

# Chapter 2

## Engineering, Daydreaming, and Control



Daniel C. Dennett

Philosophy and sailing go together beautifully in my experience. First, and most important to me, philosophers' mistakes hardly matter (which is a good thing, since we philosophers make them by the dozens); our expertise, such as it is, seldom is called upon to deal with a crisis, help someone in peril, or—most dramatically—save a life. Nobody would get rich selling malpractice insurance to philosophers! Sailing on the open ocean, in contrast, is a place where one's knowledge is put to the test. I have always relished the occasions when I and my crew and passengers would have been in mortal danger had I not known what I was doing. Calmly dealing with rough weather, a ripped sail, a deranged GPS, a seasick sailhandler is a demonstration—to oneself—that one is in charge, using knowledge to make a difference. To my taste, it is one of life's most gratifying experiences.

That may say more about me than about philosophers in general. I have always been in my own mind a sort of engineer who happened to specialize in philosophy, a deviser and critic of intuition pumps, a fixer of perspectives, a maker of gadgets that can open people's minds to new ideas. Anyone who owns a sailboat big enough for overnight cruises has to be handy—or have a first mate who can be relied on to know how to set up a jury rig, retrieve a dropped tool from the bilge, reinforce a broken stanchion, and if need be sew up a nasty gash. The resourceful Robinson Crusoe was my childhood hero (long before MacGyver came along).<sup>1</sup> But a sailboat is also an ideal floating platform for daydreaming, for extended solo musings. At the

---

<sup>1</sup>I have been informed that Robinson Crusoe is an icon of colonialism, but that doesn't change the fact that he was my childhood hero, a fact I wouldn't want to suppress. Crusoe, importantly, said he loved Friday, a remarkable sign that he wasn't such a villain.

---

D. C. Dennett (✉)  
Center for Cognitive Studies, Tufts University,  
Medford, MA, USA  
e-mail: [daniel.dennett@tufts.edu](mailto:daniel.dennett@tufts.edu)

helm of a sailboat, slower and quieter than a motorboat, I can let time pass unmeasured. I'm a gregarious chatterbox most of the time, but alone at the helm, while others are sleeping or reading or just looking at the waves, I can be content for hours, carving out a true compass course, checking my bearings and the set of the sails.

Sailboats are also ideal for long discussions, occasionally interrupted by sightings of whales, dolphins or seals, or by passing boats of interest. On my beloved Xanthippe, an old Beneteau First 42 I bought in 2004, I would invite my grad students and postdocs to choose a guest for our annual Cognitive Cruise, somebody they wanted to talk to at length about their work. We'd have no schedule, no formal talks; we'd just go sailing for a few days, anchoring in beautiful coves, exploring deserted islands, and talking, talking, talking. I know of no better way to share ideas constructively.

The country singer Ricky Skaggs has a song ("Can't Control the Wind," 1995, in the album *Solid Ground*) the lyrics of which are as maudlin as only pop religious lyrics can be, but there is one great line:

I can't control the wind, but I can adjust the sail.

Skaggs is probably not much of a sailor, since a sailor would say "but I can trim the sail," which would even fit the music better, but never mind; what he sings is an important philosophical truth. You don't have to control everything to control something. If you took some philosophers seriously, you'd think that "control" is a synonym for "cause," but that is just wrong. There is an obvious and important difference between a sailboat driven (caused to move) by the wind while being controlled by a skipper and a sailboat driven (out of control) by the wind alone (while the skipper sleeps, say). Control requires *knowledge*, as I just noted, or at any rate *knowhow*, or at any rate *information* in the form of feedback to a competent mechanism that can exploit that information in a purposeful manner. Life itself depends on information in just this sense, and the difference between being alive and not being alive is primarily a matter of being able to use available information to modulate some processes. Non-living autonomous robots and other such mechanisms can also exploit such information, such knowhow. I was once cruising with friends along the coast of Greenland, on a compass course for Disco Island, and in a choppy sea I was having a hard time keeping on course. Back and forth the boat swung, leaving a snaky wake, and wearing me out at the helm. "You're working too hard, Dan, let the autopilot do it," the skipper suggested, and he turned it on and set it for the compass course. The below-decks hydraulics kicked in, and began to move the rudder, *whrrrrr! . . . whrrrrrrrrr . . . whrr-whrrr . . . whrrrrrrr*; as the helm spun back and forth on its own. "It's working as hard as I was," I commented, and sat back to rest. In a few minutes I noticed that the sound was different: *. . . .whrr . . . . .whr . . . . .whr . . . . .whr . . . . .whr*, and after about twenty minutes it was *. . . . .r . . . . .r . . . . .r . . . . .r . . . . .r* and the helm was almost stationary for longer and longer periods. The autopilot had learned to prune back its corrections to a minimum, and the wake was a straight white line behind us as far as I could see it. It had learned in half an hour what it sometimes takes me hours or days to learn—the "feel" of a

boat's helm under particular conditions, so that the slightest veer off the compass course gets just the right correction at just the right time. What was more surprising to me at first was the realization that it was using *much less* information than I was. It couldn't see the waves cresting, or even note the heel of the boat in the wind. All it had to work with was the compass heading, swinging back and forth and the effects of its steering actions on this. I'd been overthinking the problem, apparently, and we do a lot of that when we're novices at any activity involving our arms and legs—hitting a golf ball, turning on skis, returning a tennis serve to the backhand, doing elegant calligraphy or just ladling pancake batter on a griddle. When we “automatize” such behaviors, with or without wise advice from a coach, we streamline and optimize the information-handling without knowing how we do it or why—just like the reinforcement-learning algorithm in the autopilot. Of course it couldn't steer around any icebergs or flotsam in our path and it couldn't figure out when to change course (since we hadn't programmed in a turning point on the GPS). This is a nice example of a well-known irony: when the originators of the first wave of AI (Good Old Fashioned AI or GOF AI) tried to model human thought quite explicitly and directly, they created strange brittle systems that often output bizarre “misunderstandings” and failed to execute cognitive tasks that any normal six-year-old does with ease; when the creators of the second wave of AI (so-called *deep learning* systems) turned their backs on human psychology and just treated the tasks as engineering problems, they ended up creating systems whose behavior—in general—seems much more brainlike and “natural.” But before we decide that now we have found the right way to do AI, we should note that on many problems, such as deciding whether to alter course when some novel phenomenon looms, deep learning systems are as blind and useless as an autopilot. Yes, the self-driving car projects are confronting more and more of these issues and making impressive progress on them, but human-style judgment is still a promise on the horizon, not anything we should expect in the next few years.

The “dividing line” (it isn't a bright line) between the living and the non-living might well be drawn at the motor proteins that march by the trillions in all our cells, carrying materials around where they are needed. Are proteins alive or are they “just” very competent macromolecular robots with hundreds of moving parts? In either case, it turns out that Skaggs' line might be their motto, for they exploit the energy in the chaotic storm of jostling intracellular water molecules to move so intrepidly and reliably along their tubulin highways. Without the *uncontrollable* “wind” of those water molecules bombarding them, they would not have the energy they need to locomote along their paths. They are nano-sailors steering a true course and keeping us alive. Of course their “steering” is simplified by the fact that they are on rails of sorts, like the underwater rails that keep the adventure boats at Disneyland on course no matter what the “guide” does with the helm, but there is one degree of freedom that they do control: whether to move forward or backward on the rails. Without these hypermodest deciders doing their jobs down in the engine room, there would be no more sophisticated alternative courses of action to ponder, no “free choices” to determine our trajectories.

Can an agent with (real) “free will” be made out of parts that have only the freedom of motor proteins? Can something green be made out of atoms that aren’t themselves green? Can something alive be made out of non-living parts? Can something conscious be made out of unconscious parts? Of course! That’s the *only* way something could be green or living or conscious or free (in the sense that matters). Consider in this regard the flipping of a fair coin or the rolling of a pair of dice. The trajectories of these artifacts are caused, but *practically uncontrollable*. That’s what makes them perfectly good randomizers. The causation that *determines* the outcome is utterly incalculable. Dice are *designed* to be uncontrollable by us, and of course they are too simple to control themselves. We are also designed-by natural selection—to be uncontrollable by others, but we are able to control ourselves to a significant degree, thanks to the information we embody and acquire. A main task of cognitive science is to explain in detail how we merely physical bodies embedded in a deterministic world are competent *enough*—not *perfectly* “free”—to be wisely held responsible for our actions.

It is language that most sharply distinguishes us from all other animals, largely because it makes possible the sharing of voluminous and detailed knowledge on thousands of topics of relevance to our expanded life interests, and with this explosive collection of interests and knowledge we have many more degrees of freedom—in the engineering sense—than any other animal. Controlling *all* those degrees of freedom is practically impossible, but we have discovered hundreds of ways of improving our self-control and passed them on, culturally, to our offspring. By the time most of them become normal adults, they have the cognitive and moral competence to be held responsible for their actions. That is the free will worth wanting, and it has nothing to do with causation.<sup>2</sup>

**Daniel C. Dennett** is the director of the Center for Cognitive Studies and Austin B. Fletcher Professor of Philosophy at Tufts. [daniel.dennett@tufts.edu](mailto:daniel.dennett@tufts.edu)

---

<sup>2</sup>Kristin Andrews and Susana Monso, in “Animal Moral Psychologies,” forthcoming in John M. Doris & Manuel Vargas (eds.), *The Oxford Handbook of Moral Psychology*. New York: Oxford University Press, argue that many animals exhibit signs of the psychological attributes that are held to be prerequisites for moral responsibility. They do not advance the claim that animals are, or could be, morally responsible agents, which I think is a wise forbearance. Young children exhibit many more signs of these moral attributes than any animals, and yet are not considered morally responsible agents. Animals and young children used to be tried, convicted and even executed for crimes, but we human animals have advanced beyond that barbarous practice. Human exceptionalism at its best.