Ordinary Language Philosophy

Examining What We Should Say When

Samuel Williams

Thesis submitted in partial fulfilment of the requirements for the degree of
Bachelor of Philosophy in Philosophy

Faculty of Philosophy
University of Oxford
Trinity 2021

For Grandma and Nannie.

Acknowledgements

I would like to thank my supervisor, Anita Avramides, for the interest she showed in the project of this thesis, for all her help in realizing it, and for helping me to understand how the importance of the ideas here expressed extends beyond the scope of the thesis. Thanks are also due to Ben Campbell: for his comments and conversation, without which I could not have gotten clear on what it was I wanted to say here, and for his friendship, which has made the last year such a happy one. Finally, I am grateful to my family, and particularly to my parents, for their love and support throughout the process of writing this thesis, and in all the years before.

Abstract

This thesis takes its cue from J.L. Austin ('A Plea for Excuses'): philosophy which proceeds from ordinary language (OLP) is an exercise in 'examining *what we should say when*, and so why and what we should mean by it' (*PFE*: 129). It argues that the method of OLP (The Method) is properly understood when we view language through a Cavellian lens. Linguistic competence emerges not as tacit knowledge of a system of strict rules, but as a kind of skill: a fluid ability to leverage one's attunement to others in order to make oneself understood. In light of this, The Method ought to be understood as a self-conscious deployment of this skill, an invitation to project our words into novel contexts. Such an invitation is illuminating, because it asks us to attend, carefully and critically, to *the world*, and to follow out the holistic conceptual connections revealed by our projection of a word into a given context. Because this attention is careful and critical, OLP need not be conservative; neither need The Method precipitate us into a gross relativism. Indeed, a vision of language based on attunement offers a realistic, and deeply sensitive, picture of what is involved in responsibly negotiating our differences.

Contents

Introduction

The Method of OLP

> [W]ords are our tools, and, as a minimum, we should use clean tools: we should know what we mean and what we do not, and we must forearm ourselves against the traps that language sets us.
>
> (J. L. Austin, *PFE*: 129-30)

This thesis is concerned with ordinary language philosophy. I will adopt that label for what I intend to defend, despite the fact that there are certainly questions to be asked about to what extent the supposed practitioners of ordinary language philosophy (henceforth, OLP), and the various philosophers whose work has since taken up ideas from post-war Oxford philosophy, constitute a unified "school" of philosophy, either doctrinally or methodologically.

The label has accrued some not-so-felicitous associations. (G. J. Warnock (1998: 147) suggests that '[t]he label "ordinary language philosophy" was more often used by the enemies than by the alleged practitioners of what it

was intended to designate.') Despite this, I am resigned to it, for two reasons. The first is that, whatever its origins or baggage, the label 'OLP' does not carry too much risk of misleading the reader, provided they consent to taking it at face-value. My aim here is to philosophize about ordinary language; that is, about *natural* (as opposed to artificial or formal) languages, as they are put to ordinary uses in human lives. And I hold the view that it is at least beneficial to orient our philosophy towards such ordinary uses of language; that doing so is likely to be illuminating even where the subject in question is not natural language itself, but the things about which we speak in natural language.

The second reason that I am resigned to call this thesis an exercise in OLP is that the label has stuck. And it has adhered particularly strongly to the work of J. L. Austin, work which provides the impetus for what I want to achieve. I will take as my starting point what Anthony Flew (1986: 95) identifies (however contentiously) as OLP's 'most explicit methodological statement': Austin's 'A Plea for Excuses' (PFE).

> [The study of excuses] is an attractive subject methodologically, at least if we are to proceed from 'ordinary language', that is, by examining *what we should say when*, and so why and what we should mean by it.
>
> (PFE: 129, emphasis in original)

I will take the method of examining what we should say when as definitive of what I mean by OLP. Because of its centrality to my understanding of OLP, I will refer to it as The Method. My aim is to understand the character

and philosophical value of The Method: on what is one drawing when one employs it, and what is the importance of its deliverances?

0.1 Suspicion of The Method

> Thankfully, the banalities of ordinary language philosophy are no longer with us, done to death by a thousand miserable attempts to solve philosophical problems by careful attention to upper-middle-class English usage.
>
> (Papineau, 2006: 20)

OLP is viewed with suspicion, I suggest, because The Method has not been properly understood among mainstream Anglo-American philosophers. Wikipedia's (2020) characterization of OLP is representative of the prevailing impression of OLP:

> Ordinary language philosophy is a philosophical methodology that sees traditional philosophical problems as rooted in misunderstandings philosophers develop by distorting or forgetting what words actually mean in everyday use.

OLP is often taken to recall ordinary uses of natural language in order to dismiss, out of hand, any claim which conflicts with these uses. It is true that some philosophers have tried to dismiss philosophical claims on the grounds that they conflict with ordinary usage:

> [I]f a philosophical statement is paradoxical, that is because it asserts the impropriety of an ordinary form of speech. [...] But

> it is not possible for an ordinary form of speech to be improper. That is to say, ordinary language is correct language.
>
> (Malcolm, 1981: 18)

For example, if there are situations in which it is perfectly correct, by ordinary standards, to say 'I see my pen', no philosophical argument could show that one does not *really* see one's pen (but instead, e.g., a "sense-datum"). Such an argument would have to *misuse* the natural language expression 'see'.

A notable manifestation of this sort of appeal is the Paradigm Case Argument (PCA), to which Malcolm in the same paper has recourse (*loc. cit.*). More recently, PCA has been endorsed in Oswald Hanfling's (2000) defence of OLP. The paradigm PCA opposes universal denials of existence/occurrence: that any things of type F exist, or that events of type G ever occur. As Anthony Flew puts it:

> Crudely: if there is any word the meaning of which can be taught by reference to paradigm cases, then no argument whatever could ever prove that there are *no cases whatever* of whatever it is. Thus, since the meaning of 'of his own freewill' can be taught by reference to such paradigm cases as that in which a man, under no social pressure, marries the girl he wants to marry (how else *could* it be taught?): it cannot be right, on any grounds whatsoever, to say that no one *ever* acts of his own freewill.
>
> (1955: 35, emphasis in original)

Such arguments have a strong association with the label 'OLP'.[1] This creates

[1]The major source of the association of PCA with OLP is probably the influence of

the suspicion that The Method amounts to little more than inferences of the following form:

> We (are inclined to) say that p; therefore p.

This speedy inference is backed up by an argument of the following sort: the meaning of an expression is fixed by what we, the members of a given linguistic community, do with it. Therefore, if a competent speaker of a language L describes a (fully specified) state of affairs S by saying that p, then that must be the *correct* way, in L, to describe S. That the speakers of L use the expression 'p' to describe S fixes the meaning of 'p' in such a way that it must be correct to assert 'p' of S. If someone refuses to assert 'p' of S, then there are three options. (a) They are using 'p' differently to how we do; to use Wittgensteinian terminology, they are playing a different language-game from us, with different rules. (b) They are using 'p' incorrectly; they are playing our language-game, but they have broken the rules. (c) They are incoherently trying to use 'p' in some "absolute" sense, outside of any particular language-game (it is only the rules of a language-game that furnish an expression with sense). In none of these cases does the refusal to apply 'p' to S contradict our claim that p.[2]

If this were indeed what The Method amounted to, then the oft-repeated

---

Ernest Gellner, whose polemical book *Words and Things* (1959), though rarely discussed today, was explosive in its own time. (The controversy surrounding *Words and Things* is recorded in Uschanov (2006).) Gellner claims that PCA is "essential, fundamental and pervasive in Linguistic Philosophy [his preferred label for OLP]" (36).

[2]I would not endorse this particular appropriation of Wittgenstein's terminology. There are, however, several passages in his later work from which one might garner this impression of his use of 'language-game'. Cf. *Z* §320; *PI* §§22, 44, 71, 116, 136, 156, 261, 288, 632, 656; *PPF* §§31, 34, 52, 161, 332; *OC* §§3, 18, 21, 24, 65, 371, 374, 391, 393, 411, 446, 480, 558, 579, 596, 599.

criticism that OLP is intellectually conservative, slavishly adhering to "common sense" or preconceived opinion, might be fair. (This criticism has sometimes appeared as the particular charge of *political* conservatism; most notably in Gellner, 1959.) Our *de facto* ways of speaking — and whatever uncritical attitudes towards the world, to public, political or moral life, those ways of speaking may harbour — seem to be made immune from rational criticism.

But this is not at all what The Method involves. The attention it asks us to bring to ordinary uses of language is a deeply critical attention. By its very nature, it demands that we take into account perspectives other than our own. The Method asks us to enter into a collective negotiation as to what it would mean to say that p. Undergoing this negotiation is, I will suggest, a philosophically valuable process, capable of deepening our understanding of the world in which we live our lives with others.

0.2 Thesis outline

The structure of this thesis is as follows.

In order to understand The Method, we must understand the character of the natural language with which it operates. In Chapter 1, I draw on the now-familiar rule-following investigation in Wittgenstein's *Philosophical Investigations* in order to argue for a particular vision of natural language. This vision is that which I take to be expressed in the following passage from Stanley Cavell:

> We learn and teach words in certain contexts, and then we are

> expected, and expect others, to be able to project them into further contexts. Nothing insures that this projection will take place (in particular, not the grasping of universals nor the grasping of books of rules), just as nothing insures that we will make, and understand, the same projections. That on the whole we do is a matter of our sharing routes of interest and feeling, modes of response, senses of humor and of significance and of fulfillment, of what is outrageous, of what is similar to what else, what a rebuke, what forgiveness, of when an utterance is an assertion, when an appeal, when an explanation — all the whirl of organism Wittgenstein calls "forms of life."
>
> (*AWP*: 52)

Anglo-American philosophy has tended to idealize proficiency with a natural language as knowledge of a set of rules, knowledge which allows one to compose utterances whose meanings are determined according to these rules. In the *Investigations*, as refracted through Cavell, linguistic proficiency appears in a rather different light. Language-use appears as a fluid, evolving skill analogous to that involved when a proficient musician plays their instrument. When a jazz musician performs a standard — 'Autumn Leaves', say — they solo over an established chord progression. But that progression is there for them to *make something of*, for them to (as I shall say) *leverage* or *exploit*. Similarly, when one makes use of a language, one exploits one's attunement to the other. ('Attunement' is a key (Cavellian) notion for this thesis. My inheritance of this notion is explained in §1.3. I'm not completely

happy with the (admittedly rather ugly) terminology of 'leveraging' or 'exploitating' attunement. But I haven't found a better way to express the way in which one does not merely use language, but *makes use* of it.) Crucially, there are no hard-and-fast rules for how one must leverage one's attunement to the other. Success can be the result of — can even rely on — a creative appropriation of this task.

Seeing linguistic proficiency as a skill of this sort has important consequences for our understanding of The Method. In Chapter 2, I present The Method as an invitation to *project* our words into new contexts. ('Projection' is another key Cavellian notion (§1.4).) This is a self-conscious, reflective employment of our linguistic mastery. Far from being intellectually conservative, The Method invites us to be deeply attentive and responsive to the realities of particular situations, to see how we might navigate these situations with our language, and to allow the insight thereby gained to change our preconceived opinions and *de facto* ways of speaking.

Chapter 3 concerns *failures* of attunement. It addresses the worry that relativism will result from the fact that we are only imperfectly attuned to our fellows. I argue that this worry rests on the functional conception of meaning — the idea that, in speaking, the words one utters (plus features of context) determine a unique content for what one has said, which must be grasped by anyone who is to understand you. Despite its dominance in Anglo-American philosophy, I suggest that a vision of language based on attunement has good reason to eschew the functional conception. I suggest that this provides a more realistic understanding of deep disagreement, free from the distorting influence of idealizations about language and rationality.

I should state explicitly that my aim here is not exegesis. I will frequently cite Austin, Wittgenstein, and Cavell, but I make no claim for the faithfulness of my understanding of those texts. Both Austin and Wittgenstein placed emphasis on the importance of attending to the uses to which natural language is put in ordinary situations, and this is an emphasis which Cavell's work takes up. My intention is to self-consciously *inherit* those ideas, on the understandings that to me are most interesting and seem most fertile. Inevitably, the work of Wittgenstein will appear here in a certain light; but what is presented here is Wittgenstein twice removed: once, as understood by Cavell, and again, as understood by me. Similarly, the inheritance of The Method in light of the conception of language I find in Wittgenstein and Cavell has the effect of assimilating Austin's work to theirs. At least here, the justification for that assimilation rests not on textual evidence, but on its service to the project of presenting The Method — and so, OLP — in the best possible light.

OLP is often considered dead. David Papineau (quoted above) is not alone in expressing relief at its demise. But, as Crary & de Lara note, OLP has refused to remain buried: 'OLP has shown up in recent debates in diverse areas: from philosophy of language to philosophical accounts of pornography, from epistemology to gender and queer studies, from ethics to literary theory, and from philosophical theology to discussions about the authority and critical power of the interpretive methods of the social sciences' (2019: 318). My hope is that this thesis, in drawing together what I take to be the most important features of the OLP orientation to language, will allow it to be seen with fresh eyes. It deserves to be seen as a live alternative

to established philosophical methodologies. Properly understood, the vision of language which underlies its Method is liberating. It offers us a thoughtful, sensitive picture of what is involved in responsibly negotiating our differences.

Chapter 1

Understanding Attunement

As indicated in the Introduction, The Method is a deeply critical approach to what we are in fact inclined to say. Note the wording of The Method as it appeared in *PFE*:

> [W]e are to proceed from 'ordinary language', that is, by examining *what we* ***should*** *say when*, and so why and what we should mean by it.
>
> (129, Austin's italics, my boldface)

If we read this 'should' as the normative '*ought*', rather than the factual 'would', we begin to head in the right direction. The question is not about what we would usually say about some situation. The question is rather: what *ought* we to say? What would be the apt way to describe a given situation?

But the way in which people in fact use their words is, of course, not irrelevant to what one ought to say. Words are not magical. 'Pusillanimous'

does not have the meaning it does independently of the way that a certain community — that is, English-speakers — *take up* that written shape (or uttered sound), and put it to particular uses. Words are *ours*, and they bear the meanings they do because of the ways in which we use them. It is because they have such meanings that there are things one *ought* (and things one ought not) to say about a given situation: it is because 'red' means something different to 'blue' that one ought to describe postboxes as being red, and not as being blue.

If we are to understand The Method, it is essential that we understand in what sense there are things one, as a competent speaker of a language faced with an opportunity to make use of that skill, ought or ought not to say, and how this normativity relates to *de facto* usage. That is the overarching concern of this chapter; we want to know: what is the character of the normativity that natural languages exhibit?

Much of this chapter covers well-travelled territory. In order to motivate the vision of language which informs the rest of this thesis, I turn to Wittgenstein's much-scrutinized discussion of rule-following in the *Philosophical Investigations* (*PI*, §§138-242). In §1.1 of this chapter, I outline a certain tempting picture of language: as a practice in which extant usage sets the standard of correctness for future usage by establishing a privileged pattern which future uses must continue. But, if we think about the normativity of language in this way, as primarily a matter of faithfulness to some standard implicit in extant usage, we fall into confusion: despite the fact that we evidently *can* identify given usages as incorrect, it becomes mysterious how *any* usage could count as such.

Following John McDowell (1984; 2000), I take this to be an opportunity to step back, to reconsider the phenomenon of language-use. In light of this paradox, we forego the picture of normativity — of language-use as answering to "rules-as-rails" — that sustained it (§1.2). Like McDowell, I find an alternative vision of language-use in the passage from Stanley Cavell quoted in the Introduction.

According to this vision, the sharing of a language is not an agreement on which rules to follow in the use of our words. It is instead an *attunement* between speakers, a shared inclination to take the words of others in more-or-less similar ways. 'Attunement' is an essential concept for this thesis; in §1.3 I explain that the harmony to which I am referring as 'attunement' covers a broad spectrum, from biological similarities to rich, shared senses of what it means to say something. Brute biological similarities between us are built upon in training which deepens this attunement. We are trained so that it simply becomes second nature to take certain signs in certain ways.

This training does not bring us all onto the same "rails-to-infinity"; rather, it brings us together in a shared sense of how to inherit our language, how to take it forward into new contexts. These new contexts are approached not with pre-established application-conditions for our words, but with a shared sense of their diverse and plastic expressive possibilities. A vision of language based on attunement presents linguistic mastery as fluid and ever-evolving. §1.4 further explores the emerging vision of language through a discussion of Cavell's notion of *projection*.

The sense in which there are things which one should or shouldn't say in attempting to utilize one's language emerges as more of a pragmatic con-

straint. To violate these norms is not *ipso facto* to produce nonsense; it is just that, if one chooses thus to speak, one runs the risk of not, in fact, being understood.

A (Wittgensteinian) theme in this chapter is the reappropriation of ways of speaking which are in and of themselves unobjectionable, but which become problematic when they are annexed to certain problematic fantasies of what it is that makes these ways of speaking appropriate.[1] In particular, talk of 'linguistic rules' may nevertheless remain applicable, albeit differently understood (§1.5) — the "rules" one learns in being taught a new word are better construed as rules *of thumb*.

1.1 The rule-following "paradox"

Suppose I were teaching my friend, a native French speaker, to speak English. I explain to her that this insect — which she already knows to refer to as 'un papillon' — is called (in English) 'a butterfly'. Later that evening, we watch as a moth flits about around the porch light. My friend says, 'If that butterfly isn't careful, it will get burned.'

My friend has made a mistake. Let us suppose that she is well-versed in different insect species; she could tell me the scientific name for this species of moth, and knows that it is biologically rather different from the butterfly seen earlier. She has not mistaken a moth for a butterfly. Rather, she has

[1] See particularly Wittgenstein's reappropriation of the traditional notion of *essence* at *PI* §371. This is one way of understanding what he means at *PI* §116: 'What *we* do is to bring words back from their metaphysical to their everyday use.' (Emphasis in original.) It combats 'a tendency to sublimate the logic of our language' (*PI* §38). Stephen Mulhall (2004: 83) discusses the way in which the later Wittgenstein is 'transfiguring every term of [traditional philosophy's] self-understanding'.

made a mistake in the use of English; a mistake which is understandable, because colloquial French uses 'papillon' to refer both to butterflies and to moths. We might describe the mistake in this way: English usage enshrines a rule for classifying things as butterflies, certain conditions which a thing must meet in order for it to be correct to speak of it as 'a butterfly'. These are the conditions which all and only butterflies in fact meet, and which moths (along with everything else that is not a butterfly) do not. But the rules for the use of 'papillon' in French call for both butterflies and moths to be spoken of using 'papillon'. My friend mistakenly took the rule for 'butterfly' to be that for 'papillon', and so, faced with this new thing to classify, they did so incorrectly.

The idea is this: implicit in our linguistic practice — and sometimes made explicit when we induct others into that practice — are *rules* for the use of expressions. To teach someone a word is to teach them the rules according to which we use it. If one knows these rules, together with syntactic rules for how meaningful sentences are to be constructed, then one will be able to apply them to understand, or to produce, novel utterances.

One uses a word incorrectly when one violates one of these rules, when one fails to speak as in fact we do; to respect, for example, the application-conditions that we have adopted for our common nouns. Future uses of 'butterfly' will be correct only if they continue the established pattern of accepted use. The use of 'butterfly' in English thus far is thought of as constituting the initial segments of a unique series, 'a visible section of rails laid invisibly to infinity' (PI, §218).

At PI §185, Wittgenstein problematizes this way of thinking about nor-

mativity, of thinking of "what we do" as codifiable by rules like: 'apply "butterfly" to all and only insects with features F.' The question is: what is to count as our proceeding to continue according to this rule? An ideal case of rule-following, thus construed, would seem be to be the expansion of an arithmetical series where one must "do the same" each time; the sort of exercise sometimes given to young pupils in mathematics. For example, we ask the pupil to continue the following series:

$$2, 4, 6, 8, ...$$

They are to discern that the series proceeds according to the rule '$n+2$', and continue it accordingly. If they continued it differently, they would not have continued the pattern, and their answer would be wrong.

But even this apparently ideal case is problematic (McDowell, 2000: 41). Suppose that the pupil continued as expected up to 1000, but then wrote:

$$1000, 1004, 1008, 1012$$

How can we insist that this is incorrect when there *is* a rule (indeed, many rules) which would demand this continuation? (One such rule: for $n_i < 1000$, $n_{i+1} = n_i + 2$; for $n_i \geq 1000$, $n_{i+1} = n_i + 4$.) For any finite series of integers is an initial segment of an infinite number of series.

What if we had explicitly told the pupil to proceed according to the rule '$n + 2$'? When she writes the above, we admonish her: 'No! You were supposed to keep adding two at each stage.' But suppose that, instead of acknowledging her mistake, the pupil is confused. She says, 'But I *did* continue to add two.' We might suppose that she has misunderstood the

instruction, 'Add two', taking it to mean the rule that we would express by saying: 'Add two up to 1000, then add four'.

Such a misunderstanding was (and always will be) possible because, in teaching her how to expand series according to rules such as these, we will only ever have given her a finite number of examples. And, again, for any finite series, and for any expansion of that series, there will be *some* rule according to which the expansion is correct. Furthermore, no explanation of our practice — no appeal to some more 'fundamental' rule — is immune from similar misinterpretation. For example, it will not do to insist that she must always do "the same" thing at each point, because every way of going on is the same according to *some* rule. (In the above example, she did go on in the same way, at least by the rule 'Add two up to 1000, then add four'.)

The same risk applies to teaching someone to use a word. The examples one can give in teaching are finite, and so they will exemplify infinitely many possible patterns of continuing use. There is always the risk that the pupil will "count something in" that ought (we judge) to have been excluded. Both butterflies and moths are insects, but if the question of spiders has never come up in the teaching — and *some* question will not have come up — then the pupil might wrongly count spiders among the insects. (To insist that we could explain that insects have exactly six legs is only to push the problem back further; what if I pull two legs off of a spider? This gap, too, can be plugged; but another could then be found.) Or they might "count something out" that ought to have been included; a dead butterfly is still a butterfly, but if the question has never come up, might not our pupil fail to count it as such?

The picture now becomes this one: there are many rails, infinitely many, all with the same initial segment. The pupil must, somehow, latch on to the right one; this is not something that we can be sure that she will do. More concerningly: our previous practice with the rule, itself finite, does not decide between the infinitely many ways that our practice with a word might be extended. And so it looks like nothing in how we have thus far been using, for example, 'green' can rule out someone, after some future time t, referring to the colour of a summer sky as 'green'. This would be compatible with the rule: 'green' is applicable to a if and only if the time is before t and a is green, or if the time is after t and a is blue.[2]

If we view normativity in this way — as faithfulness to how "we" use a given word, where this is construed as an extant pattern of application — then that any use should really be *wrong* seems impossible to understand. For no pattern can, by itself, determine how future "continuations" are to be judged faithful to it. As Wittgenstein puts it:

> [N]o course of action could be determined by a rule, because every course of action can be brought into accord with the rule. [...] [I]f every course of action can be brought into accord with the rule, then it can also be brought into conflict with it. And so there would be neither accord nor conflict here.
>
> (*PI* §201)

[2]My stating of this rule — based on that for Goodman's (1983) 'grue' — assumes that 'green', as used in the metalanguage (the language in which the rule is stated) is not the peculiar language for which the rule is stated.

1.2 Giving up the picture

One way of taking up the "rule-following paradox" of *PI* §201 is exemplified by Kripke (1982). Kripke finds in *PI* an argument to the effect that there really is no such thing as meaning one thing rather than another by one's words.[3] He takes the task of a 'straight' solution to the paradox of §201 — one which does not concede to the sceptic his central claim that no utterance *really* means one thing rather than another (*op. cit.*: 66) — to be to identify some fact (or facts) which constitute one's meaning one thing rather than any other by a given expression (*op. cit.*: 69). For any given future use of that expression, it would follow from these facts whether or not the expression is being used with the same meaning as before. Kripke finds extensive arguments in *PI* for the conclusion that there are no such facts.

Kripke goes on to find in *PI* a 'sceptical solution' to the paradox, an account of how talk of correctness and incorrectness can be appropriate in the absence of such facts. This sceptical solution repudiates a truth-conditional theory of meaning in favour of one based on conditions of warranted assertibility: even if there are no facts to which talk of 'meaning' must answer, the rules of language sometimes license our speaking of what someone means.

I don't have space to address this sceptical solution in detail, but here is a reason for pessimism about its prospects. To associate a word with conditions under which one may use it in assertion is still to invite the question: what are these conditions? That is again to ask for a *pattern* of use, a pattern

[3]Kripke is explicit about the limits of his exegetical aspirations; that he 'should be thought of as expounding neither "Wittgenstein's" argument nor "Kripke's": rather Wittgenstein's argument as it struck Kripke, as it presented a problem for him.' (1982: 5))

which cannot, in and of itself, demand any privileged extension. So that if one has already bitten §201's paradoxical bullet, how one can be using a word — even the word 'meaning' — with the same meaning as before will be no less mysterious when meaning is conceived as embodying not truth-conditions, but assertibility-conditions.[4]

If "solving" the paradox of §201 requires a Kripkean 'straight' solution — a way of showing how only one continuation of extant usage is *really* faithful to what came before — then I have no intention of solving the paradox. Rather, my intention is to dissolve it. We need to take a step back, and question the picture of meaning from which the paradox arose.

In the above, the meaning of an expression appears as that which is captured by the rule according to which we use that expression. The rule has been pictured as a rail, continuing into the future, along which our use of the expression must travel if it is to be correct, remaining faithful to what the expression means. It is as if one were viewing the usage from above, as a line which branches at the present moment into several — indeed, infinitely many — possible continuations:

[4]Furthermore, Kripke's casting of this as a *sceptical* solution relies on a contrast between assertions which are true, because answering to the facts, and assertions which are merely warranted. But a semantics based on assertion conditions, like that adopted by Kripke's Wittgenstein, seems unable to maintain a robust distinction here. For the conditions under which one is warranted in asserting that p is true, or that p corresponds to the facts, will be just those in which one is warranted in asserting that p. (If one takes Wittgenstein to be advocating an assertion-conditional theory of meaning, then *PI* §136 will likely be taken to be making just this point.)

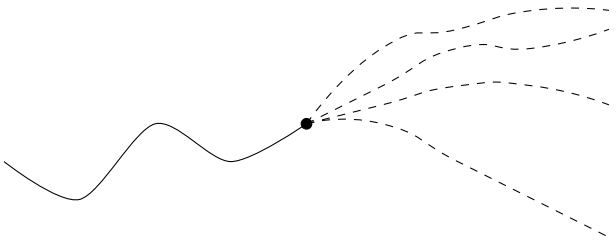

The task set by Kripke for a straight solution to the paradox of §201 would then be to explain how the initial segment can *only* be continued in one of these ways. My intention is to follow McDowell (2000) in questioning the underlying picture of normativity, in which our commerce with the language is viewed "from above".[5] We need to reorient our thinking, to escape the idea that avoiding the paradox of §201 requires us to identify something which allows us to choose, from the infinitely many possible continuations of extant usage, which one is uniquely "correct".

1.2.1 Take your time

> This is how philosophers should salute each other: 'Take your time!'
>
> (*CV*: 91)

I have stated my intention to forego any attempt to view linguistic normativity "from above". But eschewing that perspective is not in itself sufficient to understand the character of the normativity involved in language-use. I want to be careful to differentiate what I will be claiming here from a view expressed by David Egan:

[5] Or, as McDowell (*op. cit.*: 44) puts it, 'from sideways on'.

> [W]hen we see rule-following behaviour as part of a custom, the question of correctness [...] loses its grip on us. The concept of correctness has a role *within* certain language-games [...]. But that means that it makes no sense to inquire about the correctness of the given language-games themselves. [...] [T]alk of facts and correctness have no place outside specific language-games.
>
> (2019: 114-5, emphasis in original)

The normativity we are interested in here is not 'Why ought one to speak as we do?', but 'How ought one to speak, if one is to speak as "we" do?' Put in the language of language-games; it is not 'Which game should we be playing?', but 'How am I to go on, if I want to go on playing *this* game?' We have rejected a particular picture of what an answer to that question demands: that on which a solution must allow us to pick out a particular pattern of usage, "from above". But questions of correctness do not 'lose their grip on us', even from within our practices. (A friend of mine recently wondered aloud whether one could speak of the artificial scenery on a TV set as 'prosthetic'. Well, can one?) What is important to understand is not just that it is *our* concept of correctness which is involved here, but to see its contours in the right light.[6]

[6]I should make it clear that there is much in Egan's book with which I agree. In particular, there are large points of contact between what I want to say and his discussion in ch. 6 of 'Authenticity and Play'. However, even in that chapter, Egan writes that 'the answer to the rule-following paradox is not that we *lack* reasons for following rules in the ways that we do, but rather that questions of right and wrong apply only *within* a practice that we have been trained into' (*op. cit.*: 145, Egan's emphasis). This mislocates the problem, which is a problem for normativity *within* a practice, and as such mislocates the solution, which is to view correct action itself in the proper light.

1.3 Attunement

I have said that we need to reorient our perspective on language-use. To get at this reorientation, consider again the situation of the wayward pupil at §185. This passage is often taken — is taken, for example, by Kripke — to dramatize a more general problem for normativity. We are supposed to generalize from the fact that nothing we can present to the pupil is logically immune from misinterpretation to the idea that there is nothing which determines any particular continuation to be correct. We move from the particular situation, our engagement with the pupil, to an abstracted view of normativity which invites us to view language-use from above. But let's stay with the particular situation: my instruction to the pupil, her mistake, my surprise, the subsequent failures to bring her into conformity. Here, now, a particular failure has occurred. I have failed to make myself understood. In reconceptualizing linguistic normativity, we will take particular failures such as this to be of central importance.

This kind of mutual incomprehensibility is not the norm. If you use a certain expression while talking to an English speaker, you will *in fact* be taken in a certain way. If you say, 'There's a woodpecker in the garden', your audience will look out the window, in the expectation that they will see a bird of a particular kind. Such *de facto* regularities in how utterances are understood by others are, I want to suggest, the source of normativity in the use of a natural language. What we do in teaching someone our language is not to get them to grasp, from finite fragments, uses of words which *must*, on pain of contravening some standard which the community

somehow determines, be extended in a particular way. Rather, we inculcate in them *useful* habits of expression. These habits will, in fact, allow them to succeed in imparting information to others, entreating them, promising or apologizing to them, joking with them, etc.; that is, in being taken as having done these things when they attempt to do so. If they conduct themselves in a fashion which diverges from what the community would expect, their conduct cannot be stamped, as if from above, as "wrong". But if they conduct themselves deviantly, they will frustrate their own purposes. The person who uses 'pterodactyl' where others would use 'butterfly' will often meet with incomprehension; although, if this is her only idiosyncracy, it may be identified by others and overcome. The point is not that there is a privileged pattern, charted out by communal use of 'pterodactyl' that she has failed to extend, but simply that, given how her peers are in fact inclined to understand that word, she will, as a matter of fact, frustrate her own aims to be understood. My aim in the rest of this chapter is to clarify this alternative picture of normativity. It has serious ramifications for our understanding of The Method: of what is involved in asking what one *should* say when.

One speaks 'as we do' if one speaks in a way that is, in fact, and given the way that one's fellows are prepared to understand certain locutions, likely to meet with *success*. With respect to the question of whether stage props can be called prostheses, the question is whether you are likely to succeed in your purposes, a question which depends greatly on context.[7] If you ask someone where the prosthetic bookcase is, will they understand you? A

[7]Context of utterance was a matter of some emphasis for both Austin and Wittgenstein. See Chapter 2, fn. 5.

way of speaking may be called 'incorrect' if it risks *particular* failures of communication in a certain characteristic way.

I am suggesting that we conceive of a shared language as less like convergence on an agreed rail to travel in our use of each individual expression, and more like a conformity in instinct for how to take up our linguistic training in order to achieve our communicative purposes. What is required to induct somebody into a linguistic practice is not to get them to grasp a rule for the use of a word, at least if that rule is considered as a rail-to-infinity. Rather, what is required is to inculcate in that person a "feel" for the (plastic) expressive possibilities of particular expressions among a given community. Adopting a Cavellian term, I shall say that what is required is to *attune* that person to us.

I adopt the term 'attunement' for its suggestion of harmony between those who have been brought to share a language (Cavell, *CR*: 32). In a well-tuned piano, the dispositions of the strings are coordinated so that striking the keys produces consonant sounds. Similarly, the members of a linguistic community exhibit convergent instincts. (The comparison with the piano is apt for two reasons. (1) It emphasizes the contingency of attunement; there are no logical guarantees here. (2) The strings of a piano can be in tune to a greater or lesser degree. As we will see in Chapter 3, that attunement comes in degrees is essential to understanding the OLP perspective on deep disagreement.)

> "So is whatever I do compatible with the rule?" — Let me ask this: what has the expression of a rule — say a signpost — got to do with my actions? What sort of connection obtains here?

— Well, this one, for example: I have been trained to react in a particular way to this sign, and now I do so react to it.

(*PI* §198)

Those human beings who have been inducted into the practice of using signposts are attuned in the ways they are inclined to understand them. At a fork in the road, a signpost is erected. It has a panel into which is carved a depiction of a finger pointing left (☜), and the legend 'HEADINGTON'. For one such as us, it is just a matter of course for us to take this as indicating that, if one wants to get to Headington, then one should walk down the left-hand path.[8]

This convergence of instinct with respect to signposts — the particular way in which those party to this practice are attuned — places a pragmatic constraint on those who would direct them using signposts. Suppose that Headington does indeed lie along the left-hand path from the junction above. If you placed the signpost at that junction so that the finger pointed right (☞), your doing so could not be ruled out as a continuation of "the" extant pattern, if viewed "from above", for it will continue *some* extant pattern. But you will *in fact* mislead others. If people follow their *de facto* inclination — walking down the path in the direction of the arrow — they will not get to Headington. The sign will mislead them.

The instincts in which we are attuned cover a spectrum from innate to learned. Learned instincts like the one above are inculcated on the basis of instincts lying closer to innateness. *PI* §185 exhibits a case in which

[8]Cf. *PI* §238.

induction into a practice — the attempt to attune the pupil to her fellows in a learned instinct of how to continue a series — falters. But, again, this is not the norm. When Wittgenstein introduces the example of the wayward pupil at §185, he makes this suggestion:

> [W]e might perhaps say: this person finds it natural, once given our explanations, to understand our order as *we* would understand the order "Add 2 up to 1000, 4 up to 2000, 6 up to 3000, and so on".
>
> This case would have similarities to that in which it comes naturally to a person to react to the gesture of pointing with the hand by looking in the direction from fingertip to wrist, rather than from wrist to fingertip.

But, of course, a reaction like this would not be usual among our peers. We tend to find the same sorts of continuations natural, to find the same similarities salient, etc.[9] Martin Gustafsson (2005) has made use of an illuminating analogy here. Some people have the (rare) ability to identify the pitch of a sound simply by hearing it. Strike a key on the piano, and a person with 'perfect pitch' will identify it as (e.g.) a B♭. But most of us do not have perfect pitch. Instead, we are taught to identify pitches by comparing them to reference tones (such as that produced by a tuning fork) each time we hear them.

Imagine that *everyone* had perfect pitch. They hear a tone for the first time. Then, when they hear it again, no matter how long the intervening

[9] Cf. *Z* §355.

time, they hear it *as the very same tone* that they had previously heard. Because this identity of pitch strikes them, we could teach them, directly, that *this* note is a B♭. Explanations come to an end here. We say to them: 'Listen: can't you hear that this is the same tone that you heard before?' And we trust that they will hear. In this imagined situation, they *do*.[10]

In our actual practice of teaching people about pitches, explanations still come to an end. But, because most of us do not have perfect pitch, they come to an end with the spontaneous judgment that two notes played within a short interval — the one struck by the tuning fork and the one struck by the piano — are of the same pitch. If someone were not able to hear when two tones have (at least roughly) the same pitch, one could not induct them into our musical practice.[11]

I should be clear about what I mean by saying that "explanations come to an end". The explanations to which I am here referring are explanations of *meaning*, the sorts of explanations which can be given in inducting someone into a practice. I am claiming that, eventually, such explanations must come to rest at spontaneous judgments on which we converge: on a shared capacity to hear two tones as the same pitch. I do not deny that the fact that we converge on these judgments can be given a physico-biological explanation, to do with how frequencies are received by organisms like us. But note that

[10]Cf. *PI* §§217-9.

[11]Perhaps that is not quite true. There may be roundabout ways to attune a person to the rest of us in a given sphere, but doing so will rely on their being attuned to us elsewhere. (I will make use of this idea in Chapter 3.) In the musical case, one could perhaps induct the pupil into our practice via the explanations of pitch in terms of frequency. But that will rely on their having learnt the complex language of physics, an education which, if what I am suggesting in this section is correct, will ultimately come to rest on analogous spontaneous judgments.

there remains an element of arbitrariness here. Those of us immersed in a Western music tradition hear the notes produced twelve keys apart on a piano as the "same" note, with one an octave higher than the other. But they are, of course, different pitches, soundwaves of different frequencies. They are mathematically related (doubling the frequency gives the "same" note an octave above). And there might be a biological explanation for why such frequencies affect us in this way. Nevertheless, these explanations will bottom out in brute regularities: when these neurons are activated, we just do hear two tones as the same note. There is no reason why this regularity should obtain rather than any other, that we should hear notes as recurring when the frequency is doubled, rather than multiplied by 3/2 (the perfect fifth). We just do.[12]

Learning a language rests on spontaneous judgments analogous to the judgment that two tones played within a short interval are of the same pitch; on contingent similarities in how things *strike* us. One sees the various green things (with their various more specific hues) as of the *same* colour. And one naturally perceives regularities in what one's teachers do. As such, one can be brought to follow them.[13]

I want to emphasize the *continuity* between our innate attunement in, for example, the spectrum of light that is visible to us, the way in which we are inclined to group objects according to the wavelengths of light reflected, and our attunement in what we take a given utterance (e.g. 'there's a woodpecker in the garden') to express. This continuity is visible in the way that language

[12]Cf. *BB*: 140-1.
[13]Cf. *PI* §§241-2.

learning begins with *training* rather than explanation; out of this training grows a capacity to attend to explanations.

Initially, induction into one's first linguistic practice must exploit brute biological similarities through training. If my child is just beginning to learn a language, I cannot *explain* to her that: 'This sort of animal is called "a cat"'. She does not yet know what 'animal' means, or what it is for something to be called by an expression e. But I can point to the cat, and say 'kitty'. With my encouragement, the child learns to point at the cat and say 'kitty'. When we play these games, pointing and "naming", she can emulate me because she is biologically similar, she is struck by the world in a way which enables the cat to stand out to her as salient.[14] We can rely on our shared biology — our "first nature" — in order to inculcate further instinctual responses; it becomes, for the child, second nature to manipulate and respond to the spoken sound 'kitty' in certain ways. At some point, this training will progress until I can teach the child new words just by saying something like: 'A kayak is a narrow boat propelled by means of a double-bladed paddle'. Her inclinations have developed to the point that she will, in fact, take this explanation up in a manner that allows her proficiently to use the word 'kayak'. Her susceptibility

[14]The scare quotes around "naming" are intended to register that it is not obvious at what point in this process it is reasonable to say that the child *is* naming the cat, rather than merely playing at doing so; at what point she has actually learnt a word. On this, see Cavell (*EW*). (The wrong way to approach these questions would be to assume that there must be some fact of the matter one way or the other, some definite point at which we can say with absolute certainty, 'Ah! Now she has begun really to *name* the cat.' There will be points at which there will be reasons for saying that she has learned to name the cat, and countervailing reasons for saying that she is as of yet simply parroting. These reasons are as deep as the question goes; there is not, in addition, a unique "fact of the matter" as to which would be the correct thing to say. The question is whether it is *apt*, illuminating, to say that she is naming the cat, or if this is misleading, obscuring differences between her and the mature adult. See §2.1.1 below.)

to training shades into a susceptibility to explanation.

Appreciating this continuity between our attunement in language and the innate, biological attunement on which it builds is essential for understanding this vision of natural language. For that reason, it is worth emphasizing. I assume that if there is any biological tendency towards associating 'kitty' with cats, it is very remote — we might just as well have trained the child to use 'moo-moo' here, although in doing so we would have relied on a natural attunement in the salience that an independent organism like a cat has for a human being. But other customs have a less remote relation to natural inclination. Some signposts show a pointing finger which indicates the way to go. Now, suppose that a person has the propensity to follow a (real, rather than wooden) pointing finger in the direction from fingertip to wrist. Even prior to being taught how to follow such signposts, we might expect such a person to be inclined to follow it in the direction that the (wooden) finger points. It is not *entirely* arbitrary that we have settled on this particular custom, for its adoption was natural given other instincts we possess.

It is worth noting another way in which our learned attunement grows out of contingent biological similarities. For human animals like us, features of our environment press themselves into salience not just because of our particular discriminatory capacities, but because of their relevance to our particular biological endowments and vulnerabilities. (Cavell (*DD*: 42) emphasizes the biological ways in which we are attuned, in which we share a form of life. 'Here the romance of the hand and its apposable thumb comes into play, and of the upright posture and of the eyes set for heaven; but also the specific strength and scale of the human body and of the human senses

and of the human voice.') An antelope is not dangerous to a human in the way a bear is; an elephant is dangerous, but the danger is different, and should be managed differently. It is unsurprising that we should distinguish bears from antelopes, and both from elephants; that this way of classifying should seem to force itself on us. In many modern cultures, there is a strong interest in scientific explanation; in understanding, for example, why this large, furry creature is found in this area, and why it behaves in such-and-such a way. And so these more natural discriminations become the more nuanced, genetic ones suitable for such explanatory purposes.[15]

There is a *continuity* in the development of our attunement to one another. There are brute biological harmonies, like the fact that we all perceive pitch similarly. There are other factors, like those mentioned above, which may influence what kind of further trained inclinations can be made to "take", where a trained conformity is likely to come to rest. (Leibniz's analogy is apt here: the veins in a block of marble influence the shape of the statue that can be carved out of it (1996: 6).) And there are those trained

[15]These remarks have the effect of rescuing the notion of a 'natural kind' from a metaphysical role into which it has been pressed. One response to the rule-following "paradox" of *PI* §201 has been to enlist natural kinds, "utterly objective" joints in the universe, to act as reference magnets for our general terms. But that is work that they are ill-suited to do. These "objective joints" will not serve to pick out, from above, one continuation of our usage as uniquely correct. When atomic number was discovered, we did not *have* to use 'gold' to refer to the element with atomic number 79. We could have maintained a usage based more on surface similarities. Our taking up of our past with the term in the way we did reflects an interest in scientific theorizing. (Of course, our possession of this interest is not entirely arbitrary; it builds upon a distinctly *human* way of relating to — of manipulating — one's environment.)

Does this mean that the meaning of 'gold' changed with the advent of atomic number? To this — and, anticipating later developments in this thesis — I should like to say that we need to get away from the idea that there is a fact of the matter, a uniquely correct answer as to whether we mean the same thing or not. We are still pretty attuned to those people of the past; we find their surviving writings perfectly comprehensible. (On this theme, cf. Diamond, 1999.)

conformities themselves, such as following a signpost as we in fact do.

We should also expect feedback loops in the development of the inclinations in which we are attuned. The ability to perceive comparative pitch will allow one to enjoy an education in Western music. But this education will also feed back into one's ability to perceive pitch, and into one's sense of what is consonant, what is dissonant, what is in tune, etc. First nature informs second nature, which feeds back into first nature. (This idea will be important in Chapter 3, when we discuss how we might approach certain deep disagreements.)

1.4 Projection

Again, I want to avert a misunderstanding of what I am here claiming. The suggestion is not that a uniquely correct extension of usage is picked out, not just by past usage, but by the *de facto* propensities (natural and trained) of one's peers. I am not claiming that we will in fact line up along a single pattern of use, a single rail-to-infinity (even if one could not have been picked out "from above"), so that diverging from that pattern risks confusing our fellows.

A vision of language based on attunement paints a radically different picture. To understand the character of natural language as I want to present it, we need to understand what Cavell means by *projection*. There are three reasons for exploring this important concept. First (§1.4.1), the notion of attunement is further clarified by presenting it as harmony in inclinations to *project* our words into certain contexts. Learning a language is a matter of

developing finely calibrated instincts for projection; that is, *fruitful* instincts, given a desire to communicate with others who have their own trained inclinations to project our common stock of expressions in certain ways. This reveals the way in which taking up one's training in language can be — indeed, is often essentially — a creative, imaginative exercise. Furthermore, it deepens the alternative conception of language we are developing by challenging a traditional notion of literal meaning. Language-use appears as a fluid, responsive *skill*, analogous to dramatic or musical improvization. Secondly (§1.4.2), our discussion of projection will help to show how language is something that grows with us, individually and collectively. Our training in language is continuous, the basis (and possibly the extent) of our attunement constantly changing as we continue to take up our individual and collective histories with words. Thirdly (§1.4.3), it will illustrate the holistic nature of our attunement in language, the way in which conceptual connections are implicated in our inclinations to project our words into certain contexts, and to withhold them from others. (This last aspect is particularly important for understanding The Method.)

1.4.1 Projection and metaphor

Cavell introduces the notion of projection as follows:

> We learn the use of 'feed the kitty', 'feed the lion', 'feed the swans', and one day one of us says 'feed the meter', or 'feed in the film', or 'feed the machine', or 'feed his pride', or 'feed wire',

and we understand, we are not troubled.

(CR: 181)

The importance of this example is this. In order to communicate with others, one must take up and exploit one's training. One must utilize the ways in which one is (hopefully) attuned to others in order to make oneself understood. This is something one can do in all sorts of different and imaginative ways. A parking meter is not an organism which needs sustenance in order to sustain its metabolism. But, in speaking of 'feeding' the meter, one can exploit certain similarities in order to make oneself understood: when one puts coins in the meter, one gives it what it "needs". When one feeds someone's pride, one helps it to *grow*. Cavell is here drawing our attention to the fact that we do, in these ways, succeed in making ourselves understood.

Of course, not just any projection will be intelligible (though I would not wish to say that there is any that could not be made so in *some* context). If I ask someone to feed my *house* while I'm away on holiday, they are unlikely to make sense of what I've said. One must be able to see how feeding a meter is relevantly like feeding a kitten, what similarities are being exploited. This is what I take Cavell to mean by his claim that the 'inner constancy' of a concept is necessary to its being used meaningfully, to being an instrument of communication (CR: 185). What is relied upon here is again our *attunement* to one another, which is manifested in our finding the same lines of projection natural. It is only if we are relatively well-attuned in which projections we will accept that we will be able successfully to communicate with one another.[16]

[16]Referencing PI §§241-2, Cavell (CR: 30-2) apparently uses 'attunement' to mean an agreement in 'criteria': that is, agreement as to what counts as an F for a concept F,

One may be tempted to respond that these projective uses are simply *metaphorical*. But we need to exercise caution here. The intended contrast is, I take it, with literal meaning: one does not *literally*, but only figuratively, feed a parking meter. What are we to make of this notion of literal meaning? Well, certainly, to use the word 'feed' for what we do to the parking meter comes less immediately, less as a matter of course, than to use that word for what we do for swans. This is related to Cavell's claim 'that we do not learn words in all the contexts in which they could be used [...] and that not every context in which a word is used is one in which the word can be learned' (*CR*: 168-9).[17] If you wanted to teach someone what 'feed' means, you would talk about feeding animals, but not feeding parking meters. If the pupil is to get a feel for the expressive possibilities of 'feed', they will have to be introduced to the "normal" cases. But, insofar as there is a pattern to these cases, it is not one which, as it were, extends itself. The pupil has to extend it, to exploit their attunement to others to make themselves understood. They have been taught to speak of feeding a kitten — it will then be up to them to understand others who speak of feeding *swans*, despite the fact that swans won't eat the same food as kittens, you wouldn't present food to a swan in the same way you would to a kitten, and swans, if not fed, will find their own food in a way that a pet kitten may not.

In light of the conception of language which I am defending, we ought to

---

where F is a mundane, non-specialist concept like 'chair'. It is conspicuous that I have avoided speaking of criteria here. I have done so to avoid getting into thick weeds of Wittgenstein exegesis. For our purposes, it will suffice to gloss the sharing of Cavellian criteria — Cavellian attunement — as a shared readiness to project 'chair' into (near enough) the same contexts.

[17] Cf. *BB*: 9-10.

see that a projection like 'feed the meter' differs not in kind from a usage like 'feed the swans', but only in degree. That is, the connection between feeding a kitten and feeding a swan is itself *projective* in character. Our training in language is something which we must *leverage*, something which we must exploit. What is to count as doing what *we* do, in the absence of rules-as-rails, channeling our usage into the future? There will be no hard and fast classifications of utterances which do or do not "count" as continuing what we do. Rather, we make use of our training in particular situations, and do so more or less successfully. 'Feed the meter', in suitable contexts, will meet with success, even if it would be a poor choice by which to induct someone into our practices with 'feed'.[18]

Am I saying there is no distinction between literal and figurative speech? No; but we shall have to accept a spectrum, with a great many cases fitting comfortably into neither camp. A statement can be meant more, or less, literally. Cavell claims that

> what is essential to the projection of a word is that it proceeds, or can be made to proceed, *naturally*; what is essential to a functioning metaphor is that its "transfer" is *unnatural* — it breaks up the established, normal directions of projection.
>
> (*CR*: 181-2)

Here, we should read 'naturally' as turning us back to those inclinations

[18] 'But you can define "feed" in the kitten/swan cases as "to give sustenance to an organism".' 'Give' must be projected if one is to give sustenance to a kitten and also to a swan. This recurrence of the problem of definition resembles that in the rule-following considerations. That is no accident: we are concerned with how to conceptualize language without appealing to a misguided picture of rule-following.

in which we are attuned. As a result of one's training, one will be more inclined to accept some projections, less inclined to accept others. What comes naturally to us covers a spectrum. It comes *very* naturally to say of both a kitten and a swan that one can feed them, on the basis of the similarities between those activities. It would *not* be a matter of course to describe a person as the sun; when Romeo describes Juliet as such the claim is metaphorical. It does not wear its directions of projection on its sleeve; they must be worked out, explained using *more* — though not "absolutely" — literal uses of words. There is no place here for a notion of absolute literality.

In light of the projective character of our use of language, the picture of language-use as answering to rules-as-rails appears as a gross distortion. Our training in the language, the way in which our elders have worked to attune us to others in the linguistic community, is not something which simply delivers us with pre-packaged ways of reaching others. Rather, it is something we must take up. This inheritance — the particular way in which we exploit our training to make ourselves understood by, and to understand, others — has a creative dimension. The speaker using her language is like a jazz musician, taking up the chord sequence of the piece, and her knowledge of music theory, to improvise a solo.[19] If her solo is to be a success, rather than an ugly cacophony, she must not stray too far from her audience's musical comfort zone, though she may achieve great results by pushing at those boundaries. (Of course, different audiences will appreciate different solos. This foreshadows issues to which this thesis will return in the next

[19]Improvisation is an analogy of which Egan (2019: 152-5) also makes use.

two chapters.)

What the linguistic novice develops is not knowledge of a calculus, but a fluid mastery of the conversational space. They are less like the person who buys a ready-to-assemble chest of drawers from a Swedish outlet, and diligently follows the instructions to build their furniture, and more like a master carpenter who is so at home with her tools that she can make any number of robust and beautiful pieces of furniture without need of instructions. (If she makes use of rules, they will be rules *of thumb* which she learnt as an apprentice. I return to this below (§1.5).)

We now turn to the way in which our training in language is always *ongoing*.

1.4.2 An evolving practice

An interesting feature of Cavell's example is the way in which it suggests a diachronic *progression* in our modes of speech. We have been going on together, and then '*one day* one of us says "feed the meter".' This may be a more-or-less novel way of utilizing the tools we have been bequeathed in language. But in pushing forwards into this new context, the speaker makes of themselves and their usage an example which we may or may not follow. 'Feed the meter' is an example of a projection which has solidified, in the sense that it is now a familiar part of common parlance.

What we have done in the past does not *determine* how we are to bring our words to any given situation. We have to *take responsibility* for making ourselves intelligible on the basis of a shared training in language.

What is visible in Cavell's example is the way in which this training continues into adulthood. For our personal and collective histories with a word are always growing. (To invert Heidegger's turn of phrase: to our words, more and more significations accrue (1962: 204).) This speaker's projection of 'feed' is an element in our collective training, an element that becomes more central as that projection solidifies.

As Gustafsson (2011: 659) highlights, changes in our forms of life will change which projections will be intelligible. (One example he gives is 'surf the web'.) Conversely, changes in us or our world can make certain projections obselete or unintelligible. With the rise of digital media, 'film' is no longer primarily a physical medium from which moving pictures are projected. Without the analog projector, the idea of 'feeding in film' may cease to have the intelligibility that it once did.

We are always, by the nature of language, *finding our way* with our words. Making ourselves intelligible to one another, and making sense of the world in the terms which our language provides, is a continuous task, the grounds of which are constantly evolving.

1.4.3 Attunement and conceptual connections

There is one further point I want to emphasize: the ways in which we are attuned to one another are shot through with *meaning*. Cavell speaks of 'our sharing routes of interest and feeling, modes of response, senses of humor and of significance and of fulfillment, of what is outrageous, of what is similar to what else, what a rebuke, what forgiveness, of when an utterance is an

assertion, when an appeal, when an explanation' (AWP: 52). These are all aspects of our shared attunement. The sense of what counts as a rebuke is a sense for the particular *significance* of that remark. To see a remark as a rebuke is a rich conceptual achievement, requiring an understanding of how being rebuked alters someone's relationship to another, what it *means* to be rebuked.

The ways in which we are attuned are not merely instincts, natural or trained, to blurt out particular words in particular contexts. They include rich senses of what it *means* to say a particular thing. This is very visible in the projective character of language: in projecting our words into new contexts, we are constrained by holistic connections between our concepts. Our conceptual repertoire is not a collection of isolated categories; concepts are deeply interwoven. As we have seen in this section, to be ready to extend 'feeding' to some activity is to be ready to see in it connections with sustenance, growth, and the like.

A proper appreciation of this point is essential for grasping the character of The Method, of examining what we should say when. It also helps to dislodge a common, but ultimately impoverished, view of what is involved in making a descriptive statement: the view that description reduces, essentially, to classification, to the subsumption of objects under concepts. Speakers of a common language have been trained into attunement, in the sense that they will largely accept and understand one anothers' projections. But for an established language, this has grown into more than a propensity to see one thing as 'like' another. It is not just a matter of being inclined to 'apply' the same words to the same things. To claim, for example, that

'the media feeds our children morally pernicious nonsense' is not merely to subsume some things under some categories. It is richer than that: the projection, relying primarily on the idea of consumption, utilizes the ideas of sustenance and growth; the media's offerings are junk food, they are not "nutritious" to the child.[20]

Relatedly, I have spoken of conceptual *connections*; such connections should not be thought of on the model of material implication, of conceptual containment. There is a conceptual connection between the concept of 'feeding' and that of 'growth'. The connection persists even if not every context in which the concept of 'feeding' can be employed is one in which the notion of growth gets hold ('feed the meter'). Furthermore, if the connected concepts *do* get a hold, they too are *projected*: the sense in which a meter is "sustained" by pennies is not exactly that in which a swan is sustained by bread. (These ideas will be developed in the next chapter; they are central to my understanding of The Method.)

1.5 Custom

None of this is intended to deny that there are rules governing language-use. The manner of this governance needs to be reconceived, however. Mistaken

[20]On this, cf. Diamond (1988: 266): 'grasping a concept (even one like that of a human being, which is a descriptive concept if any are) is not a matter just of knowing how to group things under that concept; it is being able to participate in life-with-the-concept.' Diamond's aim here is to elucidate a difference between the concepts *member of the species Homo sapiens* and *human being* which is not merely a difference in Fregean "sense" between concepts that are (necessarily) coextensive, as if to call a person a member of the species *Homo sapiens* were to convey the "same information" as one would by calling them human, though in a different way. The place in our lives of the latter concept is *very* different from that of the former, and Diamond holds that this is essential to understanding the ethical significance of the latter concept.

uses of words can be framed as violations of rules implicit in our linguistic practice. But we ought not to think of such a violation as a failure to continue an extant pattern which, as it were, demands its own continuation. Rather, to speak "incorrectly" is to court misunderstanding in a specific sort of way; to speak in such a way that the *de facto* training of one's audience inclines them to take you other than you had hoped. As a matter of fact, if one uses 'pterodactyl' where others use 'butterfly', one will sow confusion rather than co-operation.

Making oneself intelligible to others — that is, successfully informing them, asking them, commanding them, etc. — is a *skill*. It is a skill we would not so much as be able to develop if it weren't for the fact that we are, to some extent, naturally attuned to one another, and if this attunement could not be deepened through training. When one learns a new skill, one may be given some "rules" to follow. Your windsurfing instructor might tell you to put your front foot next to the mast, to keep your weight over the centre of the board, to keep your hands a little wider than shoulder width on the boom. If you "break" these rules, you're likely to fall in. When the child develops to the point where she can take on board explanations of meaning — 'a kayak is a narrow boat propelled by means of a double-bladed paddle' — this explanation may justly be called a statement of a rule. But we ought, I suggest, to conceive of the rule as like the rules *of thumb* involved in learning to windsurf.[21]

It is natural to conceive of language-use as something conventional, some-

[21]In permitting talk of rules of language, despite the Cavellian insights discussed in this chapter, I am following Stephen Mulhall (2014).

thing contractual (McDowell, 1984: 325). Here, we are reappropriating the notion of convention, without the misleading pictures which have been associated with it. We view a convention not as an agreement on which linguistic rules-as-rails to continue along, but as a trained attunement to be exploited for certain purposes. This reappriopriated notion of convention is what I take Wittgenstein to mean by a *custom* (cf. *PI* §198). Custom is what can be *relied upon* to achieve our purposes in language: the trained inclinations to understanding which I (hopefully) share with my peers. These inclinations allow for a plasticity in the way that we approach new situations with our language, a plasticity occluded by a picture of language-use as based in rules-as-rails, in set application-conditions for our words. How these inclinations are to be exploited is encapsulated, at best, in rules *of thumb*.

1.6 Conclusion to Chapter 1

Our linguistic training begins in childhood with the familiar exercises by which parents induct their children into the language. But much of that training is also passive, imbibed by the child as they grow up around people already skilled in the manipulation of the language. This training continues into adulthood; we should expect the history an individual has with a given word to influence the way in which they are inclined to take it in the mouth of another. (For the most part, these histories are reasonably convergent. They have to be, if we are to continue to be intelligible to one another. But we should be prepared for the possibility of divergence.)

What the child learns is not like learning to operate a vending machine:

press the right buttons (say the right words) in the right order, and the desired snack will be dispensed (the desired meaning will be produced). It's more like learning to drive a racecar; successful operation is more a matter of feel, of finely-developed instinct, than the rote following of rules. She needs to get a *feel* for how her various aims in language might be achieved.

If I want to inform someone of something, I had better use words which a member of my community will, as a result of their training, instinctively take as drawing their attention to that state of affairs. If I want to apologize to someone, it will not usually be effective to say, 'Oh, grow up.' Their training and history with those words will not prepare them to hear this as an apology.[22] Rather than the pseudo-logical constraint of rules-considered-as-rails, the normative constraints on language-use emerge as *practical* constraints. Speak as we customarily speak, or risk misunderstanding, unintelligibility, isolation. (It is thus unsurprising that our histories with words are reasonably convergent. A continued conformity to a more-or-less regular custom is a practical condition of intelligibility.)

To be skilled in the use of a language is to have a refined sensitivity for the expressive possibilities of our common stock of words, expressive possibilities which they possess because of our innate and learned attunement to one another. These expressive possibilities are diverse, rich, and plastic in ways occluded by a conception of natural language as answering to rules-as-rails.

In the next chapter, I will begin to explain the importance of our alterna-

[22]Cf. *BB* 42: 'Make the following experiment: say the sentence "It is hot in this room", and mean: "it is cold".' The point to take from this is not that it is impossible to do so — if one imagines a certain history between me and my audience, I may well be able to be understood as having said that it is cold.

tive vision of language for our understanding of The Method. The Method does not involve an attempt to abstract the usual conditions for our use of a given word. Rather, it invites us to project our words into *new* contexts. This invitation is an illuminating one; it invites us to look at the world anew.

Chapter 2

The Invitation to Project

> [W]e are to proceed from 'ordinary language', that is, by examining *what we should say when*, and so why and what we should mean by it.
>
> (*PFE*: 129, emphasis in original)

Benson Mates suggests that when the ordinary language philosopher makes a claim about the "ordinary use" of a given expression, their claim is (at least covertly) a descriptive one, 'refutable by observation of the ordinary folk, magistrates, parents and teachers' (1981: 68). He distinguishes two procedures by which the claim might be established, 'though any really adequate procedure will probably have to be a combination of both' (*op. cit.*: 69):

> In the extensional approach one observes a reasonably large class of cases in which the subject applies the word, and then one "sees" or "elicits" the meaning by finding what is common to these cases.

> [...] In the [intensional approach], one asks the subject what he means by the given word or how he uses it; [then] one proceeds in Socratic fashion to test this first answer by confronting the subject with counterexamples and borderline cases, and so on until the subject settles down more or less permanently upon a definition or account.
>
> (*loc. cit.*)

From this quote it is apparent at what Mates takes the ordinary language philosopher to be aiming: the specification of some *rule*(s) according to which a given person (or a community at large) uses a given word. What is to be gleaned from the investigation, either by abstraction from particular uses or by asking subjects directly, are *conditions for the application* of the word in question. That is, Mates takes it that a successful OLP investigation would correctly discern the rails along which we travel in the use of a given expression. 'One can use expression e only if...'

This characterization of the results of an OLP investigation is at odds with the vision of natural language developed in the last chapter. To ask for the meaning of a word is to ask for a primer on how to use it. But the answer to that question is not a *condition* for successful use. It is, at best, a rule of thumb that will assist the speaker in being understood. Natural language expressions may be used in more-or-less regular ways, but this is not because to speak meaningfully is to continue along such a rail. We are viewing the use of a natural language not as the continuation, with our words, along rules-as-rails, but as a skill, the skill of leveraging one's attunement to others

in order to make oneself understood. Deploying this skill involves *projecting* our words, making use of them in novel contexts, and trusting that others, in their attunement to us, will see the similarities that we see, the lines along which we are projecting.

In this chapter, I want to present OLP's characteristic Method, of asking what we should say when, as an invitation to project our words into particular contexts. In §2.1, I want to explain what Austin is doing as 'consulting our attunement'. He invites us to exercise our skill in the use of language: to say what, on the basis of each of our individual trainings in the language, we take to be the correct way to describe this specific situation. We are consulting our sense of what it would mean to say that some *specific* thing is F. As we shall see, this unlike Mates' 'intensional approach' in that it does not aim at the specification of a general rule. In particular, it does not issue in necessary and sufficient conditions (even vague ones) for F-ness. Rather, it deepens our understanding of *this specific situation* by exploring the expressive potentialities of our words.

The kind of attention for which Austin is calling is a *critical* attention. By learning what it means — that is, what it means to *us* — to say that something was done by mistake rather than by accident, we may come to think that much of what we have said in the past elided a significant distinction that we would do well to respect in future. (We might say that, unbeknownst to us, we were not always saying what we meant in speaking indiscriminately of things done by mistake or by accident.) The Method can precipitate radical changes in *de facto* usage. Furthermore, employing The Method requires us to be deeply responsive to the *world*. This is the subject

of §2.2.

In §2.3, a worry is raised. Our judgments as to how to project our words will coincide only insofar as we are attuned to one another. What if we are *not* so attuned? That worry is taken up in Chapter 3.

2.1 Austin's donkeys

To have in mind a sample of The Method, let's take the following example of Austin's:

> You have a donkey, so have I, and they graze in the same field. The day comes when I conceive a dislike for mine. I go to shoot it, draw a bead on it, fire: the brute falls in its tracks. I inspect the victim, and find to my horror that it is *your* donkey. I appear on your doorstep with the remains and say — what? 'I say, old sport, I'm awfully sorry, &c., I've shot your donkey *by accident*'? Or '*by mistake*'? Then again, I go to shoot my donkey as before, draw a bead on it, fire — but as I do so, the beasts move, and to my horror yours falls. Again the scene on the doorstep — what do I say? 'By mistake'? Or 'by accident'?
>
> (*PFE*: 133n1, emphasis in original)

I hope I can grant that Austin succeeds here in reminding us of what we would say. In the first scenario, the neighbour's donkey was shot by mistake; in the second, by accident. That is to say, with respect to this example, we are in fact attuned in how it strikes us to project the words 'accident' and

'mistake'.

What Austin does here is to vividly describe a scenario, and then invite us to project our words into it. If we needed to describe what had occurred here to one who had not been present, how would we do so? If it were *you* on the doorstep, cap in hand, by what words would you admit what you had done? In order to answer Austin, we will need to make use of our skill in language. As with most skills, our ability to do these things is usually deployed without reflection. We follow the natural lines of projection for our words and, attuned as we are, meet (for the most part) with success. But we can also reflect on the deliverances of this skill. This is what Austin invites us to do. We all agree (don't we?) that the first donkey was shot by mistake, the second by accident. Thus, our attention is drawn to distinctions which we show ourselves unconsciously equipped to make, already unconsciously making, in our deployment of these words.

Austin invites us to say what, on the basis of each of our individual trainings in the language, we take to be the correct way to describe this specific situation, how we ought to project our words. If we are in fact attuned to one another, our voices will harmonize. Insofar as each of us is in fact attuned to our fellows, we may speak for the rest in speaking for ouselves. We are consulting our (hopefully shared) sense of what it would mean to say that the donkey had been shot by mistake, rather than by accident. (The obvious question is then: what if we are *not* attuned? I bid the reader exercise patience. We will get there.)

We can speak of The Method as 'consulting our attunement'. As we saw in the last chapter, the ways in which we are attuned cover a wide spectrum from

biological similarities of sensitivity all the way through to shared inclinations of projection that result from training in the skill of language-use. It is in self-consciously drawing on the ways in which we are attuned to one another that we see that we ought to say that first donkey was shot by mistake, the second by accident. We make use of a capacity analagous to that of perfect pitch: our sense of how these circumstances relate to the practice into which we have been trained. We make use of our 'ear' for meaning, our sense of *linguistic* pitch.

What we are *not* being asked to do is to determine how "we" use the words 'mistake' and 'accident', by surveying what we have done with them in the past, in order to extract some pattern which we can then apply to this situation in order to determine the correct way to classify each event. (There will be many such patterns, many ways of bringing what came before, considered merely as enshrining principles of classification, to the current case.) This is particularly apparent in Austin's example, which is a *novel* one, a context in which we have probably never before been asked to apply our linguistic competence. (Analogously: at some point, the pupil of *PI* §185 reaches a part of the series they have never before seen expanded.) Before considering this case, there is nothing "we" say about it, nothing which can be used as the arbiter of what it makes sense, or would be correct, to say. Thus, the particular case must be given primacy.

Am I saying that we must consult our intuitions as to whether a particular thing was done by mistake or by accident? One may call this exercise of proficiency with language 'intuition', if one likes: certainly, the judgments are not derived from prior *theoretical* knowledge. But the usual use of 'intuition'

in Anglo-American philosophy has been to make initial judgments from which can be abstracted theoretical accounts, necessary and sufficient conditions, in pursuit of an equilibrium between theoretical elegance and accommodation of intuitions. That is not at all what I am suggesting is involved in examining what we say when.

What are revealed in our self-conscious exercises of linguistic facility are conceptual connections. We are not looking to discover the necessary and sufficient conditions of F-ness. We do learn about F-ness, however. Given the holistic connections between our concepts, to project a word into a context is to invite other projections. To project 'feed' is also to invite the projection of notions of nourishment, of growth, etc. Such connections need not be material implications ('If one can apply "blame", then one can also apply "responsibility"'). The connections are more fluid than that. When one asks someone to 'feed the meter', one may also project the notions of sustenance, of nourishment, of giving the meter what it "needs". But these notions are, of course, *projected*; the way in which a meter is "nourished" by coins is unlike the way in which a swan is "nourished" by bread. Furthermore, some of the invited projections are *not* taken up; if there is something in the case of the meter that corresponds to the idea of growth, I am not sure what it is.

To accept one projection rather than another is to see certain lines of similarity as salient, and dissimilarities as receding in importance. We reflect upon our inclinations to project, inclinations developed as a result of an ongoing training, continuing into our adult lives in a world with words, a world in which mistakes are made and accidents happen; what we learn is

nothing so easily codifiable as necessary and sufficient conditions. It is something much more fluid. Our knowledge receives expression in our ongoing deployment of our facility with words.[1]

As a result of one's investigation, one might frame general statements. 'A mistake occurs when someone intends to perform an action, but acts on something other than the intended object of that act. An accident occurs when someone performs an action they did not intend to perform.' General statements like this may well succeed in getting others to see what you see in particular situations. (She put salt in my tea (instead of sugar) by mistake; she dropped the cup by accident.) But it is important to remember that whether something meets the conditions framed by the general statement will always be essentially problematic. The same questions as to whether one is right to project 'mistake' will reoccur for the terms in which the condition is framed. In both of Austin's examples, I intended to shoot my donkey, I end up shooting yours. The particulars of the case are relied upon to decide which description is apt: did I act upon something other than intended (your donkey, instead of mine), or did I perform an action (shooting your donkey) other than the one I intended (shooting my donkey)? If one tried to frame conditions to determine, in advance, how we should decide this question, similar problems would reoccur (again, recall the wayward pupil). One is ultimately thrown back upon one's own judgment, upon one's own skill with the language, upon one's always-developing sense of which similarities and differences matter for the particular case at hand. One is thrown back upon

[1]This is what I take to be the importance of Wittgenstein's observation that most interesting concepts are family-resemblance concepts; cf. *PI* §§65-70.

the ways in we which are attuned.

2.1.1 Yes and no

The character of this kind of investigation, and the way it differs from the search for necessary and sufficient conditions, is elucidated in the way it can be informatively inconclusive.

For many interesting cases, reflection might leave us with the answer: 'Yes and no'. Neither p nor not-p is the "right" thing to say. There are competing pressures to be respected; that is, to be *respected*, rather than synthesized to give a unique answer.[2] For example: do garden birds have agency? To say either 'yes' or 'no' may be illuminating at the risk of being misleading. To say 'yes' is to occlude a great difference between sparrow decision-making and human decision-making. To say 'no' is to occlude a great difference between a sparrow and a wind-up toy.

For a philosophical method based on reflective equilibrium, there may be borderline cases of F-ness. The resulting theory of F-ness may produce conditions on F-ness which some cases only *more or less* meet. For example: it is necessary and sufficient for something to be a vixen that it is a female fox. Suppose we encounter a female canine whose ancestors are *mostly* foxes; however, her maternal grandsire was not a fox, but a jackal. Is this animal a vixen? We might say that it is a borderline case.

The picture here is one of *vagueness*, as that phenomenon has been under-

[2]Compare moral views according to which there are incommensurable goods, and so the possibility of irreconcilable conflict between them (e.g. Nussbaum, 1990). My duty to my father cannot simply be weighed in the balance with my duty to my spouse, synthesized to leave me with a "resultant" duty to one or the other, and absolving me from guilt if I act according to this resultant duty. The claim I am making here is similar.

stood in analytic philosophy. We have predetermined boundaries for fox-ness, but those boundaries are blurred at the edges. We are still, in essence, simply sorting cases; although, rather than putting them into two neat boxes (the foxes and the non-foxes) we arrange them along a spectrum of foxiness.

A picture of language-use based on attunement instead pushes us to give up the idea that our language decides cases in advance, in this sort of fashion; we are not travelling along rails to infinity. The question is not 'Under which category does this situation fall', as if the answer ought to be prepared for us in advance by how we have been using our words thus far, with inconclusive cases falling into the gaps.[3] A better question: what is the *mot juste* here, given one's sense of the meanings of those words, of their expressive potentialities? What matters is not so much our final verdict, but the connections and similarities that verdict reveals.

The attitude towards language here defended has a somewhat deflationary effect on the notion of truth. Even allowing for an indeterminate answer, the question is not, 'Is our creature *really* a vixen?'. Recall that the normative

[3]Neither are we concerned with what Waismann (1945: 121) calls 'open texture'. That phenomenon could be fitted into the rules-as-rails picture by construing those rules as leaving some leeway, which is then eliminated as we take into account previously unconsidered questions (is a cat still a cat if it sprouts wings and flies away?). Diagrammatically, it would look like this:

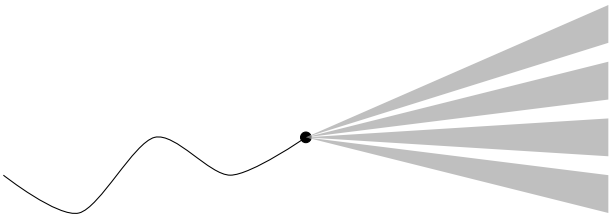

The cone which represents the *actual* continuation of the rule would narrow to a continuation of the line as we accommodate future cases. The vision of language I am defending is more radical than this.

constraints on language-use are not those of rules-as-rails, of predetermined applications-conditions, but those constraints imposed by the inclinations of our fellows, in which we are, hopefully, attuned. The constraint is more *pragmatic*. Suppose I tell you, 'Daisy gave me an apple'. When I later bite into it, I find that it's plastic. So there was something amiss in my earlier report to you, from which you would have taken that Daisy had given me a real, organic apple. How do I correct this failure?

(i) Daisy didn't give me an apple after all, it was only a model.

(ii) That apple Daisy gave me was plastic. What a joker!

Now, it looks as if, if I can say (i), then the original report — that Daisy gave me an apple — must have been false; if I can say (ii), it must have been true. But I can say either. Any misapprehensions my original assertion caused, my surprise, etc. — all this is explicable without having to decide if my original assertion ought to be listed under TRUE or FALSE.

> [F]orget, for once and for a while, that other curious question 'Is it true?' May we?
>
> (*PFE*: 133fn.)

The point here is not that there is "no such thing" as truth. It is just that there is a prevailing tendency to construe truth is if speaking truly were simply a matter of sorting cases under the right labels: vixen vs. not vixen, free vs. not free, conscious vs. not conscious, etc. For great stretches of our use of language, the question of whether what someone says is true effectively coincides with the question of whether it is *apt*, whether it is helpful. (Both

(i) and (ii) seem to be perfectly apt — perfectly *true* — ways of expressing what occurred between Daisy and me.)[4]

An example. Catharine MacKinnon (1993: 109) claims that the use of pornography amounts to 'sex between people and things, human beings and pieces of paper, real men and unreal women' (we might now include online videos among the 'things'). One will miss the importance of this fascinating claim if one approaches it simply by asking if masturbating while watching a pornographic film "counts" as having sex. What MacKinnon is doing here is exploiting expressive possibilities which she finds in the word 'sex', in order to bring out elements of what she finds objectionable in the production of violent pornography. When one follows out the conceptual connections in her projection, one may find oneself saying: 'How true'. One might equally contest her projection; but it is contested authentically only if one properly appreciates its projective character. This is not a case of her expressing something figuratively which she might instead have said literally; recall that we are giving up that dichotomy as sharp. (She might, perhaps, have said it *more* literally; but perhaps, if she is to change our way of thinking about pornography, she finds herself forced to exploit less natural lines of projection.)

Is pornography-consumption of a form of sex? Perhaps when asked to make this projection, we will, on careful consideration, want to say, 'Yes and no'. A 'yes and no' answer is informative in itself, as Cavell notes (*CR*: 13-4). Cavell asks of such questions: 'What would be the point of deciding upon

[4]The claims made here relate closely to what I will have to say about meaning-determination in the next chapter.

a more definite answer — I mean, a shorter answer — to these questions?' But, for at least some of these questions, it is worth our while to pursue a *longer* answer. For just as a conclusive answer as to whether the donkey was shot by accident or mistake helps us to make sense of responsibility in a human world where such distinctions matter, so too might an exploration of the competing pressures which lead to an inconclusive answer.[5]

I take myself to be elucidating the character of conceptual investigation here. To reflect on the concepts MISTAKE and ACCIDENT (and with them, the related concepts of RESPONSIBILITY and BLAME) is a matter of self-consciously exercising one's proficiency with them, with the *words* 'mistake', 'accident', 'responsibility', 'blame'. In consulting our attunement, we evaluate what we *should* say, what seems to us most apt, the most fruitful way to leverage our mastery of the language in this particular case. In so doing, we learn something about what it means to say that something was done by mistake, rather than by accident. Our training — a training which, importantly, includes our constantly accumulating life-experience — is taken up, brought to this situation. We do so in similar ways because we *are* similar, and have received similar training. In doing so, we learn not just about the meanings of words, but about what *mistakes* and *accidents* are in the context of a human life. But what we learn are not the rules for applying 'mistake', or necessary and sufficient conditions for being a mistake. We develop a more

[5]This phenomenon — illuminating inconclusiveness in the deliverances of The Method — is what is exploited in John Wisdom's use of contrary statements in order to illuminate his subject matter. (e.g. 'Philosophical statements are really verbal'/'Philosophical statements are not verbal'; Wisdom, 1936: 71-2.) I am rather fond of the following slogan of Wisdom's: say what you like, but *be careful* (*op. cit.*: 74). Know which similarities and differences you are emphasizing, and which you are occluding.

fluid, richer, more textured understanding of life in the world with others.

2.2 Principled reform

It has sometimes been suggested that OLP is inherently conservative or reactionary; that it cannot but reflect one's already-held beliefs back at oneself. (The claim is prominent in Gellner's *Words and Things*, 1959.) If The Method of OLP was to examine *de facto* usage in order to abstract rules which govern our use of language, in order to dismiss out of hand utterances which 'violate' those rules as false or nonsensical — the kind of approach exemplified by the Paradigm Case Argument (PCA) discussed in the Introduction — that criticism might be apt.

But OLP, as I understand it, is not in the business of abstracting rules of this sort. In fact, as I show in 2.2.1, the vision of language I have been defending provides a principled response to attempts to police language-use by reference to 'what we do'. In 2.2.2, I explain that, far from being intellectually conservative, The Method can precipitate reform in our *de facto* ways of speaking. It does so by encouraging careful attention to the reality about which we use our language to speak.

2.2.1 Policing language-use

> The PCA is an argument from ordinary language. The idea is to describe cases which would clearly be regarded as having property P in accordance with the ordinary meaning of 'P', thus refuting the claims of philosophers that nothing has (or that, for all we

know, nothing has) property P.

(Hanfling, 2000: 89)

This is an argument which Hanfling endorses on OLP's behalf, commenting with approval on Flew's use of it (noted in the Introduction):

> The man wants to marry the girl, they have been in love for some time, he is not under any kind of duress, etc. Here is a paradigm case of acting freely, and by describing such examples the meaning of 'free will' may be explained. Someone who denies that this would be a case of free will cannot be using the expression with its normal meaning, and thus the meaning of his denial is itself thrown into doubt.
>
> (*op. cit.*: 76)

The PCA demands that the philosopher speak as "we" "ordinarily" do, or else speak irrelevantly; if she does not use 'free' according to our standards, she will not succeed in denying anything we believe about freedom of the will. If she *does* use 'free' according to those standards, her claims are immediately refuted by the fact that they contradict claims which are ordinarily permitted, in which those standards are enshrined. She is faced with a dilemma: irrelevance, or falsity.

A similar appeal to ordinary language appears in Baker & Hacker's influential commentary on *PI*:

> If someone claims that colours as we apprehend them are sensations in the mind (as Descartes did) or in the brain (as contemporary psychologists and cognitive neuroscientists do), one should

> point out that he is misusing the words 'sensation' and 'colour'. [...] [O]ne can have (i.e. it makes sense to speak of) sensations in the knee or in the back, but not in the mind. It is extended things that are coloured.
>
> (Baker & Hacker, 2010: 56)

Here the second horn of the dilemma presented to the philosopher is slightly different; if she subjects herself to our standards, which she must if she is to speak to us at all, her claim is shown up not as simply false, but as *nonsensical*.[6]

What should now be apparent is that both of these appeals rely on abstracting from "ordinary" use some rule for how an expression may and may not be used. If it were the case that our previous commerce with an expression demanded a certain continuation, decided in advance what must and must not be done, then an appeal to the "rules" enshrined in our ways of speaking might suffice to rule out Descartes' usage. Similarly, if our applica-

[6]I do not think that this difference is a deep one, for if we are correct in our way of viewing language in this thesis, then the distinction between falsity and nonsense is not deep. Manifest falsity shades into nonsense. Given our vision of language, the question is not whether the rules for the language determine that I have failed to say something, but whether others can, in fact, and on the basis of their training in the language, make something of what I have said. There clearly are situations in which I can make something of the instruction, 'Go and get me three cans of red paint'. No situations readily come to mind in which I can make something of 'Ab sur ah' (cf. *WLC*: 64): nobody I know has been trained in the use of these signs. Can I really make something of my sister's claim that her pet fish is plotting to kill her? What on Earth would this look like (given my knowledge of the goldfish in question as apparently unremarkable)? There is no once-and-for-all answer to this question. For what it's worth, I take this to be the significance of *PI* §282:

> "But a fairy tale only invents what is not the case; it does not talk *nonsense* does it?" — It is not as simple as that. Is it untrue or nonsensical to say that a pot talks?

tion of 'free' to cases like that of Flew's marriage had the role of determining a meaning for 'free' which future uses must respect, then the dilemma posed by the PCA might stand. But our previous usage does not have the role of arbiter, to which any putative usage must be compared in order to determine whether it is legitimate or not. Rather, it is a basis which we must *exploit* in order to say what we mean. Our training in language prepares us to say what we mean, but one does not do so merely by producing an approved formula. The customs according to which we speak are there for us to leverage in speech, the raw material of which we must *make* something. But plenty of leeway is left for how we might do so. Our customs do not have the kind of authority over particular uses that Baker & Hacker suppose them to have.

Descartes' assertion must be approached as a *particular* utterance. Like anyone making an assertion, Descartes projects the words 'sensation' and 'colour' on the basis of a shared history with those terms and an (ongoing) training in their use. We must at least be open to the possibility that '[this] new projection, though not at first obviously appropriate, may be made appropriate by giving relevant explanations of how it is to be taken, *how* the new context *is* an instance of the old concept' (*CR*: 192, emphasis in original). We need to try to understand what Descartes is trying to do with his words at this particular moment, in this particular speech situation; how he intends to project them. What conceptual connections does Descartes' utterance hope to exploit?[7] Now, we may fail to follow Descartes' projections:

[7]Of course, Descartes may not have had an accurate view of what he was doing. He might well have thought that he was just continuing an extant pattern of usage in the only way it could be continued, that 'sensation', 'mind', 'colour' determine conditions of application which demand his speaking as he does.

in our attempt to understand what Descartes says, his words may fall apart in our hands. Such a failure would cast doubt on Descartes' having meant anything at all. Or we may find that we can only understand him as saying something with which we cannot agree. What we cannot do is defer our responsibility for understanding to "the language": we cannot assume that our customs embody rules which *determine* a content for what Descartes says (or determine *no* content for it). (I take up the subject of meaning-determination in the next chapter.)[8]

2.2.2 Language and world

OLP need not subscribe to forms of argument which maintain the intellectual status quo; which endorse, by default, the ways of speaking of the man on the Clapham omnibus. It is, in principle, open to substantially new ways of thinking.

In this section, I want to emphasize that The Method should be understood as a *critical* examination, and that, as such, it is capable of *dislodging* the intellectual status quo. A musician, through a sort of reflection on what she is doing which is normally absent in her playing, may find that some aspects of her recital are sloppy, that her performance would be improved by

[8]Even if we cannot find a way of making sense of what the philosopher claims, we should never be *sure*; perhaps someone will eventually show us how to make something of it. This is what I take to be the importance of the following remark of Cavell's:

> [I]f I say 'They are crazy' or 'incomprehensible' then that is not a fact but my fate for them. I have gone as far as my imagination, magnanimity, or anxiety will allow; or as my honor, or my standing cares and commitments, can accommodate.
>
> (*CR*: 118)

a change in the way she plays, by adjusting her *de facto* deployment of her musical skills. In critically assessing her playing, she employs the same facility with her instrument which was deployed in the playing itself. Similarly, it may be that people usually run together accidents and mistakes, using 'accident' and 'mistake' interchangeably. What Austin shows us is that there *is* a difference, a difference appreciable even for those who have been conflating the two notions.[9]

In his laudatory introduction to *Words and Things*, Bertrand Russell writes:

> The linguistic philosophy, which cares only about language, and not about the world, is like the boy who preferred the clock without the pendulum because, although it no longer told the time, it went more easily than before and at a more exhilarating pace.
>
> (Gellner, 1959: 15)

There is no such disconnect between language and world in The Method as I am defending it. It asks us how we wish to project our language into some context. In so doing, it requires us to be carefully attentive to the reality of that situation. We are to see how the particular reality of *this* situation strikes us, and how we are going to bring our linguistic training, our fluid skill in the use of language, to bear upon it. (The case to which we are asked to

[9]The analogy with the musician raises an interesting question. Musicians often play better when they are not thinking too much about what they are doing, when they exist in a state of "flow". Might the attentions of OLP be similarly distorting? That they might be is not something I wish to rule out. The kind of attention for which OLP asks is not easy. Stepping out of the practical context in order to self-consciously examine our inclinations to speak might, if we are not careful, alienate us from the very skill we hope to be employing. This might be pursued as an explanation of the well-documented philosophical impulse to scepticism, an impulse of which OLP has been very critical.

respond may be novel, like Austin's, but it may equally ask us to reconsider some past usage. We can turn The Method back on past usage, because that usage is not taken as if it were the visible segment of a rail-to-infinity, against which future uses must be compared to judge their correctness (and which would, of course, judge itself correct). We might turn The Method on past usage directly, or we might be prompted to do so by the consideration of a novel case that leads us to re-evaluate the conceptual connections in this sphere, so that we look on past usage in a newly critical light.)

The Method requires us to be deeply responsive to the *world*. Consider someone who did not feel the force of Austin's donkey example, seeing no difference between the donkey shot by mistake and the one shot by accident. This is a *failing* on their part, a failing to see a real difference in what we might call the 'texture' of the two cases. And their failure impoverishes them, because the difference is (humanly) important. Depending on the facts of the case, it will make sense to hold me responsible in different ways. What I can be trusted to do in the future may depend on this. Suppose the donkeys look very similar, but were standing far apart. If I shoot yours by mistake, then I have not been sufficiently attentive, given the seriousness of my task. I should have been more careful to check which animal was which. Would you trust such a person to fill in a complicated, but important, form for an elderly relative? Suppose instead that the donkeys don't look at all alike, but are standing close together. If I shoot yours by accident, then I have been unduly reckless, though perhaps not inattentive. I should have taken the time and effort to lead mine away from yours before shooting it. Would you trust such a person to transport a delicate antique? Mistakes really are

distinct from accidents, though the distinction is visible only from within the life of a creature whose form of life brings them to make this distinction; a creature who shares their life with others on whom they must sometimes rely, and to the consequences of whose actions they are vulnerable.

> When we examine what we should say when [...] we are looking again not *merely* at words (or 'meanings', whatever they may be) but also at the realities we use the words to talk about: we are using a sharpened awareness of words to sharpen our perception of, though not as the final arbiter of, the phenomena.
>
> (PFE: 130, emphasis in original)

By asking us to reflect on how we might approach particular situations with our language — a task that is not simply a case of subsuming situations under labels, but of capturing their rich similarities and dissimilarities to other cases, of seeing how individual cases fit into the course of a human life — The Method brings us to a deeper understanding of those situations. As a result of this deepened understanding, we may reform our use of language. We may think, for example, that, even though the groom honestly believes he wants to marry his bride, and does it for all appearances happily, there is something of a tension in calling his marriage 'free'. Our reasons need not be all that sceptical. We might suggest that his upbringing and the culture that surrounds him has limited his idea of the possibilities for a flourishing life, so that he takes it for granted that getting married and starting a family is just "what one does". Certainly, he *chose* to get married. There was no shotgun at his back. But did he really choose *freely*?

To reiterate: the suggestion here is not that we might discover necessary conditions for free action which this apparently happy bridegroom does not meet. Rather, in attending carefully to this specific situation, we discover expressive possibilities in our word 'free', we find that there may be some mileage in denying that he acts freely. It helps us, for example, to bring out an important question about societal influence on our sense of what is possible for a human life.

It might be tempting to ask, at this point: but is he *really* free? I hope that this chapter, and particularly §2.1.1, will have gently steered the reader away from asking this question; as if by doing so she could reach past the expressive possibilities of her language to a naked world which wears on its face the way cases *must* be connected, what we *must* find significant in them. (We might say, 'It depends what you mean by "free"', but this is also potentially misleading: it is the single concept FREE, with its various holistic connections to other concepts, which I intend to be utilizing here.) Responsibility to the world does not mean accepting that every well-formed claim must straightforwardly have a truth-value. Rather, it means bringing careful attention to the world, and utilizing the fluid expressive possibilities of one's language to make sense of its texture.

2.3 Failure

I have presented The Method as a consultation of those inclinations to understanding in which we are attuned. We consider, in rich detail, particular situations and bring a critical awareness to the ways in which our words

might be employed to navigate them. We must now address the elephant in the room: what if we are *not* attuned to one another? What if our reflective judgments about how to project our language diverge?

Austin may not be able to bring someone to appreciate a distinction between mistakes and accidents. Similarly, I may not be able to bring my French friend from Chapter 1 to register a difference between butterflies and moths, to see them as kinds of insect that ought to be distinguished. (One has to imagine quite a different form of life for this person; they must be someone for whom the differences between flitting between flowers in the daylight and chewing through clothes in the dark are not salient.) Or we might imagine a human community whose members are all colour-blind, and whose palates are not as refined as ours. We may be unable to bring these people to distinguish between lemons and limes.

None of these failures of attunement are particularly concerning. For even if alternative methods of bringing them into attunement fail (for example, teaching the last group about the biological distinctions between lemons and limes), we shall simply be able to say that there is something that these people are *missing*, a way in which their lives are impoverished.

But not all failures of attunement are so benign. We can distinguish two cases. The first is that in which some other, not well attuned to us, discriminates not less finely, but *differently*. Rather than occluding a distinction, they draw it in a different place. As distinctions become more esoteric, drawn less as a biological matter of course, this becomes a very real risk. Nevertheless, the disagreements precipitated may be far from trivial. One might, for example, submit the word 'murder' to The Method. Consider the following

cases:

(a) Clarissa works in construction. One day, she forgets to apply the handbrake on her digger when she goes to lunch. The digger rolls, running over a colleague and killing them.

(b) Clarissa is pregnant. The fetus is viable, and the pregnancy poses no significant risk to her health. She chooses to abort the fetus.

(c) Clarissa is pregnant. The doctor tells her that there is a very good chance that she would die during childbirth. She chooses to abort the fetus.

(d) Clarissa's partner is violent. One evening, while drunk, the partner assaults her. Clarissa grabs a kitchen knife to defend herself; in the ensuing fray, she stabs her partner, who later dies from their wounds.

(e) Clarissa is a police officer. She shoots and kills a suspect whom she mistakenly believes to be holding a firearm.

The question to be asked is: which of these situations is aptly described as involving a murder? Into which of these contexts do we want to project the word 'murder'? This is a weighty word; it matters a great deal how we employ it. Even if the cases were fleshed out in great detail, it is very likely that considered judgments on which of these cases to describe as 'murder' would differ.

The other cases in which failures of attunement are unsettling are those in which others draw a distinction which we are *not* happy simply to call a matter of a finer sensitivity. There are cases in which we feel that to draw

the distinction, rather than to respect the similarity, is to miss something important. For example, suppose someone refused to call the romantic relation that can exist between two men 'love'. This situation differs from those into which they would project 'love' in that both parties are of the same sex. That is a difference; but not one that I would wish to give any weight.

As these examples make clear, these worries are most potent in the ethical sphere. They are worries which are routinely raised for "Wittgensteinian" approaches. Addressing them is important in its own right, but it will also serve to differentiate my approach from this supposedly "Wittgensteinian" one.[10]

Here is Onora O'Neill:

> The Wittgensteinian approach to ethics by examples depends on the possibility of arriving at 'what we do want to say' in the course of reflecting on the example. This method must presuppose sufficient community of moral views — an ethical tradition, perhaps, or a shared ideology — for there to be something which 'we' (whoever 'we' may be: and this is a large question) do want to say about a given example. Where that shared tradition is lacking, some Wittgensteinian writers claim, we find ourselves confronting not moral disagreement but a breakdown in moral communication — an impasse of incommensurable moral frame-

[10]The scare-quotes are appropriate insofar as we should be wary, as I think that we should be, of ascribing such conclusions to Wittgenstein himself. But, to reiterate, our aim here is not exegesis. As the most likely misunderstanding of what I am trying to say, however, this topic must be addressed.

works.[11]

(1986: 12)

The worry is that such a view precipitates us into a particularly bald and uninviting relativism. We appear to be approaching the claim that (e.g.) homophobia is, at most, wrong *for us*.

In the next chapter, I will develop this worry in more detail, and respond to it. Doing so will further elucidate the vision of language with which we are working: that vision constitutes a radical critique of prevalent ideas about the determination of what is meant by a given utterance. I will suggest that what is said by a given utterance is something which must be *negotiated* between us. What is essential to understand about this negotiation is that it does not aim at a unique fact-of-the-matter about "what was meant"; neither does it *determine* such a fact-of-the-matter.

[11]While some of O'Neill's attributions of this view (e.g. to Phillips & Mounce, 1969) may be fair, others (e.g. to Diamond, 1991a) are not. (Cf. Diamond, 1991f.)

Chapter 3

Meaning, Negotiation, and Rational Appeal to the Other

> The philosophical appeal to what we say [is a claim] to community. And the claim to community is always a search for the basis upon which it can or has been established. I have nothing more to go on than my conviction, my sense that I make sense. It may prove to be the case that I am wrong, that my conviction isolates me, from all others, from myself. That will not be the same as a discovery that I am dogmatic or egomaniacal. The wish and search for community are the wish and search for reason.
>
> (*CR*: 20)

Although it comes in many forms, philosophy almost always presents itself as an appeal to *reason*. At the core of the self-image of our discipline is the idea of the rational appeal to the other. We are to bring others round to our point of view by the rational scrutiny of their beliefs; we are to be open

to the possibility that rational scrutiny of our own convictions might prove them groundless, or flatly incorrect.

The worry that was emerging at the end of the last chapter is that a philosophy based on The Method severely limits the power of rational appeal. The worry is not simply that we will not always agree with one another. That is a practical problem, a fact of life which any philosophical methodology will have to accommodate. (As Cora Diamond (1991a: 295) notes, even if you provide someone with a formally sound argument, you cannot *force* them to accept it. A person may persist in irrationality.) The worry is not that rational appeal may be ineffective, but that it might become *impossible*; that making a rational appeal to another is possible only insofar as they are sufficiently attuned to you.

The worry comes in an extreme form, and in a less extreme one. The extreme form of the worry (developed in §3.1) is that, if there are others who are not attuned to us, whose inclinations to project their words differ systematically from ours, then their words will mean different things from ours. We will not, in fact, be able to meet them with our words. What looks like a deep disagreement is in fact a situation in which people talk *past* each other.

As I shall explain, this worry fails to appreciate the radical nature of a vision of language based on attunement. Ideas about meaning which are almost ubiquitous in Anglo-American philosophy force on us a picture of what attunement is being enlisted to do; that is, to *determine meanings* for our (context-bound) speech acts. But the attunement conception of language-use poses a deep challenge to prevailing ideas about meaning-determination.

To speak is to enter into a (perhaps inconclusive) negotiation of meaning with others. Perfect attunement is not required for this negotiation to be entered into. §3.2 applies this idea to show how The Method can be used in productive discussion between parties that disagree in their values and commitments.

§3.3 raises a weaker form of the worry about rational disagreement, based on the fact that, if the normativity of language-use is based in the ways in which we are attuned — in inclinations to understanding that we share — then we are left without a place to stand should we wish to critically scrutinize those inclinations themselves. In §3.3.2, I argue that these inclinations are subject to a kind of *internal* scrutiny. There are no guarantees that encouraging such scrutiny in others will bear fruit. But the result is not moral relativism of a sort that ought to worry us.

It is my hope that this chapter will show that, far from trivializing deep disagreements, OLP's vision of language gives us a more realistic, and more responsible, perspective on negotiating difference.

3.1 Strong incommensurability

3.1.1 Talking past one another

If we are arguing about how to get to the bank, and it transpires that I am talking about the Barclays building, while you are talking about the bank of the river, we will not have disagreed. In the strongest form of the worry to be discussed in this chapter, the OLP vision of language turns all deep

disagreements into just this: not really *disagreements* at all.

We are trained in the use of a language, and as a result, we develop inclinations to project our words into different contexts. If the training succeeds, then our inclinations will more or less harmonize with those of our fellows. When we exploit that training, they will understand us, and vice versa. We will hear the meaning in their words.

One might take the significance of these ideas to be as follows. In Chapter 1, we appealed to the *de facto* inclinations of others, the way in which one would in fact be taken if one projected one's words in certain ways, in order to explain the normative constraints on the use of our words. The normative constraints on the use of a word characterize its *meaning*. So the inclinations to understanding around which we are in fact attuned determine the meanings of our words.

As such, a community attuned around *different* inclinations to understanding with respect to the same written symbols or audible sounds would *mean* different things by those expressions. Despite the physical similarities between our word 'love' and theirs, the two are not translatable by one another. Suppose that we insist that two men can be in love, and they insist that they cannot. We try to proceed by reflection on examples in the style of The Method, attempting to draw out the conceptual connections involved, but this does not succeed in bringing either of us round to the other's point of view. All that it does is trace the dispute to fundamental differences in inclination to project 'love' (and related concepts). Then, the worry goes, we must mean different things by 'love'; we will not in fact contradict each other when we insist that two men can be in love, and they insist that they

cannot. Our situation is analagous to the people who argue about how to get to the bank.

We could put the worry in these terms: despite everything I have said in this thesis, we are still essentially committed to the garden-variety "Wittgensteinain" idea that to speak is to participate in a language-game, conceived of as a self-contained system with its own rules and standards of correctness, and insulated from external rational criticism. All we have done is introduced some scope for creative exploitation of those rules, *within* language-games, by construing them as rules of thumb.

This worry misconstrues the importance of attunement for what one means by what one says. It is an understandable misconstrual, since it derives from ideas about the determination of the meaning of a given utterance which are prevalent in Anglo-American philosophy. If the vision of language I have been defending in this thesis is apt, however, then these ideas are distorting of the character of deep disagreement.

3.1.2 The functional conception of meaning

In introducing the idea that language-use is involved with normativity in Chapter 1, I briefly alluded to the following idea: if one knows the rules of a language, then one will be able to apply them to understand, or to produce, novel utterances. The first line of Kafka's *Metamorphosis* reads:

> As Gregor Samsa awoke one morning from uneasy dreams he found himself transformed in his bed into a gigantic insect.
>
> (2005: 95)

It is unlikely that the first-time reader has ever come across just *this* sentence before. But a fluent reader of English will have no trouble understanding it. They are able to do so because they understand the individual words (that is, they have been trained in their use), and they understand how English sentences are customarily composed out of those words (they have themselves been trained in doing so: in putting individual words together to achieve complex ends).

That we can understand novel utterances in this way — that natural language is compositional — is undeniable; I have no intention of denying it. If we understand 'the rules of language' in the way urged in Chapter 1, that is, as more like rules of thumb, it is perfectly innocuous.

But much more robust ideas about meaning have been drawn from (and perhaps not properly distinguished from) these innocuous ones. One such idea I call the *functional conception of meaning*. 'Function' is here used in its mathematical sense; a function is a mapping from one set of objects (the domain) to another (the range), with each element of the domain being mapped to exactly one element of the range; the object to which a domain-element is mapped is its *value*. (Distinct domain-elements may have the same value with respect to a given function.)

The functional conception of meaning, then, is the idea that for every meaningful utterance, there is precisely one meaning to which it is mapped. According to this idea, to speak is *ipso facto* to determine a meaning, some unique content which the utterance has. (For our purposes here, we will focus on propositional content, so that to determine a meaning is to determine what is said to be the case by the utterance.) To understand a speaker is to know

which meaning is determined by their utterance. Misunderstanding arises when audience and speaker do not take the same meaning to be determined. (In such a case, at least one of them must be mistaken.)[1]

Now, were the rules of language — of semantics and of composition — such that past commerce with language *guaranteed* how that commerce must be carried into the future, if they were aptly construed as rails-to-infinity, then it would be the case that co-ordinating one's words in a certain way (in a certain context) *determined* how they must be understood. Understanding a novel utterance would be a matter of *computing* its meaning, deriving it from a knowledge of those rules.

But, in Chapter 1, we eschewed such a picture of language. Language does not run along rails, according to patterns which demand their own continuation. Rather, our attunement to one another in language is something which we exploit, more or less creatively, in order to make ourselves understood. In speaking, one doesn't, as it were, "automatically" grind out a content.

[1]The "meaning" here determined is often treated as an entity — a Fregean "thought" (Frege, 1948: 214n.; 1956: 292), perhaps, or a "proposition" (considered as a truth-evaluable, abstract entity of some kind, or sometimes more specifically as a set of "possible worlds"). This entity serves as the value of the function from speech acts to "meanings". One might worry that this idea of meaning-as-object is misleading in a manner analogous to that in which the picture of rules-as-rails is. But my aim in this section is not to argue against meaning-entities. My targets are instead certain ideas about the *determination* of what is meant by an utterance. (My inclination is to think that talk of meaning-entities could be innocent, just as talk of rules of language can be, provided it is purged of misleading pictures.)

It has been usual in 20$^{\text{th}}$- and 21$^{\text{st}}$-century Anglo-American philosophy to treat propositional content as fundamental, and to treat other kinds of content — questions, commands, etc. — in terms of this. That assumption is one which was strongly challenged by the philosophers associated with OLP. (It is also a key target in Avner Baz's (2012) contemporary defence of OLP.) However, even if it is not the fundamental use of language, we *do* state things to be the case. (This is not something I would expect many proponents of an OLP approach to deny; it is just that making statements is one of the *many* things that we do with language.) My target here is a particular picture of what making a statement involves.

Rather, one relies on one's attunement to the other to communicate. We must rely on the training of others in the faith that they will take our speech as we intend: that the natural reactions inculcated in them will bring them to respond in a way that furthers our projects, whether that be making them aware of something, or apologizing, or entreating or commanding action, or whatever.

In the following sections, I want to expand and clarify my opposition to these ideas about meaning-determination. The key claim here is that when language-use is seen properly, what is meant by a particular utterance made at a particular time is something that must be *negotiated* between speaker and audience; this negotiation is not something which aims at a fact-of-the-matter determined independently of that negotiation. Neither does it have the role of determining such a fact-of-the-matter. These potentially inconclusive negotiations are highly illuminating of the nature of language, of deep disagreement, and of The Method.

3.1.3 Two notions of meaning

We start by distinguishing two uses of the word 'meaning'. First, there is meaning in the sense of the question: 'What does this word ('pusillanimous', say) mean?' Second, there is meaning in the sense of the question: 'What do you mean by that?', asked when someone has uttered words with which I am perfectly familiar.

The first question can be answered with a statement of a rule, in the 'rule of thumb' sense of Chapter 1. 'Meaning' in this case can be thought of as a

propensity to be put to particular uses, a suitability for certain communicative tasks. A hammer is suitable for diverse tasks in a carpentry workshop as a result of its particular physical properties. Similarly, because of the training that English speakers have had in its use, the word 'pusillanimous' is suitable for communicating certain descriptions, accusations, expressions of derision, etc.[2]

The second question asks what particular thing a person was doing with a given word, how they have taken up and leveraged that potentiality. For the avoidance of confusion, let's call meaning in this second sense, as applied to what was being done with a particular utterance, 'what was said'.

Meaning in the first sense is independent of any particular context of use, from any particular instance of the word's being put to use. A signpost lying on the carpenter's table is not directing anybody anywhere, though it has the potential to do so, given how people are in fact trained in the use of signposts. Similarly, a word as it appears in the dictionary is not being put to any particular use. It is meaningful there in the sense that one could (on the basis of the information imparted in the dictionary) put it to particular uses. (Of course, this propensity shouldn't be thought of as *logically prior* to individual uses. This potential is sustained by our continuing to go on as we do, and our training others to do so.)

Meaning in the first sense underdetermines what is said, for depending on

[2]It is in this way that I would understand the oft-cited *PI* §43: 'For a *large* class of cases of the employment of the word "meaning" — though not for *all* — this word can be explained in this way: the meaning of a word is its use in the language.' By 'use' here I understand the suitability for being put to *particular* uses; think of the "use" or function of a hammer. ('Use' translates 'Gebrauch', which would also bear translation as 'custom'; the kind of custom that can be explained by stating a rule, in the 'rule of thumb' sense.) Cf. §421: 'Regard the sentence as an instrument, and its sense as its employment.'

who produces an utterance, when, and under what circumstances, what they are reasonably understood to be doing with their words will differ. ('In this section, we will gradually crescendo', said by the conductor to the orchestra during a rehearsal, is an instruction. But if uttered by the second oboe to the first, it may instead be a *prediction* of what the conductor will instruct them to do.)

In this section, we are concerned not with meaning in the first sense, but with *what was said* by a given utterance, in a given context. Training in a language attunes one to one's fellows in inclinations to understand, to hear certain meanings in certain words. But the interaction of context with the principles of this training does not determine a unique value for what is said. This is most visible where *mis*understanding occurs.

3.1.4 Negotiating what was said

Communication is a matter of leveraging our attunement to one another. Where communication succeeds, we identify what is said with what both the speaker and the audience understood to be asserted (requested, promised, etc.). (The signpost was placed to direct people to Headington, and it in fact succeeds in doing so. It says that one is to go this way to get to Headington.) But what about where there is a failure, where the attempt to exploit our attunement to the other is less than completely successful? An example: I am mapping out some walking trails in the Alps. In order to help hikers to follow the trails, I paint marks like the following on rocks and tree stumps:

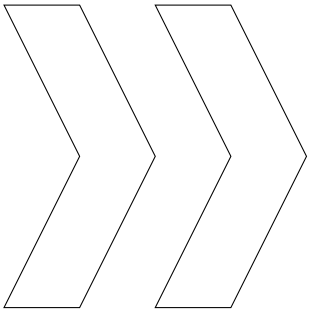

Now, my intention here is to direct hikers to the right (treating the mark as pointing like an arrow). But many of the hikers who walk my trails will have hiked in the Spanish Pyrenees. The trails there use marks much like this to represent, diagrammatically, a path bending to the *left*.[3] As a result of their training, these hikers will react to my signs in such a way that they will get lost.

On the functional conception, we are faced with a sharp divergence of options. Unless the sign is meaningless, exactly one meaning is determined: it directs the walkers either to the left or the right. So; does the pre-existing Pyreneen custom determine a meaning for my marking, so that I am thereby using it incorrectly? Or does my intention start a new custom, so that the hikers are misinterpreting me? What I want to suggest is that there is no right answer to this question. What we have here is simply a failure of attunement, resulting from a difference in training. That is sufficient to explain the breakdown in communication, without positing a unique meaning determined in one way or the other by the markings, which either I or the hikers fail to identify.

[3]FEDME, 2021: 27.

A linguistic example, taken from Bezuidenhout (2002: 108). Suppose you are in the habit of reading the daily comics in the newspaper while you have your lunch. Seeing you searching through the magazines on the coffee table, I say to you 'I left the newspaper on the dining room table'. In fact, the paper on the table is incomplete: our housemate has removed the financial section and taken it to work. I know this, but I assume that you are looking for the paper for the comics, as usual. If you are, no problem: you will find what you want, and all will be well. But suppose instead that you have a report due on trends in share prices, and need the financial section. Finding it missing, you will be disappointed.

Now, if we insist on speaking of a determination of what is said, we shall try to isolate a unique content of my context-bound speech act, and to determine whether I spoke truthfully or not. There are a number of things we *might* say. For example, we might say that, because I knew the financial section was missing, what I meant by 'newspaper' was something which the paper minus the financial section satisfies. Or we might say that, because your goal was in fact to find the financial section, this determines an application-condition-in-context for 'newspaper' which the papers on the table do not satisfy. Or we might say that, knowing that the financial section was missing, I ought not to have said that the newspaper was on the table, because what that literally means is false. But what I want to bring out is that we need not say *any* of these things. As is evident from the description given in the last paragraph, the thwarting of your purposes is intelligible without ever needing to appeal to a unique content which my utterance had in context. Because of our shared training in language and knowledge of the

world, I succeed in directing you to the table. But lacking knowledge of why you wanted the newspaper, I directed you erroneously.[4]

Once we dislodge the idea that language-use runs along "rails-to-infinity", we ought to give up the idea that it is a machine for determining specific contents-of-utterance (even in context). Determination of a unique meaning is, in Wittgenstein's idiom, a wheel that is idling, disconnected from what is needed to explain minor tragedies like that of the newspaper:

> Here I'd like to say: a wheel that can be turned though nothing else moves with it is not part of the mechanism.
>
> (*PI*, §271)

I am arguing that the functional conception of meaning is not required either to explain understanding or misunderstanding. Where there is understanding, our attunement allows us to converge on what we take to have been said, without our needing to assume that the utterance-in-context was mapped to this value independently of our converging on it. Where there is misunderstanding, a failure of attunement explains our frustrations without having to posit a unique, independently determined content-of-utterance which at least one of us has failed to grasp.[5]

[4]There are parallels between what I say here and Avner Baz's (2012) highly illuminating treatment of knowledge, of our uses of the word 'know'. I am unsure whether Bezuidenhout would accept my reading of her example.

[5]An important manifestation of the functional conception of meaning is the way that it is taken up in the field of formal semantics. (Heim & Kratzer (1998) is a representative example from that research programme, which has a distinguished Fregean heritage.) These researchers use formal resources (usually set theory) to provide what we might call *algorithmic* theories of the semantics of natural language; algorithmic, because the meaning of any given utterance is viewed not just as uniquely determined, but *computable*, in advance, from a fixed stock of semantic principles.

The algorithmic conception has come under pressure from authors (among them Travis

The functional conception is *distorting* in the way it invites us to view misunderstanding. Consider this example: suppose that I betray a friend's trust to achieve some minor personal gain. When I meet them again, I say, sincerely: 'I deeply regret the hurt my actions caused you.' Is this an apology? (Note that I did not say that I regretted acting as I did, just that it had the effect on you that it did.) I want to resist thinking about this along the following lines: either I did apologize, or I didn't; some features of the situation and what I said determine which. Instead, each of us must make a judgment as to whether I have apologized or not, as to whether my speech alters the relationship between us as an apology would. (We might say, we must judge whether to *make* of my speech an apology by allowing the relationship so to be changed.) These judgments are not of a sort that must be responsible to some independent fact, but they are nevertheless

---

(2008b), Recanati (2004, 2005), Bezuidenhout (1997, 2002) and Lahav (1989)) who emphasize the extent to which natural language expressions are radically context-sensitive. Many of these authors explicitly link their work to that of Austin and Wittgenstein. Austin, for example, writes that

> The total speech act in the total speech situation is the *only actual* phenomenon which, in the last resort, we are engaged in elucidating. (*HDTW*: 147, emphasis in original)

and that '[s]entences are not *as such* either true or false'; we must take account of 'the circumstances in which it is uttered' (*SS*: 111, emphasis in original). (Cf. *PI* §117; *OC* §§10, 348-50. This strand of Wittgenstein's thought has been traced back to the composition of the *Tractatus* (*TLP* §3.326); Conant 1998: 238.) The wrong way to construe this emphasis on context, however, would be to see meaning as determined not just by 'bottom-up', algorithmic processes, but also by 'top-down', non-algorithmic ones (as, for example, Recanati (2005) seems to; the language of 'bottom-up' vs. 'top-down' determination is his). The OLP criticism goes deeper, to the functional conception which underlies algorithmic meaning-theories. The point is not simply that the meaning-function is non-computable, but that we need not posit a meaning function *at all*.

For the way in which the OLP treatment of context-sensitivity undercuts H. P. Grice's (1991) use of his notion of 'implicature' in criticism of certain OLP arguments, see Travis (2008a).

reasoned judgments — judgments for which relevant considerations can be adduced and adjudicated.[6] We may end up reasonably disagreeing about what was said. This does not result in our talking past each other in the river bank/financial bank sense. Rather, there is a reasonable disagreement about a specific topic; about whether you ought to allow my words to restore our relationship.

To speak, then, is to enter into a *negotiation*. No one person is the sole authority on what has been said by a given utterance. Not the speaker: intention cannot absolve all sins. In most contexts, saying 'Oh grow up' just will not be taken as an apology. Not the audience: the hilarity that ensues when a British English speaker goes to Australia and is confused by talk of 'thongs' results from a failure of attunement, a *difference* in training in the use of a word. A British English-speaker's training will not prepare them to hear an Australian's talk of 'thongs' as concerning flip-flop sandals, rather than underwear. This does not determine the Australian-English remarks, when directed at an English-speaker, to be about underwear, rather than about sandals. Neither is "the community" the arbiter of what is said. In the sense in which words have a public meaning, this meaning is a *suitability* for being put to *particular* uses, and this suitability underdetermines what may be said with them on any particular occasion. Insofar as what is said

[6]These judgments of what is meant bear a resemblance to Cavell's treatment of aesthetic judgments (*AP*). The claim that something is beautiful is something for which one must be prepared to give certain sorts of recognizably relevant reasons (rather than retreating to the claims of merely subjective taste). But, in aesthetics, recognizing someone else's reasons as reasons does *not* require me to accept their conclusions. There is such a thing as agreeing, reasonably, to disagree. The similarity to what I am here suggesting about judgments of what has been said is non-accidental, and relates to Cavell's suggestion that a discussion of aesthetic judgment will illuminate OLP's appeal to what we say when. I don't have space to pursue it in detail, however.

exists, it exists between us, in the particular context in which we speak; it must be negotiated between speaker and audience. Which is not to say that there is a uniquely right answer to the question of what is said at which this negotiation should be aiming.[7]

3.2 Negotiation and responsibility

3.2.1 On the hook

The OLP vision of language radically reconceives what it is to say something. It is not simply to draw on the customs of one's community to determine a meaning. It is to navigate an (often messy) web of reasonable expectations and inclinations to understanding in order to make oneself understood.

My intention here is to shake the idea that talking of attunement must precipitate us into an 'us vs. them' mentality, a vision of rational discussion as possible only within the confines of some sufficiently attuned group, and, at best, non-rational influence on those without. We are more-or-less attuned to others in our inclinations to understanding; since each of us has an individual history with our language, we are unlikely to be perfectly attuned to *anyone*. I give no theoretical role to facts about what was meant, of the sort that would precipitate strong incommensurability. If there are two of us painting marks in the Alps, myself and a FEDME trail-keeper, then there are not two indistinguishable sign-types appearing on rocks and tree stumps, some directing people left and others directing people right. (Neither do the signs

[7]This ought to illuminate the way in which the vision of language here defended desublimes the notion of truth (cf. §2.1.1).

snap to one meaning or the other depending on who is looking at them.) There are just some marks on some rocks, which, as a matter of fact, are misleading.

The result of this is that, in speaking, one can be on the hook for more than one realizes. As explained above, where breakdowns in communication occur, we do not need to assign a unique content to what was said. Nevertheless, we often take the question of what was meant in these situations very seriously. Here, the reasons which we give for thinking that a specific thing was said are very often concerned with *responsibility*, responsibility for reasonably-incurred expectations. Take an example from Nancy Bauer (2015: 100):

> Suppose that a man has been seeing the same woman for three months. One night, while they are having a lovely time tarrying over dinner in a romantic setting, he tells her that he loves her. Now imagine that soon afterward the woman learns that the man is also seriously dating someone else. She confronts him. He says, 'Well, when I said that I loved you I just meant that I am fond of you. I didn't mean that I was in love with you.' Let us suppose that the man is absolutely sincere, that he intended to tell the woman that he was, merely, fond of her. [...] What the man intended when he said 'I love you' to the woman under these circumstances is beside the point. What he said is a function of what it is to say these words under these circumstances.

Now, I wouldn't endorse Bauer's last sentence here, which comes close to

suggesting the ideas of meaning-determination which I have been trying to avoid. (I would like to take this charitably, as an unfortunate wording, rather than an endorsement of those ideas on Bauer's behalf.) But the point I want to take from this example is the following: the woman can give a reasonable argument to the effect that what it means for the man to say 'I love you' under these circumstances is incompatible with his seriously seeing someone else. This is an argument based on the reasonable expectations that he is responsible for incurring by so speaking. Each of us must make a judgment: to what extent was the man responsible for incurring these expectations, and so what did he say?

The imperfect attunement between the man and the woman does not precipitate strong incommensurability. Rather, it is up to the man to appreciate the significance of what he said to the woman, to judge whether her understanding is reasonable, and, if so, to take responsibility for it. In the next section, I want to expand this idea to explain how The Method can bring people to a new appreciation of the significance of what they do and say.

3.2.2 Introducing Ma Beckoff

In order to connect these ideas with moral disagreement, I want to draw on an example which has been discussed by Naomi Scheman (2017).[8] In

[8]Although my way of approaching this example differs slightly from Scheman's, I do not think it is at odds with her treatment of it. Scheman makes use of the idea of different (that is, not merely additive) critical perspectives. These are perspectives for which we can, and are expected to, give reasons, reasons recognizable as such (but not necessarily universally compelling: '[s]ometimes we just "agree to disagree" recognizing the limits to our shared practices (hence, to the objectivity of our judgments)' *op. cit.*: 95). Like my

the 1988 film adaptation of Harvey Fierstein's *Torch Song Trilogy*, Arnold Beckoff is a gay man whose lover is buried alongside his father in the family's burial plot. His mother refuses to accept that what he is doing for his lover is the same as what she does for her husband: that they are both *grieving*. Ma Beckoff refuses to describe the pattern which appears in Arnold's life (for example, his visiting Alan's grave to recite the Jewish prayer for the dead) as 'grief'. And we may suppose that she does so because she thinks that the fact that Arnold's actions concern a male lover mean that his feelings cannot be sufficiently like those which she feels for her husband's death.

Suppose we hoped to bring Ma Beckoff round to a different way of seeing things. We want her to appreciate the similarities between her son's situation and her own, the parallels between their suffering for the loss of a non-fungible other, and the love for that other that this suffering discloses. We want her to appreciate these similarities, and to convince her that no weight ought to be given to the major difference around which her refusal to project 'grief' centres: that Arnold 's relationship was a homosexual one. So we employ The Method. We examine the holistic conceptual connections which are thrown up in our reflective inclination to project 'grief' onto Arnold's situation; we examine what is the significance of treating his situation as that of someone grieving. In justifying our projection to Ma, we bring out these conceptual connections.

For example, we might suggest that to deny that Arnold is grieving is to see his expressions of "grief" — what Ma would have called expressions

discussion of the negotiation of what was said, this is strongly reminiscent of Cavell's (*AP*) use of aesthetic judgment as an object of comparison for The Method, mentioned in fn. 6 above.

of grief had she not known for whom Arnold was "grieving" — as somehow counterfeit. If Ma wants to withhold 'grief' from Arnold, then she owes us an explanation of her stance towards his expressions of grief. She will have to see him as duplicitous, or as deluding himself in the sincerity of his performance.

Once we bring up these connections, it is up to Ma Beckoff to disclaim them, or to take responsibility for them. (Of course, we cannot *force* her to do so.) She must take on board our perspective on the significance of what she is saying about her son, decide if it is reasonable, and, if it is, own its consequences. She must assess whether she feels the conceptual tensions which we find in her withholding of 'grief' from her son. This is a distinctly *rational* approach to Ma; we have entered into a negotiation as to what it *means* for her to deny that Arnold is grieving.

Of course, Ma might own the implication that Arnold is either duplicitous or deluded. She might explain that, by withholding 'grieving' from her son's situation, she is expressing her judgment that his relationship with his partner was unnatural and wrong, and thus cannot have the seriousness of grief. (Consider her response to Arnold's claim that, in reading out the Jewish prayer for the dead, he is doing the same thing she is doing: 'I'm reciting Kaddish for my husband; you're blaspheming your religion.') She might say that a heroin addict in recovery can pine for the drug, but not grieve for its loss. The situation here would not be one of incommensurability; here we have something with which we can *disagree*, substantial judgments about homosexual relationships that we can argue against.

In making use of The Method, we must take on board the voices of others, in order to negotiate for what we are responsible in speaking as we do. The

question is then: can you bear that responsibility?[9]

3.3 Weak incommensurability

3.3.1 A matter of taste?

A worry about relativism persists. For there seem to be two ways in which our appeals to Ma Beckoff might fail. She might not see what we see in the situation; she may not see our perspective as reasonable. Or she might (as was suggested above) own the responsibility for the conceptual connections implicated in her refusal to project 'grief' to Arnold, fully and reflectively endorsing the suggestion that he is a second-class citizen. What are we to do, faced with such an impasse?

In emphasizing that The Method involves an invitation to project our words, an invitation which we must take up reflectively by responding carefully to the realities of particular situations, we have ensured that The Method remains responsible to reality. But many of the ways in which we are attuned concern not just the aspects of reality we are prepared to notice,

[9]The perspective we are taking on language has the effect of blurring the boundaries between linguistic and non-linguistic meaning. (Certain acts have an almost "conventional" significance; given who we are, our particular cares and commitments, they cannot fail but to *mean* something. Consider how shallow a utilitarian explanation of the wrongness of, for example, betraying a sibling can seem. What such an explanation fails to appreciate is the *meaning* of such an action in the context of a familial relationship, as such things occur in our form of life. This is a meaning which justifies the sibling's hurt feelings, even if the betrayal has no other "adverse effects". This idea, of extending OLP's characteristic Method from the explicitly linguistic to the significance of action more generally strikes me as a fruitful way of understanding what Cora Diamond is up to in her 'Eating Meat and Eating People' (1991b). There, she invites us to consider what it would mean, to us, to eat a person. On this basis, we are to consider what unacknowledged significance might attach to our eating an animal. The question one must then ask oneself is whether one can live with what it means for one to do so.

but the ways in which those differences *matter* to us. The worry is that, where there are differences in what matters to us, there will be differences in how we project certain words; in particular, weighty, morally-laden words like 'murder' or 'love'. The critical reflection of The Method seems to rely on what *does* in fact matter to us, not what *should* matter to us. On the basis of particular concerns and values, a homophobic community might refuse to speak of 'love' when the context involves two members of the same sex. From our perspective, they seem to have clustered, in their attunement, around the *wrong* lines of projection. But to insist on this would seem to apply a normative standard to the inclinations in which we are attuned. It is unclear what could provide the basis for such a standard; in Chapter 1, the ways in which are *in fact* attuned were enlisted as the source of normativity in language-use. What source could be found for a normative standard to which attunement itself is subject, a standard for which inclinations and concerns we *ought* to have?

Here we would have incommensurability in a weaker sense. It is not that our words don't engage with each other; in a sense, we can be taken to mean the same things by 'love' or 'grief' (just as people with different palates can mean the same thing by 'delicious'). Nevertheless, the way in which these words are projected are bound up with our sense of what matters, and so we do not a have a common standard for how to project (e.g.) 'love'. Simplifying, Ma will extend 'love' to relationships of which she approves, we will extend 'love' to those of which we approve. On what rational grounds can we appeal to Ma to adjust her attitude, given that this attitude seems to be part of the way in which she is attuned to others, though not to us?

Judgments about love and grief begin to look like judgments of taste; if I like strawberries, but you do not, the disagreement is not deep.

3.3.2 Deep disagreement

We should not overlook the fact that some very weighty ethical concepts do supervene on comparatively local features of culture and sentiment. A medieval notion of chivalry (or its loose Japanese analogue, *bushidō*) rests on cultural concerns which do not obviously engage with the structure of a modern European life. Conversely, ties of kin and friendship are important to us, given the sorts of lives we lead. They did not *have* to be, perhaps; we might imagine biologically rather similar beings who do not form family units. This does not diminish the importance of these ideas for those who *are* immersed in these lives.

It does not follow from this that every difference of opinion, even if practically intractable, stems from a difference in 'form of life', and so becomes insulated from rational appeal. We can approach Ma Beckoff by trying to get her to appreciate internal tensions in what matters to her, exhibited in the way that she is inclined to project her words.

Consider first this suggestion: perhaps Ma just lacks the empathetic capacity to appreciate the reality of Arnold's grief, the depth of his pain. When I introduced the idea of attunement in Chapter 1, I suggested that the training one receives in learning a language can feed back into one's discriminatory capacities (as a musical education can refine the sensitivities to pitch on which it builds). The French speaker who cannot see the point in distinguishing

between moths and butterflies might, accordingly, struggle actually to tell the difference when faced with a butterfly and a moth (as we might struggle to distinguish individual butterflies).

But perhaps we could bring them to the point in another, more roundabout way. Perhaps we can teach them, for example, about the way that moths have been known to eat clothing, in the hope that *this* way of leveraging our attunement to them allows our attunement to be deepened elsewhere, bringing them to share our capacity for distinguishing between moths and butterflies. (We should, of course, also be prepared for them to bring *us* around; they *may* have reasons for their way of speaking.)[10]

A similar approach might be engaged in with Ma Beckoff; we try to sharpen her perception of Arnold's pain by exploiting other ways in which we are attuned to her, trying to get her to see what we see. One might say: there is such a thing as being blinded by ignorance, or by prejudice. Some important distinctions cannot simply be pointed out to one who is not prepared to see them. If so, preparing them to see them, to see why their life is impoverished by failing to see them, is a *rational* appeal to that other.[11]

The worry above was not based on failures of perception, however; the

[10]They may intend to register the fact that the phylogenetic classification of moths and butterflies is far from straightforward.

[11]My remarks here might offer an interesting perspective on Iris Murdoch's arguments that moral philosophy needs to respect the possibility that *vision*, and not just choice, is morally important. Where moral differences are differences of vision, rather than differences of choice given the same facts, 'a moral concept seems less like a movable and extensible ring laid down to cover a certain area of fact, and more like a total difference of *Gestalt*' (1956: 40-1). The suggestion, as I understand it, is that one's ability to appreciate the facts of a given matter are not independent of one's values. A view of what the world is like can be a moral view, and vice versa (cf. (Diamond, 1996: 89); also (1991c; 2010)). This will be unsurprising if one's values are a way in which one may be more or less attuned to others, an attunement which another may have to work to deepen if you are to see what they see.

question it asks is: what if Ma Beckoff appreciates Arnold's pain, but *does not care*? I began by discussing failures of vision because I want to suggest that a failure of concern can be approached in the same way. We can try to get Ma to appreciate internal tensions in what matters to her. (For example: perhaps Ma condemns her son's homosexuality primarily because she cares for him as a *son*, caring for his moral well-being. We might attempt to re-orient this concern, to leverage that care for her son to bring her to appreciate the harm which she causes him by denying his grief.) As with a failure of perception, we approach Ma's failure of concern by working to *deepen* our attunement to her; exploiting the shared aspects of our sense of what matters in order to approach those for which we diverge.

Differences in attunement are bridgeable. (If they were not, the skill of using a language would not be learnable.) And this effort of bringing another into attunement need not be a *non-rational* appeal. It may draw upon ways in which we *are* attuned. (This is exactly what an explanation of meaning does, given to one whose training has prepared them for the explanation (recall §1.3).)

One can see *Torch Song Trilogy* itself as working in this way; as an attempt to make use of the ways in which its audience *is* attuned to the playwright in order to bring us into deeper attunement. Fierstein presents us with a richly detailed portrayal of a man in pain. His dramatic decisions are made in the hope that he can exploit certain aspects in which we are attuned to one another (as human beings who love, and grieve) in order to bring us to see something which, as of yet, we cannot or will not see.

Where deep ethical disagreement is construed as an impasse between

incommensurable moral frameworks, making a *rational* appeal to another will seem possible only if such a framework is shared. Anything else will be relegated to the level of 'mere persuasion'. But without that conception of ethical disagreement, such a dichotomy — internal argument, vs. external propogandizing — is not forced upon us. We are more or less attuned to one another. Where attunement fails along one dimension, we may hope to deepen it by leaning on a dimension in which we are more attuned. That, I suggest, is how one should view a film like *Torch Song Trilogy*. Exploiting one's attunement to another can be a way of *deepening* that attunement. The development of an attunement to the other is continuous with the provision of arguments of a more traditional sort.[12]

It is important here not to annex to the idea of a rational appeal a kind of supernatural force: could we just properly express our reasons, anybody would *have* to see their force, to change their minds. Our arguments may well fall on deaf ears. You cannot force someone to appreciate conceptual tensions of the sort I have been talking about. (They are not, as I argued in Chapter 2, matters of material implication, of necessary conditions.) Ma Beckoff may well be able to maintain a formal, logical consistency, since there *is* a difference between her situation and Arnold's. Even accepting that Arnold was in love would not logically compel her to accept that he was grieving: hers was a heteroromantic love, his a homoromantic one. The most

[12]Given more space, I would link these remarks in greater detail to what Harcourt (2015: 210) refers to as 'expansionist' conceptions of moral thinking. Cora Diamond (1991a; 1991c; 1991d) and Alice Crary (2007; 2016) are among a number of philosophers who argue that (some) literature has a place in moral philosophy, as more than a stock of handy examples for pumping intuitions and applying general moral principles, or as non-rational appeals to the reader where rational appeal might fail. It is no coincidence that these authors' readings of Wittgenstein are similar to that which animates this thesis.

we can do is make her aware of the conceptual space, of what her words might reasonably be taken to mean. It is up to her to shoulder that responsibility, and up to her *how* she shoulders it.

There are no guarantees here. A rational appeal may fall on deaf ears. Good reasons may go unheeded. That is not relativism or subjectivism. That's life.

3.4 Conclusion to Chapter 3

The worry registered at the end of Chapter 2 was that failures of attunement would precipitate incommensurability; a situation in which we talk past each other. In this chapter, I have tried to explain why this is not the case.

Where attunement fails between two parties, the result is not that two different contents are determined for their utterances, so that they fail to contradict each other. True, communication may falter, as it did between the couple in Bauer's example above. But the failure is more mundane than that which would be involved in strong incommensurability. The man and the woman do not speak different languages; it is just that, in speaking to people to whom one is imperfectly attuned (and there is likely nobody to whom one is *perfectly* attuned), one can reasonably be taken to have said more than one intended. One is responsible for more than one might realize; discussion with others can help one to appreciate that responsibility.

When one employs The Method, one enters into a *negotiation* as to what it would mean to say that p. This negotiation does not aim at a unique, antecedently fixed fact-of-the-matter as to what was said. Nor does it have

the role of fixing a unique fact-of-the-matter. There may be more than one reasonable way to understand the importance of what someone has said.

Far from trivializing (moral) disagreement, I suggest that OLP gives us a more responsible perspective on negotiating difference. One cannot just speak for oneself. Language is a tool for communication, and if one is to use it successfully, one must respect the reasonable understandings of those with whom one would communicate. We ought to be open to the possibility that others might appreciate, and help us to appreciate, unacknowledged significance in what we do and say.

You cannot force others to accept as important the similarities which guide your projections. You cannot, for example, force them to extend the language of suffering, with all its moral connotations, to people other than themselves, or other than their close companions, or to non-human apes. But neither, as Diamond (1991a: 295) notes, can you *force* someone to accept a formally sound argument. The refusal to accept certain conceptual connections may manifest a similar failure of rationality.

Similarly, a person might accept the connections, and yet not care for them. Perhaps we will be able to bring them around by leveraging other aspects of our attunement. Perhaps not. (Perhaps *they* will bring us round in a similar fashion, showing us why some distinction ought (or ought not) to matter to us.) But even if we do not, in the end, agree, The Method will at least allow us to see where the differences lie.

The orientation towards disagreement to be found in OLP is not relativistic, but *realistic*, in the sense of 'realistic' which contrasts with 'fantastic',

with wishful thinking.[13] It accepts that disagreement may be practically intractable. It refuses to submit to a fantasy of rational argument which sees the business of (moral) philosophy to be giving arguments which must, on pain of irrationality, persuade anyone who properly understands them. But it also avoids recoiling into a picture that trivializes moral disagreement by seeing deep disagreement as precipitating incommensurability.

[13]Diamond (1991e) discusses the importance of this and related senses of 'realistic' for realism in philosophy, for the idea that whatever conventions we have for representing some domain do not settle what is *true* in that domain.

Conclusion

A Plea for OLP

Philosophers discuss a wide variety of topics. In doing so, they employ the words of the natural languages in which they have been trained. Some of this language is decidedly technical, and training with it must wait until one has grown up and embarked upon the discipline. But even these technical terms have their connections to more usual ones. Extra-philosophical talk of what is possible is endemic; it is in terms of this that we come to get a grip on how to manipulate the philosophers' word, 'necessary'.

Philosophical discussion is not detached from the rest of our lives with language. It is a part of it. When philosophers use the words of natural language, even highly technical ones, what they say draws on a training in language which began almost at birth and which has continued into their adult lives. Furthermore, we should like, I think, for philosophical discussion of, for example, possibility, or knowledge, to be concerned with those phenomena as they appear in ordinary life, where people talk about the possibility of rain, or insist that someone knows where the keys have been

hidden.

All this is to say, the word 'knowledge', in the mouth of philosophers as much as in the mouths of everyone else, is *our* word. This is the starting point of OLP. (One would hesitate to call it an insight. It is little more than a reminder.) The result of this admission, that philosophers speak the same language as the rest of us, is not that philosophy must be shackled, held back from claiming anything which runs contrary to "common sense". The impression that it might comes from ways of thinking about language-use which, I have suggested, are subjected to deep criticism in the work of ordinary language philosophers.

It is tempting to view a language as a shared system for determining meanings according to fixed rules. To treat it as such might be illuminating, because formally tractable: it is one way of modelling a complicated human capacity. But it is also, I have argued, deeply distorting of the nature of language-use and of rational disagreement. In this thesis, I have developed an alternative vision. This vision, found in the later Wittgenstein as refracted through Cavell, is of proficiency with language as a fluid, constantly evolving skill.

In light of this, we have seen that the charactistic Method of OLP, of examining what we should say when, does not serve to abstract "the rules" according to which we use words, rules which one must follow on pain of speaking falsely or nonsensically, and to which one can appeal to in order to dismiss philosophical claims as false or nonsensical. (Whatever rules one learns in learning a language are rules *of thumb*.) Rather, The Method involves an invitation to project our words into different contexts. This is an

invitation to attend to the *world*, to the particularities of the situation at issue. By exercising one's skill in the use of the language, by seeing what one feels one *ought*, on reflection, to say, and what might (perhaps usefully) be expressed by saying something else, conceptual connections may be uncovered which illuminate the ways in which this situation is like, and unlike, others — and how it might be appropriate to respond to it. The world is endlessly various; far from being inherently conservative, The Method of OLP prepares us to have our eyes opened by the new and the surprising. It eschews a vision of language-use which sees situations as pre-packaged for easy assimilation by our linguistic rules, and gives itself over to complexity and nuance. In our approach to particular cases, we are aided by taking on board a plurality of perspectives, a consideration of all of the many things which might reasonably be understood by the projection of a particular word into a particular context.

The vision of language which stands behind The Method promises to reorient our perspective on disagreement, on the responsible ways in which one may approach difference. It presents an unflinchingly realistic picture of such disagreement. Disagreement is endlessly complicated by the fact that one's interlocutor is an *other*, inheriting an individual history with your shared language. As such, disagreement is often deep, often practically intractable, and sometimes demands that we approach the other in creative ways. These ways of approaching the other are likely to fall outside the confines of what is usually accepted as 'philosophy'; they are nevertheless *rational* appeals to the other. Like any use of a natural language, these appeals are ways of leveraging one's attunement to the other in order to bring them to see things

as you do.

Although the discussion of rule-following in Chapter 1 illustrates how one may come to the OLP vision of language, most of this thesis has been concerned with disarming worries arising from what I take to be misunderstandings of OLP's relation to natural language. As such, it would be too much to hope that OLP has here been established as a *more* fruitful approach to philosophizing than more mainstream contemporary approaches. I hope, however, that by clearing up certain misunderstandings about its Method, and by illustrating the richness of its vision of natural language and disagreement, I have made the case that an ordinary language approach at least deserves a seat at the table.

Primary texts cited

Abbreviations for works by Austin cited

PFE 'A Plea for Excuses', in Austin et al. (1961).

HDTW *How to Do Things with Words* (Austin & Urmson, 1962).

SS *Sense and Sensibilia* (Austin & Warnock, 1964).

Abbreviations for works by Cavell cited

AWP 'The Availability of Wittgenstein's Later Philosophy', in Cavell (1976).

AP 'Aesthetic Problems of Modern Philosophy', in Cavell (1976).

CR *The Claim of Reason* (Cavell, 1999).

EW 'Excursus on Wittgenstein's Vision of Language' (Cavell, 2000).

DD 'Declining Decline' (Cavell, 2013).

Abbreviations for works by Wittgenstein cited

TLP	*Tractatus Logico-Philosophicus* (Wittgenstein, 1922).
WLC	*Wittgenstein's Lectures, Cambridge* (Wittgenstein et al., 1982)
BB	*The Blue and Brown Books*, in Wittgenstein (1998).
CV	*Culture and Value*, in Wittgenstein (1998).
Z	*Zettel*, in Wittgenstein (1998).
PI	*Philosophical Investigations* (Wittgenstein et al., 2009).
PPF	*Philosophy of Psychology — A Fragment*; previously known as *Philosophical Investigations Part II*. In Wittgenstein et al. (2009).
OC	*On Certainty* (Wittgenstein et al., 1972).

Bibliography

Austin, J. L., & Urmson, J. O. (1962). *How to Do Things with Words*. William James lectures. Oxford: Clarendon Press.

Austin, J. L., Urmson, J. O., & Warnock, G. J. (1961). *Philosophical Papers*. Oxford: Clarendon Press.

Austin, J. L., & Warnock, G. J. (1964). *Sense and Sensibilia*. London: Oxford University Press.

Baker, G. P., & Hacker, P. M. S. (2010). *Wittgenstein - Rules, Grammar, and Necessity : essays and exegesis of 185-242*. Malden, MA: Wiley-Blackwell, 2nd, exten ed.

Bauer, N. (2015). What Is to Be Done with Austin? In *How to Do Things with Pornography*.

Baz, A. (2012). *When Words Are Called For : a defense of ordinary language philosophy*. Cambridge, Mass.: Harvard University Press.

Bezuidenhout, A. L. (1997). The Communication of De Re Thoughts. *Noûs*, *31*(2), 197–225.

Bezuidenhout, A. L. (2002). Truth-conditional pragmatics. *Nous-Supplement: Philosophical Perspectives*, *16*(1), 130–134.

Cavell, S. (1976). *Must We Mean What We Say?*. Cambridge: Cambridge University Press.

Cavell, S. (1999). *The Claim of Reason: Wittgenstein, skepticism, morality, and tragedy*. New York: Oxford University Press, new ed.

Cavell, S. (2000). Excursus on Wittgenstein's Vision of Language. In A. Crary, & R. J. Read (Eds.) *The New Wittgenstein*, (pp. 21–37). London: Routledge.

Cavell, S. (2013). Declining Decline. In *This New Yet Unapproachable America*, (pp. 29–76). Chicago: University of Chicago Press.

Conant, J. (1998). Wittgenstein on meaning and use. *Philosophical Investigations*, *21*(3), 222–250.

Crary, A. (2007). *Beyond Moral Judgment*. Cambridge, Mass. ; London: Harvard University.

Crary, A. (2016). *Inside Ethics : on the demands of moral thought*. Cambridge, Massachusetts.

Crary, A., & de Lara, J. (2019). Who's Afraid of Ordinary Language Philosophy? A Plea for Reviving a Wrongly Reviled Philosophical Tradition. *Graduate Faculty Philosophy Journal*, *39*(2), 317–339.

Diamond, C. (1988). Losing Your Concepts. *Ethics*, *98*(2), 255.

Diamond, C. (1991a). Anything but Argument? In *The Realistic Spirit : Wittgenstein, philosophy, and the mind*, (pp. 291–308). Cambridge, Massachusetts.

Diamond, C. (1991b). Eating Meat and Eating People. In *The Realistic Spirit : Wittgenstein, philosophy, and the mind*, (pp. 319–334). Cambridge, Massachusetts.

Diamond, C. (1991c). Having a Rough Story about What Moral Philosophy Is. In *The Realistic Spirit : Wittgenstein, philosophy, and the mind*, (pp. 367–381). Cambridge, Massachusetts.

Diamond, C. (1991d). Missing the Adventure: Reply to Martha Nussbaum. In *The Realistic Spirit : Wittgenstein, philosophy, and the mind*, (pp. 309–318). Cambridge, Massachusetts.

Diamond, C. (1991e). Realism and the Realistic Spirit. In *The Realistic Spirit : Wittgenstein, philosophy, and the mind*, (pp. 39–72). Cambridge, Massachusetts.

Diamond, C. (1991f). Wittgenstein and Metaphysics. In *The Realistic Spirit : Wittgenstein, philosophy, and the mind*, (pp. 13–38). Cambridge, Massachusetts.

Diamond, C. (1996). "We Are Perpetually Moralists": Iris Murdoch, Fact and Value. In M. Antonaccio, & W. Schweiker (Eds.) *Iris Murdoch and the Search for Human Goodness*, (pp. 79–109). Chicago: University of Chicago Press.

Diamond, C. (1999). How old are these bones? Putnam, Wittgenstein and verification. *Proceedings of the Aristotelean Society, Supplementary Volumes*, *73*(1), 99–134.

Diamond, C. (2010). Murdoch the Explorer. *Philosophical Topics*, *38*(1), 51–85.

Egan, D. (2019). *The Pursuit of an Authentic Philosophy*. Oxford: Oxford University Press.

FEDME (2021). *Manual de Señalización de Senderos*. Federación Española de Deportes de Montaña y Escalada, digital ed.
URL https://misendafedme.es/manual-de-senalizacion-de-senderos-gr-pr-y-sl/

Flew, A. (1986). Apologia pro philosophia mea. In S. Shanker (Ed.) *Philosophy in Britain Today*, (pp. 72–95). Albany: State University of New York Press.

Flew, A. G. N. (1955). Philosophy and language. *Philosophical Quarterly*, *5*(18), 21–36.

Frege, G. (1948). Sense and Reference. *The Philosophical Review*, *57*(3), 209–230.

Frege, G. (1956). The Thought: A Logical Inquiry. *Mind*, *65*(259), 289–311.

Gellner, E. (1959). *Words and Things : a critical account of linguistic philosophy and a study in ideology*. London: Gollancz.

Goodman, N. (1983). *Fact, Fiction and Forecast*. Cambridge, Massachusetts, fourth ed.

Grice, H. P. (1991). *Studies in the Way of Words*. ACLS Humanities E-Book (Series). Cambridge, Mass.: Harvard University Press, 1st harvar ed.

Gustafsson, M. (2005). Perfect Pitch and Austinian Examples: Cavell, McDowell, Wittgenstein, and the Philosophical Significance of Ordinary Language. *Inquiry*, *48*(4), 356–389.

Gustafsson, M. (2011). Familiar Words in Unfamiliar Surroundings: Davidson's Malapropisms, Cavell's Projections. *International journal of philosophical studies : IJPS*, *19*(5), 643–668.

Hanfling, O. (2000). *Philosophy and Ordinary Language : the bent and genius of our tongue*. Routledge studies in twentieth century philosophy. London ; New York: Routledge.

Harcourt, E. (2015). Literature, Moral Thinking, and Moral Philosophy. In S. G. Chappell (Ed.) *Intuition, Theory, and Anti-theory in Ethics*, Mind Association occasional series, (pp. 210–228). Oxford.

Heidegger, M., Macquarrie, J., & Robinson, E. S. (1962). *Being and Time*. Oxford: Blackwell Publishing Ltd.

Heim, I., & Kratzer, A. (1998). *Semantics in Generative Grammar*. Blackwell textbooks in linguistics ; 13. Malden, MA ; Oxford: Blackwell.

Kafka, F. (2005). The Metamorphosis. In N. N. Glatzer (Ed.) *The Complete Short Stories*, (pp. 95–149). London: Vintage.

Kripke, S. A. (1982). *Wittgenstein on Rules and Private Language : an elementary exposition*. Oxford: Blackwell.

Lahav, R. (1989). Against compositionality: The case of adjectives. *Philosophical Studies*, *57*(3), 261–279.

Leibniz, G. W., Remnant, P., & Bennett, J. (1996). *New Essays on Human Understanding*. Cambridge, second edi ed.

MacKinnon, C. A. (1993). *Only Words*. Cambridge, Mass.

Malcolm, N. (1981). Moore and Ordinary Language. In V. C. Chappell (Ed.) *Ordinary Language : essays in philosophical method*, (pp. 5–23). New York: Dover Publications.

Mates, B. (1981). On the Verification of Statements About Ordinary Language. In V. C. Chappell (Ed.) *Ordinary Language : essays in philosophical method*, (pp. 64–74). New York: Dover Publications.

McDowell, J. (1984). Wittgenstein on following a rule. *Synthese*, *58*(3), 325–363.

McDowell, J. (2000). Non-Cognitivism and Rule-Following. In A. M. Crary, & R. J. Read (Eds.) *The New Wittgenstein*, (pp. 38–52). London: Routledge.

Mulhall, S. (2004). Philosophy's Hidden Essence. In E. Ammereller, & E. Fischer (Eds.) *Wittgenstein at Work : method in the Philosophical Investigations*, (pp. 63–85). London: Routledge.

Mulhall, S. (2014). Inner Constancy, Outer Variation: Stanley Cavell on Grammar, Criteria, Rules. *Varieties of Skepticism: Essays After Kant, Wittgenstein, and Cavell*, (pp. 291–310).

Murdoch, I. (1956). Vision and Choice in Morality. *Proceedings of the Aristotelian Society, Supplementary Volumes*, *30*, 32–58.

Nussbaum, M. C. (1990). The Discernment of Perception: An Aristotelian Conception of Private and Public Rationality. In *Love's Knowledge: Essays on Philosphy and Literature.*, (pp. 54–105).

O'Neill, O. (1986). The Power of Example. *Philosophy*, *61*(235), 5.

Papineau, D. (2006). The tyranny of common sense. *The Philosophers' Magazine*, (34), 19–25.

Phillips, D. Z., & Mounce, H. O. (1969). *Moral Practices*. Studies in ethics and the philosophy of religion. London: Routledge & Kegan Paul.

Recanati, F. (2004). *Literal Meaning*. Cambridge: Cambridge University Press.

Recanati, F. (2005). Literalism and Contextualism: Some Varieties. In G. Preyer, & G. Peter (Eds.) *Contextualism in Philosophy : knowledge, meaning, and truth*, (pp. 171–196). Oxford: Clarendon Press.

Scheman, N. (2017). Forms of life : Mapping the rough ground. In H. D. Sluga, & D. G. Stern (Eds.) *The Cambridge Companion to Wittgenstein*, chap. 12, (pp. 383–410). Cambridge: Cambridge University Press.

Travis, C. (2008a). Annals of Analysis. In *Occasion-Sensitivity*, (pp. 65–93). Oxford University Press.

Travis, C. (2008b). *Occasion-sensitivity : selected essays*. Oxford: Oxford University Press.

Uschanov, T. P. (2006). The Strange Death of Ordinary Language Philosophy.
URL https://www.mv.helsinki.fi/home/tuschano/writings/strange/

Waismann, F. (1945). Verifiability. *Proceedings of the Aristotelian Society, Supplementary Volumes*, *19*, 119–150.

Warnock, G. J. (1998). Ordinary Language Philosophy. In E. Craig (Ed.) *The Concise Routledge Encyclopedia of Philosophy*, (pp. 147–153). London: Routledge.

Wikipedia Contributors (2020). Ordinary language philosophy. In *Wikipedia, The Free Encyclopedia.*.

Wisdom, J. (1936). Philosophical Perplexity. *Proceedings of the Aristotelian Society*, *37*, 71–88.

Wittgenstein, L. (1922). *Tractatus Logico-Philosophicus*. New York: Cosimo Classics.

Wittgenstein, L. (1998). *The Collected Works of Ludwig Wittgenstein*. Charlottesville: InteLex Corp., electronic ed.

Wittgenstein, L., Ambrose, A., & Macdonald, M. (1982). *Wittgenstein's Lectures, Cambridge, 1932-1935 : from the notes of Alice Ambrose and Margaret Macdonald*. Chicago: University of Chicago Press.

Wittgenstein, L., Anscombe, G. E. M., Hacker, P. M. S., & Schulte, J. (2009). *Philosophical Investigations*. Chichester: Blackwell Publishing Ltd, revised 4t ed.

Wittgenstein, L., Anscombe, G. E. M., & von Wright, G. H. (1972). *On Certainty*. New York ; London: Harper Torchbooks.