

# Serial Hyperspecializers, and How They Think\*

Elijah Millgram

Department of Philosophy  
University of Utah  
260 Central Campus Drive Rm 341  
Salt Lake City UT 84112  
lje@philosophy.utah.edu

February 11, 2006

©2006 Elijah Millgram. All rights reserved. DRAFT---Do not cite or quote without permission. Comments, criticism and suggestions welcome.

The theory of practical rationality (that is, the rationality of decision, choice, action and related matters) dates back at least to Plato and Aristotle. That makes it well over two millennia old, and over that time, we have, to all appearances, built up a substantial body of results: results that allow us to understand what we are doing, when we make up our minds what to do, and which also amount to guidelines that can help us do a better job of it. In my view, however, the body of theory we have developed in this area is unusable. It is, I am going to claim, not so much mistaken as irrelevant, because philosophers have made a very basic error, and an error that, I will suggest, has damaged much other work in moral philosophy, including work

---

\*I'm grateful to Jay Bernstein, Karin Boxer, Sarah Buss, Ben Crowe, Havi Carel, Christoph Fehige, Gabi Juvan, Kathrin Koslicki, Chandran Kukathas, Tamar Laddy, Brenda Lyshaug, Peter Momtchiloff, Huw Price, Wayne Waxman, David Wiggins and audiences at a Rosenblatt Free Lunch and a Law School Brown Bag at the University of Utah for helpful conversation. I also thank Chrisoula Andreou, Pepe Chang, Steve Downs and Anya Plutynski for comments on earlier drafts.

on agency, on personal identity, and on substantive moral and ethical theory. If I am right, at least in the theory of practical reasoning, and quite possibly also in these related areas, we are going to have to start over, almost from the very beginning.

## 1

Aristotle famously asked, in his *Nicomachean Ethics*, what the human *ergon* was: what it is that people *do*. I am going to try to answer that question, and so I will argue that we ourselves are a solution to a design problem. But I will differ from Aristotle in my account of what the human *ergon* is.

The approach is unusual enough nowadays to require a cautionary remark. Because the current fashion is to explain or replace design-based explanations of, especially, living things with evolutionary explanations,<sup>1</sup> and to accept functional descriptions of an organism only if they can be construed as shorthand for claims about identifiable pressures for that function operating within the history that produced the organism, I am likely to be understood as claiming that the design solution I will describe is an adaptation for which we have been selected. The biological world, and we ourselves as part of it, are products of natural selection, and so of course I want my claim to be compatible with a plausible evolutionary history. But I do not mean to imply that we were selected to be the design solution I will identify. (In fact, I will argue in due course that this hypothesis would be quite implausible.) My question is not, How did it happen? but, What is it good for? Compare: eyes solve the problem of seeing, whether or not that explains how they arose.

On the contrary, rather than assuming a sort of reductionism—of functionality to a history of natural selection—I think that the interesting question to ask about natural selection is what design problem *it* solves. The observation I will use to frame my discussion of humans is that natural selection is a solution to the problem of producing creatures suitable to available ecological niches.<sup>2</sup> How does it go about that?

---

<sup>1</sup>‘Especially,’ because the approach has been adapted to other areas as well; see, e.g., Boyd and Richerson, 1985, Boyd and Richerson, 2004.

<sup>2</sup>That itself can be explained as the outcome of a selection effect. Suppose that there are organisms occupying some small initial set of niches. And suppose that features that allow their descendants to adapt (or adapt more rapidly than competitors) to new niches are heritable. Then organisms that happen to have those features will have descendants occupying more new niches than organisms that do not, and the occupiers of those new niches will be more likely to have descendants occupying still further niches. The long-

Here are three large observations about evolution as it operates in the biological world. First, Darwinian selection produces species that fall along a spectrum, with the very weedy species at one end, and the very specialized species at the other. Weedy species (remember: weeds are the plants that you can't keep out of your garden) are the ones that invade niche after ecological niche, and because they travel from one to the other, they're not necessarily particularly tailored to any of them.<sup>3</sup> Specialized species, on the other hand, are often fitted to their places in an ecology with a memorable and jewel-like precision. Think of the migrating bird, spectacularly engineered to exploit resources that are located some thousands of miles apart. When it touches down in the marshlands that are the southernmost terminus of its journey, it turns out to have long stinky legs that are perfectly suited to standing in the marsh, and the scoop-like beak that is perfectly suited to preying on the fish and frogs that it will find there. When it arrives at the northernmost end, it turns out to be perfectly fitted to nest in the crevices of the cliff which it will find *there*.<sup>4</sup> Specialized species take advantage of ecological niches in ways that weedy species normally cannot, but they pay a price in fragility; too often, when the niche disappears, so does the species: as the forests of the Pacific Northwest are cut down, the spotted owl vanishes.

Second, to pull a metaphor over from another discipline, in any species that relies on a nervous system to compute its behavioral outputs, the adap-

---

term effect will be an ecosystem predominantly composed of populations of organisms with features that allow them to spread, through the operation of natural selection, into still further niches.

Examples of such features are sexual reproduction and mechanisms that contribute to the reproductive isolation of a species. (That is to say, that many organisms comprise species, in the sense of the term given by Mayr, 1984, esp. p. 539, is itself such a feature.) I expect that vulnerability to mutation is another such feature, and that there are optimum ranges of vulnerability that are selected for in this way.

Finally, assume that scientists (or more generally, theorizers) arise only in environments with high levels of biodiversity. Then scientists (and theorizers) are likely to find themselves possessed of features of the sort in question. That is, there is a point of view from which the selection effect I have just described looks like an anthropic selection effect. For some discussion of anthropic selection, see Roush, 2003.

<sup>3</sup>The contrast between weedy and specialized species should be not confused with another familiar contrast, that between R- and K-selected species; see A word of caution: not all variants of the niche concept allow for niches to be independent of the organisms that occupy them; to describe organisms as traveling from niche to niche presupposes that sort of independence. For a recent overview of niche concepts, see Odling-Smee *et al.*, 2003, pp. 37ff; for a note of caution, see Arthur, 1987, p. vii.

<sup>4</sup>For gorgeously shot illustrations, see Debats *et al.*, 2001; for a cinematic illustration of another equally extreme adaptation, see Jacquet, 2005.

tations that we see are both in hardware and software. The bird's legs and beak are special-purpose hardware designed (or, since it's *natural* selection, 'designed'... but I won't repeat this important qualification throughout) for the narrow range in which the species lives. But the bird's cognitive architecture is just as much a part of the engineering solution to the problem posed by the niche. When the seasonal cues trigger the overriding urge to fly north, that urge has an adaptive function: to get the animal from one location, where it has been using resources of one kind, to a second location, where it will avail itself of resources of a different kind. The program is, in cases like these, part of the package, one that is as invariant as the hardware—a fact we mark by calling the bird's impulse to migrate an 'instinct'.

Third, much of the time, natural selection operates on a use-it-once-and-throw-it-away model. At the level of individual organisms, the process churns out many (often vastly many) near-copies of a design solution, throws away the ones that don't work... and throws them *all* away after a lifespan.<sup>5</sup> But it is not just individual organisms that are disposable, and what is interesting to us now is a contrast we see as we move along the specialized-weedy spectrum. Proliferating throwaway species is how sexually reproducing life at the specialized end of the spectrum fills the available ecological spaces: a niche opens up, a specialized species evolves to fit it; the niche goes away, and the species is discarded. On the weedy end of the spectrum, however, while the individual organisms are just as much disposable hardware as ever, planned obsolescence is no longer the most salient feature of the species as a whole. Weedy species are reusable, and a more economical alternative to Nature's all-too-common NASA-style practice of throwing out design solutions after a single use.

Now here's a neat idea that the as-if designer might have: to build a species that was *both* weedy and specialized, that was in fact weedy *by* being specialized. The trick would be to use less specialized hardware, and put the capacity for specialization into the software. The weedy-but-specialized species would run different software in different niches, and the movie-perfect fit would be worked, not by the beaks and legs and so on, but by behaviors that changed from niche to niche. How would such a design work?

There is evidently more than one way to answer that question, and I am going to focus on two alternatives, one of which is more ambitious than the

---

<sup>5</sup>Continuing the list of features started in note 2, sexual reproduction is notoriously a solution that requires death in order to work efficiently; if older copies of gene mixes stick around, the genetic profile of a population won't adapt as efficiently as if they all die off. I will return to the topic of death in due course.

other. Let's call the less ambitious solution *Piltdown Man*, for reasons I'll explain shortly. Piltdown Men are born, identify the environment they're in, load a program appropriate to that environment from a database of available strategies, and run the program till they die. Piltdown Man can occupy what would intuitively look like many different locations in many different ecologies, and not the way that species like mint do: on the seashore, they might be fishermen; in the mountains, they might be yak herders; in the cities, they might be merchants; on the plains, they might be farmers. However, Piltdown Man can only occupy relatively stable niches, because those software libraries have to come from somewhere, and if we don't have a Kubrickian monolith hand-coding them, they will be produced by a process that is either natural selection, or that (like meme selection) resembles natural selection in being, by human standards, slow. Since stable niches are a small fraction of all the niches, Piltdown Man will be distributed relatively sparsely through the strategy space.

Piltdown Man is already a departure from the throwaway approach so typically taken by natural selection. Piltdown Men are still individually disposable, in two senses; first of all, of course, they die; second, if the niche drops out from under them, the former Piltdown Man can't become a yak herder or merchant, because his behaviors are fixed by his initial download (to drag in a biological metaphor, by his initial imprinting), and he starves (or ends up in a refugee camp). But the *species* is reusable: when one niche evaporates, new Piltdown Men are produced to occupy other niches. (The yak herders die off, but their children become farmers.) There's no need to start over and develop a new species, with hardware that is painstakingly (at great cost, and over a long period of time) redesigned for the new environment. Piltdown Man is a much less wasteful, much more efficient approach than natural selection's default strategy.<sup>6</sup>

The more ambitious version of the neat idea pushes it further in a couple of directions. First, you could make the individual creatures reusable rather than disposable. After all, why throw out perfectly good hardware when the program it's running is no longer up-to-date? It's much more efficient to design such a species to look for new and better software when it's cued that the old software is no longer getting decent results. Let's give this part of the design solution a name: creatures like this are *serial specialists*.

---

<sup>6</sup>There are halfway stages to Piltdown Man: think of dogs, which are shaped largely by artificial rather than natural selection. Here, the adaptations are still in hardware and hardwired programs (dachshunds, designed to go down burrows; boxers, designed for bear baiting; hunting dogs; tracking dogs; herding dogs), but the species as a whole is less disposable, because there are those relatively disposable breeds.

Second, instead of relying on a fixed database of programs, you could have the individual creature reprogram itself on the fly, to meet the demands of a new ecological niche. (That is, while these creatures will sometimes search for and adopt off-the-shelf software, when that is the cost-effective option, sometimes, and especially when suitable software is not already available, they will develop their own.) Doing it that way opens up the option of taking advantage of extremely transitory niches: VLSI engineer, comics inker, Cobra gunner, French professor specializing in eighteenth-century poetry, adventure travel agent, director of cinematography. . .

What's more, as the list I just began suggests, it opens up the option of more elaborate adaptations, to niches that turn out to be more exotic and delicately constructed than anything in nature. Suppose that Piltdown Man's software results from a process that, over long periods of time, shapes it to match a stable niche. Over long periods of time, niches will change at the edges, maybe quite rapidly; this means that programs which don't change as rapidly as all that need to be a little bit robust in the face of such fluctuations, that is to say, a little bit weedy: thus, not quite all the way to the specialized end of the weedy-specialized spectrum. (Piltdown Men will be farmers and merchants, but not the peculiar cattle breeders one runs into nowadays, who purchase prize cows at auctions and frozen bull semen by mail, artificially inseminate the cows, and FedEx frozen embryos to be reimplanted in cows elsewhere; and they will not be hedge fund brokers.) However, with on-the-fly reprogramming, that constraint is removed: the specialization can be as filigreed and fragile as the underlying computational resources allow. The improvement over Piltdown Man which we're now considering allows, not just specialization, but *hyperspecialization*. At this point, we've reached a design solution with which we're intimately familiar (although with important qualifications I'll get to in due course)—from first-person experience.

A moment ago, I compared Piltdown's Man's alternative source of software to the alien monolith in Kubrick's *2001: A Space Odyssey* (1968), and this gives us a way of saying why we shouldn't worry that, for all we know, we're really Piltdown Men. And it also gives us a way of saying why we wouldn't want to be. On the first count, without the ability to do on-the-fly self-reprogramming, the proliferation of specialized forms of life that we see in human societies could only be managed, given the brevity of history, through the intervention of more intelligent outside management. But there are no space aliens or Prometheuses bequeathing us the exotic skills of modern-day cattle breeding or derivatives trading. On the second count, Piltdown Man seems like a peculiarly *dependent* species. The Nor-

wich Terrier is a very recent breed of dog that gives birth only by Caesarean section, and that could not exist without ongoing human intervention; Piltdown Man, in a social world that is reasonably complex by our lights, would be a little like *that*.

If we aren't Piltdown Man, what are we? *We are serial hyperspecializers*. That's the right ecological description of the species *homo sapiens*, and our unusual implementation of the weedy-species strategy. William Tenn once wrote a novel suggesting that humans were a perfect fit for the sort of role occupied by rats and cockroaches,<sup>7</sup> and while people do seem to be able to fill this role when they have to (and historically they often *have* had to), what they do better (and what they do when they have the opportunity) is to mimic—and displace—biodiversity with cultural diversity. Humans change the ways they cope with and fit into their environments, again and again and again, and this has important consequences—consequences for how we should understand our own psychologies, for what rationality must be for creatures like us, and for the assessment of much previous philosophy.

## 2

Migrating birds come equipped with a set of desires (or anyway, that's how we might as well think of them for now<sup>8</sup>) that match the environment in which they're going to find themselves. The desires are triggered by cues in the environment, but such creatures are not equipped to think about, redesign, or resynthesize their desire set. Our imaginary Piltdown Man, who downloads one behavior-controlling software package over the course of his lifetime, works almost the same way, and the claim I want to put on the table now is that Piltdown Man represents the mainstream philosophical conception of the human being. A survey of the many philosophical views to which this accusation can be made to stick would be a survey of much of the history of philosophy, and of much contemporary analytic philosophy as well. In the limited space I have here, I am going to touch on the closely intertwined debates about practical rationality, agency, personal identity, and substantive moral theory.

The standard model of practical reasoning (or figuring out what to do) is still instrumentalism, the view that your ends or goals or desires are just

---

<sup>7</sup>Tenn, 2001; for an animated film with a very similar view of humanity, see Laloux, 1973.

<sup>8</sup>Sterelny, 2003, ch. 5, argues that calling items like these desires or preferences probably attributes too much modularization to the cognitive architecture, and while that's a qualification I'm happy to allow, it doesn't affect the present point.

somehow given to you (hardwired in, or burned in by your upbringing, or prompted by cues in your environment...), that figuring out what to do is figuring out how achieve your ends, and that there's nothing that counts as thinking about what *to* want in the first place, or what your goals *should* be.<sup>9</sup> That is, the default philosophical conception of practical rationality conceives of people as scarcely one step away from creatures like migrating birds. Instrumentalism is a good theory of practical rationality for Piltdown Man (recall: a creature distinguished from us in the first place by its inability to reprogram itself for new environments).

But let's stop and ask how creatures designed as serial hyperspecializers *would* have to think about what to do. These creatures solve a design problem, that of occupying novel and ephemeral ecological niches, by producing behavior that is specialized to the niche. Since the niches are paradigmatically novel, the problem cannot be solved by using prestored guides to behavior (such as the urges, triggered by environmental cues, that we noticed in the migrating birds, or, in human beings, stable clusters of memes transmitted culturally). Rather, the behavior-guiding goals have to be computed on the fly. Now, that point is broader than I have just made it sound, so let me adjust it before moving on. An organism's controlling software need not rely exclusively on goals (or desires or ends or preferences); desires or goals are elements of just one of the available control systems, and by no means the most sophisticated of the lot. Accordingly, I will be supplementing talk of goals with more general notions, such as standards, priorities and guidelines.

If the task has to be performed intelligently, and if we allow that intelligence requires thinking, then it has to be possible for such creatures to rethink their top-level goals, standards, priorities or guidelines. (That is, in an older philosophical jargon, they have to be capable of deliberation of ends.) The puzzle then is: How can a strategy like that be implemented?

I don't have a complete recipe, but I do have a couple of the ingredients on hand. If one is developing standards, priorities, etc., to govern activity within a niche, and if one wants the standards, etc., to be reasonable guides to action, one had better explore the niche first. A reasonable way of exploring a novel niche is to try things out, and see how well they go; but of course in assessing how well they go one cannot invoke the standards one has not yet developed. So we should expect serial hyperspecializers to come with cognitive equipment that can be used for such assessment, and we should expect it to be relatively rough-and-ready (that is, not to presuppose too

---

<sup>9</sup>For a recent and sophisticated defense of such a view, see Vogler, 2002.



much about the structure of any niche in particular). One such tool might plausibly be a signal that tells the creature when it's doing better, and when it's doing worse; armed with such a tool, a serial hyperspecializer exploring a novel niche will be able to sort the options it has tried out into an system of rankings—which perhaps can later be refined into a more elaborate system of standards.

And in fact we do find such a cognitive tool in human beings. The most generic form of pleasure, and of its contrary, displeasure,<sup>10</sup> responds to changes in one's well-being, and represents, roughly, the first derivative of one's welfare (that is, it indicates the rate of change). If you feel good, things are getting (or have just gotten) better; but after a while, that feeling will fade. If you feel bad, things are getting (or have just gotten) worse; but after a while, modulo malfunction in the signaling system, you will adjust, and the feeling will fade. These signals are an important input into the deliberation of ends: disappointment prompts people to adjust their ambitions downward, by giving up or scaling back their goals; the elation produced by success prompts them to adopt new and more ambitious goals. In other words, in human beings, goals or ends are adjusted and selected on the basis of the input provided by pleasure and displeasure.<sup>11</sup> Hedonic tone is cognitive equipment that looks to be part of a solution to a problem, that of developing appropriate standards for novel niches.

It is not the only such piece of equipment. When one is exploring a novel niche, one will do much better at it if one's attention is directed to the features that will be important for, among other things, developing a system of standards and priorities appropriate to the niche. So you would expect a serial hyperspecializer to come from the factory equipped with a suitable attention guiding device. Since niches differ from one another, the device must be quite flexible; but it must be able to pick out candidates for exploration with better-than-chance payoffs. And indeed, when you take a look at humans beings, they are equipped with a signaling system that does this, namely, interest (and its complement, boredom).<sup>12</sup> Philosophy is a domain in which one audience of this paper will have seen this piece of equipment in action. When you train a philosophy student, you train him to have a sense of what is philosophically interesting; that is, you are training his philosophical nose. Astoundingly enough, it is possible to produce reli-

---

<sup>10</sup>So-called to distinguish it from pain, the signal that one's body has been damaged; see Pitcher, 1970.

<sup>11</sup>This account of the cognitive role of pleasure and displeasure is developed and argued for in Millgram, 2005, ch. 2.

<sup>12</sup>This account of boredom and interest is suggested in Millgram, 2004.

able noses (though very hard to produce them reliably), that is, responses of interest and boredom that can be used to guide, and *successfully*, choices between avenues of intellectual exploration. (Astoundingly, because it is not as though the problems of a professionalized philosophy were the life-or-death demands made on our Pleistocene ancestors.)

On the instrumentalist way of seeing things, success or failure must be spelled out in terms of whether one is achieving already given goals or ends. So notice that on the view I am sketching, the direction of explanation (and of justification) is, for the most part, the other way around. When your hedonic responses tell you that you are doing well, or doing badly, those signals (although they can be theory-laden, and although they are, once again, trainable) can pull free of the goals and desires you already have—and must do so, at least sometimes, if they are to do their job. When you are getting everything you want, and you still feel like you are failing, that is telling you something. Signals of success and failure tell you what your ends and values need to be, and not the other way around. Thus it is a mistake (and question-begging, to boot) to ask what antecedently fixed goals the serial hyperspecializer strategy serves for a particular serial hyperspecializer. The point of the signals is to permit one to revise one's goals and other standards, and the usually-offered candidates (adaptive fitness, survival, and so on) turn out to be just as much targets for revision as any other goals and standards are: you might learn from your own practical experience that your own reproductive fitness, or even your survival, does not particularly matter to you.<sup>13</sup>

Serial hyperspecializers will be equipped with cognitive devices enabling them to explore and establish standards, priorities and goals suitable for guiding behavior in a novel ecological niche. I will eventually be in a position

---

<sup>13</sup>Specificationism is a good antidote to the temptation to think that signals which are usable in dramatically different niches must be tied to some antecedently fixed goals. (For recent survey, see Millgram, forthcoming.) Specificationist reasoning takes very thin, very abstract ends and adds further content to them. Here is the dilemma for the advocate of fixed goals, and let's take health as a representative example of such a goal. To the extent that the goal travels from niche to niche, in the way that it must if it is to serve as an anchor for the sort of signals we have been discussing, that is because it is underspecified: it does not have enough content to serve as a substantive standard. (The point is easiest to make when what is attainable is less than perfect health. What counts as health, post-masectomy, depends on whether you need to do heavy lifting. If you are an office worker, and don't, one procedure will restore you to health; if you are a photographer, and do, you will need to have a different procedure.) But to the extent that the goal of health has content, it will not travel, and one's allegiance to it can be undercut by experience. (An extreme example: the French Decadents thought that experience had taught them that health was undesirable for them.)

to raise the question of why philosophers have not produced a task analysis complete enough to tell us what the full functionality of those devices must be. But we can nonetheless say that we come with what are evidently components of the package built in. That's reassuring, as far as whether the present account is on track: a reality check on the story I'm telling is the ability to identify its elements in our own lives. But if we humans are serial hyperspecializers, and these are among the components, then part of practical reasoning, for us, is the deliberation of ends.

### 3

Serial hyperspecializers move from niche to niche. So they will need a signal that tells them when to invest more resources in a current activity, and when to abandon the activity for another—that is, when to move niches. There can certainly be different indicators of success, but one of them will be competence in the activity. If a method of doing things is working, that's a sign that the strategy of occupying this particular niche is workable; if it's not, that's a sign that it's time to move on. In our own lives (here's that reality check once again), the negative signal is frustration, the feeling that you're banging your head against a wall, and that you just can't do it; the positive signal is one aspect or specialized sort of pleasure, an often elated sense that *you can do it*: that is, the cognitive content of one sort of pleasure is competence; a related variety involves the sense that, as one puts it colloquially, it's all coming together.<sup>14</sup> These complementary signals make up a great deal of what counts as overall hedonic tone. (The early utilitarians were right about this much: what they called 'pleasure' and 'pain' are enormously important in human life. Their mistake was to think that being important had to be: being a *goal*—instead of being a *signal* used to reformulate goals, as well as other standards, guidelines and priorities.)

A successful serial hyperspecializer will not just respond directly to signals like these (*keep doing this! stop that!*). It will extract patterns from the signals that it uses in making intelligent decisions about niche switching, and in tailoring its goals as it adjusts to a new or changing niche. From

---

<sup>14</sup>That is, when Vogler, 2002, pp. 80–89, reconstructs Aquinas's definition of pleasure, "the unimpeded *operatio* of a habit that is itself in harmony with one's nature," as the sense of a nexus of "agent, action, and circumstance"—the sense that it's all coming together—she is actually describing the content of the signal I have been characterizing. It follows (and here I am taking issue with Vogler) that, when the signaling system has been corrupted, such pleasures can be mistaken (or, as the traditional vocabulary has it, 'false').

the inside, this is generalizing from experience (philosophers say: “induction”), and because the content of the conclusion isn’t just factual, but is about what *to do*, and what *to go after*, it is “practical induction.”<sup>15</sup> This sounds fancy, but it’s something we’re all familiar with from our daily lives (which is, again, the reality check). Maybe you notice that whenever you’re photographing people, you’re awkward with the models and uncomfortable with your product, and, not coincidentally, that the work is making you miserable; but that when you’re animating clay figurines, the time seems to fly, you get better and better at it, and because you’re getting better, you’re enjoying what you do. You decide that what works for you, what you should be devoting your time to, is working with clay rather than photographing human subjects. In short, serial hyperspecializers will reason inductively about what their goals and other priorities should be. They will learn what matters, and what is important to them, by generalizing from experience. If necessity is the mother of invention, practical induction is its midwife.

Now, in the instrumentalist model of practical reasoning, if your ends don’t come with a built in feature that tells you which one is more important, when they conflict, you will perhaps be helplessly paralyzed, or perhaps choose randomly; in any case, you won’t be able to make a *rational* decision. Instrumentalists respond by conceiving of desires as having strengths; the alternative—desires that can’t be weighed or measured on a uniform scale, that is, *incommensurable* ends or desires—is thought of as a threat to practical rationality.<sup>16</sup>

But let’s think about incommensurability from the point of view of serial hyperspecialization. We expect a serial hyperspecializer to come equipped with signals that tell it when it’s time to switch niches (and on a smaller scale, when it should reallocate resources among activities pursued within a niche). These signals, we suggested, will tell it something with the approximate content: *This* is working, and *that* is not. To use the signals, the response ought to be, other things being equal, to stop doing *that* (or do less of it), and allocate more resources to doing *this*.

For such signals to be available, and such an allocation decision to be available, there have to be a *this* and a *that*. What this means is that serial hyperspecializers will sometimes be (the oxymoronic sound of the phrase notwithstanding) *parallel* hyperspecializers as well. (Only sometimes, however, because while it will sometimes pay for serial hyperspecializers to hedge their bets by pursuing different activities in parallel, dividing up your re-

---

<sup>15</sup>The terminology is from Millgram, 1997.

<sup>16</sup>For an overview of this debate, see Chang, 1997.

sources between different activities means that you have less to devote to any one of them.) Let's consider an instance where this sort of strategy is being pursued, and let's take a case where the activities in question are very different from one another: someone who divides her time between medical journalism, avant-garde space installations, and a country home in a small town, to any of which she could devote more or less time and energy.

Evidently, when activities are as different as this, they will involve structurally different standards and modes of assessment. As a medical journalist, for instance, there is a fairly clear hierarchy of clients, and you're doing better as your clients come from higher in the pecking order. As an artist, there's also a (perhaps less clear) hierarchy of professional acknowledgments; for instance, if you're invited to exhibit at the Kassel *Dokumenta*, you're doing pretty well. Around the country home, your real estate values could be going up or down, and you could be getting along better or worse with your neighbors. But the standards won't normally specify how to make tradeoffs between them. Is a space installation for the municipality a step up from a survey article for such-and-such a trade journal? Neither set of professional standards has a take on this. Is it more important to improve one's standing in the neighborhood by devoting extra time to gardening, or to work on lining up a performance venue in the Congo? Again, the methods of assessment proprietary to the respective activities won't say.<sup>17</sup>

So one reasonable strategy for a hyperspecializer will be to divide its energies between activities that don't share standards and methods of assessment, and let's abbreviate that to the fuzzier, problematic, but more familiar word, "values." It will be a normal side effect of pursuing the parallel hyperspecialization strategy that its values are incommensurable. The hedonic signals that guide reallocation of resources between niches do not require that niche-bound desires or goals or standards be comparable across

---

<sup>17</sup>And when they *do* say, the answers will often look oddly off-base, because what counts as success by one standard will be given no credit by another. If you're a businessperson, and your measure of success is return on investment, technological elegance is just a distraction; but if you're the researcher, financial success is the distraction.

Some economic decisions take currencies for granted; for instance, you choose this option because it gives you more dollars than the other option. But reforms in which one chooses a currency are *economic* decisions, too. And these decisions cannot generally be made on the basis of a currency taken for granted: when we replaced gold coins with dollar bills, the justification could not have been that we would have more gold coins, or more dollar bills, than before. Models of deliberation on which there is always a currency (of desires or preferences, for example) in the background must be supplemented by models of the techniques used to choose the currency; and, on pain of an infinite regress, the decision to choose one or another currency cannot always be made on the basis of costs and benefits paid out in a meta-currency.

niches. If you are a serial (and so, often a parallel) hyperspecializer, incommensurability in your values turns out not to be a mark of practical irrationality, or even an obstacle to full practical rationality, but rather, the way your evaluative world will look to you, when you are doing your practical deliberation normally and successfully. Evaluative incommensurability is a threat to the sort of rationality suitable for Piltown Man. Serial hyperspecializers gravitate towards—and come *equipped* for—incommensurability.

#### 4

Having said that we are serial hyperspecializers, I now need to take some of it back, and while I am at it, I will discuss some of the metaethical issues which my claims raise.

I just described serial hyperspecialization as a solution to a design problem, and I claimed that we engage in it. But we're by no means optimized for it. For one thing, the hardware platform that the self-modifying software runs on is only just slightly changed from its disposable predecessors: we *die*. We do live longer than many animal species, but the hardware wears out, and still exhibits what in the artifactual world would be perceived as planned obsolescence. People can change careers midlife, but as they get older, they get less flexible, and learning new skills, new languages, and new attitudes gets harder. Even when they can do it, they often stop *liking* it. Some techniques and their associated priorities and guidelines seem to be only acquired in youth: mathematics, for example. And there's a lot of behavior (and associated desires, goals, and ends) that is relatively inflexible. As the theory of natural selection (along with some obvious extra premises) would predict, it often has to do with reproduction, and with food, and with other basic maintenance, like sleeping. Some things are just too important to be left up to the self-modifying software.<sup>18</sup>

That raises a practical question for us: *should* we be serial hyperspecializers? The question has a precedent early on in the history of philosophy. In Aristotle's *Nicomachean Ethics*, he first described the person who excels at a life of public activity. (Think of a wealthy community leader, someone who engages in politics but isn't really a professional politician.) Then he went on to describe another life of which humans are also capable, the life

---

<sup>18</sup>That said, there's surprising flexibility even about these. The variations in what people will eat, and how they will go about eating it, and the different ways humans go about reproducing, highlight just how much latitude the self-programming software really has.

of contemplation. (Think here of a mathematician, someone who spends his time, all of his time, in an office in an ivory tower thinking about hard-to-imagine mathematical objects.) We *can* aim for either sort of life, so the question is, which should it be? Aristotle thought that contemplation was a higher and more god-like activity, and so he recommended that. The analogy notwithstanding, I don't propose trying to answer our question in those terms. But we have a very similar question before us. We could work at being better serial hyperspecializers, or alternatively, we could forego that option, and try to become better Piltown Men. How ought we to think about the choice?

That it *is* a choice may assuage a worry on the part of alert readers, that I am committing what used to be called the naturalistic fallacy. That worry is misplaced. First, the form of the argument is emphatically not: We are serial hyperspecializers; serial hyperspecializers do such and such; therefore we should do such and such. Rather, it is a bit like pointing out to someone fixing his bicycle that he is consulting a toaster repair manual, and that he might do much better with an instruction sheet for the bicycle. (If he decides to use his broken bicycle as a doorstop, he is not, as far as the argument goes, doing anything unreasonable.) Observing that your theory of rationality is tailored to a different sort of creature, and that there is a theory of rationality more closely tailored to the sort of creature you are, may amount to advice—advice that you should move from the first theory to the second—but there may also be all sorts of quite reasonable ways to slough off that advice. Second, the thought that a naturalistic fallacy must be somewhere in the neighborhood is tied to the presumption that the account I am giving is a 'naturalistic' or 'naturalizing' account. Now, naturalism, these days, is a little like the flag, mom, and apple pie, circa the McCarthy period; far too many intimidated philosophers have signed loyalty oaths to it. But it is not a new observation that no one can actually say what they mean by it.<sup>19</sup> Once you delete commitments that no sane philosopher ought to take on (such as, that our current sciences are, as they stand, right about what 'nature' contains) nothing is left of naturalism but

---

<sup>19</sup>[Hempel's dilemma, cite] That discussion centered on the term 'physicalism'; the dialectical moves for naturalism are practically identical. The rough map is: physicalism and naturalism are, respectively, the views that there's nothing in the world that isn't the sort of thing our physics or our science more generally tells us there is. If we interpret that to mean the sort of thing our *current* theories tell us there is, we already know it's wrong; the history of science tells us that theories change quite dramatically from time to time. If we interpret that to mean the sort of thing our future theories will tell us there is, we don't know what we're committing to.

a fuzzy feeling of confidence in the sciences and their future. In other words, analytic philosophers who take themselves to be naturalists (or, for that matter, physicalists) have lapsed into a doctrine for which emotivism is the correct account, and I take that point in just the way their logical positivist ancestors would have: ‘naturalism’ and ‘naturalistic’ *don’t mean anything at all*. It follows that my own account is not naturalistic (and that it’s not non-naturalistic, either). And so worries about naturalistic fallacies should not be quite as pressing.

There is, however, a related complaint that does belong on the plate, namely, that observing that we implement (or approximately implement) a design solution does not yet show that what we do is *rational*. (For instance, people systematically overestimate their own looks, intelligence, driving abilities and much else; there are many benefits to doing so, not least that accurate assessment seems to be correlated with depression; but the benefits do not make the inflated self-images rational.) The argument that what I have been describing is indeed practical *rationality* is not best constructed from the point of view of an observer or as-if designer; in my view, it has to be assembled from the point of view of a deliberator of the requisite sort (from the inside out, as it were), and I defer that part of the story to another occasion.<sup>20</sup>

The choice of whether to work on being a better serial hyperspecializer, or a better Piltdown Man, is in one respect less open than it might seem, because we can only entertain it if we are pretty good approximations to being serial hyperspecializers already. (Full-fledged Piltdown Men don’t ask themselves this sort of question; someone who does is considering becoming a Piltdown emulator, rather than a full-fledged Piltdown Man.) Since we are entertaining it, we are equipped to engage in practical induction, that is, to learn, from experience, what does well for us, and what does badly. And so we should ask: which strategy has worked well in the past? Do the people who adapt to new circumstances do better, or the people who don’t? Do the people who try out new things, and explore new ways of living, do better, or the people who don’t? Do the cultures and societies that change and adapt do better, or the cultures that keep the old ways? (Once again, these questions are not to be answered in terms of prior, given goals, but by looking to the signals of success and failure that are part of the cognitive equipment of serial hyperspecializers.)

---

<sup>20</sup>For the point that rationality on the part of a designer does not mean rationality for the designed agent, see Millgram, 1991, sec. 5; for an overview of depressive realism and related topics, see ?; for the argument-from-the-inside that thinking of the sort I have been describing is rational, see Millgram, 1997.



At first glance, the answer might seem pretty unequivocal. Examples at the scale of societies—in which failure to adopt the methods and techniques of economic, political and military competitors, and the systems of standards embedded in them, gives rise to one or another form of catastrophe—tend to divide up into the politically charged and the unfamiliar, and so I will leave these to the reader as an exercise. But for a recent example at the scale of industries, recall when slide rules, not so long ago, vanished off the face of the earth. The people who had worked in slide rule manufacturing divided up into two categories: those who learned how to do something else, and those who had dead-end lives.

At second glance, the tradeoffs are more complicated. First, one possibility is that sometimes the serial hyperspecializing strategy does better, and sometimes the Piltdown Man strategy does better. If people exhibit the relevant sort of second-order plasticity, we might expect them to sometimes live one way, and sometimes the other. Which alternative makes most sense will depend largely on one's environment, and especially on how elaborately articulated one's human ecosystem is; if there isn't already a lot in the way of what I've been presenting as the cultural surrogate for biodiversity, there aren't many niches to occupy, and so there isn't a lot of advantage in being prepared to occupy novel niches. Perhaps very poor and war-torn societies today don't reward serial hyperspecializers, and it's likely that in our distant past, human societies were also not complex enough to do so. (For that reason, it is unlikely that we are serial hyperspecializers because ancestors who adopted this strategy were thereby winners of the Darwinian game: the capabilities we are discussing look to be serendipitous, from the point of view of natural selection, rather than an evolutionary adaptation.<sup>21</sup>) Which alternative pays may also depend on the mix of strategies already in play in one's society. For instance, perhaps being Piltdownish is a good strategy in an economy that already contains enough serial hyperspecializers to keep it flexible and afloat. (If the dot-commers find new jobs in other industries, you can keep doggedly working at your retail job until you die.)

Second, part of the design solution that human beings implement is the lengthy childhood that permits them to learn not just a system of standards for the environment into which they are born, but methods of learning the systems of standards appropriate for subsequent environments. That re-

---

<sup>21</sup>There is a further reason for not insisting that being a serial hyperspecializer is an adaptation: you couldn't, in principle, give a legitimate argument for the claim. The decently concrete just-so story that such an argument would be built around would require a definite phenotype to work with, and to be a serial hyperspecializer is to have, precisely, an *indefinite* phenotype.

quires long-term parenting; parenting is often tedious, frustrating, and variously unpleasant, that is, it requires overriding the signals that tell us to stop what we are doing and do something else.

Those caveats notwithstanding, as far as the big picture goes, I think the first-pass practical-inductive take on whether to invest one's resources in serial hyperspecialization, and whether to endorse the canons of practical-inductive reasoning, has it right. Piltdown Man is boring, and serial hyperspecializers are interesting. Social institutions and lives tailored to Piltdown Man are frustrating; creativity, novelty, and originality, intellectual and otherwise, *feel* much better. And to be a Piltdown Man is for changes in one's environment to be nonrecoverable catastrophes. It's a no-brainer.

## 5

How could so many philosophers have misidentified the very species for which they were philosophizing? The misidentification explains why the task analysis of adaptation to novel niches is in the most preliminary stages. But the mistake is even worse than it sounds. In producing a theory of practical reasoning that fits Piltdown Man, our human philosophers were after all philosophizing. But philosophy consists largely in the exploration of new intellectual niches, of designing standards for those niches, and setting intellectual goals to pursue in them. Therefore, Piltdown Man does not philosophize.

My own guess is that there have been several contributing factors, and I will mention just two. First, philosophers of the past, even when they were not (as Popper argued that Plato was) apologists for traditionalist societies, tended to see much more in the way of stable social roles than we do. Second, and more interestingly, we tend to forget that instrumental rationality was a hard won achievement. To reconceive human beings as instrumentally rational was a revolutionary and liberating step forward: against sheer rote and custom, it opened up possibilities for criticism, and paths to rapid social reform, that we have since come to take for granted. The Benthamite utilitarians exemplify this frame of mind, but recall as well the enthusiasm for 'rationalisation' (the British spelling is intentional) that reshaped England during the postwar period. Forcing social institutions to make instrumental sense was a heady and exhilarating enterprise, and an admirable one in its day. A commitment to treating instrumental rationality as the be-all and end-all of practical thinking was, not that long ago, a passionate political commitment.

Perhaps we are not yet finished with the possibilities opened up by instrumental rationality: lip service to the contrary, there is still much less of it around than you would think. Nonetheless, Piltdown Man was a primate species whose best-known property turned out to be *not existing*.<sup>22</sup> Instrumentalism *can't* be right. The notion that all reasoning about what to do is means-end reasoning, and that you can't reason about what ends to have—that's a theory of rationality suitable, not for humans, but for an imaginary species of dwarves, a stunted primate that might have evolved but didn't, and which, if it had evolved, would have had to occupy a very different role in the larger ecology than we do.

## 6

Work on practical reasoning has, in recent decades, been closely connected with theorizing about agency and personal identity, and with substantive moral philosophy. Having misidentified the species for and about which they were philosophizing, and so having gone wrong about what our practical rationality must be like, philosophers have gone on to make equally deep mistakes in these other areas as well.

The mainstream view in moral psychology is that agents are unified, both in fact (for the most part) and ideally (disunity of agency is regarded as a defect). Accounts of agential unity vary, and include: agents produce actions which can be in a very robust sense attributed to them, actions they *own* because the agents are identified with, rather than alienated from, their choices; agents have 'practical identities' or 'ground projects' which they may lose, but cannot disown; agents do not pursue projects at cross-purposes with one another; having made a decision, they follow through on it (and do not instead act on some contrary impulse); they possess a unified point of view from which they render judgments about what is worth doing and what they will do; they reflect on their actions, and endorse their choices when they do; when they act, they act so as to understand what they are doing and why; their choices are governed by policies which dictate how competing reasons will be taken into account.<sup>23</sup> Across these variants, we find a rough but shared picture, of a creature that has integrated its goals, evaluative judgments and other guidelines into a single and internally

---

<sup>22</sup>For the latecomers, Piltdown Man was a hoax involving a human skull, the jaw of a modern ape, and sandpaper; Millar, 1972, is a popular account.

<sup>23</sup>See, for instance, Williams, 1981, Bratman, 2001, Bratman, 2003, Velleman, 2000, Korsgaard, 1999, Korsgaard, 1996a, ch. 13, Korsgaard, 1996b Frankfurt, 1988.

consistent pattern, and whose control structures generate actions that are consistent with the pattern.

That’s a pretty good picture of Piltdown Man, provided its software download is up to scratch. But it’s not at all a satisfactory rendering of a serial hyperspecializer. Because serial hyperspecializers are often parallel hyperspecializers, they will not infrequently pursue activities that go on in different niches, with incommensurable systems of standards. Consequently, their various activities will not be governed by a unified hierarchy of evaluative judgments and goals. They may even be at cross-purposes, and although a serial hyperspecializer is to some degree equipped to manage resource competition between projects and activities, we should not expect these to be regulated by higher level goals, desires or standards on the part of the agent. (That would be to commit the error of thinking that an engineering solution which works within niches can be scaled up to solve problems that span niches, as though the world were simply a larger niche, and could be navigated by constructing a system of standards suitable to a much more constrained environment, only *bigger*.)

Because serial hyperspecializers are equipped with niche-switching prompts like frustration and boredom, they can surprise themselves by taking steps that contravene their current projects: steps they do not endorse, steps that do not ensue on their policies for weighing reasons, steps that lead them to say, “I didn’t really do it. . . it just kind of happened.” When they act, they may not know what they are doing, or why. (“I don’t know why I did that: I just *found* myself doing it. I really surprised myself.”) And when they behave in this way, they may well be evincing, not their irrationality, as the theoretical mainstream has it, but what is precisely practical rationality for serial hyperspecializers.

Mainstream theory of agency is very good at articulating the phenomenology of acting on the basis of an override—a signal that is independent of your system of standards, of your complex of goals and desires (in one famous bit of terminology, of your ‘subjective motivational set’), of your practical identity, and so on. (The accompanying utterances at which I just gestured are typical, and it can feel as though it wasn’t *you* who did it. Although of course you can often *explain* why you did it; for instance, you were bored out of your mind.) Because mainstream theorists identify the disposable personae you happen to be projecting at the moment with your *self*, they treat your responding to such signals as the unraveling of your agency—as a disaster. But if you are a serial hyperspecializer, it is not necessarily a disaster at all: this is just how you cast off devices that are suitable, temporarily, for coping with one environment, in favor of other devices that are

suitable for other environments. If you like, and granting for the sake of the argument that the mainstream is giving a successful analysis of what it is to be an agent, you are not a single agent, but, quite possibly, owner of a series of them: agents are interfaces you conjure up to meet the needs of the moment; do not make the mistake of thinking that one or another of them is *you*.

Serial hyperspecializers should not be expected to be completely fragmented; we will often see a great deal of top-down, policy-governed coordination of activity within a niche (at any rate, within an already explored niche). But we should expect to see substantial agential disunity, both synchronic and diachronic, induced by fault lines between niches, and by the cognitive devices that facilitate niche exploration and niche jumping. Recent theoretical work on agency has been rich, subtle, and interestingly argued, but also, unfortunately, philosophizing suited to a species not our own.

## 7

In the past several decades, a great deal of attention has been given to theories of personal identity, and the mainstream view among analytic philosophers today is neo-Lockean: what makes you the same person you used to be is psychological continuity, typically glossed as remembering your past, acting on your former intentions, having a similar character, and so on.

However, it should be obvious that identity concepts need to be tailored to species.<sup>24</sup> A butterfly is not psychologically continuous with the caterpillar it was. It does not—except for the *matter*, which has melted down and completely reformed itself inside the cocoon—even have the same body. But it is still the same animal, because metamorphoses are what butterflies *do*: that's their *ergon*, or, in the more modern locution, metamorphoses are part of the design solution they implement.

*We* metamorphose psychologically: that's what it *is* to be a serial hyperspecializer. In moving from niche to niche, you can quite correctly throw out your old goals, standards, preferences, intentions, and policies wholesale. (You can also, quite correctly, throw out almost all of your memories, but defending that claim will require a discussion of what a serial hyperspecializer's reasoning about matters of fact should look like, and I will leave that to a sequel.) When this happens, you are still someone in whom your former self has a prudential stake (this being what really matters, in this philosoph-

---

<sup>24</sup>This is a Wiggins-like observation, but I wouldn't want to buy into all of the details of his treatment. For the latest revision, see Wiggins, 2001.

ical subject area, about being the same person), provided the transition to your new psychology was managed using the serial hyperspecializer's characteristic mode of rationality. When boredom and frustration prompt you to move on, to forget your former life as thoroughly as possible, and to take up entirely new activities that you can find interesting, and in which you feel yourself competent and at home, that is a benefit to *you*.<sup>25</sup>

On the other hand, neo-Lockean accounts of personal identity are just about right for Piltdown Man. Piltdown Man starts life with a program that will govern his activities until the end of his days. The psychological states that the program deposits and uses—records in memory of the creature's progress and the state of its environment, overarching goals, subgoals and the like—may (and ought to be) updated constantly; but they will never be need to be deleted wholesale. And if they are deleted wholesale, an individual Piltdown Man will not normally survive the operation; so treating a Piltdown Man's life as coextensive with a continuous psychology of this kind is entirely reasonable.

A less popular but still respectable position on personal identity prefers bodily continuity as the criterion of sameness—you are the very same person you once were if you have the same body—and it might seem that my account is committed to this alternative. I am not at all certain, for two reasons. First of all, I take seriously Bernard Williams's arguments to the effect that we do not have a philosophically satisfactory account of the body, and that we have not thought through what our distinction between body and mind comes to.<sup>26</sup> 'Body' (and, more recently, 'organism'), in these discussions, is just a placeholder, a we-know-not-what. Second, we do not have an explanation for bodily continuity being the basis of an identity concept suitable for serial hyperspecializers. It is not—recall the caterpillars—the only alternative to psychological continuity.

While I do not have on hand an account of personal identity suitable to serial hyperspecializers, I expect that it will turn on the theory of practical rationality: that you will be the same person you now are if (very roughly) you get from here to there via patterns of reasoning, that is, via thinking that *counts* as reasoning for serial hyperspecializers.

We can now acknowledge an objection that has been lingering in the background for a while now. The other side of a philosophical theory of

---

<sup>25</sup>Instrumentalist theories of practical rationality are unable to account for this observation; see Williams, 1973, ch. 6.

<sup>26</sup>Williams, 1973, chs. 1–5, and esp. pp. 11f, 68ff. Oddly enough, Williams himself went on to endorse a bodily-continuity account of identity, which strikes me as an uncharacteristic failure of philosophical nerve.

personal identity is a philosophical theory of death; after all, you are dead once there is no one who is identical to you.<sup>27</sup> I have been laying out an account of what it is to be a human being, one on which the fact that you are going to die is not an essential part of the design, but rather on the order of a manufacturing flaw. And that might suggest that something is deeply amiss in the account. It is a widely held view that your death is the frame in which the elements of your life are meaningful, and that to lose sight of your death is to live ‘inauthentically’. A suspicious reader might even wonder whether the present account is not just a way of avoiding the confrontation with one’s own mortality.

Certainly there is no point in pretending that one is not going to die. But in the form of life at which I have been gesturing, the frames that make activity meaningful may well be much shorter than an entire life. Facing up to one’s death ought not to involve thinking about it in terms that would be suitable, not for us, but for Piltdown Man.

## 8

Theories of practical reasoning are closely tied to corresponding ethical or moral theories.<sup>28</sup> I don’t here want to try to delineate a moral or ethical theory that would be suitable for a community of serial hyperspecializers. But I do want to indicate, albeit tersely, why the moral theories we have won’t do. Each of the theories on the standard menu is naturally taken as according primary importance to one or another possible moral priority; without giving a great deal of argument, I am going to suggest that it would be unreasonable for serial hyperspecializers to accord these priorities the primacy they are given by the standard theories.

Utilitarianism takes the satisfaction of preferences or desires, or alternatively, hedonic tone (pleasure and the absence of pain), to be the only thing that matters intrinsically. But in the lives of serial hyperspecializers, preferences, desires and pleasure are cognitive tools; they are no more the source of all value than any instrument can be. In migrating birds, desires and ends are hardwired, part of a design solution that gets the organism to the next stop on its route. Since the bird will normally be doing well only when it does get to the next stop, desire satisfaction is a good proxy for success. Piltdown Man is only one step away from migrating birds: his

---

<sup>27</sup>For one version of this observation, see Parfit, 1987, pp. 281f; and I have heard it phrased this way by Mark Johnston.

<sup>28</sup>This claim is spelled out in Millgram, 2005, Introduction and ch. 11.

strategy for coping with his environment is preprogrammed, and once he steps off the straight and narrow path that it defines (or presupposes), he becomes dysfunctional. So for Piltdown Man, also, treating desire or preference satisfaction as success, or a proxy for success, is a reasonable ethical strategy. But because serial hyperspecializers compute desires, preferences and the like on the fly, what matters is what the desires and preferences (and signaling channels) are there to facilitate: thus, it is always an open question whether the desires or goals one has are *correct*, and always an open question whether the hedonic signals have been corrupted. Accordingly, for serial hyperspecializers, none of these are plausible candidates for being the ultimate repository of value, or even a usable proxy for it.

Kantian moral theory is famously built on three large ideas, and I'll consider just one of them now. The point of having serial hyperspecializers is that they can exploit narrow niches—niches that don't have room for many occupants. That means that, for serial hyperspecializers, the Kantian question, What would happen if everybody did that? (that is, the first formulation of the Categorical Imperative) is simply beside the point. Only the most primitive ecologies contain just one role; hyperspecialization is a strategy suitable for highly articulated ecologies. But because Piltdown Men, living on the plain, are all peasants (and more generally, because they all pick their life strategies from a short menu of fixed options, one which it's reasonable to think of as a single disjunctive strategy), the question, What would happen if everybody did that? *is* a reasonable test to impose on *them*.

Aristotelian moral theory focusses on the shape of a well-lived life (which Aristotle called *eudaemonia*, or happiness). Aristotle meant to read that shape off of an account of the human 'function,' that is, off of the functional design of the human species. So in that respect his method has a good deal in common with the present approach. But if the account I have been sketching here is correct, Aristotle was mistaken about what design solution human beings (approximately) implement, and so his substantive ethical advice can be at most accidentally right. And because, on the account I have been sketching, much of the shape of a human life, at the level of specificity which Aristotle hoped to capture, is computed on the fly, and properly varies from individual to individual, there is much less in the way of generic but substantive ethical theory to be had than Aristotle thought.

There is a further difficulty with eudaemonism. Eudaemonism assumes that there is a coherent pattern of activities that will constitute a human life's going well. Aristotle made that assumption because he really did treat ethics as continuous with biology: when you find out how squirrels live their lives (during the summer, they bury nuts, and during the winter, they



retrieve and eat those nuts), you know what it is for a squirrel's life to go well (they succeed in finding the nuts they put down); likewise, when you find out how humans live their lives, you know what it is for a human's life to go well. That approach makes sense for squirrels (and for Piltdown Man), because squirrels (and, let's suppose, Piltdown Man's programs) have been around long enough for the rough edges in their form of life to get buffed down; their lives, when they go well, do exhibit coherent patterns of activity. However, if I am right, the forms that human lives take are extremely transitory; it would be unreasonable to suppose that there has been enough time to debug them. And that expectation is confirmed by observation: for the most part, human activities exhibit—once you press on them even a little—deep incoherences. If humans are serial hyperspecializers, and if happiness is taken (as philosophers have traditionally taken it) to involve satisfying coherence conditions on one's aims, activities, and so on, then happiness is not a guiding concept appropriate for human beings (although it may well be a suitable guiding concept for Piltdown Man).

The classical moral and ethical theories are impressive intellectual accomplishments, but they are misdirected. Moral philosophers have almost without exception been in the business of designing moral theories for a different species—and they have been quite successful at it. Utilitarianism, Kantian moral theory, and Aristotelian ethics would be highly suitable guides for the communal life of one form or another of Piltdown Man. But they are not at all appropriate guides for us, if we are serial hyperspecializers—or even if we are reasonable approximations to them.

## References

- Arthur, W., 1987. *The Niche in Competition and Evolution*. John Wiley and Sons, New York.
- Boyd, R. and Richerson, P., 1985. *Culture and the Evolutionary Process*. Chicago University Press, Chicago.
- Boyd, R. and Richerson, P., 2004. *Not By Genes Alone: How Culture Transformed Human Evolution*. Chicago University Press, Chicago.
- Bratman, M., 2001. Two problems about human agency. *Proceedings of the Aristotelian Society*, 101(3), 309–326.
- Bratman, M., 2003. Autonomy and hierarchy. *Social Philosophy and Policy*, 20(2), 156–176.

- Chang, R., 1997. *Incommensurability, Incomparability, and Practical Reason*. Harvard University Press, Cambridge, Mass.
- Debats, M., Cluzard, J., and Perrin, J. 2001. Winged migration. Produced by Christoph Barratier and Jacques Perrin. Columbia Tristar Hom.
- Frankfurt, H., 1988. *The Importance of What We Care About*. Cambridge University Press, Cambridge.
- Jacquet, L. 2005. March of the penguins. Produced by Ilann Girard. Warner Independent Pictures.
- Korsgaard, C., 1996a. *Creating the Kingdom of Ends*. Cambridge University Press, New York.
- Korsgaard, C., 1996b. *The Sources of Normativity*. Cambridge University Press, Cambridge.
- Korsgaard, C., 1999. Self-constitution in the ethics of Plato and Kant. *Journal of Ethics*, 3, 1–29.
- Kubrick, S. 1968. 2001: A space odyssey. Produced by Stanley Kubrick and Victor Lyndon. Metro-Goldwyn-Mayer.
- Laloux, R. 1973. Le planète sauvage. Produced by Roger Corman. Based on novels by Stefan Wul. Anchor Bay Entertainment.
- Mayr, E., 1984. Species concepts and their application. In Sober, E., editor, *Conceptual Issues in Evolutionary Biology*, pages 531–540, MIT Press, Cambridge, Mass.
- Millar, R., 1972. *The Piltdown Men*. Victor Gollancz, London.
- Millgram, E., 1991. Harman's hardness arguments. *Pacific Philosophical Quarterly*, 72(3), 181–202.
- Millgram, E., 1997. *Practical Induction*. Harvard University Press, Cambridge, Mass.
- Millgram, E., 2004. On being bored out of your mind. *Proceedings of the Aristotelian Society*, 104(2), 163–184.
- Millgram, E., 2005. *Ethics Done Right: Practical Reasoning as a Foundation for Moral Theory*. Cambridge University Press, Cambridge.

- Millgram, E., forthcoming. Specificationism. In Adler, J. and Rips, L., editors, *Reasoning: Studies of Human Inference and its Foundations*, Cambridge University Press, Cambridge.
- Odling-Smee, F. J., Laland, K., and Feldman, M., 2003. *Niche Construction*. Princeton University Press, Princeton.
- Parfit, D., 1987. *Reasons and Persons*. Oxford University Press, Oxford.
- Pitcher, G., 1970. Pain perception. *Philosophical Review*, 79(3), 368–393.
- Roush, S., 2003. Copernicus, Kant, and the anthropic cosmological principles. *Studies in History and Philosophy of Modern Physics*, 34, 5–35.
- Sterelny, K., 2003. *Thought in a Hostile World: The Evolution of Human Cognition*. Blackwell, Oxford.
- Tenn, W., 2001. *Of Men and Monsters*. Orion, London.
- Velleman, J. D., 2000. *The Possibility of Practical Reason*. Oxford University Press, Oxford.
- Vogler, C., 2002. *Reasonably Vicious*. Harvard University Press, Cambridge, Mass.
- Wiggins, D., 2001. *Sameness and Substance Renewed*. Cambridge University Press, Cambridge.
- Williams, B., 1973. *Problems of the Self*. Cambridge University Press, Cambridge.
- Williams, B., 1981. Moral luck. In *Moral Luck*, Cambridge University Press, Cambridge.