

Epilogue: *Bewusstseinskultur*

Is minimal phenomenal experience (MPE) a *good* state of consciousness? If we move beyond phenomenology, science, and academic philosophy of mind, is the experience of pure awareness something positive, something that we should foster and cultivate in our lives? And is there a sociocultural dimension to the kind of research I want this book to encourage? What is the wider context?

At the end of our journey, it is now time to change our perspective one last time, taking three steps back and a deep breath. My two goals in this final section are to open up a broader—specifically, a normative—perspective, while at the same time offering one last philosophical concept: *Bewusstseinskultur*, a culture of consciousness. I will keep things simple. In case you want to know more, there will be a short book to consult that may be helpful.¹

Bewusstseinskultur

I introduced the idea of developing a *Bewusstseinskultur* a quarter century ago.² *Bewusstseinskultur* aims at a special form of cultural innovation. In a first-order approximation, it consists of three major elements:

1. The adoption of an ethical stance toward one's own mental states
2. The systematic cultivation of states assessed as valuable
3. A continuous process of rational, evidence-based enculturation—that is, an active embedding of such states of consciousness in culture and society

Before looking at each of these points in turn, I should make clear that parts and variants of the basic idea have been around for many centuries. Philosophers, of course, have always had long debates about what the term “philosophy” itself really means, what it *truly* means to love wisdom—because that is what *philosophia* (φιλοσοφία in

Greek) originally meant. To give just one example, Marcus Tullius Cicero (106–43 BC) said in the second book of his *Tusculan Disputations*³ that truly loving wisdom means “taking care of and cultivating one’s soul” (*cultura autem animi philosophia est*). I think this is one beautiful and highly topical variant of the general idea of creating a culture of consciousness. Unfortunately, it also points to one of the greatest weaknesses of the highly professionalized (and often truly excellent) form of academic philosophy that we have today:⁴ It is not exactly conducive to the mental health of those who participate in it, let alone to the project of cultivating interesting states of consciousness with epistemic potential, a project that goes beyond mere theoretical knowledge. But given the empirical knowledge about the evolutionary roots of the human mind, and given the power of present-day cognitive neuroscience plus exciting recent advances in computational phenomenology, Cicero’s classical motif offers a good example of what we might gain by radically reinterpreting some of our cultural goals.

Bewusstseinskultur is about a specific form of cultural innovation. It concerns a broadening of horizons: We need to develop a new philosophical perspective that not only assimilates all those exciting new scientific discoveries and insights, but also helps us deal with all the new potentials for action in an ethically sensitive manner. Therefore, developing *Bewusstseinskultur* can also be viewed as a novel form of applied ethics. It is about practical philosophy—or, more precisely, about creating a new connection between applied ethics and philosophy of mind. A host of technological possibilities for manipulating and changing the human mind have emerged—even artificial consciousness is being debated as a plausible part of our medium-term future—and we risk being overwhelmed by the psychological and sociocultural consequences of scientific progress and the technologies it generates. Therefore, we need a whole new branch of practical ethics. This branch could be called *Bewusstseinsethik*—an applied ethics of consciousness—because it specializes in all issues directly related to conscious experience itself.

The central aim of consciousness ethics would be to help us with what Cicero, more than two millennia ago, called “taking care of and cultivating one’s soul.”

Practical Philosophy of Mind

Following Aristotle’s distinction between natural philosophy (aiming at the right kind of theory) and moral philosophy (aiming at the right kind of practice), today’s academic philosophers often distinguish between theoretical philosophy and practical philosophy. Theoretical philosophy consists of subdisciplines like logic, epistemology, philosophy of science, and philosophy of mind; practical philosophy comprises fields

like ethics, decision theory, and political philosophy. *Bewusstseinskultur* will have to unite a specific subset of these two broad sets of research goals in a new way. It can be seen as the societal and political implementation of insights from what I will dub a “practical philosophy of mind.” Practical philosophy of mind starts by taking an ethical stance toward one’s own conscious mind—and also toward the minds of others.

Taking an ethical stance toward one’s own mental states first means asking what makes a mental state *good*. For example, one might ask whether certain states of consciousness possess an intrinsic value of their own. Is there a rational, well-defined way to say that some of the states and modes of conscious experience described in this book are intrinsically valuable, that they are *better* than other states and modes of conscious experience? Are there any objective criteria that we could apply?

In classical ethics, one would ask what makes an action good, or whether there are certain action goals, moral attitudes, or virtues that are intrinsically valuable. Here, the idea is to extend exactly the same strategy to our own conscious minds, by asking: What makes a state of consciousness a *good* state of consciousness? Are there “phenomenological virtues” at all? Are there intrinsically valuable attitudes, ways of life, or goal states of conscious experience? What are beneficial and wholesome ways of dealing with one’s own mind?

This emphasis on the ethical question of conscious experience itself is the main reason why the first of the three elements that I listed here could also be called the project of developing a *Bewusstseinsethik*. *Bewusstseinsethik*, however, is not meant to replace traditional ethics. On the contrary, the idea is to complement and enrich traditional ethics by drawing on empirical evidence and rational argument and focusing on cultivating conscious experience itself.

The Systematic Cultivation of Positive States of Consciousness: Concrete Examples

Let us now look at the second point in the working definition of *Bewusstseinskultur* presented here: the systematic cultivation of valuable states. If our project of developing an ethics of consciousness yields results, then we can begin to think about how to implement them. Depending on our philosophical background theory, we might even feel a moral and political *obligation* to do so. If we accept any moral obligation at all, we also have the obligation to pursue our ethical goals as efficiently as possible, to consider all possible instruments to reach these goals, and to look at everything with an open mind. If and when we have identified positive states of consciousness, what will be the most effective ways to realize them?

All of this may be starting to sound quite abstract. However, *Bewusstseinskultur* is eminently practical because it is a new branch of applied ethics that—as soon as we begin to think seriously about it—leads to a whole range of very concrete questions. To make this clear, here are some examples.

Which States of Consciousness—If Any—Should Be Illegal?

The “technological landscape” for generating altered states of consciousness is undergoing dramatic changes, and it is fair to say that the overall situation is completely out of control. In the pharmacological domain, well over 1,000 new psychoactive substances emerged in illegal markets worldwide between 2009 and 2021.⁵ Digitally, as a consequence of the fertile confluence of artificial intelligence (AI), virtual reality (VR), and neurotechnologies such as brain-computer interfaces (BCIs), new ways of “constructing” phenomenal states will keep emerging. The emergence will be driven partly by technologies that directly target the brain, but also via the creation of new medial environments and entire life-worlds, most likely owned by profit-driven private companies not oriented to either the common good or the mental health of their users.⁶

We can ask specific variants on the overall question posed here, such as:

- Given the recent explosion of new synthetic drugs and the failure of prohibition, what would a rational and ethically defensible drug policy for the future look like?
- Do we need legal regulations for advanced mind-altering technologies?
- How can we create a truly *humane* technology of consciousness? For example, how can we ethically design algorithms that are constantly creating new cognitive niches for us, as those underpinning social media, VR environments, and the “metaverse” are doing right now?⁷

A major problem for a contemporary culture of consciousness is the new “attention extraction economy,” which extracts attention from human brains and turns it into money, supported by self-learning AI and continually self-improving algorithms. In what Tristan Harris has called a “race to the bottom of the brain stem,” social media and tech firms aim to maximize user engagement (chapter 28) by creating ever-better attention sinks and developing pathological, addictive forms of media consumption.

Today, information and entertainment are now available almost everywhere for free. Attention, on the other hand, is the new scarce resource that must now be systematically exploited. Human attention has thus become a currency, a means of payment that can be traded. But what is being monetized in the new attention economy is really the destruction of mental autonomy—that is, our ability to control the focus of our

own attention in a deliberate and self-determined way (see figure 25.1 in chapter 25). We need mental autonomy, however, if we are to maintain and defend our democracies as responsible citizens.

The capacity for mental self-determination is also of crucial importance for our overall quality of life, including our mental health. Today, the question being asked is no longer whether a human or a machine is a world champion at Go or chess. The game that is being played against us right now is a very different one: Who gets to control the scarce resource of attention generated by our biological brains—us or some tech corporation trapped in its own business model? What persuades human beings—which algorithmic strategy generates feelings of intimacy and the phenomenology of trust most efficiently? Questions like these are absolutely central to the culture of consciousness of the future. Questions such as the following come up here:

- What ethical and legal principles should guide us in the attention economy of the future?
- Which parts of the new sociophenomenological infrastructure should be public goods, systematically protected because of their relevance to the common good?
- Are there media environments and business models that should be illegal?
- In an open and free society, should all citizens have the right to manipulate their own brains in whatever way they wish? Or are there states of consciousness that ought to be “off-limits,” for example because the means to achieving them cannot be safely controlled by beings like us, resulting in too high a probability of individual and/or collective suffering?
- How do we get psychoactive substances and AI to systematically increase the mental autonomy of its human users instead of destroying it?
- Which states of consciousness improve empathy and social cohesion?

What States of Consciousness (If Any) Should We Force upon Other Animals?

A growing number of us are beginning to understand that many of our ways of treating other sentient animals are ethically untenable. Animals clearly have a general preference for living as long as possible, having as many offspring as possible, and being as physically comfortable as possible. They do not want to be separated from their children. For many animals, these preferences—which we do not respect—are directly reflected in their conscious self-model; they identify with them. For self-conscious animals, frustrated subjective preferences become their own negative states, which creates the capacity for conscious suffering. Frustrated preferences don’t have to come as thoughts; they can also come as emotions, feelings, or bodily sensations.

Modern consciousness research shows that many nonhuman animals have sentience, especially the ability to feel happiness and suffering. As beings capable of suffering, they are automatically the subject of ethical considerations. We must respect not only reason and the capacity for moral insight in other beings, but also their vulnerability, capacity for suffering, and mortality. Here we may ask the following questions:

- What exactly is the relationship between animal ethics and *Bewusstseinsethik*?
- What are effective strategies for reducing animal suffering?
- Which ways of treating nonhuman animals should be illegal, and how can their interests and preferences be brought into the mainstream of society?
- To arrive at convincing ethical decisions, we need hard facts, including about which animals are likely to have a phenomenal self-model, at least some of the time. How should modern consciousness science inform the practical philosophy of *nonhuman* minds?
- Do sentient nonhuman animals possess meaningfully definable forms of dignity too?

What States of Consciousness—If Any—Can We Force upon Machines or Other Conscious Postbiotic Systems?

You may recall that this book is dedicated to the conscious postbiotic subjects of the future. One reason for this is that AI ethics, and in particular the possibility of synthetic phenomenology and machine consciousness, are linked in important—and not innocuous—ways to the intermediate results of our search for a minimal model explanation of consciousness. There may be an information hazard: If we have a mathematical model of the simplest state of consciousness, then we can implement it on machines or other types of carrier systems. If pure awareness really is the simplest state of consciousness, then it may be the easiest to implement.

When taking an ethical stance toward the conscious postbiotic subjects of the future, the term “postbiotic” refers to the conceptual point that the distinction between “artificial” and “natural” systems is not as clear cut as it often sounds. (A philosopher might point out that it is an outdated distinction because it is neither exhaustive nor exclusive.) We have intelligent systems using biologically evolved algorithms on artificial hardware. The deep neural networks that we simulate in large computers are modeled on biology—namely, on the interconnection of neurons in the nervous system of a living being. Soon, however, we will probably have systems that use algorithms developed by humans but which improve themselves independently, or even entire cognitive architectures realized on biological “hardware” (e.g., on a genetically engineered substrate).

As I have argued elsewhere,⁸ the possibility of artificial consciousness raises a host of ethical issues. We should not recklessly (whether intentionally or not) create artificial

consciousness because artificial consciousness may involve a consciously experienced sense of self that may in turn generate artificial suffering in autonomous, intelligent systems. On the other hand, what makes MPE itself so interesting from an ethical perspective is precisely the fact that it is a form of conscious experience that involves *no* psychological suffering whatsoever (chapter 33). Imagine that we could build a machine that peacefully abides in a clear and silent state of pure consciousness, nondual and without subjectivity, for as long as it exists. If we could create a conscious machine that reliably remains in a single full-absorption episode of pure awareness—would there be anything ethically wrong with creating such a machine? The four questions given here are intended merely as illustrative examples of the many ethically sensitive links between the MPE research program sketched in this book and the problem of machine consciousness:

- Is it ethical to risk the creation of artificial consciousness before we even have a convincing theory of consciousness, let alone of conscious suffering?
- Given the potentially large number of unknown unknowns, should there be a global moratorium on synthetic phenomenology until we know definitively, or at least reasonably accurately, what we are doing?
- Is there a possible information hazard⁹ involved in even trying to describe the simplest form of conscious experience or develop a mathematically precise computational phenomenology of suffering? That is, could a good caricature of the essential statistical physics be enough to conjure up the target phenomenon? Could the physical implementation of a minimal model unexpectedly create the relevant microscale causal mechanisms that are sufficient for consciousness to occur?
- If we eventually arrive at a theory that successfully describes conscious intelligence without suffering—are we then morally obliged to realize such suffering-free forms of conscious experience on artificial carrier systems, perhaps even devoting more resources to developing them than we do to propping up biological forms of intelligence for which suffering is unavoidable?

Which States of Consciousness Do We Want to Foster and Integrate into Our Society?

Bewusstseinskultur and practical philosophy of mind are concerned not only with what we ought not to do, but also with what ought to be promoted proactively—that is, purposefully and with foresight. Here is an example. Many interesting and evidence-based academic initiatives that test ways of building secular meditation practices into education exist worldwide. Against the background of the attention economy and the escalating drug problem, this seems to me to make extraordinary sense, especially for children and young people. This type of endeavor can be linked—as the reports from

fifty-seven countries presented in this book clearly show—to the experiences of millions of adult practitioners and the expertise generated within the many spiritual subcultures or countercultures that began to develop in the West after World War II.

But for every social movement and for every cultural innovation, there is an optimal time window in which its core insights must be formally integrated into broader structures or risk being lost again. Now may be the time for political institutions to finally catch up with social reality, separate the wheat from the chaff, take responsibility, and preserve what is substantial before it gets further watered down by market mechanisms. Alongside education and teaching, there is research itself. Western academia lacks a tradition of scholar–practitioners. We do not systematically combine distinct epistemic practices like science and meditation, so we miss out, for example, on the heuristic fecundity of MPE phenomenology described in this book. Given the preliminary findings from this survey and the research goals listed at the end of chapter 34, should we create such a tradition?

Again, a more systematic exploration of altered states of consciousness will have great heuristic fecundity not only for science, but for society as a whole—and there is a time window that may close for good if we do not manage to preserve what matters before it gets destroyed by McMindfulness, organized crime, and the forces of a globalized black market. We currently have no protected spaces for those who want to devote their whole lives to cultivating elusive states of consciousness, maximally free of ideological distortions.

In this domain, questions worth asking include:

- Could the proper, safe, and beneficial use of psychoactive substances or the new mind-altering technologies mentioned in this discussion generate epistemic practices that are complementary to meditation and science, cultivated by specific subtypes of scholar–practitioner or phenomenological polymath?¹⁰
- Given the findings of modern meditation research and the phenomenological data presented in this book, should we integrate secular forms of meditation training into any or all levels of our education systems?
- Could something like a secular monastic tradition be created, one much purer than what we had before because it sets aside mortality denial and narrative self-deception (chapter 17)?

What States of Consciousness Do We Want to Show Our Children?

Education is a process by which young human beings become active members of society and culture, and thus also politically mature citizens. The right to education is enshrined in article 26 of the Universal Declaration of Human Rights: Education shall

be free, at least in the elementary and fundamental stages (§ 1), and it shall be directed to the full development of the human personality (§ 2). In the educational context, questions that we would do well to ask and answer include the following:

- What regions of phenomenal state space should *every* human being get to know before becoming an adult?
- Against the backdrop of the escalating drug problem and the burgeoning attention economy, and given the promising results of modern meditation research, should we integrate secular forms of meditation instruction into some (or even all) levels of our educational system?
- Should we offer to the young a “contemplative toolbox” that enables them to systematically explore these regions, whenever they may be interested or have a specific need for them later in life?
- The participants in our study used and combined many meditation techniques as part of their own development (see figure 14.1 in chapter 14). Which of these techniques comprise the essential set that every child should get to know? Which specific meditation techniques should children learn, and during which time window of their psychological development?
- More generally, what would we consider the “indispensable core of consciousness-culture education” that should be imparted to a society? What should our future “phenomenological education canon” look like? What should be the basic set of conscious states that must be freely available to every human being to come into contact with at least once in their life?

What States of Consciousness Do We Want to Die In?

Bewusstseinskultur will obviously have to include a new culture of dying. Existence bias and the resulting culture of mortality denial are major determinants of human phenomenology (chapter 17). Many modern societies have systematically repressed death and dying, turning them into a sensitive and taboo topic. *Bewusstseinskultur* will have to break this taboo because it must pose the ancient philosophical question of what it really means to live a good life, which naturally leads us to ask what precisely a “good death” could be. More specifically, we might ask the following:

- From a rational, evidence-based, and ethically motivated perspective, are some ways of dying better than others?
- Should we create protected spaces in which some aspects of the dying process can be practiced safely, for instance with the help of meditation or certain psychoactive substances?

- Would it be useful if one could learn about the phenomenology of peaceful ego dissolution long before physical death occurs?
- What state or mode of consciousness would *you* like to die in?

These are just a few examples to give you a first impression of the wide variety of concrete issues that need to be addressed; I have written more elsewhere, so I will not repeat myself here.¹¹ Nevertheless, I hope this small selection shows how many aspects of our societal life would be directly affected if we ever decided to actually implement the provisional results yielded by a *Bewusstseinsethik*, a systematic practical ethics of consciousness. Just imagine what might happen to our society if we took the systematic cultivation of valuable states of consciousness seriously.

Enculturation

The third major aspect of the idea of *Bewusstseinskultur* relates to the creation of a new cultural context. *Kultur* does not just mean that individual citizens cultivate certain states of conscious experience rather than others; it also means organizing the social fabric in a way that supports individuals and groups in systematically investigating the types of question listed in this epilogue, exploring the landscape of possibilities for genuine cultural innovation. *Bewusstseinskultur*, therefore, refers to a new kind of sensitivity that creates a cultural context in which this very process of inquiry, investigation, and experimentation is itself valued—as a shared epistemic practice, as a pro-social activity of cultivating one’s own mind while exploring the possibilities of a new normative context for all, and as something that makes a contribution by proactively creating a new kind of common good.

Bewusstseinskultur is a common good, something that we can all share and benefit from. But at its core, it is also an emancipatory project, for it aims to increase our own intellectual autonomy. The techniques underpinning it should never be commodified and turned into a private property, exploited primarily by corporations. Indeed, its essence is fundamentally anticorporate. As an antiauthoritarian, decentralized, egalitarian, and participatory project, *Bewusstseinskultur* is essentially based on community, cooperation, and transparency, and it rejects the logic of consumerist capitalism. If you will, it is a kind of mental infrastructure: In the future, we will need not only clean air and water, schools, museums, public transportation systems, civil liberties, and public safety, but also a continuous process of rational, evidence-based enculturation of those states of consciousness that we have found to be valuable. Having such a process is *itself* a common good—an asset that a given society may either possess or not.

Bewusstseinskultur is something that comes in degrees and grows over time. There is a global context to this: One society may be more advanced in the ethical treatment of animals; another may have a better educational strategy, and yet another the most successful drug policy. We must learn from each other. Here, “enculturation” refers to a continuous process by which ancient practical wisdom, new mind-altering technologies, and all relevant scientific insights are incrementally embedded into societal practice. As I pointed out in the introduction, the problem of pure consciousness will have to be handed over to the hard sciences of the mind, to cognitive neuroscience and computational modeling, but from there, it will later have to return to philosophy of mind and applied ethics. New scientific insights always need conceptual interpretation and an ethical assessment of any new potentials for action that they may generate. In formal research and other areas as well, there is a need for intercultural learning. Helping to construct the new cultural context that allows such a continuous embedding on a global scale is itself a prosocial activity, just as cultivating one’s own mind through sustained meditation practice is. Therefore, the third major reading of *Bewusstseinskultur* refers to the proactive development of a shared sociocultural context that prioritizes the common good.

The Wider Context: Secular Spirituality and the Planetary Crisis

As I said in the introduction, if anything was ever a “big-picture issue,” then pure consciousness is. There is a quality of profundity in some of the reported experiences that directly relates to many of the deepest philosophical puzzles. Quite obviously, there is a much wider context to all of this. It goes far beyond applied ethics, and we all are living through a very special historical epoch. To get a fuller picture, it may be helpful to sketch out two of the most important dimensions of the new context that would be created by the project of a *Bewusstseinskultur*.

Bewusstseinskultur has a lot to do with finally taking responsibility for your own life and with what I would like to term “the principle of self-respect.” There is a problem to be solved: How does one preserve one’s self-respect during a historical epoch in which humankind as a whole is losing its dignity? Humanity is in the midst of a planetary crisis that is self-inflicted and historically unprecedented. It does not look good. Both political institutions and large numbers of individuals around the world are failing miserably in their management of this crisis, and they have been failing miserably for a very long time. Most people are beginning to feel that something has changed: When it comes to the climate catastrophe, it is no longer intellectually honest to be an optimist. Many of us also feel that we have been lying to ourselves for some time now.

Yes, of course, from the perspective of climate science, it is still possible for warming to stay below 1.5 or maybe 2 degrees Celsius. But from the perspective of psychology and political science, it isn't.

If we face up to the psychological and political facts available today in an unbiased way, everything points to the conclusion that humanity will fail in the face of this problem, and it will fail with its eyes open. We have not managed the transition from a growth-oriented, greed- and competition-based model of economics into a sustainable form of steady-state economics. We knew it all long ago, we decided to ignore the facts for as long as possible, and we successfully organized our own self-deception on a political level. Now the greed-based growth model has finally led us into an accelerating environmental catastrophe.

The central factor in this (which everyone also knows about) is a close correlation between economic growth and the rate of global carbon dioxide emissions. Economic growth is the dominant normative and cultural context in which we currently live our lives, and, as societies, we have not managed to evolve an alternative context that is more attractive than the neoliberal growth model. Few of us know the unboundedness of pure awareness discussed in chapters 2 and 23, but the delusional belief in unbounded economical growth still has a stranglehold on many.

If we combined gross domestic product with the general level of *Bewusstseinskultur* as a standard prosperity indicator that goes beyond mere happiness measures and subjective life satisfaction, we could create more socially just and ecologically sustainable societies. It is easy to see how we could have chosen not to aim solely at maximizing the mean standard of living, but also to value something that, in free and open discussions, we had previously defined as the "mean standard of consciousness culture." To give one example in line with the central topic of this book—what if one of our central goals had been maximizing the minutes of pure awareness per capita? What if we had valued and systematically increased the average lifetime hours that every citizen can spend enjoying MPE states, or even MPE modes? What if our goal had not been to land on Mars, but to land in pure awareness?

Given the phenomenological material presented in this book, one may plausibly assume that many of us would have had much richer lives on an individual level, and that we would have stopped the planetary crisis before the chance to do so slipped through our fingers. But everything points to the prediction that it is now too late for this. Very soon, it will no longer be possible to respect humankind's behavior because this behavior does not change even when human beings clearly see, at the level of their own conscious self-experience, that we have done and continue to do all this knowingly.

Our species does not respect humanity as a whole, neither in the present nor in the future. Collectively, we lack the ethical integrity, the quality of compassion, and the capacity for rational action that would have enabled us to avert medium- and long-term catastrophe at moderate medium- and short-term cost. We are causing an enormous amount of future suffering, and we are doing so knowingly. Very soon, therefore, it will no longer be possible to respect the behavior of large segments of humanity. We will no longer be able to take ourselves seriously, for our behavior does not change even when we clearly recognize that it must.

This is where the idea of a new, secular form of spiritual practice comes back into play. There actually *is* something that we can still respect in ourselves and others: the nonegoic form of conscious self-knowledge that gradually reveals itself in contemplative practice (chapters 29 and 30). This form of conscious self-knowledge is much deeper than the high-level dignity that rational subjects may (or may not) see and acknowledge in each other, because it creates a much more fundamental relationship between the individual and the community of all sentient beings capable of experiencing it. It is nonconceptual. It is ownerless. It has nothing to do with words or thoughts. Animals may have it. Perhaps future machines should have *only* this form of self-awareness. The capacity for conscious but nonegoic self-knowledge is something that we can respect in ourselves and others, even if humanity as a whole fails. It has worth, and it can be valued—but first, it needs to be recognized. We have to rediscover it.

Amid the rolling catastrophe, we must somehow learn to do the right thing just because it *is* the right thing, without needing to obtain positive results for ourselves. In experiencing our own failure and inadequacy, we must somehow also learn to do the right thing in a sustainable way that involves a quality of self-compassion. Interestingly, both self-respect and the capacity to fail gracefully relate directly to the phenomenology of nonidentification that we have encountered and investigated at many different places on our journey through this book (e.g., in chapters 8, 19, 29, and 33) and to the capacity to act, at least sometimes, from a position of nonegoic self-awareness (see, e.g., chapters 29–33).

To a degree, both self-respect and graceful failure can be systematically cultivated. As we have seen, the quality of nonegoic self-awareness may be closest to what makes many describe self-knowing MPE as an irreducibly *spiritual* state of consciousness. It is possibly what characterizes the deepest sense of the word “spirituality”: a form of conscious experience that is completely independent of religious belief systems. Within the space of family resemblances related to the experience of pure awareness, reflexive MPE is a prototypical core region that forms the phenomenological anchor for what many meditators, if asked, may try to describe as the spiritual “essence” of their

experience—or even as their “true self” (chapter 29). As we have seen in this book, the recognition of reflexive meta-awareness via meditative practice arises on an entirely nonconceptual and even nonegoic level. There is more than one form of self-respect, more than one form of self-knowledge, and therefore more than one form of dignity. We could even coin a fancy term for this if we wanted to, like “nonegoic dignity.”

At least in principle, discovering nonegoic dignity in themselves can help people to do what is right in a way that is not emotionally attached to achieving positive results or getting the rewarding experience of egoic self-efficacy. At least to a certain degree, therefore, contemplative practice may prevent bitterness and emotional burn-out by helping to prevent ego-defensive anger, self-condemnation, constant depressive rumination, and generalized anxiety. Discovering the silkiness of self-cognizing silence is not something that demands a special form of giftedness or an unusual intellectual capacity; it could form the experiential kernel around which the requisite cultural growth happens at scale. There may actually be a way to preserve one’s self-respect—including one’s mental health—at a time when humanity as a whole is losing its dignity. I think it could be rewarding to think further in this direction.

The wider context is deep and rich. *Bewusstseinskultur* and its role in proactive damage control and the preservation of dignity form one dimension; secular spirituality forms another. We have already seen that *Bewusstseinskultur* is not a new form of religion, but it is closely linked to the deeper philosophical question of whether a fully secularized form of spirituality is at all conceivable. What about meditation as a more substantial, genuinely epistemic practice? Can there be a new form of spirituality that explores the epistemic dimension of nonegoic self-awareness shining through in some of our reports, a slightly more radical culture of consciousness that remains entirely free of irrationality, a set of practices not driven by mortality denial and narrative self-deception (as described in chapter 17)?

Taking a closer look at the quest for a secularized form of spirituality will give us a fuller picture of what *Bewusstseinskultur* really is. To begin with, one problem is this: Millions of meditators worldwide like to call themselves “spiritual but not religious (SBNR)” or “spiritual, but not affiliated (SBNA)” —but nobody knows what these concepts actually mean.¹² Yes, there has been a successful rebellion against organized religion and a fresh global counterculture has emerged. Millions have taken things into their own hands, and this can be seen as one of the major cultural innovations of the twentieth century. But if one looks at the reality of all the different movements, traditions, and groups that have formed—especially their irrational belief systems and the resulting social dynamics, which are often pathological and cultlike—one may conclude that the original innovative impulse has long been outweighed by more mortality denial

and self-deception, just in slightly subtler forms than they used to take. Many of the spiritual movements that have developed in recent decades in Europe and the US have long lost their progressive impulse, and most sectors of spiritual counterculture seem already to have morphed into experience-based forms of privately organized religious delusion. Today, they merely stabilize or conserve the status quo and are characterized by a slightly infantile form of complacency, and often by crude forms of intellectual dishonesty. As I very briefly indicated in chapter 15, if viewed from an empirical perspective, we therefore have grounds to wonder whether the people labeling themselves as “SBNR” or “SBNA” are religious after all, just in a slightly new way.

In most of these contemporary forms of “nonreligious spirituality,” something is sorely missing: a particular kind of honesty. After all, the opposite of religion is not science, but spirituality.¹³ The ethical principle of intellectual honesty can be analyzed as a special case of the spiritual stance and, in their purest form, the scientific and the spiritual stances emerge from the same basic normative idea. But what exactly is “intellectual honesty”?

Intellectual Honesty

“Intellectual honesty” means simply not being willing to lie to oneself. It is closely related to old-fashioned values such as propriety, integrity, and sincerity, as well as to a certain kind of “inner decency.” Perhaps one could say that striving for intellectual honesty is a very conservative way of being truly subversive. Intellectual honesty might be exactly what the many “teachers” and representatives of cults, organized religions, and theologians of almost any type simply cannot have—even if they often like to make claims to the contrary. Intellectual honesty means not pretending to know or even pretending to be *able* to know the unknowable, while still having an unconditional will to truth and a commitment to the growth of knowledge, which includes being genuinely open to new scientific results. Intellectual honesty is the missing element in many of the social movements whose members think of themselves as “SBNR.” It is needed even (or especially) where self-knowledge is involved, and even where self-knowledge is not accompanied by pleasant feelings or is not in accordance with the doctrine in force at the time.

The sincere pursuit of *intellectual* integrity can be viewed as an important special case of the pursuit of *moral* integrity, and it turns out that a relationship also exists between these forms of integrity and contemplative practice. Philosophers sometimes call the type of approach that emphasizes such connections an “ethics of belief.”¹⁴ Correspondingly, *Bewusstseinskultur* also demands that we think clearly about the conditions under

which it might be not only *irrational* or *imprudent* but also *morally wrong* to hold a belief on insufficient evidence. A philosophical ethics of belief may be another element missing from large parts of the spiritual counterculture.

A flood of empirical data, as well as the mathematical models of modern computational neuroscience, show how the beliefs we hold directly influence the way that we consciously experience the world—a simple empirical fact whose importance is hard to overstate when it comes to meditation practice. Whoever wants to become whole—a person of integrity—by gradually resolving all conflict between their actions and values must pursue this harmony with their *inner* actions as well. This requirement is especially true for their “epistemic actions”—actions for the sake of knowledge. We act epistemically whenever we strive for insight, knowledge, true belief, sincerity, and also authentic self-knowledge. To the extent that meditation is an epistemic practice, it cannot work without radical honesty toward oneself, without a self-critical ethics of belief. There is a bridge between spiritual practice and the ideal of reasonable, rational thought: Both involve an ethics of inner action for the sake of knowledge. Moreover, in both cases, the goal is a systematic enhancement of mental autonomy, of inner freedom.

Immanuel Kant put this point in a completely different but particularly beautiful way. What is needed, he says, is *the sincere intention of being honest towards oneself*.¹⁵ The “sincerity” or “purity” of the desire for honesty toward oneself is, I think, the central point—and it may actually have something to do with discovering pure awareness. Without the will and the courage to really *look*, pure consciousness does not come into existence. The sincere intention of being honest toward oneself is what connects critical rationality and the nonconceptual insight of contemplative practice.

What Kant did not see is that exactly the same principle is also crucial for a somewhat radical, philosophically motivated meditation practice. This form of meditation involves a nonconceptual form of mental action for the sake of knowledge—an action that marks the beginning of the search for a nondual state of consciousness beyond the subject/object split. Without the intention to be radically honest with oneself, dissolution of the phenomenal self cannot take place. Kant’s purity of intention is the beginning of an inner movement toward a clear and effortless form of mental *inaction*.

But you cannot choicelessly observe your own thoughts as they arise and disappear again if you are not prepared to honestly face what you will now begin to see: the painful restlessness of your very own mind, your violent fantasies, your desire for retaliation, your boredom, your loneliness, your existential despair, or your envy. At the beginning of this section, I offered one example of what *Bewusstseinskultur* could refer to: Cicero’s idea that truly loving wisdom, and thus being a philosopher, mean “taking care of and cultivating one’s soul.” Now we can view Kant’s point as a second example:

In terms of *Bewusstseinskultur*, any genuine contemplative practice requires the sincere intention of being honest toward oneself.

The importance of intellectual honesty also provides another bridge connecting a secular, self-critical form of spiritual practice and the ideal of rational thought, this time on an eminently practical level. I hope that you are beginning to sense that a strict and altogether old-fashioned form of rationalism could have a lot to do with spirituality—and with what is still missing in large parts of the spiritual counterculture. To stay with Kant for a moment: For him, dishonesty is nothing other than a form of unconsciousness, a lack of awareness and conscientiousness. Here, my point is that rationality is actually a special form of mindfulness. In order not only to think clearly but also to listen properly—in order to *understand* one's opponent—critical rationality first needs the sincere and nonreactive quality of mindful attention. And then it also needs careful, rational thinking, a way of handling concepts that is as precise as possible.

The essence of being honest toward oneself while listening to others is the form of reflexively aware mindfulness that Dakpo Tashi Namgyal and others¹⁶ spoke of many centuries ago (chapter 30). There is a deep connection to the Western ideal of critical rationality here, because the quality of reflexively aware mindfulness is what really creates open-mindedness—and it is something that can be discovered and trained. Epistemic openness (chapters 4 and 5) is a precondition of genuinely rational thought, and it also has a social dimension. Our empirical data clearly show that MPE also has a social dimension; this primordial and entirely nonintellectual variety of open-mindedness creates an experience of connectedness because there is an element of compassion in it (chapter 11). Such connectedness, in turn, is a basic condition for genuine solidarity. Of course, much more could be said here, but my second point can be put very simply: *Real* meditation practice cultivates the conditions of possibility for rational discourse by cultivating the sincere intention of being honest toward oneself. If it doesn't, there is something wrong.

Spirituality, Intellectual Honesty, and *Bewusstseinskultur*

We have now briefly looked at what *Bewusstseinskultur* is, and at some aspects of the broader context in which it could come into being. As it turns out, *Bewusstseinskultur* is a slightly subversive but prosocial way of taking an ethical stance toward one's own mental states. As I have suggested, it is a “practical philosophy of mind.” *Bewusstseinskultur* involves the systematic cultivation of states assessed as valuable, aiming at a continuous process of rational, evidence-based enculturation. It is a project—something that we need, and that we certainly don't have yet.

In this book, we have looked at the experience of pure awareness carefully and from many angles. We saw that pure awareness may turn out to be the “convergence zone” where spiritual practice, cognitive neuroscience, and modern philosophy of mind finally come together. At its prototypical core, pure awareness is not a subjective phenomenon because it is not necessarily tied to a knowing self or a consciously experienced first-person perspective. But in the course of our journey, it has also become clear that what philosophers, scientists, and the more than 500 experiential reports that I have shared with you are trying to describe is not a single and unitary phenomenon. It is not some metaphysical essence. Rather, it is a specific region in phenomenal state space: The space of possible conscious experiences. Pure awareness has many aspects and dimensions. They have been ignored for too long. This region in phenomenal state space could be of interest, certainly to philosophers, computational modelers, and neuroscientists working on consciousness—but maybe also to all of us who are interested in the project of developing a *Bewusstseinskultur* that is fit for the future.

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>