

The Ordinary and the Experimental: Cook Wilson and Austin on Method in Philosophy

Guy Longworth¹
October 30th 2017

To what extent was ordinary language philosophy a precursor to experimental philosophy? Since the conditions on pursuit of either project are at best unclear, and at worst protean, the general question is hard to address. I focus instead on particular cases, seeking to uncover some central aspects of J. L. Austin's and John Cook Wilson's ordinary language based approach to philosophical method. I make a start at addressing three questions. First, what distinguishes their approach from other more traditional approaches? Second, is their approach a form of experimental philosophy? Third, given their aims, should it have been? I offer the following preliminary answers. First, their approach distinctively emphasizes attention to what we should say when. Second, their approach is closer to contemporary experimental mathematics than it is to some prominent forms of contemporary experimental philosophy. Third, some purported grounds for pursuing their aims by way of surveying what individual speakers would say when are not compelling.

...one who wants to achieve knowledge above the ordinary level should feel ashamed at having taken ordinary ways of talking as a basis for doubt.

—Descartes, *Meditations on First Philosophy*

I. Introduction.

J. L. Austin is a paradigm case of an ordinary language philosopher. That isn't yet to say very much, either about Austin or about ordinary language philosophy. As Austin stressed, co-classification often masks significant difference:

¹ I'm grateful to Michael Beaney, Stephen Butterfill, John Collins, Thomas Crowther, Naomi Eilan, Eugen Fischer, Nat Hansen, Mark Kalderon, Joshua Knobe, Hemdat Lerman, Rory Madden, Mike Martin, Josephine Salverda, Matthew Soteriou, Charles Travis, David Wiggins, and three anonymous reviewers.

I am not denying that cases...*could* be lumped together under some single name. A single name might in itself be innocent enough, provided its use was not taken to imply either (*a*) that the cases were all alike, or (*b*) that they were all in certain ways alike. What matters is that the facts should not be pre-judged and (therefore) neglected. (1962a: 14, fn.1)

Austin's emphasis on the importance of attending to distinctions is pre-eminent amongst the ideas that he inherited from John Cook Wilson. Other threads in their shared conceptions of method will be unraveled below. A major task will be to say what, if anything, marks their methods off from others.

My aim in discussing the two thinkers' conceptions of method is partly historical. A further hope is that such a discussion might enhance our understanding of the extent to which their methods were, or should have been, experimental. To foreshadow, experimental methods are sufficiently motley to allow for their intersection with the methods of Austin and Cook Wilson, but the members of that intersection are more akin to cases of experimental mathematics than to cases of contemporary experimental philosophy. Cook Wilson and Austin aimed to exploit attention to ordinary language with the end of uncovering optimal ways of characterising the facts: tracts of (typically, and approximately) extra-linguistic and extra-conceptual reality. By contrast, the end of much recent work in experimental philosophy has been limited to exploring extant conceptions of that reality.

§2 considers some of what Cook Wilson and Austin had to say about their methods and seeks to discern ways in which their methods were novel. One seemingly novel idea is that we might seek to address philosophical questions via attention to "what we should say when" (Austin 1956–7: 182). §3 begins to unpack that idea and to consider some skeptical reactions. §4 pursues such reactions further by discussing the extent to which Austin's methods involved, or should have involved, experiment—that is, "twisting the lion's tail," in Thomas Kuhn's memorable misattribution to Bacon. (Kuhn 1976: 12) (My discussion will be highly selective. For further discussion of Cook Wilson's and Austin's methods, see Birken-Bertsch 2014; Hacker 2013; Hanfling 2000: 26–37; Longworth 2012; Lyas ed. 1971; Marion 2000a, 2000b; Travis and Kalderon 2013; Urmson 1967, 1969; Warnock 1989: 1–10.)

2. What the Philosophers Say.

Let's begin by considering some of Austin's 'cackle' (1956–7: 189)—that is, some of what he has to say about his methods. We'll join Austin in focusing on the idea of proceeding from an examination of "*what we should say when*, and so why and what we should mean by it." (1956–7: 181) Austin offers three main lines of justification for this way of proceeding:

First, words 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. Secondly, words are not (except in their own little corner) facts or things: we need therefore to prise them off the world, to hold them apart from and against it, so that we can realize their inadequacies and arbitrariness, and can re-look at the world without blinkers. Thirdly, and more hopefully, our common stock of words embodies all the distinctions men have found worth drawing, and the connexions they have found worth marking, in the lifetime of many generations: these surely are likely to be more numerous, more sound, since they have stood up to the long test of the survival of the fittest, and more subtle, at least in all ordinary and reasonably practical matters, than any that you or I are likely to think up in our arm-chairs of an afternoon—the most favoured alternative method. (1956–7: 181–182)

Each of Austin's lines of justification presupposes a distinction between subject-matters: on one hand, words; and, on the other, (other) facts or things. Words are tools with which we attempt to, amongst other things, characterize facts or things. Those attempts can misfire, either (first) because we lack a clear understanding of how specific tools should be used, or (secondly) because selected tools aren't fit for the specific task at hand. However, (thirdly) there are grounds for optimism that a natural language will contain a sufficient array of tools to deal adequately with at least those distinctions and connexions amongst facts and things that it has been important, for ordinary or practical purposes, to mark. (Austin sometimes preferred to characterize words as instruments, noting our natural reaction to a surgeon's saying, "Right, I'll just go and get my tools." (Warnock 1973: 38–39; Pitcher 1973: 24))

Austin's third line of justification tracks very closely an earlier comment of Cook Wilson's:

The authority of language is too often forgotten in philosophy, with serious results. Distinctions made or applied in ordinary language are more likely to be right than wrong. Developed, as they have been, in what may be called the natural course of thinking, under the influence of experience and in the apprehension of particular truths, whether of everyday life or of science, they are not due to any preconceived theory. On the other hand, the actual fact is that a philosophical distinction is *prima facie* more likely to be wrong than what is called a popular distinction, because it is based on a philosophic theory which may be wrong in its ultimate principles. This is so far from being appreciated that the reverse opinion is held and there is a tendency to regard the linguistic distinction as the less trustworthy because it is popular and not due to reflective thought. The truth is the

other way. Reflective thought tends to be too abstract, while the experience which has developed the popular distinctions recorded in language is always in contact with the particular facts. (Cook Wilson 1926: 874–875)

Cook Wilson and Austin both claim that distinctions made in ordinary language are more likely to be right (worthwhile) than wrong (pointless), and more likely to be right than are distinctions drawn by philosophers without due attention to the functioning of ordinary language. However, it's not immediately clear at which alternative practices they are tilting.

Suppose, for example, that one hoped to attain philosophical understanding of what it is for someone to know something. One very natural way of proceeding would go via attention to the distinction, or distinctions, between cases in which instances of “*x* knows *y*” are correctly applicable and cases in which they aren't. Furthermore, given that one hopes to attain understanding of what it is for someone to know something, it's not clear how attention to any other distinction, or distinctions, could be more illuminating. In that case, it would seem that opposition to Cook Wilson's and Austin's claim would have to amount to objection to the hope of attaining understanding of what it is for someone to know something, perhaps due to skepticism about the rightness or worth of the purported distinction(s) between cases in which someone knows something and others.

More delicately, someone might preserve the hope of coming to understand knowing, and proceed with that end, but fail adequately to track the distinction between cases in which instances of “*x* knows *y*” are applicable and others. That might happen, for example, because, although their theorizing took off from attention to cases on either side of the target distinction, it was then driven by attention to salient, or especially simple, projections from those initial cases, rather than by attention to way their competence with the expression “*x* knows *y*” in fact guides its applications to particular cases. In that way, they might in effect develop an armchair-constructed distinction, rather than tracing a distinction developed in the natural course of experience.

There is room for dispute over the comparative standing of ordinary and artificial distinctions. However, I expect that many contemporary theorists will find reasonable the aim of seeking understanding of, e.g., knowing. And they are likely also to find reasonable the idea that tracking the distinction between cases in which “*x* knows *y*” is applicable and other cases is one way of tracking the distinction between cases in which *x* knows *y* and other cases. After all, that has been a central approach since Plato. (See e.g. Hanfling 2000: 15–25.) At this point, then, we face two connected questions. First, given that the core of Austin's and Cook Wilson's proposal seems broadly to fit the canon, is there anything distinctive about their methods? Secondly, given that their core proposal specifies ends, rather than means, how should we go about pursuing those ends?

Austin makes a start at addressing the two questions in the following passage:

In view of the prevalence of the slogan 'ordinary language', and of such names as 'linguistic' or 'analytic' philosophy or 'the analysis of language', one thing needs specially emphasizing to counter misunderstandings. When we examine what we should say when, what words we should use in what situations, 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. For this reason I think it might be better to use, for this way of doing philosophy, some less misleading name than those given above—for instance, 'linguistic phenomenology', only that is rather a mouthful. (Austin 1956–7: 182)

The passage contains two central ideas. The first is that an examination of "what we should say when, what words we should use in what situations" might be a means to the end of uncovering the ranges of conditions in which the target expressions are, or aren't, applicable. Although that doesn't clearly mark out a distinctive approach, it hints at one. The second idea is that attention to what we should say when isn't merely, or simply, attention to words; it is also, and at the same time, attention to the phenomena which words are used to characterize.

Again, Cook Wilson got there first, and with greater articulation. The base proposal is that we should pursue questions about, say, knowing, via reflection on ordinary talk about knowing. Two traditional concerns about this approach are, first, that, since it relies on our standing knowledge of how words should be used, it is mysterious how it could facilitate discovery and, secondly, that in pursuing this approach, we should be forced to accept as authoritative every naïve application of a word:

It seems...that we can only take as data the actual application of the names: if so, we are at the mercy of usage. We cannot criticize it: even if we found in it anything apparently contradictory we should be helpless to decide. Again, what security have we that we make the right generalisation and find out what people really mean? Would it not be the safest way, as it is also the easiest, to ask people who use the word *what they mean?*" (Cook Wilson 1926: 41)

Cook Wilson rejects the idea that ordinary speakers are straightforwardly authoritative about the principle governing the application of words with whose use they are competent:

Not uncommonly those who apply a principle rightly in particular cases are unable to give a correct account of its general character. (Cook Wilson 1926: 82–83)

His explanation of how we can nonetheless exploit our competence in order to uncover our principle is worth quoting at length:

Obviously we must start from the facts of the use of a name, and shall be guided at first certainly by the name: and so far we may appear to be examining the meaning of a name. Next we have to think about the individual instances, to see what they have in common, what it is in fact that actuated us. This seems by contrast to be the examination of a thing or reality as opposed to a name. At this stage we must take first what seems to us common in certain definite cases before us: next test what we have got by considering other instances of *our own* application of the name, other instances (more accurately) in which the principle has been working in us. Observe that in every such step we rely upon the rightness of our use of the principle in particular cases; this does not mean that we are sure of ourselves in every case, but that there *are* cases at all events about which we are sure. This explains what in the Socratic attempt to find definitions would otherwise be paradoxical and inexplicable. There is a further stage when we have, or think we have, discovered the nature of the principle which has really actuated us. We may now correct some of our applications of the name because we see that some instances do not really possess the quality which corresponds to what we now understand the principle to be. This explains how it should be possible to criticize the facts out of which we have been drawing our data.” (Cook Wilson 1926: 44–45)

Cook Wilson’s proposal is that we should begin with secure individual applications to cases of a target expression. We should then try to discern the features of those cases which actuated our applications. We can then refine, and develop, our view about those actuating features by testing them against our secure willingness (/unwillingness) to apply the target expression in cases exhibiting (/failing to exhibit) those features. In that way, our security with respect to a core range of applications can transmit to security with respect to our view of actuating features in order to support a view about the core principle, or principles, guiding our naïve applications. And our thus supported view about the core principles can then be used to support, or correct, applications about which we were initially less sure. In effect, our secure applications of an expression help to make available data for the construction of a theory about the principle governing our overall competence with the use of the expression. (We shouldn’t seek too sharp a distinction between that model and one on which exposure to, and reflection on, secure

applications is at the service of something closer to training than to theory construction.) In that way, we can exploit our competence with the use of the target expression in order to make discoveries about the principle governing that use, and so to sharpen both our competence and our conception of the subject-matter marked out by that use. (It is worth comparing the discussion of Plato's engagement with this issue in Nehamas 1986. At this level of generality, we can leave open how precisely applications are to be understood—whether, for example, they must serve descriptive functions. We can also leave open the precise delimitation of cases—whether they should be individuated only by appeal to the focal subject matters of invited judgments or also, for example, by appeal to the intents and purposes of those being invited to judge.)

We are now in a better position to understand Austin's promotion of adherence to the distinctions marked out by our common stock of words, and its connection with his recognition of the sometime arbitrariness and inadequacy of those distinctions. Consider the second stage of Cook Wilson's procedure. At that stage, having selected a range of cases with respect to which our applications of a term are especially secure, we seek to discern the features of those cases which actuated those secure applications. At that stage, we are liable to attend to the most salient features of the cases, and to project in what we take to be the most natural or straightforward ways from those features. In that way, we are liable to attend to the phenomena rather than to the operation of our competence with the target expression. If the operation of our competence tracked the most salient features of cases, in the most straightforward ways, then that approach might safely converge on the principle governing our use. However, we can only acquire evidence of convergence by checking our initial view of the principle against novel applications of the target expression. And when we do that, we find that the pattern of those novel applications sometimes appears arbitrary relative to the most natural seeming projection from salient features of our initial cases. Thus, Austin and Cook Wilson can be seen as distinctively promoting single-minded attention to the actual pattern in our applications of an expression over what might seem to be more natural projections from the subject matters of a proper subset of those applications. Their prioritizing of the actual pattern of use over the seemingly natural pattern is in turn justified by the fact that it is the actual pattern, as opposed to the seemingly natural, which has withstood the test of time. (See also Prichard 1912 and Williams 1988.)

It is important not to overlay the distinction between application and projection. For even naïve applications of an expression can involve judgement or discernment, with respect both to the phenomena and to their fit with our principle. And those exercises of judgement or discernment may be conditioned by natural projections from features of other cases in which the expression has been applied. Austin sketched

the suggested interplay between application and projection in the following way:

...[E]ssential though it is as a preliminary to track down the detail of our ordinary use of words, it seems that we shall in the end always be compelled to straighten them out to some extent.... [I]t seems hardly deniable that in...thought and language we do, for better or worse and whether consciously or unconsciously, make use of...models (not, of course, necessarily only one such). (Austin 1952-3: 134)

Though cases in which we shall have to call, describe, &c., instead of, in black-and-white terms, identifying, stating, &c.—that is, cases in which “there occur in the world to be talked about items of types which do not exactly match any of the patterns in our stock (the sense of any of our names)”—], are sure to arise continually, we feel ourselves sometimes bound to cope with them as they arise by means of fresh linguistic legislation. In calling there is indeed already implicit an element of legislation by precedent. (Austin 1952-3: 149; interposition: 146)

Crudely, the correspondence between applicability conditions and environmental distinctions—for example, between the question whether “*x* knows *y*” is correctly applicable in a particular case and the question whether, in that case, *x* knows *y*—means that there is no principled distinction between finding out about applicability conditions and finding out about the world. David Wiggins presents a closely related thought in the following way:

Not only does the real definition depend on real specimens. The question of its application and correctness or incorrectness also depends crucially on the facts about these specimens. All the same, real definitions lie within the province of semantics, as well as of empirical fact. No lifelike elucidation of ordinary discourse could be contrived without them. [Footnote suppressed.] Let us forget once and for all the very idea of some knowledge of language or meaning that is not knowledge of the world itself. (Wiggins 2001: 12. The suppressed footnote cites Austin 1952-3 as a precursor. See also Wiggins 2001: 77-106. Austin’s work on Leibniz may have figured in opening his eyes to the role of specimens in the determination of applicability conditions.)

So much, then, for the interplay between attention to applications and attention to their phenomenal targets, and for Austin’s and Cook Wilson’s distinctive emphasis on the former. Still, one might wonder about the extent to which their methods are novel. If there is novelty here, it is plausibly to be found in Austin’s proposal about the initial

source of data for theorizing about applicability conditions—that is, his idea that we might gain purchase on the applications of an expression through attention to “what we should say when, what words we should use in what situations.” (Austin 1956–7: 182) What, precisely, does Austin’s proposal amount to?

3. What We Should Say When.

Our question about Austin’s proposal can be resolved into three sub-questions. First, how should we understand his appeal to what we should say when? Is this a matter merely of our uses of words in circumstances; our presentation of words as correctly applicable or inapplicable in those circumstances; what should (or shouldn’t) be said in those circumstances; or something else? Second, how can we make the transition from claims about what we should say when, when understood in Austin’s way, to claims about the correct applicability of target expressions? Is a claim about what should be said when simply a claim about correct applicability; or is it rather a claim about data for what are then further theoretical claims about correct applicability? (We’ve already seen reasons to think that secure claims about what we should say in some particular cases might serve as the basis for more tentative, theoretical claims about what we should say in cases with respect to which our views are less secure.) Third, what, if anything, grounds our knowledge, or entitles us to our opinion, about what we should say when? Is this empirical knowledge, say, about what each of us, or others, would say when, or is it something else?

Charles Travis and Mark Kalderon take a firm stand on our first and second sub-questions:

Austin’s vocabulary here, specifically, ‘what we should say when’, can be misread. Supposing that there are things words are *for* saying, it would be natural to read this as: ‘If you (one) were to use *these* words of *this*, or in *these* circumstances, what would you say?’ What one thus asks after is how words in fact work. Austin clearly hears things this way. Grice insists on a different reading. On it, ‘what we would say’ merely reports our customs, mores, manners: ‘One wouldn’t say, “What’s the vigorish?” when the neighbor asks to borrow a cup of milk’, ‘One shouldn’t say, “That’s just autobiography” to your small niece when she says she wants another biscuit’. But asking what one would say when *can* be a way of asking what the words one uses in fact apply to, or describe *truly*—what they *are* for in their language. (Travis and Kalderon 2013: 498)

Their view is that Austin’s talk about what we should say when is, in effect, already talk about the correct applicability of expressions. That straightforward answer to the first and second sub-questions might

naturally seem to provoke a pointed version of the third: how do you (/does one) know what you (/one) should say when—that is, whether a given expression is correctly applicable here? Alternatively, the answer to the first sub-question which Travis and Kalderon attribute to Grice—on which what we should say when is a matter of our willingness, or unwillingness, to utter certain words in specific circumstances, or whether those words *seem* to us to be applicable here—might seem to make it comparatively easy to address the third sub-question. But we would then face difficulties in answering the second question. For, as Grice and others have stressed, the gap between, on one side, claims about our willingness or unwillingness to utter certain words in specific circumstances, or our opinions about whether those words seem to apply or not, and, on the other side, claims about correct applicability, can be traversed only on the basis of substantive theory, and then only tentatively. (See also White 1967; Travis 1985.)

Some of what Austin wrote might seem to suggest that he would have been willing to sacrifice reach for security. Thus, for example, he considers “a good site for *field work* in philosophy” an area “where ordinary language is rich and subtle”, and “which is not too much trodden into bogs or tracks by traditional philosophy”:

Here at last we should be able to unfreeze, to loosen up and get going on agreeing about discoveries, however small, and on agreeing about how to reach agreement. (Austin 1956–7: 182–3)

Austin’s emphasis on the possibility of reaching agreement might seem to suggest the idea of an especially secure basis for further theorizing, a sort of linguistic given. The emphasis recurs in subtly different form here:

We must...imagine some cases (imagine them carefully and in detail and comprehensively) and try to reach agreement upon what we should in fact say concerning them. If we can reach this agreement, we shall have some *data* (‘experimental’ data, in fact) which we can then go on to *explain*. Here, the explanation will be an account of the meanings of these expressions... (Austin 1966: 274)

Here, Austin suggests that we have data only where we have reached agreement. And that might be taken, in turn, to suggest that our data must be especially secure. We might, for example, take disagreement to entail error on behalf of at least one party. (At a minimum, this would require the obtaining of genuine disagreement, as opposed to mere difference. See Austin 1956–7: 183–185.) Thus, the possibility of disagreement would entail the possibility of error. And we might, in turn, hold that we can obtain knowledge only where there is no possibility of error. Thus, if non-theoretical claims about correct

applicability are possible sites of disagreement, in a way that claims about willingness or unwillingness to utter are not, then that might seem to make pressing the pointed question, how can one know whether a given expression is correctly applicable here?

Pulling in the opposing direction is that Austin was well aware that it doesn't follow from the fact that an ability is fallible that it cannot be exercised successfully. Suppose, then, that competence with a linguistic expression comprises an ability to know whether it is correctly applicable with respect to presented cases. On that supposition, successful exercises of the competence results in knowledge of whether the expression is correctly applicable. Such an ability is bound to be fallible: there will be cases in which its exercise is unsuccessful and, so, in which that exercise does not result in knowledge, and may even result in false opinion. But the fact that the ability to knowlegably apply an expression is fallible doesn't entail that it cannot be exercised successfully and, in that case, result in knowledge of correct application. Indeed, that an ability to know can be fallible is one of the central take-home messages of Austin's 1946 and 1962a. Thus, it needn't follow from the fact that there can be disagreement—or, more generally, error—about whether an expression is correctly applicable that competence with the expression isn't a capacity whose successful exercise results in knowledge of whether it is correctly applicable.

Austin's emphasis on seeking out good sites for field work, and Cook Wilson's emphasis on starting with especially secure cases, both manifest recognition that knowledge delivering capacities can be fallible. Two special sources of fallibility stressed by Austin are prejudice and what he called "looseness", where the latter seems to cover not only inattention to available distinctions, but also inadequate training:

It is likely enough that our ordinary use of the terms is fairly loose, that we do not always distinguish carefully between them, although there is a distinction which can be marked by their means. Contrast, for example, the following:

- (i) You call that crimson? But surely no crimson can have so much blue in it? That is not what crimson is at all.
You describe it as crimson? But look, it has lots of blue in it. It is not really like crimson at all....

If many such examples are studied, the watershed between calling and describing appears to take shape. (Austin 1952-3: 148-149; 1956-7: 182-185)

Furthermore, there are other natural explanations for Austin's emphasis on agreement. For one thing, although disagreement needn't support the view that neither party knows, agreement can support the view that both parties do. For another, both Austin and Cook Wilson hope to

begin with an especially secure range of applications, and it would be reasonable to think that agreement is a mark of security. For a third, the nature of the target facts might make agreement especially important. Facts about the correct application of an expression plausibly depend, in part, on facts about the plurality of competent users of that expression. In that case, it would seem reasonable to prioritize those views about correct applicability which we can bring others to share. (For an excellent discussion of related issues about reach and security, see Williamson 2004, especially 109–131.)

4. Twisting the Lion's Tail.

We saw in the previous section that Austin was prepared to characterize some of competent speakers' shared views about whether an application is correct as 'experimental' data (his quotation marks). What did he mean? And was he right?

The expression "experiment" began to diverge in meaning from "experience" only in the 17th Century. (Dear 2006: 106; Hinton 1973a: 5–8; Wootton 2015: 310–360) However, although typical experiments involve interventions with observable results, observation seems to be inessential to contemporary usage, as recorded in the *OED*. Instead, we find an emphasis on active intervention: "The action of trying anything, or putting it to proof"; "A tentative procedure"; "An action or operation undertaken in order to discover something unknown, to test a hypothesis, or establish or illustrate some known truth." (See also Field and Hole 2003; Sorensen 1992: 186–251; Tiles 1993.) So, despite the drag induced by etymology and stereotypes, there is no immediate reason to think that Austin's 'experimental' data must be got via observation. Alternatively, we've already seen a way in which Austin's method involves active intervention: we are to intervene on subjects (often, ourselves) by presenting them with various cases, together with questions about what they should say about those cases. Furthermore, we've seen one way in which the procedure is tentative: secure judgement, or knowledge, about the specific characterisation of particular cases, is used in order to develop and test less secure general hypotheses concerning the range of cases with respect to which target expressions are applicable.

As we saw in the previous section, an alternative view would locate tentativeness—the space between experimental results and hypotheses—at a different point. According to the alternative, experimental interventions on competent speakers deliver only data about competent speakers' opinions about the correctness of particular applications. That information may then be used to develop and test distinct hypotheses even about the correctness of particular applications. As we've seen, an attempt to motivate such a view might be made on the basis of reflection on individual fallibility. Connectedly, it might be motivated by the idea that facts about correct application

are facts about what *we* should say, together with skeptical concerns about individual speakers' authority to judge about what other competent speakers would say.

We can gain some purchase on Austin's attitude towards the alternative view by considering his reaction to its implementation in the work of Arne Naess. Naess pursued what he called 'experimental semantics' by surveying speakers' judgements about cases, as well as their more general views about target expressions. (For more detailed discussion of, and references to, Naess's work, see e.g. Chapman 2011 and ms; Murphy 2014.) His work was therefore similar in some respects to some prominent strains of contemporary experimental philosophy. (See e.g. Weinberg, Nichols, and Stich 2001; Hansen 2014; Hansen and Chemla 2015.) Although the evidence we have about Austin's attitude to this project is fragmentary, the fragments cast a shadow. First, we have the following reports from Geoffrey Warnock:

I remember that [Austin] once came back from America—I think in 1956—a good deal perturbed by what he thought to be the increasing prestige there of Arne Naess. This must have been because he thought he saw the right *purpose*—a more empirical, 'objective' way of doing philosophy, offering the hope of getting things actually settled by patient industry—in danger of being compromised by what he took to be radically wrong *methods*. 'It's infiltrating from the West', he said, shaking his head. (Warnock 1973: 43)

In discussion with Arne Naess at Berkeley in 1958, Austin appears to have spoken as if he still regarded some kind of systematic co-operation in philosophy as not only desirable, but also practicable. However, the record of this discussion is neither perfectly clear nor certainly reliable. (It goes without saying that Austin was careful to distinguish the programme he had in mind from the kind of Gallup-poll, empirical team-work which Naess believed in, and which Austin regarded as, in principle, misguided. (Warnock 1963: 14, fn.2)

As Warnock mentions, some records remain of Austin's discussions with Naess. They seem to indicate that Austin was not open to having his views about what we should say when overruled by other speakers' unreflective opinions:

AUSTIN: The subjects gave wrong answers concerning their own use of expressions, e.g.: when saying they would never use "he yawned voluntarily" as a description of a perfectly ordinary event of yawning *because* it is perfectly obvious that such yawnings *are* voluntary. *Actually the subjects would not say it because it cannot be said.*

NAESS: The subjects interpreted “he yawned voluntarily” as synonymous with “he was not forced to yawn” and thus conceived it as obvious that he was not forced to (and therefore *not worth while saying*).

AUSTIN: But then they do not know well enough the expression “voluntary.” It is too difficult a word, maybe. Better use “clumsy.” (Austin and Naess 1964, as reported in Murphy 2014: 354.)

Austin clearly felt that his competence with “voluntary”, as honed by reflective training, put him in a position to correct naïve opinion. Furthermore, although he appears willing to make predictions about other people’s usage, they are based on what he takes to be facts about the target expression (“he yawned voluntarily” cannot be said’), and are conditional on the predicted usage’s issuing from competence. Disagreement with the secure judgment of the competent is evidence of incompetence.

Austin’s model of experimental philosophy involves the assembly of particular pieces of knowledge to the effect that a target expression is, or isn’t applicable, in particular cases. Such knowledge is knowledge about an expression—say, about the expression “voluntary”—and, at the same time, knowledge about the world—say, about whether some particular activity is, or could have been, voluntary. Those pieces of knowledge are at the service of developing and testing hypotheses about the principle, or principles, governing application of the target expression. (Alternatively, or simultaneously, they are at the service of honing one’s competence.) Those principles, in turn, concern both the target expression and the environmental distinctions and connections which fully competent use of the expression would track.

We might usefully compare our two models of experimental ordinary language philosophy with two models of how experimental methods might be applied to mathematical practice. The first model is embodied in an extant mathematical practice, experimental mathematics. To a good first approximation, the practice involves the assembly of a large number of particular instances of a conjectured generalization—say, an equation—as a guide to the development of methods of proof or disproof of the target conjecture. Alan Baker offers the following nice example:

A real number is said to be *normal in base n* if every sequence of digits (of any given length) occurs equally often in its base- n expansion. A number is *absolutely normal* if it is normal in every base. Consider the following hypothesis:

Conjecture Every non-rational algebraic number is absolutely normal.

Borwein and Bailey used a computer to compute to 10,000 decimal digits the square roots and cube roots of the positive integers smaller than 1,000, and then they subjected these data to certain statistical tests. (Baker 2008: 333, reporting work in Borwein and Bailey 2004.)

As with Austin's model of experimental philosophy, the practice depends on the accuracy of the assembled instances. And like Austin's model, those instances are used to ground conjectures about the overall pattern of instances which may then be put to further test. (Further test may involve checking the conjecture against further instances, or—more commonly in mathematics than in philosophy—the construction of proof.) (See e.g. Baker 2008.)

The second model is more likely to be found in psychology than in mathematics. Operating in accord with this model involves the assembly of a large number of people's opinions about the outcome of particular calculations of values of a mathematical generalization. In general, there is unlikely to be a perfect correlation between people's opinions about the outcome of a calculation and the actual outcome, so it would be unreasonable to expect the data delivered by this procedure to figure centrally in the mathematical exploration of the target generalization. Rather, it would be unreasonable to expect such data to be revealing about anything other than the patterns of individual or group competence and incompetence, and perhaps about the development and structure of individual mathematical abilities. (Such results might figure peripherally in the development of mathematical practice—for example, by helping to shape optimal presentation of proofs, or pedagogical practice.) In some cases, for example in work on the Wason Selection Task, attention is also paid to subjects' reflective corrections of their initial responses, and those might sometimes be a better guide not only to the outputs of subjects' competence, but also the mathematical facts that they aim to track. (See e.g. Wason 1968.) However, the core aim of such work is to delimit subjects' abilities (the mathematical facts in question are often trivial), and there is no guarantee that subjects' most reflective judgments will reflect their core competences. (A relevant case study is provided in linguistics by the way that so-called prestige accretions can distort judgment, leading, for example, to the over-application of case marking, as when "who" becomes "whom". See e.g. Lasnik and Sobin 2000.)

It seems clear that Austin's conception of experimental ordinary language philosophy is closer to the first model—experimental mathematics proper—than it is to the second. The aim, in both ordinary language philosophy and experimental mathematics, is to exploit competence in order to develop understanding of the domain of facts that competence works. In both cases, exploitation goes via secure individual deliverances of competence—pieces of knowledge about particular applications of expressions in particular cases and individual

calculative results, respectively. (In the case of experimental mathematics, the competence in question is typically embodied, vicariously, in computers.) In both cases, those individual deliverances are used in order to gain purchase on cases with respect to which individual competence is less forthright, and perhaps silent. What makes it natural to think of both cases on the model of experiment is, first, the role of active intervention on variables (e.g. specific aspects of the case being considered) in order to assemble specific results and, second, the role played by those specific results in the development of more general conceptions of the subject area. By contrast, on the assumption that speakers and mathematicians can be more or less competent—that is, more or less able to track linguistic and mathematical fact—and that competence in both domains can be improved by reflective practice, the second model would not offer a natural way of pursuing Austin's central aims. As in the mathematical case, operating in accord with the second model might provide a guide to patterns of competence and incompetence, and to speakers' current abilities and conceptions of a subject area. Information of that sort might figure, peripherally, in helping to shape our views about the standing of particular judgments, or in the development of pedagogical strategy. However, its bearing on the nature of full competence, and the domain of facts over which that competence operates, is likely to be, at best, indirect. (On the potential bearing of psychology in the assessment of particular judgments or inferences, see e.g. Fischer 2014; Weinberg 2015.)

So much for Austin's opinion. Was he entitled to it? That is a large, open question and I can't hope to address it adequately here. (There is a large literature on this topic, the main early contributions to which are usefully collected in Lyas ed. 1971.) However, I shall comment briefly on three relevant issues.

The first issue concerns the suggestion that Austin would have viewed the role of his experiments as akin to that of those undertaken in experimental mathematics: the assembly of pieces of knowledge, and thence facts, about particular cases. Minimally, the suggestion presents the proponent of the alternative model—which seeks instead to assemble information about subjects' opinions about cases—with a challenge: either to defend the view that such information might play a fundamental role in mathematics, or to defend the differential treatment of mathematics and philosophy in this respect. Naturally, proponents of the alternative model are liable to favour the second disjunct. However, typical attempts to defend differential treatment of mathematics and philosophy point to what are, at most, differences of degree, rather than of kind, and, furthermore, differences the relevance of which is open to philosophical dispute. For example, a number of such attempts appeal to the comparative lack of consensus in philosophy. (See e.g. Chalmers 2014; Weinberg 2015.) Such attempts arguably both overstate the extent of consensus in mathematics (at least

with respect to the mathematically non-trivial) and understate the extent of consensus in philosophy (at least with respect to the philosophically trivial). (For pertinent discussion, see Clarke-Doane 2014.) Furthermore, it would need to be shown that absence of consensus is relevant to deciding whether it is reasonable for philosophers to take themselves to be in a position to acquire knowledge about particular cases. Some aspects of that issue were discussed in §3; others lead into the second issue that I'll consider.

The second issue, then, concerns Austin's plural formulation of his task: accounting for what *we* should say when. Even if we were willing to allow that an individual speaker can be authoritative with respect to the deliverances of their own competence, we might object to the idea that they can be authoritative with respect to the deliverances of other people's competence, and so to the idea that they can be authoritative about what *we*—as opposed to *they*—should say when. This is a place where the proponent of a model of experimental ordinary language philosophy like Naess's is liable to seize their chance. Taking this line, they would try to argue that in order to find out what others should say when, one would need to experiment on the others' competences.

One response to that line of attack points, again, to the possibility of variance, and development, in competence. As discussed in §2, Austin's claims about what we should say when are claims about what those who have reflectively sharpened their competence via attention to particular cases would say when, and not claims about what any minimally competent speaker would say when. (It may be relevant here that the initial claim concerns what we *should* say, rather than what we *would* say. Subjects' willingness reflectively to view their initial judgments as erroneous, as on the Wason Selection Task, might also be relevant here. The issue here connects with issues about the extent to which broadly competent thinkers and speakers can be more or less expert. See e.g. Gustafsson 2005; Horvath and Wiegmann 2016; Kauppinen 2007; Williamson 2011.)

But a more general response can be offered, albeit along the same broad line. Claims about what we would do when are ordinary plural claims and, so, are potentially open to both distributive and collective construal. On a distributive construal, a claim about what we would say when would immediately entail a claim of the same type about what each of us would say when. (Compare: on its distributive construal, a claim to the effect that we are philosophers entails that each of us is a philosopher.) By contrast, on a collective construal, a claim about what we would say when would fail immediately to entail a claim of the same type about what each of us would say when. (Compare: a claim to the effect that we are numerous does not entail that each of us is numerous.) When so construed, such claims concern the activities of a number of people without imposing specific demands on the activities of any individual amongst that number. So, if we construe claims about what we would say when collectively, they would carry no specific

commitments with respect to what each of us, so others, should say when. (See e.g. Oliver and Smiley 2013. For sympathetic models of the ways in which facts about correct applicability can depend on collective activities, see e.g. Burge 1979; Putnam 1975; Wiggins 1997; Williamson 2007: 122–127.)

The third issue is the one raised by Descartes in this paper's epigraph, what Austin calls "the Last Word":

Certainly ordinary language has no claim to be the last word, if there is such a thing. It embodies, indeed, something better than the metaphysics of the Stone Age, namely, as was said, the inherited experience and acumen of many generations of men. But then, that acumen has been concentrated primarily upon the practical business of life. If a distinction works well for practical purposes in ordinary life (no mean feat, for even ordinary life is full of hard cases), then there is sure to be something in it, it will not mark nothing; yet this is likely enough to be not the best way of arranging things if our interests are more extensive or intellectual than the ordinary. And again, that experience has been derived only from the sources available to ordinary men throughout most of civilized history: it has not been fed from the resources of the microscope and its successors. ...Certainly, then, ordinary language is *not* the last word: in principle it can everywhere be supplemented and improved upon and superseded. (Austin 1956–7: 185. See also Austin 1940: 67–69; 1956: 231–232)

We've seen that Cook Wilson and Austin both allow that competence with an ordinary expression might be developed over time, and especially that it might be developed in light of reflection on the details of cases in which the expression is, or isn't, applicable. However, they both restrict their attention to those details that could have figured in the actuation of exercises of competence, as fixed by an extant conception of the principle which governs those exercises. There is therefore a tendency in their thinking towards conservatism about the extent to which the development of one's conceptions of some subject area can amount to the improvement of one's competence with an old expression as opposed to the acquisition of a competence with a new expression. This tendency manifests itself in Austin's discussion of the Last Word in the way in which he expects the findings of experimental science ("the resources of the microscope and its successors") to figure in the development of philosophy. (See also Hinton 1973b.)

Austin's expectation is that the capture of such findings will require linguistic innovation, rather than being absorbed into the development of competence with our extant linguistic resources. (Austin 1962b contains an array of examples of Austin's own terminological innovations, including his distinction between 'locutionary', 'illocutionary', and 'perlocutionary' acts.) Now Austin is

surely right that that is a possible—indeed, actual—outcome of experimental science. However, it has been a central theme of more recent work that experimental science can figure more centrally in the development of an extant competence. That work has indicated ways in which the sorts of interplay between one’s initial conception of a subject matter and reflection on the details of that subject matter that figures in the development of an extant competence can be porous to information that didn’t figure in the initial development of the competence, or in its actuation. For example, it indicates that our competence with the expression “water” can be developed in light of experimental findings about the chemical structure of the stuff to which the expression has been applied. (See especially Putnam 1975 and, again, Wiggins 2001.) That work imposes a sharp limit on the extent to which naïve reflection on cases can fully reveal the geography of distinctions and connexions ensconced in our competence with ordinary language.

5. Conclusion.

Cook Wilson’s and Austin’s approach to philosophy aims for the reflective development of competence through the interplay between attention to language and attention to the bits of the world to which language is applied. Insofar as their approach is experimental, their views about the aims of experiment are distinctive. Most importantly, they view the aims of experiment as being the assembly of pieces of knowledge about particular cases, and so facts about those cases, rather than mere opinions. Although our ways of coming to have such knowledge are bound to be fallible, that raises no special problem. In particular, our fallibility doesn’t immediately force us to seek a more secure basis for our claims about what we should say when. Furthermore, I’ve suggested that the role played in their approach by the reflective development of competence serves to distinguish those approaches from others which aim only to discern the shapes of subjects’ extant competences. In both those ways, then, there is a principled distinction between their approach and common contemporary forms of experimental philosophy. However, I’ve also suggested that the same feature of their approach leads to a way in which our conception of the reach of ordinary language has developed beyond theirs, in order to take in the findings of experimental science. Insofar as we continue to pursue their approach, we do so with an acute sense of its limitations.

Austin summarized his thinking about method in a set of notes entitled “Something about one way of possibly doing one part of philosophy.” In those notes, Austin expressed what I’ve argued to be an appropriate modesty, both about the reach of the methods that he discusses, and also about their novelty. His methods represent only a development of *one* traditional way of doing *some* of philosophy’s work.

It seems appropriate, then, to allow Austin the last word in summarizing our main points of agreement:

Shan't learn everything, so why not do something else? Well; not even whole of philosophy but firstly always has *been* philosophy, since Socrates. And some slow successes. Advantages of slowness and cooperation. Be your size. Small men. Foolproof × genius-proof. Anyone with patience can do something. Leads to discoveries and agreement. Is amusing. Part of *personal* motive of my colleagues to avoid interminable bickering or boring points of our predecessors: also remember all brought up on classics: no quarrel with maths etc., just ignorant. (Austin "Something about one way of possibly doing one part of philosophy," as quoted in Urmson 1969: 83.)

References.

- Austin, J. L. 1940. "The Meaning of a Word." In Austin 1979: 55–75.
- Austin, J. L. 1946. "Other Minds." *Proceedings of the Aristotelian Society, Suppl. Vol. 20*: 148–187. References to the reprint in his 1979: 76–116.
- Austin, J. L. 1952–3. "How to Talk: Some Simple Ways." *Proceedings of the Aristotelian Society*. References to the reprint in his 1979: 134–153.
- Austin, J. L. 1956. "Ifs and Cans." *Proceedings of the British Academy*. References to the reprint in his 1979: 205–232.
- Austin, J. L. 1956–7. "A Plea for Excuses." *Proceedings of the Aristotelian Society*. References to the reprint in his 1979: 175–204.
- Austin, J. L. 1962a. *Sense and Sensibilia*. Reconstructed from the manuscript notes by G. J. Warnock. Oxford: Oxford University Press.
- Austin 1962b. *How to Do Things with Words*. 2nd edn. Edited by J. O. Urmson and Marina Sbisa. Oxford: Oxford University Press.
- Austin, J. L. 1979. *Philosophical Papers*. 3rd edn., J. O. Urmson and G. J. Warnock eds. Oxford: Clarendon Press.
- Austin, J. L. and Naess, A. 1964. "On Herman Tennessen's 'What Should We Say'." In H. Tennessen ed. *Eighteen Papers on Language Analysis and Empirical Semantics*. Edmonton: University of Alberta Press: 143–150.
- Baker, A. 2008. "Experimental Mathematics." *Erkenntnis* 68: 331–344.
- Beaney, M. ed. 2013. *The Oxford Handbook of the History of Analytic Philosophy*. Oxford: Oxford University Press.
- Berlin, I. ed. 1973. *Essays on J. L. Austin*. Oxford: Clarendon Press.
- Birken-Bertsch, H. 2014. "Austin's Method." In B. Garvey ed. 2014: 89–107.
- Borwein, J. and Bailey, D. 2004. *Mathematics by Experiment: Plausible Reasoning for the 21st Century*. Natick, MA: AK Peters.

- Burge, T. 1979. "Individualism and the Mental." *Midwest Studies in Philosophy* 4: 73-121.
- Chalmers, D. 2014. "Why isn't there more progress in Philosophy?" *Philosophy* 90, 1: 3-31.
- Chapman, S. 2011. "Arne Naess and Empirical Semantics." *Inquiry* 54, 1: 18-30.
- Chapman, S. ms. "The Experimental and the Empirical: Arne Naess's Statistical Approach to Philosophy."
- Clarke-Doane, J. 2014. "Moral Epistemology: The Mathematics Analogy." *Noûs* 48, 2: 238-255.
- Cook Wilson, J. 1926. *Statement and Inference*. A. S. L. Farquharson ed. Oxford: Clarendon Press.
- Dear, P. 2006. "The Meanings of Experience." In K. Park and L. Daston eds. *The Cambridge History of Science. Volume 3: Early Modern Science*. Cambridge: Cambridge University Press: 106-131.
- Fann, K. T. ed. 1969. *Symposium on J.L. Austin*, London: Routledge & Kegan Paul.
- Field, A. and Hole, G. 2003. *How to Design and Report Experiments*. London: Sage Publications, Ltd.
- Fischer, E. 2014. "Verbal Fallacies and Philosophical Intuitions: The Continuing Relevance of Ordinary Language Analysis." In B. Garvey ed. 2014: 124-140.
- Garvey, B. ed. 2014. *J. L. Austin on Language*. Basingstoke: Palgrave Macmillan.
- Gustafsson, M. 2005. "Perfect Pitch and Austinian Examples: Cavell, McDowell, Wittgenstein, and the Philosophical Significance of Ordinary Language." *Inquiry* 48, 4: 356-389.
- Hacker, P. 2013. "The Linguistic Turn in Analytic Philosophy." In Beaney ed. 2013: 926-947.
- Hanfling, O. 2000. *Philosophy and Ordinary Language: The Bent and Genius of our Tongue*. London: Routledge.
- Hansen, N. 2014. "Contemporary Ordinary Language Philosophy." *Philosophy Compass* 9/8: 556-569.
- Hansen, N. and Chemla, E. 2015. "Linguistic Experiments and Ordinary Language Philosophy." *Ratio* 28: 422-445.
- Hinton 1973a. *Experiences: An Inquiry into Some Ambiguities*. Oxford: Clarendon Press.
- Hinton 1973b. "Linguistic Philosophy, Empiricism, and the Left." *Philosophy* 48: 381-385.
- Horvath, J. and Wiegmann, A. 2016. "Intuitive expertise and intuitions about knowledge." *Philosophical Studies* 173: 2701-2726.
- Kauppinen, A. 2007. "The Rise and Fall of Experimental Philosophy." *Philosophical Explorations* 10, 2: 95-118.
- Kuhn, T. S. 1976. "Mathematical vs. Experimental Traditions in the Development of Physical Science." *Journal of Interdisciplinary History* 7: 1-31.

- Lasnik, H. and Sobin, N. 2000. "The Who/Whom Puzzle: On the Preservation of an Archaic Feature." *Natural Language & Linguistic Theory* 18, 2: 343–371.
- Longworth, G. 2012 "John Langshaw Austin." *The Stanford Encyclopedia of Philosophy* (Spring 2017 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/spr2017/entries/austin-jl/>>
- Lyas, C. ed. 1971. *Philosophy and Linguistics*. London: Macmillan.
- Marion, M. 2000a. "Oxford Realism: Knowledge and Perception I." *British Journal for the History of Philosophy* 8, 2: 299–338.
- Marion, M. 2000b. "Oxford Realism: Knowledge and Perception II." *British Journal for the History of Philosophy* 8, 3: 485–519.
- Murphy, T. S. 2014. "Experimental Philosophy 1935–1965." In T. Lombrozo, J. Knobe, and S. Nichols eds. *Oxford Studies in Experimental Philosophy. Volume 1*. Oxford: Oxford University Press: 325–368.
- Nehamas, A. 1986. "Socratic Intellectualism." In J. J. Cleary ed. *Proceedings of the Greater Boston Area Ancient Philosophy Colloquium*, vol. II. Washington, D.C.: University Press of America: 275–316. Reprinted in his 1999 *Virtues of Authenticity*, Princeton, NJ: Princeton University Press: 24–58.
- Oliver, A. and Smiley, T. 2013. *Plural Logic*. Oxford: Oxford University Press.
- Pitcher, G. 1973. "Austin: A Personal Memoir." In Berlin ed. 1973: 17–30.
- Prichard, H. A. 1912. "Does Moral Philosophy Rest on a Mistake?" *Mind* 21, 81: 21–37.
- Putnam, H. 1975. "The Meaning of 'Meaning'." In K. Gunderson ed. *Language, Mind and Knowledge. Minnesota Studies in the Philosophy of Science VII*. Minneapolis: University of Minnesota Press: 131–193.
- Sorensen, R. A. 1992. *Thought Experiments*. Oxford: Oxford University Press.
- Tiles, J. E. 1993. "Experiment as Intervention." *The British Journal for the Philosophy of Science* 44, 3: 463–475.
- Travis, C. 1985. "On What is Strictly Speaking True." *The Canadian Journal of Philosophy* 15, 2: 187–229.
- Travis, C. and Calderon, M. E. 2013. "Oxford Realism." In M. Beaney ed. 2013: 489–517.
- Urmson, J. O. 1967, "Austin, John Langshaw." In P. Edwards ed. *The Encyclopedia of Philosophy*, New York: Random House, Inc.: 211–215; reprinted as "Austin's Philosophy", in Fann 1969: 22–32.
- Urmson, J.O. 1969. "A Symposium on Austin's Methods I." In Fann ed. 1969: 76–86.
- Warnock, G. 1963. "John Langshaw Austin, A Biographical Sketch." *Proceedings of the British Academy*. References to the reprint in Fann ed. 1969: 3–21.
- Warnock, G. 1973. "Saturday Mornings." In Berlin ed. 1973: 31–45.
- Wason, P. C. 1968. "Reasoning about a rule." *Quarterly Journal of Experimental Psychology* 20, 3: 273–281.

- Weinberg, J. M., Nichols, S., and Stich, S. 2001. "Normativity and Epistemic Intuitions." *Philosophical Topics* 29, 1: 429-460.
- Weinberg, J. M. 2015. "Humans as instruments: or, the inevitability of experimental philosophy." In E. Fischer and J. Collins eds. *Experimental Philosophy, Rationalism, and Naturalism*. London: Routledge: 171-187.
- White, A. R. 1967. "Mentioning the Unmentionable." *Analysis* 27, 4: 113-118.
- Wiggins, D. 1997. "Natural Languages as Social Objects." *Philosophy* 72: 499-524.
- Wiggins, D. 2001. *Sameness and Substance Renewed*. Cambridge: Cambridge University Press.
- Williams, B. 1988. "What Does Intuitionism Imply?" In J. Dancy, J. M. E. Moravcsik, and C. C. W. Taylor eds. *Human Agency: Language, Duty, and Value: Philosophical Essays in Honour of J. O. Urmson*. Stanford, CA: Stanford University Press: 189-198.
- Williamson, T. 2004. "Philosophical 'Intuitions' and Scepticism about Judgement." *Dialectica* 58, 1: 109-153.
- Williamson, Y. 2007. *The Philosophy of Philosophy*. Oxford: Blackwell.
- Williamson, T. 2011. "Philosophical Expertise and the Burden of Proof." *Metaphilosophy* 42, 3: 215-229.
- Wootton, D. 2015. *The Invention of Science: A New History of the Scientific Revolution*. London: Penguin.