Imagining, Inside and Out

Peter Kung, Pomona College
April 27, 2009
Under review at *Erkenntnis*

For a while after he first burst onto the scene, I would imagine being Tiger Woods. Sometimes I would see in my mind's eye something like a telecast of me winning the Masters. Other times I would imagine interacting with people as Tiger would, from "inside Tiger's head." Some writers find the claim that I can imagine being Tiger Woods problematic, as problematic as the claim that I could *be* Tiger Woods. But this certainly seems to be something that we can imagine — children *pretend* to be Tiger Woods — and it would be good to accommodate that intuition in one's theory of the imagination. In this paper I argue that we can.

Imagining from the inside is a central preoccupation in philosophy of art, but it crops up across philosophy, most prominently in thought experiments.[1] I use my theory of imagination to diagnose two famous thought experiments in philosophy of mind that seem to involve imagining from the inside, though whether and how they do so will be a subject for discussion. In the spirit of Descartes, we might think we can imagine ourselves existing disembodied; I imagine myself existing without any matter existing. More recently, dualists have argued for the imaginability of zombies, beings who are microphysically identical to conscious beings but who are not themselves conscious. I analyze both cases and conclude, with some important qualifications, that we can imagine zombies and being disembodied.

Some philosophers find the imaginability of zombies and disembodiment just as problematic as the claim that I can imagine myself to be Tiger Woods. Since it is impossible for Peter Kung to be Tiger Woods, the worry goes, I cannot be imagining the identity of Peter Kung and Tiger Woods.[2] Likewise, because many philosophers believe that zombies and disembodiment are impossible, they are unwilling to admit we can imagine them.[3] Such philosophers typically take imagination to be a guide to metaphysical possibility and, given their commitment to the impossibility of zombies and disembodiment, they infer that we must not be able to imagine these impossible things. I argue that this reasoning should be rejected, for it gets the order of explanation reversed: if we want to use imagination to settle modal issues, we should not allow our modal conclusions to govern our theory of imagination. A plausible modal epistemology flowing from my theory of imagination shows us how we can avoid the mistaken reasoning, and allows us to see on independent grounds that some imaginings, like my imagining being Tiger Woods, provide no evidence for metaphysical

---

[1] In philosophy of art, see for example Currie (1995), Smith (1997), Walton (1990), Wolheim (1987). Imagination also figures prominently in modal epistemology, simulation theory in the mindreading debate, empathy, pretense, action theory, some accounts of irony and propositional content, hypothetical reasoning, dreams, and creativity.

[2] See Reynolds (1989), Mackie (1980), Walton (1990), Williams (1973).

[3] *Materialists* contend that the zombies and disembodiment are impossible. See Hill (1997), Marcus (2004), Nagel (2002), Shoemaker (1993), Tye (1983), and Zimmerman (1991). Others contend that these though experiments support arguments for *dualism*; see Hart (1988), Kripke (1980, at least tentatively), Searle (1992, modulo his unorthodox views about "brains causing minds"), Swinburne (1997), Taliaferro (1986, 1994), and Yablo (1990, 1993).

possibility.[4]

1 IMAGININGS ABOUT OURSELVES

Let me start with a series of cases inspired by Bernard Williams' "Imagination and the self" (1973). These cases will bring out two key distinctions: imagining from the inside vs. imagining from the outside (§1.1), and imagining replacing X vs. imagining being X (§1.2). We'll better understand how to analyze these distinctions with the help of some theoretical resources developed in section 2. That will put us in a position to assess how we imagine zombie and disembodiment cases (sections 2.1 and 2.2).

1.1 Inside vs. Outside

Up-and-coming golfer Anthony Kim recalls watching Tiger Woods' historic 1997 Masters victory as an eleven-year old: "I remember in my mind, putting my face on his body" (Culpepper 2008, July 16). Here's how my version of Kim's imagining would go.

> **Replacing Tiger**
>
> I imagine being good — really good — at golf. I visualize my winning the Masters at age 21; I walk up the fairway on eighteen and hug my father (my real father, Edward, not Earl Woods). My picture, a likeness of myself, not of Tiger, appears on the cover of *Sports Illustrated*. People call me "Tiger" and my dad "Earl," but I still look like I actually do, my dad looks like he does, and so on. My imagined life is similar to my actual life, modulo a lot more time on the golf course inserted into my past.

This description is incomplete in at least one crucial way; it is compatible with two different points of view I might take in my imaginative project. First, I might contemplate the epistemic position that I would be in were I to do such things, and conjure up the experiences I would have from such a point of view. It is natural to call this imagining *from the inside* (or just *inside imagining*, for short) and *first-personal*. Such an imagining takes on the perspective from which the situation is surveyed:

> **Replacing Tiger – Inside**
>
> I imagine from the inside walking up the eighteenth fairway, looking out over the adoring crowds. My hand appears in my field of view and moves back and forth as I wave to the crowd. (They love me.) The back of my hand is the tannish yellow color that it actually is. When I imagine hugging my dad near the fringe of the eighteenth green, the experience I imagine is very like the experience of actually hugging my real father. When I imagine myself taking a shower after I've completed my historic round, I imagine the feeling of warm water on my back; the image I see in the steamy locker room mirror is my familiar visage.

In contrast, I might imagine this same case *from the outside*, rather than from the inside, making it a *third-person* imagining.[5] An outside imagining of this case would be something like

[4] Henceforth I'll suppress the 'metaphysical'; by 'possibility' I will always mean metaphysical possibility.

[5] Williams occasionally calls outside imaginings *visualizations*, a usage that has been picked up by some, but not all, writers. Nagel (1974) and Hill (1997) refer to outside imaginings as *perceptual imaginings*. Both labels have connotations that I find misleading: see the discussion of verificationism in section 4.3.

watching a mind's eye home video of myself winning the 1997 Masters.

> **Replacing Tiger – Outside**
>
> I imagine a figure — a whole body, from head to toe, and not just the arms and whatever else I usually see of myself from the inside — against a backdrop of lush green, looking much like I do in photographs. I watch this figure smile, hear him shout as he hugs a man who looks like my father. The figure shakes the hand of Constantino Rocca, with whom he was paired for the final round, and I notice that he is taller and better looking than Rocca. In my imagined situation, this person is me, Peter Kung, even though I'm not, as we might say, looking out through that person's eyes and hearing through his ears.

To understand the difference between inside and outside imagining, let's consider first-person vs. third-person narrative fiction. In fictional writing we distinguish the author of a story, the story itself, and how the story is told. The same author can tell the same story in a first-person or a third-person narrative; the story will be the same in that it has the same characters undertaking the same actions in pursuit of the same goals. For example, compare Twain writing, "I took the canoe out from the shore a little piece," where Huck narrates, vs. "Huck took the canoe out from the shore a little piece."

We can make similar distinctions for imagining. The author is *who* is doing the imagining, the imaginer. The story is *what* is imagined; I'll also sometimes refer to the *content* of the imagining, what's *true* in the imagined situation. For our purposes, the number and identity of *characters* in the story figures centrally. The mode is *how* it is imagined. In this paper, the mode is a matter of whether the story is imagined from the inside or the outside.

A difference between inside and outside imagining is that, with outside imagining, there is no character *in* the imagined situation *from* whose perspective one is imagining experiencing the situation. For example, the story of Replacing Tiger – Outside contains two characters, not three. When I imagine shaking Constantino Rocca's hand, I do not imagine a story with three characters: Constantino, me (the figure shaking Constantino's hand), and an observer. That's one character too many — there is no observer in the story, hence it is not part of *what* I imagine. And just as there is no observer character imagined, no observer character's experience is imagined either. I do not imagining a story with three things, two people and an experience of those people; I am just imagining the two people. My outside imagining experience is something that I, the author, Peter Kung, *have* in order to imagine the two people, not something I take the story to *contain*.

Berkeley's famous case makes this point nicely.[6] Imagine a tree stands alone, unobserved. As with every imagining, there is an author, someone *doing* the imagining. In my case that's Peter Kung. The perspective from which the tree is surveyed is not the perspective of any character *in* the story. If it were then we'd fall into Berkeley's trap: the tree *would be* observed by a character after all, and it would be impossible to imagine a thing that exists unperceived. Most find Berkeley's suggestion implausible and I concur. We can imagine from the outside an unobserved tree, and distinguishing between author and characters in the story explains how.[7]

---

[6] Berkeley (1979), pp. 35–36.

[7] For dissenting views see Peacocke (1985) and Martin (2002), who argue that *all* imaginings are

By contrast, when you imagine from the inside, the subject through whose eyes and ears you experience the situation *is* a character in the story. The mind's eye sees what the subject whose head we are "inside" is imagined to be experiencing. Imagine from the inside picking up a golf ball. Were you to give a running commentary on your imagining, you might say, "I see my hand stretch out in front of me and grasp the ball between my fingers. The ball is a white sphere that grows larger as I bring it closer to my face…" The first personal pronouns in the quoted sentence refer to the subject whose head you are inside. "In front of me" means in front of the *character* in the imagined story, the one who is watching her hand grasp a golf ball.

In sum, it is clear that in imagining from the inside but not the outside, the imagined perspective belongs to a character *in* the story. But *who* is this character? One answer seems obvious: this character is Peter Kung. In the normal run of things I imagine Peter Kung from the inside, imagine Peter Kung having various experiences. What could be more natural? I will argue next, however, that the character need not be Peter Kung. Let's turn to that argument.

1.2 I am Tiger Woods

Henceforth we'll focus primarily on imagining from the inside. Consider Replacing Tiger – Inside. Although my little fantasy is inspired by Tiger's exploits, it does not seem that I have imagined *being* Tiger Woods. What I have imagined is "walking in Tiger's shoes," doing the things that Tiger has actually done. I have imagined myself *Replacing Tiger* (hence the label for the set of cases), rather than *being* him in some stronger sense. Here is how I would describe imagining *being* Tiger Woods.

> **Being Tiger – Inside**
>
> I imagine having experiences as of the back of a brown-skinned hand. I see (in my mind's eye) the hand reach for a mirror, and as the mirror comes up in the right hand, the image reflected in the glass resembles the guy we've all seen on television and Wheaties boxes. I imagine hugging the Earl Woods we've all seen on television, who I hear calling me "Tiger" as I hear my own voice calling him "Dad."

In this case, I imagine having certain *Tigerish experiences*, having the sorts of experiences that Tiger has had. But I claim that there is more to imagining being Tiger than imagining having Tigerish experiences. I have to imagine that *I am Tiger*. But how can I do this? Isn't it obviously impossible? If Tiger and I are distinct people, we could not be the same person. So how can I imagine being Tiger?

The solution to this puzzle is not to confuse characters in a first-person story with the author of the story. In every case that I imagine, I, Peter Kung, am the author. But Peter Kung does not necessarily figure as a character in every imagining, even in every imagining from the inside. I will argue that we need to postulate what Velleman (1996), following Williams (1973), calls a "bare Cartesian ego," the character in a first-person story from whose perspective the narrative unfolds. For convenience, let us name this character "Ego." The bare Cartesian ego has *no further identity* beyond being the character having the imagined first-

---

imaginings from the inside. They both endorse a version of the following principle: to (sensorily) imagine φ is to (sensorily) imagine *experiencing* φ. For discussion see Noordhof (2002), White (1987) and Williams (1973). A variant on this view appears in the film theory literature as the Imagined Observer Hypothesis. See Currie (1995) for objections.

person experiences. With imaginings from the inside, indexicals like "I" and "me" are ambiguous. Sometimes they refer to the author, Peter Kung. Other times they refer to Ego, the character *in* the story from whose perspective the imagined experiences unfold. It is not *necessarily* part of my imagining that Ego is Peter Kung.

The argument for Ego is that without Ego we cannot correctly distinguish between different first-person stories, stories that we plainly can imagine. We'll consider five cases (the first and last have already been described above). In each case the author is me, Peter Kung. And in each case the story is told from the inside; it's a first-person story. The stories differ in the characters involved: in some stories there is only one character; in other cases there are two. In some cases the author is a character in the story, in other cases the author is not. Let's look at the cases.

1. Replacing Tiger – Inside. See above.

In the Replacing Tiger story, there is one character, me, Peter Kung. Tiger Woods does not exist.

2. Resembling Tiger – Inside. When I look in the mirror, I see the light brown African-American face that resembles the one we've all seen on Wheaties boxes. The face I see in the mirror is *my* face. However this is science fiction spy thriller; I have had surgery to make my face resemble Tiger's. I am Peter Kung, *not* Tiger Woods. I am an imposter fooling everyone into thinking I am Tiger Woods. I have even disposed of the real Tiger Woods so that nobody discovers my deception.

In the Resembling Tiger story, there is again one character, me, Peter Kung. Though I have Tigerish experiences, I am not Tiger Woods. Tiger Woods does not exist (anymore).

3. Behind Tiger – Inside. When I look in the mirror, I see the familiar light brown African-American face that resembles the one we've all seen on Wheaties boxes. The face I see in the mirror is *not* my face however. What has happened to me is like what happens in the film *Being John Malkovich*: I crawled through a portal and found myself behind Tiger's eyes. Tiger still exists; he is really good at golf and has just won the 1997 Masters. Tiger is looking in the mirror, but even though I see the world from Tiger's perspective, I am *not* Tiger (I am terrible at golf, and I've certainly never won the Masters). Every experience that Tiger has is one that I have as well. I am Peter Kung, seeing, hearing, (feeling, …) the world from Tiger's perspective.

In the Behind Tiger story, there are two distinct characters, me, Peter Kung, and Tiger Woods. We are both having Tigerish experiences.

4. Craig Behind Tiger – Inside. Just like previous case except that I am Craig Schwartz, the "protagonist" in *Being John Malkovich*, rather than Peter Kung.[8] I am Craig Schwartz, seeing, hearing, (feeling, …) the world from Tiger's perspective.

In the Craig Behind Tiger story, there are again two characters, but this time they are Craig Schwartz and Tiger Woods. Both are having Tigerish experiences. In this story, Peter Kung does not exist.

5. Being Tiger – Inside. See above.

[8] In the movie, Craig Schwartz is portrayed by John Cusack.

In the Being Tiger story, there is only one character, Tiger Woods, who is having Tigerish experiences. Peter Kung is not a character in this story.

Each story above is told in the first-person. Ego, the character from whose perspective the narrative unfolds, is *not* the same character in each story. In some stories Ego is me, the author, Peter Kung (Replacing Tiger, Resembling Tiger, and Behind Tiger). In other stories Ego is not the author; Ego is someone else. In Craig Behind Tiger, Ego is Craig Schwartz. In Being Tiger, Ego is Tiger Woods. We can make it even plainer that Ego is not always Peter Kung: I imagine being Tiger Woods and meeting Peter Kung. I shake Peter Kung's hand and sign an autograph for him. In this case, it is clear that Ego is Tiger Woods, not Peter Kung. We need to accommodate the fact that these stories differ in who the first-person narrator character is.

Notice that we cannot identify Ego being Tiger Woods with Ego having Tigerish experiences. In every case except Replacing Tiger, Ego has Tigerish experiences. But in only one of them is Ego supposed to be Tiger Woods.

I conclude that in first-person stories there is a character, Ego. Ego is not necessarily identical to the author, Peter Kung. Sometimes — perhaps usually — Ego is Peter Kung. But other times Ego is someone else, like Tiger Woods or Craig Schwartz. Drawing this distinction between Ego and Peter Kung allows us to decipher what would otherwise be a puzzling expression, "I'm imagining that I'm not me, I'm someone else."

The only grounds I can see for objecting at this point, and claiming that we cannot imagine 1) to 5) as distinct cases, are worries about imagining the impossible. If one were to peek ahead, one might worry about the modal epistemological consequences of allowing, for example, that I can imagine being Tiger Woods. However the methodology I'm recommending suggests we not look ahead. We should settle on an independently plausible theory of imagination first — one that agrees with commonsense that we can imagine 1) to 5) — and then explore its modal epistemological implications.[9]

Here is what I have taken myself to have established thus far. There is a difference between imagining from the outside and imagining from the inside. In imagining from the inside, there is a character, Ego, who is not necessarily identical to the imaginer (in my case, Peter Kung). What I'll do now is introduce more theoretical machinery to explain how the story content of imagining from the inside and the identity of Ego works. We'll use that machinery to analyze how we imagine the zombie and disembodiment cases that appear in the philosophy of mind literature.

2 THEORY OF INSIDE IMAGINING

Some imagining is sensory imagining; it involves *mental imagery*. When you imagine seven-hundred-and-fifty trillion dollars burning or Justin Timberlake singing the Brady Bunch theme, you entertain some mental imagery — a sight or "picture" in your mind's eye, a sound in your mind's ear.[10] How should we theorize the story, *what* you imagine, the *content* of your

[9] For more on imagining impossibilities, see Kung (2009).

[10] Kind (2001) argues persuasively against those who claim that imagining is *purely* non-imagistic. I'll follow the usual practice of relying on visual examples, primarily because they are most familiar and hence easiest to describe, and also because the inside/outside distinction is clearest for the visual cases. Officially "mental imagery" should be understood to encompass all sensory modalities.

imagining? Sensory imagination highlights a key distinction between the *qualitative* and the *assigned* content of imagining.[11] The qualitative content comes from the "picture" itself: for example, the orange of the fire, the shape of Justin Timberlake's head. But there's lots of content to your imagining that isn't "pictured." In addition to the surfaces, colors, shapes, and so on that you "picture," you imagined that the bills summed to *seven-hundred-and-fifty trillion dollars* rather than some other amount. You imagined *Justin Timberlake* rather than someone wearing a very convincing Justin Timberlake mask. Those facts are not visualized, they are assigned.

To get a better sense of assigned content, imagine Justin singing a duet of the "Brady Bunch" with his doppelgänger Dustin; Justin is standing and Dustin is seated. What makes it *Justin* who is standing? Assignment. You can imagine the reverse — Dustin standing, Justin sitting — merely by changing the assignments. Though everything "looks" the same in your mind's eye, you nonetheless imagine something different by virtue of different assignments.

My view of imagination's *basic qualitative content* borrows heavily from the philosophy of perception. Perceptual experiences have representational content that present in a direct and immediate way aspects of the world around us, aspects that we are conscious of: they specify the distribution of objects and "basic observational" properties in three-dimensional (egocentric) space. Basic observational properties include at least the traditional primary and secondary properties. In vision, for example, we are consciously presented with three-dimensional space filled with objects of varying colors and shapes. Sensory imagination also has basic qualitative content. When you visually imagine huge stacks of thousand-dollar bills burning, your imaginative experience presents greenish whitish flat objects, laid out in space, some above others, some to the left, others to the right. Imaginative experience isn't presenting aspects of the actual world around us, but it is presenting "basic observational" properties in imagined space.

For our purposes what's crucial is that some assigned contents specify the *identities* of the things imagined. Think of this as the imaginative analog of *seeing as*. You imagine the yellow-orange "stuff" *as fire*, the flat greenish-whitish object *as a $100 bill*, and so on. The contents <fire> and <$100> bill are assigned. Assigned content explains *de re* imagining. In the Justin Timberlake cases, assignment makes my imagining a story about *Justin Timberlake* — that particular guy — rather than about some guy who merely resembles Justin Timberlake.

Assignments can also specify background information about the imagined situation. For instance, if you imagine that Justin and Dustin perform their duet *at the Super Bowl*, you needn't be picturing anything indicating where they are. There is simply an assignment to the effect that <It is the Super Bowl>.

The term 'assigned' has one unfortunate connotation that I want to quash. It suggests that imagining is a two-stage affair, where we first conjure up some purely qualitative mental "picture" — this picturing being the *real* imagining — and then assign various labels to the "picture" — the latter step being to describe what you have imagined. I hope the "imagining as" locution helps combat this distortion. Assigning is not a separate step from imagining;

[11] I develop the distinction at length elsewhere, but many details won't matter for our purposes. See Kung (forthcoming). Among other things, I distinguish between two kinds of assigned content, *stipulative* and *label* content. Label content captures the sense in which the mental "pictures" come precategorized, though it takes some work to articulate the correct sense of "precategorized."

'assigning' is a technical term for *part of* the imaginative process. That imaginative process generates imagery that comes with everything already assigned.

How does this machinery to explain what's happening in imagining from the inside? First, even in imagining Tigerish experiences, assigned content figures heavily. To imagine that I am looking in a *mirror*, rather than something that merely qualitatively resembles a mirror, I need to assign that it is a mirror. Otherwise it might be a high quality LCD screen, or a clever tromp l'oeil painting. Some of the Tigerish experiences are of *Earl Woods*, and not just a figure who resembles Earl Woods. It is assigned that the figure is Earl Woods.

Second, I explain the different identities of Ego with assignments. Recall that in cases 2) to 5) I imagined having the same Tigerish experiences. The qualitative content is all the same. The assigned content about what the Tigerish experiences were experiences *of* is all the same.[12] But the stories differed because they featured different characters. That difference is explained by the different identities assigned to Ego. In Replacing Resembling Tiger and Behind Tiger, I assign Ego = Peter Kung. In Craig Behind Tiger, I assign Ego = Craig Schwartz. In Being Tiger, I assign Ego = Tiger Woods. These assignments contribute crucial elements to the story.

Now that we understand imagining from the inside, and in particular the role that Ego and assignment play, we are ready to examine the zombie and disembodiment cases.

2.1 Zombies

Your zombie twin is an exact physical duplicate of you that is completely non-conscious. What is required to imagine it? Both Nagel (1974, 1998, 2002) and Hill (1997) contend that this imagining requires splicing an imagining from the inside with an imagining from the outside: "All I have to do is imagine the physical system from the outside and then imagine it from the inside — as not having any inside in the experiential sense" (Nagel 2002, p. 216). Our preceding discussion suggests that is a mistake.[13] First, imagining your zombie twin does not require combining imagining from the inside and imagining from the outside. Second, it is easy to imagine your zombie twin from the outside. Third, there is no sense in which we can imagine a zombie from the inside.

As we've already seen, when we imagine other people, or other conscious creatures, we do not have to imagine *being* those people. Recall Replacing Tiger – Outside. There we imagine from the outside Tiger Woods and Constantino Rocca shaking hands. In the imagined situation, Tiger and Constantino are conscious, though you did not imagine *being* either of them. Consider also one's typical experience in reading a novel. I imagine evil Count Olaf ecstatic and anticipating the Baudelaire fortune, Klaus frightened, Violet hatching a plan to thwart Olaf; I may picture all three characters in my mind's eye, from the outside. I do not, in my imagining, imagine being each one of them in turn — I do not pop "inside" each person's head to establish that they are conscious, or that they have the emotions or thoughts I imagine them to have. Even when I imagine from the inside, though I imagine being one

[12] In other words, the Tigerish experiences are of *hands* holding a *mirror*, and of hugging *Earl Woods*. All those facts are assigned.

[13] Here I describe what it takes to *imagine* a zombie; I disagree with Hill and Nagel that it requires imagining from the inside. I set aside the issue of whether zombies are *possible* until section 3.3. In section 4.2 I criticize Hill's argument that zombies are impossible.

subject, Ego, when I encounter other people I do not also imagine being those other people as well. In Being Tiger – Inside, Earl Woods figured as a character in the imagining, but I did not imagine being Earl Woods from the inside. Hence to imagine my zombie twin, I do not, contra Nagel and Hill, have to imagine *being* the zombie.[14]

The familiar way to imagine your zombie twin is to do so from the outside, to imagine someone who resembles you, doing what you are doing right now, and to imagine that this creature is not conscious. This is easy to do once we realize that imaginings have assignments. To imagine a genuine zombie case, you assign both: (i) that this creature is a microphysical duplicate of you; and (ii) that this creature has no conscious experience. Imagining a situation without *both* (i) and (ii) is not yet to imagine a zombie case. If, e.g., you imagine someone who looks just like you and is assigned *merely* to have no conscious experience, then you have not yet imagined a zombie case because your imagining is consistent with your imagined twin being microphysically dissimilar.

By contrast, we cannot imagine zombies from the inside.[15] Consider a neutral example: is a child who understands that rocks are non-conscious inanimate things able to imagine from the inside being a (non-conscious) rock? Seven-year-old Weezie is in the back of a station wagon, and sets out to imagine from the inside being a (non-conscious) rock.[16] She will not succeed. Here is what she can do: She can *say* she is a rock. She can *behave* like a rock, perhaps by curling herself into a ball, closing her eyes, and allowing her body to be moved by the motions of the car. She can *pretend* to be a rock (and may do the other two things in the service of so pretending). She may be disposed to answer questions or make utterances consistent with her pretense ("Mom, I'm a rock, I can't walk. You have to lift me out of the car." "That doesn't hurt. Rocks don't feel pain.") But none of this is to imagine being a rock from the inside. Since there is nothing it is like to be a rock, and she knows this, there is nothing for her to imagine. How would she imagine it? How can she imagine having *no* experiences *from the inside*? We have been developing the theory that to imagine from the inside is to imagine certain experiences. One can imagine experiences that are quite minimal: visually, nothing but blackness, aurally, nothing but silence, and so on. Imagining minimal experiences is quite different from imagining *absence* of experiences. To remove the experiences from the imagined situation is either to cease imagining altogether or to imagine from the outside.

Let's be fair to Weezie. A child may not realize that being non-conscious means more than seeing nothing, hearing nothing, smelling nothing, …, it means not having a point of view at all. And if Weezie does not realize this, then in virtue of not realizing this, perhaps she does imagine being a (non-conscious) rock. Perhaps. This raises interesting questions about the minimal accuracy requirements for imagining. "Imagine" is a word with ambiguous success grammar, which we can see by distinguishing between successfully getting yourself

[14] I may have to use a first-person *concept* to imagine zombies from the outside. But we are examining imagining, not concepts. Imagining from the inside is not the same as deploying a first-personal concept in imagination.

[15] For a similar line of reasoning see Marcus (2004).

[16] The parenthetical "non-conscious" is essential. She can imagine being a conscious singing and dancing rock — see the Being a Singing Candelabrum case in section 3.2. The issue is whether she can imagine from the inside being something such that it is *part of the imagining* that the something is non-conscious.

into an imaginative mental state (call this the imaginative *project)* and an *accurate* imagining. I imagine my great-great-grandmother. I can succeed in doing this even if the face I picture in my mind's eye does not even remotely resemble my great-great-grandmother. It is assigned that the woman I visualize is my great-great-grandmother. The imaginative project is successful, though it does not accurately depict the actual world (or any past state of the actual world). One reason this may happen is that I am misinformed about my great-great-grandmother's appearance. I mistook an old photograph of Cyd Charisse for my great-great-grandmother, so the face I visualized resembled Cyd Charisse's.[17] Another is that I have no idea what great-great-grandmother looks like, and the Charisse-like image is the one that sprang to mind. My inability to *accurately* imagine the likeness of my great-great-grandmother (except by accident) does not prevent me from successfully completing the imaginative project.

In discussing the epistemology of what it is like, it is vital not to confuse project success with accuracy. Compare the following two principles:

1. if you cannot imagine being X, then you do not know what it is like to be an X.
2. if you cannot (non-accidentally) accurately imagine being an X, then you do not know what it is like to be an X.

Nagel rests on the second, not the first, principle in "What is it like to be a bat?"[18] Weezie may imagine from the inside being a bat, succeeding in that *project*, without *accurately* imagining being a bat. Say she imagines having incredibly acute hearing, allowing her to locate objects that make the slightest sound via echolocation. She sees a bug on the window, and imagines that she has not seen it, but echolocated it. (Her "echolocation" imagery might even be visual, as it usually is in the movies.) Or she hears her cat in the next room, but imagines that the sound is too faint for human ears to detect. A great deal of this will be done by assignment, and that is the reason why it is not an accurate imagining; all the important details are filled in via assignment, rather than accurately qualitatively imagined. Still, there is some sense in which she has succeeded in imagining being a bat. Nagel need not deny any of this, because he relies on the second principle.

How much accuracy is required to succeed in the imaginative project? That is an interesting question for another occasion, but if bat example is any indication, it looks like not very much.[19] However there is a limit. To return to rocks and zombies, barring confusion on the part of the imaginer, whatever one does one is bound to fail in the project of imagining being a non-conscious rock, or being a non-conscious zombie. You cannot imagine being a zombie from the inside.

We've seen that you can imagine zombies from the *outside* but that you cannot imagine being a zombie from the *inside*. What about the other case on which dualists often rely, imagining being disembodied?

---

[17] This does not make it Cyd Charisse that I imagine. To be clear, I am not looking at the photo and imagining that *that woman* is my great-great-grandmother. I am imagining my great-great-grandmother and I conjure up an image of a woman who resembles Cyd Charisse.

[18] See Nagel (1974, pp. 520–21); see also Joyce (2003).

[19] See also the Being a Singing Candelabrum case in section 3.2 below.

2.2 Disembodiment

The Cartesian-inspired disembodiment thought experiment is perhaps the most famous thought experiment employing imagination from the inside.[20] We'll focus on the question of whether I can imagine myself existing without any material things existing.

In our discussion of imagining from the inside and assignment, we uncovered an ambiguity in phrases like "imagining myself existing." When I imagine myself doing such-and-such, this might mean that I imagine *Peter Kung* doing those things. On this reading, when I imagine being Tiger Woods (as in Being Tiger – Inside), I am not imagining myself existing. Peter Kung does not figure as a character in that story. Alternatively, it might mean that I imagine the subject of the imagined experiences, *Ego*, doing such-and-such. In my imaginings, Ego is myself, in the sense that Ego is the head I am "inside" when I imagine; the experiences that Ego is imagined to have are the very ones the author, Peter Kung, imagines. Ego's imagined thoughts involving first-person pronouns refer to Ego.[21] However, as we stressed in section 2, Ego is not necessarily identical to Peter Kung. The subject of the imagined experiences may be Peter Kung, but it also may be someone else; Tiger Woods, for instance. It depends on what assignments are in play.

Hence to imagine myself existing without any material things existing could be to imagine one of two stories. In the first, what is imagined from the inside is that some subject, Ego, has certain experiences, experiences that match the ones now being imagined. What is also true in the imagined situation, by virtue of assignment, is that there are no material objects. For example, imagine from the inside that you seem to be gazing into the face of Justin Timberlake when, in fact, there are no material objects. You imagine a particular experience; the "picture" in your mind's eye resembles Justin Timberlake. This is delivered by the qualitative content of the inside imagining. That there are no material objects is an assignment, and not part of the qualitative content of the inside imagining. (I don't have to "picture" *no material objects*.)

Notice that this first story makes no mention of Peter Kung. That's because it isn't *true* in the first story that the disembodied subject is Peter Kung. I imagine being disembodied, but Ego, the disembodied subject, has no assigned identity.

A distinct, second story would be to imagine *Peter Kung* being disembodied. It seems like I can do that as well. I imagine from the inside, as before, seeming to look at Justin Timberlake's face, and imagine that there are no material objects. As before, this involves the assignment that there are no material objects. What differentiates this story from the previous one is that I am supposed to be imagining myself, *Peter Kung*. Peter Kung is a character in this story. This requires the assignment that Ego = Peter Kung, that the subject having these imagined experiences is the actual me, Peter Kung (rather than just *some subject*, or *Tiger Woods*).

2.3 Summing Up

Let me sum up where we are. In our sensory imaginings, some things we "picture" in the

[20] Hart (1988) and Swinburne (1997) offer vivid examples. Descartes himself wasn't interested in what we could *imagine*, since he viewed imagination as too limited a faculty for the real distinction argument.

[21] See Velleman (1996) for an extended discussion of these "notional reflexive" thoughts.

mind's eye, other contents are assigned. In imagining from the inside, not only do we imagine seeing, hearing, and doing things, but we also imagine seeing, hearing, and doing things *from* the perspective of the character doing them. Call that character "Ego." We imagine Ego's *experiences*. By contrast, in imagining from the outside, we imagine worldly things like people, trees, chairs, and so on; since there is no Ego in imagining from the outside, we do not imagine experiences.

We noted that Ego's identity is flexible, and assignment establishes the identity of the Ego character (if Ego has an identity). Sometimes we imagine from the inside being ourselves — in my case, assign Ego = Peter Kung — but other times we imagine being someone else — assign Ego = Tiger Woods, or Ego = Napoleon.

This account has two benefits. Imagination is, of course, fascinating in its own right and is frequently employed in philosophical analysis on a wide range of topics.[22] We benefit from having a theoretical account of it.

The second benefit is the link to modal epistemology. As I've hinted at several points, my account allows for imagining the impossible. We haven't addressed worries about modal epistemology, but that is as it should be. There's no reason to question the commonsense view that we easily imagine the various cases described above. It's the job of a theory of imagination to explain what's happening when we do it. We should take as data that those cases are imaginable and adjust our modal epistemological theory accordingly.

Now it's time to ask whether the fact that we can imagine impossible situations dooms imagination as a guide to possibility. I think the answer is no. We'll see in the next section that the qualitative/assigned distinction allows us to explain away imagined impossibilities by showing that those imaginings provide no evidence for possibility. In section 3.3 we will use our new understanding of imagining and modal epistemology to analyze the zombie and disembodiment thought experiments. I compare my analysis to others in the literature in section 4.

3 INSIDE MODAL EPISTEMOLOGY

In discussions of zombie and disembodiment thought experiments, dualists and materialists disagree about whether those cases show (via the relevant arguments) that dualism is true. Some instances of this disagreement hinge on imagination. The issue is whether our putative imagining of those cases provides evidence that zombies and disembodiment are metaphysically possible. You might think that it's because they disagree about whether imagining provides evidence for possibility. But often that isn't the sticking point.

Suppose it is granted that imagining is evidence for possibility. If parties agree on that starting point, why the divergent opinions about the efficacy of zombie and disembodiment thought experiments? In the literature two competing themes emerge, one championed by dualist proponents of the thought experiments, the other by anti-dualist opponents. Dualists start with the assumption that imagining is *prima facie* evidence for possibility, tinker with it to handle familiar Kripke and Goldbach cases, and claim the result is a workable modal epistemology of thought experiments. Nothing in that account excludes zombie or disembodiment thought experiments, so we ought to concede that one or both succeed. In

[22] See the topics mentioned in footnote 1.

short, dualists find materialists' resistance *ad hoc*.[23]

Anti-dualists reply that, no, a workable modal epistemology of thought experiments is not so permissive. "One should be wary of intuitions based on the first-person perspective," cautions Nagel (2002), "since they can easily create illusions of conceivability" (p. 216). Anti-dualists stress *illusions of conceivability*.

Both camps have a point. Anti-dualists are correct that it is not as straightforward as dualists think to craft a successful first-person thought experiment. Dualists are correct that rejecting the zombie and disembodiment thought experiments is often *ad hoc*.

As an example of the *ad hoc* problem, take the pessimistic contention that dualist thought experiments fail because, even if the mind-body connection were metaphysically necessary, zombies and disembodiment would remain imaginable.

> I want now to argue not directly for the necessary connection between mind and brain, but rather for the position that even if there were such a necessary connection, it would still appear through this kind of [imaginability] test that there was not. (Nagel 2002, p. 216)

> The brain has physical properties we can grasp, and variations in these correlate with changes in consciousness, but we cannot draw the veil that conceals the manner of their connection. Not grasping the nature of the connection, it strikes us as deeply contingent; we cannot make the assertion of a necessary connection intelligible to ourselves. (McGinn 1989, p. 364)[24]

These brute observations should not worry dualists.[25] For any proposition we deem to be contingent as the result of a thought experiment, it is of course true that if, unbeknownst to us the proposition were necessary, then we could still imagine it false. I imagine wearing a green rather than a yellow shirt today. Even if it were necessary that I wear a yellow shirt today, I would still easily imagine not doing it. That fact does nothing to undermine my confidence that my wearing a yellow shirt today is contingent. It is crucial, then, that anti-dualists *explain why* first-person thought experiments fail, and they do so in a way that is not *ad hoc*. Absent some further story dualists will rightly take the persisting tug of contingency as evidence in their favor.

This demand for explanation dooms a whole range of otherwise carefully developed anti-dualist positions, positions that embrace what Stoljar (2005b) calls "the phenomenal concept strategy."[26] The anti-dualist authors meticulously explicate how phenomenal concepts

---

[23] Either *ad hoc* or a tacit slide into modal skepticism.

[24] I substitute "imagine" for "conceive," since my focus here is on imagination. See also Wilson (1982) and Shoemaker (1983), as cited by Taliaferro (1994, pp. 135–36), who both assert that imagining mind-body distinctness just means we haven't yet noticed a contradiction.

[25] And they don't; see Taliaferro (1994, pp. 193–96) and Yablo (1990, p. 187). Neither Nagel nor McGinn stop with this brute observation; McGinn offers version of "the phenomenal concept strategy" that I discuss next in the text, while Nagel favors the splicing view we examine in section 4.2. My point in this paragraph is polemical: we should be wary of any strategy that, when all is said and done, relies on just this brute observation.

[26] See Stoljar (2005b, 2006) for authors who adopt the phenomenal concept strategy. Stoljar offers

are unique, detail what relations phenomenal concepts bear to sensory experience, emphasize how different they are from scientific concepts generally and neurophysiological concepts in particular, and then claim that modal arguments in philosophy of mind fail to appreciate this difference.

> Given these differences between sensory concepts and physical concepts, a sensory state and its nomologically correlated brain state would *seem* contingently related even if they were necessarily one. (Hill & McLaughlin 1999, p. 449)

Notice the similarity to the brute observation above. These authors explain why, as a psychological fact, we are able to imagine the phenomenal and the physical as distinct. But they fail to satisfactorily explain — or rather, explain *away* — the modal epistemological significance of that fact. The dualist *grants* the striking differences between phenomenal and physical concepts, but she thinks that this *illuminates* certain metaphysical truths by way of first-person thought experiments.[27] The burden is on the anti-dualist to explain, given the assumption that thought experiments generally are trustworthy, why phenomenal-physical thought experiments are not to be trusted.

In the rest of this section I show how some *independently motivated* modal epistemological constraints flows from the theory of inside imagining developed in section 2. I won't be developing a full-fledged modal epistemology here; I'll focus on when imagining does *not* provide evidence for possibility.[28] I'll emphasize the epistemically illegitimate role that assignments play in some thought experiments. These constraints will allow us to see why, e.g., imagining being a singing candelabrum provides no evidence that there could be a singing candelabrum. They'll also give us a new understanding of the zombie and disembodiment thought experiments.

3.1 "Assignment makes imagining the impossible possible"

Suppose we grant, with many authors, that imagining does provide evidence for possibility.[29] Even if we grant this assumption, we don't have to treat every imagining as providing equally good evidence. (Visual perception provides evidence for actuality, but someone wearing rose-colored glasses shouldn't take it that they have visual evidence that everything is red.) A *probative* imagining is one that provides evidence for possibility. We'd like a principled way to declare cases of imagining the impossible to be non-probative. Assignments are the key. Roughly, an imagining that P will not be evidence that P is possible if P's truth in the imagined situation follows from the assignments alone; in slogan form, "assignment makes imagining the impossible possible." The reason is that assignments are virtually unconstrained, and what minimal constraints there are have no modal epistemological value.

---

effective criticism of the phenomenal concept strategy, though his argument differs from the one I sketch next in the text.

[27] See, for example, Swinburne (1997, p. 315).

[28] For the full modal epistemology, see Kung (forthcoming).

[29] See Chalmers (2002), Geirsson (2005), Gendler (2000), Gregory (2004), Hart (1988), Hill (1997), Kung (forthcoming), and Yablo (1993). Many, many others assume it without argument, e.g., from (famously) Hume (1978) to Nagel (1974) to Lewis (1986) to anyone who uses imaginary cases in philosophical argument.

I assume that qualitatively imagining that P provides evidence that P is possible.[30] A full-fledged positive modal epistemology would need to defend that assumption. I don't have the space to embark on that project here; for our purposes I'll simply grant the assumption.

Assignments, however, are a different story. Even granting our assumption, I think it's clear that assignments can play only a limited role in providing evidence for possibility. Assigned content has few constraints. If assignment has no constraints at all — if for any P, we can imagine that P via assignment — then imagining via assignment provides no evidence for possibility because it fails to discriminate between possible and impossible Ps. Imagining via assignment would be no more probative than supposition.

There may be some constraints on assignment. It is difficult to imagine (even via assignment) that 1+1=79, for example. Here is a proposal that explains our difficulty: the principal constraint on assignment is absolute certainty. By 'absolute certainty' I mean the strongest possible psychological certainty: to have absolutely no doubts at all, for there to be nothing one is more certain of.[31] This kind of absolute certainty marks propositions like *that* 2=2 and few others. I'll assume that psychological certainty confers the very best epistemic status.[32] I propose: so long as we find P *non-certain* — true for all we know with absolute certainty — we will be able to imagine P via assignment. The proposal makes intuitive sense: in being less than absolutely certain that a proposition is true, we leave a tiny bit of room to imagine a way for it to be false. For propositions that are absolutely certain, there isn't even this tiny bit of room. I am extremely confident that I have hands. But I am not absolutely certain of it; I can imagine a skeptical scenario in which I don't. On the other hand because I am absolutely certain that 2=2, I can imagine no way for it to be false.[33]

Let P be some proposition whose possibility we are trying to establish via imagining. The mere fact that we find P non-certain, and hence are capable of making the assignments required to make P true in the imagined situation, provides no evidence for P's possibility. It would be very odd if *non-certainty* counted as evidence of P's possibility. To be non-certain is to fall short of the very best epistemic position one can be in; how can failing to be in the best epistemic position be evidence for some proposition's possibility, particularly when we note that total ignorance is one way to fail to be in the best epistemic position? We seek positive evidence for our claims of possibility, but assignments do not provide it; they merely reflect our less-than-ideal epistemic position. What goes for each assignment individually goes for what follows from the assignments alone: if it is only by virtue of non-certainty that one is able to assign Q and assign R in the same imagining, and P is true in the imagining only in virtue of

[30] I assume that to qualitatively imagine that P, P must be something qualitative.

[31] See Unger (1975, ch. II).

[32] See Reed (2008) for a discussion. Rejecting the assumption only undermines imagining via assignment as evidence for possibility because there would be no *epistemic* constraints on assignments

[33] My proposal is that any *proposition* we find non-certain is one that we can imagine it via assignment. Some indexical *sentences* like "I exist" or "I am here now" are absolutely certain as well, but that is because we understand how those sentences get their truth values. The linguistic meaning of "I exist" guarantees that when a sentence with that meaning is thought or uttered by me, it will be true, even though the content expressed at that moment, that Peter Kung exists, is not certain. It is not absolutely certain that I am Peter Kung. Thus my proposal correctly predicts that I will not have any difficulty imagining that I do not exist because the propositional content that Peter Kung exists is not certain.

Q and R, then the imagining does not provide evidence for P's possibility.[34]

Let me add that if the reader rejects my proposal that assignments are constrained by absolute certainty, this does not salvage imagining via assignment as evidence for possibility. If there are *no* constraints on assignment — if we can imagine *anything* via assignment — that makes the modal epistemological situation worse not better for imagining via assignment.

We now have a test — call it the Assignment Test — for showing that an imagining does not provide evidence for possibility.

3.2 Diagnoses and Authentication

With our "assignment makes imagining the impossible possible" slogan we can now diagnose cases that are even more recherché than the Tiger Woods cases from section 2. Consider an imagining inspired by the Disney movie *Beauty and the Beast*. I happen to own a candelabrum, perched on a pile of books, and I imagine being *that* candelabrum.

> **Being a Singing Candelabrum**
>
> I imagine from the inside being a conscious dancing, singing candelabrum. I imagine moving myself (walking?) over to the mirror, peering into it, and seeing three brassy prongs sprouting from my vaguely pyramidal base. I imagine greeting Justin Timberlake, who has come to visit, with repeated choruses of "Be our guest!"

The singing candelabrum worries proponents of an imagination based-modal epistemology because, while it is plausible that we can imagine being a singing candelabrum (the movie *invites* us to do just that), it seems impossible for me to *be* a singing candelabrum. But the Assignment Test shows that the worry is overblown.

In this story I am "inside" the candelabrum's "head". To make that a fact in the story it must be assigned that Ego is *that* candelabrum, so the assignment amounts to: that *that* candelabrum is the subject of experiences, and the experiences are of *this* character (where 'this' demonstrates my current imaginative experiences). Hence the assignments establish that the candelabrum has experiences. By the Assignment Test, because of the dependence on pure assignment, this imagining provides no evidence that the candelabrum could have experiences. We can happily accept, with commonsense, that we can imagine being a singing candelabrum, and yet don't need to admit that a candelabrum could be conscious. At most,

[34] Another constraint on assignment that has been much discussed in the literature recently is imaginative resistance. A full discussion would take us away from our main point, so I'll confine myself to a few remarks. The puzzle is to explain why it is difficult to imagine, e.g., counter-moral claims, such as "In killing her baby, Giselda did the right thing; after all, it was a girl" (Gendler 2006, p. 159). See Gendler (2006) and Weatherson (2004) for description of this and various interconnected puzzles. Two broad solution strategies have emerged in the literature. On one strategy (that Gendler favors), our difficulty arises because in some sense we don't *want* to imagine what we're asked to imagine. It's clear that desiring that not-P in the way that leads to imaginative resistance fails to generate evidence that P is impossible. That is a more elaborate kind of wishful thinking; modal wishful thinking, in this case. And, more importantly, merely lacking desires that would lead to imaginative resistance obviously does not count in favor of modal belief. The second strategy (Weatherson's) is that resistance emerges when we violate certain supervenience relations. I take this strategy to be consistent with my own non-certainty constraint; when we are certain these supervenience relations hold, we cannot imagine their falsity. And the *absence* of supervenience-based resistance is not a modal epistemic credit. I discuss the puzzle and its in relation to modal epistemology at length in Kung (forthcoming).

what Being a Singing Candelabrum establishes is that there could be experiences of a certain sort, experiences of singing, of looking in a mirror and seeing a three pronged object, and so on.

What about the Being Tiger Woods case from section 1.2? When I imagine being someone else, as when I imagine being Tiger Woods from the inside, it is assigned that Ego is Tiger Woods. Now this imagining might provide evidence for possibility, because we can *authenticate* it. Consider this perceptual analogy.

> Suppose you see several figures in the distance, though they are too far away for you to discern who they are. However, if you have it on good authority that the one on the left is Justin Timberlake, you may very well be perceptually justified in believing that Justin is punching one of the other guys. Your perceptual experience alone does not provide you with evidence that Justin is punching someone. You don't see the figure on the left *as* Justin (at least not initially). But your experience coupled with the knowledge that the figure on the left is Justin Timberlake does give you reason to believe that Justin is punching someone.

Carrying over the analogy, if you have *independent* evidence that an assignment is possible then an imagining that uses that assignment can still be probative. Your independent evidence could come from (a) evidence of actuality (and hence possibility); or (b) another imagining that does not merely make the same assignment; or (c) some other source. (I'll set this third option aside, since I am not discussing other sources of modal evidence in this paper.) Your imagining coupled with this independent authenticating evidence provides evidence that P is possible.

To see how this works, let's use an example of imagining from the outside. Imagine from the outside Tiger Woods winning the 1997 Masters; imagine that instead of being paired with Constantino Rocca in the final round he is paired with Justin Timberlake. Visualize Timberlake shaking Woods' hand to congratulate him on his historic victory. In this imagining, the qualitative part fills in what the two figures look like. The assigned part fills in who the two figures *are*. The figure on the left is imagined *as* Woods, the figure on the right *as* Timberlake. Does this imagining provide evidence that Woods could have played Timberlake in his historic final round?

For this imagining to provide evidence for possibility, we must authenticate each of the assignments. One thing that is assigned is *that Timberlake exists and that he is conscious*. Can we authenticate this assignment? Yes we can. We have independent evidence that the assignment is metaphysically possible because we know that Timberlake *does* exist and *is* conscious, hence we know that he could exist and be conscious. The italicized assignment is authenticated. This example suggests that, because most *de re* imaginings are established by assignment, most *de re* imaginings will be authenticated by appeal to actuality. That makes intuitive sense: we take the particulars we've encountered in the real world and imagine them in different situations. That those particulars could exist isn't established by imagination; it's established by the fact that they actually exist.

Now that we understand how authentication works, let's turn to some more philosophically interesting cases.

In imagining being Tiger Woods, we assign that Ego = Tiger Woods. Does this provide evidence for metaphysical possibility? Since Tiger actually exists, we already know that he could be — because he *is* — the subject of experiences. Hence the Assignment Test plus authentication does not rule out imagining being Tiger Woods as providing evidence that Tiger could have experiences like the ones we imagine.

Remember, though, that my imagining being Tiger Woods is not the same thing as my imagining *Peter Kung* being Tiger Woods. An important distinction in section 2 was between Ego, the character in the imagined situation, and me, Peter Kung, the author. It's that distinction that helped us realize that imagining from the inside does not have to be imagining about Peter Kung. To imagine *Peter Kung* being Tiger Woods I would need two assignments, first that Ego = Tiger Woods and second that Ego = Peter Kung. Those two assignments would make it true in the story that Peter Kung = Tiger Woods. Hence by the Assignment Test, this imagining provides no evidence that Peter Kung could be Tiger Woods *unless* we provide independent evidence that Peter Kung = Tiger Woods is possible. Our prospects for doing that look rather dim. It is not actually true that Peter Kung = Tiger Woods, and it is hard to see how to devise another imagining to establish that Peter Kung could be Tiger Woods that wouldn't simply assign that very fact.

Similarly, we can't authenticate the assignment in the candelabrum case, that the candelabrum is the subject of experiences. We cannot appeal to actuality; the candelabrum is not actually conscious. It doesn't look like we can appeal to imagining; that is what we were *using* this imagining to establish!

Thus the modal epistemology we developed shows how some things that we can easily imagine provide no evidence for possibility. We can imagine being a singing candelabrum, but that provides no evidence that there could be a singing candelabrum. We can imagine being Tiger Woods, but that provides no evidence that Peter Kung could be Tiger Woods.

Let's now see what our analysis says about the zombie and disembodiment cases.

3.3 Zombies and Disembodiment

The zombie case is very straightforward. Recall the results of section 2.1: you cannot imagine your zombie twin from the inside. You can imagine it from the outside, and doing so requires two assignments: (i) that this creature is a microphysical duplicate of you; and (ii) that this creature has no conscious experience. (Remember imagining a situation without *both* (i) and (ii) is not yet to imagine a zombie case.) The existence of a creature that is both your microphysical duplicate and non-conscious is established by assignments alone and cannot be authenticated by independent means. Hence imagining a zombie from the outside provides no evidence that there could be such a creature.[35]

The disembodiment case requires a bit more work to analyze.

In section 2.2 I argued that we *can* imagine from the inside being disembodied; that is, we can imagine from the inside ourselves existing without any material objects existing. As I just reiterated, Ego is not necessarily identical to Peter Kung. So there are two different stories

[35] Absent a story like mine about imagining zombies from the outside, anti-dualists have not successfully blocked the dualist argument. Nagel (1974), Hill (1997), and Marcus (2004) for example, have nothing to say about imagining a zombie from the outside.

that fit the description "imagining being disembodied."

First, I might be imagining such that Peter Kung does *not* figure as a character the story. What is imagined from the inside is that *some subject*, Ego, has certain experiences, experiences that match the ones now being imagined. That there are no material objects is an assignment.

For this imagining to be probative, the assignment must be authenticated. To do so via imagining would require imagining that there are no material objects without assignment. Though I am not confident about this, I am sympathetic to the idea that this can be done. Imagine from the outside that there are no material objects, that is, imagine that space is entirely empty.[36] You might do this by visually imagining nothing but blackness, aurally imagining nothing but silence, and so on. What are you "picturing" in your mind's eye? I suggest you picture nothing: your picture is *of* emptiness. Think about a similar perceptual experience: you walk into a pitch black cave and see…absolutely nothing. Your visual experience represents no things in front of you. I suggest that in both the imagining and the perceiving, the qualitative component presents emptiness.

The issue is whether there are assignments as well. You might argue: "Compare imagining nothing to perceiving nothing again. When you open your eyes in the cave, you would have the same experience whether or not the region of space had a few electrons in it. Hence in the imagining case, you need to assign that the region really is empty." This is a mistake. In the cave, the qualitative component of your visual experience presents emptiness: it looks to you as if there is nothing in front of you. Your visual experience is wrong, there are electrons in front of you, but you don't see them. The world decides whether your perceptual experiences are right or wrong. In imagining, no external facts make it that case that your imaginative experiences are right or wrong. When you visually imagine exactly three red spheres on a table, because what you picture — the qualitative content of your imagining — presents exactly three red spheres on the table, you need not also assign that there are no invisible red spheres on the table. There is no corresponding worry about your picturing things wrongly. Hence when you visually imagine nothing, you need not also assign that there *really is* nothing there.

A more vexed question is whether every imagining comes with the assignment, "that's all." In the cave, vision tells you nothing about what is behind you. In imagining, the worry would be that what you picture — the qualitative content of your imagining — tells you nothing about the rest of the imagined world; it tells you nothing about what is "off stage." So when you visually imagine exactly three red spheres — period, not just on the table, or in the room, but three red spheres in the world — while you picture exactly three red spheres, you need the assignment "and that's all." It looks hopeless to authenticate a "that's all" assignment via more visual imagining. I myself am inclined to think that the "that's all" assignment is not required, that the imagined world contains just what you the imaginer put in it, so you don't need to add, "and there's nothing more." I confess that I do not know how to argue for this claim. But assume for the sake of argument that it is correct.[37] Assume that

[36] By "empty" I mean empty of *material* objects.

[37] Bear in mind that the dualist has other disembodiment thought experiments ready to fall back on, and many involve a "more mundane" disembodiment: you awaken and look in the mirror to find your body gradually falling away (Hart 1988, pp. 52–53), or you have an out of body experience, where you float up above your body (Taliaferro, pp. 191–92).

we can visually imagine absolutely nothing, without assignment, hence we have evidence that there could be no material objects. How far does this get the dualist?

In imagining from the inside myself disembodied, seeming to gaze into the face of Justin Timberlake, it is assigned that there are no material objects. We are assuming for the sake of argument that this assignment can be authenticated, so according to my view, we can then use this assignment in other probative imaginings. This means we can use it in imagining myself disembodied. Thus, under this assumption, imagining does provide evidence that there could be disembodied subjects of experience, that *some* subject could be having experiences of a certain character, yet there be no material objects.

This conclusion is striking; it is enough to falsify the identity theory. But it does not scuttle materialism entirely, since it is compatible with the kind of mental-physical dependence that some materialists care about.[38] The fact that there are possible worlds with disembodied subjects of experience says nothing about actual physical creatures and their mental lives. Some materialists insist only that actual mental lives be physically constituted: the physical description of me is the complete description of me. These materialists would be happy with the conclusion that our actual mental lives are physically constituted even if disembodied subjects roam other non-actual possible worlds. Something like this is what some materialists intend by the mental-physical supervenience claim. Dualists need the stronger conclusion that some *actual* conscious creature could be disembodied if they are to falsify this minimal materialism. Hence even under the arguable assumption that we can authenticate the assignment that there are no material objects, dualists do not yet have enough to reject minimal materialism.

To disprove minimal materialism, I would need to show that I, *Peter Kung*, could be a disembodied soul. Then materialism would be in serious trouble. Building on our discussion about assigning the identity of Ego, we can see that imagining provides *no* evidence that such a thing is possible. I imagine from the inside, as before, seeming to look at Justin Timberlake's face, and imagine that there are no material objects. As before, this involves the assignment that there are no material objects. What differentiates this project from the one just discussed is that I am supposed to be imagining myself, *Peter Kung*. This requires the assignment that Ego = Peter Kung, that the subject having these imagined experiences is the actual me, Peter Kung (rather than just *some subject*, or *Tiger Woods*). Now we have two assignments, that I, Peter Kung, am the subject of experiences, and that there are no material objects. These two assignments jointly entail the target Cartesian proposition: that Peter Kung is a disembodied subject. Thus my inside imagining provides no evidence that Peter Kung could be a disembodied subject unless the joint assignments can be authenticated. Since this is what we set out to do by our Cartesian imagining, we can see that this project fails. Imagining provides no evidence that I, Peter Kung, could be disembodied.

Hence even under the contestable assumption that the "no material objects" assignment can be authenticated, the dualist can draw only moderate support from disembodiment thought experiments. Imagining may provides evidence that there could be disembodied subjects of experience, but it does not provide evidence that I, Peter Kung, could be disembodied. The kind of mental-physical dependence that, according to some, defines physicalism, survives disembodiment thought experiments.

---

[38] Stoljar (2005a) discusses what constitutes "minimal physicalism."

4 ALTERNATIVE MODAL EPISTEMOLOGIES

We have just seen how my assignment-based proposal permits the plausible claim that we routinely imagine the impossible, and yet allows us to preserve imagining as evidence for possibility because it excludes these imagined impossibilities from providing (misleading) evidence for possibility. Hence the slogan, "assignment makes imagining the impossible possible." I used my view to diffuse worries about puzzling cases in the literature: being Tiger Woods, being a singing candelabrum, zombies, and being disembodied.

In the remainder of the paper I'll briefly canvass other views in the literature. Each alternative view has one of the following problems:

i. It gets ordinary cases of imagining wrong.

ii. It is *ad hoc*.

iii. It is overly permissive.

I'll argue that, by comparison, my view avoids these problems. Due to considerations of space, I will not attempt a thorough presentation and criticism of those views here; rather my aim will be to illuminate features of my own view by contrasting it with others in the literature.

4.1 Rephrasal

Kripke famously argues that often when it appears we imagine some impossible E, really we imagine some closely related and possible E* that is easily mistaken for E. We should reconstrue, or rephrase, our imagining of E as an imagining of E*. Kripkeans reject my methodological strategy of analyzing imagination prior to modal epistemology: we have to adjust, i.e., reconstrue, what we imagine to respect the modal facts.

Take a strange case, like imagining from the inside that Peter Kung is Tiger Woods. A Kripkean might insist that we cannot imagine that case; we cannot imagine that Tiger Woods and Peter Kung are one and the same person. Rather we imagine that Tiger Woods also goes by the name 'Peter Kung', that he was adopted by my actual father, Edward Kung, and that's why he hugs Edward on the 18th green. This kind of rephrasal has become standard in modal epistemology.

As I argue elsewhere, the Kripkean strategy seems misguided.[39] Kripkeans propose an error theory: in many cases, we are wrong about what we imagine. That seems implausible for something as familiar as imagining. I maintain we can imagine even *a posteriori* necessarily false propositions that Kripke discusses, such as water is XYZ. Kripke thinks we don't imagine that *water* is XYZ; rather we imagine some other clear substance — one that falls from the sky, is essential to life, is called 'water', and so on — being composed of XYZ. But can't we imagine being *wrong* about what water is? Imagine it turning out that we were *wrong* to believe that water is H_2O. The whole point of the imagining is that we have made a *mistake* about the composition of water. If we rephrase so that you are not imagining *water*, then you are not really imagining being wrong. In my view, we imaginatively assign that the stuff is water, and that it is composed of XYZ. It's these assignments that make it a story where water is composed of XYZ.

If we set aside our preoccupation with modality, we find many other cases that seem

[39] See Kung (2009).

easily imaginable and yet according to many are impossible. Kripke himself argues for the necessity of origins, yet I have no difficulty imagining the following case.[40]

> **My Real Parents**
>
> I imagine making a shocking discovery. Edward and Marcia, the people who raised me, turn out not to be my biological parents. I imagine uncovering documents in a dusty attic that reveals I was adopted as an infant, that my real biological parents, martial artist Bruce Lee and baker Sara Lee, were forced by poverty to put me up for adoption.

The rephrasal strategy is strained in this case. The Kripkean strategy requires us to claim that I am not imagining myself, Peter Kung, I am imagining someone who looks like me, who is called 'Peter Kung' and was raised by Edward and Marcia, but is the biological child of Bruce Lee and Sara Lee. We can devise such a rephrasal, but it is hard to accept that the rephrasal is what we really imagine. Am I really supposed to be confused about whether I am imagining myself? My view permits that My Real Parents is easily imagined: in my imagining, I assign that Bruce Lee is my biological father and Sara Lee is my biological mother.[41]

Assignments render Kripke's error theory unnecessary. In the water-XYZ case, it is assigned both that the substance is water and that it is XYZ. Since it follows from the assignments alone that water = XYZ, and these assignments are not plausibly authenticated, on my view imagining a water-XYZ case provides no evidence that water could be XYZ. In imagining My Real Parents, I assign Bruce Lee as my biological father. Hence on my view I have no evidence that it is possible for Bruce Lee to be my biological father.

Because we can explain both how assignments allow us to imagine the impossible and why we shouldn't take these assignment-based imaginings as evidence for possibility, we need not resort to a Kripkean error theory.

4.2 Splicing

In a famous footnote, Nagel suggests a way to defeat dualist thought experiments by appealing to something like the inside vs. outside distinction:

> Where the imagination of physical features is perceptual [roughly, from the outside] and the imagination of mental features is sympathetic [roughly, from the inside], it appears to us that we can imagine any experience occurring without its associated brain state, and vice versa. The relation between them will appear contingent even if it is necessary, because of the independence of the disparate types of imagination. (1974, pp. 526–27n11).

Nagel's quite attractive thought is that there is something fishy about imaginings that splice imagining from the inside and imagining from the outside. Hill (1997) develops Nagel's suggestion at some length.[42] Take the example of imagining C-fibers firing without pain. Hill

---

[40] Thanks to Janet Levin for helpful suggestions here.

[41] Levin (2007) explores why this kind of case is harder to reconstrue.

[42] Hill's (1997) splicing strategy is different from the one he recommends later in a coauthored paper, Hill & McLaughlin (1999). The later paper endorses a version of the phenomenal concept strategy — sensory *concepts* are importantly different from phenomenal *concepts* — a strategy I criticized at the beginning of section 3. Splicing concerns different imaginative *faculties*.

contends that we imagine pain (or its absence) from the inside, and we imagine brain processes like C-fibers firing (or their absence) from the outside. The resulting imagining is a result of splicing, and so, he claims, no evidence for possibility.

Hill recognizes that, so far, this Nagelian solution is ad hoc, because it seems specially tailored to meet the very cases dualists rely on. He needs some independent reason to think that splicing is unreliable. The dualist, after all, will claim that splicing *is* reliable,[43] and that we have latitude to combine imaginings from the inside with imaginings from the outside because such combinations are genuine metaphysical possibilities. Hill answers the ad hoc charge by pointing to more *general* imaginative faculties:

> "commonsense" imagination =df generates images of "a commonsense phenomenon (i.e., a phenomenon to which we have access by a commonsense faculty of awareness)" (p. 71).
>
> "theoretical" imagination =df generates images of "a 'theoretical phenomenon' (i.e., a phenomenon to which we have access only by theory construction and laboratory apparatus)" (p. 71).

Hill takes inside imagination to be one variety the commonsense imagination, and the kind of outside imagining required to imagine brain processes to be one variety of theoretical imagination. These more general imaginative faculties allow him, he believes, to explain why splicing is unreliable. Kripke's famous *a posteriori* necessity cases give us reason, Hill thinks, to distrust splicings of commonsense and theoretical imagination.

Hill takes as data Kripke's *a posteriori* necessities — water is necessarily H_2O, heat is necessarily mean kinetic molecular energy — and suggests that imagining an impossible situation — water without H_2O, heat with molecular motion — requires combining images from the "commonsense" imagination with images from the "theoretical" imagination. We "commonsense" imagine water and heat, and "theoretically" imagine the lack of H_2O or molecular motion. Kripke's cases demonstrate the unreliability of commonsense/theoretical splicing. Imaginings that splice inside and outside imagining are subsumed under this unreliable commonsense/theoretical splicing mechanism. Hence inside/outside splicing is unreliable.[44]

Hill's proposal fails, and the way it fails is, I think, instructive. It fails in a way that serves as a general caution for those who question the modal epistemological value of first-person thought experiments (or, indeed, any kind of thought experiment).

Any theory attempting to say something illuminating about imaginability and possibility cannot simply take Kripke's *a posteriori* necessities as data. Kripke argues that some *a posteriori* identity statements are necessary. To do so, he must explain away our apparent ability to imagine the falsity of these identity statements; he must explain away our apparent ability to imagine water composed of XYZ, Hesperus distinct from Phosphorus, or this table

---

[43] Recall that our working assumption in this paper is that imagining is in the usual run of things a good guide to possibility. Hill recognizes the need to *argue* that splicing is unrealiable.

[44] See Hill (1997, p. 71). Note that the commonsense/theoretical distinction does not divide right along the inside/outside line. Some of the things we commonsense imagine, such as water, are imagined from the outside. Hill goes on to use splicing to diffuse zombie and disembodiment thought experiments for dualism, but since as I'll argue in the text Hill's modal epistemology is unjustified, his analysis of zombie and disembodiment cases is beside the point.

made of ice. Kripke offers his rephrasal strategy, which we just discussed.

Hill faces a dilemma. Suppose on the one hand he accepts Kripke's rephrasal strategy for explaining away the apparent separability of water-H_2O, Hesperus and Phosphorus, and so on. Hill has vindicated Kripke's *a posteriori* necessities data but at the cost of rendering his own theory superfluous. Why should we think that there is any problem with *splicing*? We should not take our apparent ability to imagine water composed of XYZ as reason to believe that such a thing is genuinely possible for the reasons that Kripke gives. Splicing is not part of Kripke's explanation. We need independent reason for thinking that splicing gives us *another* reason to reject these cases.

Alternatively Hill could reject Kripke's rephrasal strategy; he could take Kripke to have given bad arguments for a true conclusion. This is what he in fact does.[45] However once he dispenses with Kripke's rephrasal strategy he needs some other way to explain away the apparent ability to imagine water without H_2O, Hesperus distinct from Phosphorus, and so on. Otherwise he should just conclude that these are not identities, or that such identities are not necessary, or that 'water' is not a rigid designator. As far as I can see, Hill offers no alternative explanation. Thus Hill is not entitled to take Kripke's *a posteriori* necessities as data. He must earn them on his own view.

This objection to Hill highlights a crucial feature of my view: I explain, on independently plausible grounds, *why* assignment-based imaginings provide no evidence for possibility. My view avoids the ad hoc charge; it is not tailored to handle, nor does it make brute assumptions about, prominent cases in the literature. Hence my view can explain Kripke's cases (better, I contended above, than Kripke's rephrasal strategy can). With any modal epistemology, the theory needs to do more than divide imaginings into those that do and those that do not count as evidence for possibility; it needs an independently plausible explain *why* that is the right division. Hill's theory cannot do that.[46]

4.3 Verificationism

A *verificationist* modal epistemology, of the sort inspired by Shoemaker (1993), answers modal epistemological questions by determining *what one would be justified in believing* if presented with an epistemic situation qualitatively like the one imagined.[47] You gather imagined evidence and decide what it is reasonable to believe on the basis of such imagined evidence. You have evidence that P is possible just in case you can imagine gathering evidence that would justify you in believing that P. Many philosophers are drawn to this view, some explicitly, but most often tacitly.[48] If you take visually imagining that P to be imagining *seeing* that P, or *hearing* that P, that encourages you to wonder whether seeing could justify you in believing that P. You take yourself to imagine perceptual experiences, and then ask what those imagined perceptual experiences justify.

---

[45] Hill regards Kripke's explanation as "fundamentally misguided…for…in non-pathological circumstances introspection gives us pretty accurate access to the contents of our own states of imagination" (1997, p. 83n10).

[46] I would level a similar objection at proponents of the phenomenal concept strategy.

[47] I say "inspired by" because I am less concerned with whether this is Shoemaker's considered view. I do not have space to weigh the alternative proposals I find in Shoemaker's text.

[48] Nagel (1998, 2002) seems drawn to this view. See also Peacocke (1985).

We can see how the verificationist modal epistemology plays out when Shoemaker analyzes imagining a zombie from the outside:

> But can one imagine from the third-person perspective a case in which someone is not in pain despite having in his brain an optimal candidate for being a total realization of pain?
>
> I think the answer is no. The reason is that playing the causal role of pain… will essentially involve producing precisely the kinds of behavior that serve as our third-person basis for ascribing pain. We cannot be in a position to judge about someone *both* that she is not in pain *and* that she is in a state that influences her behavior in just the ways we think influences behavior. (p. 506)

If we were to observe a creature who was in the usual brain state accompanying pain, who screamed and hollered when poked with a hot iron, and so on, we would of course judge that creature was in pain. In an imagined *perceiving* of such a creature, we are similarly forced to conclude that they are in pain. The person not only behaves like, but also has a brain exactly like, someone in pain, so any test run from "outside" will reach the verdict that the creature is in pain. In the imagined situation we could not justifiably think that the creature was not in pain. Hence according to verificationist modal epistemology we have no reason to think such a situation — brain state typical of pain without the painful feeling — is possible.

Note the contrast with my assignment-based view. It looks like according to the verificationist we *cannot* imagine a zombie from the outside, because no amount of imaginative investigation would reveal the creature to be a zombie. Verificationism is, in this sense, an error theory just like Kripke's rephrasal view. If we take ourselves to be imagining a zombie, we must be mistaken. We can imagine no such thing. In my view, we easily imagine zombies simply by assigning that the imagined creature feels no pain. As we have seen, imagining zombies from the outside fails the Assignment Test, meaning that it provides no evidence that zombies are possible.

Shoemaker deploys the verificationist approach in his much more elaborate attempt to defuse dualist first-person thought experiments (pp. 508–513). The approach is clearer, however, without the complications of Shoemaker's analysis of those cases, so let's confine ourselves to a more general discussion of verificationist modal epistemology.[49]

Verificationism is far too permissive. We can imagine evidence for just about any claim, no matter how outlandish or obviously impossible. For any claim that is not absolutely certain, we can imagine gathering evidence that it is false.

On one gloss of epistemic possibility, it is epistemically possible that most of our beliefs are false because, for all we know for certain, many of our beliefs *are* false. Kripke specifically cautions us against mistaking what we can imagine ourselves being convinced of for *genuine* possibility.

> …I (or some other conscious being) could have been *qualitatively in the same epistemic situation* that in fact obtains, I could have the same

[49] If it turned out that the verificationist approach had merit, it would be worth turning to the particulars of Shoemaker's analysis. However, as we will see, the general approach is problematic.

> sensory evidence that I in fact have, about *a table* which was made of ice. …[W]hen I speak of the possibility of the table turning out to be made of various things I am speaking loosely. *This* table itself could not have had an origin different from the one it in fact had, but in a situation qualitatively identical to this one…the room could have contained *a table made of ice* in place of this one. (1972, p. 142)[50]

We can imagine being convinced that *just about* any of our beliefs are false. We can imagine discovering that water is an element (imagine boldface headlines in the *New York Times* tomorrow, and subsequent hand wringing asking how we could have been so badly mistaken), or that Justin Timberlake is a robot (imagine he is accidentally unmasked by a child while filming a music video), or that there is a subtle error in the incompleteness proof. It certainly seems like I can imagine *this table* in front of me is made of specially engineered ice. Even though we can imagine having experiences such that, were we to have them, we would believe the negation of all sorts of necessary truths, we should not accept that these are genuine metaphysical possibilities.

These considerations seem to me fatal — something is wrong with the verificationism's very approach. If imagining is to be a guide to possibility, it is not because we imagine gathering evidence and decide what it is reasonable to believe on the basis of such imagined evidence. We need to focus on imagining *as* evidence, and not imagined evidence.

4.4 No Contradictions

Now it is time to examine the sort of modal epistemology non-Kripkean dualists favor. Several authors hold something like this initially plausible version of Van Cleve's (1983) strong conceivability claim: if you "non-negligently" imagine a thought experiment — you don't overlook obvious contradictions, don't misapply concepts, don't fall into *de dicto/de re* confusions, and so on — then you possess *prima facie* evidence that the thought experiment is possible.[51] Proponents of this "no contradictions" view stress fallibility, admitting that imagining sometimes leads to illusion, but hold that in the main it is reliable.

> [I]magination is our favored, and perhaps only, basic epistemic access to nonactual possibility…[W]hen one has imagined as inventively as one is able, and it has always seemed clearly to be the case that p in whatever scenario one has spun out, and one has no good reason to think otherwise, then one has good reason to think it is possible that p. In neither case is one guaranteed success: objectivity always leaves open the possibility of error. But…one has the best of the basic sort of evidence available. (Hart 1988, pp. 28–29)

These authors insert the "no negligence" qualification to handle cases acknowledged cases where we (seem) to imagine the impossible. For instance, it might initially seem that we can imagine a barber who shaves all and only those who shaves themselves, but one can do so only

[50] See also pp. 103–04.

[51] Hart (1988, ch. 2), Swinburne (1993, ch. 3), and Taliaferro (1994, pp. 134–39) and (1986, pp. 100–101) each endorse versions of this claim. The "non-negligence" locution is Taliaferro's

by neglecting the obvious contradiction waiting to be found.[52]

In light of our discussions above it should be clear why this plausible-sounding proposal is incorrect. It is *not enough* to imagine that P "non-negligently" or without apparent contradiction. Even dualists ought to admit upon reflection that assignments provide no (new) evidence for possibility. When you imagine a situation solely via assignment, then even if you are otherwise diligent, even if your imagining presents no obvious or hidden contradictions, you still have yet to produce evidence for possibility.[53]

5 CONCLUSION

To sum up: I have defended an independently plausible theory of imagining from the inside and an independently motivated modal epistemology. Because the theory of imagination is developed prior to the modal epistemology, the resulting account both handles a range of ordinary cases of imagining well and is not *ad hoc*. It offers *independent* grounds for regarding purely assignment-based imagining, such as imagined impossibilities, as non-probative. It explains familiar cases in the modal epistemology literature, and sheds new light on the zombie and disembodiment thought experiments.

The modal epistemology in this paper addresses only sensory imagining as evidence for possibility. I have not addressed other putative sources of modal evidence, such as non-imaginative conceivability or intuition. My suspicion is that sensory imagining will fare better than the alternatives. Conceiving is, I suspect, akin to sensory imagining without the qualitative component — that is, pure assignment. We elicit intuitions by constructing imaginary cases, so I hypothesize that intuition will depend on imagination. But discussion of the alternatives will have to wait for another occasion.[54]

[52] See Chalmers (2002, pp. 151–55) on prima facie without ideal conceivability, also Taliaferro (1994, p. 138) and Yablo (1990, p. 175).

[53] In a series of papers, Yablo (1990, 1993, 2006) has articulated what I regard as a combination of the "no negligence" view and the Kripkean error theory. Yablo's view is, roughly, that we can take imagining as positive evidence for possibility unless we have reason to think there is a "modal defeater"; he explains at length what does and what does not count as a modal defeater, and in his (2006) he has a lot to say about the limits of first-person imagining. I dispute Yablo's model; I believe he is mistaken about how modal defeaters work, and I believe his model commits him to implausible claims about what is unimaginable. For example, as I read him, Yablo would have to say that My Real Parents is unimaginable. But Yablo's view is sophisticated enough that it deserves separate (and fairly lengthy) treatment, so I do not discuss it here. For more detailed discussion, see my (2009).

[54] Ideas in this paper have been brewing for a while, so I have many people to acknowledge. Thanks Ned Block, Paul Boghossian, and Tom Nagel. I am also grateful to Yuval Avnur, David B. Barnett, Ray Buchanan, Don Garrett, Pete A. Graham (UMass Amherst), Paul Hurley, Brian Keeley, Janet Levin, Laura Perini, Dion Scott-Kakures, Alex Rajczi, Peter Thielke, Rivka Weinberg, and Masahiro Yamada. I received excellent comments from students in my 2009 imagination seminar at Pomona College. Thanks to you all. Special thanks to Amy Kind and Peter J. Graham (UC Riverside…yes, there are two) for incredibly helpful feedback on multiple drafts. And finally thanks to anonymous referees for two other journals.

REFERENCES

Berkeley, G. (1979). *Three dialogues between Hylas and Philonous* (R. Adams, Ed.). Indianapolis, IN: Hackett. (Original work published 1713)

Block, N., Flanagan, O., & Güzeldere, G. (Eds.). (1997). *The nature of consciousness: Philosophical debates*. Cambridge, MA: MIT Press/Bradford Books.

Chalmers, D. (2002). Does conceivability entail possibility? In Gendler & Hawthorne (Eds.), *Conceivability and possibility* (pp. 145–200). New York: Oxford University Press.

Culpepper, C. (2008, July 16). Twenty-three items about Anthony Kim. *Los Angeles Times*.

Currie, G. (1995). *Image and mind: film, philosophy and cognitive science*. Cambridge University Press.

Geirrson, H. (2005). Conceivability and defeasible modal justification. *Philosophical Studies, 122*, 279–304.

Gendler, T.S. (2000). *Thought experiments: On the powers and limits of imaginary cases*. New York: Garland.

Gendler, T. (2006). Imaginative resistance revisited. In Nichols, S. (Ed.), *The architecture of the imagination* (pp. 149–74). Oxford University Press.

Gregory, D. (2004). Imagining possibilities. *Philosophy and Phenomenological Research, 69(2)*, 327–48.

Hart, W.D. (1988). *The engines of the soul*. Cambridge University Press.

Hill, C. (1997). Imaginability, conceivability, possibility and the mind-body problem. *Philosophical Studies, 87*, 61–85.

Hill, C. & McLaughlin, B. (1999). There are fewer things in reality than are dreamt of in Chalmers's philosophy. *Philosophy and Phenomenological Research, 59(2)*, 445–54.

Hume, D. (1978). A *treatise of human nature*. 2nd ed. L. A. Selby-Bigge & P. H. Nidditch (Eds.). Oxford: Clarendon Press. (Original work published 1739–40)

Jonze, S. (Director). (1999). *Being John Malkovich* [motion picture]. United States: Gramercy Pictures.

Joyce, P. (2003). Imagining experiences correctly. *Proceedings of the Aristotelian Society, 103*, 361-370.

Kind, A. (2001). Putting the image back in imagination. *Philosophy and Phenomenological Research, 62*, 85–109.

Kripke, S. (1980). *Naming and necessity*. Cambridge, MA: Harvard University Press.

Kung, P. (2009). You really do imagine it: Against error theories of imagination. Unpublished manuscript, Pomona College, Claremont, CA.

Kung, P. (forthcoming). Imagining as a guide to possibility. *Philosophy and Phenomenological Research*.

Levin, J. (2007). Can modal intuitions be evidence for essentialist claims? *Inquiry, 50(3)*, 253–69.

Lewis, D. (1986). *On the plurality of worlds*. Oxford: Blackwell.

Mackie, J. (1980). The transcendental "I." In Van Straaton, Z. (Ed.), *Philosophical subjects: Essays presented to P.F. Strawson*. Oxford: Clarendon Press, 1980.

Marcus, E. (2004). Why zombies are inconceivable. *Australasian Journal of Philosophy, 82(3)*, 477–90.

Martin, M.G.F. (2002). The transparency of experience. *Mind and Language, 17(4)*, 376–425.

McGinn, C. (1989). Can we solve the mind–body problem? *Mind, 98(391)*, 349–66.

Nagel, T. (1974). What is it like to be a bat? *Philosophical Review, 53(4)*, 435–50. Reprinted in Block, Flanagan, & Güzeldere, (Eds.), (1997), 519–27.

Nagel, T. (1998). Conceiving the impossible and the mind-body problem. *Philosophy, 73(285)*, 337-352

Nagel, T. (2002). The psychophysical nexus. In *Concealment and exposure and other essays*. New York: Oxford University Press.

Noordhof, P. (2002). Imagining objects and imagining experiences. *Mind and Language, 17(4)*, 426–55.

Peacocke, C. (1985). Imagination, experience, and possibility: A Berkeleian view defended. In J. Foster & H. Robinson (Eds.), *Essays on Berkeley*. Oxford: Clarendon Press.

Reed, B. (2008). Certainty. In E. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy, Fall 2008 Edition*.

URL = <http://plato.stanford.edu/archives/fall2008/entries/certainty/>.

Reynolds, S. (1989). Imagining oneself to be another. *Noûs, 23(5)*, 615–33.

Searle, J. (1992). *The rediscovery of the mind*. Cambridge, MA: MIT Press/Bradford Books.

Shoemaker, S. (1993). The first-person perspective. *Proceedings of the American Philosophical Association, 68(2)*, 7–22. Reprinted in Block, Flanagan, & Güzeldere, (Eds.), (1997), 503–15.

Smith, M. (1997). Imagining from the inside. In R. Allen & M. Smith, (Eds.), *Film theory and philosophy*. Oxford University Press.

Stoljar, D. (2005a). Physicalism. In E. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy, Winter 2005 Edition*.

URL = <http://plato.stanford.edu/archives/win2005/entries/physicalism/>

Stoljar, D. (2005b). Physicalism and phenomenal concepts. *Mind & Language, 20(5)*, 469–94.

Stoljar, D. (2006). *Ignorance and imagination: The epistemic origin of the problem of consciousness*. Oxford: Oxford University Press.

Swinburne, R. (1993). *The coherence of theism* (Rev. ed.). Oxford: Clarendon Press.

Swinburne, R. (1997). *The evolution of the soul* (Rev. ed.). Oxford: Clarendon Press.

Taliaferro, C. (1986). A modal argument for dualism. *Southern Journal of Philosophy, 24(1)*, 95–108.

Taliaferro, C. (1994). *Consciousness and the mind of God*. Cambridge: Cambridge University Press.

Tye, M. (1993). On the possibility of disembodied existence. *Australasian Journal of Philosophy*, *61(3)*, 275–82.

Unger, P. (1975). *Ignorance: A case for scepticism*. Oxford: Clarendon Press.

Velleman, J.D. (1996). Self to self. *Philosophical Review*, *105(1)*, 39–76.

Walton, K. (1990). *Mimesis as make-believe*. Cambridge, MA: Harvard University Press.

Weatherson, B. (2004). Morality, fiction, and possibility. *Philosophers Imprint*, *4(3)*, 1–27.

Wollheim, R. (1987). *Painting as art*. London: Thames & Hudson.

White, A. (1987). Visualizing and imagining seeing. *Analysis*, *47(4)*, 221–24.

Williams, B. (1973). Imagination and the self. In *Problems of the self*. Cambridge University Press.

Wilson, M. (1982). *Descartes*. Boston: Routledge & Kegan Paul.

Yablo, S. (1990). The real distinction between mind and body. *Canadian Journal of Philosophy supplementary volume 16*, 149–201.

Yablo, S. (1993). Is conceivability a guide to possibility? *Philosophy and Phenomenological Research*, *53(1)*, 1–42.

Yablo, S. (2006). No fools cold: Notes on illusions of possibility. In M. Garcia-Carpintero & J. Macia (Eds.), *Two-dimensionalsSemantics*. Oxford University Press.

Zimmerman, D. (1991). Two Cartesian arguments for the simplicity of the soul. *American Philosophical Quarterly*, *28(3)*,