

*On the Hypothetical and Non-Hypothetical in Reasoning about Belief and Action**

I am sure that I do not understand the idea of a reason for acting, and I wonder whether anyone else does either.

Philippa Foot

Perhaps it shouldn't be surprising when a central notion of common sense proves elusive on reflection – that's what makes the philosophical world go round. And some aspects of practical rationality seem obvious enough: rational agents form intentions, adjust means and ends, and so on. Yet even very elementary questions can excite not only controversy between conflicting, entrenched positions, but also expert bafflement.

An indirect approach suggests itself: it might help us to understand reasons for action if we started with reasons for belief. First, action involves belief. Second, one of the most crucial and problematic notions in practical reason – the notion of *non-hypothetical* reasons or requirements (reasons or requirements not dependent upon contingent ends of the agent) – appears to be well domesticated within the literature on theoretical reasons.¹ On the usual view of things, two agents in the same epistemic situation (same evidence, same background beliefs) would have the same reasons for believing any given proposition, regardless of possible differences in their personal goals.²

Can the “usual view of things” in the theoretical realm be given a principled basis? If so, can a similarly non-hypothetical basis be found in the practical realm? In what follows, we will be developing a sequence of arguments that purport to show just this. These arguments turn on considerations concerning belief and action of a kind David Velleman has

called *constitutive*.³ By way of conclusion, we will ask what the limitations of such arguments might be.

1. REASONING ABOUT BELIEF

Let us begin with Gary, a student in our introductory course on the Theory of Knowledge one autumn term. He's confronting epistemology as a discipline for the first time, and he's been staring with silent but disarming intentness from the back of the room for several weeks. Now he's ready to speak: "These philosophers we've been reading seem to agree that there are certain standards of belief, standards we should follow even when they lead us to conclusions we don't like. They spend all their time disagreeing about exactly what these standards are, but they just seem to assume that we'll want to follow them. Suppose I don't? What can they say to me?"

One can imagine our initial response: "Well, you understand why you should eat your vegetables? You may not care about these epistemic standards as such, but you do care a lot about other things. And you're more likely to get what you want if you have warranted beliefs. Following epistemic norms won't guarantee reliability, but there isn't any better alternative short of magic or luck."

Notice, though, that this line of response is *non-epistemic* and *hypothetical*. It advertises the existence of benefits accruing to an agent who follows epistemic norms, but the values or goals in question are not distinctively epistemic, nor do we assume that they carry distinctive epistemic presuppositions. To be sure, this hypothetical justification is, or purports to be, quite *robust*. Virtually any goal Gary might have would be well served by his following epistemic norms. Indeed, we might at this point refer Gary to the Dutch Book argument, to show that he is at risk of being a sure loser if he does not conform his degrees of belief to certain probabilistic principles.

Gary has obviously been preparing his case. He counters by asking the class to imagine an individual dying of an incurable disease, to whom little or nothing matters besides peace of mind. Belief in the Hereafter would comfort him mightily and would come to him spontaneously if he could just relax his epistemic scruples. What makes us think that the balance of non-epistemic considerations will always favor keeping those scruples? Gary pushes his question: he wants to know whether there are any considerations that require or favor following epistemic standards that don't depend at all on our personal goals.

We can, of course, point out that although someone might find the thought of an afterlife reassuring, reassurance is not evidence and so does not yield *epistemic reasons for belief* in the Hereafter, only *practical reasons for being a believer in the Hereafter*. Epistemic evaluation, then, appears to be quite untouched by Gary's deathbed example.

Gary, however, finds this serene lack of regard for whether agents are comforted or tormented by their beliefs off-putting. Moreover, he cannot see how appealing to "epistemic reasons" could provide any sort of answer to his initial question. Isn't it circular to invoke epistemic reasons on behalf of epistemology? The fact that these epistemic reasons are themselves non-hypothetical is beside the point. After all, there are lots of norms that lay down standards that pay no attention to the agent's particular goals.

Gary asks us to consider "anti-epistemology," which tells us to reduce our belief in proportion as evidence increases. This norm is just as non-hypothetical as orthodox epistemology, since it prescribes degrees of belief without making any allowance for the agent's personal goals. Why should we submit ourselves to the old-fashioned rigors of epistemology rather than take on the exciting new challenges of anti-epistemology?

It plainly will not do for us to say to him that epistemic norms recommend non-hypothetically against this, since anti-epistemic norms speak non-hypothetically for it. Of course, it is unlikely that Gary or anyone else would really be prepared to abide by anti-epistemic norms. Only a very singular set of personal goals and circumstances could make anti-epistemic thinking much of a boon. But we can see that this again affords no more than a very robust hypothetical and practical justification.

After all, we reflect, Gary might find himself in some pretty unusual circumstances. Perhaps it is he who lies on the deathbed, yearning for peace of mind. Or perhaps a powerful, mind-reading anti-epistemic demon is prepared to torment him mercilessly unless his beliefs fly in the face of evidence. We consider responding, "Look. Belief isn't voluntary. You can't just decide what to believe." But this threatens to show too much. There won't be much left of normative epistemology unless we recognize some forms of control over what we come to believe. However, there is something that might help explain the oft-repeated phrase that belief isn't voluntary, and that might also help with Gary.

Consider a form of "Moore's paradox" – the extreme oddness of:

- (1) *h* is true, but I don't believe it.

According to anti-epistemology, the more one takes the evidence to favor *h*, the weaker should be one's belief that *h*. In the limit, then, we find the

anti-epistemologist saying:

- (2) I recognize that the evidence for h has become conclusive, so I don't believe that h in the least.

But (2) seems (almost?) as odd as (1). Maybe anti-epistemology isn't a real alternative after all.

What makes (1) – and perhaps by extension (2) as well – so odd? Various explanations have been proposed. One might start by noting that belief is a propositional attitude partly characterized by its representation of its object as true. “Belief is believing true,” the saying goes.

But this is too quick. For even the propositional attitude of “pretending that h ” amounts to “pretending that h is true” – such is the ‘believe’ in ‘make-believe.’ And there is nothing paradoxical about:

- (3) h is true (or: I recognize that the evidence that h is true has become conclusive) but I'm pretending otherwise.

So we must go further. We might say this: a belief that h “aims at” the truth of h .⁴ A belief that h necessarily “misses its target” when h is false, whereas a pretence that h does not. Beliefs are evaluable as true or false, and are false whenever their propositional objects are. To have mastered the distinction between belief and pretense is in part to understand this. That suggests:

- (4) A believer that h holds that, necessarily, her belief that h is false if h itself is false.

This, however, is overloaded conceptually. Most of us think that school-age children have genuine beliefs and can distinguish belief from pretense quite well, even though we suspect that they do not explicitly hold the modal attitude expressed in (4). Rather, they manifest their awareness of the special tie between belief and truth implicitly, by showing sensitivity to the distinction between what is the case (as far as they can tell) and what they would like to be the case, and through their responses to evidence for or against h . Believers in effect hold their beliefs to be *accountable* to truth.⁵

To be sure, it is not false belief as such that is paradoxical. There is nothing odd about:

- (5) h is true, but I wrongly disbelieved it at the time.

Paradox emerges in (1) – or in (2) – not because the belief in question is false or incongruous with the world, but because the belief is incongruous with something else the agent already thinks.

What is the nature of this incongruity, and what sort of problem is it for the believer? There is, I am sure, often incongruity among my beliefs. To the extent that I remain unaware of incongruity, no Moore-like paradox arises. I can, for example, unthinkingly pick up the telephone to call the repair office to report that my phone is dead. This is a state manifesting incongruous beliefs, but not one that seems unattainable. By contrast, the state of mind that would be accurately expressed by (1) – or by (2) – seems not foolish but opaque. What *could* someone who confidently uttered (1) – or (2) – have in mind?⁶

The distinctive propositional attitude of belief is therefore one that not only represents its propositional contents as true, but also one that cannot represent itself as unresponsive to – unaccountable to – their truth. This is still an unacceptably crude formulation.⁷ But fortunately for present purposes we need only a rough idea, since even the rough idea enables us to explain why anti-epistemology is untenable. In order for a propositional attitude to be an attitude of *belief*, it cannot represent itself as wholly unaccountable to truth or evidence.

We've been lost in thought for a bit, but now are in a position to respond to Gary with what looks like a non-hypothetical argument. For we can say why he must, at least in the limit, accord some deference to what he takes to be truth and evidence thereof in his belief-formation. It is part of the *price of admission* to belief as a propositional attitude that one not represent one's attitude as unaccountable to truth. Someone unwilling to pay this price – who, for example, insists that he will represent himself as accepting propositions just as it suits his fancy and without any commitment to their truth – would not succeed in *believing* these propositions at all. The special relation between belief and truth thus comes with the territory of belief, and is not hypothetical upon any contingent aim of the believer.

None of this argues against the possibility of belief that is in fact – even as the outcome of prior design – unresponsive to evidence or truth. One could, it seems, have some success in coming to believe certain convenient falsehoods through a suitable program of self-imposed indoctrination. What this argument purports to show is not the impossibility of such a program but a *design constraint* upon it: if one is to succeed, one must somehow contrive to veil the program's true nature from oneself. Transparent anti-epistemology, for example, is not an option.

Gary might, however, think that we have overstated what has been shown. He can say, "Your answer to my challenge is still hypothetical, as far as I can see. It presupposes that I am in, or plan to enter, the belief business. But what if I opt out? Why can't I just do without belief, and

manage my affairs instead with other propositional attitudes lacking its particular relation to truth?”

Here we might be tempted to reply, “Well, why not do without automobiles and manage instead with boats? Beliefs, after all, play many roles in one’s mental economy – in inference, deliberation, action, even emotion. They are evolutionarily ‘made for’ these roles, and it is by no means obvious how many of these roles could be played by propositional attitudes other than belief.⁸ Just ask whether a strong desire for self-defense, plus a *pretense* that a mortal enemy lurks behind the next hedge, would do the job in producing an action-guiding intention to engage in all-out self-defense.”

Gary is tenacious. “Still hypothetical. You’re telling me that people typically have goals that are better served by having some attitudes that play all the roles of belief – just as people typically have goals that are better served if they have vehicles that can play all the roles of cars, and don’t have only bicycles and boats.” In atypical circumstances, he observes, things might be otherwise. He can remember one incident involving a broken guardrail on the coastal highway south of Monterey when he quite suddenly found himself thinking just how much there is to be said for boats as opposed to cars.

“Very well,” we reply, “you want a non-hypothetical argument and you will have it. But remember: To show that a norm or reason is non-hypothetical is not to show that it is utterly without condition. It is only to show that it would necessarily apply to any agent as such, regardless of her contingent personal ends.

“So. Consider how deeply implicated belief is in our notion of agency. An agent acts on intentions and plans, which constitutively involve beliefs and are formed deliberatively in part on the basis of beliefs. To replace *all* belief with (say) wishing would be to form no intentions at all. Moreover, our notion of ourselves as agents extended over time constitutively involves *memories* and *expectations*. These, too, involve beliefs. There is all the difference in the world between believing that one is the father of John, or believing that one will experience the pains of an unattended-to toothache, and pretending or merely supposing these things. To delete all forms of belief from your mental repertoire would leave you with no recognizable notion of identity.

“Being ‘in the belief business’ therefore isn’t as optional as you imagine. It is a precondition of agency. So the argument is non-hypothetical in a familiar sense: as an agent you must possess beliefs; as a believe you must represent certain of your propositional attitudes as accountable to

truth and as disciplined by truth-orientated norms (at least, in the limit); therefore, as an agent you must so represent at least some of your attitudes, irrespective of what other goals this might or might not serve.”

The argument is not dispositive. But it does place a certain burden on Gary. It seems that he would have to exhibit the compatibility of our notions of practical deliberation, agency, personal identity, etc. with a mental economy that contains no beliefs. The magnitude of this burden affords a *prima facie* case for the following claim: paying the price of admission to belief is necessary to gain entry to agency. A self-representation of certain of one’s attitudes as “aiming at” truth is *partially constitutive* of belief, which in turn is *partially constitutive* of agency. Let us, then, call this sort of argument a *constitutive argument*.

Unlike our first, hypothetical, eat-your-vegetables defense of conforming to norms of theoretical reason, this constitutive argument concerns not an agent’s actual conformity (or attempt to conform) to epistemic norms but her self-representation as such. Moreover, it concerns only a limiting case, the case of deeming certain evidence to be conclusive. Not very much normative epistemology can be wrung from that. Finally, Gary might surprise us and successfully discharge his burden of proof by showing that a genuinely alternative propositional attitude – or constellation of such attitudes – could play as many of the roles of belief as one would need to attain agency.

Rather than explore these issues further at this point, let us simply note that despite its limitations, the constitutive argument provides a *prima facie* case for the non-hypothetical status of certain broad epistemic requirements.

II. REASONING ABOUT ACTION

It is now the spring term. We find Gary, undaunted as ever, in our Introduction to the Theory of Action. This time he sits in the front row; and has his question ready earlier in the term. “These philosophers,” he begins, “each has his own view about what practical reasoning requires. But what makes any of these views something I have to pay attention to? How could any of them insist that I pay attention to their favorite norms if I didn’t care to? They might not be *my* favorite norms.” In particular, he concludes, he’d like to know how the practical case compares with the discussion of theoretical reasoning last term.

Let us see how much parallelism we can find. In the autumn we began with a non-epistemic but robust hypothetical defense of familiar standards

of theoretical reason: Gary could expect to do better relative to almost any of the goals he might have if he formed warranted beliefs. In effect, this constituted a *practical* defense of theoretical reason. Would a parallel in the present case be a practical defense of paying heed to norms of practical reason? This time that would seem circular at the outset.

There is, however, at least one way of construing the question that would avoid circularity – though this might not satisfy Gary. We can distinguish two ways in which an agent's deliberations, decisions, and actions might be said to follow a norm: objectively versus subjectively.

We will say that an agent's deliberations, decisions, and actions are *in objective conformity with* a norm to the extent that he is actually succeeding in complying with whatever the norm prescribes. Consider a norm that directs one to act so as to maximize one's self-interest. An agent would be in (full) objective conformity with this norm just in case his acts were those, relative to available alternatives, that would maximally benefit him.

It would not follow, however, that this agent is acting exclusively for his own sake, or even with his own benefit in mind. That is, it would not follow that this agent's deliberation and actions are *subjectively patterned on* a norm of maximizing self-interest.⁹

Using this distinction, we can interpret the question whether there might be a practical defense for paying heed to norms of practical reasoning in the following, non-circular way. Given any particular norm of practical reason, one can ask whether in a particular instance – or in general – subjectively patterning one's deliberation on this norm would constitute behavior in objective conformity with it. For many norms much of the time, the most promising way of achieving objective conformity will indeed be subjectively to pattern one's deliberation on the norm. But not always. It seems plausible, for example, that if one were to regulate one's conduct by self-consciously and exclusively consulting one's self-interest one would be incapable of the sorts of commitment to other individuals or to groups or causes that are the source of some of life's deeper satisfactions.

If Gary wants to know whether, in his circumstances – or in general – one would have good (objective) practical reason to be (subjectively) practically rational, then he will be asking a genuine question to which the answer is not preordained. For many theories of practical reason, it is a contingent matter.

But Gary is growing restless. He meant his question to be less internal to the domain of theories of practical reason. "Objective or subjective," he says, "it matters little to me how you put it. I want to know whether

you can give a good, non-circular defense of having anything to do with norms of practical reason. Why should I be bothered?”

This is starting to look like an impossible request. For either the defense presupposes a set of norms of practical reason, which would appear to be question begging, or the defense makes no such presupposition, and there is nowhere to start.

However, something more might be said within the domain of practical reasoning. Philosophers often engage in a process that looks like a non-circular practical defense of practical norms – a process that sometimes leads them well away from their starting points.

The process begins with our many intuitive notions about which actions or principles of action make – or do not make – sense. These notions are much less coherent or articulate than a “theory of practical reason.” Here, for example, is an appeal to intuition meant to raise doubts about whether maximizing one’s expected utility really makes sense even as a basis for self-interested choice:

You are forced to play Russian roulette – but you can buy your way out. One bullet is placed in a six-cylinder revolver . . . What is the most you would pay to have the bullet removed?

Next . . . You are forced to play Russian roulette with four bullets in the revolver. Answer a new question: What is the most you would pay . . . to have *one* of the four bullets removed, leaving three? More? Or less?¹⁰

Most say “less.” Maximizing expected utility seems to say “more.” This has the form of the “Allais Paradox,” though one might say that it is no paradox at all, but merely a counterintuitive result. For our purposes what is most important is that the example does not appear to depend for its force upon people’s acceptance of some alternative theory of practical rationality. Indeed, popular intuitions have proved exceedingly difficult to describe in any general, coherent way.

In consequence, one can also have the opposite response to the case: after being told that the choice “more” would maximize expected utility, one might come to think that one *should* accept this answer as rational and simply ignore its intuitive oddness. Indeed, a considerable literature in cognitive psychology suggests that commonsense reasoning is prone to various errors and fallacies in assessing probabilities and risks.¹¹

Such dialogue – an interplay of examples, intuitive responses, empirical theory, and proposed norms – is a form of *wide reflective equilibrium*. The critique and acceptance of norms of deliberation through such dialogue is

surely a form of reasoning that deserves the name ‘practical.’ Yet reflective equilibrium arguments need not take a particular set of practical norms for granted. Inarticulate commonsense notions of practical rationality will figure in such a process and partially shape it, but they need not serve as a *constraint* on the process, or leave it intact.

A wide reflective equilibrium argument thus might answer to the demand for a non-circular practical justification, but would it be convincing to Gary? Convincing or not, he’d point out, the argument would nonetheless be *hypothetical*: which adjustments we are prepared to make in response to various intuitive tensions is sure to depend upon our particular goals and priorities.

Therefore Gary still wants to know whether there exists a non-hypothetical argument concerning fidelity to norms of practical reasoning, akin to the constitutive argument made in the case of theoretical reasoning. One established way of approaching this question is to ask whether there is anything in the realm of action that plays the role that truth plays in the realm of belief. We can think of ourselves as looking for a feature *F* such that one must represent oneself as “aiming at” *F* in action in approximately the same sense in which one must represent oneself as “aiming at” truth in belief.

Among philosophical accounts of the nature of intentional action that might be seen as offering a candidate for *F*, two have perhaps attracted the greatest interest historically. For each, I will argue, it is possible to construct a constitutive argument. Let us call the two philosophical accounts of agency we will be considering *High Brow* and *Low Brow*.

High Brow is a view with excellent pedigree, tracing its ancestry back to ancient Greece. According to High Brow, just as belief necessarily “aims at” the True, action necessarily “aims at” the Good. Deliberation seeks to identify the good, and action is guided by it. In choosing an action we place it (or find it to be) in a positive evaluative light, and deem it choiceworthy.

Note that the High Brow’s claims concern action as such, not merely *rational* action. The constitutive argument for belief held that a self-presentation as “aiming at” the truth is part of what makes a propositional attitude be one of belief, rational or irrational alike. Similarly, the High Brow claims that action that is irrational – as distinct from arational or non-rational behavior such as reflexes or kicking in one’s sleep – is also in some sense “aimed at” the good. Weakness of the will, as it is ordinarily understood, is a paradigm of practical irrationality that clearly manifests

this. The akratic agent is said to be “aiming at” the good but falling short due (say) to insufficient motivation or “willpower.”

Of course, our description of High Brow is quite vague. Action requires representing one’s choice in a positive evaluative light, but which? There are many varieties of goodness: good for oneself, good for one’s kith and kin, morally good, aesthetically good, and so on. One might formulate High Brow by identifying one of these goods, or perhaps a *summum bonum*, as the true end of all action. This would, however, be needlessly ambitious for our purposes. We seek a constitutive feature of action as uncontroversial upon reflection as the connection between belief and truth, and a generic value claim is much less controversial. Action, our High Brow will say, involves representing the act chosen or the ends for which it is done as good *in some self-acknowledged sense*.

Despite its generic character, this claim is non-trivial. It implies, among other things, that individuals incapable of representing an end or a course of conduct as good – nonhuman animals, or (perhaps) human infants – would also be incapable of agency, properly so-called. Moreover, those among our fellow adults who have a latent capacity to represent a course of conduct as good, but who fail to develop or exercise it – that is, who fail or refuse to *acknowledge* any good – would also lack agency, properly so-called. Their lot might be a kind of motivation-driven behavior that nonetheless remained in some profound sense aimless. Alternatively – and perhaps far more likely – individuals claiming not to acknowledge any good are actually kidding themselves. Their deliberation and action reveal their nihilism to be no more than a posture.

This suggests a High Brow response to Gary. Suppose that a visiting High Brow philosopher has just given a guest lecture in our class on practical reason. Gary’s eager hand is up. Why, he asks, must he pay any attention to the good when deciding what to do? What magical force would stop him from simply ignoring questions of good and bad, or flying in their face, and acting as he pleases?

The High Brow philosopher can reply: “I have claimed that deliberation and choice constitutively involve representing what you choose as in some sense good. You can no more decide to ignore questions of good and bad – of choiceworthiness – in your deliberation and action than you can decide to introspect someone else’s thoughts rather than your own. No ‘magical force’ is needed to police this constraint. Of course, you might lack the nerve, or will, or energy to follow through on your judgments of choice-worthiness. But to the extent that you are aware of this, you yourself will be sensible of it *as a lack*, a gap between what you value

and what you do. You will be in no position to say ‘Well, that’s nothing to me.’ If it weren’t *something* to you, it wouldn’t have been a choice in the first place.”¹²

Gary is being attributed paradoxical claims, which might be thought to echo – though in a practical setting – Moore:

- (6) I believe I have reason to choose act *A*, but I can’t see anything good about it.
- (7) Act *A* would be good, but that’s no reason for *me* to choose it.

One way of explaining the oddness of (6) and (7), according to the view under consideration, is that they seem to suggest the existence of a gap within practical deliberation that makes room for a purely *hypothetical* dependence of an agent’s deliberation on judgments about the good. This gap presumably would need to be filled by the agent’s possession of some independent, intermediating goal, such as that of “doing what is good.” But such a view would be deeply confused, according to the High Brow. The deliberative role of judgments of goodness is indispensable and needs no mediation – it simply comes with the territory of acting. To use our previous phrase, it is *partially constitutive* of agency that one perceive the landscape in an evaluative light, and steer toward the good as one sees it. Gary proposed to act in the face of, or indifferent to, questions of goodness. But anyone who managed to become a complete stranger to goodness would simply have dismantled his capacity for deliberate action and begun a life of merely behaving, of roaming at the behest of his appetites.

We thus have arrived at an argument in the practical realm that affords an interesting parallel to the argument made in the theoretical realm: both appeal to a constitutive condition to identify a non-hypothetical element in reasoning.

High Brow is, however, highbrow. Many philosophers, in my experience, are not. Gary, it seems, isn’t either. He points to a long line of Low Brow thinkers, beginning perhaps with Hume, who have denied that an agent engaged in deliberate action necessarily “aims at” the good.

If not the good, then what does action “aim at”? To Hume is often ascribed the view that agency aims at no more than the satisfaction of current desires. But this might be wrong about Hume, and it in any event is not necessary for a Low Brow. Just as High Brow comes in both generic and brand-name forms – the latter identify a particular sort of good as “the aim” of action – so does Low Brow. We can see Hume as endorsing in the first instance a generic Low Brow position: agents necessarily possess

and act on *ends*, and this involves both a *representational* and a *motivational* component, though neither component need involve a judgment of, or an “aiming at,” what is good. The belief/desire view often attributed to Hume is an example of this type. According to it, motivationally inert beliefs will suffice for representation, and non-evaluative, intrinsic attractions and aversions will suffice for motivation. Hume famously wrote:

Ask a man *why he uses exercise*; he will answer *because he desires to keep his health*. If you then enquire *why he desires to keep his health*, he will readily reply *because sickness is painful*. If you push your enquiries further and desire a reason *why he hates pain*, it is impossible he can ever give any. This is an ultimate end, and is never referred to any other object.¹³

Humean individuals engage in both theoretical and practical reasoning. They inquire into causes and effects; form beliefs about the conduciveness of means to ends; take into account the relative strength and independence of desires; acquire habits; form intentions to act; and formulate and respond to rules and sanctions. Their conduct therefore can, it is claimed, be given fully fledged intentional, rational-agent explanations. Why-questions about their conduct can often be answered correctly by citing *their reasons* for behaving as they do, and these will include: how they represented the situation, what their goals were, how they weighed their various ends, how they adjusted means and ends, and so on.

Our interest in Humean individuals lies precisely with the claim that they exemplify agency even though they do not by their nature “aim at” the good. We need not evaluate the stronger claim that reason giving always terminates in current intrinsic desires (or that belief must be motivationally inert).

The tenability of any Low Brow position therefore depends upon the possibility of distinguishing *the possession of ends* from *the making of judgments that certain ends are good*. Desire appears at first to afford a clear case: we often speak of acting on desire (so desire seems capable of playing the necessary role in choice) and also of desiring something that we do not take to be good (so desiring seems suitably distinct from evaluating).

But such claims might be challenged. To desire, it can be argued, is to represent as desirable, and desirability is itself a species of good. When we speak of desiring that which we take to be bad, this can be understood as (say) reflecting the difference between a *prima facie* and a conclusionary judgment of value. This is, however, a very demanding position. It would force us to deny that young children, who (it seems) lack the evaluative

concept of desirability, have desires. To reject the evaluative notion of desire need not be to treat desire as a mere animal appetite (whatever that might be). There is a great deal of psychic distance between a fish that swims to the surface “because it is hungry” and a child who responds to our question “Why did you come downstairs?” with the answer “I’m hungry.” We can begin to account for this difference by pointing out that the child “acted on a desire” in a way that the fish did not, and this despite our reluctance to suppose that the child has judged there to be something further that is good about having breakfast.¹⁴

Perhaps a more promising challenge to the Low Brow’s distinction would be to argue that even if desires (say) are non-evaluative, they cannot function as *ends* until some suitable evaluative judgment has at least tacitly been made. This view is less severe. A young child can be spoken of as having desires in the familiar sense, but even when these “mere desires” cause her behavior, a child cannot be seen as *acting on reasons* or as *possessing an end that furnishes her rationale in acting* until she is capable of exercising a certain amount of judgment as to the appropriateness or worthwhileness of acting on her desires.

Perhaps the best Low Brow defense is to give illustrative examples. First, consider a case of a kind brought to our attention by Jean Hampton.¹⁵ Our two children have been begging all week to go to the shore. Both, however, dislike long summer car rides. When the weekend comes one child absolutely refuses to get into the car. “But we’re going to the beach, which you love!” “I don’t care. I don’t want to ride in this stuffy old car. I hate it! I won’t do it.” He has to be carried bodily to the car and buckled in, thrashing. Once in the car, he still refuses to be jollied along. “It’s *your* fault I’m in this stuffy old car! I told you I hate it.” The second child confines her thrashing to loud complaints. “Not another car ride! Last time I felt sick the whole time!” But when the time comes to leave she climbs into her seat of her own accord, waiting sulkily to be buckled in. On her face is a look that says “Okay, I’ll ride in the car, but don’t expect me to like it.”

The second child possesses a capacity for self-control (relative to her weightier desires) that the first child lacks, though, according to the Low Brow, we need not also impute to the second child a judgment that being at the beach is a good thing – beyond her strong desire for it. There is a sense in which the second child’s thinking and conduct, but not the first’s, accord the strong desire to be at the beach the force of a *rationale* for the despised beach trip. The desire speaks on behalf of means toward its fulfillment, even unwanted means. By contrast, for the first child the

desire to be at the beach does not yet function as end-setting, and his conduct is not rationalized even by his own desires.

Second, consider an example inspired by a case due to Michael Smith.¹⁶ Two “unwilling addicts” – individuals who strongly desire heroin but who also very much wish they did not¹⁷ – are both beginning their day. Each has overslept – it is now too late even to consider going to work. “I’ve got to quit this stuff. It’s ruining my life. I won’t even have a job by the end of the week – if I haven’t been fired already.” This is no new resolve. Each has already judged his taking of heroin to be a bad thing on the whole. Though neither reconsiders this judgment, as the day grows longer the desire for heroin becomes fierce. By noon, each has set out to get a fix. One locates a needle and, trembling, injects himself. The other, who is just as aware as the first of how to use heroin, locates a needle and hungrily tries to eat it. He chokes to death.

We now ask why each has used the needle as he did. For the first, we are able to cite a reason, *his* reason: he strongly wanted to dose himself with heroin, and he knew that this is how to do it. For the second, we are at a loss. Without further information we must see his needle-eating conduct as inexplicable by any “rational agent” explanation. Perhaps sheer craving somehow overcame him. The difference in intelligibility between the two cases is not attributable to the addicts’ differing capacities to form and be guided by judgments of the good – they formed the same judgment on this score, and equally failed to be guided by it – but to their differing ability to adjust means to ends. Thus although we almost certainly regard the conduct of neither as truly *rational*, we do see the one’s conduct as having a *rationale* in terms of what he believes and desires that the other’s does not.

These two examples give us at least a *prima facie* case for the Low Brow’s distinguishing of the notion of an agent having and acting on ends from the High Brow notion of an agent aiming at the good.¹⁸ We now must ask whether Low Brow conceptions of agency – which ordinarily are seen as allergic to anything non-hypothetical – can support a constitutive argument of their own. Perhaps so. Consider the Moore-like statement:

(8) *E* is an end of mine, but that’s nothing to me in my deliberation.¹⁹

Our original Moore-ism

(I) *h* is true, but I don’t believe it

is a statement that could easily be true (there are many truths I do not believe) but that seemed deeply problematic for any agent to assert.

Something similar holds for (8). It is hardly odd for someone to fail in a given case to take one of his ends into account. He might not even notice its relevance. But asserting (8) would be peculiar indeed, according to the Low Brow, since we have no clear idea what it could amount to for *E* to be acknowledged by me as an end of mine if it counted for nothing in my deliberation whether or not *E* is realized. Of course, we must make room for inattention, distraction, and depression. The connection suggested in (8) is non-hypothetical not in the sense that it has no conditions, but in the sense that it does not presuppose a further, contingent desire on my part “that I realize my ends” or “that I realize end *E*.”

It should be emphasized that (8)’s oddness manifests a *structural* connection, which, though non-hypothetical, is not a device for generating univalent, non-hypothetical imperatives. Thus, if an end *E* of mine would be advanced by act *A*, this can be taken as either counting in favor of performing *A* or counting against retaining *E*. If success in the army requires unquestioning obedience, for example, I might consider giving up my military ambitions.²⁰ The oddness of (8) points to the unavailability of a third option: genuinely retaining the end while in effect setting oneself to accord it no deliberative relevance. To see one’s deliberation as guided (at least in part, in the limit, other things equal, in normal circumstances, etc.) by one’s own ends thus comes along with the mere possession of ends.

Suppose, then, that a Low Brow philosopher visits our class. At the lecture’s end, Gary raises his characteristic challenge. “Maybe I can’t have an end unless I take that to count in some way in thinking about how to act. Fair enough. But that’s still hypothetical. You yourself admit that very young children might have desires or appetites but no ends as such. Maybe they know something you don’t.”

The Low Brow philosopher can respond. “You, like most of us, have ends, desires, appetites. Nothing prevents you from becoming a being with appetites and desires but no ends. There are lots of such beings around: infants, maybe animals. You could join their ranks. But then you would cease *acting* on desire – you’d merely be behaving. If you are to remain an agent, you must have ends. And once you acknowledge ends – as you’ve conceded – you must (in the limit, under ordinary conditions, etc.) be to some degree engaged in the business of weighing courses of conduct in light of their tendency to contribute to the realization of your ends.” So, we reach a “principled basis” for a Low Brow non-hypothetical response to Gary. An agent as such must in effect see herself as deliberating in a way that gives weight (in the limit, etc.) to the realization of what she takes

to be her ends, independently of what these particular ends might be.²¹ This is so even for agents who are acting *irrationally* relative to their ends.

III. STOCKTAKING

The High road and the Low road thus both lead to non-hypothetical requirements for practical reasoning. The path in each case proceeds using a constitutive argument that has much in common with the constitutive argument made for theoretical reasoning. In all three cases a linkage is made to the (alleged) nature of agency, thereby avoiding dependence upon contingent personal goals.

But have we found convincing answers to Gary's questions? Or convincing grounds for rejecting them? To simplify exposition, I will begin by narrowing the argumentative field, focusing largely on the case of theoretical reasoning and the Low Brow version of the practical case.²²

Constitutive arguments have the strength that comes from purportedly necessary connections. And necessity is hard to argue with, even for Gary. But this strength can also be a weakness. If the necessity turns out to be linguistic, the argument may lack the power to sustain substantive conclusions. And if the necessity is of a more substantive kind, then the argument may have the unintended effect of pulling the claws of the very criticisms one wishes to make.²³ We now face both of these dangers. Let us look at them in turn.

First, the linguistic danger. Consider the Low Brow constitutive argument that connects taking oneself to have an end *E* with taking oneself to be responsive in deliberation to whether *E* is realized. Someone might see this as an analytic truth: "That's just what it *means* for *E* to be an end of yours – an end is something you see yourself as giving weight to in deliberation." Gary, however, sought answers to seemingly substantive practical and epistemic questions: "Why do things that way?", he wanted to know. It would be surprising if we could give an answer with nothing more than a few definitions. To be genuinely responsive to the concerns expressed, constitutive arguments must capture a substantive – not merely linguistic – necessity.

This brings us to the second danger, the danger of pulling the claws of criticism.²⁴ Assume, for example, that the connection between taking oneself to have an end and according that end deliberative weight is a substantively necessary, non-analytic connection of the same modality as the connection between being gold and having atomic number 79. What would we then be able to say by way of *criticism* of an agent who refused

to give deliberative weight to his own acknowledged end *E*? Would he be “necessarily deliberatively defective” or perhaps “self-defeatingly irrational”?

If the constitutive argument is right, we cannot even raise the question! To fail to take oneself as according *E* deliberative weight is to fail to acknowledge *E* as an end. But then the agent cannot merit the label ‘self-defeating’ or ‘irrational’ with respect to *E*. An analogy: to discover that the metal in the sample tray on one’s laboratory bench has atomic number 82 is not to discover that it is “defective gold,” but rather that it is not gold at all.

A similar problem confronts all constitutive arguments. Suppose, for example, that someone has a propositional attitude toward *p* that involves, among other things, her representing *p* as true. Thus far, this attitude is a candidate for belief. But suppose further that she sees no relevance to this attitude of admitted evidence against *p*, even evidence she recognizes to be conclusive. When challenged, she is not defensive and produces no elaborate rationale, but simply points out that she is quite indifferent as to whether her attitude toward *p* is responsive to the truth of *p*.

According to the constitutive argument, she does not have an irrational or epistemically defective belief that *p*; she simply fails to *believe* that *p* at all. Perhaps she instead is *supposing* that *p*. What if she nonetheless claims that her attitude toward *p* is one of belief? It would seem that, on the strength of the constitutive argument, our only criticism could be that she has *mislabeled* her propositional attitude (like our mislabeling of the metal sample). Labeling errors are not, however, defects of rationality. Once she has found the right word for her propositional attitude, the criticism would vanish.

To be sure, we could at this point invoke a more general, higher-order constitutive argument. If we were to come across someone who failed quite generally to deliberate in a way that he takes to be responsive to his ends, or who failed quite generally to form propositional attitudes that he takes to be responsive to evidence, we could argue that such an individual thereby would fail to possess ends or beliefs at all, and thus would fail to be an agent.

Indeed, we might raise the stakes still higher. Perhaps *speaking a natural language* itself presupposes the formation of beliefs and intentions, so that an individual without beliefs or intentions could not even offer an *argument* on behalf of his way of life – his vocalizations would not constitute speech. This is beginning to sound serious! Or is it? Now

when he emits the sounds *bǐ-lēf'* or *ǎk'-shən* we cannot even charge him with a linguistic mistake.

It seems that we are turning up the volume of criticism while simultaneously ensuring that the purported target of our criticism is ever more profoundly deaf. If we rely on these ascending constitutive arguments, we quickly reach a point in which the only thing left to say of someone is to dismiss him as not one of us. This is xenophobia, not criticism.

Eager for a secure justification, a knock-down answer to the likes of Gary, we sought a requirement – a “must” – that applies non-hypothetically, arising from the very conditions of agency. That now looks unwise. For then there could be no such thing as failure to conform *on the part of an agent*. Perhaps we have asked too much, or the wrong thing, of our constitutive claims.

We may begin to regroup by recognizing that we have formulated the constitutive arguments too rigidly. Having beliefs and having ends are, even in the limit, complex phenomena to which we have not done justice. Beliefs, for example, come in degrees, and are not all or nothing. Moreover, having a belief involves possessing a large bundle of dispositions – not only to represent one’s thinking in certain ways, but also to infer, to notice, to act, to avow, to assert, to claim a measure of authority, and so on. Psychological realism alone compels us to recognize that many of the attitudes in ourselves and others that we unhesitatingly call beliefs may from time to time lack one or another of the complex bundle of attitudes and dispositions paradigmatically associated with belief as an ideal type. Interpretative charity often demands that we be latitudinarian with respect to departures from the ideal type.

Now consider a person who has a propositional attitude toward *p* that he deems to be belief, but that he does not – or does not with any consistency – hold accountable to admitted evidence concerning the truth of *p*. If that person nonetheless allows this attitude to play all the other roles of belief – in assertion, in intention-formation, in expectation, etc. – then he will almost certainly find himself in a variety of difficulties, difficulties more serious than mislabeling, difficulties that mere relabeling could not remove. Given his attitude’s extensive overlap with characteristic roles of belief, we would have some interpretative justification for calling it a ‘belief’; but given its unresponsiveness to admitted evidence, we would also have some interpretative ground for calling it a ‘belief *manqué*,’ or even a “rationally defective” or “irrational” belief.

Why “rationally defective”? Consider an example. Suppose that I am a nervous flyer. I recognize there to be compelling statistical evidence that commercial airplane travel is very safe. Moreover, my frequency of electing to travel by air fits with what one would expect of an individual who deems it safe; for example, a small difference in travel time or cost will tip me toward air rather than rail or car. Yet I find that nonetheless, in the sense in which belief is connected with expectation and perception, I do not seem really to *believe* that air travel is relatively safe. This is not because my attitude is a mere supposition or pretense or of the like – it has less in common with these attitudes than it does with belief, as my travel choices show. One might with some justice interpret me as partly believing that taking wing in a commercial airliner is safe, and partly disbelieving this.

But the division between belief and disbelief here is not a simple probability distribution, the way that I “distribute” my belief over ‘Clinton will win in November’ and ‘Clinton will not.’ For in a suitably abstract context, I will sincerely and confidently assert the view or place a bet that flying on a commercial airliner is much safer than driving to work. In a different context, when I’m aboard a jet taxiing for takeoff, I may find myself irresistibly believing that I am in a very precarious situation, wishing I were anywhere else, jumping to conclusions about the meaning of small sounds or little bumps and jiggles, and so on. By way of contrast, despite what I know of statistics, I have no such belief when speeding through traffic in my rattletrap of a car, clutching a cup of coffee between my knees, late for work (again!). How would I myself describe things? I would probably say that my beliefs on such subjects as the relative safety of air versus car travel, the safety of the particular flights or car trips on which I find myself, and so forth are simply not wholly rational.²⁵

We need to effect a similar relaxation of the Low Brow constitutive argument in the case of practical reasoning. When deliberating about what to do, a rational person takes her ends into account. But to have an end paradigmatically involves possession of a complex bundle of attitudes and dispositions, involving action, perception, sentiment, belief, and so on. As in the case of belief, interpretative charity will often license attributing an end to a person even though some of these elements are missing, or inconsistently present. An agent who arranges a considerable part of her life in order to promote a long-term goal will from time to time find herself in contexts in which she is attracted to other things, feels alternative pressures exclusively, or even lacks interest in her life. In such cases, even though she may see the bearing of her long-term goal, and even though she may remain disposed to avow it, she may nonetheless

find herself giving it no weight in certain deliberations.²⁶ Do we say she no longer has the end? Or that she has the end but isn't at the moment being fully rational with respect to it? The agent herself – at least, if she is like me – will sometimes opt for the latter description.

Failures of rationality come in many shapes and sizes, and do not form a unified type. But it may be useful to think of some forms of theoretical or practical irrationality as instances of incomplete-yet-nearly-complete approximation of believing or having an end. How many of the elements in the bundle must one possess to be “nearly complete”? There are limits, but vague limits, no doubt. And dynamic, holistic limits – they concern chunks of one's thought and stretches of time. Particular elements may come and go, but may do so in mutually compensating ways or so that, at any given time, enough hold.

We thus remove from the constitutive arguments an artificial rigidity. But are we any closer to an answer to Gary? We may have made matters worse. If having a belief or possessing an end is a complex phenomenon, with vague and holistic limits, then we have left behind the Manichaean world of the original argument. In that world there seemed to be only stark choices: to be an agent or . . . infantile, or a beast. Now it seems one could pass almost imperceptibly from belief to near belief, and thus from agency to something else. And that something else therefore might not be so terribly alien.²⁷

Gary asked why he should pay attention to epistemic norms. If we reply that this is necessary in order to be a believer and thus to be an agent, he can respond: “But just how severe a cost does this threaten me with? Somewhere on the continuum between the ideal type of belief, on the one end, and clear non-belief on the other, there is a region that forms the borderland of genuine belief. I want to know why my attitudes should be on one side rather than the other of that borderland. The claim that I would cease to be an agent on one side of the region sounds dramatic. But if life on the believer side of the borderland has certain pluses and minuses, how do we know in advance that the balance must be worse on the other side? Mightn't it even be better, on the whole?”

Consider two possibilities. First, suppose that there is more that Gary would find enjoyable or valuable on the believer side of the borderland. Then we have a reply to Gary, but it once more looks hypothetical – whether Gary finds life sweet or sour seems unlikely to be independent of what he happens to desire. Second, suppose the opposite: there is more that Gary would find enjoyable or valuable on the other side. Then our reply to Gary can only be, “The enjoyment, however great, would not

be that of an agent. The value would not be the value of the life of an agent.”²⁸ This reply is indeed non-hypothetical, but Gary could be excused if he finds it unconvincing.

We might at this point be inclined to be dismissive. Surely once one has demonstrated that a condition is essential to agency one has justification enough. After all, we *are* agents and that seems to be a very deep fact about us. Justification has to start somewhere, and if it is to be justification *for us* it had better start where we are. Indeed, the mere fact of Gary’s asking such a question, posed as matter for choice, seems to presuppose that he, too, is an agent. Yet all the same can be said for two much more difficult-to-dismiss questions.

Consider first a patient with a painful, incurable disease who wonders whether to elect to end his life by euthanasia rather than live out the disease’s wretched course, destroying his family’s finances and becoming every day less the sort of person he has aspired to be. He is an agent, and moreover his very posing of the matter as a question of choice presupposes his agency. To choose euthanasia would, however, be to put an end to his agency. Does this suffice to show that there can be no question of justified voluntary euthanasia? Can we say that, since life is a necessary condition of agency, “choice of euthanasia” is ruled out as a practical contradiction? Most of us, I suspect, do not think so. The considerations on the side of ending his life, and thereby resigning agency, might be more compelling than the considerations on the side of continuing it. But then we can understand the idea of a rationally justified transition from agency to non-agency.

Second, imagine a Schellingesque case of a kind discussed by Derek Parfit.²⁹ You have been captured by mobsters. They seek revenge on members of your family who have testified against them and who now have new identities and locations thanks to the Federal Witness Protection Program. The mobsters will torture you to reveal your family members’ whereabouts. You know that you will not be able to resist this torture. If you could abolish your agency by knocking yourself senseless – perhaps irrevocably – your captors could not extract from you the information they need. Does the fact that you would be crossing the borderland into non-agency show that such a choice could not be a genuine or appropriate option for you as an agent?

Dramas aside, non-agency need not be the end of life as we know it. We all pass from non-agency to agency sometime during the first years of life, and, arguably, we all commute daily back and forth to a state in which agency is at least temporarily disengaged when we sleep and awaken.

Suppose we were to say to Gary: “The pleasures (or other advantages) on the other side of the borderland couldn’t possibly count *for you*. They wouldn’t be *yours*. You’re an agent – that’s one of the deepest facts about you – and they’d be the pleasures of a non-agent.” This way of speaking is belied by our comfort with speaking of a life that stretches from birth to death as the life of a single person, despite its various transitions to and from non-agency, active versus suspended agency, and the like. We cannot simply refuse questions of partial or even complete border crossing.

Previously we spoke of the “price of admission” to belief, action, or agency. Now we are discussing the “exit price,” temporary or permanent. Often that price will be high. Now is the time to remind ourselves, and Gary, that most of us most of the time will be in a much better position to figure out and accomplish what matters to us if we are agents. But a high-priced option is very different from an impossibility. High prices are sometimes worth paying, and circumstances (such as facing the end of one’s life) can drive down the price. We arrive, then, at a somewhat unanticipated sense in which one might intelligibly ask for “reasons for action” – reasons for agency versus non-agency as a way of being.

This very observation does show, however, that there is a sense in which Gary has not succeeded – even slightly – in suggesting anything like the possibility of an alternative to familiar forms of practical reasoning. When he asks, in effect, whether the exit price is worth paying, he is asking whether *being an agent* is the best or only way of getting what he most wants from life. This is itself a means/ends form of reasoning of the familiar, Low Brow sort. It therefore betrays Gary’s deference to (at least) Low Brow notions of agency after all. When Gary contemplates a border crossing and asks whether life on the other side might be an improvement, he is giving deliberative weight to the tendency of a means to promote an end.

Perhaps the grass really is greener on the other side. But then not only would Gary have good reason in the familiar sense for crossing over, he would also have good reason *to be* and *to stay* on the other side even after he ceased himself to be a fully fledged agent. Indeed, unless Gary takes for granted means/ends reasoning, it is unclear what bearing the (possibly) high quality of life on the other side of the borderland would have upon what he should do.

We might be able to put this point more clearly by invoking another turn-of-the-century Englishman’s paradox, not G. E. Moore’s this time,

but Lewis Carroll's.³⁰ Achilles entertains an argument:

- (9) If p then q
- (10) p
- (C) So: q .

Carroll's Tortoise asks Achilles whether there isn't a gap in this argument, a missing premiss. Couldn't one grant both premisses but fail to be driven to the conclusion unless one also granted:

- (11) If $[(\text{if } p \text{ then } q) \ \& \ p]$ then q

to effect the connection between (9) and (10) and (C)?

This seems reasonable to Achilles, on whom it only slowly dawns that he has just launched a regress. For suppose our premisses now enlarged to be (9)–(11). Tortoise will cheerfully argue that we would need a new premiss to effect *their* relevance to the conclusion, namely:

- (12) If $\{ [(\text{if } p \text{ then } q) \ \& \ p] \text{ then } q \} \ \& \ (\text{if } p \text{ then } q) \ \& \ p \}$ then q .

Were (12) added, Tortoise would notice the need for yet another premiss to link (9)–(12) with the conclusion (C). And so on.

The moral: one cannot treat rules of inference (such as *modus ponens*) as premisses, on pain of regress. Put another way (and using Carroll's own terminology): We cannot see rules of inference in logical argument as *hypotheticals*. This is not to say that we should see them as *necessary* or *non-hypothetical premisses*. Far from it. Taking premiss (12) to be a necessary truth – or as “constitutive of logical inference” – would no more enable it to stop the regress than taking it to be simply true. Rules of inference differ essentially in *role* from premisses, not in modality.

Somewhat similarly, we should not be led by questions such as Gary's to think of the mutual bearing of ends upon means as itself hypothetical, or as something like a premiss in our deliberation about action, on pain of regress. For suppose we started with the practical argument:

- (13) E is an end of mine
- (14) Means M would secure E
- (C*) So: There is that much to be said deliberatively in favor of my doing M , or against my having E .

And suppose Gary asked, “Isn't this argument missing something? – Doesn't it suppose not only that I have end E , but that I also have the further aim, call it F , of choosing so as to bring about the realization of my ends? If I didn't have that further end, couldn't I reject any relevance

of (13) and (14) to my deliberation?” It would seem that we need to add this premiss:

- (15) *F* [= choosing so as to bring about the realization of my ends] is an end of mine.

But if one did not already recognize that having an end makes deliberatively relevant questions about the means that would advance it – if, that is, (13)–(14) were insufficient to support the conclusion (C*) – then adding the further premiss (15) could hardly help. And notice that the situation would not be improved by claiming that the end *F* is somehow necessary for agents as such. For if one cannot see the bearing of having an end upon the choice of actions, then knowing an end to be necessary would not enlighten one on that score.

So we arrive again at what we want to say to Gary. Not: “Giving deliberative weight to one’s ends is constitutive of agency, and you are, after all, an agent.” He might sensibly wonder whether *that* state of affairs should continue. Rather, we want to say: “You already defer, in posing this question, to the very thing you seek to challenge. You must already see – and feel – the ‘practical logic’ of what you claim to find arbitrary or problematic: the bearing of ends upon means. If you reply, ‘Well, so that’s just another end of mine – I can change it’ then we can answer ‘No, on pain of regress, it cannot be just another end of yours.’”

Does this show that Gary cannot be raising a genuine issue about whether or not to be rational? Return to our previous distinction between subjective and objective notions of conformity to norms.

Gary was dismissive of this distinction, since his ambition was to ask a question less “internal” to the theory of practical reason in its orthodox form. “Subjective or objective,” he said in effect, “I want to know whether rationality’s worth it.” It now appears that he succeeded instead in asking a more “internal” question. Roughly: “I want to know whether subjectively patterning my thinking along means/ends lines would really be in objective conformity with realizing my ends – especially, the end of living well.” This question, more internal than he imagined, is also more real than others have imagined. Constitutive arguments of the kind considered here can neither answer it nor set it aside. The Low Brow argument, if successful, would show that subjective patterning to means/ends reasoning – that is, representing oneself as deliberatively adjusting means and ends – is partially constitutive of agency. But whether a life of subjective patterning would be in objective accord with realizing one’s most important ends is another question. Arguably, it is Gary’s real question.

Gary is asking, not “Do ends bear on means?”, but “Why be the sort of creature who asks about ends and means?”

Lewis Carroll’s paradox is sometimes used to argue that there cannot really be an “alternative logic” – we cannot drop or pick up rules of inference like premisses. But this paradox cannot really establish that orthodox logic will be adequate to – or necessary for – the fullest possible development of our thought and experience. The paradox shows instead that a certain way of thinking about how an alternative logic might be introduced or argued for is absurd. We cannot say, for example, “Just compare the implications of existing logical rules with those of my new rules. . . .” The very notion of *implication* presupposes that logical rules are already in place. Similarly, the present discussion cannot demonstrate that Low Brow means/ends reasoning will be adequate to – or necessary for – the fullest possible development of our thought and experience. We can at most show the absurdity of attempting to give a Low Brow rationale (in terms of objective conformity) for questioning whether Low Brow reasons (again, in the objective sense) are relevant for what to do or how to live. Gary seemed to be pursuing such a rationale, so his line of questioning can come to look a bit silly. One can’t lift oneself by one’s own bootstraps, but he seems to have managed to pull himself down thereby.

Yet Gary isn’t without a “less internal” response. He can, without absurdity, be seen as trying to find a way of keeping us aware that no one really knows where reflective equilibrium and our evolving experience might take us. Each step in a reflective equilibrium process is linked by intelligible forms of reasoning to the step before, but this does not mean that we could not take steps that would, in sum, yield the result that our conception of reasoning itself has changed. Experience has held some interesting surprises for those who thought certain principles – such as the Principle of Sufficient Reason – were constitutive of the entire possible domain of thought and action. Perhaps Gary just wants to remind us of this. Moreover, he could insist, it seems inevitable that wherever reflective equilibrium takes us, its route will depend upon facts about us and our contingent nature. As Gary never tires of pointing out, this shows our reasoning is never on a wholly non-hypothetical footing.

Where, then, are we left? We began with the question whether a non-hypothetical account could be given of why we must conform to certain forms of theoretical or practical reasoning. In that context, we developed several constitutive arguments that showed some prospect of returning an intelligible, positive answer. We now have also seen the

limitations of such arguments. Especially, they cannot supply a self-sufficient non-hypothetical response.³¹

But we should not imagine that this means we are left with wholly hypothetical considerations. Another conclusion to be drawn from the Low Brow constitutive argument is that each element – hypothetical and non-hypothetical, end-setting and ends/means-adjusting – has its own distinctive role to play in reasoning about action. Neither can do the other's job. To ask, of a given act of Low Brow reasoning, whether it owes its conclusion to hypothetical or non-hypothetical considerations is a bit like asking, of an act of deductive theoretical reasoning, whether it owes the belief-guiding force of its conclusion to the reasoner's hypothetical deference to the premisses or her non-hypothetical deference to the rules of logic.³² The answer, of course, is always both.³³

NOTES

Ethics and Practical Reason, Garrett Cullity and Berys Gaut, eds. (Oxford: Clarendon Press, 1997).

- * I would like to dedicate this essay to the memory of Jean Hampton, who taught us much about practical rationality.
- 1. This will function herein as something like a stipulation about the meaning of 'non-hypothetical.' Thus, reasons that depend upon ends *necessary* for agents as such would count as non-hypothetical for present purposes. One might distinguish *personal-goal* non-hypotheticalness (absence of dependence upon contingently held *ends* of the agent) from *personal-belief* non-hypotheticalness (absence of dependence upon contingently held *beliefs* of the agent). What I refer to in the text as 'well domesticated within the literature on theoretical reasons' is the idea that epistemic reasons are *personal-goal* non-hypothetical.
- 2. I write 'personal goals' because there is a school within contemporary epistemology according to which theoretical reason is *end-orientated*. But the ends in question typically are assumed to belong to a special class of *epistemic* ends that are subject to at most limited variation across rational individuals.
- 3. See David Velleman, "The Guise of the Good," *Noûs* 26 (1992), 3–26 and "The Possibility of Practical Reason," *Ethics* 106 (1996), 694–726. I am much indebted to Velleman's discussions and to his comments on earlier papers, though I do not mean to suggest that he would agree with my use of the notion. See also his "Deciding How to Decide," in G. Cullity and B. Gaut, eds., *Ethics and Practical Reason* (Oxford: Clarendon, 1997).
- 4. Bernard Williams gives a seminal discussion of belief as "aiming at" truth in "Deciding to Believe," in his *Problems of the Self* (Cambridge: Cambridge University Press, 1973) 136–51.
- 5. More precisely, a believer-that-*h* holds *this* attitude accountable to the truth of *h*. It cannot be essential to belief (in beings with finite minds, such as us) to hold that, for all *p*, if *p* is true then one should believe it.
- 6. See Section III for some further discussion.

7. For some further discussion, see P. Railton, "Truth, Reason, and the Regulation of Belief," *Philosophical Issues* 5 (1994), 71–94.
8. Those who think of *being a belief* as a purely functional property will presumably hold that any attitude playing all the roles of belief would simply *be* a belief. Here, however, we are supposing that the attitude would lack at least some of the central roles of belief, namely those involved in the "internal relation" between what one believes and what one takes to be true or evidential.
9. Subjective patterning on a norm *N* need not involve a second-order thought to the effect that "I do this in order to satisfy norm *N*."
10. The example itself is used by Daniel Kahneman and Amos Tversky, "Prospect Theory: An Analysis of Decision under Risk," *Econometrica* 47 (1979), 263–91, at 283, and attributed to Richard Zeckhauser. The version given here is quoted from Allan Gibbard, *Wise Choices, Apt Feelings* (Cambridge: Harvard University Press, 1990) 15.
11. See for example Daniel Kahneman, Paul Slovic, and Amos Tversky, eds., *Judgment Under Uncertainty: Heuristics and Biases* (New York: Cambridge University Press, 1982).
12. Sometimes, of course, we just *pick* rather than choose. This can be seen, however, as a species of action in which it is decided simply to select an option rather than deliberate further.
13. David Hume, *An Enquiry Concerning the Principles of Morals*, app. 1, "Concerning Moral Sentiment."
14. We have been working throughout with generic High Brow views because of their greater plausibility. That complicates the present discussion, however, since open-endedness about the notion of the good can make it difficult to distinguish such views from Humean Low Brow views. Unless a certain amount of substance is built into the idea of goodness, it will be rather too easy (and uninformative) to think of any sort of desiring as "deeming to be good."
15. I am grateful to Shelly Kagan for bringing this sort of example to my attention.
16. Personal communication. He is obviously not to be held responsible either for the claims made about the example or for the interpretation offered of it.
17. Psychologists would, I think, challenge the suggestion that the cravings of an addict can be understood as a subspecies of our familiar notion of desire. Let us, however, follow philosophical convention and set that concern aside.
18. The Low Brow need not rule out the possibility of an agent also inquiring into or aiming at the good. His point is simply that this is at most an *option* for agents, and perhaps also that it presupposes the means/ends relationships that are the stuff of Low Brow agency.
19. In this formula, *E* must be understood as occurrently taken by me as an end of mine. As with belief, there is nothing odd about the diachronic case:

 (8') I once deemed *E* to be an end of mine, but that now counts for nothing in my deliberation.
20. This is like the relevance of *modus ponens* to inference. If I already believe that (if *p* then *q*) and come to believe that *p*, should I conclude *q*? Perhaps, in light of *q*'s implausibility, I should question one or both of the premisses.

21. It should perhaps be emphasized again that an agent's ends need not be self-oriented. They could include the well-being of others, moral or aesthetic causes, and so on.
22. This might not be wholly unfair to High Brow theories. After all, High Brows can agree that means/ends reasoning is a central part of agency (even, of course, irrational agency).
23. For a discussion of related concerns about the critical limitations of constitutive or linguistically necessary principles, though in connection with a conception of instrumental *rationality* rather than agency, see Christine Korsgaard, "The Normativity of Instrumental Reason," in Cullity and Gaut, eds., *Ethics and Practical Reason*.
24. I am indebted here to Shelly Kagan.
25. In cases like this, we might seem to be flirting with Moore's paradox. That is, I might be tempted to say: "Yes, I grant that it is unquestionably true that commercial air travel is safer than car travel, but I don't really believe it." Once we realize all that "coming to believe" actually involves, this looks more like a needlessly paradoxical way of expressing a fairly familiar sort of imperfection in rational belief.
26. This sort of case differs from classic cases of weakness of the will, in which the agent feels the positive deliberative force of an end, but is swayed to act otherwise. The case I am imagining is less psychologically conflictual.
27. Bas van Fraassen recommends that the community of scientists take an attitude of *acceptance* rather than belief toward the truth of their theories, but should *also* behave as scientists in every other respect – inferential, experimental, etc. – as if they believed the theories to be true. If sustainable, this attitude would be a form of near belief. Could it, perhaps, even spread beyond scientific theories? See *The Scientific Image* (Oxford: Clarendon, 1980). For critical discussion, see P. Railton, "Truth, Reason, and the Regulation of Belief."
28. Note that we cannot uncontroversially say "The value would not be *true* value." We will see shortly some examples where ends of (what we agents deem) true value can be attained only at some cost to one's own agency.
29. Derek Parfit, *Reasons and Persons* (Oxford: Clarendon, 1984) 12–13.
30. See Lewis Carroll, "What the Tortoise said to Achilles," *Mind* 4 (1895), 278–80. I have altered Carroll's example slightly.
31. We have not, however, tried to demonstrate the impossibility of other types of argument that could provide the grounds for "purely" non-hypothetical justifications.
32. I am grateful to John Searle for suggesting an error in the original version of this remark. Searle's concern, however, appears to be with the *logical validity* of a deduction, rather than the phenomenon of *theoretical reasoning* as such. See John R. Searle, *Rationality in Action* (Cambridge: MIT Press, 2001) 19–20.
33. I am grateful to a number of people for helpful comments and conversation. In particular, I should mention Garrett Cullity, Stephen Darwall, Berys Gaut, Allan Gibbard, Shelly Kagan, Michael Smith, David Velleman, and an anonymous referee.