

Skookumchuck, Kiidk'yaas, Gibbard: normativity, meaning, and idealisation
 (critical notice of Allan Gibbard *Meaning and Normativity*)
draft, but essentially as in *Canadian Journal of Philosophy* 2014, including pagination

Adam Morton
 University of British Columbia
iadam.morton@ubc.ca

abstract: Gibbard argues that “meaning is normative”. He explains the claim with an account of the normative which bases it on the process of planning, taken in part as issuing instructions to oneself. It seems to entail that the right kind of plans make norms. One *ought* to continue adding with plus rather than quus in a Kripkenstein horror story. I focus on Gibbard's characterisation of normativity: it is not what one might expect. The main purpose of this review article is to present the way of understanding normativity that makes most sense of what he says, and which makes some otherwise implausible assertions defensible and perhaps even true. I give reasons for thinking that Gibbard's understanding of normativity-through-plans cannot do the work he wants it to. I also argue that he is onto something right, and it opens interesting new questions.

Keywords: Allan Gibbard, meaning, normativity, rationally ideal agents, truth

Beginning a long walk through the woods, you anticipate a fork in the trail much further in. One branch leads to rapids where tide sometimes runs against gravity, so kayakers can perform improbable feats. The other leads to a rare golden spruce tree, one of a handful in existence. You resolve to go to the rapids: your plan takes the form of an imperative to your future self "choose the rapids!". Then hours later you get to the fork. You shrug off your earlier resolve, which suddenly seems cheap to you, and set off for the spruce¹.

Ought you to have gone to the rapids? Was there any normative component to your earlier resolve? Have you broken faith with your earlier self? Affirmative answers to any of these would need a lot of defending, to say the least. Allan Gibbard's impressive, wide-ranging, and powerfully argued *Meaning and Normativity* (Gibbard 2012) articulates an account of the normative which bases it on the process of planning, taken in part as issuing instructions to oneself. It seems to entail that the right kind of plans make norms. Gibbard's theme is easy

¹ For outsiders: Skookumchuck is a reversing rapids, Kiidk'yaas was a golden spruce, both names come from Salish languages, which defy easy translation by parallel construction.

to state and hard to unravel: meaning is normative. One *ought* to continue adding with plus rather than quus in a Kripkenstein horror story, for reasons that *seem* parallel to taking the rapids trail rather than the spruce trail. In this review article I shall focus on Gibbard's characterisation of normativity, because on first reading I found it deeply mysterious. After rereading and reflecting I think I know what Gibbard means by "normative": and it's not what one might expect. The main purpose of this review article is to present the way of understanding normativity that makes most sense of what he says, and which makes some otherwise implausible assertions defensible and perhaps even true. My negative and rather grumpy purpose here is to give my reasons for thinking that Gibbard's understanding of normativity-through-plans cannot do the work he wants it to. It seems to me blinkered by implausible and arbitrary assumptions, some of which seem to return us to earlier ages of philosophy. There is a positive function, too. I think he is onto something right, and it opens interesting new questions. I'll get to that. (I should say also that there are fascinating and valuable discussions of a number of issues which I shall not mention.)

1. exposition

Gibbard initiates the reader slowly into his take on normativity. At first it seems possible that there is a single basic Ought in terms of which one ought to do this, want that, or believe these. He cites approvingly Ewing's idea that there is a single basic ought, that of rationality. At this stage it seems possible that moral, prudential, and epistemic oughts are variants of the same, to which Gibbard wants to add an expressivist understanding. Frequent comparisons with the moral ought, on which Gibbard's expressivist position is well-known, might encourage this. But it later emerges that the only ought that is relevant to the normativity of meaning is the ought of rational belief, as found in the central cases of deductive consequence and Bayesian-inductive evidence. Analogies to morality and prudence are only that: hints at a general style of analysis that suggests itself in all these domains. What one ought to mean by a word, and how one's words ought to be interpreted, is determined by the reasons one has to believe assertions involving them. And in the end this is to be consistent with an expressivist account along the general lines he would endorse for other attitudes he would also classify as normative.

Gibbard constructs a semantical theory to fill this out. (Officially a meta-theory, saying what "meaning" means, rather than one which with enough data will specify what a particular word means. But the line is often hard to discern.) Meaning is to centre on synonymy. Most of what linguists do under the heading of semantics, studying for example the different ways times or quantification are represented in the world's languages, has to be shoehorned in one way and another. There is no treatment of context-sensitivity and only stipulation about the treatment of intensional constructions, with no theory to back up the stipulation. The examples are from languages such as French and English where

shared beliefs and influences, and centuries of contact and co-evolution, have led to parallel vocabulary and deeply ingrained traditions of translation. This makes it easier for Gibbard to suppose that for each term of either language there is a series of specifications of the role that the concept plays in combination with other terms, which can be represented in the other. This is its Carnap sentence: for "cat" it might say that if there are cats then they are the species which annoy dogs, catches mice, and so on. "Dog", "mouse", and so on, will have similar Carnap sentences in interlocking ways partially grounded in bare experience. Then we can translate "cat" by "chat" because (a) the towers of Carnap sentences can be lined up and (b) we are disposed and inclined to continue to use the terms in accordance with them.

The theory of meaning is presented as a theory of truth conditions rather than as directly verificationist. The gap between these is filled with a layer of idealisation: what is true is what a thoroughly rational agent -- really stupendously impossibly rational and omniscient -- would accept if *all* the evidence were available to her.² Perfectly rational agents are central to Gibbard's account, but it is not really clear what counts as evidence. One would think it was cheating to declare that "neutrinos have mass" is true if someone who had the evidence that neutrinos have mass would accept it. But on any understanding of evidence a being to whom *all* the evidence was available, and who could process it, would not be human, thus making it less obvious that no ideally rational agent could hold "neutrinos have mass" as evidence. The idealisation is combined with a Horwich-style minimalism about truth, so that a wide range of assertions can count as true. (This complicates the eventual expressivism. "You ought to apologise" is true if and only if you ought to apologise, even when we do not take the ought as representing anything in the world.) The minimalism and the idealisation have a puzzling relation to each other. Agents with more than human capacities would be able to understand things that we cannot, and many of these things would be true, but we cannot ascribe truth to them in minimalist terms while speaking intelligibly. A version of the problem, that we can ascribe truth to an agent's assertions without being able to state the content of those assertions, arises with unidealised humans too. The ten year old child of a Nobel-prize-winning economist will be able to say, meaningfully, "Mother said that general equilibrium theory is false". But if the child tries to fill this out as "General equilibrium theory is true if and only if p ; but not p " he will find nothing that he can understand to fill the p slot.

What is normative about this? It sounds like an ingenious bridge between old-fashioned verificationism and a newer truth conditional semantics. Normativity enters in two stages. The first is simply the reliance on what agents with unlimited and unerring reasoning powers would conclude. This allows Gibbard to appeal to what a person's semantical commitments really involve, as distinct from what the person may be inclined to take them to involve. This leaves room for the possibility that what a person is inclined to say -- e.g. in adding a pair of hundred-digit even numbers and getting an odd one, easier to do in some

² Bruce Hunter points out a potential circularity here: presumably all evidence means all truths of some kind, but if truth is defined in terms of evidence we may have a problem.

notations than others -- may be simply a slip, even if it is one that emerges naturally from the person's cognition and which it would be next to impossible to avoid. To take this line we need first of all some specification of what someone's fundamental semantic commitments are.

One difficulty in specifying fundamental commitments lies in the fact that they are usually implicit in our behaviour and dispositions. I take Gibbard's reliance on Carnap sentences to be an insistence that *something* be fundamental. Different ways of grounding a person's meanings could be incorporated in the same kind of general account. It even seems possible to me that a coherentist version might be possible, where one took the best fit to a person's overall pattern of linguistic responses and then considered how an infallible perfect reasoner would project to novel cases. (This obviously would leave a lot more room for indeterminacy.)

Gibbard uses the standard example, since Hartry Field (1973), of translating "mass" from Newtonian to Einsteinian dynamics to illustrate these issues. He lists the standard translations from the language of a mid-nineteenth to that of a mid-twentieth century physicist and suggests that which of these is to deliver the translation depends on the Carnap sentences that lie behind the old physicist's usage. (And I *think* he deliberately leaves room for saying of the case where there are too many or too few of these that we can decide to interpret which ones to use in interpretation, thus making it the case that for us the original words mean this or that.) I wrote "lie behind" the usage, but it is important that this is taken as idealised: it is what really follows from or is best supported by the physicist's basic understanding of *mass* and related terms. Perhaps the physics example is misleadingly simple, as a physicist is likely to state or refer to some defining features at some point. In messier cases one may have to cope with inconsistency, gaps, and over-determination. I think over-determination is the biggest challenge in that even a very explicit scientist is likely simply to assume that some criteria coincide, which a later age will take as far from obvious.

Suppose we know what a speaker's most basic hold on the meanings of her words is, and we know how she compounds simple meanings into composite ones. Then the first take on what it is to say that one of her utterances would be interpreted, by a perfect reasoner, as synonymous with a sentence as said by the reasoner, is that the same items of evidence would be taken by that reasoner as supporting or undermining either. (And presumably an utterance of one mere mortal in one language means the same as one of another mere mortal in another when the evidence for either, as interpreted *sub specie perfectionis* would be the same.) But to say just this leaves open the question how the perfect reasoner knows what a person's basic grasp of her words is. It is here that we need the other main plank of Gibbard's account, planning. For Gibbard the important situation is when one person plans on behalf of another, and he presents an ingenious analysis of it which, whether or not it solves problems of normativity, makes one think about the nature of identification and of thinking on someone else's behalf. He deliberately considers a problematic case, one of us now

planning on Caesar's behalf not to cross the Rubicon. The analysis is resolutely internalist, so we are supposed to build in only information and preferences that Caesar would have had. Then in fact the feat of thinking and planning from Caesar's full point of view is humanly impossible – certainly for a modern human and perhaps even for a contemporary of Caesar --and the most we can say, though Gibbard does not put it this way, is that we think that if a plan were made with Caesar's full information and preferences then it would be one that refrains from crossing. Think of the obstacles to making such a plan oneself-for-Caesar. One would have to access all his culturally specific preferences: his attitudes to honour, his family name, his position in history, all different from ours to an extent that classicists struggle to get the roughest sketches of. One would have to access all his minute information about his situation: there are hundreds of relevant individual contemporaries even whose names we do not know, not to speak of the geographical and logistical facts that are lost for ever. None of this need block the opinion that all things considered and based on intuition and guesswork one thinks that for Caesar holding back is a better plan. But it does make the process of actually making the plan, from what is in fact Caesar's point of view, problematic in the extreme³.

Yet it is this process of actual accurate plan-simulation that Gibbard needs, if the judgement that Caesar ought to hold back is to be really normative, on Gibbard's understanding of the normative. For that understanding roots the normativity in a kind of performance: one makes the plan to do A on *a*'s behalf and in terms of *a*'s full complement of information, preferences, inclinations, and so on. Then one announces the result as "*a* should A". This produces a dilemma. Either one is reporting a limited and incomplete stab at making a plan on *a*'s behalf, and announcing the plan one has made, or one is asserting the opinion that if the superhuman feat of making a perfect plan for *a* were accomplished, the result would have certain features. Each has its price. If the former, there is the possibility that in fact doing A would have terrible effects, which an impossibly full analysis from *a*'s point of view would have foretold. Then *a* ought not to have done it, even from her own point of view. If the latter, one is not making a plan or issuing a resolve but making an assertion of fact, that will be true or false. I return to this dilemma in my final section, "agreement".

Back to meaning now, but stick with Caesar. "lacta alea est", he notoriously said, crossing the Rubicon. Except he didn't: like most upper-class Romans of his time much of his conversation was in Greek (he never said "et tu Brute" either⁴) and he quoted a line of the Greek poet Menander with the standard translation "let the game be ventured". Did he really mean "the die is cast", as we usually suppose? To know, we have to face the same barriers of time and culture as in the Rubicon decision. In Gibbard's terms we ask what Caesar ought to take as specifications (Carnap sentences) of his terms, and what he ought to take as relevant evidence. Suppose a centurion comes up to him and points to soldiers playing with dice in a way that is clearly disanalogous to his situation. Does he treat this as suggesting that he mis-spoke, or not, supposing that he is thinking

³ Some of the problems are reminiscent of theological worries about whether God, because of as well as in spite of his powers, can empathise with the experience of mortals.

⁴ Bruce Hunter points out the resemblance to 19th century Russian aristocrats' use of French.

objectively and widely? If we answer Yes then, on Gibbard's account of what "meaning" means, we should feel more confident that he is making a metaphor about dice rather than about games with uncertain outcomes in general.

I have been keeping the two elements, perfect reasoners and planning for another, separate. But it is essential to Gibbard's project that they be combined, even though when discussing either the other tends to be in the background. The challenges to planning for another are in part due to our limitations in putting ourselves in the situation of another. But these are accidental human limitations -- or so the project assumes -- analogous to our limitations in following correct patterns of inference and handling complex information. Postulate, then, another dimension to intellectual perfection, the capacity to put oneself entirely and accurately in someone's situation. A thinker with this capacity would be able to hold in mind all of another person's assumptions, presuppositions, and dispositions to believe, however distant the culture and however remote the science, and to think in terms of them. In particular the gaps in knowledge, out-of-character whims, and inexplicable preferences of the other would all be held in mind. Gibbard combines this form of perfection with perfect reasoning ability under an admittedly prescriptive label of "plans". One plans what to do if e.g. in Caesar's sandals, taking all these facts on board and drawing full and correct consequences from them. (I return below to the fact that this is not something that real people can do, and not in general something they should try to do.)

For the analysis of "meaning" the important plans are plans to believe that result from putting oneself in someone's exact epistemic circumstance. Belief is not generally voluntary, so that the plans in question are plans to put oneself in a situation which will result in one's acquiring the belief. (Without, presumably, abandoning any of the beliefs that constitute the person's epistemic circumstance.) There are complications from the fact that a person can misunderstand her evidence and may be ignorant of her own "linguistic proclivities". So, stirring in such considerations in the course of a very subtle and difficult passage (pages 174-187), we get this formula for "s as said by S means that p as uttered by oneself": for a sufficiently wide variety of epistemic circumstances E one plans to assert s if one is about to be S in E if and only if one plans in one's own right to accept p if in E. One plans something in a circumstance, in the relevant sense, remember, if a perfect thinker in that circumstance would do it.

We now have the heart of the account, and we can finally see what "normative" in the slogan "meaning is normative" must mean. Normativity is what you get when you combine two idealizations: the ought of procedurally perfect epistemic agency and the capacity to internalize the full situation of another. There are other legitimate uses of the word, and below I contrast some of them with Gibbard's use. But, unpacking his terms, this is what the claim is. A target assertion has the same meaning as a given assertion when an epistemically ideal agent who could put herself perfectly in the situation of the (merely human) person asserting the target would plan to assert the given one if and only if evidence supported one exactly when it supported the other.

2. criticism

In this section I raise some worries about Gibbard's account of normativity and meaning. Not all of the considerations here obviously need be fatal for his line, but they all are things for which an answer is needed.

a) *Ideal reasoners* The first is a "misleading advertising" or "bait and switch" worry. We are given an analysis in terms of epistemic ought, and the results are paraphrased in terms of what one ought to believe and the plans one ought to make. But the ought in question is what a procedurally flawless being with unlimited information-handling capacities and boundless self-control would believe. There is a big gap between this and what you and I would do well to believe. Behaving as if one is perfect is very often stupid. It's a fool who goes to apologize when he knows that when he begins he will lapse into rage and vituperation. In epistemic matters it is often best to have no opinion on matters where the evidence is too complex to handle. Intelligent scientists do not plan experiments so that the results can only be understood in terms of statistics that they cannot master. And paradigmatic human reasoners often consult one another when they are uncertain -- in fact they have a vice of irrationality if they do not -- in contrast to the splendid cognitive autonomy of Gibbard's ideal thinkers.

If by "ought" and "normative" one means anything along the lines of what would have desirable consequences, the gap here is between what a perfect agent would do or think and what one ought to do or think oneself as one actually is. If the emphasis is more on standards that we hold one another to, then it remains to be argued either that we do encourage one another to have the beliefs and plans that Gibbard's account builds on, and to sanction their neglect, or that we have adequate reason to do so. Since the reasons why it is usually counterproductive to proceed as if one was more able than one is frequently entail that such rational hubris will mess up the projects of others too, we can be pretty sure that the standards we actually hold one another to do not involve pressure to do the impossible⁵.

In fact, there is a strong reason not to assimilate the normative to the ideal, on ought/can grounds. No human person can operate in anything remotely like the degree of idealization (or of isolation) required. So no human person ought to.

These are reasons for discomfort with the slogan "meaning is normative"; they are not so evidently reasons to doubt the analysis that the slogan is short for. But they do suggest problems, I think. In particular they make it hard not to construe meaning as a more determinate business than it in fact is. Human language users do not anticipate all the situations to which their words might apply. It is not in their interest to, and besides it is beyond their cognitive powers. The troubling cases are those in which a situation is not anticipated but could have been were superhuman cognition applied to a person's basic language use. Consider a fourteen year old, to whom we give a quick explanation of ordinals beyond

⁵ For a systematic development of this point of view see Morton (2012).

omega. Then we ask her whether $\omega + 7$ is different from $\omega + 12$. She says No, since she is a very intelligent child and thinks of ω as "infinity". Her answer is wrong, on our standard understanding of ordinals. Does she have the standard understanding? Is she making a mistake on the conception of omega she has picked up? I would doubt that the concept of having the standard understanding, "meaning the same", is definite enough that there is a fact to the matter. Or consider cases like those discussed by Mark Wilson in a classic paper (Wilson 1982), in which the speakers of a language meet situations that they had never expected. (Wilson's example is of people using a word they apply to birds when they first see airplanes.) Suppose that, given the definitions and stipulations they share, there is a most strongly supported choice of terminology for the unexpected situation ("bird", say), but that none of them is remotely capable of seeing this. Is there a fact to the matter about whether the term is correctly applied in this situation?

You may reply that there is a definite meaning in such a case, but the person does not know what it is. If you say this you will have to accept that if, say, "cats are neutrinos" follows from assumptions shared by all English speakers, though none of them has any idea this is the case, then "neutrinos purr" is true in English. And, along the same lines, you will have to accept that when shared assumptions are invisibly inconsistent, as they must often be lest there be mass unemployment among philosophers, then everything means anything. Moreover, to connect back with the previous point about Gibbard-normativity, it will follow that sometimes, perhaps often, using words to mean what they really do will tell against the explanatory and communicative purposes of language.

(And it is not obviously incoherent to consider what we would do if we find enormous l and m such that $n = l + m$ calculated in standard grade school decimal fashion turns out to differ from the sum we get by evaluating long strings of units in binary fashion and reconverting, and the variant sum makes more sense of the process of computation and of counting. What "ought" we to record as the sum? Our present practice would not determine it. Of course we cannot give plausible examples of l , m , n ; but that is just to say that we are imagining facing a situation of a kind that we have not anticipated.)

b) *Truth and evidence*

A second category of worries also stems from the use of ideal cognition. As already mentioned, Gibbard assumes that what is supported by all possible evidence is true. The assumption is too anti-realist for my taste, and I have doubts about its intelligibility, but those are not issues for this piece. Consider a case with limited evidence, though. Does identity of non-conclusive evidence over a range of situations bring identity of truth conditions? For an extreme case where the evidence relevant to two sentences is identical because there is none for either, consider "God's favourite composer is Bach" and "God's favourite composer is Mozart." Assume that there is no evidence for or against either. (If you don't like

these examples you can probably supply others. There are examples from physics, as in Glymour 1972.) But they are surely not synonymous.

Is the example still trouble when it is run through the full Gibbardian machinery of planning for oneself and planning for another? I think so, particularly since self and other can be identical, but the machinery is intricate, with many places where epicycles can be inserted. Behind the particular examples there is a titanic clash of approaches to meaning. Gibbard's account is thinking-based. Meaning comes down, one way or another, to the role a linguistic item plays in a person's reasoning. He gives a substantial and in fact fascinating discussion of how to incorporate demonstratives and token reflexives into his approach, which deserves attention in its own right. The need for the discussion does underline, though, that the causal/interactional/contextual aspect of language does need some special considerations to fit it into a thinking-based picture. Contrasting with thinking-based approaches are interactional approaches, stemming from the suggestions of Kripke and Putnam in the 1970s. Bach and Mozart are different individuals and so, on this other style of theorizing, there is more scope for making assertions referring to them refer to different possibilities and, if you want to put it this way, have different meanings⁶.

(c) *a circularity* Gibbard's account relies on procedurally correct reasoning. In terms of it he gives a response to the Kripkenstein quus question that goes beyond dispositions. A person can be disposed to respond as if the extension of "+" were "quus" although her fundamental semantic commitments better support a standard "plus" extension.

(Presumably it can go the other way, too, though this is not explored. The point would be the same, that not all dispositions to respond determine extension.) Correct reasoning is rule-governed, too, and deviant disposition-compliant extensions can apply just as easily there. Consider for simplicity deductive reasoning. A person accepts something like Peano's axioms, from which it follows deductively that $n+m=s$, for some enormous n , m , s . But given these numbers, in a notation in which she can just barely handle them, some quirk of her brain will make her do the sum as resulting not in s but in a number t seventeen units higher. Moreover the quirk is not isolated, but is an effect of her basic numeral-handling routines. For Gibbard her conclusion is false, because the extension of "+" for her is given not by her dispositions but by the consequences of her semantic commitments, in this case Peano's axioms.

The extension is PLUS *if* $n+m=s$ is a consequence of the axioms. Consider a deviant notion of consequence, schmonsequence, naturally, according to which it follows from the axioms that $n+m=t$. (There will be many other discrepancies between consequence and schmonsequence, of course, all requiring mind-strainingly long deductions. See sections 2.3 and especially 2.4 of Cherniak (1986).) Suppose that were any human person to work it out, a systematic deduction-handling quirk in her brain would kick in and she would

⁶Meaning, I think, is on an interactional account a book-keeping device for explaining the processes that result in reference. If your account relies on truth conditions it is a way of getting the right hand side of T sentences. And if you appreciate Tarski's insight that satisfaction is more basic than truth then word-meaning plus facts fixes satisfaction which fixes truth, after which you do not need to worry about sentence meaning. For a related suggestion see chapter 3 of Simchen (2012).

deduce that $n+m=t$. She would be going wrong if in fact the axioms do entail that $n+m=s$. Do they? It depends on the extension of "consequence", and variants. We can say that $n+m$ is s or t depending on what really is a consequence of the axioms, but on the present supposition we humans will be wrong about which of these it is. Suppose that there is a definite answer to which is consequence, and it yields that $n+m=s$, though we are all confident of the alternative, however much we reflect. (See: I'm sportingly letting pass premises that might be challenged.) Then we have a conflict between two ways of reading Gibbard. There is the expressivist Gibbard, for whom the right answer is the one that, on as careful reflection as you can, you endorse yourself by planning to assert it on the other person's behalf. And there is the absolutist Gibbard, for whom the right answer is the right answer, dammit, given by what really follows from the other's semantic commitments. They are not the same. Will the real Allan Gibbard please stand up?

3. agreement

It must seem as if I think Gibbard's approach is completely misguided. Indeed I do think that the slogan that meaning is normative is seriously misleading, and does not fit what he is actually saying. And I think that the claim to be presenting an expressivist conception of the normative, in connection with meaning, does not fit perfectly with his actual analysis. These are problems of self-description as much as problems with his real conclusions. But I think that the second of these in fact points to a strength, and a stimulating source of possibilities for new theorising. In this last section I will explain what seems to me interesting and defensible.

On Gibbard's very revisionary construal, plans are the things one would get if one had superhuman cognitive powers which allowed one first to consider all information, second see things perfectly from a given person's perspective, and then third form a single optimal intention that one was then assured of being able to fulfil. Suppose for the sake of argument that this is an intelligible conception. The most basic thing to note is that these are not plans that anyone makes, since they are not plans that anyone is capable of making. (And, more subtly, they are not first order consequences of a second order plan to make perfect plans. Not if they are things that we ought to do, since intending to act as if one were perfect is often a stupid resolve, as shown by the examples in a) of the previous section.) The natural first thought of what they might be is counterfactual assertions of what intentions a cognitively unlimited agent -- a very ideal observer, equipped with ideal powers of empathy -- would plan in a given person's place.

This idea of a plan-based normative judgement could be defended, though there would be problems of intelligibility due to its extreme counterfactuality or counterlegality. But it will not deliver one important element of Gibbard's aims, for it does not give an expressivist understanding. On such an understanding a judgement that someone should A is not something to be compared to independent facts

-- if we distance ourselves from minimalism we could say it is not something that can be true or false -- but something one *does*. One expresses a non-cognitive attitude of planning to A on that person's behalf. But now, where his arguments eventually lead him, there is no attitude that a real person can express because there is no suitable attitude a real person can form, no plan they can have. Moreover, a real person can only hypothesise what attitude an ideal agent would have, so the attitude of judging-as-if-ideally is inevitably tentative and inevitably qualified with an "as far as I can see", and thus inevitably cognitive. It is an opinion of what an agent capable of an expressivist position would do. Connected with this are Gibbard's frequent worries about Moorean open question arguments. One can intelligibly ask why the judgement that an ideal observer would form should have any authority over what we consider ourselves bound to respect. The question could be resisted in various ways, but they will all require breaking with Gibbard's fundamental strategy of pushing the open question down to a point where it cannot intelligibly be asked, because at that point the question is addressed to the person who after all possible reflection is actually making the choice, at least in simulation.

I think we can get a good dose of expressivity by understanding Gibbard's account in terms of a three component process. Here I finally move beyond my prevalent grumpiness, to agreeing and approving of something fundamental about the approach. The reading is definitely creative: though it is consistent with much of what Gibbard says it is not anything he encourages or even suggests. One first expresses an endorsing, approving, or admiring attitude to the idealised point of view, which itself (second) possesses an identificatory perfect plan towards the person's situation which eventuates in the act in question. One also (third) expresses a conviction that the idealised point of view, inasmuch as one can get one's mind around it, would eventuate in the act. This last expression is of a truth-evaluable opinion and thus not the kind central to expressivism, though the first, endorsing, expression is. Many variations on the attitudes expressed are possible. One promising line is to take it as defining a single complex attitude to a content -- we could even call it a kind of plan -- where one considers that p via one's attitude to an evaluative point of view towards it⁷.

One advantage to the suggestion is that it captures the delicate twining of objective and subjective points of view in normative judgement that has been defended by Nagel and others. Another is that it makes a place for the element of defeasibility that even an expressivist should acknowledge: we just do often say what ought to happen and then retract it later, not because of any change of attitude but because we think we have summed up the facts wrongly. And in the evaluative life of real people no normative judgement is immune to the possibility of this kind of retraction.

But to my mind the greatest advantage of the suggestion is that it frees Gibbard's account from the straightjacket of Ought. In real life oughts usually serve as summings-up of the pros and cons of an action or attitude. Before we get to the point of considering what we ought to do we usually consider more finely

⁷ There are resemblances here to part III of Morton (2013).

grained normative considerations. And we often agree up to that point but part ways there. We agree on what would be fairest, most considerate, or in most people's best interest, but then we disagree about what we ought to do. (Note the "we": one common pattern is that we can decide the more focussed evaluations individually, but need to put our heads together to get an aggregated summing up, typically for shared action.) This fact can be accommodated on the suggestion I am now exploring, by acknowledging a plurality of idealised points of view. There is the point of view that counts everyone's interests equally, there is the point of view that cares for the most vulnerable, there is the point of view that considers social arrangements that are fair and generalizable, and more. All of these come in idealised versions: really grasping everyone's interest, really feeling everyone's suffering, or really and fully calculating all alternative social arrangements, though more than a feeble approximation is in each case beyond the scope of any real person. (If you want, there are ideal agents Aphrodite, and Athena, and Apollo.) And there is a variety of attitudes one can have to these points of view, respecting some, resonating to others, and full-heartedly endorsing others. No wonder we often agree up to the point where we have to decide what to do. (And sometimes we shudder and say: this will be an atrocity, but in the circumstances we ought to do it.)

Gibbard's aim in this book is to construct a coherent normative attitude to linguistic meaning. The three component suggestion helps here too. I suggested ways in which Gibbard's understanding of meaning seems somewhat blinkered. The first was its resistance to indeterminacy, to situations where the best conclusion seems to be to describe various semantic features of a sentence -- perhaps what range of ambiguity various words have or how well a formulation fulfills a theoretical purpose -- without having to sum this up with a minimalist "s is true if and only if p". Separable idealisable components help with this. We can describe an aspect of the semantics of a person's use of words without claiming to give a total single judgment on what the meaning is. In particular we do not have to say what utterances the person "ought" to make in specified circumstances. Without doing this we can still say what would be the enlightening, careful, or fruitful way for her to continue, or which words are ambiguous in which ways and which constructions serve which normal purposes.

But the most promising application is to semantic deference. I include cases where our grasp of a word depends on our trusting that others have a better understanding of it, as in the examples that drive semantic externalism, or where we take a word to refer to whatever another has in mind (or whatever actually explains what others use it to explain), as in the examples that drive causal accounts of reference. No doubt there are important differences between different kinds of deference. But in all of them a semantic value depends on what we think a person in the position of another would conclude if they were thinking well *in a particular respect*. (What the actual person concludes may be a mistake. What we think they should think may be a mistake too, and it is a very delicate matter how this effects what we mean by our words.)

Another blinker was the restricted account of the things we can mean, or at any rate of the terms in which meaning is to be understood. The point is best put in terms of meaningfulness rather than meanings. We routinely acknowledge that things that others say make sense, though our understanding of them is very limited. For me this includes a lot of economics and physics, and distressingly much philosophy. And in the other direction we should allow that sometimes our failure to see the details of the emperor's brocade is not our fault. Both possibilities are allowed if we accept that there are semantically ideal points of view, from which the utterances of angels, geniuses, and future humans could be assessed. And there are also points of view that though very limited are factually and semantically better informed than our own, so that the undergraduate you chat with in the coffee shop may convince you that something makes sense, though it is not worth the effort of getting to the point where you can yourself see how this is so. Semantic polytheism also allows that we can attribute, e.g., explanatory power to a claim without being sure of what direct evidence would be crucial to it or being able to paraphrase it in transparent terms. This will appeal to the expressivist component, more directly than on Gibbard's official version: you can *use* the claim to explain data, even though there are other things you cannot do with it.

For Gibbard there are two poles to normativity, ideal agency and identification. These are intimately connected, as he presents them, and, in his account of meaning at any rate, the combination undermines both real normativity – what in earlier work Gibbard has expressed as “making sense” -- and expressivism, since the ideal agent occupies the position where the norm-expressing person should be. In this final section I have been suggesting that both of these problems are mitigated if we add a third element, the real person's attitude to a potential variety of idealised points of view. Then we can have it both ways. It's a real human attitude to the ideal(s), and it's the person in question who is having it.

ACKNOWLEDGEMENTS: I have been helped by comments from Bruce Hunter and conversation with Ori Simchen.

BIBLIOGRAPHY

- Cherniak, Christopher (1986). *Minimal Rationality*. M.I.T. Press.
- Field, Hartry (1973). "Theory change and the indeterminacy of reference." *Journal of Philosophy*, 70 (14): 462-48
- Allan Gibbard (2012). *Meaning and Normativity*. New York: Oxford University Press.
- Clark Glymour (1972). "Topology, Cosmology, and Convention." *Synthese*, 24: 195-218
- Morton, Adam (2012). *Bounded Thinking: intellectual virtues for limited agents*. Oxford: Oxford University Press
- Morton, Adam (2013). *Emotion and Imagination*. Cambridge: Polity Press
- Simchen, Ori (2012). *Necessary Intentionality*. Oxford: Oxford University Press
- Wilson, Mark (1982). "Predicate meets property." *Philosophical Review*. 91 (4): 549-590