

17 Emptiness and Fullness

To be filled, you must be empty. [#654]

Many of our meditators describe a phenomenal quality of emptiness that is paradoxically rich and full. This kind of emptiness is a global quality of openness and connectedness, like a lucid void or a clear space that is nonetheless dense or even vibrant. On the other hand, some of our participants protested against the (originally Buddhist) notion of “emptiness” as a descriptor for their experiences, possibly because they did not like associating it with other phenomenological concepts such as the experience of a conscious “vacuum” or “void” used in item #13 of our survey. This may have to do with problems of translation, a conflict between metaphysical and phenomenological readings of the original notion, or both. For now, let us stay with the phenomenology itself, as reported by real-life practitioners, and look at a series of examples:

1585 Infinite spaciousness, stillness, a void that is full and vibrant [. . .]

1662 [. . .] There was a sense of quiet awe and contentment and experiencing “emptiness” (as the closest “canonical” term, i.e., *sunyata*, even though there wasn’t a feeling of emptiness, but one of aliveness and infinite interconnectedness). [. . .]

1788 [. . .] It is an emptiness that does not mean the absence of something, it is not a lack of something, but it is really elementary, very real and clear, extremely alive, and the objects chair, table, etc. still seem like chair and table, but not as real, as if they were only dream objects. They recede into the background as meaningless and a very present emptiness comes into the foreground, but it is more than the space between the objects. The living emptiness then permeates everything, concepts like here and now no longer exist. [. . .]

3173 [. . .] the experience of an absolute emptiness comparable to deep sleep. I basically experience this state only once it has stopped. Then I know that it was just there. But in this state itself there is nothing. After that there are first

determinations like silence or gentleness (while still in the meditation) and then good mood or feeling refreshed (after the meditation). [. .]

3223 It felt so empty. As if for a moment nothing was there and you were just present with yourself, undistracted.

At the same time, emptiness is often explicitly described as coemerging with the phenomenology of “fullness”:

178 [. .] Feelings of emptiness and fullness at the same time. With feelings of calmness, peacefulness, and ease.

654 I have had feelings of complete fullness. In order for that fullness to fill, I was first in a state of emptiness. To be filled, you must be empty. The more empty, the more this divine experience can fill you.

1837 The feeling of pure awareness came from meditating. This oneness from my environment and my environment with me felt like fullness and emptiness all at once. [. .]

2537 [. .] there were moments and pauses between breaths of completely filled emptiness, in vastness, openness, and fusion. [. .]

2754 It is not easy for me to describe this state. “Presence,” “silence,” “fullness,” “emptiness” are terms that encircle it. A having-arrived, no longer becoming, timeless. Peaceful, eternal. Depth, fulfilment, insight.

3440 There is a point when knowing, knowing just stops. All is same. Just there. Not real. What is. Empty and full. If peace is defined by the emptiness of movement on each side of peace.

3601 [. .] It reveals itself as pure presence, no thoughts, no physical sensations. A feeling of fullness and emptiness at the same time. [. .]

Emptiness, Meaning-Making, and Mortality Denial

Emptiness is the nature of mind, clarity is its characteristic, and [their] unification is its essence. It is free from all extremes of elaborations, such as good and bad; arising, ceasing, and abiding; existence and nonexistence; and permanence and impermanence. It transcends speech or thought and is beyond identification. And yet there is something to be experienced: it is sharp, naked clarity, which is of the nature of bliss, clarity, and nonconceptuality.

—Wangchuk Dorje (The Ninth Karmapa; 1556–1603),
Dispelling the Darkness of Ignorance (II, 8)

It is a fact, Simmias, that those who go about philosophizing correctly are in training for death, and that to them of all men death is least alarming.

—Plato (428/427 or 424/423–348/347 BC), *Phaedo* (67e)

As soon as your breath ceases, what is known as the clear light of the first intermediate state [. . .] will manifest. [. . .] At the same time, a naked awareness that is neither inside nor outside of you will dawn, clear and empty, without center or perimeter. At that moment, recognize this instant presence as your own nature and relax in that state.

—Padmasambhāva (ca. eighth–ninth centuries CE) and Karma Lingpa (1326–1386),
Bardo Thödol (The Tibetan Book of the Dead)

“Emptiness” (*suññatā* in Pali; *śūnyatā* in Sanskrit) is one of the most important concepts in Buddhist philosophy—perhaps the most important of all. It has been discussed, interpreted, and reinterpreted by scholars and practitioners for well over 2,000 years. Many of the participants in our study may have known the concept of “emptiness”; they may have heard or read about it and may have had a personal and intuitive understanding of it. If so, this will have colored many of the reports given here. We first need to distinguish between a metaphysical and a phenomenological reading of this special concept—which, again, may well be one of the most interesting concepts ever developed in the history of human philosophy. As always, I am not at all interested in the metaphysics, only in the fine structure of consciousness itself. From a metaphysical perspective, “emptiness” means that all phenomena lack substantiality or an intrinsic nature of their own—they are not “ontologically self-subsistent” and they have no “essence,” as a Western philosopher might say. But let us set all of this aside for now.

One of our new phenomenological tools is the concept of “epistemic openness.” You may recall from chapter 4 that epistemic openness involves a distinct phenomenal quality—namely, a specific form of openness to the world. Epistemic openness is openness related to knowledge, to an inner space defined by the *possibility* of acquiring knowledge. The notion of epistemic openness offers a modern reinterpretation of the ancient Buddhist notion of “emptiness.” From a phenomenological perspective, full-absorption states—in which (as we have already seen) pure awareness remains as the only reportable phenomenal character—are a prime candidate for generating experiences of emptiness. Here, epistemic openness per se is all there is. But there are many other situations in which emptiness arises. For example, in the context of mindful perception, “seeing the empty nature of phenomena” (chapter 9) can involve a specific form of conscious experience without the slightest trace of conceptual overlay: The distinct and crystal-clear phenomenology of seeing and perceiving out of timeless silence. In this state, things have no *conceptual* essence, in the sense that they are not named, categorized, or judged at all. Phenomenologically, all forms of cognitive penetration have come to rest. As we saw in chapter 9, on suchness, there is a specific, positive sense in which perceptual states are “meaningless” whenever their emptiness is phenomenally experienced. They are not even “things”; they are not reified

as mind-independent entities; and therefore their content is epistemically open, such that a multitude of conceptual interpretations and perspectives are possible.

Importantly, in global minimal phenomenal experience (MPE) modes, the lack of conceptual overlay also applies, for example, to the implicit assumption that appearances are objects, juxtaposed with and consciously experienced by a knowing self. The overlay of subject/object structure can itself disappear, creating a deeper form of openness. As we will see in chapter 28, “Transparency, Translucency, and Virtuality,” a strong experience of epistemic openness can even lead to suspension of the distinction between existing and nonexisting entities. Using slightly more contemporary terms, I will call this the “phenomenology of virtuality.”

If, as I propose, we interpret “emptiness” as referring to the conscious experience of epistemic openness and virtuality, then this helps us gain a new perspective in a multitude of phenomenological contexts. From a third-person perspective, we can now describe certain aspects more clearly. But what will happen if a meditator continuously brings her own theoretical perspective into her practice, and if in her practice she is actually *motivated* by some belief system and the goal states that it defines—including, for example, particular metaphysical assumptions about what emptiness “really is”? Will third-person expectations shape first-person experience?

“Emptiness” is the perfect illustration of what in chapter 2 I called the “problem of theory contamination”. The phenomenology of meditation is special, in that theoretical assumptions and belief systems may have colored the reports that we received in a particularly strong way. What is more, we all *embody* our beliefs and our deeply ingrained cultural background assumptions—namely, via the two-way window of the self-model created in our brain, which causally connects our inner life as “knowing selves” to the social reality around us.¹ The biological/computational process that underlies the creation not only of our low-level bodily self-model, but also of the brain’s more abstract inner image of a thinking, knowing, and socially situated self, of course will directly influence conscious experience itself. Our best current theories show not only that there is a deep continuity of mind and life but also that life, embodied self-consciousness, and cultural context form a nested hierarchy. Our minds are *nested* minds, and by functionally appropriating theories via our self-model, we can “make beliefs our own”—changing our very brains, and changing the way we see reality.

There are a few chapters in this book that are longer than the others, and there is always a reason for that. This chapter is one of them; chapters 24, 26, 27, 28, and 34 are other examples. Here, I will take a closer look at the methodological problem of theory contamination by highlighting two of its deeper roots. The first root of theory contamination is the human need for meaning-making, and the second one relates

to the constant temptation to deny our own mortality. I am fully aware that this will be a difficult, perhaps unsettling, chapter for many of my readers—but I am equally convinced that it will help us confront meaning-making and mortality denial constructively, to get the full picture and find out what the deeper value of meditation as an epistemic practice (as discussed in chapter 16) really is. As is the case with any mixed plate of hors d'oeuvres, there will be some that you like and others that you would rather leave untouched. If you want to give this one a try, let's take a deep breath and begin by slowly and carefully approaching the first topic: the process of so-called meaning-making.

It must have been around 1980 when, in Chalet Tanneg in the Swiss village of Gstaad, I had the chance to talk to Jiddu Krishnamurti all by myself.² For the previous four years or so, I had practiced yoga, breathing exercises, and meditation systematically and regularly. I told K about all the positive effects that this had had in my life, on many levels at once. Then he asked me a question that tore into my heart like a dagger—one of the deadliest, meanest questions I have ever been asked. It was a question I had never thought of, a question that at first hurt me but that also proved to be extremely helpful and reverberated through my life for many years to come: “Can you find out—not intellectually, but in your own direct experience—which of all these positive effects came from the actual meditation practice itself and which came from the fact that you *found something*?”

Many of us try to *find something*: something that allows us to live in an insane world without going insane ourselves. Something that creates coherence. Something truly sustainable that functions as a normative metacontext, some overarching theme or highest-level goal. Something that reenchants our world, a reliable long-term perspective. The process of meaning-making is this attempt to *find something*; it is the search for coherence, for something that holds steady in a world where everything else is impermanent. My first point is that, for many committed practitioners, the core motivation allowing them to sustain a regular practice over many years may consist to a considerable degree in the fact that they have *found something*, that they have adopted a certain belief system (or joined a group or spiritual movement, identified with a certain lineage or teacher, etc.), and not solely in the intrinsic force of the pure-awareness experience itself. This belief system, the sense of community, the acceptance of an authority, and the ongoing project of meaning-making may be among the major sources of motivation.

Previously, I said that the first motivational root of theory contamination is the human need for meaning-making, and that the second one relates to the constant temptation to deny our own mortality. Both these motivational roots influence our behavior on a subpersonal level, in the brain itself, but they also naturally lead to an

intellectual interest in theories of consciousness, spiritual teachings about pure awareness itself, and the like. They therefore will influence the experience, as well as the way in which it is described later. Many of our participants may be familiar with the terminology employed by certain traditional and prescientific theories of consciousness, particularly that used in the context of spiritual practice. As I have pointed out, some will also have a good understanding of the epistemological and metaphysical background assumptions of these terms, and many will also have developed their very own, more intuitive understanding of what terms like “emptiness” refer to.

One example of this theoretical influence on phenomenological reports is found among religious practitioners. When trying to make verbal sense of their own contemplative experience, Christians may gravitate toward notions like “fullness” or “grace,” while those inspired by Buddhist teachings may actively search for the experience that most directly correlates with their own intuitive understanding of “emptiness” or “spontaneous presence.” A few years ago, I was invited to give a talk at the European Summer Research Institute, an event initiated by Mind & Life Europe that took place on the beautiful Fraueninsel at Lake Chiemsee in Germany, a large freshwater lake often called “the Bavarian Sea.” (My talk later got censored by the organizers and banned from the website, along with that of the Buddhist teacher and author Stephen Batchelor.) During this meeting, Matthieu Ricard importantly pointed out how much of a difference is made by the basic assumptions from which a spiritual practice begins—for example, whether it starts from the idea of an original sin or an innate, ever-present Buddha-nature.

This is obviously right. Real spiritual practice, in my view, involves none of these things—it is the epistemic practice of *liberating* oneself from background assumptions. What we call “background assumptions” are not mysterious abstract entities, but conceptually and linguistically mediated processes that also involve subpersonal neural differences. The human self-model is the computational link connecting subpersonal and suprapersonal (i.e., social) levels of information processing.³ We are enculturated beings, and cultural background assumptions create personal-level expectations, which continuously change the embodied brain’s inner landscape of priors and unconscious beliefs. These assumptions are not merely reflected in the conscious layers of our self-model; our beliefs actually *penetrate* it from above. What is more, many of the preexisting traditional frameworks not only provide a conceptual system of phenomenological descriptions, but they also present us with a *prescriptive* element, telling the practitioner what progress really is and how the different “stages” of meditation *should* look.

This observation leads us to a second aspect of meaning-making: To practice meditation seriously, you need a certain quality of earnestness, a strong and reliable source

of motivation and commitment. Intrinsic motivation is therefore an important causal force, and not everybody has it. For example, it is empirically plausible that meditation practice has low appeal for those who are “existentially indifferent”—the roughly 35 percent of the general population who show a low commitment to all sources of meaning but at the same time have no problem at all with this general lack of meaning. Existentially indifferent people demonstrate a distinct lack of interest in self-knowledge, spirituality, and explicit religiosity. Empirical research has shown that their psychological well-being is lower, but also that they have fewer mental health problems—the lack of experienced meaning and commitment to a system of values does not really bother them or cause psychological crisis.⁴ They are just not interested in these things.

Empirically, I would predict that only very few long-term practitioners of meditation will be existentially indifferent in this sense. On the contrary, it seems likely that many of them will not only follow the philosophical ideal of living an examined intellectual life but will also have a higher-than-average interest in self-knowledge, spirituality, and explicit forms of religiosity. Committed practitioners will be those who have taken the ideal of self-knowledge to heart and transposed it into the more fundamental, nonconceptual levels of their life, thereby extending the original philosophical project well beyond the space of thought. They have an expanded interest in knowledge. As explained previously, long-term meditators will begin to embody their practice—and perhaps also to embody some of the background assumptions that guided them in the beginning, when they first joined a movement or identified with a teacher.

The qualities of intrinsic existential seriousness, sincerity, and often admirable discipline manifested by people highly preoccupied with a search for meaning can create their own pitfalls, however, because the sense of commitment may often be anchored in an ideology, authority, or other form of attachment to a metaphysical belief system that cannot be supported by rational argument or empirical evidence. In these cases, as practice progresses, personality changes will gradually ensue. The interaction of meditative practice and ideology is likely to lead to the construction of an alternative life story, and the practitioner may get entangled in an evolving narrative identity. Some give themselves a new name; others begin to wear funny clothes or even begin to act as spiritual teachers. (People have been known to actually write books.) In sum, the process may create a new kind of personal-level self-model, of the kind that Loch Kelly called a “spiritual ego-manager.”⁵ This may involve contraction into a spiritual superego—a more or less fundamentalist know-it-all who constantly labels, judges, and controls.

Strong intrinsic motivation to meditate also creates methodological problems. Many of our participants (about 77 percent) were regular practitioners who had been meditating for years; it seems plausible to assume that such strong discipline may

often be anchored in adherence to specific belief systems and the conceptual framework of a certain lineage or spiritual tradition, or of a specific teacher, movement, or organization. This throws into question the whole idea of an “expert meditator” because the so-called expertise valued by scientists recruiting for experiments may be the most theory-contaminated of all. I have been talking to meditators all over the world for more than four decades and have engaged in numerous personal communications in the context of this research. For many meditators, there clearly exist “metaphysical taboo zones”: things that remain and must remain unexamined. I also know from private correspondence that such belief systems can sometimes even lead to a decision *not* to participate in scientific projects. For phenomenological surveys like ours, this also means that there may have been not only a positive but also a negative self-selection bias at work.

It is conceivable that a relatively large percentage of *phenomenological* experts chose not to participate in our study because the idea of examining pure awareness contravened some aspect of their belief system. Perhaps the number of practice hours is not a particularly useful parameter, because in some cases, it expresses the degree of attachment to a certain belief system rather than some mysterious form of “expertise.” Here, then, is a research question for the future: At what point in the lifetime of a long-term practitioner will theory contamination reach its peak? This is an open question. It is certainly possible that some motivating belief systems are actually self-reinforcing, becoming *stronger* over time and gradually turning the meditator into a pious, intellectually narrow-minded “true believer.”⁶

Obviously, all this is not to say that ancient theories of pure awareness—based on millions of hours spent in silent meditation by the serious scholar-practitioners who came before us—do not have great value. To quote the old analogy coined by Bernard of Chartres (who died after 1124), in this domain of research, too, we are truly standing on the shoulders of giants. I think that the phenomenological depth of these traditional systems is often enormous, and it’s positively embarrassing that Western science has failed to acknowledge their value for so long. But a fresh, bottom-up approach in a new historical context, spanning many countries and invading intellectual taboo zones, will have its own value. Such an approach helps to weaken the influence of theory contamination (which creates positive biases) on the one hand and taboo zones (which create theoretical blind spots) on the other. Nevertheless, I believe that when the focus of our inquiry is the phenomenology of pure consciousness, the influence of preexisting theories is a much more serious obstacle than it is in other, related fields of research. In what follows, I want to highlight the two fundamental psychological needs that most regular practitioners will have, briefly pointing out how they relate to

intrinsic motivation and the methodological problem of theory contamination, specifically in the context of pure-awareness research.

As we have seen, the first of these deep needs is to develop a successful and sustainable strategy of meaning-making—to find a more permanent solution to the constant struggle to make sense of life events. This need is shared not only by meditators, but by all those who are not existentially indifferent. But for many of our participants, their meditation practice will have been a central part of their own strategy of meaning-making, embedded in the larger context of the attempt to make sense of their own life. “Meaning-making” is a vague and slightly superficial term that often refers to the process by which human beings internally construe or try to understand or emotionally make sense of the chain of events constituting their life, relationships, and what they take to be “their” self.

Typically, the difference between the mere subjective experience of meaningfulness and the much more difficult philosophical question of whether something like an objective “meaning” exists at all—and how it can be known, if so—is deliberately glossed over. For many people, regularly repeating experiences of meaningfulness or of “being touched” seems to be good enough; the good feeling is what counts, while the deeper issue relating to the conceivability of objective meaning in life is systematically avoided. Often, the idea of meaning-making is used as a confabulatory euphemism, a milder word for something considered to be too harsh or blunt, for a reality that we do not want to face because it is too unpleasant or embarrassing. I will call this reality our deep-rooted need for “narrative self-deception.”

One simple conceptual point to note at the outset is that words and sentences have meaning, but events don't. Events just happen. Events as such, as well as whole chains of events connecting birth and death, are intrinsically meaningless—but not in the emotionally negative sense that we may attribute to the idea of life as a whole being meaningless. What “meaning-making” really refers to is an active search for positive experiences that create certain phenomenal qualities of emotional security and coherence over time—a quest for transtemporal sameness of the self, ideally accompanied by the phenomenal signature of knowing, a positive feeling of insight. Of course, given the impermanence of everything, this coherence over time can be only limited at best. There is a quality of absurdity in the depth of our longing for coherence and the obvious futility of all our attempts to ever reach a stable state; and there may be a related form of absurdity in seeking this coherence on the level of phenomenology, in fleeting conscious experiences of “meaningfulness.” I think that to intentionally engage in narrative self-deception is a fundamentally antiphilosophical activity.

To repeat our simple conceptual point from before: There is a difference between systematically seeking mere *experiences* of meaningfulness, or cultivating the phenomenology

of existential soundness and coherence, and the more genuinely philosophical project of investigating the possibility of *epistemic* meaning-making while avoiding the trap of the E-fallacy (as explained in chapter 7). It is quite possible that many practitioners of meditation do not see this difference. Be that as it may, for many meditators their practice is something that not only takes place from moment to moment and day to day, but also is attached to a larger project of understanding life by building a coherent inner model of their own existence as a whole, of their spiritual path as it unfolds over time.

Let us flesh out the new concept of “narrative self-deception,” first by contrasting it with what it is not. A second, equally vague, and slightly superficial cluster of ideas centers on notions like “narrative self,” “narrative identity formation,” and the like. Of course, a lot of current research shows that there simply is no self that could be the active narrator of an inner life story. The whole point is that the process in question is entirely selfless but *seemingly* creates an entity that is in control and remains the same across time.⁷ There is no narrative self, only a subpersonal process generating a fictitious unit of identification. But for enculturated beings like us, the idea of a narrative identity is highly intuitive, and therefore often recurs in popular debates or phenomenological approaches to the mind. It is as if characters in a self-organizing life story suddenly assumed that they actually had a genuine form of self-awareness (which no fictitious entity can have) and even began to play a strong authorial role (which no character really has). It is as if they were trying to control the script—perhaps even trying to “wake up” to their own fictitiousness. But likening selfless, dynamical self-organization in the embodied and enculturated brain to a work of literary fiction, inspiring as it may be, creates puzzles and paradoxes. For example: Can one simultaneously be an author of and a character in the same life story? “Narration” vaguely refers to a data format that appeals to us because in a certain way, we really seem to live through our own stories—externally as well as internally. Human beings are storytellers, and they stabilize the fabric of their long-term self-model by creating a permanent inner monologue. They identify themselves with Harding’s “little one” (chapter 8), the entity constantly chattering on about all its extremely interesting and important perceptions and thoughts and emotions to keep the long-term self alive. There is no better way to understand this fact, to observe its continuous fight for survival under the microscope of mindfulness, than to participate in a silent meditation retreat.

The idea of a life story seems entirely natural at first, and it has the advantage of potentially adding not only causal but also *thematic* coherence to our long-term self-model. A story encourages us to hallucinate cause-and-effect relationships between events in our narrative (“Nothing ever happens by chance!”) and detect patterns in

the way that they affected us (“I have always been interested in this!”). Most of all, a life story offers a chance to insert a single, overarching theme that holds everything together, a normative metacontext that allows us to morally evaluate ourselves and others with the goal of enhancing and sustaining our sense of self-worth. The romanticism of “being on a path” is a good example of this. Now we can judge, negotiate, plan for the future, and thereby continuously stabilize the long-term self-model. There is a way of life, a general theme. The story has a beginning and a possible end; it has characters, episodes, a setting—and like a dream that has become fully lucid, it even has a plot that can apparently be controlled. Nonlucid dreams, in which plot control is absent, are a nice example of what “narrative self-deception” actually involves on a subpersonal level. While dreaming, you can see the storytelling mechanism in action: a mechanism desperately attempting to weave random events into a coherent narrative; an automatic brain process in constant search of a thematic red thread; the tireless creation of meaning, intelligibility, and control in an inner environment that in reality is entirely unpredictable. The result is the dream narrative you can later report.

Perhaps you know from your own dream experiences the feeling of attempting to remember something important that you apparently *just* forgot—the feeling of constantly trying to reestablish a sense of control, trying to recall the overall theme, as it were, or to understand *what it was all about*. Like Sisyphus rolling a boulder up a mountain only for it to roll down again every time it neared the top, the dreaming brain tries to stabilize the long-term self-model. Again and again, it looks for the red thread, attempting to successfully hallucinate sameness across time. A never-ending Sisyphean task, it shows not only that narrative self-deception is something that takes place on the level of whole persons and culture, but also how we embody the battle for meaning and uncertainty reduction, subpersonally, even at night. Narrative self-deception is a general computational principle that works on a number of levels simultaneously.

My main point is that, for a considerable number of meditators, their practice may play a crucial role in endowing their own lives with thematic coherence, to the point where their ongoing search for thematic consistency may actually be more important than the actual practice and any of its other causal effects. Sitting on a cushion, alone or with others, may be only one part of crafting an overarching framework, finding a more gratifying sense of community, and somehow trying to feel at home in this dangerous and unpredictable world. Meditation may be part of an attempt to “recall the overall theme,” driven by the hope of *finding something*. How the theme is formulated will influence how practitioners later describe their experiences—but, of course, the theme will also partly determine the actual content of those experiences themselves, what they actually find.

The impact of narrative context on the contents of consciousness seems crucial to acknowledge. However, as Emily Troscianko has reminded me in a personal communication, the whole idea of “the” actual content is in fact dubious because we have to get away from talking both about consciousness as a container with “contents” and about “the experience itself.” New theories of the human mind as a form of hierarchical Bayesian-updating,⁸ plus much empirical data, suggest that we should instead develop a new version of what Daniel Dennett has called the “multiple drafts model” of consciousness.⁹ In this model, the nature of the probe applied (e.g., the precise way in which we ask the question “What am I conscious of now?”; see figure 34.1 in chapter 34) determines what is retrospectively called “the experience,” with no “real phenomenology” beyond that post hoc labeling. There is a quality of groundlessness here, and this is where advanced spiritual practice comes back in: If we start to investigate the actual process of life-storytelling with fresh eyes, under the microscope of mindful attention, then it becomes obvious that what’s going on is really a competition among multiple drafts. In this new context, let me also draw attention to the fact that, quite often, narrative self-deception may have enabled human cooperation in large groups and in turn now critically depends on that social cooperation, on the kind of mutual reinforcement we find in religious communities, cults, or spiritual movements trying to construct an alternative reality for themselves. All these social factors shape the virtual stories competing for airtime, as well as dramatically raising the stakes of the competition.

The shaping of experience and recall by meaning-making endeavors also has a sociological aspect because for some meditators, the project of meaning-making is that of actively “re-enchanted their life-world.” Max Weber, when he coined the famous phrase *die Entzauberung der Welt* (the disenchantment of the world), was referring to the sense of disenchantment caused by cultural rationalization, by the unstoppable rise of science and technology that had led to a bureaucratic, “soulless,” and secularized Western society, where scientific understanding was more highly valued than faith or personal experience. More than a century ago, in 1918, Weber wrote: “The fate of our times is characterized by rationalization and intellectualization and, above all, by the ‘disenchantment of the world.’”¹⁰ This sociological context has created countermovements, and for many meditators, their own practice and the way that they interpret their experiences may be precisely a project of restoring the magic, aiming at *die Wiederverzauberung der Welt* (the re-enchancement of the world).

As our data show, this project can clearly work on the level of individual phenomenology: Just think of the experiential quality termed “nonsensational awe” (chapter 1) or our investigation of “suchness” (chapter 9). You may also think of the experience of connectedness (chapter 11) and the rich spectrum of states characterized by joy

and gratitude (chapter 15). As we move on, we will encounter more examples of MPE modes of consciousness that clearly involve a magical and enchanted quality, like timeless change (chapter 22) or translucency and virtuality (chapter 28). One important research target for the future will be to better understand the mutual relationship between the meditator's individual phenomenology and the sociological context in which it is embedded.

Our problem is that we need a better understanding of theory contamination as it applies to the pure-awareness experience. We need an understanding that goes beyond intuitive, feel-good, vague, and therefore popular notions like "meaning-making" or "narrative self." Let me offer a series of new concepts to get a clearer view of this problem. The first of these concepts is the "autobiographical self-model," which allows individuals to organize episodic memories and abstract knowledge of their past into a coherent biographical picture. The autobiographical self-model is not a thing, but rather a process. It is also not a little man in the head. This process is mostly subpersonal: To a large extent, it is not something that *you* do, but an automatic process taking place under the hood, in your unconscious brain. The process constantly tries to minimize surprise and to keep a certain layer in your self-model as coherent as possible (namely, the layer that portrays your life as a whole). At the same time, the process of the autobiographical self-model tries to make the organism more intelligent by expanding what mathematicians and computer scientists would call the "predictive horizon."

Intelligence can be described as the capacity to successfully predict the future, the continuous attempt to expand the organism's predictive horizon. Biological intelligence consists partly in maintaining a steady state in the here and now (e.g., through homeostatic control of body functions), but efficient regulation also requires anticipating future needs and preparing to satisfy them before they arise (called "allostatic control"). The human self-model does all these things at the same time. The body model helps with online motor control and with simply staying alive, while the autobiographical layer in the self-model stores a selection of past events while also planning for the future, seamlessly connecting what has already happened to possible future selves. The model expands into time, creating a past and a future. Whenever this layer is conscious, it can function as what we have termed the "unit of identification." You may think that you simply *are* your life story, which springs from the ongoing search for thematic stability. But you may also begin to hallucinate an observer, an active narrator, or a stable protagonist within the story itself—an entity that reflects on and monitors the whole process (more on this in chapter 25).

I coined the term "narrative self-deception" a few years ago, when investigating the problem of conscious suffering.¹¹ Here is a concrete example. A spiritual biography can

be a form of narrative self-deception, at least in part. It can create an illusory unit of identification, an invisible yet robust sense of self endowed with a feeling of transtemporal sameness. Scientifically, a number of different types of memory have been identified and classified, and we know that none of them is aimed at veridicality, that all are instead incrementally optimized for replication and survival. We have also learned that self-deception has evolved, making animals more efficient in the short term and also allowing them to deceive *others* more successfully.¹² Today, we must face the fact that our autobiographical self-model is not a tool for self-knowledge in a philosophical or spiritual sense, but rather something that evolved in an evolutionary arms race. As a matter of fact, it seems to be one of the major obstacles to the kind of self-knowledge sought in contemplative practice. It created an elaborate form of biological intelligence, but the formation and organization of autobiographical memories also constitute the central mechanism through which our conscious sense of identity across time is constructed. This mechanism is what makes the mind wander, what creates dual awareness, and what continuously decouples attention from the present moment, and of course, it directly influences how individuals develop their more temporally extended sense of self. Long-term goals play an important role in the embodied process of meaning-making, as do the emotional meanings that we assign to them. Thematic coherence—or its loss—is something you can *feel*.

Narrative self-deception may also be a clever strategy for distracting attention from suffering. It allows you to feel better about yourself. From an evolutionary perspective, any self-conscious system that discovered too many negatively valenced moments—too many conscious experiences that it would rather not go through if it had a choice—might become paralyzed and stop procreating. Like a Buddhist nun or monk, it might refrain from adding new human beings to this world of impermanence and delusion because its own insight into the nature and ubiquity of suffering would have become too clear. Monks and nuns who obey the rule of celibacy are nobody's ancestors: They do not copy their genes into the next generation, and therefore, they are side alleys and dead ends of biological evolution. The logic of psychological evolution clearly mandates concealment of the facts of impermanence and conscious suffering from any self-modeling system that is supposed to be an efficient copying device. As a general rule (but with plenty of exceptions), natural selection will often have favored those who were aggressive and greedy, functionally attached to the goal states that they hallucinated, eager to have children and to rise in a social hierarchy, and reliably caught on the hedonic treadmill.

Systematic distraction from the potential insight that the life of an antientropic system is one big uphill battle, a strenuous affair with minimal prospect of success in more than the very short term, may certainly not be in the interests of the individual vehicle itself, but it probably causally enhances the process of genetic evolution, ensuring that

genes keep getting replicated. Any conscious representation of negatively valenced uncertainty under the condition of identification (chapter 8) causes suffering. The evolutionary algorithm therefore has selected for genetic instructions that make it less likely for an organism's insights into the deep structure of its own mind—insights of the type just sketched—to be strongly reflected in its conscious self-model. As a result, we now have a cognitive scotoma for conscious suffering,¹³ a systematic blind spot that makes it difficult to see certain obvious truths about our own lives.

Conversely, an adaptive advantage may have applied to humans' tendency to suffer from a robust version of optimism bias as one form of narrative self-deception: If you unrealistically believe that you are much less likely to experience negative future events than anyone else, you may procreate more successfully. The primary function of the human self-model's autobiographical levels may be precisely to drive the organism forward by relentlessly expanding its sense of self in time, thereby making it lose contact with the often-unpleasant reality of the present moment. The autobiographical self-model of humans may have been so successful in part because it provides a functionally adequate form of self-deception, glossing over the ugly details of everyday life by developing a grandiose and unrealistically optimistic inner story—a “narrative self-model” with which we can identify.

What has all this got to do with meditation practice and the problem of theory contamination? Well, many forms of meditation practice involve the opposite of this future-oriented sugar-coating of the present: Mindful awareness means continuously *shrinking* the conscious brain's predictive horizon until the quality of timeless change, which we will investigate in chapter 22, begins to reveal itself. Mindful awareness is what *ends* the narrative. The bad news, however, is that for at least some meditators, the wider context of meaning-making and narrative self-deception may be at least as important as the actual practice itself—this possibility is what Krishnamurti's question pointed to. (The same may be true of other forms of epistemic practice: scientists conducting meditation research, say, or academic philosophers interested in the problem of consciousness.) Therefore, a lot of what appears to be strong intrinsic motivation may actually be extrinsic, driven by the surrounding social context by having succumbed to some spiritual authority or some set of metaphysical background beliefs. Belief systems offer labels, value-charged theories of consciousness, and taxonomies to describe inner experience. This simple point ends my sketch of the first of the two deeper roots of the methodological problem of theory contamination: narrative self-deception and mortality denial. We must face the fact that in the phenomenological material that I am presenting, all the factors that I have sketched out here will have contributed to an unknown amount of contamination.

The good news is that there is probably no better tool than the practice of meditation itself to help us *really* understand the mechanisms of self-deception that continuously

express themselves as the subtle workings of one's own mind. This is a deeper form of understanding than mere theoretical knowledge—one that can have direct causal consequences because it takes place on a nonintellectual level that goes far beyond words and concepts. The fine-grained mental mechanisms of meaning-making itself are one of the most interesting objects of mindfulness practice. Classical insight meditation, for example, consists in nothing other than observing the actual process of narrative self-deception at work: Every single thought arising is an attempt to generate a new self-model, to become “temporally thick,” to escape the wholeness of the present moment. Here, you can catch the king of escapees in the act and finally watch the autobiographical self-model under the microscope, including the fact that it is actually selfless—a self-model only in the computational sense, and not really a *self*-model when it comes to the phenomenology of actual experience. The automatic process of expanding the predictive horizon and automatically identifying with some future self is entirely empty in the sense that it is epistemically open and has no fixed meaning. A lot of what I said about the selfless and entirely nonconceptual quality of “suchness” in chapter 9 probably also applies to the act of viewing our own life as a whole. There is no tragedy in suchness. Life just happens—and meditation practice is a perfect way to become aware of this fact.

It is true that the need for meaning-making often leads to systematic and even socially organized forms of narrative self-deception, but it does not necessarily have to. Let us turn to the second fundamental psychological need that most of us share, meditators and nonmeditators alike: *mortality denial*. We all have to manage death anxiety, and we all have to come to terms with the qualities of absurdity and futility that accompany any insight into our own finitude. We may not consciously experience this need, but there are subpersonal as well as cultural mechanisms constantly trying to satisfy it for us. There are whole industries dedicated to helping us practice narrative self-deception and mortality denial—just think of the Roman Catholic Church. The fact that these industries have endured for many centuries, surviving much longer than any form of government, any political system, any tribe, and most nations, demonstrates how deep and universal these needs really are. After they appeared, they changed our inner environment, and we had to adapt to this fact.

Human beings are the products of evolution, which means that our bodies and our minds have been optimized not for happiness or self-knowledge, but for survival and procreation. At the same time, we have a new problem that no other animal on this planet seems to have: We have a direct, self-conscious insight into our own mortality, into the fact that ultimately, I will *not* survive. I call this new problem the emergence of a “toxic epistemic state,” something that interferes with the coherence of our

autobiographical self-model, as just described. Let us define a “toxic epistemic state” as any form of knowledge or insight that threatens the biological fitness of an organism, its individual reproductive success. An epistemic state that is toxic from a biological perspective may be seen as a healing or liberating form of insight from a spiritual perspective. One therefore may speculate that specific forms of self-knowledge have become much more explicit in some spiritual practitioners than in their fellow human beings, such as insights into mortality, futility, and the prevalence of unnecessary suffering. But viewed from an evolutionary, biological point of view, the knowledge that my own death is inevitable is poisonous. It threatens my mental integrity. A solution has to be found. This knowledge calls for an expansion of narrative self-deception into the time that lies beyond my physical death—or into the dimension of timelessness. In this context, again, it is interesting to note that in the history of humankind, many spiritual practitioners were nuns and monks who took a vow of chastity and had no children.

In his 1973 classic *The Denial of Death*, Ernest Becker said that mortality denial can take the form of a personal “immortality project” in which a symbolic belief system ensures that I can believe that my self is superior to physical reality. My first point is that for some meditators, their own practice may be part of such an immortality project, and this fact may at times distort how they choose to speak about their own inner experience. Here is how Wikipedia explains one of Becker’s central insights: “By successfully living under the terms of the immortality project, people feel they can become heroic and, henceforth, part of something eternal: something that will never die as compared to their physical body.”¹⁴ If Becker is right, there may be a specific form of mortality denial that we could call “contemplative heroism.” Meditation practice would be part of a larger project, the heroic battle against finitude, an epic quest to discover that part of yourself that you can firmly believe will never die.

Mortality denial has many other facets and consequences beyond personal immortality projects. For example, populism is directly related to meaning-making and mortality denial. If you are a politician, offering a tribal narrative plus a form of symbolic identity that allows people to identify with something greater, with something that extends far beyond their physical death, then you will typically be more successful than any secular or intellectually honest competitor you may have. If you are one of the many new entrepreneurs in the global attention economy (or an algorithm aimed at extracting attention from human brains), then you will always be more successful when you learn how to play to your audience’s unconscious needs by offering an interesting new strategy for narrative self-deception or a particularly clever route to mortality denial. And if you are a philosopher, you will always have more readers if you manage to subtly include some new and clever backdoors for mortality denial in

your work. A mild but strategic form of vagueness is a very useful ingredient because it helps the reader to actively project her desired experience of meaningfulness into your words. Even if people don't fully understand what you are saying, they will get an intuitive sense that "this feels right," that—even if what you really mean remains slightly obscure from a conceptual point of view—on an emotional level, it literally "makes sense," and therefore it must be relevant and significant in some deeper way to *them* and their own project of meaning-making, of successful narrative self-deception. Just like assisted suicide, assisted self-deception is a type of social cooperation: It is special in that it often involves interesting new forms of complicity and creativity in mortality denial, forms that are richly rewarding for all parties. Of course, everything that I have just said about ambitious politicians, attention entrepreneurs, and philosophers can apply equally to spiritual teachers.

In this chapter, I have already provided a number of new conceptual instruments: the idea of meaning-making as "narrative self-deception," the concept of an "autobiographical self-model" that automatically creates an image of our life as a whole, and the notion of "contemplative heroism" as a strategy for dealing with impermanence and futility. In this final part, let me offer you two more conceptual tools. The first is "absurdity management." Insight into mortality is a toxic form of self-knowledge, a permanent danger that has unexpectedly sprung up in our inner environment. Developmentally, it typically emerges at around the age of eight to ten years, and it presents an obvious danger to our mental health, including to our sense of integrity and self-worth. Not only does it create a concrete form of anxiety that needs to be managed,¹⁵ it also creates an almost all-encompassing, global feeling of futility and absurdity that has to be kept out of the conscious self-model just as urgently as coherence needs to be found in a dream. I will call this experiential quality the "Sisyphean quality." Our biological self-model has now been damaged—informationally poisoned, if you will. Therefore, it calls for a boulder-pushing adaptation, a psychological coping strategy, a way of efficiently integrating or simply denying the reality of death. The brain *must* react; it must create a sustainable new model of reality.

Here, my general point is that for many meditators, their practice will likely be part of a wider strategy of mortality denial, of "finding evidence" to build this new death-denying model of reality. The practice can create certain health-promoting, stress-reducing altered states of consciousness, but what really counts for the individual practitioner may sometimes be the metaphysical *interpretation* of these states, the adherence to a new belief system that helps with death denial. Please note that all of this is only one example of what may actually be humankind's oldest coping strategy, *the* classical form of absurdity management: First, you create altered states of consciousness,

such as by drumming, dancing, fasting, depriving yourself of sleep, performing shamanic rituals, or gathering hallucinogenic plants and mushrooms (step 1); and then you *interpret* in the desired fashion the alternative models of reality that your Sisyphean brain creates in such states (step 2). For example, you could later say that you have now “directly experienced” the fact that something can exist without the physical body, that you have communicated with spirits, the angels, the dead, or similar entities. Of course, the altered states of consciousness created by meditation practice can be abused in the same way—so long as you interpret them as a direct experience of “reality itself,” creating the right social context for cultivating certain versions of the E-fallacy (step 3).

This may also be one reason why many Asian systems of philosophy do not clearly separate the phenomenology and the metaphysics of pure consciousness: These systems originated in religious systems of mortality denial. If this is so, then there is a methodological problem, because such belief systems will not only color external experiential reports but may also (via the person’s self-model) install unconscious assumptions in the very brains that shape the experiences themselves, like an additional set of self-fulfilling prophecies acting on the subpersonal level.

Given this new context, do you remember the “dolphin model of meditation” sketched out in chapter 10? It showed how unconscious assumptions can make certain experiential contents “spontaneously appear,” as if uncaused and unborn. In short, many human beings have a deep need for absurdity management and their individual successes (and failures) in satisfying this need will color the phenomenological reports that they give. This is all too human and nothing to be condemned, but it does create a methodological problem.

Bhava-taṇhā is a term that Buddhist philosophers have known and analyzed for 2,500 years. It refers to the craving for existence, which is one of the deepest causes of conscious suffering in humans, and probably in many other self-conscious animals too. To be driven by the craving for existence, you do not have to possess explicit knowledge that you will die. *Bhava-taṇhā* can also be interpreted as the craving to be something specific or to unite with an experience, as well as the correlated striving for permanence—the phenomenology of identification discussed in chapter 8. I think that modern science is making it more and more obvious that biological systems like ourselves are actually physical *embodiments* of this craving for existence—if you will, antientropic Sisyphus-machines fighting a constant uphill battle to preserve their own boundaries. I think that this results in a deep and fundamental distortion of the way that our conscious models of reality portray the world and ourselves in it. This distortion had been nonconceptually investigated by meditators long before modern science even entered the stage. I call it the “existence bias.”

The link between our need for absurdity management and existence bias is that there is a deep absurdity in *being* an embodied craving for existence while at the same time knowing that you will have to die (i.e., having a cognitive insight into the fact that this craving will one day be frustrated). The insight is toxic; the craving becomes poisoned. So what exactly is existence bias? One of the deepest roots of human suffering and self-awareness is a top-level preference that creates a self-directed variant of existence bias: the fallacy of treating the mere existence of something as evidence of its goodness. Here, however, the concept of “existence bias” refers not to the well-documented fact, widely known in social psychology, that human beings generally favor the status quo,¹⁶ but rather to the more specific observation that beings like ourselves will almost always opt to sustain their own physical existence.¹⁷

Of course, human beings will sometimes sacrifice themselves to save their offspring or protect their tribe. We are replicator-copying survival machines that have been mercilessly optimized for millions of years to never give up, to optimize inclusive fitness, and to maximize our contribution to the gene pool. But on an individual, physical level, we are antientropic systems fighting an uphill battle in a constant attempt to reduce uncertainty and “understand ourselves” by finding a viable strategy of self-modeling—we are free-energy machines who, in the words of neuroscientist Karl Friston, continuously “maximize the evidence for their own existence.”¹⁸ As the philosopher Jakob Hohwy has made admirably clear, we are not just biological agents endowed with information-hungry brains relentlessly gathering more data to produce ever more evidence for our own existence.¹⁹ We can also be viewed as self-organizing systems trying to continuously *expand* our predictive horizons, desperately sustaining our existence in a dynamical environment by following an inner norm of tracking the very conditions of possibility for existence themselves.²⁰ Our phenomenology deeply reflects this computational imperative for constant self-evidencing, and on many hierarchical levels. As a matter of fact, pure awareness itself may be interestingly related to the conditions of possibility for existence themselves (more on this in chapter 32). What makes it so interesting is that it nevertheless lacks the experiential qualities of absurdity and futility.

Armed with the concept of existence bias, we can be more precise about the phenomenology of absurdity and futility. The quality of absurdity consists in the fact that the conflict between the embodied existence bias (which we *are*, from which we originated) and the high-level cognitive insight into our mortality has seemingly become a permanent functional feature of our self-model, something built in. If you will, it is now not a bug but a feature—and that is what brings the quality of absurdity into human existence. As already noted, *bhava-taṇhā*, the automatic craving for existence, is one of the deepest causes of conscious suffering in humans, and probably in many

other animals too. It also led to the evolution of religion, because what is special for humans is that we have to deal with the additional challenge of “toxic self-knowledge” threatening the integrity of our self-model. The quality of futility originates in the fact that although we want to avoid nonexistence completely, we now have explicit intellectual knowledge that, ultimately, this will be impossible. Futility is one major component of absurdity. No other animal has this kind of problem. As human beings, we need a strategy for absurdity management because we explicitly know that every single individual will eventually lose the uphill battle sketched here—that our predictive horizon will eventually shrink to zero simply because in biological evolution, “passengers are not carried.”²¹

Bhava-taṇhā and the concept of an embodied existence bias also yield a new perspective on our deep-seated tendency toward mortality denial: Mortality denial is not simply some sort of false belief or a misguided intellectual attitude; it is a computational principle. This principle consists in automatically neutralizing “mortality-salient information”—all the toxic and potentially demotivating information that leads to the individual’s awareness that death is inevitable—by explaining it away in the context of a new, alternative model of reality. We embody this principle. This is why it surfaces on the level of psychology and culture, and why it creates a systematic blind spot in our conscious self-model.²²

Of course, much more could be said at this point about different levels and units of selection in the evolution of death denial, particularly when it comes to the interplay between biology and culture. But figuratively, we were never *supposed* to mentally simulate possible worlds in which we have died, just as we were never *supposed* to use up too many precious mental resources to imagine situations in which we are sick or have no offspring—except, of course, in acutely dangerous situations where fast, intelligent planning is needed.

Friston says that “biological systems move around in their state space, but revisit a limited number of states”—namely, those in which we are safe and healthy, situations that do not present us with ugly surprises but that actually correspond to stable, survival-friendly homeostatic steady states far from equilibrium.²³ Please note that the insight into mortality is exactly one such ugly surprise, but one that occurs in our *inner* environment. Biological organisms like ourselves must maintain certain vital parameters, certain critical physiological variables, within particular bounds; we must avoid unexpected phase transitions and protect a state far from thermodynamic equilibrium.

Technically, any successful biological organism will achieve in its lifetime only a very small number of all possible obtaining states. Every member of this small group of states has a high probability, and being dead is not one of them. Unless forced

by circumstances, we will never plan for or imagine all those other states that have low probability, or else we wouldn't be here. Morbid mopers were rarely among our ancestors. The simple reason is that beings like us needed all their limited mental resources to continuously develop new existence-sustaining action policies. They used their brains to simulate paths into physically improbable worlds where they and their children were still alive and healthy. This is an important part of what it means to act intelligently and adaptively within our environment. Ultimately, our tendency toward mortality denial falls out of statistical physics: Beings like ourselves survive by predicting our own existence, not by predicting nonexistence. Previously, we saw that human consciousness can be described as being governed by an embodied existence bias. Now we can extend this point: In a sense that could be technically well defined, we actually *are* an embodied form of mortality denial. We are resistance itself.

But we are also social beings. In dealing with toxic self-knowledge, we have had to develop enculturated strategies for mortality denial and narrative self-deception, which in turn shape the structure of our conscious self-model in a top-down manner. We invented organized religion—and we became dangerous beings who were religious but not spiritual (our snappy acronym here might be RBNS). Our immortality projects and the faith-based religious heroism of absurdity management have created an endless chain of wars and social conflicts. Ultimately, death anxiety has created more death anxiety.

Again, let us ask: What does all this have to do with meditation practice and the problem of theory contamination? The bad news is that for many meditators, their practice is likely tied to a wider context of mortality denial: using the phenomenology as proof of the existence of an afterlife, interpreting it in a way that confirms some organized system of mortality denial. To illustrate, here is an interesting detail: Humankind has evolved five major religions, but only two of them—Hinduism and Buddhism—are strongly and unequivocally related to a truly systematic cultivation of contemplative practice and the experience of pure awareness. At the same time, Hinduism and Buddhism are the only two of the major religious systems that involve an explicit belief in reincarnation. Whereas in monotheistic religions like Catholicism, mortality denial is primitive, blunt, and direct (according to the Nicene Creed, there will simply be a resurrection of the dead), many ancient and all major Indian religions (namely, Buddhism, Hinduism, Jainism, and Sikhism) promote belief in another metaphysical model, according to which at least one part or aspect of the currently living human being will start a new life in another physical body after death.

The Buddhist and Hindu models for liberation and mortality denial are slightly more sophisticated because they involve a false alternative: Of course, we do not *want* to be reborn because our true goal is to liberate ourselves from the cycle of birth and death. But if we do not have the good fortune to attain enlightenment in this life,

the poor ego will simply have to continue in another physical form. We could even call this the “binary model for death denial.” The false dichotomy behind the binary model is easy to see through, but at least it offers two possible trajectories for successful absurdity management by narrative self-deception: enlightenment or reincarnation—giga-bingo or another life.

I am aware that all this may sound harsh. But the good news is that there is probably no better tool for *really* confronting one’s own mortality and the fear of death than meditation practice itself. This is what goes to the root of the problem. It takes more courage, but it may also be more efficient than any form of intellectual gymnastics. Over the centuries, many Western philosophers have pointed out in one way or another that living an examined life has a lot to do with learning to die, and as a matter of fact, confronting one’s own mortality as directly as possible may turn out to be the deeper core of any more serious spiritual practice. As the quotations presented in this section’s epigraph show, the Western tradition of learning to die has focused on using reason to transcend individuality and the passions and to move toward universal principles like rationality and intellectual insight, while the Eastern tradition operated on a more direct and practical level, also attempting to transcend individuality and the passions toward something universal, but by cultivating pure awareness and peaceful ego dissolution. If there are toxic epistemic states, maybe there are also *liberating* epistemic states?

Science and meditation are both epistemic practices; narrative self-deception and mortality denial aren’t. We have explored some of the ways in which meditation proper is an epistemic practice, and how it goes far beyond stress reduction, well-being, capitalist imperatives of self-optimization, or vague ideas of “self-actualization.” Cultivating epistemic openness and surrendering to the real possibility of ego dissolution in more advanced meditative practice may constitute a radical version of what in early Western traditions was termed “living an examined life” and “training for death.” It is the epistemic practice of investigating and confronting existence bias head on (e.g., by viewing it through the lens of nonegoic self-knowledge)—and in this sense, it is philosophical practice and consciousness research taken slightly more seriously.

There is also a deeper phenomenological discovery to be made, which I would say that you can verify through the practice of meditation: Reflexive MPE, the timeless experience of nonegoic self-awareness (see chapter 30 for more), is free of the almost all-pervading, global affective tone of futility and absurdity that was mentioned earlier in this chapter and that needs to be repressed in the egoic self-model. The Sisyphean quality is *not* part of its phenomenal character. On the other hand, MPE is positively characterized by what in chapter 1 was termed the experiential quality of “existential ease.” Ultimately, therefore, meditation may be capable of dissolving the need for narrative self-deception, mortality denial, and all heroic forms of absurdity management.

This is a section of [doi:10.7551/mitpress/15196.001.0001](https://doi.org/10.7551/mitpress/15196.001.0001)

The Elephant and the Blind

The Experience of Pure Consciousness: Philosophy,
Science, and 500+ Experiential Reports

By: Thomas Metzinger

Citation:

*The Elephant and the Blind: The Experience of Pure Consciousness: Philosophy,
Science, and 500+ Experiential Reports*

By: Thomas Metzinger


DOI: 10.7551/mitpress/15196.001.0001

ISBN (electronic): 9780262377287

Publisher: The MIT Press

Published: 2024

The open access edition of this book was made possible by generous funding
and support from MIT Press Direct to Open



The MIT Press

© 2024 Massachusetts Institute of Technology

This work is subject to a Creative Commons CC-BY-ND-NC license.

This license applies only to the work in full and not to any components included with permission. Subject to such license, all rights are reserved. No part of this book may be used to train artificial intelligence systems without permission in writing from the MIT Press.



The MIT Press would like to thank the anonymous peer reviewers who provided comments on drafts of this book. The generous work of academic experts is essential for establishing the authority and quality of our publications. We acknowledge with gratitude the contributions of these otherwise uncredited readers.

This book was set in Stone Serif and Stone Sans by Westchester Publishing Services.

Library of Congress Cataloging-in-Publication Data

Names: Metzinger, Thomas, 1958– author.

Title: The elephant and the blind : the experience of pure consciousness: philosophy, science, and 500+ experiential reports / Thomas Metzinger.

Description: Cambridge, Massachusetts : The MIT Press, [2024] | Includes bibliographical references and index.

Identifiers: LCCN 2023012135 (print) | LCCN 2023012136 (ebook) | ISBN 9780262547109 (paperback) | ISBN 9780262377294 (epub) | ISBN 9780262377287 (pdf)

Subjects: LCSH: Consciousness.

Classification: LCC BF311 .M4725 2024 (print) | LCC BF311 (ebook) | DDC 153—dc23/eng/20230830

LC record available at <https://lcn.loc.gov/2023012135>

LC ebook record available at <https://lcn.loc.gov/2023012136>