Précis of *Thinking How to Live*[*]

ALLAN GIBBARD
University of Michigan, Ann Arbor

There doesn't seem to be a fact of what to do. What one *ought* to do is another matter: whether there is such a fact is contentious. The hypothesis of the book is that the two amount to the same thing: what one ought to do is a question of what to do. Thinking what one ought to do is thinking what to do. The concept of *ought*, I propose, can be explained on this pattern—not for every sense of the term, but for a crucial sense that figures in normative concepts: moral concepts, for example, and the concepts of being rational, credible, shameful, or enviable. Thinking what's admirable, for instance, is thinking what to admire, and thinking what's credible is thinking what to believe. There is no special mystery, then, in normative concepts; we needn't appeal to "non-natural qualities" to explain them. If we understand concluding what to do, we understand concluding what one ought to do.

Are oughts, then, matters of fact? In a minimalist sense of the term 'fact', there are of course facts of what a person ought to do. In this book I am agnostic on whether there is a more demanding sense of 'fact', or whether the word 'true' has a sense that is more than minimal. The book aims to show, though, that matters of what to do act in many ways like familiar, uncontentious facts. Are these facts, then, non-natural? A straight answer to this last question would tell us little; the question cries out for distinctions.

From my earlier book *Wise Choices, Apt Feelings* (1990), I retain central theses. I start, though, with more meager resources, and I draw further consequences. My earlier account I situated in a speculative try at a realistic, naturalistic human psychology. The newer book starts simply from us as planners, able to think what to do—now and in contingencies. From this bare starting point, familiar normative phenomena emerge. We see how *oughts* "supervene" on natural *iss*. A kind of naturalism is a further upshot: there is a broadly natural property, I conclude, that constitutes being the thing to do. The resulting system mimics most closely the non-naturalism of G.E. Moore and A.C. Ewing. I think of the book, then, as realizing a form of *quasi-realism*, as Simon Blackburn calls the program we share: from a

---

[*]   Alan Gibbard, *Thinking How to Live* (Cambridge, MA: Harvard University Press, 2003).

basis that excludes normative facts and treats humanity as part of the natural world, it explains why we would have normative concepts that act much as normative realists claim. Among other things, I address objections that critics have had to this way of proceeding.

I. Preliminaries

1. Introduction: A Possibility Proof. Do moral assertions state facts—facts of a special, non-natural kind? Or do they express attitudes or some other such states of mind that aren't straight beliefs of fact? We could imagine a parallel debate about decisions: A decision can be couched as an imperative, as when Holmes, to prepare to elude Moriarty, tells himself, "Pack the bags now!" Or it can be couched as a straight assertion, "I must pack now." The assertion, we might worry with J. L. Mackie, ascribes to packing a mysterious property of "to be doneness"—whereas the imperative does no such thing. Are there special "must" facts, we could now debate, or do "must" assertions simply express conclusions on what to do? Starting with this parallel debate avoids one of the sticky questions of moral theory. For moral right and wrong, many philosophers are "externalists": concluding what one morally ought to do, they think, leaves open the question of what to do. We can't be externalists on what to do; conclusions of what to do have an automatic bearing on action. We can ask, then, whether terms like "must" or "ought" are best explained as attributing some special kind of property with built-in to-be-doneness, or as terms for saying what to do. The book develops the latter hypothesis.

The initial project is a possibility proof: I show that a language with a term for expressing decisions is possible. I show too that such a language would work much as familiar normative language works. Some philosophers argue that factual-seeming language can't possibly work in the way I describe. Negation, disjunction, and the like, they say, can't be accounted for in such a language. Whether or not our actual language fits the account I offer, though, a language that does is at least possible. I spend a number of chapters demonstrating this. For my stipulated language, I appropriate the term "the thing to do", and let it expresses conclusions on what to do. To conclude, for instance, that right now packing is the thing to do, I stipulate, is to decide on packing. My stipulations extend to planning for contingencies—even wildly hypothetical contingencies like being Holmes in his plight. To conclude that packing was the thing for Holmes to do, I stipulate, is to adopt a plan for the contingency of being in Holmes's exact circumstances.

Such a language, I set out to show, is at least possible, and it will operate in a remarkably fact-like way. The states of mind in the account are, in a broad sense, states of planning: conclusions on what to do in various contin-

gencies, foreseen or hypothetical. Talk of "the thing to do", then, expresses planning states. The account is *expressivistic*, in that it explains "thing to do" language as expressing states of mind that aren't, at the outset, explained as beliefs in some special kind of fact.

2. *Intuitionism as Template: Emending Moore*. The ultimate aim of the book is, of course, to help explain our actual language and thought. For normative concepts, then, I need a template to match, an account of how our normative language and concepts work. My template is an emended Moore. His "What's at issue?" argument refutes analytic naturalism for ethical and other normative concepts. It indicates that these concepts are distinct from the various naturalistic concepts that people have claimed they amount to. The naturalism it refutes, though, is for normative concepts, not for the properties that normative terms signify. Moore's argument doesn't show that normative *properties* are distinct from natural properties. Indeed we can say that what Moore called "*the* good", the property signified by the term 'good', is a natural property. It is the property that necessarily, all and only good things have. The template, then, is naturalism for properties and non-naturalism for concepts. (I also depart from Moore and follow Ewing in thinking that the primitive normative concept isn't *good* but a primitive *ought*.)

II. The Thing to Do

3. *Planning and Ruling Out: The Frege-Geach Problem*. I stipulate the phrase 'the thing to do' as a device for expressing plans. This raises the "Frege-Geach" problem: how can such a phrase embed in larger contexts, in such a way as to make standard logical inferences valid? A good solution has already emerged in the philosophical literature, and I offer a version of it. Take, as an example, a plan-laden disjunction, "Either Moriarty is here already or the thing to do is to pack." This derives its meaning from the combinations of fact and plan that it rules out—in this case, a single such combination: thinking that Moriarty isn't already here and yet rejecting packing. In general, I propose, plan-laden constructions rule out one or more combinations of plan and prosaic belief. An inference is logically valid if accepting the premises and rejecting the conclusion rules out every such combination.

The plans in question may be for contingencies that are wildly hypothetical. Holmes, imagine, thinks that for Mrs. Hudson, going to the police is not the thing to do. He thereby rules out a plan for the contingency of being Mrs. Hudson in her exact situation. (Necessarily, to be sure, he isn't Mrs. Hudson, but as David Lewis taught us, plans consist in the self-attribution of properties, and Holmes can hypothetically self-attribute the property of being Mrs. Hudson.) One feature of Mrs. Hudson's situation is that in fact she will

go to the police, but Holmes can still intelligibly ask himself what to do if in that situation, and reject going to the police.

We could idealize all this by imagining a plan that was utterly complete, providing for every conceivable contingency. Call this a *hyperplan*. It joins with a completely opinionated state of mind on how things stand to form a *hyperstate*. The content of a plan-laden belief, we can now say, is given by the set of hyperstates it rules out.

At this point, we need a crucial refinement. We must distinguish two ways of choosing: from indifference and from preference. If a wiser ass than Buridan's had chosen one of the bales of hay, she still wouldn't, in an important sense, have *ruled out* or *rejected* choosing the other. A plan, then, permits some alternatives and rules out others, where permitting is a kind of ruling out. What permitting rules out is this: rejecting an alternative from preference. I stipulate the term 'okay' for permitting: To believe that an alternative is okay is to permit it to oneself.

4. Judgment, Disagreement, Negation. This chapter addresses objections that have been raised to such a program, and asks what features of planning suit it to quasi-realism. First, truth: Plan-laden statements will be true or false in a minimal sense. It's false that the thing to do is to jump out the window—and to say that this is false is just to say that jumping out the window is not the thing to do. In the book I use "true" in this minimal sense. As for whether there is a more demanding sense of the term, I'm agnostic in this book. I agree with Horwich that logic is not to be explained in terms of truth, and my solution to the "Frege-Geach problem" offers an alternative basis for logic.

Horwich, though, argues that no such solution is needed: I could instead just say that ascriptions of being the thing to do express contingency plans, and that "is the thing to do" is a predicate. Dreier argues that we can't always just stipulate that a term is a predicate. We couldn't, for instance, give meaning to a term 'hiyo' like this: specify that "Bob is hiyo" is, like "Hey Bob!" used to accost, and then just declare the phrase 'is hiyo' a predicate. The difference between this and what I do, I argue, lies in two features. First, the predicate must be used to express a state of mind, such as, for instance, belief or approval. Second, the state of mind must be one we can agree with or disagree with. One can disagree with a plan, and this is what allows for plan-laden predicates. The notion I treat as primitive, then, is the mental operation of *disagreement*. We can recognize this as something we do as we elaborate our plans. To take it as undefined might seem philosophical theft, but the same goes for the usual practice of taking negation as primitive. Why would the one be illegitimate and the other not?

What state of mind does one express with an assertion? How is it that "Packing is okay to do" expresses the state of mind of permitting packing,

and not of *believing* that one permits packing? Here, disagreement holds the key. To disagree with the assertion "Packing is okay," we say "No it isn't," and thereby disagree with permitting packing. To disagree with this assertion is not, though, to disagree with believing that the speaker permits packing. The assertion "I permit packing", in contrast, expresses the belief that one permits packing—in that to deny it we say "No you don't," thereby disagreeing with *believing* that the speaker permits packing. The test of what state of mind is expressed is disagreement: In denying the assertion, what state of mind does one disagree with?

5. Supervenience and Constitution. This chapter develops further aspects of what Blackburn calls "quasi-realism" for talk of "the thing to do". Anyone who thinks what to do, I argue, is committed to realist-sounding theses of Supervenience and Factual Constitution. These follow from a *Principle of Commitment* that applies to all reasoning: A person is committed to a claim if in every hyperstate he could reach without changing his mind, he would accept the claim. It follows from this, I first argue, that planners are committed to supervenience, to accepting that being the thing to do supervenes on prosaic fact. It supervenes in the sense that necessarily, if two situations differ in what's the thing to do, they differ in some matter of prosaic fact. It follows too, I argue, that planners are committed to a claim of Factual Constitution: that there is a prosaically factual property that constitutes being the thing to do. To *constitute* being the thing to do is to be necessarily equivalent to being the thing to do, so that in any possible circumstance, an act is the thing to do just in case it has that property. The property, I argue, will be recognitionally grounded, and so be a natural property in a liberal sense of the term. Still, to call an act "the thing to do" is not to ascribe that property as such. I discuss whether names and "thick" apprehension fit this thesis, and find that they do.

Planners, then, are committed to a naturalism for properties and a non-naturalism for concepts. For the thing to do, whereas the concept is not naturalistic, the property is broadly natural. To this claim, I demonstrate, anyone who thinks what to do is committed—and as a planner, I voice the claim to which I am committed, saying: The property of being the thing to do is a broadly natural property. As for which broadly natural property it is, the question is not settled by meanings alone. It is a substantive question of how to live. Egoistic hedonists, for instance, think that it is the property of being *egohedonic*, of holding out greatest prospects for the agent's net pleasure. To think this is at least coherent—and it is coherent too to think it mistaken. One of these two views must be wrong, but the mistake, whichever it is, is not conceptual. The claims of Supervenience and Constitution pertain not directly to properties but to concepts. No property is plan-laden as such, but some concepts are.

6. *Character and Import*. The previous chapter treated the objects of thought as propositions, understood as ways the world might have been. This is inadequate, as familiar examples show: Water = H_2O is the same proposition as water = water, but it is coherent to believe the second and deny the first. In this chapter, I briefly systematize a "two-dimensional semantics" to represent objects of thought, and apply it to plan-laden concepts. Consider a plan-laden concept like *thing to do* and a descriptive concept of the property that constitutes being the thing to do. (Hedonistic egoists think that the property is that of being egohedonic.) It might be thought that the difference between the two was just another instance of the kind of difference we find between the concept water and the concept H_2O. I indicate that Moore was pointing to a further dimension on which concepts can differ.

To represent all this, I proceed to develop a set of technical terms, distinguishing the *state of affairs* that a sentence *signifies*, the belief that it *conveys*, and the *proposition* that it *invokes*. Each is represented by features of a *character matrix* for the thought. Plans then add a further dimension, and we can now speak of the *extended character* and the consequent *extended import* of a thought or other concept. The extended character of a term is settled by the term's meaning, determined by rules of language, whereas its (plain) character depends not only on its meaning but on how to live. If, for instance, egoistic hedonists are right, "Packing is egohedonic" and "Packing is the thing to do" have the same character but different extended characters. If perfectionists are right, the two have not only different extended characters but different characters.

III. Normative Concepts

7. *Ordinary Oughts: Meaning and Motivation*. Thus far, the book has been devoted to plan-laden judgments. What does this have to do with judgments we ourselves make and voice? A hypothesis to explore might be that normative terms like 'ought' straightforwardly express plans. I press a more qualified hypothesis: With normative language, we do mix fact with plan, but just how won't always be determinate. In this chapter and the next, I suggest patterns for how we might mix plan with fact in our judgments to form plan-laden concepts—some of which might be normative concepts that we actually have.

Again, disagreement is the key. We "track" a conversation by keeping track of what agrees or disagrees with what, and accordingly, I use patterns of agreement and disagreement as my instrument of analysis. We must expect a kind of slack in language: when two claims are equivalent under the presuppositions of an inquiry, there may be no clear fact of the matter which claim is being expressed. Imagine, as an example, that we presuppose, in our normative talk, that only delight is to be sought for its own sake. There *might*

then be no clear fact whether 'good' means "to be sought", or rather, "is to be sought because it is delight". The second meaning would make 'Only delight is good' analytic, whereas the first would make it synthetic. As for a person who thought that delight is not always the only thing to seek, she would thereby reject a presupposition of the term 'good'. Such a person should then say, "'Good' is not one of my words." Our ethical and other normative discussions proceed against a background of presuppositions on questions of how to live, and so there may often be no fact of the matter how naturalistic and planning elements interact in our concepts and language. This may be the case for terms of ours like 'good' and 'ought'.

Still, if 'ought' in some ordinary sense implies *is the thing to do*, implausible things follow, we may worry. I don't then believe that in this sense I ought right now to defy a bully unless I do defy him. A person, though, often isn't of one mind in his plans. If push comes to shove and you don't defy the bully, you don't, at that very moment, plan with every aspect of your mind to defy him right then. In a crucial sense of 'ought', I say, you are of more than one mind, right then, on whether you ought at that very moment to defy the bully. Unlike my 1990 book *Wise Choices, Apt Feelings*, this book is not about the psychology of *ought* but about its logic and epistemology. There may often be no clear, sharp fact of what constitutes accepting a plan, but a state of mind wouldn't amount to planning if it weren't of a kind that normally plays the right systematic role in leading to action. Likewise with judgments of "ought". Let Ira match a rigorous moralist in his "ought" assertions but match an egoistic hedonist in his actions and motivations. He's then just aping having the moralist's beliefs on what we ought to do. He might even use a different term—'should', for instance—to match the egoist's use of 'ought', and guide himself by what he concludes he "should" do. In that case, his term 'should' means what the rest of us mean interchangeably by 'ought' and 'should', and his term 'ought' doesn't. His dispute with the egoist is only mouthed, whereas his real dispute is with the moralist—a dispute on how to live.

8. Normative Kinds: Patterns of Engagement. This chapter explores another form that plan-laden concepts might take, and uses the apparatus of hyperdecided states to canvas ways we can engage or fail to engage with the plan-laden thoughts of others. I offer an alternative to some aspects of Bernard Williams's treatment of "thick concepts".

Geoffrey Sayre-McCord proposes a theory of "moral kind terms" which parallel natural kind terms in their behavior. I broaden his proposal to speak of "normative kind terms", and explore how they might work as plan-laden terms. A normative kind theorist should not be a hard-line externalist who thinks that a "sensible knave" or a complacent "irrationalist" might fully share our normative concepts but not at all be guided in terms of them. We

treat a kind—brutality, for instance—as normatively "high grade", as a good candidate for what a term in a language might signify, when it figures centrally in our planning explanations. Brutal actions, we hold, are to be shunned and condemned because they are brutal. The meaning of 'brutal' is not purely descriptive: an exotic term whose use is not to condemn but to celebrate actions won't mean "brutal", even if its users apply it mostly to brutal actions. A term that means brutal must play the right role in its users' explanations of what to oppose.

Williams spoke of a "relativism of distance". I explore how we might sometimes be able to engage the thoughts of peoples distant in time, place, and ethos, treating their thoughts as subject to agreement or disagreement, and how we might sometimes be unable to treat their thoughts this way. When must we demur, neither agreeing nor disagreeing with a thought expressed in exotic "thick" terms, even when we have left no questions of our own unanswered? We mostly of course don't inquire what to do and why in distant times and places—in the shoes, say, of a bronze age chieftain. Even with great cultural distance, though, we might sometimes be able to do so, if we tried and if we knew enough about the people whose thinking we contemplate. We might then reject their normative kind concepts as they apply to their circumstances. That is the most likely possibility. But we also might, in some such cases, adopt their concepts as our own—probably not to apply to our own circumstances, but to join them in applying their concepts to their circumstances, agreeing with their claims or disagreeing with them. Whether or not to treat their concepts as applicable to their own choices is a hypothetical question of how to live and why. It is a question of how to live in the contingency of being just like them, with their ethos and in their cultural circumstances.

9. What to Say about the Thing to Do: The Expressivistic Turn and What It Gains Us.
With patterns on the table for various ways that complex, plan-laden concepts might work, my hypothesis is that our normative concepts do work in some such ways. They are plan-laden, but perhaps in ways that are convoluted.

Plan-laden concepts will in many ways mimic concepts that are plainly descriptive. That has been the main point of the book so far. An account of normative thinking is *internally adequate*, we can say, if it gets right all matters internal to adequate normative thinking, such matters as what entails what and what the conceptually coherent possibilities are. Ronald Dworkin asks whether, if an expressivistic account is internally adequate, it could tell us anything we didn't already know. How could it then tell us anything that "moral realists" like Moore, Ross, and Ewing ever denied? If an internally adequate account is right, then the moral realist's mistakes can only be "external"—and no sense can be made, he argues, of purely external claims

about morality. I answer first that my account explains internal features of moral thinking that the moral realist takes as brute unexplained fact. Second, at least one external mistake is intelligible: thinking that an expressivistic account, even if internally adequate, irremediably leaves out part of what we are claiming when we make normative assertions. This claim both Dworkin and I deny, but it is an existential claim and hence intelligible—and other non-naturalists might make this claim once they understood it.

The book assumes throughout that claims about meaning can make sense, but it takes no position on the meaning of "meaning". It allows, for instance, that Robert Brandom might be right in thinking that the concept of meaning is a normative concept, so that with meaning, it's "norms all the way down."

IV. Knowing What to Do

10. Explaining with Plans. We often explain events in normative terms: a general's sagacity might explain his success. The concept of sagacity is normative, and being normative consists, I hypothesize, in being plan-laden. How, though, if normative concepts figure in causal explanations, could they be anything but descriptive and naturalistic? How could they be plan-laden?

The apparatus of hyperdecided states explains why such causal explanations are both plan-laden and intelligible. It explains too why, in a sense, only natural facts causally explain things that happen. We ourselves may have no fully naturalistic way to explain a happening, and yet still accept a plan-laden explanation—and the explanation indeed may be true. A hyperdecided thinker-planner who had everything right, however, would have a purely naturalistic explanation of the event, an explanation distinct from her relevant plans. This is not a matter of levels of explanation: A hyperdecided thinker-planner who got everything right might see the forest along with the trees and have purely naturalistic explanations at every level where a true causal explanation was to be had. Still, her explanations of what happened and why would be distinct from her plans.

When I argue that all properties are "natural", I use the term in the sense that most bears on Moore. My arguments don't rule out extranatural epistemic possibilities such as ghosts, demons, and vibes. These things don't exist, so far as we can tell, but that they don't is another story from the one I tell in this book. If they do exist, they are part of the causal order, and so they don't refute my arguments. The sharp distinction to draw is between kinds of concepts, not properties. The properties signified by plan-laden and naturalistic terms alike are broadly natural. Some concepts, though, are plan-laden, and plan-laden concepts are not purely naturalistic.

11. Knowing What to Do. Can there be such a thing as plan-laden knowledge? Could a person intelligibly be said to know how to live, to know the intrinsic worth of living a certain kind of life? A full answer would require an

account of what 'know' means, and the concept of knowledge has proved philosophically elusive. Roughly, I try saying, knowledge is reliable belief—meaning, belief formed in a way to rely on. To regard a person as knowing is thus to plan to rely on the kind of state the person is in. If such an account of the concept of knowledge works for prosaically factual beliefs, then it should work for plan-laden beliefs: it will allow claims of plan-laden knowledge as intelligible. Such a concept of knowledge is itself plan-laden: If, for instance, I credit you with knowing whether accomplishment has intrinsic worth, then for a case where I myself am uncertain whether it does, I plan to emulate your plans regarding accomplishment.

12. Ideal Response Concepts. The most plausible alternatives to an expressivistic treatment of normative concepts are "ideal response" analyses in the spirit of Firth, Brandt, and Michael Smith. What one ought to do, according to such an analysis, is what one would do in a frame of mind that is *ideal* for decision. (In one variant, the ideal state will be one of full and vivid awareness of everything involved.) A *descriptive* ideal response definition fills in this schema with a naturalistic definition of this term 'ideal'. Such a definition, though, is subject to Moore-like refutation. People may differ, after all, on what frame of mind makes for doing what one really ought to do, and we can ask what's at issue. Their dispute needn't stem from a sheer conceptual mistake; rather, the concept of an ideal frame of mind for decision seems itself to be normative. Hence we might best treat this concept as plan-laden: an *ideal* frame of mind, we can say, means one to trust. To regard a frame of mind as ideal for planning is to plan to trust planning done in that frame of mind. Reading the term 'ideal' in this way eventuates in ideal response definitions that are plan-laden.

Equipped with such a schema for definitions, we can distinguish two plan-laden concepts: (i) being okay to do, and (ii) being an act that one would permit oneself in an ideal frame of mind. The two are in a way equivalent: a thinker-planner can't coherently think that what's okay to do is not what she would now permit herself were she in an ideal frame of mind. Otherwise she accepts what she plans to reject. Applied to the actions of others, however, these two plan-laden qualities can come apart: Xanthippe may coherently think that for the case of being Socrates in prison, the thing to do is to flee, but that the thing for him to *think* the thing to do is to stay put. Unless a person has the right sense of plausibility, she may think, even perfection in all the epistemic virtues (apart from sheer good judgment) may not lead to judgment that's correct. Socrates, she may think, just lacks a reliable sense of plausibility. I consider an absolutist alternative to this treatment of the matter, deferring, in judgments of what an agent ought to do, to judgments that the agent himself would make in ideal conditions. Accepting this, I argue, would be incompatible with a kind of trust one needs in one's own direct

thinking on how to live and why. We still need the concept *okay* as I have characterized it.

Philosophers sometimes speak of "response-dependent properties", but when they do, what's really in play is not a special kind of property but a special kind of concept—the kind we might call *response-invoking*. I propose, in this vein, a plan-laden, response-invoking interpretation of talk of "what we would accept at the end of inquiry".

13. Deep Vindication and Practical Confidence. After spending the book stressing how planning, in its realist-seeming features, parallels science and ordinary belief in prosaic fact, I now take up a way in which the two kinds of inquiry contrast. Take as an example our capacities for vision: Darwin offers a vindication of them; he explains why these capacities would evolve to get matters right. Such a vindication is deep, not in the sense that it proceeds from an Archimedean standpoint independent of sense perception, but that the circle from sense perception to its vindication is non-trivial. It doesn't just take the form, "This is a hand, and behold, vision presents it as a hand." I scrutinize prospects for such a deep, internal vindication of our capacities to discern value, to know what ultimately to plan for—but I find no plausible way to achieve such a vindication. In one sense, then, there is such a thing as knowing how things stand in our surroundings, but not such a thing as knowing what ultimately to aim for. We must settle instead for a cautious practical confidence, for a kind of guarded faith in our capacities to find value in life.

14. Impasse and Dissent. We have no logical guarantee that exhaustive discussion will lead to agreement. Even if we were each fully coherent in our thinking and planning, we might still reach a conversational impasse. I examine the forms that such coherent fundamental impasse might take, beginning with two pure forms: The impasse is *constitutional* if both parties think that each is in ideal conditions for judgment, but they differ in their judgments because of different senses of plausibility—differing basic mental constitutions. It is a *multi-equilibrium* impasse if both parties share a basic mental constitution, but they differ in their views of what conditions are ideal for judgment, and each is in conditions that are ideal by her own lights but not by the other's. Coherent hyperdecided thinker-planners might also differ in ways that combined these two pure forms. As for parties who, though coherent, are not hyperdecided, there will be alternative ways of sharpening their views, and different sharpenings might exemplify different forms of disagreement. In that case, there will be no straight answer to what grounds their disagreement.

When two people differ in their plans for the very same contingency, though, why regard this difference as true disagreement? If nothing in their

respective views could allow for reaching agreement, doesn't this render the dispute illusory? This question is central to the project of the book. A like question can be raised for an individual over time, if he swings with his moods and enthusiasms from plan to plan. For an individual, though, we clearly treat a change of mind as coming to disagree with one's previous self—because, for one thing, planning over time requires this stance. In the interpersonal case, we perhaps have the choice of regarding a difference in plan as a mere difference and not as a disagreement. Alternatively, though, we may "put our heads together" in thinking how to live—meaning that each of us treats the thoughts the other voices as he treats his own recent thoughts, as thoughts to accept or reject. You might, though, find my thoughts not worth heeding in this way. In that case, regarding us as agreeing and disagreeing with each other would still be coherent, but our practice of joint thinking might well, in your eyes, lack a point.

I myself urge us to pursue joint thinking on how to live, putting our heads together to think through questions of what to do and why, proceeding with some tentative faith in each other's judgments. In urging this, I put forward a planning judgment for our joint scrutiny. Our normative language, I suggest, embodies a tentative faith in such an undertaking.