36

Life of science, life of faith

WILLIAM T. NEWSOME

Like Charles Townes, I am a practicing scientist and a practicing Christian – occupations that are antithetical in the minds of many. In common caricature, the practice of science is frequently portrayed as objective, comprehensive, and truthful in contrast to religious practice, which is frequently perceived as subjective, parochial, and superstitious. In 1966, Charles Townes took this caricature to task in an insightful article in IBM's *Think* magazine (Townes, 1966). There he argued that science and religion share many common characteristics and are, in fact, destined to converge in the long run. Townes pointed out striking commonalities between scientific and religious practice, including the centrality of faith, the essential importance of intuition ("revelation"), the acceptance of paradox, the testing of working hypotheses against experience, and the provisional nature of knowledge (Townes, 2006, pp. 28–43). He further argued that the differences between science and religion are not qualitative, but rather a matter of degree resulting from the far more complex subject matter that is the focus of religious inquiry (Townes, 2003, pp. 154–8).[1]

Charles Townes's efforts to synthesize his scientific and religious practices in a thoughtful, constructive manner were *timely* in 1966, but such efforts are in fact *urgent* today. The polarization of "scientific" and "religious" worldviews in the current century, driven by immodest agendas on both sides, fuels dangerous conflicts within and between cultures. This does not *have* to be the case. My own view, like Townes's, is that both science and faith contribute importantly to a meaningful, fully experienced human life. Giving up either would result in a regrettable loss of understanding, depth of experience, and simple joy. I am convinced, then, that much of the perceived incompatibility between science and religion is specious, although real tensions do exist. My purpose in this paper is to lay out the central issues from my point of view – both the real and the false sources of tension between science and religious faith as I have experienced them.

36.1 Religion and the *findings* of science

Specious conflicts between science and religion stem from the perception that the discoveries of modern science have rendered traditional religious beliefs untenable. I would argue

[1] Also see Townes's essay on free will in this volume.

Visions of Discovery: New Light on Physics, Cosmology, and Consciousness, ed. R.Y. Chiao, M.L. Cohen, A.J. Leggett, W.D. Phillips, and C.L. Harper, Jr. Published by Cambridge University Press.

that just the opposite is true – that the major discoveries of modern science are remarkably compatible with the central religious insights of the monotheistic traditions.[2] The creation stories in Genesis, for example, assert that God created the universe, that the emergence of humankind in God's image was an intentional result of this creative act, that a nurturing relationship with God is the ultimate aim and reward of our existence, and that we humans have a terrible freedom to enter into or reject this relationship as we see fit.

Although the creation stories are presented in poetic, mythical images, their central religious insights are not in conflict with any specific finding of modern cosmology, physics, or biology. The assertion of a central creative act, for example, resonates nicely with the Big Bang, the prevailing scientific theory of the origin of our universe. According to this theory, for which there is considerable empirical evidence, our universe – and space and time themselves – began in a primordial explosion of energy that occurred at a precise moment. Remarkably, cosmologists can plot the initial sequence of events associated with the Big Bang on a second-by-second basis, yet there appear to be impenetrable barriers to understanding what, if anything, existed "before" the Big Bang. In other words, our universe had a definable beginning roughly 15 billion years ago, and a dense curtain of mystery veils anything before this moment of origin. This basic picture, which is the result of dazzling theoretical insights and empirical measurements of modern astrophysics, would not be at all alien to the writers of Genesis. Robert Jastrow, prominent astronomer and former head of NASA's Goddard Institute for Space Studies, cogently summarized this ironic turn of events in his short book, *God and the Astronomers*:

> For the scientist who has lived by his faith in the power of reason, the story ends like a bad dream. He has scaled the mountains of ignorance; he is about to conquer the highest peak; as he pulls himself over the final rock, he is greeted by a band of theologians who have been sitting there for centuries. (Jastrow, 1978, p. 116)

In terms of universal origins, then, one could scarcely imagine a scientific theory more compatible with the core beliefs of the biblical writers.

Similarly, an increasing cadre of scientists is recognizing that the laws and constants that make up the fundamental physical reality of our universe are improbably hospitable to the emergence of life. These critical constants determine, among other things, the rate of expansion of the universe, the strength of interactions of subatomic particles within the nucleus, and the unique chemical-bonding properties of carbon, oxygen, nitrogen, and hydrogen – the fundamental atomic constituents of organic life. If any of these physical constants had been different by an infinitesimally small amount, the emergence of life in our universe would have been impossible. From such scientific observations has emerged the notion of an "anthropic principle," which asserts that the fundamental nature of our physical universe is peculiarly well suited to the emergence of intelligent life (Carr and Rees, 1979, pp. 605–12; Barrow and Tipler, 1986). As Stephen Hawking put it,

[2] I omit mention of other religious traditions here not out of prejudice, but out of ignorance. The three "Abrahamic" traditions – Judaism, Christianity, and Islam – are the ones that I am actually qualified to comment on.

Nevertheless, it seems clear that there are relatively few ranges of values for the numbers that would allow the development of any form of intelligent life. Most sets of values would give rise to universes that, although they might be very beautiful, would contain no one able to wonder at that beauty. (Hawking, 1988, p. 125)

Of course, these intriguing observations do not *prove* anything one way or the other, religiously speaking. The fact that our universe is improbably hospitable to intelligent life might be sheer coincidence, or there might be some explanation that we cannot comprehend at this point. Nevertheless, many thoughtful observers find these scientific observations to be religiously provocative. Hawking again:

The odds against a universe like ours emerging out of something like the Big Bang are enormous. I think there clearly are religious implications. (Boslough, 1985, p. 121)

My central point here is simply that the findings of modern science are quite compatible with a traditional religious viewpoint, reasonably interpreted.[3] To be crystal clear, I am *not* saying that such considerations *prove* anything about the existence of God one way or the other. I am making a much more modest point – that the results of contemporary science do not, by any stretch of the imagination, render religious faith untenable intellectually, despite a relentless litany of claims to the contrary. The writers of Genesis, given sufficient time to catch up scientifically (!), would find contemporary cosmology quite congenial to their central religious insights.

The scientific discovery that has proven most contentious in certain religious and scientific circles during the past century has been Darwin's theory of evolution by natural selection. At first glance, the central Darwinian vision of the gradual evolution of life (through the interaction of random biological variation, selective environmental pressure, and sexual selection) appears to be at odds with the biblical assertion that God created human life intentionally in His image. How can a process that depends on chance and is fundamentally unpredictable be an intentional, creative act of God?

I would argue, first, that dependence of a process on random events does not speak to the matter of intentional creation one way or the other. The use of random events is sufficiently important in modern scientific investigation that the design of computer algorithms for generating random numbers has become a high art. Random, or probabilistic, events are intentionally harnessed for scientific purposes in innumerable contexts, including neural-network design (i.e., Boltzmann machines[4]), chemical engineering (directed evolution of enzymes[5]), experimental psychology, and quantum computing, among many others. The field of genetic programming provides a particularly striking example (Koza, 1992). In this approach, many new variants of an existing computer program are created by probabilistic recombination of segments of computer code, and the new programs are

[3] This obviously excludes fundamentalist interpretations of a literal six-day creation, a "young" Earth, and so forth.

[4] For more on Boltzmann machines, see http://en.wikipedia.org/wiki/Boltzmann_machine. The article is good even though it is Wikipedia!

[5] See http://cheme.che.caltech.edu/groups/fha/.

evaluated by a "selection" mechanism that is related to each program's "fitness" for solving a specific task. Programs created by this method, which was directly inspired by biological mutation, recombination, and selection, are now competitive with programs created by traditional methods.[6] Even a casual understanding of the scientific landscape reveals that random events are used over and over again for purposeful ends. If human scientists can do this, why can't God? My point, of course, is that no deep contradiction exists between the evolutionary mechanism of chance mutation (coupled with selection) and the religious notion that God intentionally created life – and, among the varied forms of life, us.

Second, we must be clear concerning what, exactly, is "unpredictable" about evolution. What biologists generally mean by this is that evolution is, in the words of Stephen Jay Gould, highly "contingent." A specific random mutation in one organism in a particular environmental context can have a major impact on survival and, thus, on the future development of entire ecosystems. The identical mutation occurring in an organism of a different species, or in a different environmental context, may have little or no impact on survival, reproduction, and the ensuing ecosystem. The potential interactions among chance mutation, environmental pressure, and individual survival and reproduction are so numerous and complex as to constitute a system in which future states are impossible to predict. Gould argues that contingent events exert such an enormous effect in evolution that, if the history of the Earth could be rewound to a point, say, three billion years ago, and played out all over again, it is grossly unlikely that a creature exactly like *Homo sapiens* would emerge – a predator with frontally directed eyes, bilaterally symmetric body plan, and a central nervous system organized on the current mammalian configuration (Gould, 1989).

Two aspects of Gould's argument deserve comment. First, the scientific evidence itself provides grounds for doubting the argument. The counterargument has been made particularly forcefully by the Cambridge paleontologist Simon Conway Morris, who changed his original views, which were similar to Gould's, after life-long study of the fossil record of the Cambrian explosion (Conway Morris, 1998). Succinctly, Conway Morris is far more impressed by the "convergence" that occurs within evolutionary history than he is by "contingency." He argues that certain body plans and adaptive features (particular points in the entire space of possible animal morphologies) recur independently in evolution with sufficient frequency that they must be regarded as uniquely adaptive to life on this planet. For example, Conway Morris argues that dolphins, which evolved from dog-like mammals, are shaped similarly to fish because there is an optimal shape and strategy for moving through water.[7] Thus, he argues, if the tape of evolution were to be rewound and allowed to play out again, it is likely that these evolutionary "solutions" to the challenge of living on Earth, or very similar ones, would emerge once again (Conway Morris, 2003). If Conway Morris

[6] See http://en.wikipedia.org/wiki/Genetic_Programming.

[7] See the fascinating exchange between Conway Morris and Gould, "Showdown on the Burgess Shale," available at www.stephenjaygould.org/library/naturalhistory_cambrian.html.

is right – and he presents very compelling evidence in favor of his view – our Earth was, in a real sense, "pregnant" with certain basic life forms from its very inception.

Because the critical scientific experiments that could test this hypothesis are impossible to perform, this assertion cannot be proven empirically. From a religious point of view, however, arguments about the details of physical morphology are not critical; it is more important to consider what it means to be created in the "image of God." Is God a visually directed predator with a bilaterally symmetric body plan and a mammalian central nervous system? I doubt it. Rather, the religious insight of Genesis is directed toward the emergence of a creature with intelligence, with sensitivity to right and wrong, and with the freedom to choose between them. Would such a creature likely re-emerge if the Earth's history were rewound by three billion years, even if this creature's physical appearance were extremely different from that of *Homo sapiens*? I, along with Conway Morris, think that the answer is "yes." My view on this matter might be chalked up to religiously motivated wishful thinking, but, ironically, I claim as my ally the noted evolutionary biologist and theorist Richard Dawkins, who is certainly no friend to religion. In his book *The Blind Watchmaker*, Dawkins states

> My personal feeling is that once cumulative selection has got itself properly started, we need to postulate only a relatively small amount of luck in the subsequent evolution of life and intelligence. Cumulative selection, once it has begun, seems to me powerful enough to make the evolution of intelligence probable, if not inevitable. (Dawkins, 1987, p. 146)

These are strong words, but my gut feeling is that Conway Morris and Dawkins, who appear to be in essential agreement on this point, are correct. The selective advantages of advanced intelligence are so vast that its emergence in this particular universe, which itself appears uniquely hospitable to life, may indeed have been inevitable once the evolutionary process was started. In this sense, then, the emergence of intelligent, morally responsive life can reasonably be thought to have been an integral feature of our universe from its inception. Certainly no scientific findings argue compellingly against this point of view. For those of us who take both our science and our religion very seriously, then, even the theory of evolution resonates powerfully with the core religious insights of the writers of Genesis.

As Kenneth Miller has eloquently argued in his recent book, evolution, properly understood, is no enemy of religion (Miller, 1999). Despite the continued objections of a vocal minority, most Christians do not see evolution as a major point of dispute between science and religion. As Richard Dawkins has observed, the emergence of the theory of evolution in the nineteenth century "made it possible to be an intellectually fulfilled atheist" (Dawkins, 1987, p. 6). It had relatively little effect, I think, on the possibility of being an intellectually fulfilled theist.

To summarize, the actual findings of modern science are notably congenial to traditional religious belief: a universe with a well-defined beginning that is, against all odds, favorable to the emergence of life and an evolutionary process that may well have favored the emergence of intelligent, morally sensitive beings. While these observations and reflections *prove* nothing, they are entirely consistent with religious beliefs. No scientific result makes it unreasonable to believe that the universe is our "home," in a profoundly meaningful sense of the word, and that in some real way our existence was anticipated from the beginning of it all.[8]

36.2 Religion and the *assumptions* of science

In contrast to the *findings* of science, certain *assumptions* frequently associated with science can cause genuine tension with religion. The core assumption – or faith – underlying natural science is that the universe operates according to orderly, reliable mechanisms rooted in physical "laws" that can be discovered and described by humans. The famous "scientific method," which was born of this foundational assumption, typically involves repetitive testing of specific measurable phenomena, tweaking the conditions this way and that in each repetition, to gain insight into the physical mechanisms that mediate each phenomenon.[9] When science operates at its best, the knowledge derived from the artful combination of theory and experiment is genuinely universal. The basic observations and the theory that ties them together are accessible to, and can be confirmed by, any scientist, anywhere in the world, given the proper equipment and technical expertise. The scientific enterprise has been extraordinarily successful at understanding and gaining control over the physical world, as the history of the past four centuries amply demonstrates. By any account, natural science – both the process and the body of results – is one of the most brilliant achievements in the history of our species.

From a religious point of view, both the core assumption of natural science and the resulting method are fine as far as they go. Conflict arises when additional, extra-scientific assumptions are introduced, the most immodest of which is that the only real phenomena in our universe are those that are susceptible to study by the scientific method. A common corollary of this materialist assumption is that the scientific method provides the only secure path to truth that is meaningful and universal. These materialist assumptions are, of course, fundamentally incompatible with most forms of religious belief and practice, dismissing in one fell swoop the notion of a supreme being establishing a universal grounding for personal meaning and right action.

It is essential to realize that these radically materialist assumptions, which are critical factors in so many clashes between science and religion, are extra-scientific. They are not *findings* of science (try to locate a scientific study that proves these assumptions!), nor are they *logically necessary* to the scientific process (many excellent scientists do not

[8] In making this statement, I do not wish to be gratuitously anthropocentric. To say that the universe is our "home" is not to exclude that it may be "home" in a similarly sacred sense to other intelligent beings elsewhere in the universe, or perhaps to nonhuman animals on earth.

[9] Note, however, that the wellspring of scientific innovation lies not in the scientific method, but in the intensely personal realm of human genius and imagination. As noted by Charles Townes, Michael Polanyi, and Ian Barbour, among a host of others, genuine scientific inspiration – and the faith to pursue that inspiration through years of frustrating inquiry – bears much more than a passing resemblance to religious inspiration. The scientific method provides a reliable way to test, elaborate, and apply scientific ideas, but the method itself does not beget scientific creativity. For reflections in this vein on his discovery of the maser and the laser, see Townes (2003). For a lucid introduction to Polanyi, see Gelwick (1977). Ian Barbour's thoughtful comparison of scientific and religious inquiry is nicely summarized in Barbour (1966).

share these assumptions). Rather, these extra-scientific assumptions form the basis of an ideological position that some academics adopt for their own personal reasons. This, in itself, is fair enough. Everyone, after all, must view the world through interpretive lenses that conform to their own experience, reflection, and conscience. Danger arises, however, when this materialist ideology is packaged for public consumption, either implicitly or explicitly, as part and parcel of science itself. This is certainly not good science, and I doubt that it is good philosophy either. Yet this uncritical and sometimes unconscious marriage of science with materialist ideology pervades the scientific community, appearing frequently in classroom teaching and in the public commentaries of many scientists. For example, William Provine, biologist and historian of science, has said that

> Modern science directly implies that there are no inherent moral or ethical laws, no absolute guiding principles for human society . . . There is no way that the evolutionary process as currently conceived can produce a being that is truly free to make moral choices. (Provine, 1988, pp. 25–9)

To my mind, modern science implies no such thing. Provine's point of view is critically dependent on additional materialist assumptions that are not a part of science itself.

In a similar vein, Richard Dawkins remarks that,

> In a universe of physical forces and genetic replication, some people are going to get hurt, other people are going to get lucky, and you won't find any rhyme or reason in it, nor any justice. The universe that we observe has precisely the properties we should expect if there is, at bottom, no design, no purpose, no evil and no good, nothing but blind, pitiless indifference. (Dawkins, 1975, pp. 132–3)

In a passage remarkable for its anti-religious zealotry,[10] even in the context of contemporary "scientific" writing, Richard Lewontin, a geneticist, states that the primary goal of scientists in communicating with the public is

> . . . to get them to reject irrational and supernatural explanations of the world, the demons that exist only in their imaginations, and to accept a social and intellectual apparatus, *Science*, as the only begetter of truth . . . We take the side of science . . . because we have a prior commitment, a commitment to materialism . . . Moreover, that materialism is absolute, for we cannot allow a Divine Foot in the door. (Lewontin, 1997)[11]

Each of these views, presented by its author as science or the logical consequence of science, in fact depends for its credibility on specific extra-scientific assumptions that make up a personal ideology. If one is skeptical of the prior assumptions adopted by the authors, each of the statements loses much of its logical force.

[10] Persons unfamiliar with academic sensibilities should understand that there *are* legitimate historical reasons for the knee-jerk hostility of many academics to the influence of organized religious institutions. For long stretches of the modern period, academic inquiry was held in thrall to ecclesiastical authority. The freedom to inquire, to think and argue freely, to teach to the best of one's lights was often suppressed by church authorities, who were ever-vigilant to detect conflicts with current dogma. The dismal history of the church's dealings with Copernicus and Galileo provides only two of the most egregious examples. I feel sure that some religious authorities would do the same today if it were within their power, and academics tend to be very sensitive about this for obvious reasons. This history does not excuse the tendency within contemporary academia to treat Christianity with singular contempt, but that is a topic for another day (see Marsden, 1994).

[11] I first became aware of Lewontin's review through the extensive quotes in K.R. Miller's *Finding Darwin's God* (Miller, 1999).

This ideologically loaded interpretation of science incorporates a nontrivial amount of circularity in its reasoning. If we assume from the beginning that reality consists exclusively of what can be demonstrated by the scientific method, then, of course, we will conclude that "science" directly implies a universe in which any other source of knowledge or value is without a compelling foundation. That such circularity continues to permeate public scientific discussion borders on intellectual irresponsibility, in my opinion. Let me be very clear about one thing: I am *not* saying that it is inappropriate to have and argue a strong point of view (as I am doing in this paper). I *am* saying that it is intellectually irresponsible to present a particular point of view as a result or direct implication of science when it is, in fact, no such thing.

In general, science is simply mute concerning the ultimate questions of meaning, purpose, and value. Any being, event, or insight that lies outside the realm of cause-and-effect mechanism is not approachable by natural science. This is not a logical declaration about the limits of reality; it is a declaration about the limits of natural science as a system of knowing. As I argue in the next section, our judgments on matters of meaning, purpose, and value almost always come from sources other than science and are formed in ways that necessarily depart radically from the scientific method.

36.3 A different part of the brain?

Several years ago, I made an important career decision, largely on the basis of my wife's employment opportunities as a Protestant minister. As a result, many of my professional colleagues around the country discovered, for the first time, that I am a Christian, which subsequently led to many interesting conversations! For example, I once shared an airport taxi with a prominent neuroscience colleague, whom I also consider a casual friend, following a scientific conference that both of us had attended. We respect each other's science, we had talked casually on previous occasions about science and family life, but we had never been closely involved in any way. It is one of those relationships in which one senses the potential for real friendship if the vagaries of time and space were to allow more interaction.

In a progression that has become familiar in recent years, the conversation proceeded from my recent career decision, to my wife's "interesting" professional occupation, to my own religious sensibilities. During the conversation, it became clear that my religious beliefs contrasted greatly with my colleague's pronounced skepticism in religious matters. This conversation, like many others I have had, was both personally warm and thought-provoking. We were genuinely interested in each other's stories, and the conversation was rewardingly unperturbed by any hostility or condescension. Near the end of the conversation, however, my colleague peered intently at me with a very puzzled expression on his face and said the following about my religious commitment: "I just don't see how you get there; you must use a different part of your brain when you do that."[12]

[12] In talking about a "different part of the brain," my colleague was making a serious point in whimsical neuroscience-speak. He was not actually proposing a theory that specialized brain circuits are responsible for religious belief and behavior, although this is a possibility that some neuroscientists take seriously.

This remark deserves careful consideration because its essence is expressed repeatedly in conversations that I have with academic colleagues. It is the same question as was posed by a postdoctoral fellow in my laboratory following a rare discussion of religious matters over lunch: "But, Bill, this way of thinking is so different from your *normal* way." Both my colleague and my postdoc, of course, were contrasting the modes of thought and belief that underlie my religious commitments to the modes that prevail in my scientific endeavors. In science, I am relentlessly critical, demanding high standards of evidence before accepting any scientific "result" into the canon of what I believe to be true about the world. Both my colleague and my postdoc were struck by the apparent inconsistency in my adoption of religious beliefs without similarly rigorous standards of proof.

My reply is that, yes, the modes of thought in the two domains can be quite different. This is one of the genuine points of tension between science and religion. Importantly, however, the mode that predominates in religious life *is* the *normal* mode of evaluation and decision making in the overall context of human experience. The *scientific* mode, in contrast, is quite peculiar: it is applicable to a rather narrow range of experience and is generally practiced by a rather small community of professionals. My central argument here is almost obvious, but I find that it needs to be aired repeatedly in the professional circles in which I move: *the most important questions in life are not susceptible to solution by the scientific method*. In fact, I tend to believe that the importance of a question is inversely proportional to the certainty with which it can be answered. How, for example, does one design an experiment to answer the question "Is it better to live or to die?" This certainly qualifies as an important question and will have been (or is) a live issue for some who read this paper. Or what laboratory procedures can one perform to address the question "Should I uproot my family, all of whom are deeply enmeshed in their own social networks, in pursuit of a new professional opportunity elsewhere in the country?" Most would agree that this is an important question – of much more intense concern to most people than the value of the universal gravitational constant.

These kinds of questions, which we all face routinely, simply do not submit to scientific solutions. We cannot make one choice and see how the experiment comes out, then rewind the tape and make the other choice to determine the outcome in the alternative scenario. Rather, we have a one-time shot at our most important decisions. We are forced to rely on intuition, on experience, on the advice of friends, on precarious projections into the future, and, in the end, on our gut feelings about what is likely to prove "right" in a given situation. Anyone who has been a parent, particularly a parent of teenagers, knows that excruciating decisions must be made on the basis of distressingly little "hard" data about likely outcomes!

Simply put, this is the human condition. It is *life*, and our most consequential decisions in life have little or nothing to do with science. This does not mean that we cannot bring rational analysis to bear on the issues. Thoughtful people reflect carefully on important decisions and try to take into account as much evidence as is reasonably available at all times. Nevertheless, rational analysis rarely compels a particular choice and certainly does not guarantee any particular result.

At the risk of belaboring the obvious, it is worth considering a particular decision that people commonly confront: the decision to marry a specific person. Difficult and highly consequential judgments must be made. Do I love this person in a deeply authentic, sacrificial way that can sustain a lifetime relationship? Or is my desire to marry based on less worthy forms of self-interest, whether related to money, status, infatuation, sex, or certain notions of compatibility? Does this person, in turn, love me genuinely? Do we have what it takes to weather the storms that life will inevitably bring our way?

Certainly, experience in the relationship counts for a lot in making such judgments. Knowledge of the values provided by the potential spouse's family of origin can yield significant insight. The advice of friends and mentors carries weight. Yet all of these sources of information can be flawed and deeply misleading. In the end, a considerable amount of faith is involved in the commitment to marry. The commitment carries substantial risk, as anyone who has been through a divorce (and many who have not) can attest. But it also offers the opportunity for the most rewarding of human relationships, as many can also testify. If one waits for compelling evidence (in the scientific sense) before marrying, one will never marry. One might say, regarding my colleague, that this sort of decision making occurs in a "different part of the brain" than scientific decision-making, yet it is common to all of us, including my colleague and my postdoc.

I believe that the religious quest involves exactly the same mode of thought (i.e., "part of the brain") as is involved in the marriage example above. Reduced to its most basic level, the religious quest hinges on a gut-level judgment about what sort of universe we really inhabit. Do we accept the "indifferent" universe of Richard Dawkins, or can we perceive with Teilhard de Chardin that "there is something afoot in the universe, something that looks a lot like gestation and birth"? Can we observe with Paul Tillich that "here and there in the world and now and then in ourselves is a New Creation"? Our actions, our hopes, and our aspirations pivot critically on the answer to this single question. Sources of evidence are available to guide my judgments: my own primary experience in my relationship with God (worship, prayer, and at least something like halting obedience), my experience in my religious community, the testimony of scriptural writers and other authentic seekers through the ages, and the critical reflections of fellow pilgrims whom I meet along the journey. Nevertheless, the evidence in the end is not compelling in a scientific sense. As in marriage, faith accompanied by commitment must play a foundational role in the religious quest; considerable risk is involved, and the stakes are high. I might make a complete fool of myself, or I might, as crazy as it sometimes seems, come into contact with the central reality of our universe, which I believe is more wonderful than we usually dare dream.

This tension between scientific and religious judgment was captured pungently in a brief conversation I once had with a faculty colleague at Stanford. At the center of Stanford's beautiful old quads lies Memorial Church, a Romanesque masterpiece dearly loved by many members of the Stanford community. My faculty colleague thought differently, however, and once exclaimed to me, not even slightly in jest, "That church pisses me off; I think we should bomb the thing!" When I asked why he felt that way, he replied "It is a monument to irrationality; it doesn't belong on a university campus." As it happens, at the time I lived in a home on the Stanford campus, and my reply to my colleague was "By far, the most irrational thing I have ever done was to marry and have children. If we are going to bomb campus monuments to irrationality, we had better start with my home!" The point of my reply was the same as the point of the discussion above: religious judgment and decision making are remarkably similar to the modes of judgment and decision making that we all employ and rely on countless times throughout our daily lives. In a very real sense, science – not religion – is the odd man out.

The tendency of some scientists (and perhaps of academics in general?) to quarantine religion into a uniquely irrational category of human behavior appears to me to be profoundly mistaken. It fails to grapple honestly with the complexity of the human condition and with the highly varied forms of thought and judgment that are required of us all as we navigate our ways through life.

36.4 Human freedom

A central tenet of Christianity and most other religions is that human beings have a meaningful degree of freedom to make moral choices. We can make loving, sacrificial choices in how we interact with others, or we can act in ways that are exploitative, or, at worst, overtly hateful and destructive. The issue of human freedom is an increasingly vexing point of tension between religious and scientific worldviews. What are we to make of human freedom when, from a scientific point of view, all forms of behavior are increasingly seen as the causal products of cellular interactions within the central nervous system, which themselves are substantially influenced by the toss of genetic dice that occurred when each of us was conceived? To frame the issue in an everyday context, can I really "choose" to have fish or chicken for dinner this evening, or do events already in motion reduce me to a predetermined course of action? More disturbing yet is, if our sense of choice is illusory, can anyone reasonably be held responsible for his or her actions?

The issue of human freedom is a tricky one. Some modern thinkers find refuge from strict determinism in quantum mechanics (QM), which describes events probabilistically rather than deterministically. While QM does imply that we live in a fundamentally unpredictable world,[13] I am not yet convinced that it offers substantial insight into human freedom. It can establish probabilities for the occurrence of specific events, but, within the constraints of those probabilities, events occur randomly. It is not clear to me that randomness provides an understanding of human freedom that is any more meaningful than that of strict determinism.[14] Our intuitive understanding of human freedom is that we have some meaningful

[13] Consider, for example, a classic quantum-mechanical phenomenon – the absorption of photons by matter. Absorption of high-energy photons by DNA can lead to genetic damage that results in cancer (e.g., melanoma). If I die of cancer, my life and the lives of my family, friends, and colleagues are drastically and irreversibly changed. Yet the triggering event – photon absorption – is fundamentally random and unpredictable, even in principle. The fully deterministic world can be set aside.

[14] A deeper argument for the relevance of quantum theory to the notion of human freedom is provided by Henry Stapp (1993). In a nutshell, the argument is that quantum theory, which is our most sophisticated and far-reaching physical theory of the universe, requires the existence of an "observer" who lies outside the causal system of physics to ask questions of nature

degree of *autonomy*, or *self-determination*. While we are certainly influenced by random events (in the quantum-mechanical sense) and by strictly determined events (in the Newtonian sense), we are at the complete mercy of neither.

Some of my scientific colleagues seem to feel that the notion of human freedom must be tolerated as a practical matter in order to maintain a functioning society, but that human freedom is likely to prove illusory in the final analysis. From their perspective, brains are extremely complex neurochemical machines, and their behavior will ultimately be understood in the same mechanical terms as those in which any other machine is understood. While notions of human freedom are convenient and probably even necessary in order for one to get along in everyday life, our subjective experience of freedom itself is no more than the result of machine-like activity within specific regions of the central nervous system.[15]

What this point of view fails to realize, however, is that the sense of human freedom, or autonomy, is just as important for scientific understanding as for everyday understanding of the world. Thorough-going determinism becomes entangled in profound logical difficulties in science no less than in everyday life. J.B.S. Haldane put the matter succinctly:

> If my mental processes are determined wholly by the motions of the atoms in my brain, I have no reason to suppose that my beliefs are true . . . and hence I have no reason for supposing my brain to be composed of atoms. (Haldane, 1927, p. 209)

Haldane's point is that the entire enterprise of science depends on the assumption that scientists have freedom to evaluate evidence rationally and make reasoned judgments about the truthfulness of particular hypotheses and results. If, however, the scientist's rational judgments and his or her beliefs about the validity of the scientific method simply reflect an inevitable outcome of the atomic, molecular, and cellular interactions within a particular physical system, how can we take seriously the notion that his or her conclusions about the world bear any relation to objective truth? (Ironically, the ardent determinist becomes an intellectual bedfellow of the ardent deconstructionist.) Furthermore, if we cannot believe that the scientific approach leads to some approximation of truth, how can we take seriously the scientifically based assertion that mechanical determinism is the correct way to think about the world? The attempt to adopt a thorough-going determinism is like sawing off the limb of a tree on which one is sitting; the result is intellectual freefall. Like it or not, then, achieving a meaningful understanding of human freedom is profoundly important for science, for society, and for each individual person.

(i.e., to propose experiments). Because the observer lies outside the causal system described by the wave equations of QM, the observer is free of constraints in a manner that conforms to intuitive ideas of human freedom. I am intrigued by this argument because it appears to be profoundly required by the best available physical theory of the universe, but the biologist in me finds in difficult to swallow. For the biologist, the "observer" asking questions of nature cannot lie outside of nature; he or she is a human being who operates wholly within the natural system of life on Earth. The nagging suspicion of the biologist is that quantum theory (or at least some prevailing interpretations thereof) just doesn't have it right yet!

[15] But, as Charles Jennings has observed, throw a rock through the living-room window of the most reductionistic neurophilosopher, and you will probably find out just how quickly the dispassionate notion of behavioral determinism evaporates! See Jennings (1998, pp. 535–6.)

How are we to reconcile the "autonomy" of a reasoning intellect with our scientific conviction that all behavior is mediated by mechanistic interactions between cells of the central nervous system? Although I have no certain answer to this question, I suspect that answers will ultimately lie in a deeper understanding of emergent phenomena in complex systems. This is a somewhat slippery concept and has been used in different ways by different authors.[16] By "emergence," I mean that complex assemblies of simpler components can generate behaviors that are not predictable from knowledge of the components alone and are governed by logic and rules that are independent of (although constrained by) those that govern the components. Furthermore, the intrinsic logic that emerges at higher levels of the system exerts "downward control" over the low-level components. To foreshadow my ultimate argument, it is the phenomenon of downward control that endows a system with a behavioral autonomy, which in the case of biological organisms can be regarded as meaningful choice.

Many authors have cited examples of emergent behavior in complex systems, a favorite example being the unicellular organism. The existence of unicellular organisms permits an enormous number of new phenomena that could not be predicted from knowledge of macromolecules alone and that operate on principles that go well beyond those that govern macromolecules: cellular motility, foraging for resources, competition with other organisms, and adaptation to environmental pressure by means of mutation, to name but a few. Each of these phenomena must be identified and described in and of themselves and their internal logical rules worked out before rigorous links to lower-level mechanisms can be made. Competitive interactions between species, for example, are comprehended by observation at the behavioral level, not by inference from the molecular level. The behavior of the unicellular organism, in turn, exerts downward control over its constituent molecules. The motion of an organelle within the cell depends, in one sense, on pressure exerted from the cytoplasm as the organism moves. But in another, equally valid sense, the motion of the organelle depends on the immediate behavioral goal of the organism.

It is critical to be very clear on one point: the concept of "emergence" does not imply magic or mysticism.[17] As far as we know, nothing about the life of unicellular organisms violates the laws of physics or the chemical laws that govern the behavior of macromolecules. The cell cannot behave in any way that is not permitted by the lower levels of organization of its constituent parts; the behavior of the cell is thus *constrained*, but not *determined*, by the lower levels.

[16] There exists a large literature, both formal and informal, on the theme of emergence in complex systems. For recent examples, see Clayton and Davies (2006) and Clayton (2006).

[17] My discussion here will not invoke brain events that violate known physical principles. More than anything else, this reflects my biological intuition that the human brain, as a product of the natural evolution of the universe in general and life on Earth in particular, will operate in a manner consistent with (i.e., constrained by) known physical laws. It is certainly conceivable, and perhaps even likely, that some aspects of human and animal consciousness will never be satisfactorily understood from the point of view of the reductive sciences (e.g., Nagel, 1974), but one doesn't want to throw in the towel until absolutely forced. If Copernicus and Galileo, for example, had shrugged their shoulders and accepted contemporary theological explanations of celestial motion, progress in understanding our solar system would have been severely stunted. As many writers have pointed out, acceptance of extra-physical accounts for a particular phenomenon is "giving up" from a scientific point of view, and it is far too early in the history of neurobehavioral science to entertain the thought of giving up.

Obviously, the crucial distinction here is between the words “constrained” and “determined.” This distinction becomes clear for me in considering the operation of the computer program that is running right now on my laptop computer. If I want to understand how Microsoft Word operates, I can tackle the problem at the mechanistic level of transistors, resistors, capacitors, and power supplies; or I can tackle the problem at the level of the software – the logical instructions that lie at the heart of the process of computing. It seems clear to me that the most incisive understanding of Microsoft Word lies at the higher level of organization of the software. One wants to understand the logical relationships involved in computation: for-loops, if-statements, and the like. The logic of the computation exists independently of the physical system of electronics that make up the computer (the software can be transferred to another computer) and operates according to its own rules that cannot be predicted from knowledge of the hardware alone. The rules of computation logic, in turn, orchestrate (in a real, causal sense) the currents flowing through the myriad individual components that constitute the computer. Again, nothing magical or mystical is occurring here. The software is constrained by the hardware; the software cannot abrogate the laws of physics or the principles that govern the behavior of electronic circuits. Nevertheless, the behavior of the computer as I type this text is *determined* at a higher level of organization – the software – not by the laws of physics or the principles of electronic circuitry.

Although this computer example emphasizes the critical distinction between “constraint” and “determination,” it is *not* an example of emergence because the software did not evolve from a natural process of self-assembly, but was designed by human programmers. A better example of emergence in the computing world lies in the relatively new field of neural networks. In the neural net illustrated in Fig. 36.1(a), multiple layers of “neuron-like” computing units are linked to one another in a hierarchical manner such that the behavior of each unit in a lower layer influences each unit in the next-highest layer (arrows). The strength of the influence of any given lower-level unit on units in the next-higher level is governed by a set of “weights” that determines the effectiveness of the link between each pair of units. In the initial state of the network, the weights governing the many links are chosen randomly; some are positive, some are negative, some are strong, some are weak. An input is then provided to the lowest level of the network, and an output emerges at the highest level. In a backpropagation network (one of several types of neural network), a software entity called a “teacher” then recognizes whether the actual output is similar to the desired output and adjusts all of the weights of the links between computing units accordingly. After many iterations of the input–output–adjustment cycle, the network “learns” to produce the correct output for a given input.

Neural networks can perform remarkable feats that are extremely difficult to accomplish by traditional computing methods, which employ mathematically precise algorithms specified by a programmer. Some of the most impressive examples lie in the arenas of voice and pattern recognition and of robotics. Yet a remarkable intellectual quandary is often encountered in the neural-network field: a network can be trained to solve a fiendishly difficult problem, and, in the end, the human programmer who designed the network and orchestrated the training procedure may have little or no insight into *how* the problem has

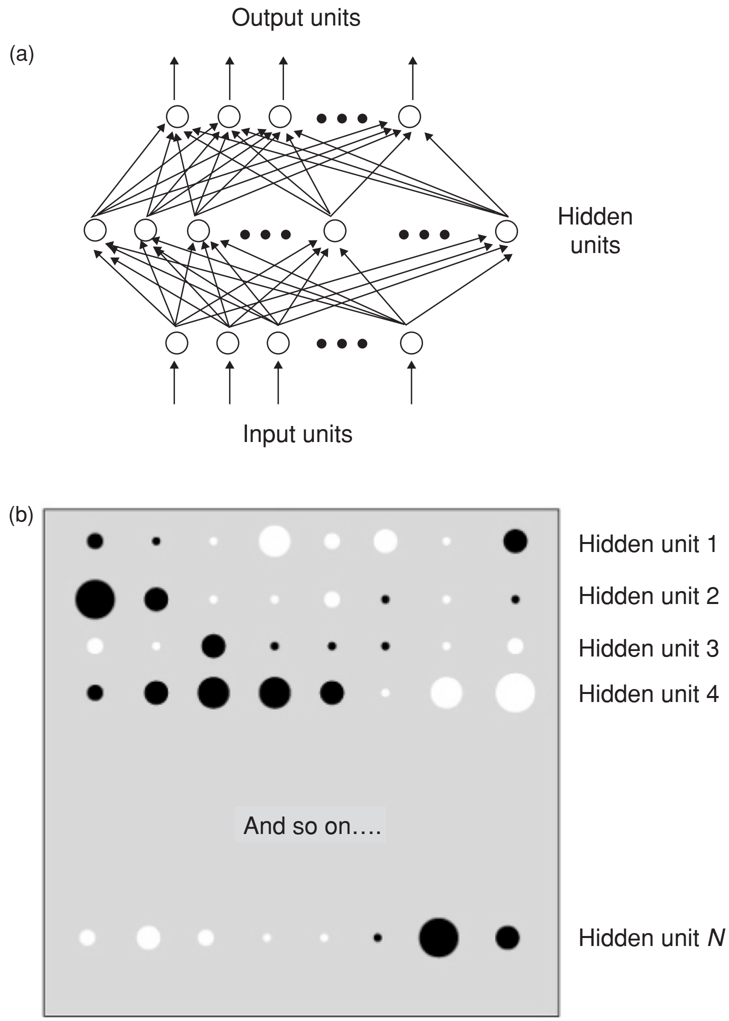

Fig. 36.1. A schematic diagram of a common multilayer neural network. (a) Network architecture. Each circle represents a computing "unit." The units are arranged in three hierarchical layers: from bottom to top, "input" layer, "hidden" layer, and "output" layer. Signal flow is "feedforward" in the sense that a given layer exerts causal influences only on the next-highest layer. Each unit in a given layer influences the activity of each unit in the next layer as illustrated by the arrows. The initial strengths, or "weights," of the connections between units are random. Some are positive (activity in the "sending" unit *increases* activity in the "receiving" unit), and others are negative. Some weights are strong (the sending unit has a *large* impact on the receiving unit), while others are weak. These weights are adjusted during the learning process according to the similarity of the actual outputs to the desired outputs. (Diagram adapted from Rummelhart, *et al*., in *Parallel Distributed Processing: Explorations in the Microstructures of Cognition*, Vol. 1 (Cambridge: MIT Press, 1986).) (b) After the learning process, the final configuration of the network, which embodies the learned solution to the problem, is depicted as the final set of weights between the various units. If, for example, there are eight input units and twenty hidden units in the network

actually been solved! The programmer can show us the final pattern of weights between the individual computing units that somehow embodies the solution, as in Fig. 36.1(b), for example; but we frequently remain embarrassingly ignorant concerning the algorithmic principle(s) the network has "discovered" in solving the problem.[18]

This example comes closer to the meaning of emergent order in complex systems. At a "low" level, we know everything there is to know about the neural network and the digital computer on which it runs. We fully understand the physical principles underlying the operation of the computer, as well as the learning algorithm that enables the network to modify its connections as it interacts with the environment. Furthermore, at the end of the learning exercise the programmer has full knowledge of the learned connection weights, and he or she may transmit the "solution" in the form of connection weights to anyone in the world who would like to implement it for their own purposes. This point cannot be emphasized too strongly: at a mechanistic level, there is no causal gap in our understanding; we know *everything* that matters about the neural network – both its final state and precisely how it got there. Paradoxically, however, we are frequently unable to state, or write an equation for, the algorithmic principle that lies at the heart of the learned solution. Our situation resembles that of an electronics assembly technician who can solder components together to create a functioning, causally complete, electronic circuit, yet has little or no idea how the thing actually works at a high level.[19]

How can this be? In the case of the electronics technician, the answer is clear: the technician simply follows a design created by another intelligence – the circuit engineer. The circuit engineer is not imaginary or epiphenomenal, but rather is a critical locus of "downward" causal control in producing a functioning circuit. In the case of the neural network, users of the network exploit a design created by a learning interaction between the network and its environment. As is typical of systems that learn, the actual structure of the

[18] A computationally savvy colleague of mine at Stanford refers to these networks, with a mixture of humor and derision, as "know-nothing networks" because at the end of the exercise the scientist still might not *understand* the solution that has been achieved.

[19] This is a somewhat humiliating situation for a scientist to be in – understanding a system essentially completely at a "low" level, but being quite ignorant of how it operates at a "high" level. Most of us feel intrinsically that we *must* understand the higher level of organization, which in the case of neural networks involves formal computational logic, if we are to be intellectually satisfied with the result. One possible reaction to this dilemma is to deny that any "higher level" exists in the network. If we know the transfer function of each individual computing unit and the weights of all the connections, we can calculate the output for any given input, and there is nothing else to know scientifically. For me, this is not a sustainable point of view. It brings to mind Thomas Nagel's observation: "To deny the reality or logical significance of what we can never describe or understand is the crudest form of cognitive dissonance" (Nagel, 1974).

in (a), the weights of the connections between each input unit and each hidden unit can be depicted as in (b). White dots indicate positive weights; block dots depict negative weights. The size of the dot is proportional to the strength (or weight) or the input. The top row of dots depicts the weights from all eight input units onto hidden unit 1, and the second row depicts the weights from all eight input units onto hidden unit 2. The rows are iterated until the weights to all twenty hidden units are represented. A similar diagram (not shown) depicts the weights from each hidden unit onto each output unit. These "Hinton" diagrams (named after their originator, G.E. Hinton) fully describe the final state of the network and can be reproduced at will on any suitable digital computer.

network changes as a result of new information acquired during the learning process, and the new (emergent) structure of the network embodies the learned solution to the problem.[20] As with the circuit engineer, the intelligent solution embodied in the emergent structure of the network is not imaginary or epiphenomenal, but rather is a critical locus of downward causal control, implementing very practical solutions to complex problems.

Having wandered a bit from my original topic, let me now state exactly how my "toy" example of neural networks is – and is not – relevant to understanding human autonomy, which I take to be the essence of freedom. The most relevant lesson is this: a complex system endowed with the ability to learn possesses the autonomy to discover solutions (to problems) that cannot be captured by, or predicted in advance from, lower-level descriptions, including the learning algorithm itself.[21] Information embedded at higher organizational levels is the most important locus of causal control of the system. A skeptic might argue that this toy example provides no understanding of autonomy (or freedom) whatsoever because every aspect of the network, including each step of the learning process, is causally determined. Given the same original set of weights between the computing units, the same learning algorithm, and the same set of inputs from the environment, the network would produce exactly the same solution by exactly the same series of steps each time it was run. My reply to this objection – which should be clear by now – is that a breach of causality is not a requirement for "autonomy"; a central point in my discussion of neural networks is that their autonomy is real even though their function is entirely causal.[22] My fundamental argument lies at a deeper level – I suspect strongly that our standard notions of causation in physical systems are impoverished. At certain levels of complexity, the primary drivers of system behavior are the logical rules of operation intrinsic to higher levels of the system; no other level of explanation satisfactorily captures the nature of the system.

Although I plainly have no elegant solution to the problem of free will, I believe that understanding human freedom is the most important and most difficult long-term challenge facing the neurobehavioral sciences. Our freedom is certainly restricted by our biology (more so than most of us would like to admit),[23] but a meaningful capacity for

[20] As mathematician George Ellis points out, the higher-order interactions of a complex system formally resemble Darwinian selection mechanisms: hugely variable events in the world impact each organism through the selective filter of the organism's behavioral goals. See Ellis (2006a). Ellis provides a lucid scientific account of emergence and top-down causality in complex systems. He considers very generally how top-down effects should properly be viewed within physics, and he elucidates nicely the effects of goals and selection in creating nonreducible, high-level information in biological systems. Also see Ellis (2006b).

[21] As a computational neuroscience colleague at MIT once said to me, "If we could figure out the solution in advance, we wouldn't have to throw a network at the problem."

[22] I emphasize that the neural-net heuristic is only that – a heuristic. It allows us to appreciate important points about complex systems, but it does *not* necessarily provide deep insight into the nature of human cognition, *per se*. Human brains, and those of other animals as well, are vastly more complex than the neural nets that we employ in our most advanced sciences, and new phenomena with their own intrinsic logic will certainly emerge at every added level of complexity within the nervous system. Of particular importance are the abilities of humans to reason with symbols and to reason recursively about our own reasoning (Deacon, 1997). With these evolutionary accomplishments, the relationship between our highest-level behaviors and the underlying "wetware" (ion channels, membranes, single neurons) becomes even more indirect. The relationship exists, of course, which is a major reason why neuroscientists such as myself have jobs. But the relationship is more a matter of *constraint* than of *generation*. As always, the biophysics of the constituent wetware constrains the phenomena that are possible at higher levels, but the behavioral possibilities that are actually realized are determined by higher-order interactions of an organism with its environment.

[23] The remarkable "identical twins raised apart" studies emphasize the pervasive influence of our genetic composition on surprisingly varied aspects of behavior, from basic temperament to small behavioral tics; we are not free to escape many aspects of our genetic heritage. See for example Kendler (1993) and McCleam *et al*. (1997).

self-determination (autonomy) is an irreplaceable foundation for taking seriously the notions of scientific truth, religious truth, and individual moral responsibility. I have argued that a satisfactory understanding of this capacity will ultimately lie in the concepts of emergence and downward causality within complex systems. Emergent behaviors of even simple learning systems are often surprising and deeply perplexing, yet they can "get in touch with" realities whose deeper foundations are difficult to discern long after we accept the validity of the behavior. Thus "emergence" becomes a pivotal concept for interpreting the reality of human life in all its complexity, from scientific endeavor to personal morality to religious understanding. Although emergence is a notoriously difficult phenomenon to study rigorously, few areas of study are likely to prove as intellectually and practically consequential in the long run.

36.5 Concluding remarks

My purpose here has been to examine the interaction of science and religion, following a lively tradition exemplified in our own era by Charles Townes, by several authors who have contributed to this book (including Townes), and by a host of others. This chapter has focused more on formal topics in the science–religion dialogue than on my personal religious experience.[24] Perhaps a few words about the latter are in order now.

Across thirty-five or so years of adult life, I have tried to discern for myself whether there is anything in the universe worth having faith in, what it means for me personally to live in faith, and how my faith is related to all other facets of my life, including the science that I do. My search for an authentic faith is as much a part of me as eating, sleeping, and breathing, and it is certainly more fundamental to who I am than is the science I do. I have often told friends and audiences that if I were forced to give up either my science or my faith, it would be the science that would have to go. Fortunately, I am not faced with such a dreadful choice. As I indicated in the first section of this chapter, I find no conflict between my faith and the findings of contemporary science; in fact, I find them to be remarkably consonant. For me, both my religious faith and the science I do are integral, highly valued parts of my life.

Are aspects of my faith "irrational"? If committing myself to important beliefs that are beyond science is irrational, then of course the answer is "yes." But, as I argued earlier in this chapter, religious commitment is similar, in this respect, to almost all truly important commitments that we make in life. Our most consequential beliefs and actions are rooted in value-laden interpretations of existence, none of which are testable scientifically. When examined carefully, the blend of intuition, analysis, and hopeful commitment that characterizes religious life is not peculiar at all; rather, it is science that is uniquely peculiar in restricting its subject matter and its way of knowing as much as possible to a slice of reality that is most amenable to mechanistic analysis.

[24] A previous publication focused directly on the personal aspects of my religious faith and religious history (Newsome, 2000).

Does my religious faith actually change anything? Does faith influence my behavior positively, or would I be essentially the same person without it? I am willing to believe that much of my behavior would be indistinguishable with or without my faith. Basic aspects of my personality were genetically shaped; many of my abilities for coping with life's ups and downs were nourished or stunted, as the case may be, by early experiences in my family of origin. Even without my faith, I would for the most part avoid overt nastiness to people with whom I interact because nastiness is – well – unpleasant. Where the rubber really hits the road, however, is how much I am willing to be involved with, and sacrifice for, people who might otherwise be faceless to me. This is where religious faith makes a real difference for me. Put simply, belief matters. If I believe that human life, in general, and my own life, in particular, are cosmic accidents, then I am definitely less inclined to struggle for higher ideals. There are winners in life and there are losers; ultimately, none of it means anything anyway. As one of my scientific colleagues earthily said to me, "We are farts in the wind, Bill." If, on the other hand, the central reality of our universe is a loving Creator, who binds us together into a single family and ultimately draws us gently to himself or herself through life's great joy and great suffering, then the game is entirely different. Our struggles for kindness, justice, and mercy are not a pointless cry against the void, nor are they simply a utilitarian adaptation for getting through life in a reasonably orderly manner. Rather, they point us toward the very heart of reality, a reality that is more blessed and more sustaining than we typically dare to hope.

Whatever else this kind of faith might be, it seems to me both coherent and beautiful – two qualities that are highly prized by scientists. It is coherent in the sense that what one believes about the central reality of our universe (ontology) is consistent with how one chooses to act (ethics). It is beautiful in the sense that it evokes, nourishes, and sustains our highest ideals and aspirations. Is this faith in the end True? I cannot know for certain. What I do know is that my faith makes more sense to me than any other system of belief (or non-belief) that I have found. I, like many other folks, have spent a lot of time trying to figure out what sort of mess I landed in by being born. Most of the time, I am convinced that it is a Holy mess. I struggle for coherence and consistency, and this Holy view of existence is the one that accounts best for life as I experience it, both with my mind and with my heart.

Acknowledgments

I need to thank several colleagues who provided critical input during the development of this chapter: George Ellis, Stephen Kosslyn, Brie Linkenhoker, Anthony Movshon, Michael Shadlen, and Robert Wurtz. The input of these reviewers should not be construed as approval of the contents; several of them, in fact, disagree fundamentally with my theistic worldview. But these colleagues did exactly what good colleagues always do – they provided frank, constructive feedback so that I could test my own ideas against the reflections of respected "others." I also thank Phillip Clayton, without whom this chapter would probably not have been written. Through his project Science and the Spiritual Quest (funded by the John

Templeton Foundation), Phil convinced me to begin putting some of my thoughts about science and religion into writing. He was a source of encouragement during the entire process.

References

Barbour, I.G. (1966). *Issues in Science and Religion* (New York: Harper & Row).

Barrow, J.D. and Tipler, F.J. (1986). *The Anthropic Cosmological Principle* (New York: Oxford University Press).

Boslough, J. (1985). *Stephen Hawking's Universe* (New York: William Morrow), p. 121.

Carr, B.J. and Rees, M.J. (1979). The anthropic principle and the structure of the physical world. *Nature*, **278**, 605–12.

Clayton, P. (2006). *Mind and Emergence* (Oxford: Oxford University Press).

Clayton, P. and Davies, P.C.W. (2006). *The Re-emergence of Emergence* (Oxford: Oxford University Press).

Conway Morris, S. (1998). *The Crucible of Creation: The Burgess Shale and the Rise of Animals* (Oxford: Oxford University Press).

Conway Morris, S. (2003). *Life's Solution: Inevitable Humans in a Lonely Universe* (Cambridge: Cambridge University Press).

Dawkins, R. (1975). *River Out of Eden* (New York: HarperCollins), pp. 132–3.

Dawkins, R. (1987). *The Blind Watchmaker* (New York: W.W. Norton) (paperback edition).

Deacon, T. (1997). *The Symbolic Species: The Co-evolution of Language and the Human Brain* (London: Penguin).

Ellis, G.F.R. (2006a). Physics and the real world. *Foundations Phys.*, **36**, 1–36. Also available at http://www.mth.uct.ac.za/∼ellis/realworld.pdf.

Ellis, G.F.R. (2006b). On the nature of emergent reality, in *The Re-emergence of Emergence*, ed. P. Clayton and P.C.W. Davies (Oxford: Oxford University Press), pp. 79–107.

Gelwick, R. (1977). *The Way of Discovery: An Introduction to the Thought of Michael Polanyi* (New York: Oxford University Press).

Gould, S.J. (1989). *Wonderful Life: The Burgess Shale and the Nature of History* (New York: Norton).

Haldane, J.B.S. (1927). *Possible Worlds* (London: Transaction Publishers), p. 209. See http://www.daviddarling.info/encyclopedia/H/Haldane.html.

Hawking, S.W. (1988). *A Brief History of Time* (Toronto: Bantam Books), p. 125.

Jastrow, R. (1978). *God and the Astronomers* (New York: W.W. Norton), p. 116.

Jennings, C. (1998). *Nature Neurosci.*, **1**, 535–6.

Kendler, K.S. (1993). Twin studies of psychiatric illness: current status and future directions. *Arch. General Psychiatry*, **50**, 905–15.

Koza, J. (1992). *Genetic Programming: On the Programming of Computers by Means of Natural Selection* (Cambridge: MIT Press).

Lewontin, R. (1997). Review of *The Demon-Haunted World: Science as a Candle in the Dark*, by C. Sagan. *New York Review of Books*, January 9.

Marsden, G.M. (1994). *The Soul of the American University: From Protestant Establishment to Established Disbelief* (New York: Oxford University Press).

McClearn, G.E., Johansson, B., Berg, S., *et al.* (1997). Substantial genetic influence on cognitive abilities in twins 80 or more years old. *Science*, **276**, 1560–3.

Miller, K.R. (1999). *Finding Darwin's God* (New York: HarperCollins).

Nagel, T. (1974). What is it like to be a bat? *Phil. Rev.*, **83**, 435–50.
Newsome, W.T. (2000). Science and faith: a personal view. *The Ellul Forum*, January (24), pp. 2–8.
Provine, W. (1988). Evolution and the foundation of ethics. *MBL Sci.*, **3**, 25–9.
Stapp, H.P. (1993). *Mind, Matter and Quantum Mechanics* (Berlin: Springer).
Townes, C. (1966). The convergence of science and religion. *Think*, **32**, 2–7.
Townes, C. (2003). The convergence of science and religion. *ASA J.: Perspectives Sci. Faith*, **55**, 154–8.
Townes, C. (2006). Marriage of two minds. *Science and Spirit Magazine*, January–February, pp. 28–43.