

A Life Hack with a Future

Annotated Table of Contents

1. *Introductory Overview*

What is a self, and what is it to be a person? We can make sense of our selves as devices, overlaid on the human animals that we are, that solve a cluster of related problems, having to do with what it takes to draw conclusions, and to compensate for our propensity to overlook the reasons to balk at drawing them. That is, being configured as persons enables us to infer, to reason, and more broadly, to figure things out. But that remarkable life hack is threatened by the division of intellectual and evaluative labor; the flip side of specialization is ignorance of everything it's not your business to know, and the rapidly-growing pool of what there is to overlook is well on the way to swamping the capacities of the device—that is, of what makes us ourselves.

2. *Private Persons and Minimal Selves*

Only you have your own thoughts: that's *logical* privacy. If you don't tell anyone, only you have to know what they are: that's *personal* privacy. Both logical and personal privacy are side effects of an administrative technique: the selection of consistency constraints—such as the principle of noncontradiction—to drive behavior in hard-to-predict environments. First-person authority, that is, your privileged awareness of what you think, turns out to be an implementation requirement on such a system. Those constraints, on the one hand, and the boundaries of the person, on the other, can be adjusted in tandem; to make the payoffs visible, we can consider a form of extended mind on which it is possible to impose additional laws of logic. The issue of whether the laws of logic can be renegotiated in this way leads to a reinterpretation of Moore's Open Question Argument.

(Previously published as “Private Persons and Minimal Persons,” *Journal of Social Philosophy* 45(3), Fall 2014: 323–347.)

3. *Outsourcing the Mind: Extended Cognition and First-Person Authority*

The problem of first-person authority is that of explaining how it is that, normally, you know what you believe, what you feel, what you're

in the middle of doing, what you want and so on—and know in a way that’s different than other people know what you believe, etc. Extended cognition is introduced by the observation that if a mind is a computational system, there’s no principled reason for the borders of your brain to be the borders of your mind. Treating these two philosophical discussions together brings into focus an engineering problem, that of the design and implementation of reliable low-overhead monitoring for inferential processes. First-person authority constrains the reach or extent of the reasoning mind, and thus what it makes sense to treat as the limits of extended cognition.

4. *Excursus: Philosophy and Hypophilosophy*

Philosophy is in the first instance the machine tool industry of the intellect: its task is making the intellectual tools that make the intellectual tools. If you think about how you’d equip creatures you were building to get along in hard-to-anticipate environments, you realize you’d want them to design and implement—and then theorize about and reassess—their own consistency regimes. We are those sorts of creatures, and the way we go in for logic, philosophy of logic, and metaphysics makes sense from a design standpoint. However, as philosophers lose track of the point of their activities, their moral immune systems are disabled, and they lapse into *hypophilosophy*: theorizing, but about nothing at all.

(A revised version of “Hypophilosophy,” *Social Philosophy and Policy* 35(2), Winter 2018: 138–157.)

5. *TBA*

Almost all of our reasoning is *defeasible*: that is, our inferences go through only other things equal. And there is always more to take account of; the list of things that might go wrong is open-ended, and doesn’t run out. Most work on defeasibility is focused on how to represent it, but here we’ll take a step back, and ask *why it’s there*. I will argue that defeasible inference is a hard-to-avoid design feature of certain kinds of boundedly rational agents, that the open-endedness is genuine, and that we need to understand defeasibility from an engineering—rather than a formal—perspective. Long story short, we reason defeasibly because we are the sort of creatures for whom the issues with which they have to cope are, very, very often, TBA.

Interlude: Mimesis and the Philosophy of Filing

The analysis of selfhood being advanced construes it as an administrative device with information management benefits, and we need to address the almost inevitable resistance that the approach is bound to prompt: that your *self* is what you *are*, and the metaphysics of the self can't amount to anything as mundane as bureaucratic procedure. However, there's a parallel to be drawn between the progress of serious literature, as recounted by Erich Auerbach, and the progress of philosophy: rather as our literary tradition over its history became able to give thoughtful and honest treatment to ever lower subject matter, philosophy improves itself by its willingness to take on ever more quotidian explanations.

6. ...

When division of labor goes sufficiently deep, it entails division of intellectual and evaluative labor. The flip side of division of intellectual labor is ignorance, of almost all those things it's not your job to know; the flip side of division of evaluative labor is not being competent to assess most of what lies outside your area of expertise. So the bounded-rationality strategy of division of labor ends up including defeasible inference and vagueness, as ways of preemptively making room for all those things you don't know and can't assess. When defeasibility and vagueness are introduced as this sort of strategic choice, to treat a class of inferences as deductive becomes a strategic choice also; this raises the question of how such a choice could be motivated.

7. *Practical Nihilism*

Nihilism about practical reasoning is the thesis that there's no such thing as practical rationality—as rationally figuring out what to do. Other philosophers have defended a theoretically oriented version of the thesis, usually called “error theory.” Here I take a shot at making the best case I can for a fully practical version of it: that we are *so bad* at figuring out what to do that we don't really know what doing it right would so much as look like. Much of our control of instrumental (or means-end) rationality is illusory, and we are almost entirely incompetent at managing the defeating conditions of our practical inferences—that is, of knowing when not to draw an apparently acceptable conclusion. There's a case to be made that instead of trying to reason more successfully, we should be trying to make failure pay.

(Previously published in *Philosophies* 8(2), 2023.)

8. *Custom and Customization*

Instrumental rationality has an aesthetic dimension; competence in practical deliberation presupposes command of a largely overlooked range of aesthetic sensibilities. Only in a human or artifactual environment in which others opt for normalcy in their instrumental deliberation are the defeasibility conditions of one's practical inferences intellectually tractable. It seems to follow that the regimentation of others' aesthetic sensibilities is a precondition of one's own practical rationality. However, if thinking for yourself, that is, your own autonomy, presupposes the heteronomy of others around you, how can we retain the Enlightenment aspiration of autonomy for all?

9. *Ifs, Ands, and Buts—and Whether You Can Average Them Away*

In a defeasible (or nonmonotonic) inference, an otherwise warranted conclusion can be defeated by an indefinite number of further considerations. Especially when potential defeaters for a line of reasoning are drawn from many different specializations, assessing whether an inference goes through can seem like an impossible task; if only specialists understand their own disciplines, no one could be competent to perform those assessments. Here we consider an end run around the problem: whether the hierarchy-of-hypotheses approach, originally developed for evaluating the state of play in invasion biology, allows for defeasibility management that does not depend on understanding defeaters drawn from different specializations.

10. *Selecting for Defeasibility*

Almost all of our reasoning is defeasible—that is, you're ready to take back the conclusion when it's brought to your attention that there's a difficulty you've overlooked. But we engage not only in division of labor, but in division of intellectual and even evaluative labor; when it's your job to know *these* things, but not *those*, there's a lot for you to overlook (namely, all those things it's not your job to know). How can our reasoning, in our highly specialized society, be good enough to get by on? Part of what picks up the slack is selection effects operating on our institutions and practices: when a practice or a technology doesn't make it easy enough to anticipate what might go wrong when you interact with it, it is likely to be avoided and drop out of the social mix.

11. *Why Are Persons Ends in Themselves?*

The so-called Austrian economists provide us the wherewithal to make sense of Kantian personhood: persons, understood as ends in themselves, solve a social coordination problem, in roughly the way that market actors solve resource allocation problems. People are ends in themselves—they must be treated not as mere means, but as autonomous agents whose choices constrain one’s own decisions—in *that* our existence as persons is a means to the very large end of solving an extremely complicated collective choice problem. You have unconditional value *only conditionally*: on a particular distributed construction procedure being the means of handling the coordination of activities—the means that we have in fact collectively adopted.

12. *Moral Education, Moral Standing, and the Persistence of Moral Skepticism*

“I see what morality requires me to do, but I don’t know why I should do it.” Moral skepticism is as old as philosophy, and over the past 2000 or so years, philosophers have sought both the moral education able to forestall the skepticism, and arguments showing the skepticism to be, simply, a mistake. Robert Nozick’s account of moral standing allows us to assemble an explanation for skepticism’s tenacity, one which has the consequence that, no matter how good it is, theory will not do away with moral skeptics. The account speaks to a question that the other chapters of this book leave open, namely, *who* the selves and persons are.

(Previously published as “The Persistence of Moral Skepticism and the Limits of Moral Education,” in H. Siegel, *The Oxford Handbook of Philosophy of Education* (New York: Oxford University Press, 2009): 245–259.)

13. *A Life Hack with a Future*

Personal identity—what it takes to be the same person over time—is best understood from what Dennett called the design stance. Rather than pursue traditional conceptual analysis—a list of necessary and sufficient conditions for being who you now are, later on—a creature-construction argument turning on bounded-rationality constraints recovers central features of personal identity. Our focus is on the time it takes to complete investigations, the need to monitor other-things-equal inferences for correctness, and the imperative, within the context

of protracted and openended investigations, for generic resources and strategic positioning.

14. *Conclusion*

Looking out for Our Selves

Selfhood is an administrative device; persons are built on top of selves, and presuppose the functionality of the underlying apparatus. But the techniques that selves are there to make possible are being rendered ineffective by the advance of specialization. Too much is at stake in selves and personhood simply to surrender them, and in the course of the book we have considered several ways of bolstering their performance. Even so, largish revisions to how selves work are evidently in the offing. How should we make decisions about those sorts of changes—changes to what constitutes our very selves?