Rethinking Moral Luck: What Conditions are Necessary for Moral Responsibility?

Katelyn Petersen

In a Tortoiseshell: *While the entirety of Katie's "Rethinking Moral Luck: What Conditions are Necessary for Moral Responsibility?" is an excellent showcase of how to navigate key terms, this section is particularly special. Here, not only does Katie introduce her own key term (which skillfully arises from the specific problems she identifies with the key terms that already exist in the scholarly conversation) but she also goes on to give a carefully crafted analysis of the key terms that appear within that overarching key term she proposes! This section not only allows Katie's readers to fully understand what her term means but more importantly allows us to really see how her "Revised Control Condition" is in direct conversation with the concerns she addresses in Nagel's "Strong Control Condition" and Rosen's "Moderate Control Condition."*

Excerpt

My Revised Control Condition seeks to narrow the Strong Control Condition while staying on-topic and allowing for these real-world expectations of scalable responsibility:

Revised Control Condition: X is properly blamed for A only inasmuch as X's action(s)

> was/were under X's control and X culpably should have foreseen the possible consequences (A) of those actions, where X's actions are wrongful and contributive to A in the relevant sense.

I'll now proceed to clarify and justify the key terms of this proposed condition: By "properly blamed," I mean *appropriately held morally accountable*[1]. The phrase "only inasmuch" is meant to account for the idea of a sliding scale of responsibility, based on how many people are acting and how much X's actions contribute to the overall consequence, A, among other factors. I will elucidate this concept of scalable responsibility later on in the paper. The phrase "only inasmuch" also indicates that the condition is not meant to be all-

[1] This could be a positive connotation, like "praised," if "wrongful" later on was changed to something like "righteous." However, there is also possibly a different condition for praiseworthiness. In this paper, I will focus on blame for the sake of simplicity.

encompassing; there are likely other conditions for properly applied blame.[2] "An action that is under X's control"[3] means that X is able to impose his will in doing or not doing said action.

By "*culpably should* have foreseen the possible consequences," I mean X morally ought to have foreseen the possible consequences[4] and it would have been reasonable to expect this, so much so that Person X can be held responsible for failing to foresee them[5]. This is different than saying that someone *did* foresee the possible consequences. For example, a person cannot escape judgment via the Revised Control Condition by intentionally keeping their faculties from operating responsibly, thereby avoiding relevant foresight.

Someone who skips school and misses a test is responsible for the failing grade whether she thought through the possible ramifications of her actions or not. For example, she might have simply neglected to consider anything other than what she wanted (the new purse at the mall), which most would agree is wrongful.[6] This inconsideration would not get her out of culpability just because she didn't *actually* foresee the possible consequences.

Moreover, a student who skips school (arguably a wrong action in itself, given her responsibilities as a student and her instructions from authority) on the day that she knows CPR will be taught does not obtain the knowledge necessary to perform the life-saving maneuver correctly. However, when an emergency situation arises later in which she must perform CPR, she is still somewhat blameworthy for killing the patient because she *culpably should have* foreseen that pressing in spot A would help, but pressing in spot B would have harmful effects. She culpably should have foreseen this because they taught it in the class that she skipped despite knowing that she would be taught life-saving skills if she attended it. This is an example of a once-removed (and therefore mitigated) culpability. This removal can feasibly go on indefinitely, with a lessening degree of culpability attributed with each removal.

There's also an importance in probability here – if a particular outcome is very unlikely, and so very easily overlooked as a possible consequence of an action, a person is less culpable for failing to foresee it. For example, her culpability would be greater if she were a lifeguard

---

[2] One example of this might be the Strawsonian condition that an individual must be expressing ill will in his actions in order to be blameworthy. Another important note is that the condition is not meant to be sufficient and there may be exceptions, meaning that perhaps not every individual who satisfies this condition is automatically blameworthy. An example of such an exception may be a child.

[3] While an analysis of control would be an interesting discussion relevant to this paper, it is not what I aim to accomplish here; this paper is instead an analysis *using* control.

[4] I am assuming a morally objectivist premise, reflected in this "should." It can be said that a person should do something whether or not that person knows that he or she should.

[5] Culpable: meets both my revised control condition and whatever other conditions are sufficient for culpability. "Can be held responsible" sounds very similar to "is properly blamed." This is a hint that the Revised Control Condition is self-iterating: it will likely be used to determine culpability of foresight, which can then effect a larger question of responsibility.

[6] It seems that neglecting to think about anything besides what one wants is wrongful at least in part *because* it often leads to negative consequences. See note on wrongfulness' dependence on possible effects on the next page.

skipping CPR training for lifeguards; in that case, the likelihood of her needing that knowledge to save someone's life later is higher, and she would be more blameworthy for the ultimate death because of this increased culpability of her lack of foresight.

The first clarification about X's actions, that they must be wrongful, hinges on the premise that there are true moral facts which are independent of an individual's inclinations or beliefs. In other words, I am assuming that the moral status of an action can be objectively identified.[7] The second clarification, that X's actions must be contributive to A in the relevant sense, means that X's actions must be one or more of the causes which directly lead to A, or one of the causes of these primary causes, and so on.

Both of these qualifications are necessary – if X commits a wrongful action, but that action does not contribute to A in any way, then X cannot be held responsible for A. For example, if Greg ran a red light in Princeton, NJ, and, with no connection whatsoever, Sally shot Bob in Santa Barbara, CA, then Greg could not rightfully be blamed for Bob's death. Alternatively, if X's actions are rightful but they contribute to A, then X cannot be held responsible for A. This is trickier to illustrate, but the most compelling possibility I can think of makes use of the first "skipping school" scenario and relies on the concept that wrongfulness and rightfulness depend on possible outcomes. If a person skipped school because her unique skillset was required to save the world from an approaching asteroid, then skipping school is no longer a wrongful action for her (or, perhaps, is no longer "net" wrongful – the rightfulness of saving the world cancels out the wrongfulness of skipping school). So even though skipping school contributed to her missing the test, she would arguably not deserve a failing grade or any other negative reaction.

The concept of *scalable responsibility* is an important component of the Revised Control Condition, as a solution to problems that arise in Nagel's Strong Condition. It's true that it doesn't make sense to hold someone accountable for something they had *no* control over. But what about something that a person had *some* control over, or contribution to?

For example, if John buys a pack of salami at the grocery store completely at the behest of his own will, he contributes a small amount of profit to a large meat-producing company, relative to their annual profit. However, the proceeds from that small sale combine with proceeds of sales all around the country for the entire year, ultimately facilitating the company's continued abuse of animals through factory farming.[8] John is not then blameworthy for the

[7] While I can't possibly say everything there is to say on moral objectivity here, I'll offer this relevant thought: it seems to me that an action's moral status depends on the likely outcome of the action. This is perhaps a hint of a self-iteration of the Revised Control Condition or its application.

[8] This assumes that animal abuse through factory farming is an objective moral wrong.

millions of cases of animal cruelty that his money went toward (billions, if we're discussing an entire industry). However, he is properly held responsible for *his contribution's consequences*. While the Strong appears to make no allowance for this circumstance, the Revised Control Condition does (in its qualification that someone is properly blamed *only inasmuch* as their wrongful and relevantly contributive actions were under their control).

Mathematically speaking, this instance of scalable responsibility would be a function of the agent's relevant contribution on a scale of 0 to 1 (determined by how many people were acting and how much each person's action contributed, as well, perhaps, as other factors) multiplied by how wrongful the individual's action was. It seems this might be measured, among other indications, by the initial objective immorality of the action (for example, sparking a nuclear war by murdering someone is more blameworthy than sparking the same war by being irresponsible and losing important government paperwork) and the ultimate impact of the said action (how many lives lost, animals abused, etc.).

The second qualification (that someone is properly blamed *inasmuch as* they culpably should have foreseen their actions' possible consequences) can be understood more fully in the following example: when my friend Mary ask me to go on a hike with her yesterday, I declined because I wasn't particularly in the mood to hike, and I wanted to sit in my room and play video games all day. Because I did not go on the hike with her, she did not have the chance to tell me some important information that, unbeknownst to me, she had recently discovered: an evil mastermind was planning a nuclear bomb explosion to end the world. Furthermore, the bomb had been programmed to detonate at the touch of a seemingly random button on a wall in our school's library. Later, I was trying to figure out how to turn on the lights in one of those library classrooms and experimentally pressed that very button. A string of nuclear bombs exploded as a result, bringing the end of the world as we know it. Although my actions had devastatingly far-reaching consequences, my blame in this disaster is mitigated by the fact that I was just one person who did comparatively little to set off the bombs. The mastermind of the plot, who schemed and organized and funded it all, is probably the most to blame. His advisors, engineers, and spies also share a larger portion of the judgment than I do. In addition, my responsibility is also significantly lessened, if not eliminated, by the fact that I had no idea that flipping the light switches in that particular pattern would end the world, *and* I did not know when Mary asked me to go on a hike with her that she would be telling me pertinent and pressing information about a threat to the future of the world. One cannot say that I *culpably should* have foreseen a

nuclear bomb detonation as one of the consequences of my actions. According to the Revised Condition, then, I am not blameworthy for the disaster.[9]

[9] If, however, I had reason to think that she was going to share something important about nuclear bombs and decided not to go on the hike anyway, my culpability would increase (still being mitigated by the fact that I was one small part in a large and complex operation). It is unclear, but interesting to consider, how blameworthy for this incident I would be if I did not know that she was intending to tell me something about nuclear security, but I knew that she was having a difficult time and could really have used a friend's caring presence.

Author Commentary

Katelyn Petersen

One of my favorite things about philosophy is its collaborative nature. Feedback (questions, objections, or challenges) and responses to such feedback are means to the philosopher's ends of progress and clarity. Such back-and-forth, involving explorations of some implications and examples of my position, is how this paper and its key terms developed. Entitled, "Rethinking *Moral Luck*: What Conditions are Necessary for Moral Responsibility?," the paper is a response to two philosophers' positions on the conditions of moral responsibility. In it, I critique Thomas Nagel's and Gideon Rosen's proposed conditions and present my own condition: the Revised Control Condition for blameworthiness, a subset of responsibility. The back-and-forth, then, was in a sense between Nagel and Rosen and me, and more literally between me, my preceptor, and other students.

When I brought the thought-seeds of this paper to my preceptor, they were mere objections to Nagel's and Rosen's positions. Planted in the rich soil of the literature (both Nagel's and Rosen's, with which I disagreed, but which were still compelling and helpful sources), watered and strengthened by contributions and challenges from other interlocutors, the seeds of my objections to others' positions sprouted into a positive position of my own.

My key terms and their explanations were particularly shaped by this collaborative process. I began with objections to Nagel's and Rosen's positions, so in order to enter into the existing conversation, I looked to how Nagel and Rosen were using their terms. Where possible (that is, unless I felt the need to explicitly critique or qualify their usage), I adopted theirs for the sake of consistency. For example, when Rosen presents a restatement of Nagel's Control Condition, he simplifies Nagel's language of "proper" moral judgments to "responsibility," and that is the sense in which I used "responsibility" throughout the paper.

When I presented my own Revised Control Condition, I introduced a few new terms that I thought were relevant to the discussion (for example, "culpable" foresight and "scalable responsibility"). Talking with others about plans and drafts of the paper helped me to see what in these terms was clear in my mind but opaque to a reader, or what I was assuming but in fact needed to be defended. In these discussions of seeking to clarify my intentions, I found real-world examples to be helpful for both me and my interlocutors. Thus John and his salami, my (fictional) friend Mary who wanted company on her hike, and others found their way into the paper's explanation of my key terms.

As I wrote in the paper's conclusion, "I have learned in philosophy that considering the possible implications of an argument is how one discovers counter-examples and refines it." Critically examining Nagel's and Rosen's positions was how I countered theirs and arrived at my own; receiving feedback and critically discussing my own position was how I found examples and counterexamples, clarified my key terms, and refined my argument. From the beginning of brainstorming to the end of the writing process, this paper was a product of discussion and collaboration. Many thanks to all, including Nagel, Rosen, and Samuel Preston, who made that possible.

Editor Commentary

Danielle Hoffman

Katie's essay "Rethinking *Moral Luck*: What Conditions are Necessary for Moral Responsibility?" skillfully navigates key terms on a deeply detail-driven and multifaceted level. We have chosen to excerpt a section from Katie's essay in which she makes the impressive argumentative move of introducing her own key term into the scholarly conversation. However, to truly understand how relevant and successful Katie's treatment of her own key term is, we must first turn to the opening moments of her essay, in which Katie carefully lays out and defines the theories that already exist in relation to moral responsibility that her proposed theory will then directly respond to.

Katie begins her paper by discussing Thomas Nagel's "Strong Control Condition," which he lays out in his *Moral Luck,* and Gideon Rosen's "Moderate Control Condition." It is in Katie's analysis of these terms, specifically in her careful articulation of the problems she thinks exist with these key terms, that she is able to develop such a focused, meaningful motive for her essay's proposal of a revised theory. Katie does not solely define these key terms. Rather, she engages with these key terms by delineating the qualms she has with each of them. For instance, she makes the claim and provides examples to support her claim that Nagel's "Strong Control Condition" is too "broad and general," as when considered in conjunction with "moral luck," it "seems to prevent virtually anyone from being responsible for virtually anything." Katie then enters a discussion of Rosen's theory that seeks to address this problem of Nagel's account, carefully arguing (again through an analysis of key terms!) that Rosen's "Moderate Control Condition" problematically shifts the focus away from "control" and onto "alternative possibilities." I think it is important to focus on this section of Katie's essay, as it is her analysis of and objection to Nagel's and Rosen's theories that allow Katie's own key term to feel so grounded in the specific scholarly debate her essay is aiming to engage with. I think students often think of the various aspects of the lexicon as being separate from one another, in that it is easy to just chronologically go from orientation, to motive, to thesis without really thinking through the connections between these things. However, Katie's essay masterfully conveys the interconnectedness of these components, creating a motive that directly arises out of her analysis of the key terms of her essay! I think this connection between the key terms and motive of Katie's paper is nicely encapsulated by a question she poses: "How can we narrow the control condition so as to salvage the concept of responsibility and avoid the skeptical conclusions that "moral luck" seems to support?"

This is exactly what Katie's "Revised Control Condition" aims to do. As she very explicitly states, "My Revised Control Condition seeks to narrow the Strong Control Condition while staying on-topic and allowing for these real-world expectations of scalable responsibility." The success of this term, though, is not only rooted in how clearly it responds to the weaknesses she has articulated in the already existing terms but also in how deliberately Katie proceeds to go through and fully define and defend each term that appears within this overarching key term. Katie explicitly defines what she means by "properly blamed," "only inasmuch," "an action that is under X's control," and "culpably should have foreseen the possible consequences," enabling her to clarify the exact contexts in which she is employing these terms. And in this layered process of definition and clarification, Katie elucidates the exact meaning of her theory while highlighting how that theory addresses the problems she identifies in Nagel's and Rosen's theories. She also brilliantly anticipates and defends her theory against possible points of objection by conveying exactly how the language of her term is functioning.

From start to finish, the effectiveness of Katie's paper is engrained in her careful treatment of key terms, both in that treatment of others' key terms and of her own. Her essay is a masterful display of how central the key terms can and ultimately should be in the overarching framework of an argument.

Works Cited

Nagel, Thomas. "Moral Luck." Article

Rosen, Gideon. "Moral Luck." Lecture, Introduction to Metaphysics and Epistemology from Princeton University, Princeton, NJ, December 14, 2016.

Bios

Katie Petersen '19 is from Charlotte, North Carolina. She is concentrating in philosophy, and she is particularly interested in the philosophy of religion and its intersections with moral philosophy and epistemology. She is also interested in theology and literature, and there is just not enough time to read all the books. She wrote this paper as a sophomore.

Danielle Hoffman '20 is so excited to be a new member of *Tortoise's* team! A sophomore from Connecticut, Danielle is pursuing a concentration in Philosophy and certificates in Linguistics and Music Theatre. She is a Writing Center Fellow, editor for the Street section of T*he Daily Princetonian*, and a dancer in diSiac and Blasé Dance Company.