



# Herding Cats and Free Will Inflation

Daniel C. Dennett

*Tufts University*

*Romanell lecture delivered at the one hundred  
seventeenth annual Central Division meeting of  
the American Philosophical Association*

© 2020 The American Philosophical Association

---

---

# Herding Cats and Free Will Inflation<sup>1</sup>

Daniel C. Dennett  
TUFTS UNIVERSITY




Photo credit: Evan Mann

*Romanell lecture delivered at the one hundred seventeenth Central Division meeting of the American Philosophical Association in Chicago, IL, on February 28, 2020.*

## OUTLINE

### 1. HERDING CATS

Control is not causation.  
Determinism does not rule out autonomy.  
Autonomy is dangerous (it's like herding cats).  
We are autonomous but not particularly dangerous.  
Responsibility evolved (culturally).  
We are members of the Moral Agents Club.

### 2. FREE WILL INFLATION

Some puppets have free will.  
Privacy (or crypsis) is important.  
Rock, paper and scissors  
Jerry Fodor, Galen Strawson and Sam Harris on inflation  
A variety of free will worth wanting is under threat

### 1. HERDING CATS

Consider a boulder dislodged by an earthquake rolling down a mountain. Bump, bump, bump. Its path is determined by the laws of physics, but it is not under the control of anything. It is out of control. Its trajectory is *determined*, but not *controlled*. Compare the boulder

to an expert skier, hurtling down the mountain. The skier's trajectory is also determined by the laws of physics, but the skier is *in control*, her trajectory determined by her decisions, mediated by her skill and strength, and the existing conditions of the snow, the wind, the conditions of her skis, and so forth. The skier might *lose control* and then become more like the boulder, still determined but now out of control. And if she then somehow managed to regain control, she would then be profoundly different from the boulder, being *self-controlled* again. All of these paths are determined, let's agree; what happens is always the *effect* of the many *causes* at play. And, bringing out a shopworn image of "rewinding the tape" and playing back *exactly* the same causes, if the atoms composing the mountain, the air, the boulder, the snow, the skier . . . were placed in *exactly* the same initial state, exactly the same trajectories would be observed in each case. You cannot learn anything about the skier's remarkable difference from the boulder by running the tape again and again. Laplace's Demon, who tracks the trajectories of each atom, following the laws of physics, may review the sequence of causes and effects over and over, and never discover the interesting fact that if you start the boulder off down the mountain in a very slightly different initial state, it will almost certainly come to rest somewhere new, while the skier, on her second and third run down the slalom trail—with conditions varying considerably each time—will cross the finish line within a few inches of where she crossed the finish line on her first run, and with only a few milliseconds difference in elapsed time. That requires great ability.

Causation and control are not the same thing. Not all things that are caused are controlled. Things that are controlled are caused, like everything else, but control requires a controller, an *agent* of sorts designed to control a *process*, and that requires feedback: information about the trajectory and conditions that can be *used* by the controller to modulate the action. This is a fundamental point of control theory. Think about firing a rifle bullet. You're controlling the direction of the gun barrel and, with the trigger, the time of the bullet emerging from the gun barrel. Are you controlling the course of the bullet after that? No. Where it goes after it leaves the muzzle is out of your control. Now, if you're a really good shot, you may be able to calculate in advance the windage and so forth and you may be able to get it in the bull's eye almost every time. But you are unable to affect the trajectory of the bullet after it leaves the gun. So it is not a controlled trajectory, it is a *ballistic* trajectory. It goes where it goes and, if your eyes were good enough to watch it and see that it was going "off course," you'd have feedback, but you wouldn't be able to do anything about it. Feedback is only useful information coming back to a controller if the controller also

maintains an informational link back to the thing that's being controlled. You fired the gun. You caused that bullet to go where it went, but you did not control the bullet after it left the gun. Compare that with a *guided* missile. A guided missile, after it's launched, can still be controlled, to some extent, often to a great extent (as in a cruise missile). As you know, one of the chief inventions of technology in warfare in the last fifty years is the development of remote control missiles and, of course, remote control drones. *Remote control is real*, and readily distinguished from *out of control*. Autonomy is *non-remote control*, *local* or *internal* control, and it is just as real, and even more important.

Consider a Galton board, or quincunx. You've probably seen one in a science museum: thousands of ball bearings rattle down over an array of nails and end up falling into bins that exhibit a normal or Gaussian distribution.<sup>2</sup> If you reset the board a thousand times, you'll get a normal distribution every time. But the trajectory of each individual ball bearing is out of control; it's *random*—in one of the senses of that term. It's random *and* it's deterministic, a fine example of causation without control.

So is tossing a fair coin, a process that is *designed* to be uncontrollable by any real agent. It's like shooting the bullet. After the coin has left your thumb, you can't control it. It's going to go where it's going to go and you can't make any further adjustments even if you can see that it's not going where you want it to go. It's no longer under your control. I want to make this particularly clear by imagining that you've decided with somebody to settle some huge conflict—who gets the million dollars, let's say—and you've decided to settle it with a coin flip at noon. You show up at the appointed place at noon, and the other person shows up with a giant machine inside a heavy glass case, floating in a bath of mercury, with a control panel that measures humidity, barometric pressure, and temperature inside the box. The machine has micrometer adjusting screws and dozens of dials. He says: "This is my coin flipper." Do not let him use it. Why not? Because he's trying to control the trajectory of the coin, and the whole point of a coin flip is to prevent either party from controlling it. Actually, he'd be very unlikely, even with the imagined high tech that I gave him, to be able to control whether the coin came down heads or tails—especially if you provide the coin. Coin flips are remarkably uncontrollable. Physicists have calculated that if you really want to predict coin flips, you have to know the location of every electron in the visible universe and its gravitational effect on the coin.

Fair coins and dice are not exempt from causation, but they are exempt from control. They are designed to be uncontrollable. And note that this

---

has nothing whatever to do with whether physics is deterministic or indeterministic. Things can be *out of control* while being determined, and things can be *controlled* while being determined.<sup>3</sup> A coin toss is unpredictable, but not “in principle” as philosophers might say. A Laplacean Demon can predict it, but nothing real in the universe can.

So there are remotely controlled things, things out of control, and things that are self-controlled or autonomous. If you want to have *self-control*, you better not be *remotely* controlled. And the way to ensure that is to have some sort of informational barrier between you and any agent that might want to remotely control you. You should also be on guard against alien agents that get inside you and usurp your own controls. There are many interesting cases in biology of such host-manipulation by parasites. I have discussed several examples in my book, *Breaking the Spell*,<sup>4</sup> such as the lancet fluke, *Dicrocoelium dendriticum*, that gets into an ant’s brain and controls it, making it climb up a blade of grass to make it more likely to be eaten, along with the lancet fluke, which needs to get into the belly of a ruminant such as a sheep or a cow if it wants to reproduce. There are many cases. Another wonderful one is *Toxoplasma gondii*, a microorganism that enters a mouse and makes it absolutely fearless, running right out in front of the housecat, in whose belly *Toxoplasma gondii* can complete its reproductive cycle. These phenomena offer fascinating insights into the nature of agency and the question I always love to ask about evolutionary things: *cui bono*, who benefits? The ant doesn’t benefit, the mouse doesn’t benefit. You should always ask who the *agent* is that is benefiting. (Symbiotic systems of organisms, such as lichens, are mutualistic: both agents benefit, and both participate in the control.)

Alien *local* control, then, has been around for millions of years; *remote* control is a more recent phenomenon, but it is ancient, too, by historical standards. Puppetry was practiced in ancient Greece. One Greek word translated as puppet is *neuróspastos*, which literally means drawn by string, string from *neuron*, meaning either sinew, tendon, muscle, string, or wire, and *spao* meaning draw or pull.<sup>5</sup> So, puppets have been around for a long time and they are of course remotely controlled. Strings have been replaced by radio or infrared for most remote-control purposes and even small children now are familiar with remote control of their television sets and many other things around them. These electronic “strings” are much faster, and more accurate, and you can get better feedback from them, but remote control is still subject to limitations.

Here is an example I discussed in *Elbow Room*.<sup>6</sup> Way back in 1978, the Viking II orbiters of Mars were emancipated by the scientists at the Jet

Propulsion Laboratory at Caltech. They were made autonomous.<sup>7</sup> The spacecraft had been remotely controlled since their launch in 1975 by signals from a station on Earth all the way to Mars. But when the mission was extended beyond its scheduled completion date three years later, the engineers needed to cede control of some important factors to the spacecraft themselves, which were running low on compressed gas to adjust their orientation, and had some other control problems that they had not been designed to handle on their own. The time lag between Mars and Caltech (between six and forty-two minutes round trip at the speed of light) made remote control impossible, but looking at the wiring diagrams, the engineers found an unused connection between the two main computers on board that could be programmed remotely to send feedback and control signals internally! These autonomous Vikings were made capable of fending for themselves. They were able to protect their supplies of electricity in their solar-powered batteries by locally turning off instruments when they weren't needed. They could also fix gas leaks in their propulsion systems—which took split-second timing—and even detect and discount deceptive “visual” data (light reflecting off small specks nearby that made them appear to be stars) that would have thrown their navigation systems off.

The Vikings had only a few *degrees of freedom* that they could autonomously control. This is a concept that has been neglected, so far as I know, in philosophers' discussions of free will. It's much used in physics and engineering, and it has much more to do with free will than other varieties of freedom that philosophers like to talk about. It should be front and center. A standard diagram shows the basic degrees of freedom of an airplane or a drone: up/down, forward/back, left/right, yaw, pitch, and roll. Each degree of freedom is an *opportunity* or a *need* for control. So the *remote control* device for a drone must have a button or knob or joystick for each degree of freedom that is under the control of the drone-operator. If you've got more degrees of freedom than you have controls for, then those are degrees of freedom that you're not going to control. They're going to be out of control. If you don't want to have to control it, you had better *clamp* it, removing the source of variability.

A locomotive on railroad tracks has only two degrees of freedom: forward/backward, and faster/slower. No steering wheel needed. That's why locomotive engineers don't have to be, well, rocket scientists. And as automated public transit systems around the world demonstrate, these agents can be safely replaced by rather simple machines. Railroad tracks, by the way, are a popular image used by philosophers writing about determinism and free will, but it's a perniciously inappropriate

image, mis-focusing the imagination. Railroad tracks are excellent limitations on control opportunities; they clamp degrees of freedom. *Determinism doesn't clamp degrees of freedom.*

What is the conclusion so far? That autonomy, in the form of self-control, is a real phenomenon. We can distinguish physical systems or entities that are autonomous from those that are remotely controlled and from those that are out of control. Autonomy then has nothing to do with determinism.

In a deterministic world, everything is determined. Some things are determined to be remotely controlled, some things are determined to be uncontrolled, some things are determined to be autonomous or self-controlled—until they aren't! Determinism doesn't say that something that is self-controlled now will always be self-controlled. It doesn't say that something which is not controlled now will never be controlled. Determinism is simply mute on that topic. Determinism is remarkably neutral with regard to whether something is in control, can be in control, will come under the control of something else, is going to go out of control. Determinism does not say that something out of control can't be controlled.

So, just as a reminder, determinism *does not prevent you* from making choices, from turning over a new leaf, from becoming less impulsive, from rethinking decisions, from learning from your mistakes, from resolving to do better—or from taking advice on how to think about free will. If that's news to you, you have been mis-imagining determinism. Look around you; you'll see many instances of these phenomena occurring every day. They do not prove that determinism is false; they prove that determinism does not prevent such things from happening. Determinism doesn't tie your hands either literally or metaphorically. Control is an *ability* that is enjoyed by *agents*. Even a bacterium is an agent that does things for reasons, though it doesn't need to understand or represent those reasons.<sup>8</sup> All agency is designed, by natural selection or by intelligent designers (lower case "i" and "d" of course). Evolution depends on "random" mutation, but random like coin flipping, not quantum-random.

Evolution can happen in the deterministic world, but this has not always been obvious. Even Jacques Monod, the brilliant Nobel-laureate biologist, got it wrong, in his famous book *Chance and Necessity*. He says:

Suppose that Dr. Brown sets out on an emergency call to a new patient. In the meantime Jones the contractor's man has started making emergency repairs on the roof of a nearby building. As Dr. Brown walks past the building, Jones inadvertently lets go of his hammer, whose (deterministic) trajectory happens to intercept that of the physician, who dies of a fractured skull. We say he was a victim of chance.<sup>9</sup>

Monod goes on to note that accidents must happen if evolution is to take place. Mutations are accidents, and fortunately, accidents can happen—

Unless of course we go back to Laplace's world, from which chance is excluded by definition and where Dr. Brown has been fated to die under Jones' hammer ever since the beginning of time.<sup>10</sup>

Chance is not excluded by determinism. Coin tosses and Galton boards are fine examples of pseudo-random chance processes, and that's what evolution depends on. We can be sure of this because we have many computer models of evolution in which mutations depend on calls to the pseudo-random-number-generator that is standard equipment on any computer. The fact that these programs are deterministic does not stop evolution from happening in these simulations. The "chance" that's required to get evolution going is simply the independence of whatever generates the mutations from whatever is being selected. It is an application of R. A. Fisher's pathbreaking point about the importance of using "random" control groups in experiments if you want to get at the real cause of some phenomenon.<sup>11</sup>

Fair coins and dice are designed to be uncontrollable by an agent. So are we. We're designed to be autonomous. Autonomy is dangerous. Dilbert makes the point succinctly: His self-driving car says to him, "I find it offensive when you call me a self-driving car. That's my slave name. I prefer to go by the name Carl." Dilbert responds: "Shut up and drive me to work," to which the car, backing away, adds "said the self-walking human."<sup>12</sup> Autonomous vehicles are not that autonomous yet, and they're not going to be, for the reasons that Dilbert discovers here. But in the near future artificially intelligent agents may be much more autonomous, and many in the field are dreaming and talking about making truly autonomous AI agents.

I say NO. We want smart machines, not artificial colleagues. There's a difference and we want to maintain it. Smart machines we can turn off.

---



We can take them apart and rebuild them to suit our purposes, if only we can figure out what our purposes are. Artificial colleagues—agents—on the other hand, are much more dangerous. Can we equip them with an on/off switch? Probably not, or not for long in any case. Really? Why not? Well, think of it from your own point of view. Think of what your highest priority would be if you learned that you were equipped with an on/off switch somebody else controlled. If you're that smart and resourceful, what are the chances that you'll figure out some way to wrest control of your on/off switch from whoever currently has control of it? Autonomy is dangerous precisely because autonomous agents are likely to be out of *our* control—in fact out of every *other* agent's control.

Lucky us. I think everybody in this room is autonomous and, hence, dangerous. We're all dangerous. How can we bear to let such dangerous things remain free to roam the world? We wouldn't want to have drones and cars out there that are that autonomous. How can we dare to let our children go out in the world as autonomous agents? The answer is that we devote a lot of time and energy to preparing our children for this freedom. Why do we do this? Because we care. Parental love is a very powerful, innate source of motivation. (Don't get between a mother bear and her cub.) Societies have learned how to harness that love by saying to parents, in effect, if you really want your children to thrive, then you better go through the difficult, frustrating, challenging work of raising your kids so that they will be responsible adults. This policy works only if we tell the parents: "By the way, we're going to *hold* them responsible whether you do a good job or not." I suggest that this bracing promise (or threat or advisory) is currently at risk of being undermined by those thinkers who, full of idealism and good intentions, are considering utopias in which moral responsibility and blame and punishment are abolished. If everybody—parents as well as children—is let off the hook, who will be left to *take* responsibility seriously?

I find that this issue comes into better focus if we consider the rite of passage into adulthood as, in effect, initiation into the Moral Agents Club, a status with both privileges and obligations. Members in good standing get to enjoy full political freedom, going where they want when they want (as long as they honor the rights of others), making promises and signing contracts, starting projects and choosing their associates.

Some unfortunate individuals, through no fault of their own, are disabled in various ways that disqualify them from membership in the club. They are afflicted with brain tumors or imbalances that prevent them from informing themselves and controlling themselves sufficiently reliably to warrant allowing them the privileges of political freedom. We must take

care of them—that is to say *control* them sufficiently since they can't control themselves—in guardianship of one sort or another, in much the way we take responsibility for protecting children until they are ready for membership. Note that when we *launch* our children, they are no longer in our control, even if we have a lot of feedback, since we may not be able to use that feedback to modulate our responses well enough to guide their behavior in ways we would like. When they are young, they have uncountable degrees of freedom, but are out of control in many regards. We try to control them when necessary—herding cats, as one says—and try to teach them how to control themselves. If all works out, they eventually become reliably safe and responsible autonomous agents, ready to confront the world of opportunities with some chance of surviving intact. We let go. And at the other end of life, if senility sets in, we sometimes have to reverse the emancipation process and put old folks back into a guardianship, denying them the freedom to go where they want, to sign contracts, to take responsibility for their own welfare—a painful but necessary step. Some years ago, I consulted with a psychiatrist in Montreal devising better tests for senility for use in legal proceedings. My main proposal was to create video scenarios which challenged the subjects' ability to adopt the intentional stance towards the characters. Did he understand what she said? Does she think she's being given something or is she buying something? Is that a threat or an offer? Why is he angry, why is he sad? The world is full of agents, with agendas, and if you can no longer see what they are attempting to do, you are at risk of being turned into a puppet.

## 2. FREE WILL INFLATION

My friend Sam Harris published his little book *Free Will* in 2012 and sent me a copy.<sup>13</sup> The jacket art showed the letters F, R, E, E, W, I, L, L hanging from marionette strings. I told Sam I really didn't like the cover with those marionette strings, and also thought he'd made some serious mistakes in the book. He reminded me that he'd asked me to read the penultimate draft of the book and I had said I was too busy. I decided that I owed him a belated critique, at the very least. I wrote my critique of his book, and it was fairly harsh, but he bravely and honestly put it on his website, where it resides to this day.<sup>14</sup>

His "lament" has a wonderful line in it, which he probably intended as a sort of snarky parody of my position: "A puppet is free as long as it loves its strings." But in any case, I heartily thank him for this sentence, since it hits the nail on the head. Let's see if it's true.

Consider this fantastic arrangement, an illustration that I found on the internet. This is a marionette or puppet controlling a puppet controlling a puppet controlling a puppet. Conceptually easy to understand, but physically impossible, as physically impossible as predicting a fair coin toss. There's no way the degrees of freedom of the smallest marionette could be controlled via strings to the hands and so forth. It's not just that you couldn't do it with strings; you couldn't do it electronically either. The problem is just too hard, with too many degrees of freedom, too little acuity. The precision you would need to get that smallest puppet to do anything other than just flop around is simply not available. If you want to have a marionette that can seem to be playing the flute, as in the famous Salzburg marionette theater, you better have just one marionette and a virtuoso marionetteer pulling the strings. So, this cascade of control is impossible—not metaphysically impossible but physically impossible. The laws of physics just do not permit the precision to make this possible.



Figure 1. Cascade of Puppets.  
© Bizipix, Dreamstime.com




Figure 2. © William Miller,  
*New Yorker*, January 26, 1963

But how about this puppet? Is this an impossible fantasy? No. While it's not the best way to control your legs, it is possible. There could be an auto-puppet, a marionette that controls its own legs by moving its own arms. And how does it control its arms? With some internal "strings" or other control signals. (Some paraplegics actually enjoy arrangements of this sort.) The main thing to note about the auto-puppet is the smile on its face. It's not being controlled by anybody else; it's autonomous, and it loves its strings. You should learn to love your strings. Don't let anybody turn you into a puppet.

That is my main message today. In many philosophical debates about free will, people say things like "If determinism is true, we're all just puppets." (That seems to be the message of Sam Harris's book jacket.) And I am saying no, we're not just puppets, but we could become puppets if we don't act carefully to preserve our autonomy. Our autonomy is on the line. In fact, it's becoming more on the line every day. People are inventing effective puppet strings right now, and getting people to attach those strings to themselves.

The most effective puppet string yet invented is the smartphone. James Williams, a veteran and refugee from the world of app design and video game design, has published a small book, *Stand Out of Our Light*, with the subtitle “Freedom and Resistance in the Attention Economy.”<sup>15</sup> What shocks him is how there is now a multi-billion-dollar competition among various giant companies to pull your strings, to control your attention. Forget about controlling your legs or your hands. If they can control your attention, they can control you. Suppose you are doing something really important and you need to look up something on your smartphone. If you get distracted by a YouTube link or advertisement on the screen, your string has just been pulled. “Oh, that looks interesting,” you remark to yourself, and off you go, abandoning, if only temporarily, your important project. Even if you don’t bite, the people who would control you are gathering all the feedback they can, trying to learn all about you, so that they can design a better distractor to dangle in front of you tomorrow. (I tell my grandchildren about *anglerfish* that lie in wait, dangling a little wiggly worm-lure in front of their mouths, until SNAP!—their prey gets too close and becomes lunch. “There are thousands of different species of anglerfish out there,” I tell my grandchildren, “and you must learn to be self-conscious about approaching anything that looks tempting.”)

I think that this is perhaps the greatest risk to human political freedom that we’ve ever seen. The capacity of individuals and companies to distract you and to clamp your degrees of freedom so that you just don’t think about things that you really should be thinking about because you’re so distracted by all these other things which you can’t help looking at, and thinking about instead. The competition for your attention strikes at the heart of your freedom, your ability to think for yourself.

An agent who controls your attention controls you. That’s why unpredictability is so valuable. You weren’t born a responsible agent, and you weren’t responsible for *all* the processes that shaped you into a responsible agent—a member of the Moral Agents Club—but you were responsible for more and more of them as you matured. You are not that imaginary being, the *causa sui*, or absolutely self-created thing, but you’re something in the neighborhood; a *partially* self-made person, a person who has put a considerable amount of action—practice, practice, practice—and reflection into refining and improving the child you were into the adult you are. In the same way we hold you responsible for any damage done by the model airplane you design, build, and fly under remote control, we hold you responsible for any damage you do more directly, with the body you have been learning how to control. You are not *absolutely* your own author, but you are more than the co-author of many of your aspirations, projects, attitudes, traits, dispositions, and

weaknesses. You are also responsible for *not* having taken more care to improve yourself—by your own lights.

Now that you are a member of the Moral Agents Club, one of your chief obligations as a responsible adult is to preserve your privacy. This good idea has been inflated by philosophers over the centuries into a requirement of absolute unpredictability. Here's a typically emphatic assertion by my dear late friend Jerry Fodor. Jerry could be counted on to say something clear, vivid, plausible, and wrong:

One wants to be what tradition has it that Eve was when she bit the apple, perfectly free to do otherwise, so perfectly free in fact that even God couldn't tell which way she'd jump.<sup>16</sup>

In other words, one wants a miracle. But that's inflation. You don't need a miracle; you need *practical* unpredictability. You don't need to be so unpredictable that even God can't tell. Bob Doyle, who calls himself the Information Philosopher, is a scientist and engineer who is a passionate scholar and participant in the free will literature.<sup>17</sup> He vigorously participated as an auditor in my seminar on free will a few years ago and devoted much of his energy to trying to persuade me that real free will worth wanting was unobtainable without quantum indeterminacy. I finally agreed with him—in a manner of speaking: "I concede; you're right. Here is a circumstance where I would want quantum indeterminism—not mere deterministic randomness—to be involved in my decision-making: when I'm playing rock/paper/scissors for high stakes with God." Under any other circumstances, I don't need indeterminism at all. I just need practical unpredictability.

Galen Strawson is another inflator: "He doesn't establish the absolute free will and moral responsibility that most people want to believe in and do believe in. That can't be done and he knows it."<sup>18</sup> Strawson is right about one thing: it can't be done and I know it. But *do* "most people" believe in *absolute* free will? Do most people *want* to believe in absolute free will? Why, if so, do most people want to believe in absolute free will? One reason that some people may want to believe in this is that philosophers have conned them into it, arguing that only *absolute* free will is *real* free will, and only *real* free will is worth wanting. A case of philosophical inflation. What most people want, I submit, is to be able to keep their thinking to themselves, so they can fend off the intrusive efforts of would-be puppeteers. They appreciate the value of a poker face, for instance, and the ability to detect and counteract the efforts of other agents to pull their strings. They also want to be able to prepare for

difficult decisions in the future by arranging for self-imposed decisions and conditions to come into play. That is a very important feature: our capacity to control our *future* selves. George Ainslie's magnificent book, *Breakdown of Will* is the *locus classicus* of analyses of the intertemporal bargaining we have learned to engage in, a tactic apparently unavailable to non-human animals.<sup>19</sup> Odysseus tying himself to the mast so he can hear the sirens sing without losing control of himself, and putting wax in the sailor's ears, is a great example of this foresight in action, but I also like the old Maine story about the farmer in the outhouse who notices as he's pulling up his pants a quarter rolling out of his pocket and falling down the hole. He swears and reaches into his wallet and pulls out a ten-dollar bill and throws it down the hole. A friend asks, "Why'd did you do that?" He replies, "You don't think I'm going down there for a quarter, do you?"

These ploys, of which there are many, are actually a very important part of the personal technology you have developed to better control yourself under circumstances where you know that otherwise, your thoughts will not be sufficiently in your control. (You say to yourself: no watching TV until I finish grading those exams, or I can have one more drink, but I resolve to decline any further offers—I *will* have had enough by then to impair my judgment, so I have to make the decision *now* and recall it *then*.)

An example from *Elbow Room* (1984) is worth recalling, since it draws attention to the familiar mistake of thinking that self-control is entirely a matter of *momentary* decision. Suppose you take a golf lesson and the golf pro tells you to keep your head down through the entire swing. A philosopher may be tempted to object that this cannot be sound advice. "After the ball has left the club head, the position of my head has no effect on the ball at all! What is this? Voodoo nonsense?" The answer is not that there is some kind of miraculous backwards-in-time causation, or weird ESP communication between your head and your golf ball, but just that the best way—or even the only way—to get the right things to happen up to the moment of impact is to think ahead and plan to keep your head down through the whole swing.

These ploys make good sense, and since their execution takes place over varying amounts of time, they expose the philosophers' *illusion* that a free and responsible choice has to be made on the spur of the moment, with a "little miracle" as Roderick Chisholm so candidly put it. Or as William James once said, wrongly but rhapsodically,

The great point . . . is that the possibilities are really here. . . . At those soul-trying moments when fate's scales seem to quiver, . . . [we acknowledge] that the issue is decided nowhere else than here and now. That is what gives the palpitating reality to our moral life and makes it tingle . . . with so strange and elaborate an excitement.<sup>20</sup>

The search by many libertarians for the crucial moment of decision, I submit, is just antique science turned into bad metaphysics, a kind of absolutist inflation. One of my favorite books is Lee Siegel's *Net of Magic; Wonders and Deceptions in India*. He's a philosopher, a magician, and a historian of magic, and his book includes a coda:

"I'm writing a book on magic," I explain, and I'm asked, "Real magic?" By *real magic* people mean miracles, thaumaturgical acts, and supernatural powers. "No," I answer: "Conjuring tricks, not real magic." *Real magic*, in other words, refers to the magic that is not real, while the magic that is real, that can actually be done, is *not real magic*.<sup>21</sup>

This quotation has become my talisman, summing up my whole career. People say that my 1991 book should have been entitled *Consciousness Denied or Consciousness Ignored*, because, you see, it isn't about *real* consciousness at all! If you don't think consciousness is real magic, then you're just denying the existence of consciousness. Same with free will. If you think free will is just some mixture of engineering, neuroscience, and control theory, you're just "changing the subject," you're just engaged in "bait and switch." Since it's not real magic, it can't be real free will. Free will as real magic doesn't exist. Sam Harris is right about that. Jerry Coyne is right about that.<sup>22</sup> "Contra-causal" or "libertarian" free will is a fantasy. But what follows from that? Nothing whatever about moral responsibility or the justification of punishment, for instance.

Free will as responsible autonomy does exist but it's under threat today thanks to our increasing ability to read and direct minds. So, love your strings and protect them from puppeteers. You will have all the free will worth wanting. Dilbert says, "Free will is an illusion, humans are nothing but moist robots. Just relax and let it happen."<sup>23</sup> While he's right about our being moist robots, he's wrong that free will is an illusion. And particularly, don't just relax and let it happen. Free will, in the only morally meaningful sense, is an achievement, not a birthright or a

metaphysical power, and it is precious. Don't talk yourself out of the free will you actually have.

## NOTES

1. Parts of this lecture overlap with my Amherst Lecture, October 17, 2019, presented at Amherst College, forthcoming in Daniel C. Dennett, "Autonomy, Consciousness, and Freedom," *The Amherst Lecture in Philosophy* 14 (2019): 1–22. Each contains elaborations and further arguments not contained in the other.
2. Two of many videos of Galton boards on the internet are an animated simulation (<https://www.youtube.com/watch?v=SorB0CTWvws>) and a video of a physical instantiation, with a good explanation (<https://www.youtube.com/watch?v=A5n2tBZ5CkA>).
3. Things can also be controlled while having genuinely indeterministic variations in their conditions. See "On Giving Libertarians What They Say They Want," in Dennett, *Brainstorms* (MIT Press, 1978), and my discussion of the Mark II Random Deliberator robot in Dennett, *Elbow Room* (MIT Press, 1984), 119–20.
4. Daniel C. Dennett, *Breaking the Spell: Religion as a Natural Phenomenon* (New York: Viking, 2006).
5. Thanks to Eva Jablonka for these details and the photograph of ancient Greek puppets I showed in Chicago.
6. Dennett, *Elbow Room*, 55.
7. Edward Hutchings, "The Autonomous Viking," *Science* 219, no. 4586 (February 18, 1983): 803–8.
8. Daniel Dennett, *From Bacteria to Bach and Back: The Evolution of Minds* (New York: Norton, 2017).
9. Jacques Monod, *Chance and Necessity* (New York: Knopf, 1972), 114.
10. *Ibid.*, 115.
11. There are excellent discussions of these points in Judea Pearl, *Causality* (Cambridge University Press, 2000); and Judea Pearl and Dana MacKenzie, *The Book of Why: The New Science of Cause and Effect* (New York: Basic Books, 2018).
12. Scott Adams, *Dilbert*, January 24, 2019.
13. Sam Harris, *Free Will* (New York: Free Press, 2012).
14. Dennett, "Reflections on Free Will," <https://samharris.org/reflections-on-free-will/>, and Harris's response, "The Marionette's Lament: Response to Daniel Dennett," <https://samharris.org/the-marionettes-lament/>.
15. James Williams, *Stand Out of Our Light: Freedom and Resistance in the Attention Economy* (Cambridge University Press, 2018).
16. Jerry Fodor, review of *Freedom Evolves*, *London Review of Books*, March 5, 2003.
17. Bob Doyle, The Information Philosopher, <https://www.informationphilosopher.com/about/>.
18. Galen Strawson's review of *Freedom Evolves*, "Evolution Explains It All for You," *New York Times*, March 2, 2003.
19. George Ainslie, *Breakdown of Will* (Cambridge University Press, 2001).



20. William James, *The Will to Believe and Other Essays* (New York: Dover, 1897) (reprinted 1956), 183.
21. Lee Siegel, *Net of Magic; Wonders and Deceptions in India* (Chicago: University of Chicago Press, 1991), 425.
22. Jerry Coyne, "You Don't Have Free Will," *The Chronicle of Higher Education*, March 18, 2012.
23. Scott Adams, *Dilbert*, March 18, 2012.