

On the Relevance of Ignorance to the Demands of Morality¹

Geoffrey Sayre-McCord

“It is impossible to overestimate the amount of stupidity in the world.”
Bernard Gert²

Introduction

In *Morality*, Bernard Gert argues that the fundamental demands of morality are well articulated by ten distinct, and relatively simple, rules. These rules, he holds, are such that any person, no matter what her circumstances or interests, would be rational in accepting, and guiding her choices by, them. The rules themselves are comfortably familiar (e.g. “Do not kill,” “Do not deceive,” “Keep your promises”) and sit well as intuitively plausible. Yet the rules are not, Gert argues, to be accepted merely because they are intuitively attractive, nor because they are already widely recognized, but because they stand as the only set of rules that can qualify as appropriately acceptable to all rational beings.

Two questions naturally arise when confronted with a proposed account of morality’s standards. The first question is: Why believe the account is true? The second question is: Supposing it is true, what reason do people have to accept those standards as guides either for themselves or others?

Before turning directly to the details of Gert’s defense of the rules and the accompanying answers he offers to these two questions, it is worth highlighting the different attitudes people have taken toward these questions and the relation answers to them might bear to one another.

A fair number of people think that the two questions are completely independent. One question focuses on what morality demands while the other focuses on the apparently distinct issue of whether, and why, one might rationally accept morality’s demands either as a standard for oneself or for others. In general, it seems just a matter of being clear-headed to recognize these as different issues. After all, it is one thing to identify what a particular legal system might demand, for instance, and something else entirely to decide whether it is rational to obey those demands. Similarly, isn’t it one thing to identify what morality might demand and something else entirely to decide whether it is rational to

obey those demands? It seems pretty clear that what is demanded, and what one might rationally do, can come apart quite radically.

Still, at least when it comes to the demands of morality, there are also a fair number of people who think that once one has established demands as the demands of *morality* there is no real question as to whether it is rational to accept or obey them. Morality's demands, they argue, set an ultimate standard and enjoy an authority that provides decisive reason to act. That one has a moral (as opposed to a legal, or professional, or even prudential) duty to act in some way, they hold, settles the issue of whether one has adequate and even decisive reason to do it. So in determining that some set of standards are in fact those of morality one has settled whether or not people should accept them as guides to their own and other peoples' behavior. The two questions are not really distinct, so not independent either.

Not surprisingly, there are people who fall in between these two extremes, people who hold that while both questions stand as distinct, an answer to one of the questions depends, in some interesting way, on the answer to the other. Determining what might be rationally accepted, some of them suggest, requires first determining what morality demands, permits, and forbids. Flipping the dependence around, others suggest that figuring out what morality demands requires first figuring out what might be rationally accepted. Either way, the idea is that successfully defending an answer to one of the questions will require grappling with the other.

Fitting squarely into this third group, Gert sees the two questions as distinct but not independent. He holds that no proposed account of morality is even in the running unless its rules or standards are such that those to whom they apply can rationally accept them. As a result, successfully defending his own account as true, and so answering the first question, involves answering the second.

Gert takes this position because he thinks morality is essentially "an informal public system," where a set of standards counts as an informal public system only if it is both known by,³ and rationally acceptable to, all to whom it applies.⁴ And he argues that, since moral judgments may be made about any and all rational agents, the rules or standards on offer must constitute an informal public system that applies to all rational beings, so "every feature of morality must be known to and rationally chosen by any rational person."⁵ As a result, if not all rational agents could rationally accept some proposed moral system, the system would be (according to Gert) a non-starter.

One thing to ask, then, about any proposed moral system, Gert's included, is whether it satisfies the second condition on being an informal public system—whether it is rational for people to accept its rules as setting a standard for themselves and others. It is important, that is, to pursue an answer to the second question I mentioned at the beginning.⁶

In pursuing it, Gert offers a particular account of when and why a system would count as rationally acceptable and argues that his account—and any plausible account—must prove rationally acceptable to all to whom it applies (which is all rational agents). I am skeptical, for reasons I will try to bring out, concerning both his account of rational acceptability and his contention that an account of morality, to be plausible, must be an account of a system that is rationally acceptable to all to whom it applies.

Rational Acceptance

When is a system of rules such that all rational beings could rationally accept it? Gert's answer to this question falls within a familiar framework of justification. The framework is one in which someone is taken to be rational in accepting (or choosing, or doing) something if and only if a fully rational person, appropriately situated, would favor accepting (or choosing, or doing) it.⁷ Real people, of course, are often irrational, woefully ignorant, horridly biased, etc. What real people might approve of frequently carries pitifully little justificatory weight. Yet those same people, cured of their ills, appropriately informed, freed of partiality, would, it seems, set a perfect standard. What grounds could there be for discounting their approval? Seemingly: None. If none, then, the thought goes, what such people—made fully rational and appropriately situated—would favor should set the standard for what we might rationally accept.⁸

This general line of thought, and the resulting framework for thinking about justification, has informed a number of otherwise quite different approaches to moral theory—Ideal Observer theories⁹, informed desire theories,¹⁰ and theories that appeal to what rational beings would agree to in a suitably described contract situation.¹¹ All of these views set out to link the rational acceptability of a set of rules or standards to the approval of rational beings appropriately situated. Different versions of each view advance different characterizations of what goes into being a rational being, of what sort of approval is relevant, and of what is involved in these rational beings being appropriately situated, when it comes to setting the standard for real people. Some think that the rational beings whose approval matters to the justification of moral standards are those—and only those—who are fully informed. Ignorance, on this view, disqualifies. Others, John Harsanyi and John Rawls, for instance, agree that all general knowledge is required, but hold that the sort of impartiality important to the justification of moral standards requires some ignorance, specifically ignorance of one's identity.

Gert goes further by requiring (in effect) almost total ignorance—to the point where those whose approval sets the standard for rational acceptability of a proposed set of moral rules are to rely *solely* on the beliefs it would be “irrational not to hold.”¹² All and only such “rationally required beliefs” are supposed to figure in determining the approvals that matter, according to Gert.¹³ Otherwise, he points out, the acceptability of the proposed rules would turn on considerations available not to all rational beings, but only to those who are better informed than rationality itself requires.¹⁴ The need to show that the rules are acceptable to all rational agents thus imposes a striking constraint, in Gert's hands, on what might be appealed to in determining the acceptability of a proposed set of rules as the rules of morality.

Among rational agents, those who possess only rationally required beliefs fall at the very bottom of the heap. Any belief that some rational agent or other might lack without irrationality, they lack. Consequently, they know extraordinarily little about themselves, about others, and about the world they inhabit.¹⁵ It would not be unfair to say that, among rational agents, they constitute the lowest common denominator. Of course, in some ways they might still be better off than many real people who fail to be rational. But even compared to such real people these rational agents know amazingly little. And if we compare these rational agents (who hold only rationally required beliefs) with others who have learned more about how the world around them works, they will look poorly off

indeed. Nonetheless, Gert argues, the fact that morality is meant to apply, as an informal public system, to all rational agents means that the acceptability of the rules that constitute the system cannot depend on anything not shared by all rational agents.¹⁶

On Gert's view, then, a proposed moral system is rationally acceptable, in the relevant sense, if, but only if, those who are extraordinarily ignorant—and others relying solely on rationally required beliefs even the ignorant would have—would favor that system. Only then, he maintains, would the system have a claim to being such that *all* rational beings, and so all to whom the system applies, could rationally accept it.

Yet, as Gert notes, even the system of rules he defends can be shown to secure the approval he thinks necessary only if, in addition to constricting the relevant beliefs to those that are shared by all rational beings, some assumption is made concerning the desires that might influence the attitudes of the people in question. What rational beings might favor, in light of their beliefs, will turn on what they happen to value or care about. Whether any, let alone all, rational agents who rely solely on rationally required beliefs might favor something or not is up in the air until their desires, values, and concerns, are specified.

Just as there are certain rationally required beliefs, Gert holds that there are some desires that are rationally required, for example, the desire to avoid pain, and some that are irrational (i.e., irrational to act on absent adequate reason), for example, the desire to die.¹⁷ And just as hope for securing a consensus among all rational beings appropriately situated involved seeing the appropriate situation as involving a reliance solely on rationally required beliefs, so too it might seem that the appropriate situation would likewise require that only rationally required desires be allowed to have sway. As it turns out, though, limiting the approval that matters to approval prompted solely by rational desires will still not secure unanimous approval for the system of rules Gert proposes. Indeed, as he notes, a rational person may consistently be an egoist, and in being an egoist will favor rules that constrain others, but not herself, whereas the rules Gert proposes are rules that are to constrain all.¹⁸

"The problem," as Gert sees it, "is how to replace the egocentric attitude . . . with the appropriate moral attitude while at the same time keeping it an attitude that would be taken by all rational persons."¹⁹ It is, he recognizes, not a solvable problem. There is no such attitude. Gert admits that, given certain commitments that are rationally permitted, a rational person will not adopt anything like the sort of attitude that Gert needs to suppose is shared in order to secure the relevant approval for the rules he recommends. In fact, as he recognizes, some rational beings might, without irrationality, acquire convictions, say religious convictions, that lead them to reject not only the rules as applied to themselves but also as applied as egoistically motivated constraints on others. "A rational person," he notes, "might hold some religious belief that was in conflict with taking the egocentric attitude toward these rules, and he might give priority to the religious belief."²⁰

Still, while there is no attitude all rational persons must take, the fact that we are considering the rules as proposed moral rules—as part of an informal public system that applies to all rational beings—means, Gert maintains, "that a rational person cannot adopt any attitude toward the rules, except one like that which would be adopted by a rational person who has an impartial concern for all persons."²¹ Considering the rules as moral rules constrains us, that is, to relying only on attitudes that all other rational beings might share. In one sense, of

course, they might share an egoistic attitude, but the up-shot of that attitude is for them each in effect to differ, sometimes radically, in what they actually favor. But they can really only share an attitude that makes no egocentric reference, either implicitly or explicitly. In other words, considering the rules as moral rules leads one inevitably to adopt both an impartial attitude and a commitment to mutual acceptability.

If we can stipulate, as this argument suggests, EITHER that the rational beings in question are concerned to favor moral rules OR, what turns out to be functionally equivalent, that they are committed to impartiality, OR, functionally equivalent again, that they are willing to favor only what could be favored by all rational beings, the argument for Gert's proposed system seems to go through.²² For any one of these concerns, in contrast with an egoistic attitude, will lead someone relying solely on rationally required beliefs to favor the rules Gert recommends.

At the same time, only with one or another of these additional commitments on board will the system of rules Gert endorses be favored by those who rely only on rationally required beliefs. Accordingly, Gert ends up arguing that a candidate moral system is justified in the relevant way—is such that all to whom it applies could rationally accept it—if and only if all rational people who rely only on rationally required beliefs (and suffer no irrational desires) would, if considering the system as a moral system, or if committed to impartiality, or favoring only what could be favored by all, all favor it. Showing that a system enjoys this sort of justification—showing that it would be favored by *all* rational beings appropriately situated—is establishing that the system in question is (in Gert's terms) strongly justified.²³

As it happens, Gert holds that any one of these three commitments—to morality, or impartiality, or mutually acceptability—is sufficient to secure the favor of all appropriately situated rational beings. And he takes the fact that each commitment will, independently of the others, secure the relevant approval reveals the deep sense in which the justification of a moral system is tied both to its being impartial and to its being acceptable to all. If the arguments work, one gets a lovely convergence.

I am not here going to spend any time on the question of whether the arguments work. For the sake of this discussion, we can simply grant that all who rely solely on rationally required beliefs and are concerned either with morality, or with impartiality, or with what all could favor, would favor the set of rules Gert advances.

Leaving these claims in place, I am instead going to focus primarily on Gert's contention that the rational acceptability of a set of rules that are to apply to all rational beings depends upon their being favored by those who are influenced solely by rationally required beliefs in a way that leaves out of account whatever might be known to some but not all. In other words, my interest is in Gert's view that the appropriate standard of justification is set by what the lowest common denominator might approve.

Why Only Rationally Required Beliefs?

Why accept this view? Why take what people who are woefully ignorant would approve of as setting the standard for whether a system is rationally acceptable in the appropriate way? Why think that what they might favor has any implications whatsoever for what people like us might rationally accept?

One might think that an appeal to their approval works as a kind of argumentative *tour de force*. After all, if even those who are pitifully ignorant would favor the system of rules Gert advocates—if even those who know very little and care not at all about others would endorse rules that constrain everyone from inflicting harm on others—that would seem a strikingly impressive discovery. Moreover, one might think, if even the ignorant and uncaring would favor the rules, then surely all other rational beings would as well. So, by limiting the premises of the argument to claims about what even the worst among us might favor, it may seem the argument avoids appealing to anything either morally or rationally controversial. In relying on the shared subset of beliefs that are rationally required we seem to have what is sufficient to ensure the favoring of the rules (assuming either a concern for morality, or for impartiality, or for securing the favor of all), so the argument apparently needs to rely on nothing more.

Now this is not Gert's strategy, and for good reason. To understand why, it is important to note that reasoning in a practical sphere is non-monotonic (i.e. defeasible)—adding premises to what was, initially, a good argument can undermine the conclusion. For this reason, given an initial inference that is acceptable, it is a mistake to assume that adding information will leave the inference untainted. If all I know of someone is that she is your beloved sister, I might reasonably trust her . . . until I learn that you do not trust her yourself, in which case I might reasonably be suspicious of her . . . until I discover that you trust no one, under any circumstance, in which case it would be unreasonable for me to suspect her . . . until I learn that she has lied to you. . . . Similarly, what people might favor if they rely on certain beliefs (e.g. concerning your sister's trustworthiness) may well be something they quickly come to oppose in the face of new information. People who have more than the beliefs that are rationally required may well not favor things that they would have favored had they not acquired extra information.

Indeed, as Gert notes explicitly in *Morality*, what people limited to rationally required beliefs would favor are things that people who rely on other rationally permitted beliefs—say scientifically founded beliefs, or various moral or religious beliefs—would firmly oppose.²⁴

This fact about how beliefs might influence what rational people favor means that limiting the argument to the attitudes of those relying solely on rationally required beliefs is a seriously *non-trivial* assumption. Far from ensuring the power of the argument by bringing all rational beings within its scope (which is what it may have initially seemed to be doing in relying on rationally required beliefs), the limitation dramatically constricts its appeal. For what such people would favor would be rationally rejected by others who are more informed. Most significantly, actual people (virtually all of whom have more than the beliefs rationally required of all) will often be among those who will rationally not favor what those who are less informed would.²⁵ Actual people who are rational and informed about the world might well find themselves committed to standards that require people to help others in need—a requirement rejected by Gert. Alternatively, they might come to reject (perfectly rationally, it seems) standards that require obeying the civil law—a requirement imposed by Gert.

As far as I can tell, for all Gert argues, actual people may come to have certain beliefs that are rational for them to hold and *in light of which* accepting the standards he defends as guides to their own behavior and the behavior of others would be positively irrational. In the cases I am imagining, people have, without irrationality, come to hold standards that differ in some significant way from the

ones Gert proposes. The standards might, on these peoples' view, have their source in reason, or in religion, or in nature. Whatever the supposed source, given that the people actually have the convictions, and assuming that, given their evidence, they have positive grounds for embracing the convictions, if their standards conflict with the rules Gert advocates, they may well be in a situation in which obeying the rules he proposes goes against all that they reasonably see themselves as having reason to do. This is true regardless of whether their views are actually correct, as long as the views are rationally held and provide the people with reason to accept the standards they do.

There are two kinds of cases to keep in mind here, I think. In some cases, the people in question might find themselves with convictions that require them to accept standards—say requiring benevolence or self-sacrifice—that go beyond Gert's. In others, they might find themselves with reason to reject one or another of the rules Gert advocates.

Cases of the first kind are interesting because Gert is committed to rejecting any rules that would not be favored by *all* rational beings who rely solely on rationally required beliefs and a concern for establishing rules all could accept. To the extent a requirement of benevolence (for instance) would be grounded in reflections that go beyond what such people would engage in, Gert would reject the requirement as one of morality. In requiring more than what rational beings situated in the way he specifies would favor, the rule goes beyond morality. That is not to say, of course, that Gert is committed to opposing beneficence, nor is it to say that he is committed to thinking that beneficence is not valuable. It is to say, though, that he is committed to rejecting it as required by morality. Thus he is at odds, in a deep and interesting way, not only with utilitarians, but also with Kantians and with many others who see having a concern for the welfare of others as a fundamental moral requirement.

Cases of the second kind, in which people are rationally required to reject (as opposed to add on to) Gert's rules, introduce some nice complexities. The rules Gert proposes are all very general and each allows exceptions. So there is some reason to wonder whether a plausible view could reject any of them without, in effect, mobilizing the sort of considerations the rules, in Gert's hands, already recognize as justifying exceptions. Thus, for instance, while Gert thinks that among the fundamental rules is a requirement that one obey the law, he does not suppose that people are never justified in violating the law. Under certain circumstances, he notes, they clearly are.²⁶ In a similar way, the prohibitions against lying and stealing and killing and adultery are all prohibitions that admit of exceptions.

To get a case in which people might rationally reject the rules Gert advocates, one needs a case in which the peoples' convictions would require them to reject the rules, taking into account the rules' ability to accommodate exceptions.

According to Gert all the exceptions come under the same broad rubric—they are all such that an impartial rational person can publicly allow their violation. Thus, for instance, the general rule that one should obey the law does admit of exceptions—exactly when, and only if, “an impartial rational person can publicly allow the law to be ignored or broken”.²⁷ This means that the alternative view to focus on is one in which the legitimacy of an exception to one of the general rules is not seen as dependent upon its being publicly allowable. Various consequentialist views come to mind as candidates that recommend, in the world as it is, that we promulgate and follow rules that largely coincide with

Gert's, even as it holds that exceptions are legitimately made whenever in fact (or in prospect) the consequences are better—whether or not the exceptions might be publicly allowed.

There are familiar problems with such a view, but it does not seem right to think that a person who accepted consequentialism would thereby be irrational. And it does seem right to think that under some circumstances, in the face of certain bits of evidence and with a certain background in place, it would be positively irrational for someone to reject her consequentialism just because those who know little and do not care about the welfare of others would not favor its requirements.

Even as the limitation to rationally required beliefs is non-trivial, it is also *indispensable* for Gert's defense of his system. For that defense goes through, if it goes through at all, only if a whole slew of rationally permissible beliefs are put aside—among them beliefs many people would be irrational to reject, given their evidence. If additional beliefs were allowed in, there would be no grounds at all for thinking all rational agents would approve of what they would approve if they knew much less. Indeed, there is good reason to think they would not, as Gert himself acknowledges: "All rational persons have rationally allowed beliefs, based on some combination of their present circumstances and their training or education, such that it would not be irrational for them not to favor impartial obedience to these rules."²⁸ It is worth noting that the beliefs that count here as (merely) rationally allowed might be such that, given "their present circumstances and their training or education," people would have to be irrational not to hold them. Thus, in context, the beliefs might be rationally required of the person in question. And in light of those beliefs, it might also turn out that the person would be irrational to favor impartial obedience to the rules, as when, for instance, favoring such obedience would require failing to do something they rationally believe to be required.

There are, then, two different contexts in which additional information could influence what someone might rationally favor and so influence the arguments. One context is that in which real people find themselves, the other is the context in which those who are supposed to set the standard for actual people find themselves.

In the first context, actual people who hold beliefs that go beyond those that count as rationally required might well find themselves with beliefs that would make it irrational for them to constrain themselves by the rules Gert proposes. Even if they would favor the rules were they to be influenced solely by rationally required beliefs and a commitment to impartiality, they are in fact influenced—and seemingly appropriately—by other beliefs and commitments that stand in the way of their accepting the rules. Of course, supposing their failure to accept the rules, under the circumstances, is appropriate is to reject the view that the standard of rational acceptability for them is set by the approval of those who constitute the lowest common denominator of rational agents. For the problem emerges precisely because what real people might rationally do, or accept, seems to be highly sensitive to what they, in their context, actually believe and have evidence concerning. So responding to this sort of worry requires showing, somehow, that actual people all have reason to believe or accept that the rules of morality (at least) are exactly those rules that would be approved of by people responding solely in light of rationally required beliefs and a concern with impartiality.

In the second context, whether the people who are supposed to set the standard for actual people would approve of just the rules Gert defends turns crucially on what beliefs and commitments are genuinely rationally required. Even if the standard of rational acceptability, when it comes to proposed rules of morality, is set by what would be approved of by people influenced solely by rationally required beliefs and a concern for impartiality, that standard is highly sensitive to what goes into the set of beliefs that count as rationally required. The fact that these considerations are non-monotonic means that discovering that an additional belief is rationally required can make a dramatic difference to whether a set of rules would secure the relevant approval. One cannot assume that, having been favored in light of some beliefs, a set of rules will continue to be favored as more beliefs are added in.

In both contexts, the threat to Gert's arguments is most clearly posed by moral and religious commitments people might rationally acquire that would lead them to approve of sets of rules Gert rejects or disapprove of the set of rules Gert advances. Such commitments most clearly have direct implications for whether a person might rationally approve of some set of rules advanced as standards for behavior.²⁹ For instance, on Gert's view, neither beneficence nor self-sacrifice are morally required, since rules requiring them would not be favored by those relying only on the beliefs and desires that are rationally required. Yet people might, so it seems, rationally accept utilitarianism (even if it is not, after all, the right moral standard). Such people will (rationally, though perhaps not correctly) see beneficence and self-sacrifice as morally required. At the same time, they will (again rationally) see rules Gert endorses, for instance the one requiring that promises be kept, as not setting an independent standard of behavior but as only capturing the general fact that breaking a promise usually has comparatively bad consequences.³⁰ Similarly, one might rationally accept a theological view in light of which certain kinds of acts—of devotion, or atonement, or obedience—are morally required. If these acts are incompatible with rules of the sort Gert advances, then a person who held such a view would find herself with a commitment not to accept those rules as authoritative.

One might argue, of course, that when we are trying to determine the correct standards for morality it is a mistake to give weight, in the first instance, to peoples' current evaluative commitments. After all, these are precisely what are at issue. We need, instead, to bracket such commitments. What would people rationally accept (in the first context) or rationally favor (in the second) were they to leave out of account their prior evaluative commitments concerning morality? This looks to be the relevant question, and it would clear away the problematic cases that appeal to those who already, as it were, accept utilitarianism, or Kantianism, or some other substantive normative view. In order to decide for ourselves which, if any, of these views are correct we need to ask, it seems, which such view would gain the appropriate approval.

This is reasonable enough, but no real help with the problems at hand. For the problems arise not on the assumption that people hold distinctively moral commitments that clash with the rules Gert is defending. Rather they emerge once we recognize that people might, for various reasons, or no reason at all, have commitments—say to the welfare of others or the significance of a deity—that make it irrational for them to accept a proposed set of rules and that influence what rules they might rationally favor. These commitments need neither be, nor be held to be, *true or correct* for them to have the relevant impact. So, the question of whether the commitments are true or correct is not, at this stage,

at issue. Instead, when focusing on the first context, where we are considering whether actual people might be rationally required to reject Gert's proposed rules, the question is whether it is always irrational for actual people to have the commitments that would make their accepting Gert's rules irrational. If it is not always irrational to have such commitments, then actual people will sometimes find themselves rationally required to reject the rules Gert proposes (despite their supposedly being acceptable by all rational beings). And when focusing on the second context, where we are considering what those who set the standard for rational acceptability might favor, the question is whether they are appropriately situated only if they lack the commitments that would lead them not to favor the rules Gert proposes. If those who set the standard might be influenced by these commitments, then they might find themselves not favoring the rules Gert proposes.

The question remains: Why restrict our attention, in trying to determine what it would be rational for actual people to accept, to the attitudes of those who rely solely on the "rationally required beliefs" and a commitment to impartiality (or its functional equivalent)? Why think the approval of such people establishes what it would be rational for us to accept, given that their approvals depend on ignoring things we know (and on failing to value things we value)?

There is an answer to be found in *Morality*, one that brings us back directly to the claim that morality is, essentially, an informal public system that applies to all rational beings. Assuming that morality is such a system, it has to be both (i) known by, and (ii) rationally acceptable to, everyone to whom it applies.³¹ If morality applies to all rational beings, it has to be known by and acceptable to all rational beings. Thus any system that could be justified only by appealing to beliefs not all rational beings might have would be a system that risked having aspects to it that not all rational beings could know and would rest on a justification some could neither know nor appreciate in a way required in order for them rationally to accept it. "[S]ince morality applies to all rational persons," Gert argues, "all of the essential features of morality must be understood and acceptable to all rational persons, and hence no religious, metaphysical, or even scientific belief that is not shared by all rational persons can be used to determine the essential features of morality."³²

I should emphasize that Gert neither argues nor thinks that religious, metaphysical, and scientific beliefs are irrationally held. Indeed, in certain contexts—albeit contexts that not every rational being finds herself in—many such beliefs might be rationally required (so that one would not be justified in failing to have them). It is worth saying too that apparently, once such beliefs are on-board, particular people may, on the basis of such beliefs, rationally reject beliefs that, absent the new beliefs, would have been rationally required.

Of course, this introduces a serious complication when it comes to appealing to morality's standing as a system of rules acceptable to all. Those who do possess the religious, metaphysical, or scientific beliefs may, as I have suggested, find themselves unable rationally to accept the rules they would favor, were they relying solely on rationally required beliefs. Apparently, then, those who are more informed may find that, for them, accepting the rules is actually irrational. If so, that would undermine once again the claim that the rules would be acceptable to all to whom they apply.

Perhaps, though, if such people realize that morality is essentially an informal public system and that people can rationally fail to have the beliefs they do, then they may have reason themselves to bracket those beliefs in thinking about

what all rational agents might accept. In the name of identifying rules all could understand, they might put aside much of what they know. Perhaps. Nonetheless, rules that would be understood by and acceptable to those who know very little may turn out to be rules that it would be positively irrational for those who know more to accept—even if they know that what makes it irrational for them to accept the rules are considerations not available to those who are more ignorant than they are themselves. The ignorance of others gives them reason to wish the ignorance away, not to ignore what they know.

We seem to face the following situation. On Gert's view, morality is an informal public system that must be both (i) known by, and (ii) rationally acceptable to, all rational agents. Any argument for a system of rules that required an appeal to beliefs not all could understand would fail to establish the rules as the sort of informal public system it would have to be (on Gert's view) in order to constitute morality. Thus, it appears, the justification for a set of rules—if it is to establish them as the rules of morality—must appeal to what each and every rational agent could be expected to understand. Yet, since some rational agents know, and so understand, a great deal more, they may find themselves unable rationally to accept what they could were they more ignorant. Any argument for a system of rules that requires an appeal *solely* to beliefs that all could understand would be an argument that would fail to establish the rules as the sort of informal public system it would have to be (again, on Gert's view) in order to constitute morality, since those who are more knowledgeable would be unable rationally to accept it (even though they would understand it). Either way, the rules will not stand as rationally acceptable to all, and so will not constitute an informal public system that applies to all rational beings. The natural question to raise is: Why would someone think morality really is such a system?

Before turning to this question, let me flag one detail worth further attention: It might well be that one could identify a system of rules (Gert's ten rules, or Kant's single "treat yourself and others as ends and never solely as means", or Mills' "promote the greatest happiness") that are quite understandable to virtually anyone, although a compelling justification for accepting it would require an elaborate appeal to relatively esoteric beliefs that not all rational beings would have access to or understand. This suggests that we should distinguish (i) asking what it takes for a set of rules to be understandable from (ii) asking what it takes for it to be rationally acceptable as a standard for their own behavior, from (iii) asking what it takes to understand what makes it rationally acceptable. The three may be linked in interesting ways, especially if rationally accepting something turns, in important respects, on what a person understands about her situation and options, even when she knows not every rational being would have the same understanding. I will not pursue this issue here.³³

Why Think Morality is an Informal Public System that Applies to All Rational Beings?

Gert's defense of the appeal solely to those beliefs that are rationally required turns crucially on his claim that morality is an informal public system that applies to all rational beings. Why think this is true? There seem to be two sorts of argument offered in *Morality* for this assumption. One highlights the assumption as so obviously true that it doesn't need defense. The other holds that morality plays a distinctive functional role in our social life that can only be played by an informal public system. I will take them in order.

Frequently, especially early on in *Morality*, Gert simply asserts this as “an essential feature” or claims that “Nowadays everyone recognizes” morality’s standing as an informal public system, or maintains that “It is uncontroversial that” it is such a system. Needless to say, every argument needs to rest on some assumption or other that is taken for granted. So an appeal to the obviousness of some assumption must, in more than a few cases, be perfectly appropriate, not least of all if, indeed, the assumption is uncontroversial or at least widely recognized. Nonetheless, this appeal to obviousness is highly dubious in the present case. There are two grounds for doubt, over and above the fact that it ends up playing such a crucial role in defending what is otherwise such an initially un-intuitive limitation on whose approval matters morally.

One ground for doubt can be found in the observation that a number of Gert’s criticisms of competing moral theories have as one crucial element the fact that they fail to recognize morality as an informal public system—appealing as they do to considerations predictably beyond the ken of many rational beings. This is, in effect, one of Gert’s major criticisms of consequentialism: If consequentialism is true then what morality might require of people may be unknown by and even rationally unacceptable to those to whom it applies—an implication, Gert points out, that flies in the face of thinking morality is an informal public system.³⁴ Certain understandings of Kantianism have the same up-shot when they treat application of the categorical imperative as a difficult exercise in practical reasoning.³⁵ For again, what morality actually requires ends up being elusive, to say the least. And certainly familiar versions of Aristotelian ethics have the same implications concerning the extent to which people may, as it happens, fail to know what morality demands of them. All of these substantive moral theories turn out to fly in the face of a putatively essential feature of morality. I infer from this that they, at least implicitly, reject the assumption that morality is essentially an informal public system. That is not to say morality is not such a system. It may be. The point is only that in this context it seems wrong to claim that the assumption is obvious or uncontroversial.

The other ground of doubt can be found in the fact that there are several attractive views that overlap with, yet differ in significant respects from, a view according to which morality is an informal public system that applies to all rational beings. The attractiveness of these alternative views might explain—without vindicating—Gert’s contention that morality is such a system.

For instance, one might hold that morality is a public system that applies only to rational agents that have a fair bit of knowledge (knowledge substantially beyond that had by someone possessing only rationally required beliefs), although it extends its protection to all rational beings. On this view it is inappropriate to apply moral standards to people who have grown up in certain environments or lack knowledge of a certain kind, while it is appropriate to recognize such people as rational agents toward whom one has moral obligations (including, perhaps, obligations to educate). Thus one might agree with Gert that morality is an informal public system, but reject his view that it applies to all rational beings. Or one might think that morality is an informal system of rules (though not a public one) that applies to all rational beings and is such that they all have reason to accept the system, even though that system is not known to them all. On this view, there are compelling reasons to accept morality’s demand though coming genuinely to understand and appreciate morality requires a kind of effort, insight, or education that cannot be expected of everyone. Or one might think that morality is a public system that is known to everyone but is

such that not everyone is rational in accepting its demands. On this view, certain people might find themselves in situations in which it would be irrational for them to accept morality's standards whether or not others might rationally work to enforce the standards.

These alternatives bring out that in thinking about whether it really is obvious that morality is an informal public system that applies to all rational beings, we need to note that one might back off of this crucial claim by rejecting the claim that morality applies to all rational beings, or by rejecting the claim that it is known by all to whom it applies, or by rejecting the claim that all are rational in accepting its demands. Backing off any one of these claims is backing off something that is necessary to the initial argument for embracing the lowest common denominator standard of rational acceptability. I think it is not obvious that these alternatives are all wrong.

All told, then, there seems to me more than a little reason to think that morality's status as an informal public system that applies to all rational beings requires a substantive argument and cannot plausibly be offered as an uncontroversial or obviously true claim. Before turning to the substantive—functionalist—argument Gert might have on offer, I want to take a tour briefly through one of the features of Gert's view that I think is morally important, though puzzling: The requirement that moral standards must be known by those to whom they apply.

This knowledge requirement clearly plays a central role in Gert's system by constituting a condition on something counting as a *public* system. Yet, at least initially, the requirement seems to carry a strange implication. Gert holds that knowing morality's standards is a matter of knowing everything that they require and recommend.³⁶ Anything less and one doesn't yet have the requisite knowledge. At the same time, Gert acknowledges that determining what morality demands of one requires knowing all sorts of facts above and beyond those that would be known by those limited to rationally required beliefs.³⁷ At first this looks like a pretty serious tension. A system doesn't count as an informal public system unless it is known by all to whom it applies, and knowing it is knowing what it requires, yet it requires things in light of facts not known by all. As things turn out, we can alleviate the tension by distinguishing the fundamental demands of morality and the system of standards that constitute morality from morality's derivative demands. The fundamental demands apply to all rational beings and must be known by all to whom they apply, while the derivative demands apply only in particular situations, in light of the more fundamental demands, and they apply on the condition that the facts relevant for the derivation of the demands are known by the relevant people (rather than by all rational people).³⁸

There is, at this point, an interesting twist in the view, for it turns out that, according to Gert, the facts that are relevant to the derivation count as known by the relevant people not because the agent to whom the demands apply knows them, but because they are known "by all qualified people"³⁹—where being a qualified person is sometimes an esoteric qualification. This twist is at least a little peculiar, since it means that people may, after all, be subject to moral demands *they* may not know about or be able to understand. As long as these agents are able to recognize who is appropriately qualified, one might argue that they are, in the relevant sense, able to learn about the demands that apply to them even if they are not able to understand the considerations that make it the case that such demands do come into play. But this seems quite a stretch to the

extent that actual understanding of morality's demands is seen as important to the legitimacy of those demands.

In any case, however that works out, there is a disturbing up-shot of this general view. Sometimes, presumably, who is qualified will be clear to all, and all who are qualified will agree concerning the relevant facts. Yet at other times things will not work out so neatly and either there will be disagreement concerning who is qualified or disagreement among those who are qualified as to the relevant facts. In fact, this sort of disagreement is familiar and more or less rampant. As Gert would have it, under these circumstances, there is no correct judgment concerning the matter at hand and so no derivative demand. "If there is unresolvable disagreement among qualified people, or if there is unresolvable disagreement about who counts as a qualified person, then there is no right answer to the problem."⁴⁰ Thus it follows from the knowledge requirement, as Gert develops it, that if there is irresolvable debate among qualified people about facts relevant to whether abortion is wrong, or whether euthanasia is wrong, or whether capital punishment is wrong, or . . . then it follows that one side of each of these issues is right—the side that holds the practice in question is not morally wrong.

According to Gert, a "morally wrong action is one that all impartial rational persons would favor not doing."⁴¹ In the case of derivative demands, this will require both that all impartial rational persons agree in favoring not doing a certain kind of action and that there is agreement concerning who counts as qualified concerning the facts at issue that underwrite the derivation and that they all agree. Moreover, "what is morally wrong is *always* determined," Gert maintains, "by what all qualified persons would decide at the time of acting or deciding"⁴² which means that there will be no relevant right answer in all the cases in which given merely the current state of information qualified people do not agree.

This generates a general argument—assuming there is the appropriate sort of disagreement either among the qualified or about who is qualified—to the effect that, say, abortion is not wrong. And this general argument might in turn be used to dismiss those who hold that it is wrong, on the grounds that they evidently do not understand the moral system. There is something suspicious going on here.

One might respond to this problem by pointing out that disagreements concerning the relevant facts cannot occur unless there is already substantial agreement concerning the fundamental rules in light of which the facts are relevant. What disagreement there is, then, will concern not the rules themselves but only their scope (do fetuses count as people?), or about when exceptions are justified (are they justified when a mother's life is endangered?), or about the interpretation of the rules as they apply in particular circumstances (is her life endangered?), and none of these, someone might hold, call the rules themselves into question. Would insulating the rules from disagreement in this way help to avoid the peculiar result that in controversial cases the right answer is always that there is nothing wrong? I don't myself see how, since the line of argument that generates this result turns not on whether the disagreement is ultimately a disagreement about the rules themselves but only on the result of what disagreement there is being that those who count as qualified do not agree or people do not agree as to who is qualified.

In any case, it is not at all clear that there is a sharp difference between disagreeing about the rules, on the one hand, and disagreeing about their scope, or

about justified exceptions, or about their interpretation, on the other. If there is no sharp distinction to be drawn, then presumably it cannot be relied upon to avoid the peculiar consequences of holding that the fundamental principles establish derivative duties only in contexts where there is no reasonable disagreement concerning the relevant facts.

There are, then, two worries I have about the knowledge requirement, as Gert ends up developing it. First, it seems seriously problematic if the knowledge requirement has it turn out that when there is either disagreement concerning who is qualified when it comes to determining the relevant facts, or disagreement among those who are qualified about those facts, there is no relevant right answer and so no derivative moral demand that turns on the answer. In such cases, I think we ought to see the debates at stake as ones about which each side is committed to thinking there is an answer (perhaps even committed to thinking there might be some procedure by which the disagreement could, at least in principle, be resolved) but not committed to thinking that in cases of such disagreement those who think something wrong was being done are simply mistaken. To acknowledge our ignorance concerning relevant facts should not of itself resolve the moral debate in favor of those who claim there is nothing wrong. Second, it doesn't seem that one can successfully defend the theory by appealing to the distinction between disagreements concerning the rules and disagreements about, say, their scope, or legitimate exceptions, or their interpretations. Even if one grants the distinction, it doesn't seem to get around the untoward consequences. The distinction would allow us to treat some disagreements as disagreements not about the rules but about their scope (or whatever), but pretty clearly some serious disagreements would concern the rules themselves and others would concern the facts that the rules would render relevant, and in these cases the knowledge requirement (as Gert develops it) would still entail the implausible claim that the very fact of disagreement itself establishes that those who are arguing that some action is morally wrong are mistaken.

Let me get back to the question: Why think that morality is an informal public system that applies to all rational beings? As I have suggested, this claim is not plausibly advanced as so obvious and uncontroversial as to need no defense. Gert does however suggest an attractive line of thought that might seem to recommend the claim. It starts by noting that morality, whatever else it might be, is a system of rules or principles that works to reduce harms and (perhaps) facilitate peace.⁴³ Thus morality performs a crucial function in our public life—most especially, it lets us substitute argument for arms in contexts where otherwise conflict would likely emerge. This role looks to be one that a system of rules could play, but only if the rules were widely known by those to whom they applied. Moreover, it could play the role only if, in addition to being known, it was rationally acceptable to those same people. So, one might think, the very function of morality could be fulfilled only by a system that was a public system—a system known to and rationally accepted by those governed by it.

Now my own sense is that this argument depends on false empirical claims. I see no reason to think that the system could be one of the best solutions to the problem of reducing harms and facilitating peace only if everyone to whom it applies could rationally accept it. The legal system provides a relevant example. It sure seems as if there are people subject to its rules who have, themselves, no reason to accept it. Nonetheless, it is a system that is better than many alternatives when it comes to regulating interactions among people in a way that reduces harms and facilitates peace.⁴⁴ At the same time, it seems that a system of

rules might fulfill the crucial function even if it was not known by all to whom it applies—say because its requirements were available only on careful reflection of a sort that people do not often engage in (even as, perhaps, it encourages attempts to get people to reflect appropriately so that they could see for themselves the reasons they have to act as they are required to). So, even if one holds that morality should be seen as a system of rules that plays a crucial role in our social life, it is not at all clear that the only systems that might play that role well are public systems.

Moreover, it seems that which system of rules could effectively play the role might well vary dramatically depending on facts about the people to be governed by the system. So, one worry is that if the moral system is designed especially to be acceptable to people who do not know things, people who do know a lot are going to find themselves in the awkward business—because of their commitment to morality (or impartiality)—of emphasizing rules even as, having more information, they themselves favor some other rules. They don't favor these other rules as a moral system because they would constitute a moral system only if they could be a public system. But they favor people acting in another way—a way that does not accord with the moral rules they are constrained—by the ignorance of others—to offer. This is a framework that risks being designed so that rationally informed people all the time have reason to violate (unless they happen to be committed to morality, or to regulating their behavior by rules that are impartial in Gert's sense). The upshot is that if one emphasizes the functional role of a system of rules it may end up that the systems of rules that work well will differ dramatically from community to community as the nature and extent of the knowledge of the people changes. And just as the Bauhaus theory of design holds that what makes a chair, for instance, a good chair, depends on the structure of the bodies that will use it, so too a good moral theory, on this view, should be structured to fit the body public (even if the theory is not public) in ways that accommodate it.

As I have said, I am puzzled by the knowledge requirement. Yet I am pretty convinced that something like a knowability requirement has got to be right. If an account of morality makes morality in principle unknowable, then it makes it in principle uninteresting. And if it is in principle knowable, then there is an issue about to whom it needs to be knowable and how. Moreover, it certainly seems (as Gert emphasizes) that it had better be knowable to all who are held responsible for acting according to its demands. After all, if one relies on the standards of morality as grounds for condemning people or punishing them, and those people couldn't have known about the standards (and so known they were violating them), then there is something seriously wrong—morally wrong—with condemning and punishing them. As a result, any account of morality that allowed as permissible such a practice would be, it seems clear, misguided. My sense, though, is that it would be useful here to distinguish two distinct roles knowability might play: It might serve as a constraint on what counts as morality (as Gert proposes) or it might serve as a substantive moral constraint. In the first role, the fact that some standard is unknowable to some rational people would establish that it does not constitute a moral standard. In the second role, in contrast, the fact that it is unknowable to some rational people would leave the standard's status as a moral standard unchallenged, even as it meant that such people could not legitimately be held responsible for violating it.

Any plausible moral theory, it seems, will need to include a prohibition (at least under normal conditions) on blaming, punishing, and holding people re-

sponsible for violating norms they justifiably didn't know about. Still, such a prohibition does not constitute a constraint on what might count as a moral system. Instead, it serves as a constraint, within the moral system, on various ways of reacting to violations of the system's requirements. Indeed, in recognizing it as a constraint on praise and blame, etc. *in the face of violations of morality's standards*, one is implicitly recognizing the standards as genuine despite their being unknown by at least some to whom they apply. Thus in accepting this substantive constraint one might, and may even have to, acknowledge that the standards in question—that were, by assumption, not known—are nonetheless standards of morality. In any case, I suspect that a lot of what might seem to recommend thinking of morality that it is a public system comes from the moral force of this substantive principle. A lot of people might embrace the substantive constraint without any commitment to the constraint on moral systems. And they may even find that the appeal of this substantive principle actually stands in the way of thinking that all the rules of morality must be known by those to whom they apply.

Notes

1. Thanks are due to Walter Sinnott-Armstrong and Bernard Gert for suggestions on earlier drafts of this paper and for greatly appreciated help in working through the subtleties of Gert's moral theory.

2. Bernard Gert, *Morality: Its Nature and Justification* (New York: Oxford University Press, 1998), 316.

3. "It is . . . an essential feature of morality in all of its variations that everyone who is judged by it knows what morality prohibits, requires, encourages, and allows." As a result, "[s]howing that a proposed account of morality contains some part that is justifiably unknown to any people about whom moral judgments are made shows that the proposed account of morality is inadequate." (Gert, *Morality*, 6).

4. I should note in passing that Gert is not requiring that all rational people have positive reason to accept the system of rules he identifies. He only insists that the adequacy of the rules depends upon agents being able rationally to accept them. The difference is real, on his account, because rational acceptance is not a function of the weight of the reasons a person might have one way or the other. For the purposes of this paper, however, one could just as well hold that, to be adequate, an account of morality must offer rules or standards that all rational beings actually have reason, on balance, to accept, since, as it happens, accepting rules that in any way limit one's freedom, is (on Gert's view) rational only if one has adequate (positive) reason. In other contexts, though, Gert is concerned to emphasize that one may rationally do some things for which one has no reason and also that one can, in some contexts, be acting rationally despite acting contrary to what one has most reason to do. For instance, he thinks it is always rational to act in one's own interest no matter what reasons one might have for acting contrary to one's interest, even as he also thinks that when one has strong reasons to act contrary to one's interests one may rationally do so as well.

5. Gert, *Morality*, 6. See also 10. Later he claims that "Satisfying both of these requirements is necessary and sufficient for considering these rules as moral rules, that is, as part of a public system that applies to all rational persons." *Morality*, 168.

6. It is useful to note the differences between Gert's proposal—that a system of rules is a candidate moral system only if the rules could be rationally accepted—and Thomas Scanlon's suggestion that the system must be such that "no one could reasonably

reject" it. [See Scanlon's "Contractualism and Utilitarianism," in *Utilitarianism and Beyond*, ed. Amartya Sen and Bernard Williams (New York: Cambridge University Press, 1981), 103-28].

7. Significantly, on Gert's version of this approach, not all fully rational beings need favor the same things. Thus the proposal is not that a person is rational in doing (or accepting, or choosing) something if and only if *all* fully rational people, appropriately situated, would favor it, but only if and only if *a* fully rational person would. For some accounts this difference disappears thanks to the specific characterizations offered of what it is to be fully rational and to be appropriately situated. However, on Gert's view, even once these are specified, there is room for fully rational beings so situated to differ in what they favor. One is behaving rationally, Gert's thought is, as long as one has at least one fully rational being, appropriately situated, favoring one's behaving in that way. That others might not favor it reflects a difference among rational agents, but it doesn't impugn one's rationality.

8. It is worth distinguishing the question of whether it would be rational for one to accept a set of rules as a limitation on, and guide for, others' behavior from the question of whether it would be rational for one to accept the set as a limitation on, and guide for, one's own behavior. Potentially, a set of rules that would be rationally acceptable for one role would not be so for the other. See chapter 7 of *Morality*, where this difference in role is especially important.

9. See, for example, Adam Smith's *A Theory of the Moral Sentiments* (New York: Cambridge University Press, 2001), and Roderick Firth's "Ethical Absolutism and the Ideal Observer," *Philosophy and Phenomenological Research* 12 (1952): 317-45.

10. See, for example, Richard Brandt's *A Theory of the Good and the Right* (Oxford: Clarendon, 1984), and Michael Smith's *The Moral Problem* (Oxford: Blackwell, 1994).

11. See John C. Harsanyi's *Essays on Ethics, Social Behavior, and Scientific Explanation* (Dordrecht: Reidel, 1976) and John Rawls' *A Theory of Justice* (Cambridge: Harvard University Press, 1971).

12. *Morality*, 36. He does acknowledge, though, that under some circumstances even the rationally required might without irrationality be rejected, but not under normal circumstances. This leads me to think that perhaps the "rationally required beliefs" Gert is after are best characterized as beliefs every rational agent holds, absent contrary evidence.

13. Such beliefs are, according to Gert, completely uncontroversial precisely because "they are the only beliefs that no rational person doubts." *Morality*, 37.

14. *Morality*, 37. See also 158 and 167. Interestingly, Gert counts among the rationally required beliefs things like "some moral agents care about some sentient beings" (fn., p. 54). It is a little unclear to me whether Gert is assuming that all rational beings must have the concept of morality. The example makes it look as if he is, since without it they could not have the above belief. Nor is it clear whether, absent a fairly thick characterization of what it takes to count as a moral agent, the belief would be uncontroversial. After all, more than a few have claimed that no one genuinely cares about others. (Whether they would say that means there are no moral agents, or only that the moral agents there are do not, after all, care about others, will turn on what is built into the concept of a moral agent.) I agree that the view is implausible, but it doesn't appear to be one that a person would have to be irrational to accept it, except when in the possession of special evidence. On the face of it, absent evidence one way or the other the belief that

no one cares about others seems roughly on a footing with the belief that some do care about others (although, of course, the latter is a bit weaker thanks to the “some”).

15. *Morality*, 36.

16. *Morality*, 170

17. *Morality*, 44-45. According to Gert, a desire is rational if a person would be irrational were she to lack it and a desire is irrational if it would be irrational to act on it absent adequate reason. I should emphasize that in calling a desire irrational Gert is not maintaining that it is always, or even usually, irrational to act on it—only that in order for it to be rational to act on it, one must have adequate reason. Other desires, Gert holds, are such that one needs no reason in order for one’s acting on the desire to be rational. He is not, I think, holding that these not irrational desires actually themselves provide reason, only that, absent countervailing reason, acting on them is not irrational. In acting on them one may not be acting with a reason at all, even though, in so acting, one is not being irrational.

18. *Morality*, 161.

19. *Morality*, 167.

20. *Morality*, 167.

21. *Morality*, 168-69.

22. *Morality*, 170.

23. A system is shown to be weakly justified, in contrast, if it would be favored by *some* rational beings appropriately situated. See *Morality*, 18. According to Gert, that some would favor a system shows that it is not irrational to accept it; that all would favor it shows that it would be irrational not to accept it.

24. *Morality*, 167. And even the beliefs that are rationally required (absent contrary evidence) might well be rationally rejected by—and so have no influence on—those who have adequate contrary evidence.

25. In speaking of what they would rationally favor, I am imagining that they are relying solely on those beliefs of theirs that are rationally permitted (many of which will be rationally required of them, given their circumstances, even though those same beliefs would not be required of others, and so do not count as beliefs all rational agents are required to have.)

26. *Morality*, 202. Gert “is not advocating an end to all civil disobedience.” But he does hold that civil disobedience “is only justified when one has some reason to believe that disobeying the law will do something toward lessening that evil [which the law is causing].”

27. *Morality*, 202.

28. *Morality*, 167.

29. It is worth noting, however, that scientific beliefs—say that humans evolved from other mammals, or resemble non-human animals in many ways, or that people suffer psychological disabilities of various kinds—might significantly alter the sort of rules rational beings would accept or favor as well. Especially when we are asking what might cause a rational being to favor one thing or another we seem to be asking a question that, absent some quite substantive assumptions by the psychological impact of various sorts of beliefs cannot be answered in the abstract.

30. The rules Gert proposes all, as I have mentioned, admit of exceptions. So it is important here to mark the difference between embracing rules along the lines of those proposed by Gert and, in effect, abandoning the rules as independent standards of behavior—holding them merely as derivative rules of thumb, or summary rules, that are useful

to promulgate. Someone who accepts utilitarianism may well find the rules Gert proposes roughly acceptable, but only as they have a utilitarian backing.

31. This follows directly from Gert's definition of informal public systems.

32. *Morality*, 144.

33. But it is worth noting that it raises an interesting issue about the sort of rational acceptance that might be required for a system to count as an informal public system. If a person counts as rationally accepting some set of rules as long as someone differently situated would favor accepting the rules, then a person's current commitments and what she happens to believe might only tenuously be connected to what it is rational for her to accept. Just how tenuously will depend upon the extent to which the favor of those who set the standard is sensitive to a person's actual situation, commitments, and beliefs.

34. G. E. Moore speaks to this view directly, if not all that persuasively, when he writes, "It might be urged, with more plausibility, that we mean by a man's duty only the best of those actions of which he might have thought. And it is true that we do not blame any man very severely for omitting an action of which, as we say, 'he could not be expected to think.' But even here it is plain that we recognise a distinction between what he might have done and what he *might* have thought of doing: We regard it as a pity that he did not do otherwise. And 'duty' is certainly used in such a sense, that it would be a contradiction in terms to say it was a pity that a man did his duty." *Principia Ethica* (Cambridge: Cambridge University Press, 1903), 150-51.

35. Kant himself seems, with Gert, to think the epistemic accessibility of morality's demands—even if not the justification of those demands—is crucial to their legitimacy. Thus, in the *Critique of Practical Reason*, he argues that "The moral law commands the most unhesitating obedience from everyone; consequently, the decision as to what is to be done in accordance with it must not be so difficult that even the commonest and most unpracticed understanding without any worldly prudence should go wrong in making it." Kant, *Critique of Practical Reason*, trans. Lewis White Beck (Indianapolis: Liberal Arts Press, Bobbs-Merrill, 1956), 36.

36. See, *Morality*, 4, 7, 10, 26, and 150.

37. See *Morality*, 227 ff. and 236-37.

38. Alternatively, instead of distinguishing in this way between fundamental and derivative demands, treating only the former as part of the informal public system, one might hold that the supposedly derivative duties are actually conditional duties that, no less than the fundamental duties, must be known by all rational beings, even if they might not know whether the conditions are satisfied. One worry about this suggestion is that the conditional duties will require, for their articulation, the deployment of concepts that not all rational beings need possess. So there will remain a set of duties that are not knowable by all rational beings. See *Morality*, 151.

39. *Morality*, 324.

40. *Morality*, 324. When the problems are practical, as opposed to theoretical, the conditions on there being a right answer are often even more stringent: All qualified people must agree given only what is known at the time of decision or action. See 325-26.

41. *Morality*, 325.

42. *Morality*, 326.

43. This is a feature that Gert himself highlights in characterizing the distinctive nature of morality. See *Morality*, 13.

44. Of course I am not suggesting that the legal system is a kind of morality, only that it is a system of rules that has some claim to fulfilling the role the argument in question maintains can only be fulfilled by a public system, which the legal system is not.