

AN ANTI-ANTI-FUNCTIONALIST ACCOUNT OF CONSCIOUSNESS

Brian D. Earp

Department of Cognitive Science

Yale University, United States

brian.earp@gmail.com

Abstract. *In this paper I scrutinize the so-called China Brain thought experiment famously articulated by Ned Block (1978) to see whether it refutes functionalism as a theory of consciousness. I argue that it does not. Block's case rests on a single premise, P, in the following argument: (P) a creature with a China Brain would lack qualitative experience, despite its being (R) a functional replica of you or me, and working under the assumption that we have (M) rich mental life complete with qualia. Yet if a China Brain-creature (Mr. Li) were truly and completely R, and yet P were true as the thought experiment purports to show, functionalism would have to be false. Since I want to reject that conclusion, I must do one of the following: (1) deny that Mr. Li is (or, stronger, could even in principle be) truly R in the first place; (2) grant that Mr. Li is (or could in principle be) truly R and grant that P, but deny that we have M; or (3) grant that Mr. Li is (or could in principle be) truly R and grant that we have M, but deny that P. I opt for the third option. First, I paint a picture according to which Mr. Li's having M would not seem as implausible as it does at first. That is, I try to make the premise P seem less intuitive. But in case I am unpersuasive in this first approach, I employ a second tactic as well. This one is to suggest that we cannot know, even in principle whether P, so it is an irrelevant intuition upon which to base the China Brain attack. In any case, the China Brain fails to provide a true counterexample to functionalism.*

Keywords: functionalism, consciousness, China Brain, qualia

Introduction¹

The aim of this paper is to scrutinize the so-called China Brain thought experiment famously articulated by Ned Block (1978) to see whether it refutes functionalism as a theory of consciousness². My title undermines any chance for suspense: I will argue that it does not.

¹ *Caveat lector*—my introduction is not short and my footnotes are not few. I hope that by laying out the territory in advance, I will have less work to do to be convincing later on.

² For a specific account of functionalism, see Part I.

While I am neither the first to defend functionalism, nor the first to criticize Block's particular attempt to undermine it, I hope to offer a fresh perspective on an old debate.

What am I defending here? Functionalism in philosophy of mind is "the doctrine that what makes something a mental state of a particular type does not depend on its internal constitution, but rather on the way it functions, or the role it plays, in the system of which it is a part."³ If this view is correct, there are exciting implications for our world. It would mean that rich mental states could⁴ be instantiated in something other than a human brain. Robots, for example, with first-person, qualitative experience of the world could conceivably be built. Functionalism is worth making a case for.

This is how I will construct my argument. In Part I, I will give an account of functionalism sufficient to see what it is the China Brain is designed to refute. Then in Part II, I will explain the thought experiment in the clearest terms and most charitable manner I am able—that is, in what I take to be its strongest form. In the course of doing so, I hope to show that the scenario rests crucially upon a single premise, P, in the following (bare bones-version) argument⁵:

P a creature with a China Brain would lack qualitative experience⁶ // despite its being
 R a *functional replica* of you or me // and working under the assumption that we have
 M rich mental life complete with qualia.

Since functionalism proposes to give a full account of mental states⁷, its being true would require that any creature meeting R also replicate (our) M. Therefore, if a China Brain-creature (Mr. Li for short) were truly and completely R, and yet P were true as the thought experiment purports to show, functionalism would have to be false. Since I want to reject that conclusion, I must do one of the following:

- (1) *deny* that Mr. Li is (or, stronger, could even in principle be) truly R in the first place—making P either irrelevant or meaningless
- (2) *grant* that Mr. Li is* truly R *and* grant that P, but deny that we have M
- (3) *grant* that Mr. Li is* truly R *and* grant that we have M, but deny that P

* or could at least in principle be

³ Levin, Janet, "Functionalism", *The Stanford Encyclopedia of Philosophy* (Winter 2007 Edition), Edward N. Zalta (ed.), forthcoming URL = <http://plato.stanford.edu/archives/win2007/entries/functionalism/>.

⁴ In principle, at least.

⁵ In what follows, only the text to the left of the '/' is the claim expressed by the capital letter

⁶ Some formulations of the argument seem to suggest that a China Brain-creature would lack *all* mental life, period, including propositional attitudes wrt its environment. I think this is going too far. Even a simple robot has certain limited "beliefs" about its environment ("the robot *believes* there is an object in its path, so it is modifying its go-to-target algorithm"); it is only a robot's *feeling* anything in some qualitative sense that should seem potentially out of whack. Since I want to give the China Brain thought experiment its best showing, I will interpret it narrowly as (just) an Absent Qualia argument—for my money, that is where its prima facie intuitive appeal is strongest.

⁷ Meaning, it holds that *all there is* to being in a given mental state is to be in a certain functional state.

I could also deny that P *and* withhold the status of being R from Mr. Li, but that would be overkill. I also decline to pursue (2) for what I hope are obvious reasons⁸. And because I want to give the best account of the China Brain as I can before attacking it, I will refrain from (1) as well⁹. This leaves me with (3)—the denial of P, granting R-status to Mr. Li and M-status to you and me. Accordingly, the task of Part III will be to cast serious doubt on P from a number of angles. Taken together, these tactics are meant to put the China Brain thought experiment—vis-à-vis what I claim is its crucial premise—on very shaky footing indeed.

As a preview, here are the "angles" I shall employ. First, I will attempt to paint a picture according to which Mr. Li's having M—or "being conscious" to put it that way—would not seem as implausible as it does at first. That is, I will try to make the premise P seem less intuitively obvious. I can do this because I believe the intuition that P is more a function of rhetorical set up than a reflection of undeniable metaphysical truth. But in case I am unpersuasive and you remain convinced that P is the intuitive position, I will employ a second tactic. This tactic, mirroring Brian Loar's phenomenal concept strategy¹⁰, is to suggest that *we cannot know, even in principle* whether P. This maneuver, which might also be termed an Explanatory Gap argument, is more complicated—so I will wait until Part III to spell it out in detail. In broadest outline, however, it asserts that *phenomenal concepts*¹¹—the concepts by which we come to know (of) qualia in first place and which exhaust our first-person¹² phenomenal knowledge about them—cannot by themselves rule out the possibility that qualia *are* in fact functional properties.¹³ But conversely, there is no third-person¹⁴ way to identify qualia that could even in principle "link up"¹⁵ to our original (first person) phenomenal

⁸ Unless you are Daniel Dennett (1988), I trust you will agree that we have M.

⁹ It is hard to see in any case why *in principle* my functional set-up could not be duplicated by a something like a China Brain, properly articulated. That such a scenario is *metaphysically* possible seems in fact to be an explicit entailment of functionalism's prized multiple-realizability thesis. My argument, then, is not that it is in principle impossible to realize my mind in a China Brain, but that (a) counter to the intuition, such a brain *just might* also replicate my mental life—up to and including qualia (at best) or (b) that, at the very least, we can't *know* that it wouldn't.

¹⁰ in "Phenomenal States" (1990)

¹¹ Phenomenal concepts are a special breed of what Brian Loar calls "recognitional concepts." Recognitional concepts in general are concepts we have when we recognize something as being one (or another-one) *of those* (whatever it is)—without relying on any background information. Phenomenal concepts are defined as recognitional concepts that specifically pick out *qualia* or '*what it is like*' experiences. Importantly, they do so *directly*, via "non-contingent modes of presentation." Here is an example to explain. Suppose you experience a pain. There is something it is like to have that experience, a certain *quale*. Moreover, you recognize it as being "one of those," that is, a pain-experience, without reference to any theory, list of properties, or background knowledge whatsoever. The concepts you have in virtue of this experience are *phenomenal concepts*.

¹² Direct, introspective, from-our-own-case...

¹³ Or: properties of fine-enough grained functional systems—like a Mr. Li.

¹⁴ It doesn't matter whether the "third person" account is theoretical/analytical or more empirical/physical. It refers to any attempt to pick out qualia besides one's own, directly.

¹⁵ Colloquialisms are messy. What I mean by 'could not link up' is that, as Loar (1990) puts it, "recognitional concepts and theoretical concepts are in general conceptually independent" (p. 299). Each mode of recognition is 'untranslatable' in terms of the other. This is the nature of the explanatory gap. It states that if physicalism were true, there would be *in principle* no way to *explain* qualia in theoretical terms, yet they would nevertheless be physical or physical-functional properties.

concepts—so it would be impossible to determine which third-person account is right. Either way, qualia's being *functional properties* (and therefore something present in Mr. Li) remains a live option. The argument that it would seem intuitively *weird* if Mr. Li were conscious, then, would be quite beside the point. We simply cannot *know* short of being Mr. Li himself.

Now I can survey my entire argument in just a few lines:

1. Only Mr. Li can even in principle know whether he is conscious
2. We are not Mr. Li, so our intuitions in the matter are irrelevant¹⁶
3. The China Brain thought experiment rests entirely upon our having certain intuitions

4. So the China Brain thought experiment is irrelevant

—perhaps that is worse than being false! In any case the China Brain fails to provide a true counterexample to functionalism. If you still aren't convinced, I do have one last nail for its coffin, but that is for the hardy reader who reaches my conclusion.

Nota bene. I will assume in this paper that physicalism is true and dualism false. This is a big deal. For if dualistic states were on the table, qualia would be first in line for membership of that class—and I'd have no explanatory gap to use for my Explanatory Gap argument¹⁷. But this assumption is not disingenuous for the debate at hand. While I won't articulate them here, there are many independent reasons to reject dualism. The best Chinese Brain argument, then, at least in my view, would be the one that tries to disprove functionalism *given physicalism*. For more on this, see footnote 6.

Part I. Functionalism

The China Brain thought experiment purports to provide a counterexample to functionalism. But what is 'functionalism' exactly? Here I provide a quick general account of the theory and then give some attention to the particular features its China Brain opponent (at any rate) assumes it has. This section will be quite brief. In its best formulation, the China Brain argument (if it worked) should be effective against just about *any* plausible account of a functionalist theory of mind—so I can work in broad strokes.

Functionalism in general holds that a thing is defined not by its internal constitution but by the role or *function* it plays a given system¹⁸. Functionalism as a theory of mind, then, holds that *mental states* (such as "being in pain") are defined in just that way—as functional states of a (cognitive) system. So what is a "functional state"?

A functional state is defined by three types of clauses: inputs, outputs, and internal interactions. Input clauses denote the kind of event typically *causing* a given mental state; output clauses denote the kind of behavior typically *caused by* a given mental state; and internal clauses denote and describe the internal *interactions* between mental states. Together, these three clauses form a sort of "sentence" which, if satisfied by some system, tells you that the system is "in" a (given) mental state. Take the mental state "believing there is a tiger in front of one." A reasonable functional sentence for this state would read like this: *Input* =

¹⁶ This means, of course, that I have to sacrifice my first "angle" to my second. But that's only if you are stubborn. If I can convince you that P is a bad intuition, then please—read no further!

¹⁷ In other words, the explanation of qualia would be simply "they are dualistic states of mind."

¹⁸ Take as a (rough) analogy a burglar alarm. What it means to be a burglar alarm (as opposed to a baseball bat, a refrigerator, or what have you) is to serve a certain function, namely advertising the presence of an unexpected visitor. What it is that fills this role could be any number of things: a stack of tin cans behind the door, a tripwire attached to a siren, a string of bells, etc. So long as my contraption serves (or tends to serve) its function, it counts as a burglar alarm.

tiger causes *Internal* = *belief that tiger* causes *Output* = *running away*. Perhaps that's too simplistic, but it gives the idea. Actual functionalists may try to define these clauses—in other words, say what "counts" as, e.g., an input—more precisely. Some, known as "common sense" functionalists, will provide sentences not very unlike my example. They argue that all it *means* to be in a given mental state is to realize a very common-sense functional sentence that even an average Joe could grasp. Other functionalists prefer to concoct more and more fine-grained sentences, narrowing the scope of possible clause-fillers to systems within the reach of empirical science. But really any type of functionalism will do so far as the China Brain is concerned. Here is why.

Functionalism holds that mental states *are* functional states. If I want to refute the theory, all I have to do is come up with¹⁹ two systems with identical functional states but non-identical mental states. So *however* you decide to define the relevant functional structure-plus-roles of some fully conscious creature, I will just dream up a system that could replicate that structure and fill those roles but be nevertheless unconscious or otherwise lacking in mental life. The China Brain argument uses this method. It starts with a paradigm example of a conscious creature (you or I) and says, in effect, "I can *meet* your functional requirements—*however* you choose to define them—with some bizarre creature that just *can't* be conscious. See how I do it."

Let us turn to that argument now.

Part II. The China Brain

I promised to give the best account of the China Brain thought experiment I could manage. Here is my try.

Suppose you are sitting across from an old friend, your best buddy Mr. Li, in a coffee shop downtown. You're interacting in all the ways typical for old friends. Talking, laughing, trading jokes. Reminiscing about old times. You're getting along quite well! After a bit, though, Mr. Li sets down the coffee he's been sipping and gives you a funny look. He leans in a bit. Listen, he says, I have to tell you something. I can't keep living this lie. Alright, you say, suddenly concerned. What is it? Mr. Li pauses, sighs. My friend—I am not what you think. Take a look...

He reaches for his neck and makes a motion that can only be described as an "unlatching gesture." Suddenly, grotesquely, a large part of his skull—your best friend's skull—swings open. You cannot believe your eyes.

The first thing you notice is what you don't notice. A brain. There is no brain in Mr. Li's skull. Instead you see two stiff antennae with a knob at the end of each—a wireless transmitter and receiver, Mr. Li explains—attached at the base of the skull and protruding up into the cavity where a brain should have been. There *are* some remaining knots and sheets of organic brain-looking tissue leading (internally) from the eyes, nose, ears, tongue, and spinal cord—all of which seem to connect up with a mass of

¹⁹ I'm vague here on purpose. "Come up with" could mean "conceive"—which I think is not enough—*or* "provide an intuitively plausible account"—which then forces the anti-functionalist to show *why* the intuition is good enough for knowledge. My whole argument is that she can't.

hair-thin wires at the base of the antennae. (It's crazier than you think, Mr. Li remarks. Those hair-sized wires are just the major link-ups. *Each neuron* is connected to a silicon pseudo-neuron leading ultimately to a transistor inside each antenna—but most of *those* connections are invisible to the naked eye.)

But where is your brain!?, you scream. Mr. Li nudges his skull shut and gazes at you sadly. Let me show you, he whispers after a moment, his voice a bit coarse.

He leads you to a warehouse—a warehouse that is literally as big China—and opens the door. This is my brain, he says...

Enough drama. Let me cut to the chase. Inside the warehouse are two towering antennae—a receiver hooked up to an enormous monitor and a transmitter beaming data back to Mr. Li. There are also a billion Chinese people²⁰ yapping away on cell phones, eyes glued to the giant screen. The monitor is updated in real time with output-data from Mr. Li's antenna, as well as with instructions telling the Chinese people what they each must do. Their task as individuals is to replicate the action of the smallest functional part of a human brain, let's say a single neuron. A neuron receives a finite set of inputs (in the form of electrical charges) and dishes out the appropriate outputs according to its specific role in the whole network²¹ of neurons in a brain. Likewise, each Chinese person receives a finite set of inputs (in the form of phone calls) and dials the appropriate numbers according to the instructions on the monitor. All this data is gathered to the transmitter and beamed to Mr. Li—resulting in all the ordinary movements, gestures, and speech acts with which the scene began.

While such a set up is surely physically unbuildable, you can see how—at least in principle, and more fully articulated—this massive network of individuals, sending and receiving simple signals, could realize the functional structure of a brain²². Perhaps you'll want to modify the thought experiment. A single neuron, for example, probably plays several or many functional roles in a brain-system—not just one. In that case you'll need tens of billions, or maybe trillions of Chinese people, and a warehouse the size of a galaxy. It's also hard to see how all these signals are supposed to "hook up" to the colossal transmitter such that they could be properly beamed to the mini- one in Mr. Li's skull. There may be other issues as well. But in order to avoid option (1) from the introduction (denying that Mr. Li either is or could be truly R in the first place), I ask that you use your imagination to fill in any remaining gaps. Now, give yourself a moment to think—does Mr. Li have qualia? Consciousness? Phenomenal states? Is there "something it is like" to be Mr. Li?

That would be absurd! (you're supposed to say). Mr. Li's "brain"—and I don't care how well organized it is—is just a bunch of Chinese people gabbing away on cell phones! How could he *possibly* be conscious in the same way you and I are?

Part III. Denying that P

I think he is. This is where I try to show that the intuition that P—that Mr. Li lacks qualitative experience—isn't so obvious upon careful reflection. But first:

²⁰ China once had about a billion people, and it used to be thought that brain had a billion neurons. Hence Block's "China brain."

²¹ Or, more probably, a merely local network, itself playing some role in a larger network, etc.

²² The idea is that the whole functional structure of a brain, on this view, could be reduced to the sending and receiving of simple signals by individual simple units—neurons or whatever.

If I want to undermine the claim that P, I should start by identifying its grounding, or the reason for its initial appeal. There is a weaker and a stronger grounding. The weaker grounding is a "mere conceivability" basis for the claim. In this version, all the anti-functionalists has to say is, "I can imagine—without contradicting myself or disobeying any other laws of logic—a creature that is functionally identical to you but non-identical in mental life. If I can so conceive of such a creature, it must be possible. Therefore, functionalism is false." In the China Brain argument, of course, Mr. Li would play the role of a creature so conceived. But as I said, this version is weak. It rests entirely upon one of the messiest notions in philosophy, the conditional "if conceivable, then possible." Without going too much into it, I will provide a counterexample in this footnote.²³ Let us just say that this is not the way to go.

What is the stronger way to argue for P? It would be to carefully construct a plausible-sounding account of an R-status Mr. Li, and then make a positive claim: *we have intuitive evidence that P*. The next step would be to show that that intuitive evidence counts for something—that we can *know* in some sense that it points to the truth. At the very least we must show that the intuitive evidence E increases the probability of P by *some* amount—and to do that we had better be certain that the intuition is not outright false.

This is how I will proceed. I will first acknowledge that the China Brain thought experiment does better than "mere conceivability"—it makes a positive claim regarding the *intuition that P*. Now I shall attack this intuition according to the plan laid out in the Introduction. First—some reasons why the intuition isn't so convincing to begin with. Second—an argument to show that it wouldn't even matter if it were, since we cannot *know* whether it "points to the truth" or doesn't. Finally, in the conclusion, a hint that it might even be "outright false."

Why is the intuition that P not that convincing to begin with?

Think about it this way. The intuition is this: "It would be *completely crazy* to believe that a creature whose 'brain' is essentially billions of Chinese people talking on cell phones could *possibly* have qualitative experience like you or I." But I would ask you to consider whether is it any *less* crazy to think that a brain comprising billions of little *neurons* somehow "communicating" with each other would somehow "give rise to" the qualitative consciousness we actually do experience. If we believe the latter²⁴, why not the former?

I suggest it is a matter of scale more than anything²⁵. We can't coherently imagine (form a clear mental picture of) a warehouse the size of a galaxy, nor trillions of temporarily

²³ Suppose you live in the 17th century—a time I picked because it predates the age of molecular chemical analysis. Some clever Joe asks you whether *water* could conceivably be made up of some underlying, invisible-to-the-naked-eye structure, namely XYZ—any structure you like, just don't pick H₂O. "Sure," you say, "that seems perfectly *conceivable* to me." Is it therefore *possible*? No! Not if Hilary Putnam's Twin Earth (1975) thought experiment shows us anything! Water is *necessarily* H₂O—so to conceive otherwise is to conceive of an impossibility. This refutes "if conceivable, then possible"—and it is just one of countless conceivable (!) counterexamples. The anti-functionalists who wishes to retain this messy conditional for her Absent Qualia argument must modify the antecedent. "If *ideally* conceivable, then possible" she'll say. Okay, fine. But I don't know anyone excepting God who can meet that "if"—so this seems the wrong direction to go.

²⁴ Assuming we do – for this essay, remember, we are physicalists.

²⁵ Maybe it's also weird because it involves the realization of a brain by objects which *themselves* have brains—but this doesn't strike me as the right account of the intuition. *Neurons* could, for all it matters

"dumb" humans performing simple communicative tasks in an excruciatingly organized fashion. Our capacity to visualize as well as our corresponding intuitions are all out of whack at that scale. So it's no surprise that it's even *harder* to conceive of some property—qualia or anything—somehow "emerging" from or supervening on such a system. We can't grasp the system itself in the first place.

Yet we have no trouble grasping our own qualia on the emerging-from-neurons picture. Why? Again, scale²⁶. When we think of neurons, we think of them as microscopic entities all bundled together inside our heads. This is very convenient since that is "where" we experience qualia as well. In other words, our brains (in a strict sense) and our qualitative experiences (in a looser sense) are both "located" within our skulls, so they "go together" more intuitively—even though (or perhaps *because*) we don't spend too much time worrying about *how*, exactly, qualia and neurons might actually be related. But just imagine:

Suppose we could go with Ms. Frizzle on her Magic School Bus²⁷ on a journey—inside a human brain. We shrink down to the size of a neuron, say, or even smaller, and we take a look around. What do we see? Frank Jackson and David Braddon-Mitchell²⁸ describe the situation pretty well: "a whole mass of widely spaced entities interacting with each other in a way that [makes] no sense...that [forms] no intelligible over-all pattern from [our] perspective... All that [seems to be] happening is an inscrutable sending back and forth of simple signals" (p. 109). Suppose we'd never been exposed to a scientific theory about our own brains and thus were ignorant about neurons and the like. It would be very tempting indeed to say: "whatever it is we're looking at could not *possibly* be conscious²⁹ in the same way we are!" But that's just what we said about the China Brain. We'd be wrong if we said it now (talking about *neurons* at a gargantuan scale)—do we still think we were right to say it then?

If I've managed to convince you that P is not the intuitive position, then I've already defeated the China Brain, for:

1. The China Brain argument rests on the intuition that P
2. But P is not intuitive
3. So the China Brain argument fails.

If you are not convinced, however, I will move to my second strategy, the Phenomenal Concept Strategy/ Explanatory Gap Argument.

to the argument, actually have little "brains" comprised of, say, mini-neurons, or electrons, or whatever—but so long as all they *do* is what they *do in fact do*—that is, fulfill their functional role—it makes no difference what's inside. Suppose in the China Brain example, instead of people and cell phones when we opened the warehouse door, we saw blobs that looked like jelly executing some analogously simple communicative task. This would remove the applicability of any intuitions about "what sort of things are possible with respect to groups of human people"—but essentially the exact same thing would be going on. Whether it's people doing it, blobs of jelly, or neurons—should make no difference. This is precisely the consequence of the multiple-realizability thesis, and for my money it's rather intuitive.

²⁶ Also because we unconsciously conflate/confuse the direct access we have in virtue of phenomenal concepts with the theoretical account of brains given by science—but that shall be the focus of my second argument.

²⁷ If you don't know this reference, go to <http://www.scholastic.com/magicschoolbus/>. You don't need to know it to get what I'm saying.

²⁸ In *Philosophy of Mind and Cognition*. Malden: Blackwell Publishing, 2007.

²⁹ Or, 'be the conscious-making component of a physico-functionally defined creature-system'.

The Phenomenal Concept Strategy/ Explanatory Gap Argument

As a refresher, here is the outline of the argument I gave in the Introduction.

1. *Phenomenal concepts*—the concepts by which we come to know (of) qualia in first place and which exhaust our first-person phenomenal knowledge about them—cannot by themselves rule out the possibility that qualia *are* in fact functional properties.
2. But conversely, there is no third-person way to identify qualia that could even in principle "link up" to our original (first person) phenomenal concepts—so it would be impossible to determine which third-person account is right.
3. Therefore, either way, qualia's being *functional properties* (and therefore something present in Mr. Li) remains a live option.
4. Thus the argument that it would seem intuitively *weird* if Mr. Li were conscious would be quite beside the point. We simply cannot *know* short of being Mr. Li himself.

What does all this mean in plain(er) language? This:

Mr. Li aside, we know that *we* are conscious, qualia-experiencing creatures simply because we ourselves actually experience qualia—namely our own—directly. At the same time, however, we cannot experience anyone else's qualia besides our own. This is true by definition, but also obvious from experience. If we want to determine whether some *other* person has qualia, then, we must devise a third-person "objective" account of what qualia are identical to in the world as described by science, and then try to see whether the person (say, Mr. Li) has them under that description. But this is problematic. The very way in which we experience qualia (that is, from a first-person, direct perspective) *rules out* our being able to say with certainty what they are identical to in the physical-functional world. How could that be so? It is because any physical/ functional account we might devise is from a *third-person* perspective. By necessity, to move towards a third-person perspective is to give up the first-person perspective—you cannot have it both ways. Thus while the two perspectives (first- and third-person) can potentially pick out the same object, they cannot translate into each other's terms. There is no way to say *from a first person perspective*—the only perspective we *know* refers to qualia as such—which third person perspective actually picks out qualia as opposed to something else. Therefore, there is no way we can know that qualia are *not* properties of the third-person perspective account called functionalism. And if they *are* properties of functionalism, then Mr. Li would *have* to have them no matter what our intuitions say. We cannot *know* that he does not.

Conclusion

This is an exciting conclusion because I still have one last move to make. If I may use idioms gratuitously, this is where I "tie it all together" and "put the last nail in the coffin" for the China Brain argument. So far all I've shown is that we cannot know that Mr. Li doesn't have qualia, so intuitions to the contrary are irrelevant. But let us suppose for a moment that qualia are in fact properties of a fine-grained functional system within a physicalist conception of the world. If this is true (and we have no way of knowing it is not) then we would be right to expect that Mr. Li's having qualia would seem intuitively weird. In other words, if functionalism is true (contra the China Brain proponent) we have the tools to explain both why the China Brain's central premise (that Mr. Li would lack qualia because it is un-intuitive that he could have them) is so appealing and why it is false. It would be false since qualia would have to be functional properties, thus R-status Mr. Li would have to have them. And it

would be non-intuitive that he would have them—making P the intuitive view—because of the very nature of phenomenal (recognition) concepts and theoretical concepts, namely that they cannot translate into each others' terms. When we look at Mr. Li, we can only see him from a third-person theoretical perspective. This perspective does not match up well with the first-person direct knowledge we have of our own qualia. And since Mr. Li is so different from us physically (his brain is comprised of billions of Chinese people in a warehouse), it is non-intuitive to ascribe qualia to him by empathy, which is the method we normally use to ascribe mental states to each other.

Just think, we did ascribe qualia to Mr. Li when we didn't know his brain was made out of Chinese people—that is, when we thought he was just another one of us. My argument has been that we were right to do so.

References

- Block, N. (1978). Troubles with functionalism. *Minnesota Studies in The Philosophy of Science*, 9, 261–325.
- Braddon-Mitchell, D., & Jackson, F. (2007). *Philosophy of Mind and Cognition*. Malden: Blackwell Publishing.
- Dennett, D. (1988). Quining qualia. In A. Marcel and E. Bisiach (Eds.), *Consciousness in Modern Science*. Oxford: Oxford University.
- Levin, J. (2007). Functionalism. *The Stanford Encyclopedia of Philosophy*. Accessed at <http://plato.stanford.edu/archives/win2007/entries/functionalism/>.
- Loar, B. (1990). Phenomenal states. *Philosophical Perspectives*, 4, 81-108.
- Putnam, H. (1975). The meaning of 'meaning'. In K. Gunderson (Ed.), *Language, Mind, and Knowledge: Minnesota Studies in the Philosophy of Science*, Vol. VII, University of Minnesota Press, Minneapolis, pp. 131-93.