

Moderate Idealization and Information Acquisition Responsibilities

JASON TYNDAL

Tulane University

Abstract

I argue that advocates of moderate epistemic idealization need some standards against which they can determine whether a particular individual P has a responsibility to acquire some specific piece of information α . Such a specification is necessary for the purpose of determining whether a reason R, the recognition of which depends on accounting for α , can legitimately be ascribed to P. To this end, I propose an initial sketch of a criterion that may be helpful in illuminating the conditions in which P might be said to have a responsibility to take α into account when searching for reasons to act or believe. I worry that reason ascriptions that are based on information acquisition expectations which are not captured by this criterion (and, hence, not legitimate reason ascriptions) may be open to charges of authoritarianism, especially when such a reason ascription prompts one person to demand that another act in a way that is responsive to the ascribed reason(s).

1. INTRODUCTION

It is a common practice in moral and political life to ascribe a reason to an individual (whether implicitly or explicitly) and then proceed to treat that individual as if he or she in fact *has* the ascribed reason. This happens in a wide range of cases, from everyday cases of social morality, where one person demands that another act in a certain way, to debates in political philosophy (e.g., debates about the public justification of political power). It is often argued, though sometimes simply assumed, that the ascribing party is justified in attributing one or more reasons to the individual in question – even if that individual does not recognize or endorse those reasons. One well-trodden way to justify such an ascription of reasons is to appeal to the notion of idealization. To engage in idealization (specifically, epistemic idealization) is to consider how a person would act if that person were deliberating under ideal circumstances.¹ Such circumstances are usually taken to entail proper reasoning and the possession of adequate

¹ I use the term “act” in its broadest sense. I take it to include non-physical acts such as changing what one believes or values.

information. As such, epistemic idealization has both a rationality component and an information component.²

We can distinguish generally between accounts of *full* idealization and accounts of *moderate* idealization. Full idealization entails construing idealized agents as flawless deliberators and as possessors of perfect information. Moderate idealization entails construing idealized agents as less than fully idealized. My purpose in this paper is to articulate a problem which confronts any account of moderate idealization that permits idealization with respect to the information component.³ The specific problem that concerns me is this: advocates of moderate idealization need some standards against which they can determine whether a particular individual P has a responsibility to acquire some specific piece of information α . Such a specification is necessary for the purpose of determining whether a reason R, the recognition of which depends on accounting for α , can legitimately be ascribed to P. As I shall argue, the practice of ascribing reasons to agents is a practice that often involves the importing of unwarranted expectations regarding the information acquisition responsibilities of those agents.

² This point is nicely expressed by Kevin Vallier in his discussion of what he calls the “standard conception of idealization” (2014, p. 145).

³ I acknowledge that an account of moderate idealization might significantly restrain what it means to idealize the information component. For example, one might argue that the information that one obtains by merely deliberating or drawing inferences from one’s existing information base counts as new information. Such information is just a by-product of further “internal” deliberation. That is, such idealization of the information component is just a by-product of idealization of the rationality component. However, when I speak of moderate idealization throughout this paper, I have in mind only those accounts which permit the idealization of an agent’s knowledge base through supplementation of information which is “external” to the agent – i.e., information which cannot be acquired by the agent through purely deliberative means, such as simply drawing inferences from her existing information base.

Yet, sorting out what these responsibilities ought to be raises numerous difficulties which do not confront accounts of full idealization.

I begin in §2 by briefly discussing Gerald Gaus' recent account of moderate idealization, which exemplifies the sort of account that is vulnerable to the problem of information acquisition responsibilities that I raise here. While I do not have the space here to sufficiently argue that full idealization should be rejected in favor of moderate idealization (though I believe that to be the case), I will provide a brief explanation for why one might be inclined to reject full idealization. In §3, I explain the problem of information acquisition responsibilities that troubles accounts of moderate idealization. This problem arises out of a need to clarify why some information (which we do not currently possess, but which we have access to) is sometimes factored into determinations about what reasons we can be said to have while other information (which is also not possessed, but still accessible) is not. As I see it, disagreements about the ascription of one or more reasons to an agent are often disagreements about what information should ground the ascription. Such disagreements indicate a need for standards regarding the information acquisition responsibilities of agents. Towards this end, I identify in §4 a series of considerations which may help to address the problem raised in §3 and I provide an initial sketch of a criterion which may prove helpful when attempting to determine whether a given agent has a responsibility to acquire some particular piece of information. I worry that reason ascriptions that are based on information acquisition expectations which are not captured by the criterion (and, hence, not legitimate reason ascriptions) may be open to charges of authoritarianism, especially

when such a reason ascription prompts one person to demand that another act in a way that is responsive to the (illegitimately) ascribed reason(s).⁴

2. MODERATE IDEALIZATION

It is crucial to clarify the distinction between *having a reason* and *there being a reason*.⁵ To say that P *has* reason R is not just to say that P acknowledges R (although that may be the case), but to say that R can legitimately be ascribed to P independently of whether P acknowledges R.⁶

This of course means that any plausible account of moderate idealization must explain when it is proper to ascribe a reason to an individual in cases where the individual does not acknowledge that reason. Moreover, as will be elaborated below, advocates of moderate idealization stress that there may be more reasons that apply to P's circumstances (reasons that there are) than can legitimately be ascribed to P. In light of the foregoing, we can distinguish between (1) reasons that are recognized by P; (2) reasons that may be *legitimately* ascribed to P (i.e., reasons that we can say P has); and (3) reasons that *apply* to P (which include reasons that may be legitimately ascribed to P as well as reasons that may not be legitimately ascribed to P).

⁴ As I use the term here, to act in an 'authoritarian' manner is to demand that a person adhere to some principle, policy, or guiding belief despite the fact that the person in question does not and cannot rationally endorse whatever principle, policy, or guiding belief is in question.

⁵ Some deny that such a distinction exists. For example, Maria Alvarez, in *Kinds of Reasons*, asserts that "I take the claim that there is a reason for A to ϕ to be equivalent to the claim that A has a reason to ϕ " (2010, p. 22). However, once the distinction between full and moderate idealization is considered, it becomes apparent that the two phrases are not equivalent.

⁶ We of course need an account of what it means to legitimately ascribe a reason to an individual. Such an account will be given in Gaussian terms in §3.

Gerald Gaus (2011) has presented an attractive version of an account of moderate idealization (see also Vallier 2014). Gaus rejects the notion of full idealization on the grounds that if idealized P is too far removed from actual P in term of access to information and adeptness at reasoning, then idealized P will affirm a set of reasons that we cannot reasonably expect actual P to arrive at.⁷ One worry, then, for advocates of moderate idealization may be this: when we attribute non-reasonably-accessible reasons to an agent in an attempt to justify some action to them, we fail to justify that action *to the affected agent*. Instead, we only justify the action to some far-removed idealized counterpart of the affected agent – a counterpart which does not suffer from the cognitive limitations that frustrate the reasoning of non-idealized agents.⁸ At any rate, those that reject full idealization in favor of moderate idealization typically do so on the grounds that we cannot expect actual P to have perfect information or to reason

⁷ As part of his argument against full rationality, Gaus argues that, contrary to popular belief, full rationality does not guarantee convergence between fully rational persons, nor does it guarantee determinacy regarding what any one individual has reason to do (2011, pp. 239-243).

⁸ It is not uncommon for theorists working in value theory and epistemology to discourage over-idealization (or at least full idealization). Robert Pasnau, in a discussion of ideal theory and epistemology, has said that “[i]dealized epistemology... does not involve abstracting away from all human cognitive limitations, but instead focuses on what would count as perfection for beings such as us, in a world such as ours” (2013, p. 1005). Russell Hardin has argued that moral theorists have failed to adequately appreciate how “fundamentally important” the fact of human cognitive limitations is to the construction of moral theories (1988, p. 1, 26). With respect to the notion of full idealization, Hardin does not mince words: “any argument that turns on perfect information, perfect calculations, and perfect theory is a house of cards” (1988, p. 17). Christopher Cherniak, in *Minimal Rationality*, argues that “the pervasively and tacitly assumed conception of rationality in philosophy is so idealized that it cannot apply in an interesting way to actual human beings” (1986, p. 5).

flawlessly at all times.⁹ Yet, it is still conceded that if we give no weight to what idealized P would affirm, then actual P is more likely to affirm reasons which he or she does not in fact have. To see this, consider the thesis that P's affirming R is both necessary and sufficient for R to be a reason of P's. One problem with such a thesis is that people can be mistaken about the reasons that they take themselves to have (perhaps due to false beliefs or faulty reasoning).¹⁰ For example, consider Bernard Williams' famous gin/petrol case: "[t]he agent believes that this stuff is gin, when it is in fact petrol. He wants a gin and tonic. Has he reason, or a reason, to mix this stuff with tonic and drink it?" (Williams 1981, p. 102) As Williams remarks, "it is just very odd to say that he has a reason to drink this stuff, and natural to say that he has no reason to drink it, although he thinks that he has" (1981, p. 102). Ultimately, Williams does not think we should say that the person in question *has* a reason to mix the petrol with tonic and drink it.¹¹ Yet, adopting such a position amounts to rejecting a thoroughly subjective account of reasons for action.¹² If it is right to reject such a view, and hence deny that agents in these sorts of situations always have the reasons that they take themselves to have, then at least some idealization is warranted when talking about what reasons a person can be said to have.¹³ The challenge for

⁹ For an extended critique of full idealization, see chapter 5 of Vallier 2014.

¹⁰ As Gaus frames it, "if people have crazy beliefs, they do not have reasons to act as the crazy beliefs indicate, though they may take themselves to have such reasons" (2011, p. 236).

¹¹ While the false belief might help *explain* the action, it would not *justify* it.

¹² Or what Gaus refers to as "a *purely* subjective understanding of one's reasons..." (2011, p. 236, emphasis added)

¹³ Even if one does not share Williams' intuition regarding the gin/petrol case, the need for some idealization is motivated by the idea that there are at least some situations in which agents are mistaken (sometimes deeply) about the reasons they think they have.

Gaus and other advocates of moderate idealization, then, is to show how P can be idealized without being fully idealized.

In determining how much idealization of an agent is called for, Gaus starts with actual agents who exhibit “real rationality,” and then seeks to specify what an idealized version of those agents would look like without losing sight of the human cognitive limitations which actual agents possess.¹⁴ According to Gaus, “[a] rational agent is one who is competent at following a set of norms about how to go about reasoning...” (2011, pp. 244-245). That is, rational persons are capable of adhering to the rules of good reasoning.¹⁵ On this account, a *good* reason for P will be any reason that P arrives at by adhering to the rules of good reasoning. But it does not follow that a good reason, R, counts as a *sufficient* reason. For R to be sufficient (justified) it must (1) be arrived at by adhering to the rules of good reasoning; and (2) not be defeated by a roughly “equally accessible” reason (Gaus 2011, p. 246).

But how much reasoning is enough? In particular, if R_1 is a good reason, how much more reasoning can I be expected to do – in search of a defeater – before R_1 will count as sufficient? As Gaus puts it, P has “a sufficient reason R if and only if a ‘respectable amount’ of good reasoning by [P] would conclude that R is an undefeated reason (to act or believe)” (2011, p.

¹⁴ For more on the notion of “real rationality,” see Pollock (2006, pp. 3-6). With regard to the notion of human cognitive limitations, Pollock notes that “[h]uman beings, and any real cognitive agents, are subject to cognitive resource constraints. They have limited reasoning power, in the form of limited computational capacity and limited computational speed. This makes it impossible, for example, for them to survey all of the logical consequences of their beliefs, or to compare infinitely many alternatives. This is a fundamental computational fact about real agents in the real world...” (2006, p. 3) For a more elaborate discussion of some of these limitations, see chapters two and three of Cherniak (1986) and Simon (1976).

¹⁵ Gaus takes these rules to be universal, though he does remark that they may be more easily retrieved via some languages rather than others (2011, p. 244).

250). It is important for Gaus that actual P is *capable* of deliberating at the level of idealized P. As such, the reasons that P can be said to have will only be those reasons that are accessible to her. Accessibility, for Gaus, is construed in terms of the relationship between P and idealized P. Specifically, P can be said to *have* a reason – that is, P can be legitimately ascribed a reason – if and only if (moderately) idealized P would affirm that reason to actual P.

One feature of Gaus' view that some may find troubling is that it will entail an inherent level of vagueness with respect to what counts as a respectable amount of deliberation. As Gaus admits, determinations about whether this standard is met in particular instances will be “often vague and always contextual” (2011, p. 254).¹⁶ I do not take this feature to be a weakness of Gaus' theory,¹⁷ for non-arbitrary determinations about what counts as the *right* amount of deliberation – i.e., the precise point at which one has deliberated adequately – may not be possible in the absence of full idealization. If this is right, and if there is a convincing case for rejecting full idealization, then the vagueness inherent in accounts of moderate idealization (at least with respect to the element of deliberation expectations) appears to be unavoidable. That said, Gaus' method of addressing the issue of deliberation (i.e., the rationality component) brings to life a new problem when the issue of information acquisition expectations is brought into the discussion. This is the problem that I will discuss next.

¹⁶ Pollock has argued that since reasoning is generally non-terminating (i.e., “[t]here will never be a point at which the agent has completed all the reasoning that could possibly be relevant to a decision”) the most we can expect from a “real agent” is that “the agent perform a ‘respectable amount’ of reasoning, and then base [his or her] choice on that” (2006, pp. 4-5).

¹⁷ So long as the norms of good reasoning can be determined (an issue I will not take up here).

3. EXPECTATIONS AND INFORMATION ACQUISITION

When Gaus discusses how P should be idealized, he focuses on adeptness at reasoning. Yet, whenever he discusses a change in an agent's moral judgment, what fundamentally drives that change is not more reasoning, but rather the acquisition of more information. This is precisely what happens in Gaus' variation on J.S. Mill's bridge case. Gaus imagines a person who

does not have a reason to cross the bridge [which appears unsafe]... If, however, he was a civil engineer and had just concluded a study showing the bridge was much safer than it seemed, we will retract our claim.

Although there is a sound deliberative route to the conclusion not to cross, there is a more warranted conclusion *to* cross of which we are now aware.

(2011, p. 252).

What grounds the change in the reasons the agent has is the acquisition of new information – in the form of facts about the bridge. This is not to deny the role of deliberation. An engineer certainly has to be able to draw accurate conclusions from the available data. However, as I will argue shortly, it is important that we carefully distinguish between the elements of deliberation and information acquisition when engaging in the practice of idealization.

The issue is that some reasons may be accessible simply in virtue of engaging in more deliberation, while other reasons may be accessible in virtue of receipt of new information (which likely also entails additional reasoning). It is not until late in Gaus' discussion of reasons that we get a more explicit glimpse of the role of information with regard to meeting the standard of a respectable amount of deliberation.

The reasons you have must be accessible to you, and as a real rational agent in a world in which cognitive activity has significant costs, rationality does not demand one keep on with the quest to discover less and less accessible reasons. I have stressed that this is by no means a debilitating limit on the reasons we can acknowledge, and indeed in many ways we move surprisingly close to a highly idealized rationality, for expert advice and the growth of social knowledge allows increasingly sophisticated and complex conclusions to be accessible as reasons to all with simply an adequate amount of deliberation. Think about all the reasons to believe and act that one has after twenty minutes on WebMD (2011, p. 253).

While Gaus appears more concerned with the act of deliberation, it is evident that what feeds that deliberation is the acquisition of new information. Additionally, in light of Gaus' discussion of the bridge case, we see that P can be said to have a reason if, when engaging "in a reasonable amount of rational deliberation *with the available information*," he or she would arrive at that reason (Gaus 2011, p. 250). This brings us to the crux of the issue: what information is to count as available? More important, what information is moderately idealized P expected to take into account? Unfortunately, Gaus does not clarify what information an agent is expected to take into account in searching for reasons to act or believe.¹⁸

¹⁸ In *Justificatory Liberalism*, Gaus advocates for a modest form of open justification in which "we treat [one's system of beliefs and reasons] as open to new information and arguments and, from this external perspective, make judgments about what would then be justified in [that system of beliefs and reasons]... The core idea of open justification is that, at any given time, a justified belief system is, ideally, stable in the face of acute and sustained criticism by others and of new information" (1996, p. 31). Achieving stability in one's belief system means coming to terms (probably slowly) with the reasons that *there are*. But P non-judgmentally informing Q of information α ,

We can conjure up numerous scenarios in which P would have a reason to Φ if only she had more information – e.g., if she knew the actual time, or knew the current circumstances her father was in, or simply read from some book sitting beside her, or looked at the current weather report, or finally got around to that unread mail. These are the sorts of situations in which people are inclined to ascribe a reason to P even if P *merely has access* to some information which might have led P to recognize that reason.¹⁹ But why might one be inclined to ascribe a reason to P in such situations (despite P merely having access to that reason), but not, for example, in a case where access to a reason requires more time and energy on the part of P – such as P needing to attend some recent public lecture on estate planning, or to be well-versed in electromagnetic theory, in order to have reason R?

The point is that, in deciding what reasons a person can be said to have, people often import expectations about what information the person in question has a responsibility to investigate. The challenge is to explain why some information (which we do not currently possess, but which is accessible) should be factored into determinations about what reasons a person can be said to have while other information (which is also not possessed, but still accessible) should not. It is not enough to say that idealized P should take into account any

which will help to stabilize Q's system of beliefs, is a distinct matter from P demanding Q to act in a particular way – or judging Q to have acted wrongly – because P believes that Q *ought* to have taken α into account. It is the latter issue that I am most concerned with. That said, an agent's system of beliefs and reasons may generate certain information acquisition responsibilities that would indeed be captured by Gaus' account of open justification. But, as will hopefully become clear below, not all information acquisition responsibilities can be generated by one's own system of beliefs and reasons. Some information acquisition responsibilities are foundational – they are meant to help form one's system of beliefs and reasons. If this is right, the system itself cannot fully set one's information acquisition responsibilities.

¹⁹ The idea here is that the reason itself was accessible in virtue of the necessary information being accessible.

information which he or she has access to. In the current technological age, most of us have access to far more information than we can possibly process. Just consider all of the information available to a person who has internet access (whether at home, work, school, or a public library). It is vastly more information than that which was available to the average person just a few decades ago. As such, the view at hand would move us towards an account of full idealization. So, the issue of what information should factor into any given reason ascription is not simply a matter of distinguishing between information which is accessible and that which is inaccessible.

With regard to information acquisition expectations (as opposed to deliberation expectations), it is unsatisfactory to say that one must take into account a “respectable amount” of information. For one thing, the issue is not so much about *how much* information we take into account, but *what* information we take into account. If someone says I have a reason to Φ based on the content of my unchecked voicemail, then it will not suffice to tell them that the fact that I have spent six hours on Wikipedia and did not find a reason to Φ constitutes a justification for my not Φ -ing (despite the fact that I took into account a lot of information during those six hours). When we talk about *what* information as opposed to *how much* information is to be taken into account, we leave little room for the notion of vagueness. If a precise piece of information should be factored into a reason ascription, then we can point directly to that piece of information and, hence, remove the vagueness that features inherently in the issue of deliberation expectations.²⁰

²⁰ With regard to the issue of deliberation, it is unclear to what extent it might be possible to point to particular reasoning processes in the same way that we might point to a particular piece of information. In some situations, we might be able to assert “If only P had completed inference I, then P would have recognized reason R.” It is difficult, however, to know how widespread such assertions might be. Moreover, it is unclear exactly what sorts of inferences

Disagreements about the ascription of one or more reasons to an agent are often disagreements about what information should ground the ascription and not how much information is being considered. This is why the elements of deliberation and information acquisition need to be treated separately in any moderately-idealized account of how we come to have reasons. Our standard of what counts as a reasonable in terms of information acquisition cannot be subsumed under our standard of what counts as a reasonable amount of deliberation. As such, an explanation is needed for why some information (which we do not currently possess, but which we have access to) should be factored into determinations about what reasons a person can be said to have while other information (which is also not possessed, but still accessible) is not. This is the part of the story that we do not get from Gaus. But it is a part that we need given the emphasis he places on ensuring that idealized P is not too far removed from actual P in terms of adeptness at reasoning and information-processing capacities.

As I have noted, one feature of accounts of moderate idealization is that they entail vagueness with respect to determinations about what counts as an adequate amount of deliberation. This vagueness appears to be unavoidable (and acceptable) insofar as a non-arbitrary determination about what counts as the *right* amount of deliberation is impossible.²¹

we might be able to expect others to make. As Cherniak argues, very little of any particular ordering of deductive inferences from “easy” to “difficult” “can be shown a priori to apply to all agents” (1986, p. 31). Considerations like this suggest that the most imposing obstacle to eliminating the vagueness issue from the deliberation component is developing a set of standards (along the lines of those offered below in §4.6 with respect to the information component) by which a third party might be able to non-arbitrarily assess moderately idealized P’s reasoning in a particular situation.

²¹ This is not to argue that non-arbitrary determinations about what counts as the right amount of deliberation are in principle impossible. It is only to recognize the challenges that confront such a view. Importantly, if it is possible to

However, such vagueness is clearly problematic with regard to the issue of information acquisition expectations. What possibilities, then, are there for dealing with this problem? That is the question I take up in §4.

4. ON THE PRACTICE OF IDEALIZATION

It is often not clear what idealization entails. It may entail good reasoning skills, knowledge of relevant facts, a requirement of impartiality, the ability to recognize and eliminate neurotic or poorly-generated urges and desires, the ability to identify the various moral considerations that pertain to a particular circumstance, and perhaps others. While it is not within the scope of this paper to question the legitimacy of the general practice of idealization, I do intend to discuss one common aspect of the practice – the inclusion of a requirement of knowledge of relevant facts. Such a requirement demarcates a domain of responsibility with respect to information acquisition.

It may be thought that one way to avoid full idealization, with respect to the issue of information acquisition, is to emphasize that agents are only expected to take into account *relevant* information. What makes a piece of information relevant to one's pending decision is debatable, but as a first gloss we might say that information α is relevant to pending decision D just in case α is necessary to discover the reasons *there are* that would impact D. For example, with regard to Gaus' bridge case, P should not have to acquire information about, say, effective

remove the vagueness that I take to be inherent in the deliberation component, this does not undermine my claim that the elements of deliberation and information acquisition need to be treated separately in developing an account of moderate idealization. After all, the standards by which we would judge whether a person has deliberated responsibly will diverge from the standards by which we would judge whether he or she has acquired information in a responsible manner.

goat herding practices in order to decide if she has a reason to cross the bridge. Such information is irrelevant to her pending decision. As such, this qualification prevents the idealization theorist from construing idealized agents as having an exhaustive information set – a set that is certainly inaccessible to non-idealized agents. Nevertheless, one still over-idealizes when idealized agents are construed as possessing knowledge of all relevant facts. This is because we cannot always reasonably require an agent to consider all of the relevant facts. It is often the case that some relevant facts will be inaccessible to P (or, more properly, reasonably inaccessible). If, in Gaus’ bridge case, P needed to be a civil engineer in order to possess all of the relevant facts, then those facts are reasonably inaccessible to P. Admittedly, it remains true that such facts are not in principle inaccessible to P. After all, P could enroll in a university, obtain a civil engineering degree, and then return to the bridge to decide if she has a reason to cross.²² Most people, I assume, would agree that this is too much to ask of P. In light of these considerations, we should say that a condition of thinking that P has a responsibility to acquire some piece of information α is that α be both relevant and reasonably accessible.

I do not mean to imply that P should always – in actuality – achieve perfect synchronicity with idealized P’s set of recognized reasons. I would still expect P to often fall short of idealized P in terms of the reasons that are recognized. But where P falls short, the reasons that she fails to

²² One commentator on this paper has suggested that I have gone too far in considering a case in which P might be expected to obtain a civil engineering degree since P could simply hire a civil engineer in order to obtain all of the relevant facts. This comment illustrates the very issue at hand: *why* think it is reasonable to expect that P hire a civil engineer in order to retrieve all of the relevant facts? While I do not dispute that doing so would likely be more cost and time efficient for P, I rely on the hypothetical situation (as indicated by the use of a conditional), in which a degree would be necessary to obtain the relevant facts, in order to demonstrate how some relevant facts might be considered reasonably inaccessible.

recognize should be – if they are to be legitimately ascribable to her – derivable from information that is relevant to the deliberative situation at hand and also reasonably accessible to her (in conjunction with a reasonable amount of deliberation).²³ With regard to Gaus' bridge case, P should not have to acquire information about effective goat herding practices or obtain a civil engineering degree in order for her set of recognized reasons to perfectly synchronize with idealized P's. But what about cases in which P simply needs to check an outstanding voicemail, or look at the current weather report in order to achieve synchronicity with idealized P? These are the sorts of everyday cases that are troubling for the moderate idealizer. Yet, it is not clear how one might go about demarcating a domain of individual responsibility with respect to information acquisition that qualifies such information as both relevant and reasonably accessible.

In what follows I will present a series of considerations that I think shed light on the complexities surrounding the notion of information acquisition responsibilities. The first consideration focuses on one possible way of determining whether an agent has satisfied his or her information acquisition responsibilities. The remaining considerations are concerned more directly with the notions of relevancy and reasonable accessibility. I take each consideration to yield some insight that is helpful in developing an initial criterion for determining whether a given agent has a responsibility to take some particular piece of information into account. After discussing each of the considerations, I offer one plausible way of organizing the individual insights into such a criterion.

²³ This point assumes the reason under consideration cannot be arrived at through purely deliberative means alone.

4.1 Consideration for Effort

One possible approach to determining whether an agent has met her information acquisition responsibilities relies on gauging the level of effort that she exerts in terms of information acquisition. Call this the *effort view*. In Gaus' bridge case, we might say that P does not *have* a reason to cross – despite *there being* a reason to cross – if P did not arrive at a reason to cross and it is the case that P both engaged in a respectable amount of deliberation and put a reasonable amount of effort into acquiring as many of the relevant facts as possible. While this view entails an element of vagueness, the vagueness applies not to how much information is taken into account, but to how much effort the individual has put into canvassing the relevant facts. This may be helpful with regards to the bridge case, but that is only because P is confronted with a clear question to deliberate on: is this bridge safe to cross? Because P is acutely aware of the deliberative situation that confronts her, we can rightly expect P to search out facts relevant to the bridge case.²⁴ If P had spent two hours researching South American banana farming practices, we could legitimately claim that P had not put in a reasonable amount of effort into discovering the relevant facts (for P's efforts were misdirected). But many of the cases which challenge the moderate idealizer do not fit this picture. For example, it might be the case that P would have a reason to go to the hospital if only she listened to the unchecked voicemail left by her father. The issue here is that P is not directly confronted by the question of whether she needs to visit her father in the hospital. As it turns out, her father had been in great

²⁴ We are not always aware of the deliberative situations which require our attention. For example, while Vance may in fact have dangerously high cholesterol, it may be the case he is not aware of that fact. Yet, a deliberative situation exists with regard to Vance's cholesterol problem – he needs to decide on a management plan. The point is that troubling cases of reason ascriptions are cases where an external party recognizes a deliberative situation in which the agent in question does not.

health recently, so P was not “on standby” with respect to her father’s health. As such, the standard advocate of moderate idealization appears unable to specify whether P can be said to have a reason to visit her father in the hospital.

It may be the case then that we can only attribute a reason to an agent in cases where that agent is directly confronted with the same question (though in its first-person form) that we are attempting to answer (e.g., Does P have a reason to Φ ?). This is to say that P is *aware* of the deliberative situation – a situation that calls for P to weigh all reasonably-discoverable reasons $R_1 \dots R_n$ in order to make the most appropriate decision. Call this the *direct-confrontation requirement*. For it is when the agent in question is directly confronted with the same question that we are, that we can make claims about whether the agent has invested enough effort into discovering the relevant facts. The agent’s awareness of the question serves to focus the discussion on *what*, rather than how much, information is being investigated. This is the feature that is present in the bridge case, but not in the hospital case.

While the direct-confrontation requirement is an important qualification to the effort view, it does not rescue the effort view from what may be its most significant problem. The effort view is restricted to providing an answer to the question of whether an agent has put enough effort into meeting her information acquisition responsibilities. But this is not the precise problem that most troubles advocates of moderate idealization. The difficult cases facing such advocates turn on claims of the following sort: P had a responsibility to take into account information α . Notice, however, that the effort view cannot specify whether an agent had a responsibility to take into account a particular piece of information. For this view is not concerned with particular bits of information, but only with the agent’s effort. Because of this, an agent may be able to satisfy the requirements of the effort view without taking into account the

one piece of information that turns out to be crucial. As such, the effort view is insufficient, even when it incorporates the direct-confrontation requirement that serves to focus the agent's effort.²⁵

What the shortcoming of the effort view teaches us is that we need an account of information acquisition responsibility that can tell us whether an agent has a responsibility to take into account a *particular* piece of information. If the account can tell us that, then it will have the resources to also tell us whether the agent has met his or her information acquisition responsibilities with respect to a particular deliberative situation. In light of this, the remaining considerations will focus on various instances in which an agent appears to have a responsibility to obtain certain pieces of information.

4.2 Consideration for Societal Standards

One possible way of arguing that an agent has a responsibility to obtain information α is to argue that α qualifies as having a *sui generis* status in the society or local community. For example, it is commonly thought that agents (*qua* citizens) are responsible for having a grasp of any well-publicized societal standards, including legal and social rules in addition to other well-publicized societal or communal standards.²⁶ A standard which does not qualify as a legal or social rule is one which is neither a formal law nor a part of the social morality, but still considered necessary to achieve a basic level of "societal literacy." As I construe it, societal literacy entails a basic

²⁵ The effort view cannot be salvaged by arguing that one has put in enough effort only if one has discovered all of the relevant and reasonably accessible facts. Such an account just leaves us where we are now – trying to clarify the notions of relevancy and reasonable accessibility.

²⁶ To my knowledge, not much has been said about what grounds the validity of expectations regarding the acquisition of knowledge of certain epistemic norms. I suspect, however, that such grounds may at least be found in some social contract accounts. Nevertheless, this is not something I take up in this paper.

understanding of (1) how the society is formally structured; (2) how members of the society are expected to conduct themselves; and (3) the most common skills of survival and risk aversion within the society. With regard to (3), I have in mind a range of skills that runs (very roughly) from being able to buy food and write a check to being able to file one's taxes, purchase a home, and avoid common scams.²⁷

Many of the skills that fall under (3) will entail a heightened commitment to information acquisition. For instance, it is standard in communities across the U.S. to hire a licensed inspector prior to finalizing the purchase of a home. Any person who buys a home prior to having the home properly inspected is usually thought to have not performed his or her due diligence. For example, imagine Susan who is trying to decide between buying house A or house B. After weighing all the pros and cons of each house, Susan finds that she does not have a decisive preference for one house over the other. If only there were a good reason to prefer one house over the other her decision would be much easier, but as far as she knows such a reason does not exist. So, Susan flips a coin to decide and ends up buying house A. It turns out that there was a very good reason not to buy house A – it has a mold problem which would have been easily discovered by a licensed inspector. In such a case, Susan has failed to investigate responsibly, at least according to well-publicized home-buying standards in the U.S.

Alternatively, imagine that Susan did hire an inspector, but instead of having a mold problem, house A suffers from a leaking water line running to the house. The line leaks after

²⁷ I acknowledge that there is plenty of room for debate about which exact skills would be required for one to achieve societal literacy. As one reviewer has noted, familiarity with home buying standards might extend beyond what we would consider to be the most common skills of survival and risk aversion in our society (since, for example, not every person will be so situated (e.g., economically) that such information is needed). It is perhaps the case then that this set of skills will be even smaller than I have anticipated.

being corroded by harsh chemicals that entered the ground after a nearby laundry endured a chemical spill years ago. While a licensed inspector probably would not have discovered the problem, Susan would have known to further investigate the issue if she (1) had purchased an Environmental Data Resource (EDR) report in order to learn of the spill; and (2) surveyed members of the community to see how they were affected by the spill. According to well-publicized home buying standards in the U.S., Susan could have performed her due diligence and still not discovered the faulty water line. This is because obtaining EDR reports is not necessary to performing one's due diligence with respect to home buying. Thus, in the first case, but not the second, Susan had certain information acquisition responsibilities, grounded in community standards, which were not met. Consequently, we could legitimately ascribe a reason to Susan to not buy house A in the first case, but not the second.

Consideration for societal standards need not be limited to the standards of the broader social-political society. We may include here the standards of specialist communities. Cases involving specialist knowledge show us that our expectations about the reasons which an agent has a responsibility to discover are affected by certain considerations about the agent's particular background. If P is a civil engineer and is trying to decide whether a bridge is safe to cross for the general public, we may expect P to take account of certain facts that we would not expect a typical agent to take account of. That is, beyond that information which is given a *sui generis* status by the broader society, P may be expected to take into account additional information in virtue of his or her specialized background or unique circumstances.

Holly Smith begins her paper "Culpable Ignorance" with a polarizing case that deals with this very issue. The case is presented in light of philosophers who distinguish "between the

objective *wrongness* of an action and the agent's *blameworthiness*" (Smith 1983, p. 543). Smith continues to summarize the position:

No matter how terrible the act may be from an objective point of view, the agent is not blameworthy if he had an excuse for what he did. Thus a doctor who treated a premature infant's respiratory distress in 1954 by exposing her to unnecessarily high concentrations of oxygen, and so caused severe eye damage, did something terrible. But if the doctor did not realize that oxygen enrichment would have this effect, he is not to blame for the baby's blindness. Ignorance of the nature of one's act is the pre-eminent example of an excuse that forestalls blame. However, there are occasions when a person's ignorance is itself criticizable – when he *should have realized* what he was doing. Perhaps the doctor should have known that high oxygen enrichment would induce blindness: the latest issue of his medical journal described a study establishing this effect and recommending the use of lower concentrations as equally effective for respiratory problems. The doctor should have read his journal, and if he had done so, would have realized he ought to use less oxygen. In cases such as these, opinion is sharply divided: many people feel that such an agent must be blamed for using excessive oxygen, while others insist that his ignorance, however criticizable in itself, nonetheless excuses him (1983, p. 543).

Though Smith does not frame the case in such terms, the divisive responses that she envisions to the case will arguably be fixed by expectations about what information the doctor had a responsibility to take into account. According to Smith, the case is one of *deficient*

investigation: the doctor “cannot be excused for his act because he should have realized what he was doing” (1983, p. 544). But this is only to say that the doctor *had* a reason to use a lower concentration of oxygen. If he had no such reason, then the claim that he cannot be excused is ungrounded. The reason would have been discovered if only the doctor would have fulfilled what are presumed to be his information acquisition responsibilities. This is apparent in Smith’s claim that “he ought to have read [the journal], and if he had, he would have discovered the use of high oxygen concentrations to be unnecessarily harmful to the infant” (1983, p. 544).

But Smith does not consider whether there are legitimate grounds for the claim that the doctor ought to have read his journal. It is assumed that he “could have acquired the requisite information, ought to have done so, but failed to fulfill this obligation” (Smith 1983, p. 547). As such, Smith does not engage the issue of information acquisition responsibilities which lurks in the background of her paper.²⁸ Once the issue is raised, however, one might ask: is this a journal that the American Medical Association expects all doctors to read? If the answer is yes, then standards particular to the doctor’s circumstances (that of being a licensed physician) may dictate that he had a responsibility to take such information into account.

4.3 Consideration for Generally-Reliable Sources

There are cases in which the agent in question may not have any grounds for thinking that the mechanism by which information α is acquired is a mechanism which he or she should investigate. Returning to the case of Susan’s pending home purchase: suppose that she has never heard of the EDR report and has no idea that such a report could be useful in deciding between

²⁸ Smith does consider a case in which the doctor’s nurse fails to inform him that the journal had arrived (see 1983, p. 548). On account of this, the doctor’s failure to read the journal is presumably excused. This, however, still assumes that the doctor had a standing obligation to read the journal.

the two houses. This fact, in combination with the fact that obtaining EDR reports is not necessary to meet home-buying standards in the U.S., leads us to believe that Susan did not fail to meet her information acquisition responsibilities. But one's evaluation of Susan's situation may change if it were the case that the home inspector advised Susan to obtain the EDR report given that a recent home buyer in the area encountered the same problem. The inspection report, he might have said, would inform Susan whether or not house A falls inside the area affected by the chemical spill. So, societal standards alone may not be enough to ground one's full information acquisition responsibilities. Generally-reliable sources – such as a licensed home inspector advising a potential home buyer – may provide an agent with information that further expands one's information acquisition responsibilities. In the case at hand, Susan now knows that there is a method of information acquisition (the EDR report) that is relevant to her pending home purchase.

4.4 Consideration for Special Commitments

In cases where an agent has agreed to some special commitment, such a commitment may entail additional information acquisition responsibilities. For example, if Joey and his wife have agreed that Joey will be responsible for managing the family finances, then Joey has a responsibility to know the days of the month in which the family's bills are due. Such a responsibility comes in virtue of one's formal and informal contracts.

Importantly, one's special commitments may entail information acquisition responsibilities that help to inform deliberative situations that are related to one's special commitments. For example, Joey's responsibility to know the days of the month in which the family's bills are due will certainly inform the following deliberative situation: Does Joey have a

reason to pay the phone bill this month? In such a case, we need not worry about whether Joey is directly confronted with the deliberative situation at hand before providing an affirmative answer to this question. His agreement to manage the family finances requires a commitment to an ongoing awareness of this particular deliberative situation.

It is also possible that there may be cases in which one's special commitments entail a responsibility to acquire some information α which will prove useful in a deliberative situation that is not directly related to one's special commitments. For example, there may be some (fanciful) case involving Joey in which knowledge of when his family's phone bill is due would be relevant to a deliberative situation that is not directly related to his special commitment to manage the family finances.

4.5 Consideration for Significant Costs

One might think that when a decision under consideration is of exceptional significance, and the agent has adequate time to conduct a highly-thorough investigation, that obtaining information α is reasonably expected, even if the acquisition of α would not be required in light of the previous considerations. The sort of cases that I envision here are ones in which the agent encounters a deliberative situation of atypical significance – e.g., involving a significant amount of risk or carrying a heavy penalty for a poorly-chosen course of action. Whether to invest one's entire life-savings into an unproven investment vehicle or whether to opt for a risky treatment in an effort to cure a not-well-understood medical condition are examples of deliberative situations of atypical significance.²⁹

²⁹ It is unclear just how common these sorts of cases might be. After all, with regard to some particular (hard) case c , while a community may not prescribe any standards for how to deliberate about c in particular, it may prescribe

4.6 Summary

I have discussed some considerations regarding how one might go about responding to the everyday cases that accounts of moderate idealization are challenged by. I believe that these considerations help to demonstrate that our expectations about what reasons an agent can be said to have are significantly influenced by our expectations about what information acquisition responsibilities that agent can be said to have. Clarifying exactly what those responsibilities are is the challenge that confronts us. As I hope to have shown, incorporating an account of information acquisition responsibility into an account of the reasons one can be said to have raises difficulties. This is especially true of cases where one person demands that another adhere to some moral imperative Φ . For instance, imagine P who had access to a public lecture which, if attended, would have given P a reason R to act according to some moral imperative Φ (in virtue of the acquisition of new information). In the case of Q, who demands that P act according to Φ , what does Q's demand entail? It seems that Q's demand is grounded in the claim that P *ought* to have attended (or had a responsibility to attend) the public lecture (and, hence, ought to having acquired the relevant information). But why think that P had such a responsibility? Moreover, if P does not in fact have this responsibility, then Q's demand is void of legitimate authority. That is, so long as we cannot explain why P ought to have attended the lecture in order to acquire the relevant information, then Q's demand is purely authoritarian in nature (i.e., it is a demand that persists despite not being legitimate). This is the concern that motivates the need to clarify the idea of information acquisition responsibility.

standards which advise a general set of cases of which *c* is a member. If so, expectations about the agent's deliberation on *c* may be captured (at least partially) by considerations for community standards.

Perhaps a criterion for determining whether a given agent has a responsibility to take some particular piece of information into account will help to summarize the discussion up to this point. Given the claim that P has a responsibility to take into account information α if and only if α is relevant and reasonably accessible to P, something like the following criterion may help to express the notions of relevancy and reasonable accessibility: information α is relevant and reasonably accessible to P (i.e., P is responsible for acquiring α) if and only if (1) A obtains; and (2) at least one of B-E obtains.

(A) P is *aware* of deliberative situation S – a situation that calls for P to recognize α as part of weighing all reasonably-discoverable reasons $R_1 \dots R_n$ in order to make the most appropriate decision.

(B) α is considered to have a *sui generis* status by the society or specialist community of which P is a recognized member and α would lead P to discover³⁰ one or more reasons crucial to S. (This may include legal as well as social rules in addition to other well-publicized community standards.)

(C) P has learned through generally-reliable means³¹ that method M will lead to α which is relevant to S and, were P to discover α , P would discover one or more reasons relevant to S.

(D) P has agreed to Φ and Φ -ing entails a commitment to acquire α which, when taken into account, would lead P to discover one or more reasons

³⁰ Requiring no more than a reasonable level of deliberation.

³¹ The generally-reliable means may be partially determined by the community standards.

relevant to S (where S may or may not be directly related to the act of Φ -ing).

(E) S is of such exceptional significance – and the agent has adequate time to conduct a highly-thorough investigation – that obtaining α is reasonably expected, even if the acquisition of α would not be required on the grounds specified in (B) – (D).

This criterion is intended only to incorporate what I take to be some of the less controversial expectations regarding the information acquisition responsibilities of typical agents. It construes an agent's information acquisition responsibilities in terms of information that is relevant and reasonably accessible to the agent. I do not claim this criterion to be exhaustive; nor is it without difficulties. With regard to potential problems, one might expect for disagreement to arise at a number of points. For instance, people might disagree about whether a particular piece of information should be considered a part of societal literacy. Or they might disagree about whether a particular source of information should be considered generally reliable. Or they might disagree about whether a particular piece of information should be required by one's special commitments.

While each of these issues raises their own considerations, they do not undermine the usefulness of a criterion for determining whether a given agent has a responsibility to take some particular piece of information into account. Such a criterion will still reduce the domain of borderline cases by honing in on *what* information can legitimately be expected to be taken into account by a typical agent. This is important since determining the information acquisition responsibilities of an agent is a crucial part of any legitimate reason ascription to that agent –

insofar as an ascription of reason R to P is justified only if P has a responsibility to acquire that information which is necessary in order to recognize R.

5. CONCLUDING REMARKS

I have primarily been concerned with trying to articulate the troubles that surround the idea that P can be ascribed a reason to Φ just as long as if only P were to take into account some relevant piece of information (that P has access to), then P would recognize for herself that she has a reason to Φ . While such an idea may seem innocuous at first pass, I have argued that it assumes an account of responsibility with respect to information acquisition that is easily overlooked. But without a proper grounding, expectations about information acquisition responsibilities can lead to illegitimate reason ascriptions.

Acknowledgements

I wish to thank Bruce Brower, Gerald Gaus, Eric Mack, David Shoemaker, and the two anonymous referees for *Res Publica* for their many helpful comments on earlier versions of this article.

References

- Alvarez, Maria. (2010). *Kinds of Reasons*. New York, NY: Oxford University Press.
- Cherniak, Christopher. (1986). *Minimal Rationality*. Cambridge, MA: MIT Press.
- Gaus, Gerald. (1996). *Justificatory Liberalism*. New York, NY: Oxford University Press.
- (2011). *The Order of Public Reason*. New York, NY: Cambridge University Press.
- Hardin, Russell. (1988). *Morality within the Limits of Reason*. Chicago, IL: University of Chicago Press.
- Pasnau, Robert. (2013). “Epistemology Idealized.” *Mind* 122, 488: 987-1021.
- Pollock, John L. (2006). *Thinking About Acting*. New York, NY: Oxford University Press.

Simon, Herbert A. (1976). "From Substantive to Procedural Rationality." In *25 Years of Economic Theory*. New York, NY: Springer US, 65-86.

Smith, Holly. (1983). "Culpable Ignorance." *The Philosophical Review* 92, 4: 543-571.

Vallier, Kevin. (2014). *Liberal Politics and Public Faith*. New York, NY: Routledge.

Williams, Bernard. (1981). "Internal and External Reasons." In *Moral Luck*. New York, NY: Cambridge University Press, 101-113.