Mind & Life Podcast Transcript
Norm Farb - Meditation and the Brain

Original Air Date: October 7, 2022
Retrieved from: https://podcast.mindandlife.org/norm-farb/

Opening Quote – Norm Farb (00:00:03): *We are in some way responsible for how we perceive things and for how we react to things. And yes, we inherit all this conditioning—just like I inherit a certain physical body, I inherit my social body, myself, my model from my culture and my ancestors. But I'm also responsible for what I do with it, and I ought to be taught and empowered to affect it. And then I ought to use that power to make life better. And so when you find a place where the model isn't working, what can you do to leverage a little bit of plasticity in that model and update it? And contemplative practices are one access point by which to challenge and update these models.*

Intro – Wendy Hasenkamp (00:00:45): Welcome to Mind & Life. I'm Wendy Hasenkamp. My guest today is neuroscientist and contemplative researcher, Norm Farb. Norm is Associate Professor of Psychology at the University of Toronto where he studies the psychology of well-being, focusing on mental habits. He was one of the first people to study how meditation impacts our brains, and his broader work incorporates emotions, body sensations, and present moment awareness. He's been a key player in the field of contemplative science for more than a decade.

(00:01:20) I spoke with Norm last summer and our conversation, as usual, covers a lot of ground. We start with his winding path into contemplative neuroscience—integrating computer engineering, philosophy and psychology along the way. And then he shares a little about his initial study on meditation, which was quite seminal for the field, looking at how different ways of thinking about the self show up in the brain, and how meditation changes that. Then Norm goes over some basics of predictive models of mind (something we talk about on the show quite frequently) and he unpacks how we're always either reinforcing or updating our model of the world. Here we get into some interesting implications for political polarization, and he even talks about the phenomenon of "mansplaining." We discuss the default mode network in the brain and its role in cognition and maintaining these models of the world, and how meditation can help shift our habitual patterns. Then we get into interoception, or internal bodily sensations, and Norm shares how sensing the body might reduce conceptual processing in the brain. And we talk about how this plays into depression, and how we can get stuck protecting ourselves from difficult feelings. This takes us into a discussion about the self as the current model of the world, and the ethics of updating our self model in a responsible way. And we end with Norm's perspective on a number of brain-related topics, including how functions are localized versus distributed in the brain, and a really insightful commentary on the current state of brain research on meditation and how we should be thinking about that.

(00:03:08) There's actually lots more in here too. This might be an episode worth a second listen if you're really into these topics. I so appreciate the breadth and depth of Norm's work and also his ability

to distill these complex ideas down to what really matters. I'm reminded of a comment from a very well-established and brilliant neuroscience colleague, when they were working with Norm on a big paper, they said, "After listening to Norm explain it, I feel like I finally understand how brains work!" That's an experience I can also relate to, including after having this conversation.

(00:03:42) You can find a link to that paper I just mentioned in the show notes along with many other resources. And I'll just highlight a wonderful talk from Norm that nicely summarizes a lot of what we dig into in today's episode. This was his lecture from 2018 when he received the Cathy Kerr Award for Courageous and Compassionate Science. I highly recommend it—it's called *How to Choose Between Beautiful Stories*. OK I hope you enjoy this conversation as much as I did. It's my great pleasure to share with you, Norm Farb.

Wendy Hasenkamp (00:04:09): I'm so happy to be joined today by Norm Farb. Norm, thanks so much for being here.

Norm Farb (00:04:19): Yeah, thanks for having me, Wendy. It's a pleasure, and [I'm] a big fan of the podcast.

Wendy Hasenkamp (00:04:23): Oh, thank you. Well, we're so happy to have you. I always like to start with a little background and understanding a little bit of how people came to be doing the work that they're doing today. So I'm curious, for you, how you came into the contemplative space, and your interest in psychological and neuroscience research.

Norm Farb (00:04:44): Yeah, I guess my super spider-man origin story was I got bit by a radioactive nerd. *[laughter]* I think I always, growing up as a sort of introspective kid, and from an early age I was quite curious about how the mind worked and why it was that we had even awareness at all. And I actually went into computer engineering after high school, to try to get into AI. This would have been in the end of the '90s. But it was a ton of math and I was also learning about not living at home, and drinking and other substances, and it wasn't the right time to be learning both things. *[laughter]*

(00:05:25) So I ended up dropping out of computer engineering after a year and thinking, "Well, I'm still interested in this mind stuff, but I don't know if I can hack the math side of it. And so why don't I get more into the philosophy psychology sort of stream?" And so I ended up switching into arts program and doing philosophy and psychology. And I got caught up into the existentialist line of philosophy, which is all about the primacy of experience, and that really resonated with me. And at the same time, I was learning psych methods for how to maybe tractably test some of the ideas that come out of inquiring into your own first-person experience, which is sort of the beginning and end of existentialism (which was a little unsatisfying). And I also saw, from my studies, how it was an incomplete methodology because you could never actually move beyond personal experience to things like social obligation from a pure existentialist perspective.

(00:06:22) Anyway, so I kind of lucked into finding an interesting intersection that could still apply to my interest in the mind that let me ask big questions through philosophy, but then try to think about, how do you even test this? Is this testable through psychology? And by the time I was done my undergrad studies, I was advised (more strongly by the philosophy profs) that psychology was a better option career-wise. And I was also sort coming around to, yeah, I'd like to actually be testing things in the world. And after taking a gap year to work as a research assistant, I ended up going to grad school in an

affective neuroscience lab, where I could study some of the hardware, as it were, that instantiates a lot of our mental programming while still keeping an eye on things like emotions and well-being and the feeling of the self, and asking some of those questions.

(00:07:15) And it wasn't until my first year of my PhD... So I did a Master's year first, studying how the body responds to our personality styles and our motivation. But it wasn't until my PhD that I sort got into thinking about meditation at all. As I've said many times, I was the guy in the lab who did yoga when my supervisor, Adam Anderson, forged a friendship or collaboration with Zindel Segal, who's one of the developers of Mindfulness-Based Cognitive Therapy. And they decided that they were going to run a study on what happens to people when they do Mindfulness-Based Stress Reduction, the most popular of the manualized mindfulness training courses. And they were going to pay to do neuroimaging. I had just been doing physio recording so far, and I really wanted to learn FMRI. So this was a chance for "the guy in the lab who does yoga anyway" to maybe dip a toe into meditation research. And what I got out of it initially—the sell was like, we'll pay to do an MRI study if you learn how to analyze it. And that quickly became my entire PhD thesis and subsequent career. I'm still referring back to the results of those first studies in planning what I do now.

(00:08:30) And it wasn't really until, I would say, a couple years into that endeavor when I was starting to actually see that mindfulness training was leading to differences in brain activity in circumscribed contexts, that I thought, "Oh, I should probably maybe try meditating to see what's going on." And coming back to the kind of existentialist roots. And sometimes presenting brain data and people asking, "Well, what was it like in the scanner for the participants?" And I'd say, "I don't know. I don't know what... I'm just a benign skeptic. I'm not actually a practitioner." And going to some of the Mind & Life Summer Institutes was really formative for me. I think I went to one right as I was just beginning this work at the start of my PhD—way back in 2006 or '07. And there, Varela's idea of the scientist practitioner was, I would say, hammered into our skulls. It was the one thing that seemed like the organizers wanted us to come away with, was this idea that we should be practicing personally in some sort of contemplative practice, and that should be informing our science. And vice versa—our scientific findings should be informing how we conceptualize and how we're motivated to do contemplative practice.

(00:09:44) So I found that really rich, and it was a bit of that "best of both worlds" of the engineering hardware and neuroscience side for me, but bringing back some appeal to get into the existentialist or phenomenological perspective of what it's like to be you, and what happens when you pay attention to your experience. And so that marriage really worked for me, even though it meant making up all of our methods because there wasn't someone to follow necessarily. But before I knew it, I was halfway through my PhD and just starting to publish from this fMRI study and I was like, "That's my thesis." And your thesis is your brand, for better or for worse, unless you just walk away from it and start something new. It sets the trajectory for what you're going to do next when you're a graduate student. So I think it was really an amazing conflation of opportunities that got me into this field, and let me feel like I really found some kindred spirits here, and that the questions excited me then and continue to excite me now.

Wendy Hasenkamp (00:10:41): That's awesome, and it sounds like a really great fit between your interests coming into this, and then engaging with Francisco Varela's ideas and this whole space of putting together first- and third-person, because it sounds like that's kind of what you were interested in all along the way. So you were really one of the first people to try to do this from the neuroscience side. Your paper, I guess it was from 2007, is that the paper that you were saying became your thesis?

Norm Farb (00:11:10): Uh-huh.

Wendy Hasenkamp (00:11:11): That was a really seminal paper at the time, and I remember it being the first time that I had ever seen anyone publish and speak about the self with meditation, too. And it sounds like that was part of your interest all along. So do you want to talk a little bit about that paper, since you said it has informed so much of your future work?

Norm Farb (00:11:30): Yeah, I mean, it's still the distinction of that paper, which is the idea that we can shift into a mode of relating to our experience that's primarily about sensory input, and that is important because it's not what we normally do. And that might actually take some practice to shift out of a mode of conceptual elaboration—we called it "judgment" at the time. I think of it now as almost just like an output mode. That we're always trying to figure out how we're supposed to respond, and shifting from a place where you're trying to respond into actually taking in information, information that could change you. It was a radical kind of move. (So I love that idea. It's like a radical act. That's a Jon Kabat-Zinn phrase, but it really sticks with me as well.)

(00:12:15) That's still the game, right? And I think I'm really clear, for myself at least, that we're probably not going to discover a new thing the mind does that no other human has discovered. It's more of a question of translating and updating our concepts to include what's essentially a fundamental human capacity. And in some ways, to correct for an imbalance that we currently experience, where we spend our whole childhoods trying to occupy a mode of responding and acting like we know. And somewhere along the way, we lose track of how to switch back into the sensing mode because we never had to learn how to be in that mode. We started in that mode. Where we knew nothing and we were just trying to find patterns in our experience, so that we could scaffold our way up and understand the world and how to survive in it.

(00:13:11) And just last night, I was watching with the kids, Hook, the Robin Williams sort of... Peter Pan grows up and he forgets how to play. He's like this high tension lawyer and his marriage is in trouble, and he doesn't have time for his kids. And it was really, it keeps showing up. Every day, it shows up if you look for it. That the "hook" in that movie (other than Captain Hook's hook), is the idea that it's really, really important to find a way to still experience play, and still experience what it's like to be exploring and to not know, and to not always have to respond.

(00:13:44) So I would be fine never moving past that distinction as the main thing that motivates me. And it was really amazing seeing the response to the paper. I mean to be clear, there were definitely architects in the generation of researchers before me, like Sara Lazar doing stuff on longitudinal brain structure, or Richie Davidson studying long-term meditators. But the idea of just taking people off the street and making them do some meditation practice, and seeing whether an eight-week training effect was even publishable—we didn't know that at the time. But yeah, very quickly it set precedent, and really felt like it empowered or inspired, along with the other researchers working in parallel with me at the time, this kind of new generation of contemplative science, which has just been an awesome wave to ride.

Wendy Hasenkamp (00:14:32): I love how you were just speaking about, when we're very young and infants, and we come into the world and we're searching for patterns that will help us understand and navigate the world. And you've spoken and written so beautifully about predictive models of mind, and it's really helped me understand the relevance of those theories. So I was hoping we could dig into that a little bit, and explore the models that we hold, and then how we respond to sensory information

coming in, whether it matches that model or not... So I think this has a lot of really interesting implications. Maybe just to start, could you share a bit more about how we build this model of the world, and the role of sensory information?

Norm Farb (00:15:17): Yeah, I think the predictive coding model is de rigueur these days, and especially in systems neuroscience... The idea that one principle could explain how a lot of the brain works—which is just trying not to get surprised, and trying to reduce not having a good model for what's coming in from the world. You know, we just get nailed with sensory information and it's upsetting not to understand what's going on. And it's upsetting at a physiological, at a basic neural level. It creates a lot of activity and it's unclear how to resolve it.

(00:15:50) So what our brains seem to do is try to find regularities or patterns in sensory information, and we can then leverage taking a familiar set of inputs, and almost predict what goes with those inputs when we start getting enough experience, enough exposure to the same thing happening. So in the example of recognizing that there's a table in a room, it requires understanding that certain gradations of light hitting our retinas represent edges of objects. And that takes a long time to form. If you've been around young babies, they don't really know that there are objects for months. But it's not that they're doing nothing, they're just building a corpus of information to the point where they start to realize that there are certain familiar patterns within the sensory tableau.

(00:16:48) And over time, we build upon the inferences we make. So an inference would be like, "Oh, there is a table there." But before you can make a table there, you can say, "Oh, there's something darker there and something lighter around it." And then, "Oh, maybe there are actually edges here." And, "Oh, there are features, there are legs, there's a flat thing on top." And, "Oh, this is the kind of thing that turned out to be a table last time." And so you're not doing this through language, but by making these inferences on sensory information, we learn how to package up sensory inputs, and we automate that process. And so by the time you've seen your millionth view of a table or 10 millionth—I don't know exactly how many views and how to break up a moment in time, but—now it's just a table.

(00:17:32) And then we do the same thing for understanding feelings in our body and eventually we do the same things for understanding the broader world and how we fit in, and how we should communicate to others, is we just look for patterns and we get used to "that's the way things go." And we only pay attention really when something doesn't already match with our expectations. So it's sort of like expectations all the way down.

(00:17:55) And that's the predictive coding model of the brain—from just even understanding where an edge is, or what color is hitting the retina, all the way up to like, was that action just or unjust? This is all about matching what the incoming information is, compared to our model in similar contexts. And the contexts are also queued by their familiarity to past experiences. And to do a comparison between, well, what do we expect to get in this kind of situation, and what are we actually receiving? And to the extent that we're off, it draws our attention because we need to try to resolve that discrepancy so that we have an appropriate model for responding to the world again.

(00:18:35) We had a Mind & Life Think Tank in 2015 where we wrote... or, prior to 2015. In 2015, we published a paper talking about how there's two ways to stop being surprised by the world. One is to go address the causes of your surprise, to go try to fix things in the world or in your body. You feel something you shouldn't. You feel something in your stomach, you take an antacid because your stomach ought not to feel that way. And the other way, which is the kind of contemplative way, is to

think like, "Well, maybe this is the way my stomach feels sometimes." And not just say that in a kind of gaslighting way, to make yourself wrong for being upset about it, but to really be like, "Yeah, it's okay sometimes that it feels different." And then that also takes away the surprise and the upset and the need to do something.

(00:19:19) So that's sort of the predictive coding model in a nutshell. And most of the time what we talk about is self-regulation, in terms of like—I'll take some overt action to resolve the discrepancy. If a sound is hurting my ears, I will try to plug my ears, or move away from the sound, or yell at the person who's making the noise, or whatever it is, because it ought not to be this loud. It's not usually this loud, and I'm uncomfortable with it.

(00:19:47) And what you see as we move from children, who are just constantly asking, "What is this? Why? What is this? Why?" over and over again, to adults where it's like, "Let me tell you what this is." We try to move into this presentation as though our models are complete. And certain phenomena, like mansplaining, is the need to always have a model, and you need to prove that your model is good, even when someone else isn't really asking what your model is. *[laughter]* Like, "Let me tell you about the thing you're an expert in, because my models are already complete." And the idea that we could still be updating our models is very similar to having a growth mindset. That it's okay that your models are wrong. In fact, it's inevitable. And so when you're upset by your model being wrong, sometimes you don't need to beat the world back into shape, which is often not just the physical world, but our social world. You could actually just try to learn something, and learn that that's the way things go sometimes. And knowing whether to try to fix the world, or to update your models for the world, is the main game that we ought to be playing as human beings, rather than getting stuck in one mode or another.

Wendy Hasenkamp (00:20:58): Yeah. And I think we can see the kind of default tendency to confirm the model just writ large on the world stage. With political polarization, for example, if any contradictory information comes in about someone that you have put in a certain category, there's all sorts of ways of discounting that information. Is that another example of how, in a very social way, we can keep confirming the model, as opposed to being open to actually new information, and potentially updating the model of the world?

Norm Farb (00:21:33): Yeah. And I think it's a point of contention where one ought to push back against model-upsetting information versus consider updating the model. So the idea that if someone is coming to you and not actually discussing something in good faith, and they're actually just trying to undermine you, and there's no way that they're going to ever change their own models, that might not be really a great time to just start updating. So it's like, who do you trust to even engage with as a source of information to potentially update your models? But it can go too far where you say, "Well, I'm never going to update my models because those people aren't trustworthy." And then of course, you will never be able to appreciate their model of the world.

(00:22:16) So I think it's risky to lean too heavily either way. And we see this sort of conflict playing out right now as to, should we even try to understand this person's point of view or not, if they have a really different political attitude than I, for instance.

(00:22:31) – *musical interlude* –

Wendy Hasenkamp (00:22:31): So when it comes to the role of contemplative practice or meditation, you were saying before that maybe that shifts our tendency to update models, or pay more attention to the sensory information that's coming in?

Norm Farb (00:23:15): Yeah, I mean I think maybe it's a radical proposition, but I would argue that every time you engage in a true act of sensation, an authentic act of sensation, you are letting yourself be changed by the world. And every time you respond, you're trying to hold yourself together and make the world change.

(00:23:33) And the idea that there's a skillfulness in letting yourself be changed, and titrating how much sensation should I let in, and also being selective in what type of information... And sensation isn't just like the feeling of the wind on my skin. Sensation is also the language I sub-vocalize inside my head. That's also a sensory event. That's the thoughts I have, the feelings I have. So what do we actually want to attend to and not modify, and in doing so, let it change us? We're either holding sensation as the thing that's not supposed to change or we're holding the model as the thing that's not supposed to change. And we're always holding one of those two things. And so contemplative practices, one way to look at it I think is the art of playing with, well, what are you holding? And can you let go of holding one thing and hold the other? And the fascinating thing for us today is that it's really, really hard to let go of the model.

(00:24:33) And it's a cartoon distinction, but this idea of the default mode network in the brain—that tends to be most active when we're doing things that are most familiar to us, most habitual—shows us that the way we're wired is to move from a place of changing into a place of relying on models. It's sort of like, how built-in that is. Whether human beings could be raised in a way where the default network does not become the primary contextualizer of our environments and producer of our responses is sort of an open question that I don't think is going to be easy to answer. But what seems clear is that it's very, very easy for people to fall into reporting on the model, and holding the model as the thing that that's important. And it takes some practice to let go of that a little bit, and let yourself be changed. And possibly damage or knock up a few models in the way, and then have to integrate that and reconcile, "Okay, so now what do I believe, what do I know? Who am I?" Depending on what model is being impacted.

Wendy Hasenkamp (00:25:38): Yeah, I love that you brought up the default mode network, because I feel like that is a brain system that comes up a lot in contemplative space and also in cognitive science. I'm wondering if it's accepted, or generally thought now, that the default mode network is the neural representation of the "model of the world" that the brain is holding.

Norm Farb (00:26:03): Well, I think yes and no. At every level of neural representation, there's modeling happening, and things only get sent further up the chain if the neural level gets surprised, as part of the predictive coding model. But at the highest level, at the whole system level, it seems like it's the default mode network. I know from personal experience, if you get someone to just start a new task in the scanner that's easy, like, "Press one finger when the arrow is pointing left, press another finger when the arrow is pointing right," you get a lot of brain activity for the first 10, 20 seconds and then it kind of quiets down, and the default mode network spins up as it becomes familiar and routine. And as far as we can tell, even basic physiology—how your cortex responds to changes in need for consuming oxygen or other metabolic resources—that stuff all needs to be deeply automated to keep us alive, and it seems to have some brain stem stuff happening, but also feed directly into the back parts of the default mode network.

(00:27:08) So saying the huge swath of brain is just doing models is sort of slamming abstractions together. But at the same time, I think it's a useful heuristic, at least in comparing it to other brain networks, to other big constellations of brain regions. And it leads to some interesting implications. For instance, we know that self-referential thought is one of the few explicit tasks that you can do that will fire up the default mode network, whereas most other formal reasoning or effortful processes turn it off as you start to become engaged and do something that's a little challenging or less routine.

(00:27:49) So yeah, I think it's a nice foil for what meditation may be doing, in terms of either weakening the dominance of habit or creating communicative pathways into the default network that modify... So it's not necessarily about "tear it all down" but like, wouldn't it be nice to be able to update your models? So I think we're getting more nuanced than just default mode network = conditioning—and then, that's bad, let's get rid of all conditioning—into more nuanced accounts of well, what gets access to the network? And what can actually feed in and influence the network rather than just be outputs of it?

Wendy Hasenkamp (00:28:30): Yeah, I love this idea of the ability of these networks to incorporate more information and be updated. It feels like, you were speaking earlier about, it's not just that you dismiss new information and it's not just that you necessarily accept it without thinking about it, but there's a lot of discernment that happens in that space. And I think, as you said, that's really the whole game is when to do one or the other. And I'm wondering whether allowing more information in—and maybe through practice if that's what meditation is in fact doing, is allowing things to be updated more easily—that also makes space for that discernment.

Norm Farb (00:29:07): Yeah, I think there's a bunch of things that happen in early meditation practice. The first is just giving yourself a bit of space to notice things. And so that is really interfering with business as usual. And it was the first, it's the thin edge of the wedge for meditation in the West was the relaxation response. Like Herbert Benson being like, "Hey, you know what happens if you really focus on something for a while that's not just business as usual? Your whole body just kind of relaxes. Isn't that weird? It's probably good for you. Here's a book."

(00:29:41) So in the 1950s (this predates MBSR or anything like that), we have *The Relaxation Response*. And I really think the idea that it's not linear in contemplative practice is true, but there's a reason why certain practices are laid out in the way they are, across a variety of traditions. And there are certain, let's say trends, in terms of what has to happen first. And so by attending to sensory inputs in a somewhat... definitely intentional way and somewhat effortful way, you really are pulling resources away from the habit network—from the default mode network, if we're going to load it up as being where habits are instantiated and expressed.

(00:30:23) So that's a really important early step. We're not talking about any kind of higher level discernment, like a moral judgment—on purpose, because that gets really complicated. It's more like, can you even discern when you're sensing versus when you're back in the model mode? And can you get some skillfulness in efficiently moving back to sensing when you realize the model took over again? Because it's the dominant force. So that'd be really early practice, because over time the dominance of the default mode network will start to decrease as it's starved from resources by the act of attending to sensory information.

(00:31:07) And then things get really interesting, because now you have a little bit of space where you're not just elastic banding back into habit every moment. It's maybe every other moment, two out of three moments. But you have these little moments where you're not just sucked in back into the narrative gravity of what's familiar. And then in that space, now you can make some truer choices around what to attend to, in a way. Which of course, will be still motivated by your past experience. But you can do something different than the most obvious thing, in terms of where you attend. And now, there's a chance to notice things that you don't expect. And from there, you can scaffold on... You can develop a model for noticing things that you didn't expect. So you're still model building. And from there, you could notice not just physical sensations that are unexpected, or you're at least open to whatever's coming, but you could notice feelings, thoughts, interpretations, judgements, meanings, and work your way up from there.

(00:32:10) But the idea of discerning and selecting, like is this the kind of thing I'm trying to notice, and am I going to keep turning back to it, becomes a much... It becomes an everything game. You can pick what you want to "forage" for—is a term we're using more and more, as Zindel and I kind of workshop these ideas. You can start foraging for particular things. And we really like the term foraging because there's no guarantee you're going to find it and you don't even know exactly what it's going to look like. But you're out there, sort of like, "I wonder what's out here? I wonder if there's opportunities to notice playfulness, or kindness, or another way to solve this equation," or whatever it is. You can forage whatever you want. Though of course, people are influenced by the traditions they're trained in, as what the likely targets should be.

(00:32:57) – *musical interlude* –

Wendy Hasenkamp (00:33:18): I feel like a lot of your work has, since the beginning, focused on the body and sensations of the body, which is where a lot of meditation practice begins. This is a process that's often referred to as interoception in the research field. So, I'm wondering how you think about how that plays into all these things we've been speaking about.

Norm Farb (00:33:42): Yeah, I think it's still an open question whether interoception, or awareness of the body's internal state, which includes the basis for feelings or emotions, is privileged in a way above other sensory targets. Like, could you just do a visualization practice, or a listening practice? And clearly there are traditions that don't begin with the breath or the body. It's just the mindfulness tradition, those sutras start with the breath. And so I think on one hand the ability just to break out of the habit network—and get into a place of not defending your model, but defending the ability to take in information—that could be done with any sense modality.

(00:34:29) But then, as I was trying to point out before, there is something very special about interoception, because that's where we can see our conditioned emotional responses. We can see the arising of the feeling tone that is thought to accompany every sense percept. Every moment of sensation, every little bit of sensation has a little bit of like, "Ooh, this looks good," or like, "Oh, I don't know about this." And that's deep conditioning that is relevant for our sense of well-being, and purpose and safety in the world, in a way that noticing, "Oh, this is what it's like when the table ends and air begins," doesn't have that same type of conditioning that is bent towards self-relevance. Or the feelings we get are all about, what should the organism do in response to the sensation? So there's something really rich there.

(00:35:29) But we've done studies—a while ago with a graduate student, Thomas Anderson—where we give people three different types of meditations: a breath focus, a visualization, and an auditory sort of mantra practice, and ask them which one they think they'll prefer. And then they do each one for five minutes. And then we ask them, "Okay, how was it? Which one do you prefer now?" And half of everyone switches, it doesn't matter where they start. So people don't even really necessarily know what their primary access point is. (And this is only after five minutes, so maybe things switch in a more predictable fashion after you do it for a few weeks or months.) But these are just access points to getting into this kind of sensory foraging state, where you really are shifting yourself into an input state.

(00:36:12) But the types of insights when you get to what I call like "stage two" meditation, like now you can get into the sense state. I'm not going to give you 10 stages or anything, but let's say there's at least two stages—like a quieting down of habit, and then a looking around. The types of things that you can notice when you look around are probably really dependent on what you're paying attention to. And to the extent that you're doing meditation to understand your own conditioned aversion, or craving, or what's running behind the scenes emotionally, there's something really, really important about the body. It would be very hard to... You'd be seeing it at least in one order removed to understand how your vision is literally colored by some kind of deeper conditioning, or your hearing. If anything, the things you might notice then that are self-relevant are your impatience or attitudes, so like the internal feeling state that's still arising even though you're ostensibly focusing out in the world.

(00:37:05) And we have some pretty cutting edge neuroimaging findings that suggest that there really are these kind of two stages that everyone can... When they get into sensory absorption, particularly of body sensation (this particular study I'm referencing is breath monitoring), there is a massive deactivation of a lot of the conceptual elaboration parts of our brain, even when it's compared to a very closely matched visual target. So comparing vision to breath monitoring that is not—I've just got to run some more control analyses, but as best as I can tell at this point, it's not—just predicated on breathing slower or deeper, or something like that, that's just affecting brain chemistry. Turning inside is like a bit of a sensory deprivation chamber, where feelings and other things might be seen in a higher contrast.

(00:38:00) In that study, actually we don't see activation of interoceptive regions. What we just see is a lot of the other parts of the brain—that are probably also contributing noise, and contributing to awareness—are quieting down. Because of course, the body must always be representing itself to keep your heart and your breathing and your blood pressure and all those things within safe boundaries. And so the idea of getting out of your own way to notice some of that stuff seems like it might be playing out at a neural systems level—that to pay attention internally is actually a more radical act of quieting down than much else.

(00:38:40) And then what happens is that people with higher expertise, or comfort, or trust in body sensations—we use Wolf Mehling's scale, the MAIA, the Multidimensional Assessment of Interoceptive Awareness, to look at people's comfort and trust in body sensations—higher scores on the MAIA suggest a sparing of a total cortex shutdown and you keep some of the salience network, the anterior cingulate, that sort of monitoring discerning faculty alive even while a lot of other things are shutting down. So it's getting into this kind of quieter space and then in this sort of low-energy mode, still being able to discern things. That's something that's really, I think, somewhat unique about breath monitoring, at least compared to visual attention.

(00:39:29) So to answer your question, I think there is something really cool about sensory attention in general, but there may be something privileged about internal awareness for some of the reasons I enumerated before. But it's still a work in progress to show that that's really a valid distinction, I'd say.

Wendy Hasenkamp (00:39:46): Yeah. I'm wondering too—I know you've done a lot of work around depression. And I feel like there's a complex relationship to internal bodily states, and the ability to sense one's own body in depression. Just wondering if you could speak to that a little bit. Has that been explored? Is there a difference...?

Norm Farb (00:40:10): In depression vulnerability?

Wendy Hasenkamp (00:40:11): Yeah, in depressive states, with interoception.

Norm Farb (00:40:15): Well, I mean it's a huge field, of course. It's the leading, I think it's just getting to being the leading cause of disability around the world, is depression and depression symptom burden. So a huge field. We know from resting-state MRI that depressed people tend to show a sort of out of control prefrontal pattern of activity. This is Yvette Sheline's work. And so you end up seeing networks that are supposed to kind of push against each other, all activating at the same time. So it's like, the default mode network is firing at the same time as the planning, control, executive network—so that you can bear down and think even harder about "why your life sucks" kind of thing. And then you're noticing how bad it is at the same time. So the salience network (to complete the trio of major prefrontal networks) are all being co-activated. And this created this dorsal nexus hypothesis of depression. This place where these three networks should all have well-ordered fences, in a way, the storm came through, ripped those fences out. And now your habit is to attend to your thoughts about your depression, and then to notice how bad it is, in a constant cycle.

(00:41:28) And I think that's the dominant view on depression, is that it's a problem of model implementation, of model defense. It's a prefrontal problem. And this whole other stream of my work that does tie into interoception was this discovery that while all this is going on—in the same way that paying attention to sensations might starve the default network, might starve these prefrontal networks of metabolic resources—in depression, and even in the local response to seeing something sad, even if you're not particularly depressed, in pushing so much effort into understanding and modeling and figuring out, "Well what does this mean for me," and all this stuff, we are actually starving resources from our sensory networks. And from the studies that I've been a part of, it's actually how much the sensory or sensorimotor parts of the brain are turning off in response to a negative emotion that's really predictive of long-term vulnerability to depression, or concurrent depressive symptoms, or the scars of past episodes of depression. Everyone is turning on the front of their brain to understand and unpack their sadness. But not everyone is doing that to the exclusion of still letting new information in to the same extent, even though we all turn away a little bit from our senses when we feel bad.

(00:42:48) And so it's the difference in the "turning away" that is actually most health relevant, at least in the context of depression vulnerability. It's not the fact that we ruminate and get a little self-evaluative, and delve maybe unintentionally into related memories, and things like that. That actually seems like something that both healthy and unhealthy people do. But to do it to the exclusion of letting new information in, is what's pathological.

(00:43:18) And so we've just published the first neuroimaging paper from a really big clinical trial in remitted depression where we showed that it's the somatosensory cortex (which is sort of the hairband

on top of your head that has a map of your body), and a bit of the motor cortex (which is an output map to go along with the input map), turning off that is most related to past episodes of depression, residual symptoms in people who are now remitted from those past episodes, and future vulnerability over the next two years. Not so much what the front of the brain is doing.

(00:43:54) And very few people are talking about non "prefrontal cortex affecting the amygdala," with the emotional salience detector accounts of depression. It's all very front-of-the-brain-centric. But there are starting to be large data-driven studies of depression that are now starting to point out that it's some of these other connections—maybe to sensory regions, maybe to motor regions—that are actually the best big-brain-system indicators of depression or depression vulnerability. It's not just about the threat detector firing too much, or something like that. So things are maybe slowly shifting.

Wendy Hasenkamp (00:44:31): Yeah. That makes a lot of sense. And also then points to what we were speaking about earlier... If the role of meditation is to maybe try to help learn to integrate, or pay more attention to sensory information, that's a kind of whole other lens on why certain kinds of meditation might be helpful for depression—which has been shown a lot, and is actually one of the first realms that you got into this field about, right?

Norm Farb (00:44:55): Yeah, for sure. And it's sort of like an understandable but ironic consequence of people's attempts to deal with their negative emotions is to say, "Well, I'm going to turn away from feeling, because the feeling feels bad, and I'll just try to understand it, and understand the conditions that it arises in, and what I have to do about it—because all of that planning and interpreting is safer for me right now than feeling the bad feeling that led me to get into all this planning and elaborating."

(00:45:19) And so the irony is, by doing that, what you've essentially done is locked yourself in a room with your depression thoughts, because now you're staying away from feeling. And feeling is—as I said before, sensation and the feelings that arise out of that, is literally how the world changes you. So, you don't feel good. So you shut the doors and you stay in the room just with the knowledge that you don't feel good. But then the question is like, "Well, what's the exit strategy?" You're like, "Well, I'll feel better again." How will you know? You won't know because you're afraid that it's just going to be more depression at the door.

(00:45:54) And so it is unsafe to open the door again. It is literally unsafe, but in a delicious way, in that you may no longer be safe in knowing that you're depressed. That's really threatening to your ability of knowing what's going on, and having a model. Your model might get kicked down by the fact that, like what does it mean about you if you're identified as a depressed person, and you feel great one day? You have to rebuild your whole identity.

(00:46:19) So I think it's understandable that people get stuck by trying to protect themselves from feeling. But if that's your problem, and I'm not saying all... sometimes people feel depressed because something really bad is happening in the world, and just feeling it more isn't going to fix it. But for people who are depressed because they're locked in this kind of echo chamber of their own making to avoid feeling more—and that tends to be the case, the more often you become depressed, the less it tends to be related to a major life catastrophe—for those people, learning how to let sensation in, and letting it be not safe, and letting their ideas about themselves, and the world, and the future get knocked up a little by sensation that doesn't care that that's what the model is, and might disrupt that model, is beautiful and necessary.

(00:47:04) And so yeah, contemplative practice is, ideally I think, learning how to do that safely. And they should be adapted and titrated for clinical situations where the level of conditioning in response to sensation may be really different, and much more intense than the average practitioner has to deal with.

Wendy Hasenkamp (00:47:24): You just touched on the kind of central piece of our core identity. And in that case of depression, it's really difficult to have that challenged, and maybe have to rebuild that. I'm wondering, your original study was also related to the self, and the way that we think about the self—either in the moment or in this kind of narrative storied way. So I'm wondering just how you've come to think about the self in this whole space, with contemplative work and predictive models and habit systems, and all of that.

Norm Farb (00:48:01): I've come to think of the self as the current model—for life, universe, and everything—that we're walking around with. And it's continuously being constructed, in response to both the context that we physically inhabit, our environment, and in response to the imagined context, the things we're remembering or expecting or imagining. And it is not reasonable to think that we would be better off without a model. So the idea of that we're trying to get rid of the self completely is I think a misunderstanding. And even the terms like *anatta*—no self, the self is an illusion—is not the same thing as saying that we would be better without it. That's actually a near miss, a near enemy, I think, of right understanding of this whole enterprise.

(00:48:55) The questions that are maybe more useful or skillful to ask are, what are the dominant types of patterns in perceiving and responding that characterize this model right now, the self right now? If I could, would I choose them for myself? (Because they've already been chosen for me.) And then if the answer is yes, it's like the Nietzsche's sacred yes, like right on! Do it over and over again, forever!

(00:49:21) But when you find a place where the model isn't working—understanding that the model's job is to perpetuate itself—what can you do to leverage a little bit of plasticity in that model, and update it? And so, it is probably the fundamental question we're always, again, grappling with is, is the model good right now and should I let it ride? Or should I challenge it in some way? And contemplative practices are one access point, of perhaps many, by which to challenge and update these models.

(00:49:55) So I don't know if everyone walks around thinking that they're responsible for the health of their big model, their self model, in the same way that we might think I'm responsible for how flabby my abdomen is getting. I think there was a bit of a revolution in the late '70s and '80s where we started saying, "Hey, your physical health actually matters a lot. And there's things you can do about it, like you as an individual, can do about it. So we all need to start being a little bit more responsible." And then gyms and yoga studios and nutrition and all that stuff... If you woke up one day in the 1970s, none of that stuff was there. I think it's hard for us to imagine. I mean it's not that no one understood fitness was important, but the idea that everyone has an individual duty to curate their physical fitness to the best of their abilities, I don't think was really endemic, even 50 years ago.

(00:50:50) And the promise of the modern contemplative movement is, what if people also thought they were responsible for their conceptual models, or their big model, like their self? And the only way you can hold people accountable or responsible for something is if they have the power to affect it. And so the reason why the contemplative movement has a chance to raise this level of responsibility is that it offers to empower people to update their models.

(00:51:22) And if now you know that you could update your model and it would have a big impact on your well-being, or the well-being of others around you, and you have the tools like the technologies of self (to borrow a term from Michel Foucault)—we have these technologies of self where we can update our models. So it does make a difference, and you can make a difference. Somewhere along the way it becomes like, well, you should make a difference then. And then that's the idea of us sort of getting to another level of maturity, and realizing that we are in some way responsible for how we perceive things, and for how we react to things. And yes, we inherit all this conditioning—just like I inherit a certain physical body from my ancestors, I inherit my social body, my self, my model from my culture and my ancestors. But I'm also responsible for what I do with it. And I ought to be taught and empowered to affect it, and then I ought to use that power to make life better. And then who you're trying to make it better for is, again, a contentious topic. And whether there's a natural slide into caring about others once you are good at taking care of yourself, I think is still an open question.

(00:52:33) – *musical interlude* –

Wendy Hasenkamp (00:52:56): You've been in this field pretty much as long as anyone, thinking about the brain and meditation. And there I think have been a lot of shifts within neuroscience and cognitive science, moving from thinking about individual brain regions as having certain functions, to a much more kind of distributed processing model. But I think that understanding hasn't necessarily filtered into the media, and representations in the public understanding of how brains work. So I feel like... Just curious how you are thinking about, or how you tend to describe localization of functions in the brain these days. Like you mentioned, people are still talking a lot about prefrontal and amygdala relationships in depression. So yeah, I'm just wondering how you're thinking about that.

Norm Farb (00:53:48): Yeah, I think the more universal the thing is, the safer it is to localize it. And by universal, I mean like if you grew up in the jungle, you would still be able to do it. It probably has a pretty solid set of regions responsible. So things that are safe to localize are like, where different sensory inputs first hit the brain—like primary representation cortices. Yes, of course, there's going to be the odd mutation or accident where things have to restructure, but everyone gets vision in the back of the occipital lobe in the back of the brain. Everyone smells things through little bulbs hanging down, way up where a COVID swab goes. *[laughter]* So primary representation of sensory afferents—so signals coming from the world, or from our bodies, and into the brain—those all hit in exactly the same places. Similarly, the things that keep you alive, and keep all mammals alive, in the brain stem and those nuclei, tend to be pretty safe. And the same is also true on the output side. So the motor map that matches up with the somatosensory map, in terms of how you send signals back out into the spinal cord—given that we all have roughly the same kind of shape, bodies, and skeletons, and number of limbs and things like that—that stuff is really safe to localize. It's really safe.

(00:55:04) Other things that are fairly universal are being able to encode and retrieve memories. So the machinery for doing that, like in the hippocampus. Or the ability to detect if something motivationally relevant is in the field of your experience, like the amygdala—and the salience network, but the amygdala as an early warning system... Those seem fairly safe.

(00:55:27) And then as soon as you have to learn something, as soon as you have to start doing model-building that's scaffolded on top of that, all bets are off. And so can you, in an individual, figure out what their representational pattern is that distinguishes looking out in the world versus in at your breath? Totally. And we have some other stuff, and other researchers have done machine learning, to prove the information is there in the brain. But it's not in the same place for every person.

(00:55:53) But there are still some generalities—and that's led into this golden age of network neuroimaging—to say that there are parts of the brain that tend to talk more with each other because they're sensible scaffoldings on the sensory representations. So like default mode network—probably because it starts off with habits of homeostasis, of physiology, and then sort of scaffolded itself up to more abstract habits—occupies very consistent territory. The salience network, because it responds directly to the proximal alerting sensory signals, tends to have the same kind of general territory. The higher cognition, planning, and goal, "executive" sort of network tends to occupy the same kind of general territory. We know that these networks must be in play, in a way, because we need to be alerted to things. We need to know whether we can rest on habit or develop a new intentional plan. But they won't be the same for each person.

(00:56:55) And I think actually a really big paper came out this week that I haven't read in depth that actually documents the variability in brain—not just structural morphology but functional morphology. Like activation that shows that there's a huge variety in exactly where these networks will be. And that attempts to say, "Oh, I found the region that does X"—if it's not a completely universal type of experience like recognizing color from the visual field or something (and even that's not totally universal, but mostly universal)—that's sort of fallen out of favor. It's more like, can we see these broad trends in these associative or learned networks? And can we in particular see the trade-off between which of these networks tends to be dominant at any time, without worrying about exactly where its boundaries are? That's become a lot more popular. And trying to match up particular ways of relating to experience with networks is somewhat reminiscent of the sense/judge, narrative/experiential self distinction, but taken to a whole brain level, and not just looking at this one distinction.

Wendy Hasenkamp (00:58:05): So that sheds light, maybe, on what we know about meditation and its effects on the brain, and whether those are consistent or not. So what's your sense of that space right now? I know there was recently a big paper that was not able to replicate any consistent structural changes after an MBSR course. But then of course there's always papers that do find changes. So how do you make sense of that?

Norm Farb (00:58:29): There's a lot of things going on. Anytime you experience something that's new or different, your brain has changed. And it's not useful to say, "I need to wait to see a brain scan," to prove to me that that's changed. Every moment of experiencing something is either strengthening or weakening trillions of connections across the brain. So the idea that there's a certain threshold and then neuroplasticity happens, is just a false idea.

(00:59:00) So, of course the brain is changing after every meditation program. And of course the brain is changing after every banana you eat, and every time you stub your toe, and every time you cry at a really moving piece of art. And every time you just fall back into habit, the brain is also changing.

(00:59:19) So, studies saying the brain is or is not changing in response to X, I think, are not very interesting. More interesting I think, or more on point, is the idea that there's been quite a few papers showing that there might be regular players in what's changing with meditation practice. And I think if we were able to characterize beforehand exactly which parts of the brain are reacting in this particular context, and tested that context after meditation, we would probably be able to see some systemic changes on the functional level, in terms of brain metabolism, and on the structural level.

(00:59:59) But at the same time, I gotta say, seeing a different paper come out every year that talked about a different set of brain regions in response to an eight-week intervention, to me sort of reeked of false positive psychology. And I would say that in my own lab, I've found many little bits of statistically significant brain changes in different places sometimes, and just not felt confident enough to try to report them. It might just be that in our zeal to discover things that the consensus standards for saying we found something that's unlikely to occur by chance—those standards are too low, is probably true.

(01:00:35) But that's a thorny issue, because you would have to restructure the entire incentive system in academia and beyond, to reward people for not finding things. And so, you pay people to find things and you reject them for not finding things. It's not a surprise that everyone—and not just the researchers but also the editors, and then the public and the media, and agencies who consume the proceeds of science—are all biased towards only caring about the things you find.

(01:01:07) And so, if there's any one paper, or a couple of papers that say, "Oh, this is the one part of their brain that changes with meditation," is it likely that they've outstripped their reach and said that this generalizes to all people, but it might have just been true of the couple of people in the sample who showed a really big effect? Yeah, that's probably true. Does it mean that meditation wasn't changing the brain? No. It meant that that was the best way, that was the best model for characterizing what happened to that group, in that study.

(01:01:32) And as someone who's led a handful at least of MBSR, full eight-week courses myself, and have been a participant observer and a participant in MBSR and in mindfulness-based cognitive therapy, and some other courses, every program works different. And if you try to catalog what people are getting out of a program—I did a study with older adults because I was a postdoc at a geriatric hospital for three years before getting a faculty position—you ask 10 people what the course was about and what they got out of it, that you might get 10 different answers. And if you had a profile for that type of brain change, and you matched it up with the answers and had a real rich neurophenomenology, and also a much more extensive corpus of information around what each type of change looks like in the brain, you might actually be able to match it up. And maybe that'll happen in the next 10 years. We'll be like, "Here's five common outcomes of meditation. The models fit way better when we divide the group effects into five substrata and match them to people's experience, than if we just assume it's one-size-fits-all."

(01:02:36) So science in general has a problem in rewarding when you find things, and also rewarding when you find the universal thing. And if one person is using meditation to fall asleep at night, and someone else is using it to renegotiate their relationship with their estranged son, and someone else is using it to find deep metaphysical insights, and someone else is using it to appreciate the beauty of the outdoors more—I think it would be weird to predict that they would all experience the same changes through their practice. So I don't think it's... I'm just not that bothered by these kinds of things. I think they're sort of like, it's like gotcha journalism. Like, "Oh, you said you did this thing on the 17th, but it turns out you took the trip on the 18th!" And it's just like, okay, but who cares?

(01:03:17) So I'm not too worried about that. I think it's more important to try to understand why one size doesn't fit all. And then, also even ask, does it matter which part of the brain changes if someone is reporting that they've had this big experience? Is that the most important question to ask? I think a lot of people are more interested in, at the time someone shows up to you and they're suffering, how you could figure out what the right type of training or what type of reconfiguration that would actually help this person would be?

(01:03:55) And I think maybe that's the promise of studying these big changes across groups is—now if we know that that's the process, I wonder if we could go back to baseline with a new group of people, and then say, "Hey, you're deficient in this thing that would be corrected by a process that we can reliably produce." And that's like the holy grail for clinical science in general. We're still really bad at matching people's baseline conditions to appropriate treatments.

Wendy Hasenkamp (01:04:21): Yeah, that was really well said. Thank you for all of that. I wish you could speak directly to the media on a lot of those points. *[laughter]* Well, this has really been wonderful. Thank you so much, Norm. I'm so glad that you've followed your initial passions—you've given so much to this field. So yeah, thank you for all of your work, and thanks for taking the time today.

Norm Farb (01:04:44): Thanks for having me. It was really fun to chat.

Outro – Wendy Hasenkamp (01:04:51): *This episode was edited and produced by me and Phil Walker, and music on the show is from Blue Dot Sessions and Universal. Show notes and resources for this and other episodes can be found at podcast.mindandlife.org. If you enjoyed this episode, please rate and review us on Apple Podcasts, and share it with a friend. And if something in this conversation sparked insight for you, let us know. You can send an email or a voice memo to podcast@mineandlife.org.*

Mind & Life is a production of the Mind & Life Institute. Visit us at mindandlife.org, where you can learn more about how we bridge science and contemplative wisdom to foster insight and inspire action towards flourishing. If you value these conversations, please consider supporting the show. You can make a donation at mindandlife.org, under Support. Any amount is so appreciated, and it really helps us create this show. Thank you for listening.