

“How is it possible to be more or less culpable depending on whether a child gets into the path of one's car, or a bird into the path of one's bullet?”

August 31, 2024

By Caroline and Michael

Introduction (What Is It)

In ‘Moral Luck,’ philosopher Thomas Nagel examines the problematic implications of our common moral assessments. In everyday life, we tend to place responsibility on people based on their controllable actions, but the overwhelming influence of external factors on said actions means that this responsibility is heavily subject to chance — what Nagel terms “moral luck.” Given how moral luck represents the tenuous line between, say, mere driving under influence and vehicular homicide, we address and give further thought to the pragmatic implications of Nagel’s arguments. Much of the discussion was spent trying to examine whether intention could serve as a less contingent criterion of moral judgment or whether the subject’s intentions are also dependent on moral luck.

Nagel discusses the philosophical consequences of cognitive studies performed on patients whose *corpus callosum* have been surgically removed, isolating the two hemispheres of their brains. We try to make sense of the neurological evidence, which suggests that both isolated hemispheres independently and at times conflictingly exhibit consciousness. This raises a crucial question — does our mind consist of more than one conscious agent, and, if so, how can we quantify or attribute them? How then could we have a unified self? We examine how Thomas Nagel uses neuroscience to shed light on philosophy in ‘Brain Bisection and the Unity of Consciousness’ and the implications of his conclusions.

Transcript

Caroline: Let's start with sharing our take on the first essay of interest, 'Moral Luck'. Were there any places that were particularly confusing or interesting? We can also consider if, and how, Nagel's argument in 'Moral Luck' led us to change anything about how we deliver moral judgment.

Michael: What is Nagel's verdict in moral luck? How does he conclude it? It seems like he didn't resolve the issue. In particular, when reading, he did raise a lot of questions about, you know — he gave us such a predicament and it's not really clear how to proceed. Do we just accept this paradox, or this illusion of, you know, moral judgment and responsibility, and kind of go with it? Or, is there any way that this can really factor out in our lives? Those are the main questions that I get while reading Nagel.

Caroline: I see what you mean. Nagel doesn't give us a firm conclusion to such predicaments, in the way that would satisfy us. He tends to question, and makes us question, our intuitions on how we attribute moral responsibility in our everyday lives. So do you find his argument about the role of moral luck reasonable and sound, or does it have holes in it that may prevent you from accepting it?

Michael: I mean, his point that moral responsibility isn't very ... it's very contingent and not really up to the doer of the action, which is very valid. I mean, there's not much that can be. It's sound, but I don't think it could be used as an argument in a courtroom case. Or rather, I don't think it'll be heated very much in practical, pragmatic situations anytime soon. Nor anytime at all for that matter. I mean, that's my take, if anyone else wants to butt in on that.

Caroline: Yes, so if anyone has any objections or agreements, we could jump off of that. Similar to how we did the discussions at Stanford, if you all recall.

Kolb: I can do an agreement. So I agree mostly in part that this is, I mean, I think everything Nagel writes is very agreeable. Like usually he just takes that common sense view to put the proof behind it. And again, I don't think any philosophy is usually, or would be used in a court. But some parts of it just feels like he's basically attributing everything and anything to luck. I think it was Parfit who said this, but he hasn't really answered anything. He has just given a one-for-all answer that isn't gonna work in detail at any point.

Michael: Indeed.

Kolb: And you see that especially in the fourth kind of luck where basically anything that doesn't happen is because of luck.

Caroline: Nagel does say that the factor of luck changes the way we ought to morally judge people. So what are the implications of this? Do you think that luck has such a significant influence? Or is there a degree of luck necessary to cause an impact?

Kolb: I think for anything with luck, it's proven that you can't really quantify its influence — which means you can't quantify anything else either. So yeah, I think he does say you can't judge people's morality just because there's a factor that you can't truly take into consideration.

Caroline: Do we think there is, or could be, any sort of similarity between the way Nagel presents his ideas and skeptical arguments? Like is he deliberately suspending belief or calling for something similar to that?

Daniel: I don't think his argument was really strong enough to do that though. Because he talks for a long time about how it seems like what we do is inseparable from the standard understanding of the consequences of our actions. But then, it wasn't very clearly formulated because he didn't really define it exactly. So are we just saying that an action is defined as the action itself plus how much of its consequences? For instance, Nagel brings up that if you're a negligent driver and you accidentally hit someone on the road. And what if that guy turns out to be like, you know, really bad or something? Where do we separate the action from the consequence? Because if there's not a very clear formulation of that, I don't think we can make a moral judgment at all.

Caroline: It does seem hard, then, to make moral judgments while also considering moral luck. Does Nagel offer a way to potentially resolve this predicament, or does it become more difficult with his proposed cases of moral luck?

Junyi: His definition of moral luck is fuzzy. Especially the fourth case about moral luck and moral assessments — how the consequence can affect the moral judgment. Like arguably, that's kind of weird. For instance, if you try to kill someone, whether that be successfully or not, we would tend to think that both are done with bad intent. It doesn't

really matter if you successfully kill someone. So I think his definition of moral luck is based on the first three types (especially the first or second): some circumstance. Maybe, you are put in a bad situation and then that makes you morally bad. Like maybe the Nazi example will be a good example. In that case, you might have bad moral luck. But I think intent is more of it, not necessarily the consequence.

Michael: You bring up the first or the second case, and that while it's fuzzily defined, what if you murder someone with the intent of murdering someone. But Nazis committed war crimes, anti-Semitism with the intention of doing it. And Nagel's whole argument is that if they weren't Germans and if they hadn't been brainwashed by the Nazi party and been subject to the bystander effect, then they would not have committed those war crimes or they would not even have had the intentions of committing those war crimes. And the term he uses is constitutive luck. And you're right, it's a fuzzy definition, but like you point that out. But the fact that it is fuzzy is not what Nagel's trying to highlight.

Junyi: Yes, if for the Nazi example that would be bad moral luck, if someone became a Nazi. But I'm trying to say that the fourth case of moral luck is not necessarily a factor. Like it's not really about the consequences. Like in the courtroom situation, you would care more about the intention. Well, of course the result matters as well, but arguably you judge someone a lot based on their intention, not necessarily based on the results.

Caroline: But wouldn't the result also matter as much? Take the Nazi example again. Say that there are two individuals living in Nazi Germany who are anti-Semitic and hold the same beliefs. But the difference is that one is a high-ranking military official who is in a position to execute their beliefs and one is an ordinary person. Don't we judge based on what the Nazi official has done, the results, the real life manifestations of those anti-Semitic views and beliefs?

Junyi: Isn't that just circumstance? Circumstantial luck?

Michael: I think it very much is. With the Nazi example, we're trying to illustrate that intention can also be, as Junyi and Daniel said, contingent and circumstantial. I don't know where that brings us.

Kolb: I think even if it is circumstantial, we can't really figure anything out, we can't really measure anything. You can still use intuition to think about morality and relativity. So you could intuit that someone who does charity is going to be more moral than a Nazi.

But just because you can't measure that because of the luck factor doesn't mean that you can't compare them whatsoever. And so in this situation, just because you can't measure the difference between all the various factors — those causes or intent or whatnot — that goes into the difference between a Nazi officer and a Nazi soldier, you can still reasonably say that one is more moral than the other. You just can't say how much more.

Junyi: That's probably why we have ethicists. Like the football example we read on the first day, the people who judge the ethics of something.

Caroline: I agree that it's very necessary to have such people. But returning to the notion that moral responsibility is dependent on external factors beyond our control, I'd like to know whether you all find this a comforting or distressing thing.

Daniel: I don't see how it's distressing though. Like how could it have happened any other way than it happened? To me it doesn't seem very clear what Nagel's just priorities are at all. He has a section on free will at the end. He's like, 'it's really hard to positively formulate free will'. And then he just leaves it at that.

Kolb: So you're calling him a bad writer.

Daniel: Basically. He doesn't establish why it's important for me to hold this idea of free will, too. And going back to the intuition thing, he was talking about the relationship between theory and intuition, how we shouldn't just have a theory and then have it not match with our intuition at all. But then again, I don't think he properly establishes the importance of that. So maybe it's just that he didn't really convince me with the weight of moral luck. But again, it seems that, under certain circumstances, even if it was unfortunate that they would have made that judgment under other circumstances, they would have made some other judgment and then some other outcome would have happened. I don't really see how that meaningfully affects the judgment of the present situation of what actually happened.

Caroline: Because, what Nagel is saying is that the circumstances, or some other external factors, are out of your control. What actually happened could simply be attributed to luck, something that you yourself couldn't have controlled or had a say in. And Nagel says that it makes sense for us to make judgments on people based on what they can control, not on what they can't. So then moral luck would be significant and meaningful to consider.

Daniel: I'm saying that if you're a Nazi, you're a Nazi. And even if under such circumstances like, 90% or whatever of normal people would also become a Nazi, they're still a Nazi. So it does have to do with intent too, because even if that intent is shaped with moral luck, we could still judge based on intent. Which connects to what Kolb was saying.

Caroline: So you think that we should judge based on intent?

Daniel: Ah yes, I think I tend to lean towards the intent side, like Kant. I think that it holds a lot of intuitive sway.

Michael: Would you say intent depends on moral luck?

Daniel: Sure, the beliefs you hold or the intent, the will that you exercise, is dependent on your personality or the circumstances or any of these things. Clearly a large part of that is outside of your control. Genetics, and all that stuff. But I don't see how that is significant in making a moral judgment.

Kolb: With the Nazi thing, I feel like a lot of what Nagel said is contingent on the fact that, because of certain moral luck factors, you don't have a choice on your moral actions. But I mean, to a certain degree, you could say that you always have a choice. If you were born into Nazi Germany and you were asked to be an officer, you could refuse. Like they might kill you, but you could refuse.

Michael: You would be ostracized. The point that I took from Nagel's exposition of Nazi Germany is that intent itself is also something that is dependent on moral luck, not just the consequence. How so? The Nazi party originally did not even enjoy majority support. They won a third of the Reichstag when they came to power in 1933, but because of Goebbels, who instigated a bunch of propaganda campaigns, essentially everyone grew to love Hitler and persecute Jews and all of that. And if you were a male, you had to be introduced into the army and if your half family was of high standing, a contingent factor, then you are to become an officer and you would therefore have significantly more guilt or more blame placed on your shoulders — simply because of your family standing. And if you take into account the aggregate war crimes committed by everyone in Austria at the time, then I would say that very much what affects your intents is not just your genetics and your condition, or your personality as you say, Daniel, but also the Zeitgeist,

the conditions of society in the world at your time. And why this is so poignant for Nagel is that at the same time in Britain, people weren't doing what the Nazis did, but they could just as well have if they had been brought under Nazi Germany, like France was. And a lot of people in France were also being very anti-Semitic. Because they were occupied by Germany. But this wasn't necessarily remembered as much as the Nazis in Germany. But that's a digression.

Kolb: But in theory, you could — it doesn't matter how much you can influence, you can influence everything. I mean, that doesn't really make sense. But in theory, if your social intelligence was good enough, you could be the person to overthrow the Nazi regime from the inside, right? So you always have a choice.

Caroline: Doesn't social intelligence also depend on your circumstance, how you were raised, or your personality and abilities, and all those other factors that are out of your direct control? Like, not really something that you are born with, but a part of your background and growth as a person.

Kolb: Like being a constitutive luck? Yeah, I guess.

Caroline: How do you suppose we make moral judgments then?

Daniel: Using Michael's example of the high class Germans becoming higher ranking officers in the Nazi party, the account of making moral judgments based on intention should still be OK. Because it's not like they were just sitting there as high-ranking officers. They were going to, you know, probably carry out things of higher magnitude of higher consequence — or, heavier responsibility and obviously with the intent of being responsible for those. So I think we can still make a judgment, even if that was largely based on luck that they were in that position to make that choice that they did, but they told me that choice, you know.

Michael: Yeah, I agree with you Daniel. And I also think intention, as far as we as outsiders can gauge the intention of other people, is a fair metric or criterion for moral judgments. But then we'd be making a concession to Nagel, as we are recognizing that intention is very much based on moral luck. And that's exactly what Nagel's criticizing. That, at the end, it's all very shaky, no matter what standard we go by or what we judge by.

Caroline: On that final note, let's move on to the next essay on our plate, 'Brain Bisection and the Unity of Consciousness'. What did we all think of this one? Were there any confusions, interesting finds, possible objections or agreements?

Junyi: I think this one's harder to refute. All of his points seem sound, not gonna lie.

Caroline: Yes, Nagel does base this one on scientific findings from the split-brain operation. Does it help him resolve the issue of consciousness? Was there anything interesting about how he used science in conjunction with philosophy?

Andrew: I think it's interesting how the two different sides of the brain had specific functions. I wasn't really paying attention in biology but I just found it interesting that certain parts of the brain had different relations to what a person does — that one side of the brain does this type of thing and the other side does another type of thing. And it was more interesting to also see how Nagel applied scientific studies of the brain for his argument.

Michael: We should also read over all of the findings that Nagel presents and summarize them, because I feel like Nagel's findings can be refuted if you consider the fact that there have been new findings after he wrote it in 1979 and that such findings revise or qualify the results known to Nagel. I can start with summarizing the results. So essentially Nagel says that brain bisection is the operation where you remove the *corpus callosum*, which is this network of fibers connecting the two hemispheres of the brain. The operation is done only to patients with severe epilepsy and extreme cases. What happens after is that the two hemispheres of the brain are essentially isolated for a while. Now onto all of the tests they're running. So if they flash a light or a stimulus to the left visual hemisphere of the patient, the patient cannot report it verbally but they can identify the stimulus with their left hand. But what's the connection to the disunity of the minds here?

Caroline: The different parts of the brain, then, seem to be more distinct and independent after this operation. They also conflict with each other sometimes. This all paints a picture of not one, single, cohesive mind. Beyond that, how exactly do you guys think that this experimental evidence of bisecting the brain presented by Nagel implies the plurality of minds?

Kolb: Yeah, I could answer that. I think you could theoretically bisect everything infinitely, right? So you could say that the different, separate parts of the brain perform

separate functions. But you also can't just leave it at that and say the brain is just the make-up of all its separate parts, because the brain can only perform some specific functions well only if its parts are together and work together. For example, playing music. You can't do that with half a brain. So there are still aspects that make sure the brain needs to be a cohesive whole. You can't separate it that easily.

Michael: Well, yes, but they did bisect the brain, and it still remained a relatively cohesive whole, but with like slight defects, right?

Caroline: Right, so after the operation, the patient could perform human functions normally and could go about their daily lives without much, significant impact. Though it was proven that the two hemispheres conflicted with each other at times, this has a minuscule impact. The patient was still able to regain full function, even with separated, independent hemispheres. So Kolb, do you still think that the brain isn't a cohesive whole?

Kolb: I still do. This was obviously a scientific report. It doesn't seem like a very traditional, philosophical rebuttal, but this is a scientific report and you can't expect them to report fully accurate findings. So when they say that they retained full function, it's not like they tested for every single possible thing that a human can do. So they're obviously incentivized to write it in a way that's more beneficial to their outlook, to their reputation. I think it's unlikely that when you bisect your brain, nothing changes. Even if it's a minor change, something must have changed, and if it's minor enough, they'd be likely to leave it out of a research report.

Junyi: Basically there's limitations to it?

Michael: Well it's reporting exactly on what has changed, so I don't think there'll be any bias in that respect. So the surgery is a cure to stop very severe epilepsy patients from having fits and seizures like 50 times every day. And they do it. They bisect the brain of the patient and they release them back into society, and they're normal for the most part. That we all agree on, but we never said that they were completely normal. There are side effects, such as being able to recognize a stimulus with your hands but not with your mouth. And those are very obvious side effects that research reports have been—

Kolb: I get that. That's not what I meant though. What you're saying is that certain functions are impaired and restricted to certain sides of the brain, right? So like, one side

can only recognize the light, one side can only react to it. It's more a matter of abilities that they've just lost.

Junyi: There's actually a new study done that relates to what you said.

Kolb: What is it?

Junyi: So basically, you can completely remove one hemisphere of someone's brain, and it's been found that the brain can recover some, or most of the functions of the lost hemisphere.

Kolb: Again, most of the functions. They're not going to tell you the functions that were lost.

Daniel: But why would it matter if it's really, really small?

Kolb: Because that means the consciousness is changed.

Daniel: Sure, but is that really your criteria for having a consciousness? Like is it just a set list of minimum, say, requirements that your brain can do? What if somebody loses one?

Kolb: I feel like going back to the identity question then. You're not your identity if you're the most of what you are. You have to be all of what you are.

Daniel: That seems like a pretty strict requirement though.

Kolb: I mean, does anyone else think it's unreasonable? It's strict, yeah.

Junyi: Do you think consciousness is one to one with the number of minds?

Kolb: No, I think your consciousness is reduced.

Junyi: That was kind of difficult, right? Because it's proven that the halves are independent, so you can say that they have two consciousnesses. But it's a question of whether they have more than one mind. That, I'd say, is the harder question.

Caroline: Yes, how can we quantify consciousness? That's a pretty important issue raised. Would consciousness be proportional to the number of minds? Or can we even count the number of consciousnesses — is consciousness quantifiable?

Junyi: The two halves are independent, so that suggests that there's two consciousnesses. Or maybe normally, we have two consciousnesses in the first place. There's just a bridge between the two hemispheres.

Kolb: You could attribute consciousness to one side of the hemisphere. That's also an option: to say the other side doesn't have a consciousness. ... [cut-off] ...one side of his brain when half of his body would act in complete accordance with the other side of his body. The thoughts going through the two sides of his brain will no longer cohesive. And so one side may like something but the other side may dislike something. So what he wants is not what one of his sides is doing. So you could say consciousness is just in one side of the brain. Then you would have to do enough bisections to figure out what side of the brain that is, or where exactly that is. They should do that, actually.

Junyi: But I think Nagel already disproves that. One of the first claims he says is that it's not true. Because even if one side of the brain can't give testimony, that doesn't mean it can't have a consciousness. The right hemisphere can't speak, but it doesn't mean that it can't think for itself. Arguably it has been shown that it can, in fact, think.

Kolb: Define what's subconscious then.

Michael: I think that, no matter how you define consciousness, or even whether or not you define consciousness, it's pretty easy. I'm pretty sure there have been results verifying that neither side is lacking consciousness — it's just a simple matter of isolating one side and giving one side isolated visual stimuli, such as a dog or a pencil, and asking them to respond to that in a way that would manifest consciousness, such as identifying a pencil or the dog. That's been talked about in Nagel's essay.

Kolb: Fair enough.

Michael: But on the matter of multiple consciousnesses and what Kolb was saying. I've read some important findings or conclusions about split brain research that people came up with after Nagel. And I think they might be pertinent to what we're trying to get at now. With 6 minutes left, I can present them really briefly. So essentially, after split brain

surgery occurs, the two hemispheres of the brain do not necessarily remain isolated. And that's a key finding that only recently has been gaining credibility. They might be, for the most part, isolated because there is no *cortical callosum* anymore, but there are processes on subcortical connections or cross-hearing — I'm not a neuroscientist, but they can communicate without having that bundle of neurons with them. So they're still in contact and not completely isolated. And that also explains the results of 'how do you respond?' and how different patients respond to different stimuli. Results vary a lot according to individuals. The fact that after surgery occurs, the brain can find ways to try to become a whole again. That can explain the individual variations. Kolb raised the question, what is consciousness? We haven't really defined it here, which is a good thing because Nagel is anti-reductionist, and he thinks consciousness can't be defined. But neuroscience right now is trying to define consciousness. And I've read three, main scientific attempts to explain consciousness. I can go over them if anyone wants to, but we're running out of time.

Caroline: Go ahead, that would be a good place to leave us off of.

Michael: I want to raise an interesting question and it's whether consciousness could be statistically defined. So how many consciousnesses do we have in our brain? Is it one, or half, or a non whole number, as is the case for, maybe lobotomy patients? Or is it several? There's a thing called integrated information theory, which equates consciousness to a mathematical value, a phi value. It's the measure of how integrated the level of information or the ability of the brain to process information is. So if you get a brain bisection and that limits your brain's ability to process information, that would have your consciousness quotients.

Caroline: These are great considerations for us, to either use as a starting block for our next meeting or simply, to think about. For our next session, which will also be around this time at the end of the month, we'll go over the essays 'What is it Like to be a Bat?' and 'Subjective and Objective'.

Kolb: I feel like we could revisit the bat thing next time, but see what it looks like without the science, because I think there's a reason why people don't really use scientific evidence in philosophy. And so I'm just curious what we think because we've been really caught up on the science and science is always changing and there's a lot of conflicting studies.

Caroline: It is quite interesting how science can work with philosophy, especially considering philosophy of the mind. So yes, we could do that next session, especially as 'What is it Like to be a Bat?' happens to be my favorite one. Thank you all for this lovely discussion, and I hope to see you next time!

Conclusion

This first philosophy meeting was a rather successful one, with thoughtful and productive discussions centered around Thomas Nagel's 'Moral Luck' and 'Brain Bisection and the Unity of Consciousness.' With respect to 'Moral Luck', we found that the existence of moral luck is problematic for our capacity to exercise moral judgment, as it implies that many actions are out of the immediate control of the doer and thus cannot serve as a basis for assigning moral responsibility. We particularly focused on extracting from the example of Nazi Germany, employing a *Reducto ad Hitlerum* and analyzing whether the German population can be held responsible. The general consensus was that, intuitively, moral judgments ought to depend on intention. However, what remains unresolved is the issue that intentions are dependent on moral luck.

Regarding 'Brain Bisection and the Unity of Consciousness', we determined that the brain retains full, normal function after a split-brain operation, although its hemispheres have been separated. We utilized our knowledge of new scientific discoveries and raised the important question of how a scientific study can conflict or work in parallel with philosophy. Such a question connects to academic incestuous auto-glorification and the 'publish or perish' culture, leading to some (non-Pyrrhonian) skepticism on the scientific study vital to Nagel's argument. These considerations deepened our conversation, and complicated the question of whether consciousness could be quantified.

Attendees

Andrew

Ben

Bridget

Caroline

Daniel

Junyi

Kolb

Lihi

Michael

Tommy