

Epistemic Conditions for Moral Responsibility and Information Manipulation

Todd R. Long

1. Introduction. In previous work, I (Long 2004) have argued that we can be morally responsible for our actions in cases in which we deliberate or engage in practical reasoning on the basis of manipulated information; and, this is so, I have argued, even if the information manipulation is so radical that the agent would not have performed the relevant action had the manipulation not occurred. In so arguing, I have focused primarily on the *metaphysical conditions* for moral responsibility. In this paper, I reveal what I take to be important *epistemic conditions* for moral responsibility and then to apply them to cases, including those involving information manipulation. My goal is advance our knowledge about moral responsibility and manipulation in a way that yields philosophically satisfying responses to some criticisms of my previous work.¹

In section 2, I provide two cases to anchor subsequent discussion of my claim that information manipulation does not preclude moral responsibility. In section 3, I entertain and reject an initially appealing objection, according to which relying on false information prevents moral responsibility. My response to the objection reveals the need to clarify terms; so, in section 4, I argue that ‘morally responsible’ is a term of art, I explain and defend my use of the term, and I offer an account of moral responsibility, which makes use of Aristotle’s epistemic conditions for voluntariness, and I explain what I take to be the relation between moral responsibility and praiseworthiness/blameworthiness. In section 5, I provide test cases for my account of moral responsibility. I argue that, using voluntariness as a guide to moral responsibility, we can plausibly explain all that needs explaining about the relation between moral responsibility and praiseworthiness/blameworthiness by appealing to nothing more than the relevant metaphysical and epistemic facts that obtain at the time of a person’s basic action. In section 6, I apply results from previous sections to cases of information manipulation, and I show that there are good reasons to resist criticism of my view: information manipulation does not preclude moral responsibility.

¹ For criticism of Long (2004), see Fischer (2010), Franklin (2006), Nahmias (2005), and Steadman (forthcoming).

2. The Family Policy Cases. Consider the following case (which I will refer to as *Family Policy*): You live in a suburban subdivided neighborhood, and you are the father or mother of a four-year-old daughter. Your neighbors, who also have a four-year-old daughter, have been good friends during the past two years. The two girls have been accustomed to playing together, and they have, with their parents' permission, freely moved between the two houses during afternoon playtime. Sadly, your neighbors have had to move away. One morning a police officer comes to your door and explains that a convicted sex offender will begin renting the neighboring house in one week. The police officer gives you all the details that she is legally allowed to divulge, including the claim that a police officer will drive by the street daily and a parole officer will check in weekly with the sex offender. Dismayed by the news, your spouse and you deliberate about what to do about the situation. You realize that your financial and work situation will not allow you to move right away. After careful, prolonged deliberation, you enact a family policy according to which your daughter will henceforth not be allowed to go outside of the house for any reason without overt parental supervision. I take it that your action of enacting this family policy is an uncontroversial example of an action that you are morally responsible for. After all, you do not engage in the action as a result of a psychological compulsion, a drug addiction, or any other kind of obvious responsibility precluding condition. It is a rational decision based on careful, reflective deliberation and it is consonant with the kinds of preferences, desires, and values that make up your character. I intend for this example to strike all non-skeptics about moral responsibility as a paradigm example of morally responsible action.²

Now consider a variation (which I will refer to as *Manipulated Family Policy*): This new example is exactly like the first, except for the following facts: the person who comes to your house and explains that a sex offender will be moving next door is not a police officer. She is a fake who has been paid by an enemy of yours to impersonate a police officer for the purpose of causing you emotional distress. However, you have no reason at all to think that she is a fake. Indeed, she is actually wearing a police

² Feel free to add whatever details are needed for your favorite theory of moral responsibility to entail that this action is a morally responsible one, and add whatever details you think are required for the action to be one for which you are praiseworthy.

uniform issued by your city, and her conversation with you is completely convincing. To focus on the key issue, suppose that the fake police officer says exactly the same words, with the same inflection and body language, as does the actual police officer in *Family Policy*. She is such an excellent actor that only technology such as a lie detector test would reveal her deception. In this *Manipulated Family Policy* example, you deliberate just as you deliberated in *Family Policy*, and the outcome of your deliberation includes your enacting a family policy of not allowing your daughter to leave the house without overt parental supervision.

Are you morally responsible for enacting the family policy in this second example? I think that you are. After all, you do not enact the policy as a result of a psychological compulsion, or a drug addiction, or a neurological disorder, or any other obvious responsibility precluding condition. In both cases, you deliberate on the basis of information that you have non-culpably acquired and which you have excellent reason to believe is true. In both cases, your action is consonant with the desires, preferences, and values that make up your character. In both cases, the way you deliberate to your decision is the same. Thus, your way of deliberating and acting on the basis of that deliberation is not manipulated at all. All that is manipulated is the information that you have to go on in your deliberation. In my view, *information manipulation does not, by itself, preclude morally responsible action*. An interesting consequence of this view is that *we can be morally responsible for actions that we would not have performed if the information manipulation had not occurred*.

3. No-Falsehood Objections. I think it is pretty obvious that you are morally responsible for enacting the family policy in *Manipulated Family Policy*. However, a number of philosophers have resisted my view. As I explain in section 4, it is very likely that much of the resistance depends entirely on differing assumptions about application of the term ‘morally responsible’, which is a term of art. Nevertheless, there are examples in the literature that either explicitly state or suggest an objection to my view, regardless of how we plausibly understand ‘morally responsible’. The idea of this general objection may be categorized as an epistemic matter understood along the lines of what I call the ‘No Falsehood’ principle:

No Falsehood: if one deliberates on the basis of *false* information, then one is not responsible for events that result from that deliberation.

Variations on *No Falsehood* have appeared in the literature. For instance, James Steadman (forthcoming) has us consider a case in which “Linda votes for McCain instead of Obama, because she possesses a number of false beliefs about where he stands on the issues (perhaps there was a misprint in her local newspaper, so that remarks by Obama were inadvertently attributed to McCain)”. Steadman’s idea is that, since Linda engages in practical reasoning on the basis of many false beliefs, we should think that the psychological mechanism constituting her practical reasoning process is defective enough that “we would also not hesitate to absolve Linda of any responsibility for her (deluded) choice”.³

Perhaps your deliberation in *Manipulated Family Policy* is based on so many false beliefs (e.g., a police officer testified to you, a sex offender will soon move next door, your daughter is in significant danger if your current family policy endures, etc.) that your choice to enact the family policy counts as what Steadman calls a ‘deluded choice’ such that we would not hesitate to absolve you of any responsibility for it. But, if that is so, that fact would no more show that you lack moral responsibility for enacting the family policy than it would show that Linda lacks moral responsibility for voting for McCain. Note first that it would be possible to *absolve* someone of responsibility for an action only if the person *were responsible* for the action. This just follows from the fact that absolution is the act of freeing one from an obligation or of remitting a sin. Just as I must have an obligation in order to be freed from it, so must I have committed a sin for it to be remitted. But, perhaps by ‘absolved’ Steadman intended something like, ‘we wouldn’t *blame* Linda under the circumstances if we knew about them’ (any more than we would think you were blameworthy in *Manipulated Family Policy*). It may well be true that neither Linda nor you are blameworthy for your respective actions. Indeed, given additional epistemic

³ These claims from Steadman are in the context of a discussion of Fischer and Ravizza’s mechanism-based, moderate reasons-responsiveness theory of moral responsibility. Steadman says that his conclusion about absolving Linda of responsibility also depends on what he says is “the (plausible) assumption that such a defect would make the mechanism unresponsive to sufficient reasons to act otherwise (in alternative scenarios discussed by Fischer and Ravizza)”. Nevertheless, as I explain later in this section, Steadman makes clear that he thinks that an agent’s engaging in practical reasoning on the basis of an abnormal number of false beliefs can preclude moral responsibility. For details of the theory under discussion, see Fischer and Ravizza (1998).

details I assume about these cases, they are cases in which the agents are not blameworthy. It doesn't follow that they are not morally responsible for their actions. After all, it is extremely plausible that, if you are morally praiseworthy for your action in *Family Policy*, then you are morally praiseworthy in *Manipulated Family Policy*; for we may suppose that in *each case* you did what the most morally upstanding, perfectly rational person would do in light of the information you had to go on. But, of course, you are morally praiseworthy for an action only if you are morally responsible for that action.

Recall that my reason for thinking that you are morally responsible in *Manipulated Family Policy* is that there is no relevant difference between that case and *Family Policy*: in each case your way of deliberating or engaging in practical reasoning is exactly the same; the information you utilize in each case is, as John Martin Fischer (2010) would put it, filtered through your own 'normative orientation' (i.e., your standing desires, preferences, and values). Fischer's idea, roughly, is that when a person's new information about the world is filtered through that person's (more or less) sane normative orientation, then actions that flow from that process are related in the right sort of way to the person such that it is correct to think of those actions as being under the *voluntary control* of the person; thus, such actions are plausible candidates for actions for which the person is morally responsible. Although imprecise, Fischer's idea is useful enough for present purposes. Consider Linda's case: remarks by Obama were inadvertently attributed to McCain in her newspaper, and this false information was filtered through her normative orientation just as (we may suppose) the true information would have been filtered through her normative orientation had the newspaper correctly attributed the remarks to Obama. Being an ordinary adult, Linda is the kind of person who has some general principles, desires, preferences, and values. The false information that she relies on in her practical reasoning either does not alter her normative orientation at all (the issue depends on the details of what is meant by 'normative orientation'), or, if it does alter her normative orientation, then the crucial point to realize is that the true information, had she gained it, might also have altered her normative orientation. But no non-skeptic about moral responsibility is worried about cases (such as *Family Policy*) in which the agent is deliberating on the basis of true information; so, why worry about cases in which the agent is deliberating on the basis of

false information? Either the new information we receive about the world (via our experiences) does not change our normative orientation, or it does change our normative orientation. If it *does not* change our normative orientation, then the falsity (or truth) of the new information does not, by itself, preclude moral responsibility for actions that flow from the process of filtering our new information through our normative orientations. If our new information *does* change our normative orientation, then, since new information can change our normative orientation whether the information is true or false, it is not the *falsity* of the new information that (if anything) precludes moral responsibility.⁴

Whence come self-proclaimed ‘intuitive judgments’ that persons such as Linda and you (in *Manipulated Family Policy*) are not morally responsible for their actions under their kinds of circumstances? I suspect various motivations. I have already mentioned one: if we knew the details in a case of information manipulation, we would be willing to *absolve* the agent of moral responsibility. One might infer that, if we would be willing to absolve the agent of moral responsibility, then the person was not morally responsible for the action. But, as I have noted, we could absolve the person of responsibility only if the person were responsible. Moreover, suppose that in a criminal trial, the jury, deciding in full compliance with the law, convicts an innocent person of burglary. *Perhaps* it is true that if we knew, say, that the prosecutor had knowingly withheld from the court incriminating evidence he had about a different suspect, we would be willing to absolve each jury member of that member’s vote for a guilty verdict. It doesn’t follow that the jury members lacked moral responsibility for the vote they made in light of their reasons. To highlight the point, suppose that John, a jury member, voted for a conviction, and suppose that John’s main motivation was that he so disliked the defendant’s demeanor during the trial that John wanted to see the defendant suffer. Surely such a jury member could be morally responsible for his voting in this case, despite the information manipulation.

Others may be motivated by an inference from no-moral-culpability to no-moral-responsibility. For instance, Christopher Evan Franklin (2006) says about a case in which Judith acts on the basis of evidence induced in Judith by a manipulator: “Such a situation would seem to render Judith more pitiable than

⁴ I further discuss this point in section 6.

morally culpable” (187), a judgment that is supposed to provide the intuitive motivation to think that Judith is not morally responsible for her action. It is certainly true that in many (but certainly not all) cases of information manipulation in which the agent does something on the basis of that information that she would not have done had the manipulation not occurred, we would not, if we knew about the manipulation, think that the agent was morally *culpable* for the action, since ‘culpable’ means ‘deserving blame’. And, it is certainly true that in many such cases, if we knew about the manipulation we would pity the relevant agent. But, it does not follow that the agent lacks moral responsibility for what she does in light of the information she has. After all, in some cases of information manipulation in which we would pity the agent because of the manipulation (e.g., *Manipulated Family Policy*), the agent plausibly deserves *praise* for what she does in light of her information; but, she can deserve praise only if she is responsible for what she does.

Consider now a legal motivation for *No Falsehood*. In discussing the voluntariness required for legal consent for a contract, Joel Feinberg (1986) says that the voluntariness of one’s agreement is always lessened by an induced false belief that is germane to the agreement. Of a case in which a seller induces a false belief about a product in a buyer, Feinberg writes: “Depending on how vital a role the belief in question plays in the buyer’s motivation, its falsity will diminish to a proportional degree the voluntariness of his consent” (Feinberg, 274). Imagine a case in which the seller of a house lies to a buyer about very important features of the house (say, features of such great importance that the buyer would not have bought the house had the seller not lied about those features). We intuitively think that the fraud involved in the case is severe enough that the buyer deserves significant compensation, and we have a legal system that provides mechanisms for such compensation to be sought and dispensed. And we have noted reasons to think that a renowned philosopher of law would describe the buyer’s consent to the deal as being so low on the spectrum of voluntariness that it falls below what ought to be the legal standard for lawful consent. One might easily infer that the buyer did not voluntarily sign the home ownership papers; and, because involuntary action is a sign of an action for which one is not morally responsible, one might easily infer that such a case reveals a reason to think that acting on the basis of

false information can preclude moral responsibility for the action. Such a thought might well motivate a view along the lines of the *No Falsehood* principle.

However, reflection reveals that Feinberg is not using ‘voluntariness’ (and its cognates) in a way that is directly relevant for assessing claims of moral responsibility. The relevant question for us is whether the buyer is morally responsible for buying the house. Suppose that the buyer has excellent reasons to think that, if he buys the house, he will thereby achieve his selfish desire to force his wife to get a well-paying job against her will and despite her deeply held, epistemically rational belief that her staying home with her young children is best for the family. I take it that, regardless of whether the buyer is duped by the seller’s lies about the condition of the house, the buyer is *morally blameworthy* for signing the ownership papers and thus for buying the house (if you think additional details about the story are needed for the buyer to be blameworthy, add them). But, he is morally blameworthy for buying the house only if he is *morally responsible* for buying the house; and, he is morally responsible for buying the house only if he *voluntarily* bought the house (in whatever sense of ‘voluntariness’ pertains to moral responsibility). Clearly, the notion of voluntariness we assume when we think about the buyer’s blameworthiness is a different notion than the one that Feinberg says pertains to legal consent. It remains very plausible that the fraud perpetrated by the seller is severe enough to justify legal compensation to the buyer. If we want to explain this justification by means of a concept expressed by phrases such as ‘the voluntariness required for legal consent’, then so be it. But, the example shows that it would be a mistake to suppose that legally useful senses of ‘voluntariness’ and ‘responsibility’ are directly relevant for assessing claims of moral responsibility.

Whatever the sources of various self-proclaimed intuitive judgments about ‘moral responsibility’, I think that non-skeptics about moral responsibility should resist any principle along the lines of *No Falsehood* for the general reason that it cuts against their non-skepticism. Jury trials provide cases in point. It is no secret that criminal trial lawyers, both for the defense and for the prosecution, are trained to provide misleading information, designed to manipulate jurors, as it serves their clients. The goal of each side’s lawyer is to spin the trial evidence in order to get a desired jury judgment. In closing statements,

the prosecution routinely tells the jury that the defendant is guilty, and the defense routinely tells the jury that the defendant is innocent. It follows that one of them tells the jury a falsehood of great importance. Thus, criminal trial jurors routinely deliberate on the basis of manipulated, false information. Does it follow that jurors routinely lack moral responsibility for their judgments? I don't think that any non-skeptic about moral responsibility thinks so. In typical cases, jurors deliberate on the basis of the information they receive at the trial and the judge's orders, doing their best to make a fair, just judgment. Their judgments are typically not the result of psychological compulsions, physical addictions, irresistible impulses, or the like. They are typically the result of more or less rational deliberation, and thus bear the hallmark of voluntary, responsible action.

There is another reason for non-skeptics about moral responsibility to deny a principle along the line of *No Falsehood*. It is epistemically possible that a radical skeptical scenario obtains. Perhaps we are in the Matrix, or we are brains-in-vats, or a Cartesian evil genius is radically deceiving us. If any such skeptical scenario were to obtain, then we would almost always decide to act on the basis of false information. However, even if we are in such a scenario so that we often deliberate on the basis of false or manipulated information, that fact itself is no more a reason to deny that we are morally responsible for our decisions than it is to think that we would lack epistemic justification for our typical beliefs if we were in a radical skeptical scenario (the latter is a lesson from what epistemologists call the *New Evil Demon Problem* for reliabilism).⁵

⁵ Suppose that there is world with agents who are introspectively identical to all of us but whose information about the world is caused by a Cartesian evil demon. Because it is plausible that our evidence typically indicates that the world is more or less the way we take it to be, we intuitively think that our common-sense beliefs about the physical world are epistemically justified. But, by hypothesis, our demon-world twins have the same evidence that we have; thus, they have the same epistemically justified beliefs that we have, despite the fact that almost all of their beliefs are false. The intuitive judgment is that misleading information does not preclude epistemically justified beliefs: one can be epistemically justified in believing propositions on the basis of misleading information. Although beliefs are not actions, I can think of no principled reason for holding that one can have an epistemically justified belief on the basis of false information but one cannot perform a morally responsible action on the basis of false information. Indeed, reflection on the example of misleading information is instructive. We are in no position to tell how often our information about the world is misleading, but it is always epistemically rational to believe what our evidence indicates. Likewise, we are in no position to tell how often we deliberate on the basis of false information. It is epistemically possible that we always (or often) deliberate on the basis of false information: that is one of the lessons of the skeptical scenarios; but, why suppose that skeptical scenarios would preclude morally responsible action?

This argument from the possibility of a radical skeptical scenario is germane to a principle in the literature that is similar to *No-Falsehood*. Recall Steadman’s central idea discussed above. Utilizing Fischer and Ravizza’s mechanism-based understanding of the kind of practical reasoning that can produce actions for which the agent is morally responsible, Steadman (forthcoming) says that a “mechanism [that] has too many false beliefs about the choice at hand” is “a deluded mechanism”, which is to say that it has a particular kind of defect or abnormality that renders it incapable of yielding decisions and actions that the agent is morally responsible for. Steadman’s idea is that we can distinguish normal practical reasoning mechanisms from abnormal/defective ones in the following way:

... we can say that a practical-reasoning mechanism is normal (at least with respect to its rational reasons for action) if the number of occurrent belief-inputs relevant to the choice at hand (both true and false) is at least half the number that occurs in the life of an agent on normal occasions of action. Given this, we can say that the mechanism is abnormal when the number of such inputs falls either above or below the normal range.... We can say that a mechanism ... containing a number of *false* beliefs *above* this range is abnormally *delusive*.

This idea assumes the following *No-Falsehood*-like principle: *if the number of false beliefs involved in one’s practical reasoning is abnormally high, then one is not morally responsible for actions that are outputs of that practical reasoning*. Although it is certainly worthwhile to look for ways to distinguish between responsibility-producing psychological mechanisms and responsibility-precluding psychological mechanisms, reflection on the argument from the possibility of a radical skeptical scenario gives us a good reason to deny this *No-Falsehood*-like principle. For if we happen to be in a radical skeptical scenario, then we routinely engage in practical reasoning on the basis of a massive number of false beliefs. But, that fact would not prevent our being morally responsible for what we do in light of the misleading information we do have. Consider Neo in the movie *The Matrix*. Before Neo is removed from the Matrix, virtually all of his decisions are made on the basis of a massive amount of false information. On Steadman’s way of individuating mechanisms, Neo’s mechanism of practical reasoning almost always contains a super high level of false beliefs relative to what we take our typical practical reasoning mechanisms to contain. Thus, Neo’s mechanism is, according to Steadman’s idea, severely

abnormal and defective.⁶ But, I find no reason to suppose that Neo is not morally responsible for the voluntary decisions he makes in light of the misleading information he has to go on. Of course, if the Matrix is directly causing Neo's normative orientation (i.e., his standing desires, preferences, and values), then it is plausible that he is not morally responsible for the actions that flow from his practical reasoning; but, the mere fact that he engages in practical reasoning on the basis of false beliefs, even if the number is highly abnormal, does not, by itself prevent his being morally responsible for the decisions he voluntarily makes in light of the information he has to go on.

4. On The Term 'Morally Responsible'. I have made claims that are in tension. In response to some resisters to some of my views about moral responsibility, I have criticized the reliance on self-proclaimed intuitive judgments about 'moral responsibility', but in defense of some of my own views I have implicitly, at least, appealed to intuitive judgments about moral responsibility. Although I think that my commentary so far forth has been worthwhile and instructive inasmuch as I have appealed to reasons that I expect to have widespread appeal among theorists about moral responsibility, I now want to acknowledge that I do not think that my appeals to intuition are decisive. Here is why: The term 'morally responsible' is a term of art, which has been used in a variety of ways (a fact that explains *many* misunderstandings among philosophers working in this area).⁷

Ordinary people do not go around talking about who is *morally responsible* for this or who has *moral responsibility* for that. They do sometimes talk about *responsibility*, they sometimes say things like, 'what he did was *morally wrong*', they are sometimes concerned with whether someone should be *praised* or *blamed*, and they are sometimes interested in basic *desert* ('She might have received the maximum legal penalty, but she deserved far worse', someone might say). These observations show that ordinary people care about matters that may sensibly be referred to under the label of 'moral

⁶ That is, if normal, non-defective mechanisms are indexed to what we take to be the relevant facts about our world (a non-Matrix or otherwise non-skeptical scenario world). But, if normal, non-defective mechanisms are indexed to, say, the Matrix world, then deciding on the basis of massive falsehood is the norm.

⁷ This section and section 5 is a summary of a central part of my "Voluntariness as a Guide to Moral Responsibility" (unpublished manuscript). I include it in order to provide some theoretical grounding for subsequent claims I will make about cases of information manipulation in section 6.

responsibility’, but the term is surely not in ordinary use as are other terms of philosophical interest, such as ‘knows’, ‘beautiful’, and ‘good’. For this reason, it is just a mistake to think that one has settled some philosophical dispute by appealing to some intuition about ‘moral responsibility’, for there is no such ordinary concept to which to appeal. However, we can expect to make some philosophical progress by stipulating the conditions for the term’s application, in light of theoretical needs and concerns.

The term ‘morally responsible’ has been employed by philosophers as a means of expressing the claim that one’s actions are caused by, or explained by, or otherwise suitably related to, central features of one’s personhood such that one is “an apt candidate for the reactive attitudes” and associated reactive responses such as praise/blame and reward/punishment (Fischer and Ravizza, 1998, 7)⁸ or such that “there is an “entry” in one’s “moral ledger”” (Zimmerman, 2002, 555).⁹ Understood in these ways, ‘morally responsible’ is used to evaluate persons in some way. Which way? The philosophical literature features a variety of answers from a variety of perspectives. I am motivated by a guiding concern among quite a few philosophers. It is this: quite apart from considerations about legality, the needs of society, our actual epistemic limitations when we consider real-world cases, and what makes us justified in praising/blaming or rewarding/punishing others, *what are the conditions for evaluating persons fairly in light of what they actually deserve?* On my view, this issue of basic desert explains what is plausible about both the apt-candidate-for-reactive-attitudes’ approach, as well as the entry-in-one’s-moral-ledger’ approach, to fixing a sense of the term ‘morally responsible’. To simplify matters here, I will usually utilize the apt-candidate-for-reactive-attitudes approach.

My understanding of ‘moral responsibility’ is motivated by this question: What could make something a candidate for deserving moral praise or moral blame? As Aristotle (Bk III, ch. 1) pointed out, one deserves praise or blame only for actions that one performs *voluntarily*. Contemporary theorists of metaphysical freedom or moral responsibility can be taken to be attempting to provide conditions for the relevant kind of voluntariness. No doubt we need *metaphysical conditions* to do this work. But,

⁸ Other views in the neighborhood include those by Copp (1997) and Wallace (1994).

⁹ For examples of views in the neighborhood, see Feinberg (1970), Glover (1970), Morris (1976), and Zimmerman (2002, 1998).

Aristotle rightly pointed out that there are also *epistemic conditions* for voluntariness: “Everything that is done by reason of ignorance is non-voluntary” (III. 1. 1110^b.20). To act voluntarily, Aristotle said, one must not be ignorant “of who he is, what he is doing, what or whom he is acting on, and sometimes also what (e.g., what instrument) he is doing it with, and to what end (e.g. for safety), and how he is doing it (e.g. whether gently or violently)” (III.1.1111a.3-6).

As these claims strike me as correct, I conclude that necessary conditions for being a candidate for deserving praise or blame include some freedom-relevant metaphysical conditions as well as some Aristotelian epistemic conditions. But, such conditions are not sufficient. Here is why: it is possible for one to satisfy the metaphysical (non-epistemic freedom-relevant) conditions of many popular theories of free will/action, and also to satisfy the Aristotelian epistemic conditions, but not to have a ‘moral sense’ of the sort that courts try to establish the presence of in cases in which the sanity of a defendant is in question. Typical four-year-old children are plausible examples of persons who routinely perform voluntary actions, but, lacking sufficient understanding of the concepts of moral right and moral wrong, they do not satisfy even the weakest standard for competence to stand trial in a criminal case. Consequently, it seems entirely appropriate to say that such persons lack *moral responsibility* for what they do. The same goes, it has been argued, for psychopaths.¹⁰ Thus, it seems to me that in order for one to be morally responsible for X, one must satisfy both metaphysical and epistemic conditions for voluntariness and one must satisfy an additional epistemic condition having to do with one’s understanding that X is either morally good or bad.

With these considerations in mind, I offer the following account of moral responsibility:

S is **morally responsible** for event E, which is a consequence of a basic action A performed by S, if and only if

(i) S satisfies the metaphysical (non-epistemic) conditions for free action with respect to A;¹¹
and

(ii) at the time of S’s performance of A, S is epistemically rational in believing, that (a) E will occur as a result of S’s performing A, and that (b) E is causally sensitive to S’s performing A;
and

¹⁰ For example, see Levy (2011, 119-120).

¹¹ I am, of course, waving my hand here in an attempt to remain as theoretically neutral as I can on this issue.

(iii) at the time of S's performance of A, S believes, or S is epistemically rational in believing, that E is either morally good or morally bad.¹²

What is the relation between 'moral responsibility' and praiseworthiness/blameworthiness? If you are morally responsible for X, does it follow that you are either praiseworthy or blameworthy for X? Some theorists think so, but others think (as I do) that to be morally responsible for X is to have satisfied a very important necessary, but not sufficient, condition for being praiseworthy/blameworthy for X. Here is a theoretical motivation for my view: suppose that we know that your action satisfies the conditions of well-known theories of free will/action or moral responsibility. Do we thereby know whether you deserve moral praise or moral blame for your action? Not on theories such as Timothy O'Connor's libertarian agent-causal theory (O'Connor, 2000) or Fischer and Ravizza's moderate reasons-responsiveness theory (Fischer and Ravizza, 1998). For one can satisfy those conditions whether or not one has on-balance epistemic justification for believing that the action is good and whether or not one has on-balance epistemic justification for believing that the action is bad. But, surely a person's being *deserving* of blame, for instance, essentially depends on whether the person had, at the time of action, good reason to believe, that the action was bad; after all, if the person had excellent on-balance reason to think that the action was good, then the person wouldn't *deserve* blame for the action even if that action were objectively bad. This seems to be a lesson from the observation that one can do the objectively wrong thing for the right reason (and thereby not deserve blame for the action), and one can do the objectively right thing for the wrong reason (and thereby not deserve praise for the action). I infer from these considerations that praiseworthiness or blameworthiness depends on epistemic facts about one's reasons for believing that a given action (or consequence thereof) is good or bad, epistemic facts which go beyond the epistemic facts that are relevant for assessing whether one satisfies the conditions of many leading theories of free will/action that seek to explain an action's being under the freedom-relevant, voluntary control of an agent.

¹² These conditions are consistent with the issues I discussed above: the metaphysical (non-epistemic conditions) are captured in (i), the Aristotelian epistemic conditions for voluntariness are captured in (ii), and the 'moral sense' condition is captured in (iii).

What are the necessary and sufficient conditions (which need to be added to those for moral responsibility) for praiseworthiness/blameworthiness? The question is *very* difficult to answer, and I will not make an attempt here. Instead, I will provide, in my commentary on some test cases, what I think is sufficient to show that the kind of distinction I have in mind between moral responsibility and moral praiseworthiness/blameworthiness is adequate to provide simple, elegant explanations that jibe well with reflective judgments about what persons morally deserve.

5. Test Cases for My Necessary Epistemic Conditions. In my account of moral responsibility, condition (ii) is the most interesting epistemic condition. To focus on the most interesting part of that condition, let us consider the following principles about *direct consequences* (such as putting cash into the hands of a homeless person) of *basic actions* (such as deciding to give cash to a homeless person), and *later consequences* (such as the homeless person's using the money to buy food) of basic actions.¹³

DC1: If person S is morally responsible for direct consequence E, which follows from a basic action by S, then at the time of that basic action S has on-balance good reason to believe that E will follow from that basic action by S.

DC2: If event E directly follows from a basic action by S and S, at the time of the decision, lacks on-balance good reason to believe that E will follow from that basic action by S, then S is not morally responsible for E.

LC1: If S is morally responsible for event E, which does not directly follow from a basic action by S but occurs as a later consequence of a basic action by S, then at the time of that basic action S has on-balance good reason to believe that E will occur as a consequence of that basic action by S.

LC2: If S lacks on-balance good reason to believe that E will occur as a later consequence of a basic action by S, then S is not morally responsible for E.

I will test these conditions using the well-known example from W.K. Clifford's "The Ethics of Belief" involving the ship owner who voluntarily sends his ship out to sea despite his having significant reservations about the ship's seaworthiness; consequently, the ship sinks at sea, killing all aboard (Clifford, 2003). Intuitively, the ship owner is morally responsible for a direct consequence of his

¹³ The first part of this section is a summary of a section in my "Voluntariness as a Guide to Moral Responsibility" (unpublished manuscript).

decision to send the ship to sea (namely, the event of his *sending the ship to sea*, an event to which DC1 and DC2 are relevant).¹⁴ My conditions very plausibly get the intuitively correct result about this event.

Clifford strongly suggests (by way of saying that the shipowner is “guilty of the death of those men”) that the ship owner is also morally responsible for a later consequence of his sending the ship out to sea, namely, *the dying of the sailors* (an event to which LC1 and LC2 are relevant). For current purposes, let us consider the case to be one in which, at the time of his decision, the ship owner has some evidence that the ship will sink but his total evidence doesn’t, on balance, support thinking that it will sink; thus, the proposition *that the sailors will die* was, at the time of his basic action (his decision to send the ship to sea) *epistemically possible* for him but not on-balance epistemically probable for him. [I hereby coin the term ‘epistemically possible’ to mean ‘a proposition’s having higher epistemic probability for one than epistemic possibility but lower than on-balance epistemic probability of the kind that makes *believing* the proposition the epistemically rational doxastic attitude to take toward the proposition’.] Because the ship owner had some reason to believe that the ship would sink, but his overall evidence did not make believing that the ship would sink epistemically rational for him, the proposition *that the ship will sink* was, at the time of his decision, epistemically possible for the ship owner. Thus, LC2 entails that the ship owner is *not* morally responsible for *the dying of the sailors*.

This result may initially strike readers as incorrect. A critic may say that the ship owner is morally responsible and blameworthy for *the dying of the sailors* due to the fact that the ship owner had serious reservations about the seaworthiness of the ship, he knew that the decision to send the ship rested entirely with him, he knew that lives were at stake, etc., even if his total evidence at the time of the decision did not make epistemically probable for him that the sailors would die as a result of his decision. However, this critical judgment is precisely what I deny: we can explain all that intuitively needs explaining by means of my epistemic principles and a plausible distinction between moral responsibility, on the one hand, and moral praiseworthiness/blameworthiness, on the other. For instance, we can very plausibly

¹⁴ We can describe the event as a direct bodily consequence of the ship owner’s decision, by specifying it as *the ship owner’s saying ‘let her sail’ to the ship’s captain* or as *the ship owner’s signing voyage release documents*.

explain our intuition that the ship owner deserves significant moral blame solely by reference to the metaphysical and epistemic facts that obtained at the time of his decision to send the ship to sea. By hypothesis of the example, the ship owner had, at the relevant time, significant reservations about the ship's seaworthiness, he knew that lives were at stake, and he knew that he was in control of whether the ship would go to sea. Given those facts, it is very plausible that he was both morally responsible and morally blameworthy for *sending the ship to sea*. This result is consistent with my account of 'moral responsibility' and a plausible view (to be discussed) about moral blameworthiness.

Now, consider an alternate case in which things go exactly as they did in the original case except that the ship does not sink. Considering this alternate case, Clifford himself says that the ship owner's guilt would be *no less* if the ship had not sunk. Surely Clifford is right about that, and the point is instructive, and, I think, decisive for a point I wish to make: the ship owner would be just as *worthy* of moral blame for sending the ship out to sea had the ship not sunk as he was in the actual case. Perhaps he would have been lucky had the ship not sunk, but if he was morally blameworthy for the decision he made in light of his evidence in the actual case, then surely he would have been equally morally blameworthy for his decision in the alternate case in which he made his decision under the same circumstances and for the same reasons. Thus, his *worthiness* of moral blame is not explained, or even heightened, by the ship's sinking or the sailors' death. It is fully explained by his satisfying the metaphysical and epistemic conditions for his being morally responsible for *sending the ship to sea*, together with his significant doubts about the ship's seaworthiness, his knowledge that lives were at stake, and his knowledge that the decision whether to send the ship to sea rested entirely with him.

Since *the ship's sinking* and *the sailors' dying* were not events that were on balance epistemically probable for the ship owner, but were instead merely epistemically possible for him, how can I explain the fact that the ship owner was *morally blameworthy* in a way that jibes with our intuitions? I can do it by specifying a plausible sufficient condition for moral blameworthiness. This seems to be one:

S is morally blameworthy for S's basic action A if S is morally responsible for A and, at the time S performs A, a later consequence C of A is epistemically possible for S, and S is epistemically rational in believing (i) that C's occurring is significantly worse than any other

consequence of A that S is epistemically rational in believing will occur, and (ii) that C's occurring is significantly worse than any consequence C* that would occur if S were to perform some other basic action A* that S is epistemically rational in believing is in S's control.

We may assume that, at the time he decided to send the ship to sea, the ship owner had excellent reason to believe that *the sinking of the ship* would be significantly worse than any other consequence of the ship owner's decision that he was epistemically rational in believing would occur, and we may assume that the ship owner had excellent reason to believe that the ship's sinking would be significantly worse than anything that would happen if the ship owner had made a different decision (such as having the ship checked first, or making a sandwich, etc.). These facts seem sufficient to explain why the ship owner is morally blameworthy for deciding to send the ship to sea.¹⁵

I conclude that the distinction I am drawing between the conditions for moral responsibility and those for praiseworthiness/blameworthiness allows us to explain our intuitive judgments about the ship owner in a simpler, more elegant way than is usually attempted.¹⁶ He is morally responsible and morally blameworthy only for what was in his voluntary control. The moral blameworthiness of his voluntary action of sending the ship to sea is fully explained by the metaphysical and epistemic facts that obtained at the time of his decision to send the ship to sea. On my view, he is just as *morally blameworthy* in the case as we thought he was, but if I am right, then he is not morally blameworthy (because he is not morally responsible) for the dying of the sailors.

My guiding idea is that one can be morally praiseworthy/blameworthy for X only if one is morally responsible for X, and one can be morally responsible for X only if X is in one's voluntary control.

¹⁵ Furthermore, we can use the case to explain, at least in part, our intuitive judgment that blameworthiness comes in degrees. Because the sinking of the ship was something that the ship owner had reason to believe would be a *very bad* thing, and because he had good reason to believe that there was a significant probability that it would sink, it is plausible that the ship owner deserves a high degree of moral blame, certainly more than he would deserve if he merely had excellent reason to think that the sinking of the ship would produce only bumps and bruises.

¹⁶ Michael Zimmerman (2002) is an example of a theorist who would share Clifford's and my view that the ship owner is equally blameworthy, whether the ship sinks or not, but Zimmerman thinks that a person can be morally responsible, period, without having performed any action, basic or otherwise. Accordingly, a person can be morally responsible for being the sort of person who *would* send the ship to sea *if* he were to have the actual ship owner's opportunity, even if the person never sends a ship to sea. Because 'moral responsibility' is a term of art, I cannot refute Zimmerman's view, and I do not deny that the specific kind of 'hypological' judgment that Zimmerman is drawing attention to is of philosophical interest (I think it is). All I can say is that I am inclined to think that moral responsibility is most usefully understood as a property that a person can have with respect to a suitably related basic action.

Theorists about moral responsibility falter on the Aristotelian point about voluntariness, for they tell us, in effect, that a person is morally responsible for basic actions (such as decisions) and direct bodily consequences of those actions because they are in the *voluntary control* of the person, but then they struggle to accommodate what are supposed to be intuitive judgments about moral responsibility for later consequences of basic action, for such later consequences are often not under the voluntary control of the person. My view has the resources to tie all instances of morally responsible action to voluntariness and to do this in a way that is consistent with reflective judgments about praiseworthiness and blameworthiness .

Let's apply these distinctions to a few possible consequences of your decision in the *Family Policy* example. Suppose that one consequence of your decision is that your daughter is never violated by the sex-offender who lives next door. It is plausible that you are morally responsible for this consequence on the assumption that it was epistemically probable for you that enacting the family policy would prevent your daughter from being violated. After all, you enacted the family policy in part for the purpose of preventing your daughter from being violated by the neighbor, and since, we may suppose, you know that child sex offenders rarely violate children that the offenders lack easy access to, it's very plausible that it was epistemically probable for you that enacting the family policy would have the consequence of preventing the relevant abuse. You voluntarily enacted the family policy for the purpose of protecting your daughter from the threat of sexual abuse, and you did it rationally believing that your action would result in your daughter's being protected. So, I think it's plausible that you bear moral responsibility for your daughter's not being sexually abused by the neighbor. This result is consistent with my epistemic conditions.

Now consider a different consequence of your decision. Suppose that, as a result of your decision, your daughter becomes a more socially withdrawn person than she would have become had you not enacted the family policy. Are you morally responsible for this consequence? I think that you could be. Suppose that, when you were deliberating, you weighed the good of doing something that would likely prevent your daughter's being sexually abused against the bad of sheltering her more than you would

have otherwise thought was good for her. And suppose that you came to the conclusion, for good reason, that the overall value of doing something you had good reason to think would result in preventing sexual abuse was significantly greater than the disvalue of doing something that you were epistemically rational in thinking would inhibit your daughter's social development. Given the facts of the case as described, it's plausible that this consequence is under your voluntary control; thus, it is plausible that you are morally responsible for this consequence of your decision. It doesn't follow that you are morally blameworthy for enacting the family policy; indeed, it could be that you are morally praiseworthy for enacting the family policy in this situation. After all, it is plausible that one is morally praiseworthy when one voluntarily makes the best of what one has good reason to think is a bad situation.

Now consider a third possible consequence. Suppose that a result of your decision is that your daughter develops a very debilitating mental illness, which, as a matter of contingent fact, she would not have developed if you had not enacted the family policy; but suppose that, at the time of your decision, you lacked any reason to believe that your enacting the family policy would result in this consequence. It seems to me that you are not morally responsible for that consequence, even though you are definitely part of the cause of that consequence. One more or less neutral reason to think that you are not morally responsible for the consequence is that you are clearly not blameworthy for that consequence (and neither are you praiseworthy for it). And part of the reason that you are not blameworthy (or praiseworthy) for that consequence is that you lacked on-balance good reason to think that such a consequence would occur. Indeed, you had no reason to think it would occur. Thus, it was not under your voluntary control.

6. Applying the Conditions of Moral Responsibility to Information Manipulation Cases.

Family Policy is a paradigm example of morally responsible action. Popular theorists of metaphysical freedom and moral responsibility all want their theories—whether they are libertarian, indeterministic theories or determinism-friendly compatibilist theories—to explain how such examples of voluntary action are instances of morally responsible action. Thus, they should agree that conditions (i) and (ii) of my account of moral responsibility is satisfied. So is condition (iii), for at the time you decided to enact the family policy, you were epistemically rational in believing that enacting the family policy was either

morally good or morally bad. It follows on my view that you are morally responsible for your action. Furthermore, given some of the further details I have discussed about the case, it is plausible that you are praiseworthy for your decision and for your enacting the family policy. For we may suppose that you had excellent reason, all of your evidence considered, to believe that enacting the family policy was the best thing to do in your circumstances. Indeed, we can stipulate that your enacting the family policy in your circumstances was the best decision that would be made by the most rational, morally upstanding human being.

Now, suppose, as I have intended for you to do all along, that you hold fixed all these facts when you turn your attention to *Manipulated Family Policy*. Clearly, if my conditions for moral responsibility are satisfied in *Family Policy*, then they are satisfied in *Manipulated Family Policy*. The only difference in the cases is that in the latter, you deliberate on the basis of false, manipulated information. I have been arguing that false information, by itself, is not a relevant difference between the cases. None of the reasons discussed for thinking that it is a relevant difference have survived scrutiny.

That leaves just the bare difference of information manipulation. Why would the mere fact that the information we rely on is manipulated prevent us from being morally responsible for what we do in light of it? Is it because *we are not in control of the information* when it is manipulated? Surely not. Note that in our everyday lives we are very rarely in control of the information we rely on in our practical reasoning. I swerve my car to the left in part because I get visual information that you have entered the right side of my lane. I don't control this visual information I rely on. It comes to me by virtue of my happening to be where I am in the road at the same time that you happen to be where you are in the road, and my having my eyes open and focused on the road. I go to the faculty meeting partly because my boss told me that I had a faculty meeting. I was not in control of this information. My boss sought me out, said it, and I heard it, but I didn't control that. I press the 'p' key on my keyboard partly because of my visual and memory information about where the 'p' key is located and I want to type 'press', but I do not control the information I have about where the 'p' key is located on my keyboard. Indeed, every day, virtually all day long while we are awake, we are engaging in practical reasoning on the basis of

information that we do not control, and yet we are not inclined to think that these facts make any difference in whether we are morally responsible for what we do in those situations. But, if we can be morally responsible for what we do in ordinary cases when we do not control the information we rely on, then why think information manipulation prevents moral responsibility?

Christopher Evan Franklin (2006) responds to my line of reasoning where he writes about ‘PAs’ (an abbreviation for ‘pro-attitudes’, which are supposed to be attitudes that can motivate us to action):

Long is correct that we are not often in direct control of our acquisition of PAs. Nevertheless, it is plausible to think that agents exercise a sort of indirect control over the acquisition of these PAs. It is implausible to think that agents can form PAs simply in virtue of an act of will, but this is not to say that agents cannot exercise control insofar as choosing what sort of environments or contexts to place themselves in. Again, consider John who may not be able to prevent his forming desires and beliefs about stealing by a sheer act of will, but he can perform an act of will to avoid circumstances that allow or make probable the forming of such a desire or belief. (181-182)

Franklin’s idea is that a necessary condition for moral responsibility is that we have at least indirect control over the pro-attitudes that have some role in what we do.

I do not deny that we can exercise indirect control over some of our pro-attitudes, such as some of our beliefs and some of our desires.¹⁷ Nevertheless, there is a serious problem with Franklin’s proposal. To appreciate it, let us consider the context of Franklin’s commentary. He is attempting to explain why agents whose pro-attitudes are induced (in a certain way) by an external manipulator are not morally responsible for their actions. Toward that end, he claims that, in order for us to be morally responsible for our actions that result from practical reasoning, we must be morally responsible for having the pro-attitudes that play a role in our practical reasoning; although we may rarely have *direct control* over our having the pro-attitudes that we rely on, we can exercise *indirect control* over them in a way that renders

¹⁷ For instance, I don’t now believe that the lights in my office are off, but I can now move toward my light switch and then push it into the ‘off’ position; consequently I will believe that the lights are off (and then, because I want to continue working, I’ll rely on this new information about being in a dark room as a motivation to turn the light switch to the ‘on’ position). This shows that I can exercise some indirect control over whether I will believe in the future that the lights in my office are off. Also, I can exercise some indirect control over some of my desires. For instance, suppose that, worried by my lack of patience, I form an intention to do something that will cause me to have a strong desire for patience; consequently, I frequently go to church or to counseling in order to develop such a desire. If I succeed, then it is plausible that I exercised some indirect control over my acquiring a strong desire for patience.

us morally responsible for them. Here is Franklin's account of moral responsibility for acquiring a pro-attitude:

EC: Agent A, who is in context C, is morally responsible for acquiring pro-attitude PA, iff A could reasonably come to have known that the forming of PA is probable given C. (181)

My first criticism is that it is entirely mysterious how the truth of a counterfactual about what I would be reasonable in believing could, by itself, be *sufficient* for being *morally responsible* for anything. Recall that Franklin's idea is that we exercise *indirect control* over our having the pro-attitudes that play a role in our practical reasoning just when EC is satisfied with respect to all the relevant pro-attitudes. But, satisfying EC is woefully insufficient for exercising *any* control over what we acquire.

In light of this criticism, one might want to hold that satisfying the condition of EC is only a *necessary*, but not sufficient, condition for being morally responsible for acquiring a pro-attitude. Perhaps there is some additional condition that would be sufficient for that purpose. But, even if that is so, there is a serious problem with the general strategy, and it is this: a person can be morally responsible for an action even when the person does not indirectly control the pro-attitudes that have some role in that person's practical reasoning. Consider a variation on John Locke's locked room example.¹⁸ Suppose that you are taken while sleeping and without your knowledge into a room in which there is a person you have been longing to talk to. Upon awaking, you open your eyes and see the other person. You do not even indirectly control your seeing the other person when you open your eyes. Nevertheless, the information you acquired by seeing the other person was a crucial input to your process of practical reasoning that resulted in your starting a friendly conversation with the other person. Surely you can be morally responsible for starting such a conversation even though you did not exercise control, direct or indirect, over your acquiring that pro-attitude.

Another general problem concerns distinguishing the pro-attitudes that constitute, or partly constitute, one's normative orientation, and the kind of pro-attitudes that I have been drawing attention to. Critics of my claims about my *Schmidt* Frankfurt-type example tend not to appreciate the relevance of the

¹⁸ See Locke (Bk II, ch. 21, sec. 10).

distinction.¹⁹ When I say that input (information) manipulation does not preclude moral responsibility for action, I am not talking about the manipulation of pro-attitudes that are part of one's normative orientation. To put the point in terms of Fischer and Ravizza's moderate reasons-responsiveness theory, the inputs that are manipulated in my *Schmidt* test cases are not pro-attitudes that (partly) constitute Schmidt's psychological mechanism of practical reasoning. Although Fischer and Ravizza have been criticized for not providing details about how to individuate these psychological mechanisms, they have been perfectly explicit about the fact that these psychological mechanisms do not themselves include the inputs. They are thinking of the entire process of practical reasoning that produces an action as a causal chain that runs from one's new experiences in the world, through one's own way of deciding what to do, to a decision and sometimes to an associated bodily action. One's new experiences in the world yield the 'reasons' that they say are the 'inputs' to the psychological mechanism. The psychological mechanism itself works on the new inputs. That mechanism is some complex involving one's 'normative orientation' (one's standing desires, preferences, values, etc.) and sometimes one's own way of reasoning. Whatever imprecision there is in Fischer and Ravizza's notion of the relevant kind of mechanism, the crucial point to see is that, when they talk about an agent's reasons as inputs, they do not intend to include as inputs the agent's long-standing desires, preferences, and values;²⁰ so, if such things count as pro-attitudes, then they are the kinds of pro-attitudes that help to constitute the person's psychological mechanism.

Now we are in a position to understand the problem with Franklin's proposal. Perhaps there is some plausibility to the claim that, in order to be morally responsible for an action, a person must exercise some control over the pro-attitudes that partly constitute the person's relevant psychological mechanism, but it is entirely implausible to suppose that a person must also control the inputs to that psychological

¹⁹ In Long (2004) I present a case of information manipulation in which a manipulator manipulates some of the information that Schmidt has to go on as Schmidt is deliberating about whether to cast a deciding vote either to give Hitler dictatorial power or to keep him from power. The information manipulation gives Schmidt good reasons to think that Hitler would be an excellent leader of Germany, but the manipulation does not affect Schmidt's way of deliberating, nor does it affect his long-standing general preferences, desires, and values. I conclude that Schmidt is morally responsible, but not blameworthy, for voting for Hitler.

²⁰ If they do intend to include as inputs the agent's long-standing desires, preferences, and values, then the response by Fischer (2010) to Long (2004) is inexplicable.

mechanism. After all, the inputs are typically new beliefs one has as a result of one's new experiences in the world. I am currently not moving toward you with a bandage. But, were I to have the experience of seeing blood running profusely from your head, I might (given the kind of person I am) then decide to move toward you with a bandage and try to stop the bleeding. It is true that I can control many of my decisions about where I will go and what I will do, but if I am living a more or less normal human life, I cannot control all the new information (the inputs) to my practical reasoning no matter where I go or what I do. Nevertheless, I can be morally responsible for whatever I do in light of that new information.

Indeed, I think that we are never off the moral hook just because the new information we are relying on is manipulated, even if we know nothing at all about that manipulation. Those of us who are persons who have not lost control of ourselves or become insane (or the like) do not escape moral responsibility for what we do in light of our epistemic condition, whatever it is. The information manipulation may be so vicious that we are not blameworthy, or it might be so virtuous that we are not praiseworthy. It may be so devious or fraudulent that we do not satisfy a reasonable standard for legal consent. But, the fact remains that, for most of us, we are morally responsible for what we do in light of the evidence we have to go on. Information manipulation does not preclude moral responsibility.

Works Cited

Aristotle, *Nicomachean Ethics*, In *A New Aristotle Reader*, ed. L.L. Ackrill (Princeton University Press, 1987).

Clifford, W.K., "The Ethics of Belief", In *The Theory of Knowledge: Classic and Contemporary Readings*, 3rd ed., Ed. Pojman (Wadsworth, 2003), ch. 10, sec. 2.

Copp, David, "Defending the Principle of Alternate Possibilities: Blameworthiness and Moral Responsibility" *Nous* XXXI:4 (1997) 441-456.

Feinberg, Joel, *Harm to Self*, vol. 3 of *The Moral Limits of the Criminal Law* (New York: Oxford University Press, 1986).

Feinberg, Joel, *Doing and Deserving* (Princeton University Press, 1970).

Fischer, John Martin, "Manipulation and Guidance Control: A Reply to Long", In *Action, Ethics, and Responsibility*, ed. Campbell, O'Rourke, and Silverstein (MIT Press, 2010) 175-186.

- Fischer, John Martin, and Ravizza, Mark, *Responsibility and Control* (Cambridge University Press, 1998).
- Franklin, Christopher Evan, "Plausibility, Manipulation, and Fischer and Ravizza", *The Southern Journal of Philosophy* 44:2 (2006) 173-192
- Glover, Jonathan, *Responsibility* (New York: Humanities, 1970).
- Levy, Neil, *Hard Luck* (Oxford University Press, 2011).
- Long, Todd R., "Moderate Reasons-Responsiveness, Moral Responsibility, and Manipulation", In *Freedom and Determinism*, eds. Campbell, O'Rourke, and Shier (MIT Press, 2004) 151-172.
- Morris, Herbert, *On Guilt and Innocence* (California University Press, 1976).
- Nahmias, Eddy, "Review of Freedom and Determinism", *Notre Dame Philosophical Reviews*, June 11, 2005 (<http://ndpr.nd.edu/review.cfm?id=2841>).
- O'Connor, Timothy, *Persons and Causes* (Oxford University Press, 2000).
- Steadman, James, "Moral Responsibility and Motivational Mechanisms", *Ethical Theory and Moral Practice* (forthcoming).
- Wallace, Jay, *Responsibility and the Moral Sentiments* (Harvard University Press, 1994).
- Zimmerman, Michael, "Taking Luck Seriously", *Journal of Philosophy* 99:11 (2002) 553-576.
- Zimmerman, Michael, *An Essay on Moral Responsibility* (Totowa, NJ: Rowman and Littlefield, 1988).