Brute Luck and Responsibility

Peter Vallentyne, University of Missouri-Columbia

April 6, 2013

The issue of responsibility for the outcomes of one's choices is relevant to both ethics and political philosophy. It seems at least initially plausible, for example, that an agent should be held morally accountable for an outcome of her choice (e.g., liable to punishment or with an obligation to compensate victims) only if she is responsible for that outcome. The concept of responsibility for outcomes, however, remains radically under-explored. In particular, the issue of *partial* responsibility for an outcome (rather than it being all or nothing) needs further development. Even the most sophisticated account of responsibility—that of John Martin Fischer and Mark Ravizza[1]—does not address this issue systematically. In this paper, I propose an account of partial responsibility based on partial causal contribution.

In the process of developing the framework, I shall make many controversial assumptions. I do this so that I can focus on the central concern—partial responsibility—in a relatively concrete context. My hope is that my key claims can be adopted (perhaps with some modification) even by those who reject some of the assumptions made. In any case, the paper is highly exploratory and only meant to motivate (rather than defend) the possibility of a certain way of conceptualizing responsibility for outcomes.

1. Brute Luck

Brute luck egalitarianism holds that individual advantage (wellbeing, resources, whatever) from brute luck should be equalized.[2] Brute luck for an agent is typically understood as something like

those states and events that the agent could not deliberatively influence (or control). Brute luck egalitarianism holds that such effects on individual advantage should be equalized, but that the effects that are attributable to the agent's choices need not be equalized.

Susan Hurley[3] has recently suggested that brute luck egalitarians should adopt a *thin* conception of brute luck rather than any specific *thick* conception. A thin conception defines brute luck as that for which the individual is not responsible while leaving open what the correct conception of responsibility is. A thick conception, by contrast, gives a specific definition of brute luck (e.g., inability to control) that is not sensitive to the general issues of responsibility.

Hurley is, I believe, correct. Brute luck egalitarians should simply appeal to the correct account of responsibility and define brute luck as the absence of such responsibility. Although my main interest in responsibility for outcomes is for its implications for brute luck egalitarianism, I shall address responsibility in general terms rather than focus on an issue purely internal to brute luck egalitarianism.

2. The Concept of Agent-Responsibility

So far, I have written of responsibility, but the term "responsibility" is used in various ways and the concept on which I shall focus might be best called *agent-responsibility*. This is sometimes called *attributive responsibility.* The rough idea is that an individual is agent-responsible for a choice or outcome to the extent that it suitably reflects the exercise of her agency. Being *causally responsible* for an outcome, for example, is not sufficient for agent-responsibility, since the agent may reasonably have been unaware that her choice had the effect in question (e.g., when one flicks a light switch that has been secretly rigged to set off a bomb). One may be agent-responsible for the foreseeable causal effects of one's choices, but one is not agent-responsible

for all the causal effects of one's actions.

Moreover, an agent's being agent-responsible for a choice or outcome does not entail that she is *morally accountable* for the choice or outcome in the sense that she has certain substantive moral duties, liabilities, etc. in virtue of the occurrence of the outcome (e.g., the duty to compensate those harmed by one's actions, liability to punishment, or the absence of duties of others to bear the cost to oneself). That is a substantive moral question that will be answered in different ways by different moral theories. Agent-responsibility merely establishes that the choice or outcome flows in the right kind of ways from the individual's agency. It leaves open what the substantive moral implications are. Although most substantive moral theories hold agents morally accountable only for that for which they are agent-responsible, some do not.[4]

In what follows, I shall focus on agent-responsibility for outcomes. Responsibility for *choice* is, at least on the view I propose, simply a special case—the one where the relevant outcome is simply the occurrence of the choice itself.

At a conceptual level, there are at least three distinct conceptions of agent-responsibility. On the (Strawsonian) *social practice conception*, an agent is agent-responsible for some outcome if and only if some kind of reactive attitude (blame, praise, etc.) towards the agent with respect to the outcome is appropriate relative to the norms of our social practice.[5] This view correctly links agent-responsibility with reactive attitudes. I believe, however, that it incorrectly links agent-responsibility to our social practices. We know that in the past our social practices were flawed in all kinds of way and there is little reason to think our social practices are now perfect. Of course, we cannot step entirely outside our social practices, but we can criticize some of them—such as those involving certain reactive attitudes—on the basis of theoretical reflection upon our

social practices as whole. Obviously, the issue is complex and controversial, and I will not attempt to resolve this issue here. I merely flag this conception as problematic and set it aside.

A closely related conception of agent-responsibility is *the normative conception*, according to which an agent is agent-responsible for some outcome if and only if some kind of reactive attitude (blame, praise, etc.) towards her with respect to the outcome is appropriate *from some correct normative perspective* (morality, prudence, etc.). This is like the social practice conception in that it focuses on reactive attitudes. It does not, however, privilege our current social practice with respect to reactive attitudes. Instead, it appeals to what is appropriate relative to some correct normative perspective.

A third conception of agent-responsibility is the *metaphysical (or ledger)* conception according to which an agent is agent-responsible for some outcome to the extant that it suitably reflects her exercise of agency.[6] Unlike the two previous conceptions, this has no conceptual link with the appropriateness of reactive attitudes. Nonetheless, the following plausible substantive assumption makes such a link: Some kind of reactive attitude towards an agent with respect to an outcome is *appropriate from a correct normative perspective* if and only if the occurrence of the outcome is suitably reflective of the individual's agency. If one grants, as I do, the substantive assumption, then the metaphysical and the normative conceptions are equivalent. I shall be working with the metaphysical conception—although, given the above substantive assumption, this is equivalent to the normative conception.

It's important to note that agent-responsibility is a broader notion than *moral* (attributive) responsibility. The latter arguably holds only if the agent has the capacity to be aware of moral considerations.[7] An individual with no such capacity may still be capable of reflectively assessing and modifying her beliefs, desires, and intentions. Moreover, she may be quite capable

of choosing so as to promote her prudential interests. Such an agent can be agent-responsible for an outcome (e.g., for her sunburn after lying on the beach all day)—even if she is completely insensitive to moral considerations. Moral responsibility entails agent-responsibility, but the converse does not hold. My topic is the more general topic of agent-responsibility.

Finally, and importantly, the fact that an individual is not agent-responsible for some outcome does not entail that her choice is not wrong or otherwise morally flawed. If wrongness is, as many claim, based on the objective consequences of one's choices, then an action can be wrong in virtue of consequences for which the individual is not agent-responsible (e.g., because she could not reasonably have known that they would follow). Moreover, an individual who is highly disposed to perform morally atrocious acts may have a morally vicious character, but she may bear little agent-responsibility for her choices if they are primarily a result of her genetic endowment and early childhood socialization. (This last claim is controversial. My point here is simply that this is a possible view.)

For brevity, I shall use "responsibility" as short for "agent-responsibility".

3. Narrow and Broad Agent-Responsibility

Before developing a framework for thinking about responsibility, one further distinction is needed. Consider two identical and identically situated agents, each of whom throws a rock at my window with the intention of breaking it, where the objective, single case chance (propensity) of breaking the window is less than one. One agent succeeds in breaking the window and the other does not. Each is, let us suppose, responsible for intending to break my window. If we are concerned with responsibility in the strictest sense, neither agent, I would argue, is responsible for breaking the window. One agent did not break my window at all. The

other agent did break my window, but that event depended in part on factors that do not reflect her exercise of agency. For example, perhaps the only reason the second agent was successful, and the first agent was not, was because a sudden gust of wind caused the first agent to miss. That difference is not a reflection of the agency of either individual. Moreover, the *normatively appropriate* reactive attitudes (blame, praise, etc.), I would argue, are not sensitive to the difference that a gust of wind makes (although, admittedly, our actual attitudes are so sensitive). How could it be appropriate to blame a person for something that in no way reflects her exercise of agency? In the strictest sense, I claim, agents are responsible, at most, only for their mental states. Everything else depends on factors that do not suitably reflect their agency.[8]

Responsibility for outcomes in the strictest sense, I claim, is *narrow responsibility* in the sense that it depends solely on the agent's mental states.[9] This is, of course, controversial and I shall not attempt to defend this claim here. The important point is that, if I am correct, then the notion of responsibility relevant for substantive theories of moral accountability (e.g., moral duties to compensate) is not responsibility in the strictest sense. Those who believe in moral accountability believe that agents are at least sometimes morally accountable for at least some aspects of the external outcomes of their choices and not merely for their *believed* outcomes. For example, a person who intentionally breaks my window is not merely accountable for so intending. She is also accountable for breaking my window.

Intermediate between narrow responsibility (based solely on the agent's mental states) and causal responsibility (based solely on the outcomes of choices, even those aspects of which the agent was unaware) is *broad responsibility*. Roughly speaking, agents are broadly responsible for an outcome to the extent that they are causally responsible for it on the basis of an autonomous choice and foresaw (or should reasonably have foreseen) that their choice would

have that causal impact. Like casual responsibility, it is concerned with responsibility for events in the external world (as well as mental states). Like narrow responsibility, it holds that an agent's beliefs about the outcomes of her choices limit that for which she is responsible. In the above example, both agents made an autonomous choice to throw a rock at my window with the intention of breaking it. The first agent broke my window, foresaw this result, and is thus broadly responsible for the breaking of my window. The second agent did not break my window and thus is not broadly responsible for its breaking. Suppose now that a third similarly-situated agent throws a rock at a tree, completely unaware (and reasonably so) that my house was behind the tree, and breaks my window. Although she is causally responsible for my window breaking, she is not broadly responsible.[10]

Broad responsibility is the notion that most have in mind when they claim that individuals are morally accountable for that for which they are responsible. It is certainly what brute luck egalitarians have in mind.[11] In what follows, I focus on broad responsibility. All unqualified references to responsibility or agent-responsibility should be so understood.

4. Some Background Assumptions

We are finally ready to begin sketching a framework for assessing (broad) responsibility for outcomes. For simplicity, I shall assume throughout that we are dealing with an agent who is making her *first* autonomous choice. I do this to eliminate any question of responsibility due to *past choices.* A plausible notion of responsibility will be historical in the sense that one can be responsible for an outcome in virtue of past choices (e.g., for injuring someone while totally drunk in virtue of one's previous choice to have a drink).[12] Here, for simplicity, I set this issue aside.

In what follows, I shall focus on the mental action of choosing that typically results in physical action. I do this because I believe that choices are the basic object of normative assessment. Such choices need not issue in physical action. Some choices are choice to "do nothing". Although I shall not explicitly address responsibility for omissions, these are, I believe, covered implicitly. Responsibility for omitting to make some choice (or action) is covered by responsibility for the alternative choice made.

In order for an individual to be responsible for an outcome, three general conditions must be satisfied: (1) an autonomy condition: she must make a suitably autonomous choice, (2) a causal condition: the outcome must be suitably causally related to the choice, and (3) a belief condition: she must have (or should reasonably have) a suitable belief that the outcome is so related. I shall briefly comment on each condition.

In order for a choice to be autonomous, the choice must issue from the agent's capacity to reflect upon her beliefs, desires, and intentions and to exercise some kind of control over her intentions. The choice, that is, must be reasonably reasons-responsive. Choices that are manipulated by someone else stimulating one's brain in certain ways are not autonomous. Choices of a three year old and choices of an agent in an extreme panic are probably not suitably autonomous. Beyond that, much is controversial and I take no stance. Perhaps the agent also needs to have some kind of higher order reflective endorsement of her first order beliefs, desires, and intentions. Perhaps she also needs to "take ownership" of her choice in some other manner. Perhaps autonomous choice requires the ability to choose otherwise. Because my focus is on responsibility for outcomes given an autonomous choice, I leave open what exactly is required for autonomous choice. My topic is responsibility for outcomes (given autonomous choice)—not free will or autonomous choice.

Consider now the second condition—the causal condition. It requires that the outcome be suitably causally related to the choice. One possible requirement is that, given the circumstances, the choice is nomically sufficient for the outcome (i.e., the laws of nature plus the circumstances plus the choice entail the outcome). This, however, is implausible because it is not appropriately sensitive to the extent to which the choice *makes a difference* with respect to the chance of the outcome occurring. Suppose, for example, that it is nomically *inevitable* that it will rain in a given spot one minute from now (e.g., it will rain no matter what I do or anyone else does). Given the circumstances, my opening my umbrella now is (vacuously) nomically sufficient for it raining one minute later. Responsibility for outcomes, however, does not include—at least not always—outcomes that will occur no matter what anyone does. Thus, mere nomic sufficiency in given circumstances is too weak a connection.[13]

Another aspect of the problem is that, even where the outcome is not inevitable for the agent, the nomic sufficiency condition fails to distinguish a case where an agent ensures that an antecedently *unlikely* outcome occurs from one where she ensures that an antecedently *likely* (but avoidable) outcome occurs. It fails, for example to distinguish a case of ensuring a flat tire by puncturing a tire that is in excellent condition from a case of ensuring a flat tire by puncturing a tire that is in very poor condition. In both cases, the agent is to some extent responsible for the flat tire, but she is responsible to a greater extent where her choice produced a greater increase in the probability of a flat tire. An adequate specification of the relevant causal connection must, I claim, be sensitive to the *difference* that a choice makes to chances and not be based merely on what those chances are given the choice.

I shall assume that the relevant causal connection is that the choice *increases the objective chance* that the outcome will occur—where objective chances are understood as

objective probabilities in the sense of single case propensities. So understood, objective chances (or probabilities) are objective dispositions (or propensities) of the world to go a certain way as determined by the laws of nature and the entire state of the world.[14] The requirement that the choice increase the chances of the outcome can probably be understood in different ways, but I shall understand it as follows: an agent's choice, c, increases the chance of outcome, o, just in case, just prior to the choice (given the entire state of the world and the laws of nature), the objective chance of o *given c* is higher than the objective chance of o *given that she makes an autonomous choice (of some sort)*.[15] (This will be explained more below.) So understood, if a choice was determined to occur by prior events (e.g., if determinism is true), then the agent is not responsible for any outcome in virtue of that choice. The laws of nature plus the state of the world at any point in the past fully determine what will happen in the future. The occurrence of the choice does not change any probabilities (they all remain 0 or 1).

The chance-increasing conception of the causal condition is, of course, highly controversial. It entails incompatibilism with respect to responsibility (the view that responsibility is incompatible with causal determinism). Harry Frankfurt, John Martin Fischer and Mark Ravizza, and many other compatibilists have are argued that one can, under certain circumstances, be responsible for outcomes that will occur no matter what one chooses (and even if one could not choose otherwise).[16]

Those who want to defend a compatibilist account of responsibility will reject the chance-increasing requirement and replace it with some weaker causal condition. For the reasons given above (e.g., to avoid responsibility for inevitable rain), the weaker condition must require more than nomic sufficiency given the circumstances. Fischer and Ravizza, for example, require that the occurrence of the outcome be suitably sensitive to the choice. Their test for suitable

sensitivity, however, appeals to what would happen if a different choice were made under certain *counterfactual* conditions. Like most incompatibilists, I am deeply suspicious of accounts of responsibility that appeal to counterfactual circumstances. I believe that an adequate account of responsibility for outcomes must take the full circumstances of choice as they are. This topic, however, is beyond the scope of this essay. I shall simply assume that an agent is responsible for an outcome only to the extent that her choice increased the chance of the outcome occurring. My hope is that the framework that I develop can be generalized to handle accounts that impose different requirements for the causal connection between choice and outcome. Here, however, I simply restrict my attention to chance-increasing accounts.

It's worth noting that the restriction of responsibility to outcomes for which the agent's choice increased the chance does not in any way preclude several other kinds of normative assessment. Agents may be (as I believe) *strictly responsible* (in the narrow sense) for making a choice that they *believe* will increase the chance of harm even if that chance is not in fact increased (and even if the harm does not occur).[17] This is an assessment of the purely internal aspects of their agency and not based on any causal condition. In addition, the *characters of agents* who are unable (or insufficiently disposed) to avoid harming others may be judged morally vicious. The moral assessment of character is concerned with the agent's disposition to choose in various ways and not with causal impact.

The third condition for responsibility for outcomes is the belief condition. It requires that the agent have a suitable belief that the outcome is suitably causally connected to her choice. I shall not address this condition in this paper. Instead, I will assume that agents have true beliefs about the world—and that any other relevant conditions on belief are met. (This does not require that, where there is objective uncertainty, the agent know the eventual outcome. It only requires

that she know the objective chances.)

The rough model, then, is this: We consider only cases where agents have made an autonomous choice (where the exact content of this notion is left open). Roughly speaking, an agent is (broadly) responsible for an outcome to the extent that her choice increased the chance of the outcome and she believed this to be so. The rest of the paper is devoted to refining this account. In the remainder of this section, I shall comment briefly on the appeal to probabilities.

I have assumed that, in any given choice situation, there are *objective probabilities* governing events in the world, including what choices agents make. This is compatible with determinism, since the probabilities may be limited to zero or one. Although I shall assume for illustration that the probabilities are perfectly determinate, this is not essential to the framework. The probabilities in question may be vague intervals. Alternatively, the probabilities may be radically indeterminate in the sense that the only determinate facts are facts about whether a given event is possible or not. The framework can be generalized to accommodate these cases.

The existence of objective (non-trivial) probabilities is, of course, controversial. The existence of objective probabilities governing the *choices* of autonomous agents is even more controversial. I shall not, however, defend these assumptions. My goal here is to develop an account of responsibility for outcomes on the assumption that there are such probabilities. If there are no relevant (perhaps vague) probabilities governing choice (and other events), then the framework below is a non-starter.[18]

5. The Basic Model

Suppose, then, that an agent is in her first choice situation (so that we need not worry about responsibility in virtue of past choices). There were, we are assuming, objective probabilities

(including conditional probabilities) governing, in that situation, how the world might turn out. These probabilities reflect the choice disposition of the agent to choose in that situation and in future situations, the choice dispositions of other agents (now and in the future), and nature's "choice" dispositions (i.e., the probabilities of various non-choice events occurring).[19] We now want to determine which outcomes (events and states of the world) are attributable to the agent's choice and which are a matter of brute luck. The model below will appeal to four distinct objective probability functions:

P_i: the *initial* objective probability distribution over events (including choices) just prior to the agent's choice.

P_a: the initial objective probability distribution over events (including choices) just prior to the agent's choice *conditional on the agent making an autonomous choice*.

P_c: the initial objective probability distribution over events (including choices) just prior to the agent's choice *conditional on the agent making the specific autonomous choice* that she made.

P_s: the objective probability distribution over events (including choices) at some *subsequent* time, s, relative to which responsibility is being assessed.

P_i gives the initial probabilities of outcomes. In a given situation, there may be some chance that the agent will not make a choice at all (e.g., have a seizure) or make a choice that is not autonomous (e.g., as the result of overwhelming panic or direct mental manipulation by someone else). If this chance is realized, then the agent bears no responsibility for outcomes in virtue of a choice at that time. Responsibility for outcomes arises only in virtue of autonomous

choices. We shall therefore restrict our attention throughout to situations in which an autonomous choice is made. P_a gives the initial probabilities conditional on an autonomous choice being made. It seems natural to assume that, when an agent makes an autonomous choice, the fact that an autonomous choice was made (rather than no such choice being made) is not attributable to that very choice. For example, the fact that there was no seizure, no panic, and no direct mental manipulation by another is not attributable to the choice that the agent made (although it may be attributable to prior choices, if there were any). Given this assumption, $P_a - P_i$ is the shift in probabilities that is attributable to the brute luck event of the agent making an autonomous choice (of some sort). P_c gives the initial probabilities conditionalized on the autonomous choice that the agent actually made.[20] $P_c - P_a$ is thus the shift in probabilities that is attributable to the specific choice. It is the (probabilistic) impact of that choice. P_s gives the probabilities at some subsequent time relative to which responsibility is to be assessed. For past events, it ascribes either zero or one, since they are then fixed. For future events, it ascribes the probabilities of how things might go (which will also be zero or one, if determinism is true).

Our question is the following: Assuming that an autonomous choice was made, which aspects of the change from P_i to P_s are attributable to the agent's choice (with the agent being responsible for those changes) and which aspects are a matter of brute luck (with the agent not being responsible for those changes)?

The following gives the schematic answer, which will be developed in the rest of the paper.

	Total Situation	Attributable to Choice	Attributable to Brute Luck
Initial Situation	P_i	n/a	initial brute luck

Occurrence of Autonomous Choice			
	$P_a - P_i$	n/a	occurrence of autonomy
Impact of Choice	$P_c - P_a$	foreseen impact	unforeseen impact
Outcome Luck	$P_s - P_c$	option luck	brute outcome luck
Total Outcome (at s)	P_s	Attributable to Choice	Brute Luck

The initial situation is P_i and, given that the agent made no prior choices (which we are assuming for simplicity), this is entirely attributable to brute luck. This is *initial brute luck* (brute luck of initial circumstances, including the agent's initial endowment). If the agent does not make an autonomous choice, then no outcomes are attributable to choice in that situation and we are done (and can ignore the above table). If the agent does make an autonomous choice, c, then the shift in probabilities from P_i to P_c can be accounted for as follows. The impact of the occurrence of autonomous choice (e.g., as opposed to a seizure) is attributable to brute luck and is $P_a - P_i$. The impact of the specific choice (given autonomous choice) is the shift in probabilities from P_a to P_c. The portion of this impact that was foreseen by the agent (satisfying the belief condition) is attributable to choice, whereas the portion that was unforeseen is attributable to brute luck. Finally, there is a shift in probabilities from P_c (that induced by the choice) to P_s (the subsequent probability distribution). As we shall see below, some of this shift may be suitably related to the foreseen impact of choice and thus attributable to choice (known as "option luck"). Some of the shift may be unrelated and thus attributable to brute luck (brute outcome luck).

In order to simplify matters in what follows, I shall ignore the brute luck of the autonomous choice occurring (i.e., $P_a - P_i$). Instead, I shall focus on the issue of allocating responsibility for outcomes, given that an autonomous choice is made (i.e., allocating

responsibility for the shift from P_a to P_s). For simplicity, we shall assume that, in the given situation, it is certain that the agent will make an autonomous choice of some sort and thus $P_a=P_i$.

The remainder of this paper is concerned with determining how exactly the above attributions are made. We shall consider two cases: (1) choice under certainty (i.e., choices fully determine outcomes) with complete and true beliefs and (2) choice under risk (i.e., choices determine less than certain probabilities of outcomes) with complete and true beliefs. I shall not attempt here to deal with responsibility when the agent has false beliefs.[21]

6. Choice under Certainty with Complete and True Beliefs

Suppose that an agent has complete and true beliefs about the world and that her choices fully determine the outcomes. She has a choice of making a loan and ending with $2 the next day or not making the loan and ending with $1 the next. For simplicity, throughout, I assume that outcomes consist simply of dollar allocations to individuals. She is 90% disposed to make the loan and 10% disposed not to. She makes the loan and ends up with $2. To what extent is this outcome attributable to her choice and to what extent is it attributable to her brute luck?

The relevant probability functions and probability shifts are as follows, where <.9, $2; .1, $1> represents a 90% chance of $2 and 10% chance of $1:

P_i = <.9, $2; .1, $1>

P_c = <1.0, $2; 0.0, $1>

$P_c - P_i$ = <.1, $2; −.1, $1>

P_s = <1.0, $2; 0.0, $1>

$P_s - P_c = 0$

The initial situation, that is, was <.9, $2; .1, $1>. The agent's choice shifted this to <1.0, $2; 0.0, $1> (i.e., $2 for sure). Because this is a case of choice under certainty, there is no luck in how things turn out once the choice is made. Thus, P_s is the same as P_c. When we examine choice under risk, this will cease to be so.

The model analyzes this as follows:

	Attributable to Choice	Attributable to Brute Luck
Initial Situation	n/a	<.9, $2; .1, $1>
Impact of Choice	<.1, $2; −.1, $1>	n/a
Outcome Luck	n/a	n/a
Total Outcome (at s)	<.1, $2; −.1, $1>	<.9, $2; .1, $1>

Example 1

The initial situation is clearly a matter of brute luck, since the agent has made no prior choices. Given that the agent has complete and true beliefs, the impact of choice—the shift in the probabilities due to choice—is clearly attributable to choice. Given that there is no outcome luck (choice under certainty), that's all there is. The final outcome of $2 thus consists of brute luck of <.9, $2; .1, $1> (the initial probability distribution) and the impact of choice of <.1, $2; −.1, $1>. The two add up to <1.0, $2; 0.0, $1>, which is the final outcome of $2 for certain.[22]

In this case, the impact of choice is relatively small (.1 shift), since the agent was, as a matter of initial brute luck, highly disposed to make the choice in question. By contrast, if, in the

above situation, the agent chose *not* to make the loan, the impact of choice would be <−.9, $2; .9, $1>, which is a large shift from the initial situation. In this case, the outcome of $1 consists of brute luck of <.9, $2; .1, $1> (same as before) and the impact of choice of <−.9, $2; .9, $1>. Note that the two add up to <0.0, $2; 1.0, $1>, which is the final outcome of $1 for certain.

The basic idea is that the agent's initial situation (including her initial choice disposition) is a matter of brute luck. The difference that her choice makes to the situation (and only that difference) is attributable to her choice.

An important aspect of this model is that it views responsibility for an outcome as a *matter of degree*—as reflected in probability shifts. Responsibility for an outcome is not an all-or-nothing affair. An agent can be partly responsible for her lung cancer (due to her choice to smoke) without being fully responsible (e.g., due to genetic vulnerability). Indeed, although it has generally been recognized that the responsibility/brute luck distinction must come in degrees, to the best of my knowledge, the proposed model is the first philosophical account to capture this feature. I claim that it does so in an appropriate manner. (An agent can also be partially responsible for an outcome due to partial autonomy of choice or less then full beliefs. I do not address these issues here.)

Given our (controversial) framework assumption that agents are responsible only for outcomes for which their choice increases the probability, the above model may seem relatively uncontroversial in the case of choice under certainty with complete and true beliefs. Let us, however, pause to consider some of its (sometimes controversial) implications.

The model assumes that an agent is responsible for an outcome only to the extent that her choice increased the probability of the outcome. This entails that an agent is to some extent responsible for an outcome only if her choice was not determined by prior events. For, if her

choice was already determined by prior events, its occurrence did not alter the probabilities of any events (p_c would equal p_i). Both the choice and the outcome are simply common effects of the prior events, given the laws of nature. The occurrence of the choice does not alter the probability of the outcome. This is not to deny that the outcomes might have been quite different had the agent made a different choice. This thought, however, appeals to a possibility that is, given the circumstances, nomically impossible. If a choice is determined by prior events, given the laws of nature, then there is no nomic possibility, just prior to her choice, of her choosing otherwise. As indicated above, I here assume that responsibility in virtue of a given choice is based on the full facts just prior to choice. This is, of course, a controversial assumption—generally endorsed by incompatibilists and generally rejected by compatibilists, but once it is made, the incompatibility of responsibility with full prior determination of choice follows immediately. I shall not here attempt to defend the more general assumption.

As just noted, the model rules out responsibility in virtue of a choice when the choice was fully determined by prior events. A second feature of the model is that it holds that an agent *can* be responsible when the choice is *not* fully determined by prior events. It thus faces the following standard objection. Given that the choice was not determined by prior events, its occurrence seems random or arbitrary. It is implausible, it is claimed, that an agent can be responsible for an outcome in virtue of some random or arbitrary event.[23] The objection is well known—as is its reply. The reply is that the fact that a choice was not fully determined by prior events does not entail that it is arbitrary or random in any sense that undermines responsibility. One of the conditions for responsibility is that the choice be made autonomously. I have left open the details of what this requires, but at a bare minimum it requires that the choice be made on the basis of a (perhaps implicit) reflective assessment of the agent's beliefs, values, and

intentions. A choice made in this way is not arbitrary or random, even if it is not determined by prior events. It is made deliberately and for reasons.[24] The issue is complex and there is much disagreement among experts on this issue. Because I have nothing to add to the debate, I merely flag it to set it aside.

A third feature of the model concerns the manner in which it appeals to shifts in probabilities as a way of accounting for partial responsibility. Even if one accepts this general idea, one might question whether the appropriate measure is the shift in responsibility from the prior probabilities *conditional on autonomous choice*. A common alternative view is that the baseline should be something like the probabilities *conditional on the agent making no choice* (or taking no action). This view holds agents responsible only for the impact of their choices compared with the probabilities of events where the individual does not exercise her agency at all. There is admittedly something initially intuitively attractive about this view. The problem is that agents have different choice-making and choice implementing capacities and potentials—and often (but not always) this is not attributable to any prior choice that they have made. The above absence-of-choice model of responsibility is not sensitive to these differences—since it focuses on the probabilities of events in the absence of the agent making a choice and ignores the very real and different probabilities of the agents making various choices. Consider, for example, two identical agents in identical circumstances, except that, due to a congenital disorder, one is much less able to defer gratification. Suppose they make the same choice. The above absence of choice model holds that they have the same responsibility for the outcome (in virtue of those choices). It thereby fails to recognize that a significant component of brute luck is one's initial choice-making and choice-implementing capacities and potentials. By contrast, the proposed model correctly recognizes these facts. It is sensitive to the way in which the brute luck of

choice-making abilities influences the probabilities of various choices. It holds that, when two agents fail to defer gratification, an agent who is less disposed to defer gratification is less responsible for the outcome that results.

Although the proposed model treats the agent's initial situation—and her choice-making and choice-implementing capacities, in particular—as a matter of brute luck, this leaves room for outcomes being primarily attributable to choices. It is true that the model has the implication that where the agent was antecedently highly disposed to make a certain choice, she bears only a little responsibility for the choice and its outcomes. This may seem to leave too little room for responsibility, but this is not so. First, as indicated earlier, where the agent is antecedently highly disposed to make a certain choice and she makes a *different* choice, the model entails that she is highly responsible. Second, in real life, agents have lots of options, and the probabilities may not be typically highly focused on any one choice (e.g., there may be 10 choices with probability 10% each). This increases the room for attributability to choice. Finally, and most importantly, even if an agent always chooses that which she is antecedently highly disposed to choose, she may still bear, on the proposed model, significant responsibility for the outcomes. This is because responsibility is *historical*: one bears responsibility for outcomes—not only in virtue of one's current choice but also—in virtue of one's past choices. For simplicity in this paper, we are focusing on cases where the agent has made no prior choices. Here I temporarily drop this assumption so as to comment on the role that history of choice can play.

Suppose that, for 32 seconds, at each second an agent faces a choice between pushing a small button and pushing a big button. If she pushes the small button at any point, she receives $1 at the end of the game and faces no further choices until then. If she pushes the big button 32 times in a row, then she wins $1000. At each point, her disposition is 90% to push the big button

and 10% to push the small button. Suppose that she pushes the big button each time and wins $1000. Here her initial brute luck gives her a $(.9)^{32}$ chance of $1000 and a $1-(.9)^{32}$ chance of $1. This is roughly a 4% chance of $1000 and a 96% chance of $1. Thus, the net result of her 32 choices is to shift the probability of $1000 from 4% to 100%. This is a significant attribution of responsibility, even though each of the 32 choices had only a small impact (e.g., on average involved a shift of only 3%). The general point is that, even if the initial situation highly disposes an agent to choose in certain ways and the agent so chooses (which minimizes the impact of choice), if there are enough choices of the right sort in the past, the net impact of all choices can swamp the impact of initial brute luck. That is, although the agent may bear only slight responsibility for an outcome in virtue of any given choice, there may be sufficiently many such choices so that the agent bears significant responsibility for the outcome (in virtue of all her choices). Moreover, in the example just given, for simplicity, I did not appeal to the way in which one's past choices can *influence one's later choice dispositions* (by helping shape one's character). When this is factored in, the effect of past choices can be even more pronounced.[25]

A fourth feature of the model concerns how it handles *threats and offers*—such as the threat that one will be killed if one does not hand over one's wallet. Sometimes, threats may be so overwhelming that they prevent the possibility of an autonomous choice (e.g., where the agent goes into an overwhelming panic). In such cases, threats do indeed eliminate any responsibility for the outcome and the model recognizes this. In many cases, however, they do not. Even here, however, sometimes, a credible extreme threat (e.g., to kill one) may make it virtually certain that the agent perform the desired action. When this is so, and the agent complies, she is responsible at most for only a very small shift of probabilities. Not all cases of threats, however, undermine autonomous choice or make it virtually certain what the agent will choose. For

example, a credible threat that one will be slapped if one does not hand over one's wallet typically has little impact on one's choice disposition and the threatened agent can be responsible for perhaps a significant shift in probabilities if she hands over her wallet. This is not, of course, to deny that the slapped agent was wronged by the slapping threatener and is owed compensation. It is just to insist that the agent bears some responsibility for the outcome as compared with a similarly situated agent who did not hand over his wallet. In short, the proposed model views the shift in the disposition to choose as a result of a coercive threat as a matter of brute luck. The agent is nonetheless responsible (given we are here assuming complete and true beliefs) for the shift in probabilities from that new disposition that her choice induces. This seems to recognize appropriately that threats typically reduce responsibility without eliminating it. Similar comments apply to offers. On the proposed model, there is no difference between the two with respect to responsibility (which is not to claim that there is no difference with respect to moral permissibility and moral accountability).

A fifth feature of the model concerns how it handles causal overdetermination—where one or more agents independently make choices each of which independently makes it a certainty that some outcome will occur. Suppose, for example, that you and I each fire a gun, at close range, at Victim. Each of our bullets strikes his heart, thereby ensuring his death. Who is responsible for his death and to what extent? If our *choices were fully determined* (i.e., with probability 1) by the laws of nature and past events (e.g., if determinism is true), then neither of our choices had any effect on the probability of Victim's death. The proposed model thus holds that neither of us is responsible for the death. Let us therefore suppose that neither of our choices was determined by prior events. On the proposed model, we shall now see, the *temporal order* of the choices is crucial for assessing responsibility.

Suppose that I choose to fire first and increase the probability of Victim's death by gunshot at this time from 1% to 100%. I am responsible for this 99% increase. When you later choose to fire, it is already causally determined that Victim will die by gunshot at this time. Thus, you bear no responsibility for that event. Of course, you bear responsibility for intending to kill him, for firing your gun at him, and for *your* bullet striking his heart. You are not, however, responsible for his death.[26]

Suppose, now, that we choose to fire our guns *simultaneously*. Here, the proposed model holds that we may *each* be responsible for an increase from, for example, roughly 2% to 100% in the probability of Victim now dying from gunshot. Suppose prior to my (and your) choice to fire, there was a 1% chance that I would shoot Victim, a 1% chance that you would choose to fire, and no chance of Victim being shot by other means. Thus, just prior to our choosing to fire, there was a 1.99% chance that Victim would then be shot (.99% chance of being shot only by me, .99% chance of being shot only by you, and .01% chance of being shot by both)[27]. We each fire and Smith dies. We are each responsible for increasing the chance of his dying now by gunshot from 1.99% to 100%. This may seem like double-counting, but it is not. We do not add the roughly 98 percentage point increases. Rather, it is simply that, for each of us, relative to the facts just prior to our choice (which include the very low disposition of each to choose to fire), our choice to fire increases the chance of Victim then dying by gunshot by about 98 percentage points. We are each almost fully responsible for his death, and that seems quite right.[28]

Several of these features of the model are controversial. If, however, one accepts that the relevant causal condition for responsibility is the probability-increasing requirement, then the implications seem quite plausible. My goal here is to develop an account of partial responsibility for outcomes and for concreteness I assume the probability-increasing requirement (which I

believe plausible). My hope is that similar ideas about partial responsibility can be developed for other causal requirements, but that project is beyond the scope of this paper.

So far, we have only considered the case of choice under certainty (no outcome luck). Let us now turn to the more complex case involving choice under risk (where the outcome is not fully determined by choice). We continue to assume full true beliefs.

7. Choice under Risk with Complete and True Beliefs

Suppose that an agent has, at the time of her choice, complete and true beliefs about the world but her choices do not fully determine the outcomes. Instead, her choices only determine the probabilities of outcomes. The final outcome depends on what other events and choices occur. Because, at the time of her choice, the only facts about how things will turn out, if she makes a given choice, are probabilistic facts, the agent knows these facts, but she does not know how things will turn out *tout court*.

Suppose, for example, that the agent has a choice between (1) buying a $1 lottery ticket that gives her a 1% chance of ending with $200 and a 99% chance of ending with $0 and (2) keeping her $1. Suppose that she is 90% disposed to buy the ticket and 10% disposed not to. She buys the tickets and ends up with $200. To what extent is this outcome attributable to her choice and to what extent is it attributable to her brute luck?

The relevant probability functions and probability shift functions are as follows:

P_i = <.9, <.01, $200; .99, $0>; .1, $1>

P_c = <.01, $200; .99, $0>

$P_c – P_i$ = <.1, <.01, $200; .99, $0>; −.1, $1>

P_s = <1.0, $200; 0.0, $0>

$P_s - P_c$ = <.99, $200; −.99, $0>

Unlike the case under certainty, P_s (the subsequent situation) is different from P_c (the probability distribution induced by the choice). This is because there is *outcome luck* in how things turn out relative to the choice. The impact of the choice, $P_c - P_i$, was to increase the chance of buying a lottery ticket from 90% (the initial disposition) to 100%.[29] The resulting probability distribution, P_c, was the 1% chance of ending with $200 and the 99% chance of ending with $0. The outcome luck, $P_s - P_c$, increased the chance of ending with $200 from 1% to 100%. The total impact of the initial situation, P_i, the impact of choice, $P_c - P_i$, and outcome luck, $P_s - P_c$, is P_s, the subsequent situation (i.e., $200 for certain).

We must now ask the question of how the total outcome luck—<.99, $200; −.99, $0>—is allocated between choice and brute luck. Given that exposure to this outcome luck—<.01, $200; .99, $0>—is the result of both the initial situation, <.9, <.01, $200; .99, $0>; .1, $1>, and the foreseen impact of choice, <.1, <.01, $200; .99, $0>; −.1, $1>, it is implausible that the outcome luck should be entirely attributable to either. Instead, it seems reasonable that it should be allocated between choice and brute luck in proportion to their contribution to the chance that the outcome would occur. In this example (where there is a 90% initial disposition to buy the ticket), the initial situation contributes 90%, and choice 10%, of the 1% chance of winning $200. Thus, 90% of the outcome luck of <.99, $200; −.99, $0> is attributable to brute luck and 10% is attributable to choice. That is, brute outcome luck is <.891, $200; −.891, $0> and the outcome luck attributable to choice is <.099, $200; −.099, $0>.

The model thus analyzes the situation as follows:

	Choice	Brute Luck
Initial Situation	n/a	<.9, <.01, $200; .99, $0>; .1, $1>
Impact of Choice	<.1, <.01, $200; .99, $0>; −.1, $1>	n/a
Outcome Luck	<.099, $200; −.099, $0>	<.891, $200; −.891, $0>
Total Outcome (at s)	<.1, $200; −.1, $1>	<.9, $200; .1, $1>

Example 2

Both the foreseen impact of choice and the associated outcome luck are attributed to choice (<.1, $200; −.1, $1>). The initial situation and the associated outcome luck are attributed to brute luck <.9, $200; .1, $1>. The total effect of the two is <1.0, $200; 0.0, $1>, or $200 for certain.

The proportional allocation of outcome luck to choice and brute luck is, I believe, highly plausible. It has, however, one implication that is generally rejected. Suppose that there are two identical agents in the above situation, both choose to buy lottery tickets, but only one wins. The generally accepted view is that, given that there was background equality, the outcome luck is entirely attributable to choice (option luck). The above model, however, says that the outcome luck is only partially attributable to choice. Indeed, in the above example (where the individuals were 90% disposed to buy the lottery tickets), it says that only 10% of the outcome luck is attributable to choice.[30]

We can make this more concrete by applying brute luck egalitarianism to the example. It holds that the net advantage of brute luck should be equalized but that the net advantage from choice need not be equalized. For illustration, I will assume that the value of probability distributions and of probability shifts is based on their expected monetary value. The two agents had the same initial brute luck; so we can focus on outcome brute luck. Because 90% of the

outcome luck is attributed to brute luck (given the initial 90% disposition to buy the ticket), the first agent (as noted above) has brute outcome luck of <.891, $200; −.891, $0>, which has an expected value of $178.2. The second agent (who bought a ticket but won nothing) had total outcome luck of <−.01, $200; .01, $0>. Given the 90% initial choice disposition, 90% of this is attributable to brute luck, which yields <−.009, $200; .009, $0>, which has an expected value of −$1.8. Thus, the total outcome brute luck for the two of them is $176.4 and a transfer of $90 from the first to the second agent is required to equalize their outcome brute luck advantage (giving them each $88.2.) The common view of brute luck egalitarians, however, is that, given the background equality, the difference in outcome is *entirely* attributable to choice and no transfers are required to equalize brute luck.

I believe that the model is correct and that the common view is mistaken. Both views agree that, for choices under *certainty*, if the initial situations are equal, then all differences in outcome are attributable to choice. For example, in the example of the previous section (where each of two agents are 90% disposed to make a loan and end with $2 and 10% disposed to simply keep their $1), their initial situation is equal (<.9, $2; .1, $1>, the impact of making the loan is <.1, $2; −.1, $1>, and the impact of not making the loan is <−.9, $2; .9, $1>. One agent makes the loan (and ends with $2) and the other doesn't (and ends with $1). The difference between the two impacts of choice is <1.0, $2; −1.0, $1>, which fully accounts for the $1 difference in the two outcomes. Thus, the model fully agrees with the common view when choice is under certainty. The issue here concerns choice under *risk*. The question is how *outcome luck* should be apportioned between choice and brute luck. Given that in this section we are here assuming complete and true beliefs, the question is how outcome luck should be apportioned on the basis of the foreseen impact of choice and the brute luck initial situation.

To help see the superior plausibility of the model over the common view, consider a different example. Here two identically situated agents (with complete and true beliefs) choose which of two paths to take, each is 99% disposed to take the short path, and each chooses the short path. Moreover, on the short path, there is a .01% chance of being attacked by thugs but no chance of such an attack on the long path. One agent is attacked, and the other is not. According to the common view, because there was background equality and each agent chose in full awareness of the chance of being attacked by thugs, the difference in outcome luck is *entirely* attributable to choice. If, however, each agent was 100% (as opposed to 99%) disposed to take the short path (e.g., because the longer path was blocked by a landslide), then the common view (at least among incompatibilists) is that the entire outcome (including outcome luck) is due to brute luck (since there was no real choice). It is, however, quite implausible, that a small difference in dispositions to choose should convert outcome luck from being entirely attributable to choice to being entirely attributable to brute luck.

The proposed account does not have this problem. It agrees that, where the initial choice disposition was 100%, then the outcome luck (and indeed the entire outcome) is entirely attributable to brute luck. It holds, however, (1) that where the initial choice disposition was 99%, outcome luck is 99% attributable to brute luck, and (2) that where the disposition was 1%, outcome luck is 1% attributable to brute luck. In sum, the proposed view agrees with the common view that, where there is equality of initial brute luck and there is no outcome luck, then the differences in outcome are entirely attributable to choice. Where there is outcome luck, however, it holds that the outcome luck must be apportioned between choice and brute luck and not merely attributed entirely to choice—even where there is background brute luck equality.

Although I have not said enough to establish the superiority of the proposed account with

respect to this issue, I hope that I have said enough to establish that it is not clearly inferior and is worthy of further consideration.

8. Conclusion

I have addressed the issue of partial responsibility for outcomes where an agent's choice is just one of the factors leading to the outcome. In order to make the discussion manageable, I made a number of specific, but controversial, assumptions. In particular, I assumed that one is responsible for an outcome only to the extent that one increases its objective probability. Obviously, many people will reject this or other of my assumptions. My hope, however, is that some of the central claims of the paper (perhaps with modification) will remain valid in the context of alternative assumptions. I have not, however, argued for this.

Partial responsibility due to partial causal responsibility, I have suggested, can be accounted for by the extent of the shift in the objective probability of the outcome in question. The baseline for the shift, on the proposed model, is the prior probability distribution over outcomes conditional on an autonomous choice being made (as opposed to no choice and as opposed to some [perhaps, non-autonomous] choice being made). Such a view, I argued, is appropriately sensitive to the agent's prior choice disposition, which, for the first choice, is a matter of brute luck. Choice dispositions, of course, can be influenced by prior choices and a full account of responsibility for outcomes must factor that in. For simplicity, I have focused on cases where the agent is making her first choice.

My goal has been to open up a framework for investigation and not to conclusively defend it. The proposed model is schematic and full of unresolved issues. Still, I hope that it has enough promise to warrant further investigation.

Notes

For helpful comments, I thank Chris Beltram, Matthew Braham, Sara Chant, Zac Ernst, Paul Litton, John Martin Fischer, Kasper Lippert-Rasmussen, Kevin McCain, Mike Otuska, Garrett Pendergraft, Jon Quong, Patrick Todd, Fritz Warfield, Paul Weirich, John Weymark, Karl Widerquist, Andrew Williams, Michael Zimmerman, and the audiences at the *Politics, Philosophy, & Economics* conference sponsored by Tulane University, University of Manchester conference on responsibility, University of Bristol, Oxford University, University College London, Reading University, University of York, and the University of Edinburgh.

---

[1] John Martin Fischer and Mark Ravizza, *Responsibility and Control: A Theory of Moral Responsibility* (Cambridge: Cambridge University Press, 1999). See also Michael J. Zimmerman, *An Essay on Moral Responsibility* (Totowa: Rowman and Littlefield, 1988).

[2] See, for example, Ronald Dworkin, "What is Equality? Part 2: Equality of Resources," *Philosophy and Public Affairs* 10 (1981): 283-345 (reprinted with minor changes in Ronald Dworkin, *Sovereign Virtue*, Cambridge, MA: Harvard University Press, 2000); Richard Arneson, "Equality and Equal Opportunity for Welfare," *Philosophical Studies* 56 (1989): 77-93; G.A. Cohen, "On the Currency of Egalitarian Justice," *Ethics* 99 (1989): 906-44; Kasper Lippert-Rasmussen, "Egalitarianism, Option Luck, and Responsibility," *Ethics* 111 (2001): 548-579; Michael Otsuka, "Luck, Insurance, and Equality," *Ethics* 112 (2002): 40-54; and Peter Vallentyne, "Brute Luck, Option Luck, and Equality of Initial Opportunities," *Ethics* 112 (2002): 529-557.

[3] Susan Hurley, *Justice, Luck, and Knowledge* (Cambridge, MA: Harvard University Press, 2003).

[4] For example, T.M. Scanlon, "The Significance of Choice," in *The Tanner Lectures on Human*

Value VIII, ed. Sterling McMurrin (Cambridge: Cambridge University Press, 1988), pp. 151-216; T.M. Scanlon, *What We Owe to Each Other* (Cambridge, Mass.: Belknap Press of Harvard University Press, 1998), ch. 6.

[5] See for example, Peter Strawson, "Freedom and Resentment," *Proceedings of the British Academy* 48 (1962): 187-211; and R. Jay Wallace, *Responsibility and the Moral Sentiments* (Cambridge, MA: Harvard University Press, 1994)—although his view seems to have elements of the normative conception, discussed next. They both focus on responsibility for choices, but the general views can be naturally extended to responsibility for outcomes.

[6] See, for example, Joel Feinberg, "Action and Responsibility," *Philosophy in America* edited by Max Black (London: George Allen & Unwin, 1965), pp. 134-60 [Reprinted in Joel Feinberg, *Doing and Deserving: Essays in the Theory of Responsibility* (Princeton: Princeton University Press, 1970), pp. 119-51]; Jonathan Glover, *Responsibility* (London: Routledge and Kegan Paul, 1970), p. 44; Thomas Nagel, "Moral Luck," *Proceedings of the Aristotelian Society*, Supplementary Volume 50 (1976): 137-51; Michael J. Zimmerman, *An Essay on Moral Responsibility* (Totowa: Rowman and Littlefield, 1988); and Ishtiyaque Haji, *Moral Appraisability: Puzzles, Proposals, and Perplexities* (Oxford: Oxford University Press, 1998).

[7] See, Fischer and Ravizza, *Responsibility and Control*, pp. 76-81; and Wallace, *Responsibility and the Moral Sentiments*, pp. 154-91.

[8] This view is also held by Michael J. Zimmerman, "Taking Luck Seriously," *Journal of Philosophy*, 99 (2002): 553-576; Michael J. Zimmerman, "Luck and Responsibility," *Ethics* 97 (1987): 374-86; and Judith Jarvis Thomson, "Morality and Bad Luck," *Metaphilosophy* 20 (1989): 203-221 [Reprinted in *Moral Luck*, edited by Daniel Statman (Albany: State University of New

York, 1993)].

[9] Narrow responsibility is at least roughly what is called “direct responsibility” in Michael Zimmerman, *An Essay on Moral Responsibility*.

[10] Broad responsibility is at least roughly what is called “indirect responsibility” in Michael Zimmerman, *An Essay on Moral Responsibility*.

[11] It is also the notion of responsibility that Fischer and Ravizza address in *Responsibility and Control.* They call it an “externalist” account (pp. 252-253).

[12] See, for example, Fischer and Ravizza, *Responsibility and Control*, ch. 7; Robert Kane, *The Significance of Free Will* (Oxford: Oxford University Press, 1998), ch. 3.

[13] Even compatibilist Fischer and Ravizza in *Responsibility and Control* (e.g., p. 102) agree that agents are not responsible for certain kinds of nomically inevitable outcomes.

[14] Appealing to propensities that are probabilities is, of course, controversial. For a defense, see the papers in D.M. Mellor, *The Facts of Causation* (Cambridge: Cambridge University, 1995) and ch. 9 of D.W. Miller, *Critical Rationalism: A Restatement and Defence* (Chicago and Lasalle, Il: Open Court, 1994). For a summary of the criticisms, and further references, see Alan Hájek, "Interpretations of Probability", *The Stanford Encyclopedia of Philosophy (*Summer 2003 Edition*)*, Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/sum2003/entries/probability-interpret/>.

[15] About a year after the ideas of this paper were first presented, I discovered that a similar appeal to the marginal probabilistic impact is made by Mario J. Rizzo and Frank S. Arnold, “Causal Apportionment in the Law of Torts: An Economic Theory,” *Columbia Law Review* 80 (1980): 1399-429. There are, however, several important differences: For example, (1) they are

concerned with liability to compensate for harms rather than for agent-responsibility. Moreover,they assume strict liability (according to which the agent's mental state is irrelevant). (2) They invoke epistemic probabilities rather than objective probabilities. (3) They take absence of choice as the relevant probability baseline rather than reflecting the agent's prior choice disposition. For more on probabilistic causation generally, see Ellery Eells, *Probabilistic Causality* (Cambridge, U.K.: Cambridge University Press, 1991); and Christopher Hitchcock, "Probabilistic Causation", *Stanford Encyclopedia of Philosophy* (Fall 2006 edition), URL = http://plato.stanford.edu/entries/causation-probabilistic/.

[16] Harry Frankfurt, "Alternative Possibilities and Moral Responsibility," *Journal of Philosophy* 66 (1969): 829-39; and Fischer and Ravizza, *Responsibility and Control.* The most powerful argument for the view that one is not responsible for that which one cannot influence is given by Peter Van Inwagen, "Ability and Responsibility," *The Philosophical Review* 87 (1978): 201-24; and Peter Van Inwagen, *An Essay on Free Will* (Oxford: Oxford University Press, 1986), ch. 6. A sophisticated (but mistaken, in my view) account of how, under very special circumstances, one could be responsible for something that one cannot influence is given by William Rowe, "Causing and Being Responsible for What Is Inevitable," *American Philosophical Quarterly* 26: 153-59.

[17] I would argue, however, that agents are not strictly responsible for harms that they (reasonably) *believe* are unavoidable.

[18] Of course, many theories of free will or responsibility assume, as I do, that there are objective probabilities governing choice. See, for example, Robert Kane, *The Significance of Free Will*, ch. 10. On the subject of subjective probabilities concerning one's choices, Isaac Levi has argued that

agents cannot assign any such probabilities—at least not to their current choices while deliberating about them. See, Isaac Levi, "Rationality, Prediction, and Autonomous Choice," *Canadian Journal of Philosophy*, 19 (1989): 339-362 [reprinted in Isaac Levi, *The Covenant of Reason* (Cambridge: Cambridge University Press, 1997), pp. 19-39]. For an opposing (and convincing, to my mind) view, see Wlodek Rabinowicz, "Does Practical Deliberation Crowd Out Self-Prediction?" *Erkenntnis*, 57 (2002): 91-122. For formal models of decision-making involving probabilities of choice, see Charles, F. Manski, "The Structure of Random Utility Models," *Theory and Decision*, 8 (1977): 229-254; and R.D. Luce R. D. and Patrick Suppes, "Preference, Utility, and Subjective Probability," in R. D. Luce, B. Bush, and E. Galanter, Eds., *Handbook of Mathematical Psychology Vol. 3*, (New York: John Wiley and Son, 1965), p. 249-410.

[19] The appeal to the agent's disposition to choose is similar to that invoked in John Roemer "A Pragmatic Theory of Responsibility for the Egalitarian Planner," *Philosophy and Public Affairs* 22 (1993): 146-66 and in John Roemer, *Equality of Opportunity* (Cambridge, MA: Harvard University Press, 1998). There is, however, an important difference. Roemer appeals to the distribution of choices actually made by agents of the same "type" (i.e., same brute luck features) as the agent in question, whereas I appeal to the objective propensity of the agent in question to choose. Roemer, that is, appeals to the relative frequency of choice for a population of individuals of the same brute luck type, whereas I appeal to the single case propensity to choose for the individual at issue.

[20] P_c is *not* the probability distribution over events *just after* the agent's choice. That distribution inappropriately reflects the influence of other concurrent events.

[21] I address the issues of false belief and responsibility in "Responsibility and False Belief",

unpublished.

[22] Note that the impact of choice is a probability shift function, i.e., a function that shifts probability from one outcome to another. A probability function plus a probability shift function is a probability function as long as all resulting probabilities are inclusively between zero and one.

[23] This objection is made by (among others) Alfred R. Mele, "Book Review of *The Significance of Free Will*," *Journal of Philosophy* 95 (1998): 581-84; Alfred R. Mele, "Kane, Luck, and the Significance of Free Will," *Philosophical Explorations* 2 (1999): 96-101; Alfred R. Mele, *Free Will and Luck* (Oxford: Oxford University Press, 2006); and Ishtiyaque Haji, "Libertarianism and the Luck Objection," *Journal of Ethics* 4 (2000): 329-37. See also Carl Ginet, *On Action* (Cambridge: Cambridge University Press, 1990), ch.6; and Randolph Clarke, "Free Choice, Effort, and Wanting More," *Philosophical Explorations* 2 (1999): 20-41.

[24] For the most elaborate (and effective) reply, see Robert Kane, "Responsibility, Luck, and Chance: Reflections on Free Will and Indeterminism," *Journal of Philosophy* 96 (1999): 217-40. See also, Robert Kane, *The Significance of Free Will*, ch. 10; Storrs McCall, *A Model of the Universe* (New York: Oxford University Press, 1994), ch. 9; Peter Van Inwagen, *An Essay on Free Will* (Oxford: Oxford University Press, 1983), ch. 4; and Randolph Clarke, *Libertarian Accounts of Free Will* (Oxford: Oxford University Press, 2003), chs. 4 and 5.

[25] See, for example, Robert Kane, *The Significance of Free Will*, ch. 10; and Timothy O'Connor, "Freedom with a Human Face," *Midwest Studies in Philosophy* 29 (2005): 2007-2027.

[26] The account of liability given by Rizza and Arnold in "Causal Apportionment in the Law of Torts: An Economic Theory" agrees that the temporal order of *choice* matters in the way that I

suggest. The compatibilist account defended by Fischer and Ravizza also makes responsibility depend on temporal order, but (implausibly in my view) it makes it depend on the order in which the bullets (for example) strike the victim—as opposed to the order in which *choices* were made to fire the bullets. See John Martin Fischer and Mark Ravizza, *Responsibility and Control: A Theory of Moral Responsibility*, ch. 4.

[27] The chance of being shot by only me is the probability of my firing (.01) times the probability of your not firing (.99), or .0099 (i.e., .99%). The calculation is the same for the chance of being shot only by you. The probability of being shot by both is the probability of my firing (.01) times the probability your firing (.01), or .0001 (i.e., .01%).

[28] Thus, where the choices are simultaneous and the outcomes are simultaneous, the model agrees with the view of Fischer and Ravizza that both agents can bear near-full responsibility for the outcome. See, John Martin Fischer and Mark Ravizza, *Responsibility and Control: A Theory of Moral Responsibility*, ch. 4. It also agrees with the assessment of Rizzo and Arnold in "Causal Apportionment in the Law of Torts: An Economic Theory".

[29] It's worth noting that if the agent makes a choice that has no impact on the probability of outcomes, i.e., where P_c=P_i, then the agent has no causal impact and bears no responsibility for any outcome.

[30] I thank Chris Beltram for identifying this implication.