

Practical Rationality and Division of Labor

An Interview with Elijah Millgram

*Professor Millgram, you're an internationally noted philosopher and you are well-known inter alia for your thoughts on the theory of rationality. The first book you published was *Practical Induction*; it was translated into German in 2010 ("*Praktische Induktion*", Paderborn, Mentis Verlag). How did you hit on that topic?*

I think one of the most interesting topics in philosophy is practical reasoning, which is the question of how you should figure out what to do. And if you look around, it's surprising what people think about this. They think, as far as I can tell, almost uniformly through the society—you see this as a presupposition of the way bureaucracies work, for example—that figuring out what to do is figuring out how to attain a goal. Sometimes it's put in a psychological register: it's figuring out how to satisfy a desire. But if you think about that—and John Dewey, the American pragmatist, was very clear about this—"I want it!" is a *childish* reason; in the strictest sense, it's a childish reason. And I think we can do better. In fact, grown-ups learn what matters and what is important from their own experience. They learn by being disappointed, they learn by being pleasantly surprised, and they generalize from those experiences; their concerns change; if all goes well their lives improve, and their lives collectively improve. But because we don't understand this in any theoretical way, people have to find this out in their own lives, each one the hard way.

So this is about making lives better?

Thinking about practical reasoning has the promise of improving decision-making; with any luck it can also improve institutional decision making, and it has upshots for lots of other areas of philosophy. This next is a more specialist's interest, but I think that when your views about rationality change, and especially your views about practical rationality change, there are ripple effects for other problems.

So you say that it's not only the case that there are things we want and that we have to find out how to get them, but that we are able to find new things that are worth pursuing. Can you show us an example of how I can find out what is worthwhile?

Here's a bit of a specialist's example. When I went into philosophy, I thought that the point of doing it was to establish philosophical results. You'd solve problems, the problem would stay solved, and then you could build on and use those solutions to solve further problems. And then at some point I realized that actually it doesn't work that way. There aren't any results. People who publish their results, or think they have solved a problem, are actually just providing material for other people to dismantle. But while I was making this discovery, I also discovered that it was enormously rewarding, intellectually, in terms of my improved understanding of these different problems, both to try to solve these problems and to establish results, and also to dismantle other people's. For what it's worth, my sense is not only that I've come to understand much of philosophy much better than I did when I started out – although perhaps some people will think that I haven't – but I've also discovered, by trying it out, that it was worthwhile to do that.

Suppose I have a child and she is in front of me and says “but I want!”. Is it a realistic approach to talk to her about her values or about the goals she has? Can I really accomplish anything by questioning them, can I get anywhere by saying “no, this is not a good goal to have”? Again, what makes something worth pursuing?

I can't answer the question as you asked it, because it's not like there is *one* thing that makes everything worthwhile. That's actually a way of approaching things which is entirely typical of philosophers, but which to me seems like a deep mistake. For example, G. E. Moore famously asked, “okay, we want to know what ‘good’ is”, as though there were this one property of all the things that are worth pursuing, and we just had to understand that. Often in philosophy, it's insisting that there will be *one* answer to many different questions that is the first wrong step. And it's very hard to recover from that.

So here is what I think: When you have that small child, and the child says “I want it!”, if you're a responsible parent you don't take the child's desires particularly seriously at all, mostly because you think they're not well grounded. So, for example, the child may strongly prefer to throw her food. That she strongly prefers to throw her food does not mean that there is a good reason for the food to be thrown. And if you as a parent think otherwise, it's just not good parenting. Now that doesn't mean that there is anything that you could say to the four year old who wants to throw her food. She first has to be trained to reason herself – and if you think about it, we don't train, not just children, we don't train the members of our society to be strong reasoners. It's just not part of what we do. And so, actually, when you look around, you'll find many adults who simply don't know how to tell what's wrong with an argument, or to construct an argument of their own. But that's the skill that's required for her to stand back and think about all the other things that she thinks she wants – of course by

the time she has the skill she won't want to throw her food anymore – and then she will be in a position to decide for herself, “It seems to me that I want this, but maybe it will turn out that that was a mistake.” There are different ways she could discover this: she could discover it by getting what she wants and being disappointed, and sometimes that's a **big** deal. Here is how she could do better: she could identify elements, components of the big thing that she wants that will disappoint her, and experiment with those elements first, and arrive at the conclusion that maybe she shouldn't be pursuing that big thing.

Would it be right to say that morality is just one part of practical reasoning?

Yes, that's right. Many decisions don't look like they have moral subject matter to anyone. Morality looks to me like a special case. I don't want to sound dismissive, but... so this will sound more dismissive than I mean. Nobody really knows how to say what counts as moral, what the moral subject matter is as opposed to other subject matter. But a sort of crude gesture might be this; this will capture what a lot of people think: morality is about being nice to people who you don't necessarily like. And if you think that's the subject matter, well, you have to figure out what to do, when the question is, how to be nice to people, and whether you should be nice to people who you don't like.

So work on practical reasoning will improve our moral theory, but only in the way it ought to improve decision making on any topic?

There is another way to think about morality, where thinking about rationality would give us real leverage. So again, this will not quite match up with the usual definitions of morality, but one of the things we expect morality to do is provide us with rules for the road for our society. It governs interactions between people, when they're anyway generic interactions. And of course not all of those, because, for example, we don't actually think of the literal rules of the roads anymore as a moral matter.

Now, think about the financial industry. As the financial industry has evolved, and there has been one crash after another, each time people think “we need to rewrite the regulations for the financial industry, now that the way it works has changed, and now that we recognize that there is a new problem”. The social environment, for which morality is the rules of the road, also changes on a regular basis, not least because division of labor is becoming ever more highly articulated. And so we have to be prepared to rewrite the moral regulations, in something like the way that we have to be prepared to rewrite the regulations for the financial industry.

To do this it really helps to understand what will count as good arguments for producing a new regulation. But also, because the users, the end-users of these

regulations, the people who have to apply them, have to be able to reason using these regulations as starting points. Since morality governs the public space, everybody has to be able to apply the regulations. And knowing how their reasoning is conducted, and how we can train them to perform that reasoning, is an essential precondition for figuring out what regulations it makes sense to provide them with. There is no point in giving them regulations they can't use. So for example, not that I mean to diss Aristotle—I think Aristotle was a deep and powerful thinker—we would not provide people with Aristotelian regulations. That is, we would not tell people “do as you think the wise person would do”. Those aren't good rules of the road, because people can't apply those regulations on their own. But morality has to be something that pretty much can guide just about everybody. They know what it tells them to do.

Can you say in a few sentences what gives practical induction its interest? What advantages does it have, compared to instrumentalism, for example?

Let me answer that question in two different ways. So – here is what I actually think is the most interesting thing about this idea: it's not already on a standard list. There's a short list of views that philosophers have had about how you conduct “practical reasoning,” and they treat it as though were a done deal. There are instrumentalists, and then there are Kantians, who think that you use the Categorical Imperative as your test, and then there are Aristotelians, who think of themselves as specificationists, that's the right way to say this, they ask “given a sort of loosely specified end or rule, how do I tighten it up to get concrete advice out of it?” And this list has looked to many people like it's stable, so they think it has to be this or that or the other. By putting something new on the list, what's interesting in the first place is not this particular option, but that there can be *more* options. What I most hope for is that other people go on to invent new ways of deciding what to do. I think the best thing that can happen is that this list of different ways of deciding what to do is indefinitely extended, and that there are regularly new additions to it.

You said you would answer the question in two different ways?

I can say something about what I think is the driving idea of this approach, of ***Practical Induction***. You know, moral philosophers are almost uniformly apriorist. And to put that into ordinary English: they act as though you ***already knew*** what morality was or what was important. And if you think about instrumentalism, when you say that what you desire is what gives you reasons, your desire is just your view about what's important right now. It's as though that you already have it ***makes*** it important. You don't have to look. But practical induction, learning from experience, is just the opposite. It's a posteriori, it's empirical. It doesn't have the certainty and the assurance that you ***already*** know, without looking, what you should be doing and what's important

and what matters. You're willing to find out. And I think introducing that is a shift of key, both with respect to decision-making, whether moral or not moral. And it would be a real step forward in our specifically moral thinking, to allow that we have to *find out* what morality is, what it requires.

If I see it right, practical reasoning hasn't lost any interest for you over these many years. You've worked on theoretical reasoning, too; among other things, in Hard Truths, you framed metaphysics as intellectual ergonomics. Your historical interests include John Stuart Mill, Friedrich Nietzsche and Oscar Wilde. You've written on David Hume, Bernard Williams and Robert Nozick. Is there one underlying idea that unifies these research interests? Or do you have a philosophical project in which you engage with several of them?

I've written about metaphysics as intellectual ergonomics, and philosophy of logic as intellectual ergonomics. That means finding intellectual devices that will help us solve the problems that we face. But the question is, which problems are those? Or what problems should we be focussing on? And I think there is a family of problems that philosophers need to take a hard look at, that have to do with this: we philosophers, and people generally, have systematically misunderstood which species we are. We have misidentified our species.

How is that possible? You know, Michael Thompson, at Pittsburgh, thinks that "human" is a first-species pronoun, in the way that "I" is the first-person pronoun. It should be as hard to misidentify your species as it is to misidentify yourself.

Here is a way to see this: Philosophers ever since Plato have understood that human beings exhibit division of labor. That's why Plato thinks we live in cities. But then he goes on to list the specialisations that you would find in a well-rounded city; it's a short list. He also understands that having different specialists live together will produce social coordination problems. And his solution to these problems is that there is one person, the Philosopher King, who by virtue of his training in what we would now think of as metaphysics, is going to be equipped to write the rules that will run the city, and to administer them. In retrospect this is a completely typical instance of how philosophers have failed to appreciate the depth of division of labor within a society.

There are two features of division of labor that we need to pay close attention to. One is that, nowadays at any rate, when you specialize, you acquire a representational vocabulary; you master a representational system that only those who have paid the overhead costs of becoming specialists understand. That means that nobody outside your field can understand you, literally. You also internalize a system of standards for your disciplinary pursuits. And of

course the system of standards is couched in specialized vocabulary, so no outsider can understand your standards and apply them, and of course, when you form your own preferences, these preferences ... you have internalized the standards, so the preferences are couched in terms of the standards, and so no outsider can understand what you want or prefer. And these standards almost always include standards of argumentation: what counts as a good argument varies from discipline to discipline. What this means is that the Philosopher King, Plato's Philosopher King, has to administer a society of people who are logical aliens to one another, whose preferences he cannot understand, whose standards he can't understand—who he can't understand at all. And we haven't even begun to think about how to do this. But this is the society we live in.

So, your picture is like multiculturalism, but for professions?

The problem isn't just that other people in the society think differently from us; the problem is that in a highly articulated specialized society everyone is everybody else's client. You have to rely on the products, epistemic and otherwise, produced by people whom you don't understand, whose work you cannot assess, and who are producing work to standards that you do not yourself accept. Nobody has even begun to think about how this is to be well done. It's pretty clear that what we manage now is pure improvisation and it's not well thought out.

Now the second aspect of specialization that we need to remember is that it's fluid. Bees specialize too, but the way bees specialize doesn't change over time. Human specializations change rapidly: new specializations appear, the way specialization works changes over time, and the configuration of specializations which provides services to one or another client changes over time. Here is a trivial example: It wasn't that long ago that there were people who specialized in providing whale blubber that was used to light people's houses. That came and it went. The configuration is ephemeral. And that's why, as I suggested earlier, moral theory is always a work in progress. Because the underlying configuration of specializations keeps changing, the rules of the road that are used to manage all of that will have to keep changing, in something very much like the way that the regulations that are developed for the financial industry have to keep changing as that industry changes.

Okay. That sounds interesting and important. But – are you optimistic?

You mean about solving these problems? Well, I think, the problems are hard, but I think we don't know how tractable they are simply because no one has ever tried to solve them. So if you look... I'll just stick with philosophers. If you look at philosophical theory, a surprisingly large part of it, and not just in ethics, is produced on the assumption that we are members of a simpler species, whose

environment does not change, whose social environment is not fluid in this way and whose members are mutually comprehensible to one another.

I'm going to produce a name for the species of creatures that I think we are, who are "Serial Hyperspecializers". We're *hyperspecializers* because we are extremely specialized, and we're *serial* hyperspecializers because none of us are hardwired to belong to one specialization; we are not bees. In fact, over the course of a human life, some people can move from one specialization to another several times, in extreme cases. Not twenty times, not a hundred times, but we see people who started out for example as chemists and then became factory floor organizers. These people can't come pre-equipped with goals that tell them what to do. Because natural selection can't have known ahead of time what goals the particular specialization will require. And somebody has to figure these out. So serial hyperspecializers have to learn what matters from experience as they explore new specialized forms of activity.

So practical induction is the form of practical reasoning that serial hyperspecializers need?

Yes. Now think about a much simpler animal. An animal that lives in a relatively unchanging environment, say a bird. Winter comes at the same time every year, and spring comes at the same time every year. And so the bird can be programmed with a trigger. When winter is on the way, the bird suddenly desires to fly south. And when spring is once again on the way, the bird suddenly has the desire to fly north. Instrumentalism works for creatures that are about as simple as migrating birds. And it doesn't work for us. When philosophers think that instrumentalism is a plausible theory, they're actually making a mistake about what species they belong to. That's a remarkable mistake to make. It's like mistaking yourself for a dog.

It's not just instrumentalism. Think about Kantian Categories. Something like Kantian Categories—which are necessary, you can't change them, they are a priori—are suitable equipment for a creature that lives in an environment that never changes. Maybe cavemen, human cavemen can have gotten by with Kantian Categories. But our environments change all the time, they change as deeply as you like; we create new environments that are *formally* as different from each other as you like. Now I have enormous admiration for Kant, but people who end up producing philosophical theories like this are producing theories for a different species of animal—an intelligent animal, not necessarily a bird or a dog, but a different species of animal. And once again: they have made the amazing mistake of getting wrong which kind of animal they are, which species they belong to. I'm not saying that a Kantian is mistaking us for dogs; but the kind of mistake they're making is *like* mistaking yourself for a dog.

When you describe the specialization of the agents who have to work together or live together, and where we must find rules of the road, that reminds me that in fact between two individuals it's the same thing, right? We could say, with everything you've lived through and all the experiences I've had and all the differences between our experiences and lives, we will never understand each other until it's over. But nonetheless sometimes we find people with whom we agree, one way or another, on how to get along. And if there is someone I'm not able to be in peaceful contact with, often I have the impression that it's only a question of distance. I have to work out how close I want to be with that person and if I go three steps back, there is no longer a problem because we don't collide. Sometimes it's impossible to keep the distance – and that's where moral philosophy comes into play; it's here that a normative theory is in demand.

I think that's absolutely right.

So we're searching for the same thing on another level?

Well, think about these differences. So this is the old way of thinking, and appropriate maybe for an earlier time. We say, look, people have different concerns and also have different private histories, but we need to find rules of the road for them. So we abstract away from these differences; maybe the solution we will find won't be suitable for everyone, but we'll treat those exceptions as noise, or the price we pay. And we formulate rules of the road that everybody can understand and master and follow in their interactions, and of course some interactions will be closer than that, but those will be optional. And so the rules of the road tend to be kind of simple. And we make them as infeasible as we can; we try to be absolutist about them. You know, don't kill the next person over, don't steal his stuff. But now we have a society where the people are much more specialized. And now the rules of the road have not simply to abstract away from personal idiosyncrasy, they have to accommodate the deeply different specializations. And that's a harder and more demanding task.

You mean, morality can't step back from the differences between people, and treat them all alike?

I'll give you a toy example. Actually, this is an example that you could have used even before today's level of specialization. John Rawls tells a legitimizing story for a liberal state. Citizens, in a well-organized liberal state, will understand that if their basic institutions were organized any differently, somebody in their society would do worse than the worst off person that they actually have in their society, and that would be unfair. And that is supposed to

give people a kind of intellectual and emotional stake in the institutions that they have. Now when Rawls is explaining the nature of the stake that people understand themselves to have, it's uniform. He says it consists in access to primary goods, roughly a claim on generic social resources, and what we might think of as... I forget what his term is, but it's something like "basic human rights". And he also thinks that people will, with respect to those basic rights and to the primary goods, people will take a maximin approach. They want to make sure that the worst they can come out is as good as it can be.

That sounds familiar. But why is specialization a problem?

No state can survive without a military, without security services: these are both internal security services and for external security. And security services that really work require a distinctive personality type. A good place to see what this personality type looks like is in the novels and in an autobiographical piece by an American popular author of a half a century back, called James Jones. He was not a great writer, but he was well-known for a while. He fought in the Pacific theatre in World War II. And he, his unit moved from Pacific island to Pacific island fighting against the Japanese. And Jones points out that... first, he describes a psychological transformation that happened to the soldiers in his unit, and presumably in every similar unit: at some point they realized that they simply were not going to survive the war. Although in his case of course that realization was a mistake—he couldn't have written a book if he hadn't survived—but it seemed to him, at the time, and it was true for many of them, that they wouldn't survive the war, and also that that was not that big a deal. It didn't really *matter*.

And also, obviously, if you're fighting your way through rainforests on Pacific islands, the level of creature comforts is not high. So, for a military to function, the people in it have to accept, as these soldiers did, that what we would think of as extreme discomfort simply doesn't really matter. And in fact a functioning military requires, if it's going to win a war, it has to be made up of people with that attitude. That is, people who don't care about primary goods or their basic rights, anyway, not in the maximin way.

So it's a precondition of the political state, of the sovereign political state, that there will be a class of people to whom Rawls's legitimizing story will seem irrelevant, and actually kind of childish. And this is just one instance of specialization. Which shows how it's harder to make our philosophical stories about this and that work than we had thought.

And I do think that these problems crop up again and again throughout not just moral philosophy but throughout metaphysics.

Okay. I can see it for moral or political philosophy, but how is specialization supposed to make a difference to metaphysics?

I'll give you a different kind of example in metaphysics. One of the tasks when you're doing metaphysics, thought of as intellectual ergonomics, is to characterize the function of the intellectual devices we already have. We need to understand them before we think about replacing them. And I think philosophers systematically misunderstand many of these intellectual devices. So think about necessity. When analytic philosophers approach the topic of necessity, they produce theories of invisible objects: possible worlds, in the first place. But if you imagine, if you remember David Lewis's version of this: the invisible worlds are very large, and no matter how long you travel, you'll never arrive at even the closest one. And when we talk about something being necessary, what we mean is that it's true in all of the possible worlds.

That's a characteristic way of approaching one of these problems, and it is meant as a semantic theory. It's the mode of approach to a problem that's traditional and characteristic of analytic philosophy. When analytic philosophers think about modality they're looking for a theory of the meaning of the modal expressions.

Now you ask: what could the *real* cognitive function be of saying that something is necessary? It can't be to talk about far away places of the David Lewis kind. From the point of view of evolutionary psychology—this is a point made by Robert Nozick, in his last book, *Invariances*—it's simply unbelievable that we would be selected for the ability to think correctly about other parallel universes. Here is a much more plausible suggestion, in my view anyway. When you say that something is necessary, that's signage that's used to manage attention. "Necessary" means "Don't waste your attention on alternatives to this". Attention is a scarce resource, we need signage like this.

But now think about two contrasting demands on attention-management guidance. Suppose you're within a specialization, talking to the other specialists. Say philosophers, in our case. When I tell somebody that he shouldn't be spending his time on some philosophical view or some philosophical approach, I can explain why. And actually within a field, I'm going to give fairly nuanced guidance as to what kind of attention to bestow and what kind of attention to withhold. I don't need signage like "necessary". "Necessary" is for outsiders. I can't *explain* to outsiders what they should pay attention to or not, and in a world of serial hyperspecializers there are always outsiders. "Necessary" is signage that you put up for people who are not specialists, and what it's there for is just to tell them to ignore things.

What you've been saying about serial hyperspecialization goes well beyond anything in Practical Induction. Are you writing about this now?

Yes, actually there's going to be a book, and with any luck it will be out in about a year; the working title is "**Serial Hyperspecializers and How They Think**".

Let's wrap up on that note. Thank you very much for your time.