field, H. (1980). Science without Numbers: a Defense of Nominalism, Princeton University Press.

Fodor, J. (1991). Propositional Attitudes. The Nature of Mind. D. Rosenthal, Oxford.

Fodor, J. (1998). Concepts: Where Cognitive Science Went Wrong, Oxford University Press.

Foot, P. (1958). "Moral Beliefs." *Proceedings of the Aristotelian Society* 59: 83-104.

Frankena, W. (1955). "Hutcheson's Moral Sense Theory." *Journal of the History of Ideas* 16(3): 356-375.

Frankena, W. (1963). Ethics. Englewood Cliffs, New Jersey, Prentice Hall.

Geach, P. (1965). "Assertion." *Philosophical Review* **74**: 449-465.

Gibbard, A. (1986). "An Expressivistic Theory of Normative Discourse." *Ethics* **96**: 472-485.

Gibbard, A. (1990). Wise Choices, Apt Feelings: A Theory of Normative Judgment, Harvard University Press.

Gill, M. B. (2007). "Moral Rationalism vs. Moral Sentimentalism: Is Morality More Like Math or Beauty?" *Philosophy Compass* **2**(1): 16-30.

Goldman, A. I. (1993). "The Psychology of Folk Psychology". Readings in the Philosophy of Cognitive Science. A. I. Goldman, The M.I.T Press.

Greene, J. and J. Haidt (2002). "How (and Where) does Moral Judgment Work?" *Trends in Cognitive Sciences* **6**(12): 517-523.

Grice, H. P. (1957). "Meaning." *The Philosophical Review* **66**: 377-388.

Grice, H. P. (1975). "Logic and Conversation". The Logic of Grammar. D. Davidson and G. Harmen. Encino, Ca, Dickenson Publishing Co.

Grice, P. and P. F. Strawson (1957). "In Defense of a Dogma." *The Philosophical Review* **65**: 141-158.

Hare, R. M. (1952). The Language of Morals, Oxford University Press.

Held, V. (1996). Whose Agenda? Ethics versus Cognitive Science. Mind and Morals: Essays on Ethics and Cognitive Science. L. May, M. Friedman and A. Clark. Cambridge, The MIT Press.

Horwhich, P. (2005). Reflections on Meaning, Oxford University Press.

Hussain, N. and N. Shah (2006). Misunderstanding Metaethics: Korsgaard's Rejection of Realism. Oxford Studies in Metaethics. R. Shafer-Landau, Oxford: Clarendon Press: 265-94.

Jackson, F. and P. Petit (1998). "A Problem for Expressivism." *Analysis* **58**(4): 239-252.

Johnston, M. (1998). Ethics. A Companion to Cognitive Science. W. Bechtel and G. Graham, Blackwell Publishers.

Kant, I. (1785/1997). Groundwork of the Metaphysics of Morals, Cambridge University Press.

Kripke, S. (1963). "Semantical Considerations on Modal Logic." *Acta Philosophica Fennica* **16**: 83-94.

Kripke, S. (1972). Naming and Necessity, Harvard University Press.

Kripke, S. (1977/1991). Speaker's Reference and Semantic Reference. Pragmatics. S. Davies, Oxford University Press.

Locke, J. (1690/1975). An Essay concerning Human Understanding. P. H. Nidditch. New York, Clarendon Press: Oxford.

Menzel, C. (2005). "Actualism." *The Stanford Encyclopedia of Philosophy* **Summer 2007 Edition**(Edward N. Zalta (ed.)): URL=http://plato.stanford.edu/archives/sum2007/entries/actualism/.

Nowell-Smith, P. H. (1954). Ethics, Pelican Books.

Place, U. T. (2004). Is Consciousness a Brain Process? Identifying the Mind: Selected Papers of U. T. Place. G. Graham and E. Valentine, Oxford University Press.

Polger, T. W. (2004). Natural Minds, The MIT Press.

Prior, A. N. (1956). "Modality and Quantification in S5." *Journal of Symbolic Logic* **21**(60-2).

Putnam, H. (1967/1991). The Nature of Mental States. The Nature of Mind. D. Rosenthal, Oxford University Press.

Quine, W. V. O. (1940). Mathematical Logic, Harvard University Press.

Quine, W. V. O. (1966/1976). The Vagaries of Definition. The Ways of Paradox and other essays (revised and enlarged edition), Harvard University Press.

Rawls, J. (1980). "Kantian Constructivism in Moral Theory." *The Journal of Philosophy* **77**(9): 515-572.

Robinson, R. (1950). Definitions. London, Oxford University Press.

Ross, S. (1991). "The Nature of Moral Facts." *The Philosophical Forum* **XXII**(3): 243-268.

Ross, S. (1998). "The Nature and Limit of Moral Objectivism." *The Philosophical Forum* **XXIX**(2).

Ross, W. D. (1939). Foundations of Ethics, Oxford University Press.

Sayre-McCord, G. (1986). "The Many Moral Realisms." *Southern Journal of Philosophy* **suppl. vol. 24**: 1-22.

Searle, J. (1964). "How to Derive 'Ought' from 'Is'." *The Philosophical Review*.

Searle, J. (1969/2001). "The Structure of Illocutionary Acts". The Philosophy of Language. A. P. Martinich, Oxford.

Sellars, W. (1956). Empiricism and the Philosophy of Mind. Minnesota Studies in the Philosophy of Science. H. Feigl and M. Scriven, University of Minnesota Press. **Volume 1:The Foundations of Science and the Concepts of Psychology and Psychoanalysis**: 253-329.

Shafer-Landau, R. (2003). Moral Realism: A Defence, Oxford University Press.

Sidgwick, H. (1874/1952). The Methods of Ethics. Reading in Ethical Theory. Sellars and Hospers, Appleton-Century-Crofts, Inc.

Smart, J. J. C. (1991). Sensations and Brain Processes. The Nature of Mind. D. M. Rosenthal, Oxford University Press.

Soames, S. (2006). "Philosophical Analysis". Encyclopedia of Philosophy, Macmillan Publishing Co.

Soames, S. (forthcoming). "Analytic Philosophy in America". The Oxford Handbook of American Philosophy. C. Misak, Oxford University Press.

Stevenson, C. L. (1937/1963). "The Emotive Meaning of Ethical Terms". Facts and Values: Studies in Ethical Analysis, Yale University Press.

Stevenson, C. L. (1944). Ethics and Language, Yale University Press.

Stevenson, C. L. (1963). Facts and Values: Studies in Ethical Analysis, Yale University Press.

Strawson, P. F. (1950/1985). On Referring. The Philosophy of Language. A. p. Martinich, Oxford University Press.

Strawson, P. F. (1964/1994). "Intention and Convention in Speech Acts". Basic Topics in the Philosophy of Language. R. Harnish, Prentice Hall.

Toulmin, S. (1950/1960). An Examination of the Place of Reason in Ethics, Cambridge University Press.

Urmson, J. O. (1950). "On Grading." *Mind* **59**(234): 145-169.

Urmson, J. O. (1968). The Emotive Theory of Ethics, Oxford University Press.

Vendler, Z. (1972). Res Cogitans: an Essay in Rational Psychology, Cornell University Press.

Warnock, M. (1960). Ethics Since 1900, Oxford University Press.

Williamson, T. (2002). Necessary Existents. Logic, Thought, and Language. A. O'Hear. Cambridge, Cambridge University Press.

Wright, C. (1992). Truth and Objectivity, Harvard University Press.