

Electronic Conformism and Logical Elitism*

Elijah Millgram

Department of Philosophy

University of Utah

Salt Lake City UT 84112

elijah.millgram@gmail.com

May 11, 2025

The logic of defeasibility is not well understood, and its importance is scarcely appreciated. Here I'm going to try to convince you that getting a handle on it is urgent, because it is coming to have unpalatable—but hitherto unrecognized—political consequences. In particular, we are seeing the emergence of a new and very small elite, with the rest of us condemned to what I'll call electronic conformism.

I'll begin by introducing the concept of defeasibility. Then I'll explain how the optimization of administrative procedures enabled by a widespread shift to networked computing is making defeasibility management a very different matter than it had been. And then I'll tell you about an older East Bloc concept, *protektsia*, about how such now-familiar elements of day-to-day life as online forms are making it all-important, and about how there is suddenly much less of this social resource to go around. We have inadvertently been creating an unprecedented political reality for ourselves, in which a new elite

*I'm grateful to Chrisoula Andreou, Mel Andrews, Kevin Baum, Phoebe Chan, Svan-tje Guinebert, Eliahu Horowitz, Eric Hutton, Philip Klein, Jonne Maas, Hillel Millgram, Katharina Nieswandt, Thi Nguyen, Amelia Spolidoro, Laura Wharton and Michael Wilhelm for comments and helpful conversation, to an audience at the University of Konstanz, to the Australian National University's Research School of Social Sciences, for a place to work, and to the University of Utah, for a Sterling M. McMurrin Esteemed Faculty Award.

is marked out by a new form of privilege, which it's not unreasonable to describe as *logical* privilege. I'll conclude by describing some of its effects, and showing how it is bound to be broadly and deeply resented.

1

As we tell our students in their introductory logic classes, when an inference is deductive, the truth of the premises guarantees the truth of the conclusion. But most inference, or reasoning, isn't like that: typically, an argument from true premises can be perfectly alright, and get you to its conclusion—but there are further premises you could add in, which would make you take the conclusion back. Those extra pending facts, and there are indefinitely many of them, *defeat* the otherwise-alright inference, and so get called *defeaters*; inference that isn't deductive, we say, is *defeasible*.¹

For instance, and I'll make it an illustration that was recently topical, if you had your plane ticket in hand, you'd be on your way to Zurich next week. That was good enough to go with, but not if your PCR (remember: a test for covid?) came back the wrong way, and you were sent to quarantine instead; and not if your coronavirus entry form was denied; and not if you were hospitalized and being oxygenated; but also not if your alarm didn't go off that morning; and not if the train to the airport got stuck on the tracks by the conductors' work-to-rule stoppage; and not if you'd lost your passport. . . and it's clear you can keep adding entries to this list for as long as you like. And notice that the concept generalizes, past the standard truth-based introduction we gave it. Defeaters can be priorities rather than facts; perhaps you won't be on your way if a competing objective or ensuing side effect is *more important*: say, if it turns out you need to attend a funeral at home. Likewise, the premises and conclusions of an argument, when it's practical rather than theoretical reasoning, might be goals or evaluations or decisions. *Most* conclusions only follow *ceteris paribus*.

Contrast that with a quick illustration of a deductive inference, the sort of argument you might encounter in a math class: say, that for every context-

¹The terminology varies from field to field; in AI, the keyword is *nonmonotonicity*; in philosophy of science, it's *ceteris paribus laws*. For recent and representative discussions, see Hlobil, 2018, or Horty, 2012; Reutlinger *et al.*, 2017, surveys the topic as it is approached by philosophers of science and causation. For an older approach, see Adams, 1966.

free language, there's a push-down automaton that recognizes it. Because the truth of the conclusion really *is* nailed down by the argument, adding additional premises won't undermine it; it won't matter what else you find out about a particular context-free language, or about push-down automata, or for that matter anything else. Once the conclusion is nailed down, it's nailed down.

Deduction has held center stage in logic for as long as we know about (all the way back to Aristotle, in fact), and perhaps for that reason—and take my word for it just now—no one knows how to explain, in a reasonably general manner, why it is that an inference that goes through only “other things equal” goes through when it does, or doesn't go through when it doesn't. Because the list of possible exceptions is endless, you can't simply check them off and be done with it, and here's the deep reason for that. Treating inferences as defeasible is a way of anticipating the likelihood that there are things you've overlooked. Because you overlook much of what there is to overlook—that is, overlooking also happens one level up, at the meta-level so to speak, and you're generally unaware that there are issues of one kind or another to *be* overlooked—you're not in a position to draw boundaries around what needs doublechecking. You should always allow that there are more things you might have overlooked.

2

In the old days, defeasibility would be acknowledged but met with a shrug: sure, a lot of reasoning comes with exceptions, and when the time comes, the appropriate person will have to make them, but when we're doing logical theory, we don't much care about second-class inference.

However, the world is changing: in particular, more and more administrative and bureaucratic procedures are being reembodied as online forms. Such a form is a flow chart turned into code, and remember that to execute a program is to perform a series of completely explicitly specified steps; as computer science instructors endlessly remind their classes, DWIM—‘Do What I Mean’—is not a construct in any programming language. This entails that anything the web designer or programmer hasn't thought of isn't accommodated: to make a procedure electronic is to make the inferences embedded in it deductive, rather than defeasible.²

²However, it's no longer the case that all computation is procedural; I'll address the

Still sticking with those recently topical examples: a friend of mine visiting the US needed a negative PCR to get on her plane home; signing up and paying for it required her to fill out a form online; that form required her home address, which had to match her billing address; that address required her to select one of the 50 American states from a dropdown menu. It simply hadn't occurred to the programmer that someone from elsewhere might need testing. Returning home, on the mandatory entry form, she was required to choose between declaring her vaccinations as having been given in her country of arrival and having been given elsewhere; that next programmer hadn't noticed the possibility that the correct answer might be "both". To add injury to insult, the bottom of the form threatened criminal penalties for entering false information.

Bringing to bear the logical vocabulary we've just introduced, underlying the forms are inferences—say, from "It's an address" to "Included as a necessary component is one of the following fifty States"—and those inferences are normally properly thought of as defeasible. And once upon a time, when forms were filled out on paper, you could treat them that way: leave a field blank, or fill it in with an unexpected but appropriate value, or cross something out. But as processes and procedures are made electronic, the ability to register an exception, and *a fortiori* to have it accommodated, is being excised. It sounds precious when it's put abstractly: what's being imposed on us is having to treat properly defeasible inferences as deductive. But think back: how many times have *you* filled out a form online, found that it had no room for an unusual and so unanticipated circumstance, and left you unable to do something you had to do?

Once upon a time, when forms were filled out on paper, human beings processed those forms, and often enough it was done locally. Maybe your bank's forms didn't make sense in your circumstances, but the branch manager was right there, and you could talk to him about it. Nowadays, while you can still talk to a human being if you are willing to wait on hold for long enough, that person is sitting in a call center—or maybe is a temp working from home—looking at more or less the same form on their screen as you are, and is equally unable to make room for options it doesn't anticipate. The ability to grant people the exceptions they need to function as effective agents is gradually being squeezed out of our society.³ How often have *you*

impact of the competing paradigm below, in sec. 7.

³Do people really need exceptions in regular doses in order to function as effective

tried to explain your nonstandard circumstances to that voice on the other end of the line, only to realize that they have no more freedom of action than you do?

Even if defeasible inference looks second-class to a logician—a frame of mind that in any case seems unreasonable to me—in the increasingly electronic world we can no longer afford not to care about it.

3

In some of the highly centralized societies of the twentieth century—I mean, what used to be called the Second World: the Soviet Union and its satellites, as well as some other socialist states—there was a word for the social resource that mediated defeasibility management: *protektsia*, sometimes also referred to as “Vitamin P”. To have *protektsia* was to know someone able to intervene and solve your problem for you. Maybe you didn’t fall into a category that allowed the health service to administer that vaccine, but if you could call up the right person (or maybe if you knew someone who could call them up), you could get your shot anyway. Maybe the application paperwork for your community college assumed you were just now graduating from high school, whereas what you needed to do was pick up a prerequisite for your graduate program; but *protektsia*, if you had it, could have solved the problem. A highly proceduralized society divides up the population into the privileged elite, namely, those with *protektsia*, who are able to get the exceptions they need, and the *protektsia* have-nots.

That divide unsurprisingly became a source of deep resentment on the part of all those people who weren’t well connected. It was not typically understood in the West, where what in English get called “connections” have a lesser and qualitatively different kind of importance, but one of the primary attractions of capitalism to the less privileged strata of socialist societies was the appeal of a system in which *protektsia* was much less necessary. In a market, if you have the money, you can just *buy* what you need, no connections needed. Now remember how all that was in societies where administrative and bureaucratic procedures hadn’t yet been turned into web pages, and into apps you’re required to download to your phone.

When they have—that is, as we saw, when the underlying inference rules force themselves on you as deductive rather than defeasible—and when cir-

agents? Millgram, 2005, ch. 3, makes a case for the claim.

cumstances arise in which your own case involves what ought to count as a defeating condition, having an exception made requires someone to override the procedure. Unlike rules directly administered by human beings (remember that local bank manager), and with an exception I'll get to momentarily, in the electronic world overrides come in the first place only from those persons who are directly involved in specifying or implementing the relevant procedures—and in the second place, from the people placed to tell those directly involved what to do. But the financial incentive for converting paper procedures into electronic ones was that software scales almost costlessly: if you have the application up and running, the incremental cost of serving the web page one more time, or downloading one more instance of the program, is effectively zero. And for this reason, code written and maintained by an extremely small group of individuals can be applied arbitrarily widely. You can no longer solve your problem by calling up your local branch manager; how likely is it that you know one of the handful of upper-echelon executives or programmers off in the bank's San Francisco (or Manhattan, or wherever) headquarters—that is, the people who *are* able to fix things? *Protektsia* has become a resource that is simultaneously much more important and extremely scarce.

I mentioned an exception, and let's register it before proceeding: perhaps hackers can fix problems for themselves.⁴ Even if that's occasionally true, it's too rare to change the shape of the issue we're considering. How often will the system that a cybercriminal has breached be the one that allows him to tweak a procedure affecting his own life? There is a market for exploits on the dark web in good part due to such mismatches being the normal state of affairs. And because such markets are highly illegal, most members of the public are not participants. The chances that you know one of the vice-presidents of the financial institution you use well enough to make that 'can you fix it for me?' call were vanishingly small. Are the chances that you know someone who can clandestinely edit your financial institution's customer database any better?

The newly electronic world stratifies its inhabitants into populations above and below a *logical* divide: the privileged elite, who are in a position to treat defeasible inferences as defeasible, because they have Soviet-style *protektsia*,

⁴Although this is now standard usage, let's recall that it's very different from the original use of the word. For the former, current, and perhaps future senses of "hacker," see Schneier, 2023, pp. 9, 12.

and the have-nots, for whom the inferences underlying the computerized texture of their lives behave deductively, willy-nilly. And that elite is relatively much smaller than elites of the past.

It is also harder to see. We are used to thinking of privilege in terms of wealth or material goods or social status. But while wealth, for instance, may help one get access to the *proteksia* that allows defeaters for inferences to be broached and taken into account, it is not wealth itself, or social status itself, in which membership in this new elite consists. (There are a great many rich and powerful people who have to do things by the book, and sysadmins, who are positioned to make all the exceptions they want, are rarely rich and powerful.) This is a *logical elite*, and that's something our eyes aren't yet used to making out.

Is it really a privileged *class*? After all, the alleged logical elite, it might be objected, doesn't come with across-the-board privileges: someone who is able to edit your phone company's customer database is unlikely to be a superuser on any other system. Less directly, the supervisor who can tell him to make those edits isn't the boss of anyone in any *other* organization. If the ability to exempt from a rule is only local, isn't it too spotty to be the basis of a social partition?

In earlier *proteksia*-based societies, we saw local networks appear, roughly of individuals who were in a position to trade such favors for similar favors. Even the locally logically empowered are, we observed, going to be a relatively much smaller group than in the past; the safe assumption should be that, for the most part, they will find one another. And so, to a first approximation, we will see a division emerge, between those who can trade those favors and those who can't. In that case, even if that logically privileged class has not yet recognized itself, we can expect the sorting to happen before too long.

4

At this point I need to address an objection by an advocate of the new ways. To be sure, any procedure might reasonably run up against cases that ought to count as exceptions to the inferences it implements. But in a large-enough class of instances, an exception to a rule for which a case can be made will appear again and again and again, and eventually be noticed. When it is, the inference rules underlying the apparently rigid procedure can be adjusted. And so, gradually, the rules and procedures, and the online forms that are so

often their public face, will become contoured and articulated so as to remove any cause for complaint—or anyway, any cause for significant complaint.

Here we need to emphasize the logical differences between, on the one hand, replacing rules with more complicated rules, and on the other, exercising discretion. The inference rules we’re thinking about are implemented, that is to say, made real as procedures which have been concretized as, among other things, code. First of all, before adjusting such a rule, it’s necessary to understand the problematic cases, which in a complicated world populated by experts and specialists, trained in very different ways, is often nontrivial.⁵ After the decision to adjust it, it’s necessary to figure out how the replacement rule will work. And then there are the implementation costs, starting off with the coding, debugging, and trainings. Going back to an earlier example, once it’s decided to make room for courses of vaccination administered in more than one country, there has to be a procedure for invoking that condition; it will involve submitting documentation, so now it’s necessary to decide what documentation will do. (What certifications will you accept, from a long list of countries, including countries in which there’s deep political resistance to producing a standardized mode of verifying vaccination? Can you bypass all of that by just requiring serology? Then there has to be a list of facilities whose blood work you’ll accept, along with thresholds for titres.) And now you’re generating occasions for further defeaters: what if the individual covered by the clause also has, and not coincidentally, nonstandard documentation? Briefly, there’s overhead, and typically hefty overhead, when you adjust rules.

The objection we’re considering implicitly assumes that, even if the reasonable exceptions aren’t exactly finite, they’ll sooner or later run out or anyway taper off. But would-be (and *should*-be) defeaters can be as qualitatively different as you like: they can come from all over the place, and from any direction, so to speak; they *don’t* run out. This means that even if potential defeaters arise fairly regularly, any *particular* type of defeater may well be encountered only very sporadically.

A manager’s incentive structure is driven both by costs and by institutional convenience; when the overhead involved in making a change is hefty, that means ignoring outliers and tolerating churn. (And personnel with whom one will interact are given a matching incentive structure: when you’re talking to call center staff, and you have a problem that doesn’t fit

⁵See Millgram, 2015, ch. 2.

into the form in front of them, have you noticed that what they most want to do is transfer you to someone *else*, as quickly as possible?) So rules can't be adjusted often, and they certainly can't be adjusted on the fly.

And so the particular types of defeaters that arise infrequently won't be addressed by adjusted rules; it just doesn't pay. But cumulatively, such defeaters might well make up the bulk of difficulties a rule produces. *Most* of what should be exceptions, which taken together might well be common enough, might well go unaccommodated. And so you would expect the overall experience with a deductively-implemented and rigidly executed rule to be that it is insensitive even to a plurality of the would-be/should-be exceptions, and that it's recalcitrant to change.

5

Suppose you're amongst the large majority without much *proteksia*—what then? Presumably, every now and again you'll find yourself navigating a procedure built on inference rules for which your circumstances should supply defeaters, but which don't behave defeasibly. Over and above the frustration you'll experience, at being faced with tasks that are both mandatory and impossible, or at any rate with not being able to get done things you need to get done, there will be outcomes: that prerequisite course at the community college you were unable to enroll in, that flight you weren't able to get on, or rebook, that vaccination certificate you couldn't produce, the charge you weren't able to have reversed. After sufficiently many life disruptions, and after you tire of simple things becoming endlessly complicated, you will internalize the need never to be unusual or nonstandard anywhere you might encounter such an automated procedure, and not to consume exceptions. And then you will have learned *electronic conformism*.

That *sounds* bad, but let me take the time to point out why it actually *is* bad. About a century and a half back, John Stuart Mill argued for liberty—for the right to do as you see fit, provided you're not actually harming anyone else—on grounds that are perhaps still not fully appreciated. Whatever the merits of those arguments, those grounds provide a convincing account of the effects of organizing one's life so as to avoid occasions for defeasibility management.⁶

⁶His *On Liberty* can be found at Mill, 1967–1991, vol. XVIII, pp. 213–310; and his essay, “On Genius,” at Mill, 1967–1991, vol. I, pp. 327–339; Millgram, 2019, sec. 9.2, lays

Mill argued that without liberty, people would end up with lives that fit them badly, in something like the way that badly-sized coats don't fit. If you're not willing to do anything that will make you need an exception to the rules, that may well be true of you as well. He argued that being able to do things your own way, to act on what you have understood for yourself—in his lexicon, being 'original' and acting on the basis of your 'genius'—was the precondition of having an active mind; if you conform, you will over time become more and more habit-bound, which means getting through more and more of the day on autopilot. That's a plausible effect of standardizing one's life to fit into the dense mesh of self-executing procedures that one encounters at almost every turn. As a utilitarian, Mill noticed that the pleasures of life depend on *not* being on autopilot, and so there are straightforward hedonic consequences. As a youth, Mill had lived through a further psychological effect of living into patterns of action determined from without: the progressive disengagement of what one does from what one cares about. He had found those latter effects intolerable himself, and we can expect life under electronic conformism to be stultifying and dismal in much the way characterized by Mill, describing the conformists of his own day.⁷

But the costs are not only personal. The early decades of electronic proceduralization have been marked by rapid innovation and social change, but this may well prove to be a self-terminating process. Innovation, in its early stages, generally requires the ability to get exceptions to the procedures in place; exceptions to the rules are needed, because the rules are written for how things are done *already*. And as people who have started up businesses know, whether it's two hobbyists in a garage building the first personal com-

out the arguments; see also ch. 7, for the lessons of his famous 'Mental Crisis'.

⁷The life of the logical lower classes, in which you are motivated first and foremost to do things *the way you do them*, looks a great deal like David Velleman's dystopian "kinda Kantian" future (2009, pp. 149–157). And there is even a way the incentives that are supposed to produce it overlap with those we are considering here.

In Velleman's approach to moral philosophy, what is supposed to produce convergence on a society in which everyone does things in the same straightforward and standardized manner is the pressure to understand what one is doing. Now notice that when reasoning is defeasible, it involves judgment: why go ahead with the conclusion of the inference, when there are always more of the indefinitely many things to 'be equal' in its *ceteris paribus* clause still to consider? And those judgment calls are often, and typically, hard to explain. (Why are you ignoring the indefinitely many cases you haven't mentioned or even thought of?) If you can treat your inferences as deductive, you understand what you're doing better.

puter, or the new restaurant trying to get permits on a shoestring, or the gallery that hopes to display its sculptures outdoors, where they can be seen from the street, it won't work out unless you can get one person after another to cut you some slack. In the world of electronic conformism, that sort of slack is unavailable, to anyone outside the *proteksia*-enabled elite. And so innovation will have to come only from that elite. But, as we have argued, that elite is going to be remarkably small: *much* smaller, relatively, than dominant classes of the past. And so electronic conformism gives rise to another effect that Mill used in a different argument for liberty, namely, social stagnation.

6

Is the contrast really as sharp as that? I earlier on marked a small class of apparent exceptions to the generalizations I have been broaching—hackers, that is, people who are able to breach the security of a system and compromise it—but there is a somewhat larger and related group of individuals who end up treated more forgivingly by our electronic procedures. Even if they are not members of the core elite, perhaps they constitute its penumbra. So let's consider to what extent they complicate the picture we are sketching.

Even when you don't know someone who can turn the deductive inference you are facing into one that is defeasible, you might still *know your way around* that world of self-executing procedures. Being knowledgeable about algorithms and machine learning, and having the curiosity and willingness to poke around may allow you to figure out how an application or web page works. And then, even someone who can't change the rules, or claim exceptions which they don't allow, can, often enough, *game* them.⁸ For instance, and now anticipating a variant of the issue I'll take up in the next section, a little savvy experimentation with a preinterviewing avatar suffices to catapult you into the next tier of job candidates (the system is sensitive to how you comb your hair). Some fooling around with a transponder, and you can have your car tell the traffic light it's actually a bus, and should be given priority. You can't turn the dynamic pricing off, but if you understand how it works, you can use a VPN to give out that you're not a customer from a wealthy country but, perhaps, India; along similar lines, you can install an app that makes your iPhone—which shows you to be willing to spend the

⁸I'm grateful to Kevin Baum for pointing this out to me.

extra dollar—look to the outside world like a thriftier Android device. Building an app to make vaccination reservations—to retry and retry and retry, and snap up slots as they free up—is not the same as asking an insider to move you up the queue, but the two approaches get much the same results.

Or do they? The stances it proves possible for the penumbral elite to adopt are evidently akin to those that force themselves on the disempowered masses, and do not reflect the possibilities of the logical elite. Understanding the hidden contours of a rule you must treat as deductive is formally and substantively a very different matter than being able to suspend its operation and thereby exempt yourself. Being able to snag an earlier appointment for your vaccination is a very different matter from overriding an inference whose conclusion is that you are not eligible for it. The ability to game the system is the ability to treat its real rather than its announced rules as given—but it is to treat them as given, that is, as written in stone. Once it is clear what those rules are, members of the penumbral elite are, like the outsiders, strongly incentivized not to put themselves in the position of needing exceptions to them.

And the attitude natural to those availing themselves of their location in the penumbra could be described with a Britishism: as *dodgy*. If you are gaming the system, you are not exactly conformist, but members of the IT-literate penumbral elite must normally avoid having their stratagems noticed. That is to say, they must be outwardly conformist.

Moreover, although knowledge can become widespread in ways that connections cannot, the main lines of the society we are envisaging, and in which we have been coming to live, likely remain largely as I have been suggesting. Yes, there are a good many more computer science majors than there used to be, but most undergraduates are not a good fit for the major; despite the sizable financial incentives, relatively few people acquire these skills. The extended elite is somewhat broader than the core, but against the larger society, the penumbra is tiny.

7

As I remarked in passing, procedural programming and computation are no longer coextensive concepts. Although the substrate which modern AI runs on remains algorithmic, the virtual machine is trained rather than programmed; subsequent programming-by-prompt is very unlike specifying the

steps of a procedure. Big-Data machine-learning applications are also on track to providing a great deal of the substance and texture of our society. So here I want to do just a bit of compare and contrast, flagging a few similarities and differences.

When it's procedural coding, you know what your application is going to do—or rather, when you don't, you've found a bug—and you can explain why it did what it did. When you're deploying a neural net you've trained, you can't generally predict what it will do in particular instances, and you can't necessarily explain its outputs. (Although XAI—Explanatory AI—is a research program meant to address that second concern.) This means that framing the issues we're concerned with—or anyway their analogs—in terms of defeaters for inference rules would be a misstep.

Nevertheless, we can expect to see analogous social effects: when the case being handled doesn't look like the training set, outputs may well need to be overridden, and then, once again, the people in those circumstances divide up into the ones who can summon up those overrides, and the ones who can't. The ones who can't will, as before, experience the pressure not to stand out, to be as typical as possible, and never to need an exception.

But since modern AI isn't rule-based, what the logical lower classes will try to conform to is, rather, a *sensibility*. (The vocabulary we will need to draw on to reason about this technology is that of dispositions.) And that points us toward a novel challenge.

Aficionados of the later Wittgenstein will remember those tales of recalcitrant pupils: taught to count by twos, for instance, they make it safely to 1000, but then continue to count 1004, 1008... where the mistake can't be cleared up, because it seems to them that they're going on in the same way as before. They'll remember his Woodcutters, who sell piles of wood, not by weight or volume, but the area covered by the pile; when you point out to them that an unscrupulous retailer could up the price of a pile of wood by taking planks off the top and placing them around its sides, they look at you blankly and reply: "But then it's *more* wood, and it *should* cost more." And they'll remember Stanley Cavell's followon gesture, at the friend asked to feed the parking meter, who steps outside and mashes snacks into the coin slot. The point, Wittgenstein's exegetes tell us, is that in matters great and small we not only rely on sharing a sensibility with those around us; we expect that as we encounter variously novel cases, that sensibility will extend to cover our responses as before, and so that we will stay, as the idiom had

it, on the same wavelength.⁹

In those Wittgensteinian tales, his readers generally imagine the characters who don't see things our way as outliers, to be managed, ignored, or taken out of circulation: tacitly, the question is always, what do we do about *them*? And because they are just outliers, answering the question isn't very important. But now imagine the coming world of large-scale AI applications, and the likelihood that as cases unlike those in the training set appear, human and AI sensibilities will diverge. The investment in such systems will make their owners unwilling to walk away from them; but they also can't be reprogrammed in the usual sense. In that case, *we* will be cast in the roles of Wittgenstein's failed pupils and lunatic tribesmen—a concern we needn't expect to arise out of old-school procedural programming.

8

Karl Marx insightfully observed that social structures and ways of thinking fall into place around whatever is, in a given era, the economically central 'means of production'. For instance, when the agricultural production of foodstuffs was the central economic activity, control—not exactly ownership, by today's lights—of farmland was managed by an aristocracy.¹⁰ In the early industrial period, when manufacturing took center stage, ownership of investment capital became critically important; it was not that people ever stopped growing food, but political and cultural institutions reshaped themselves around the newly important sector of the economy. What means of production and what 'superstructure' do we need to open our eyes to, once we have recognized the inferential divide between those who are able to treat reasoning as defeasible, and those for whom it is, willy-nilly, deductive?

"Information technology," or "IT," is evidently the driver of that divide, and so one might expect that what is being produced is, first and foremost, information; we are sometimes said to be living in a "knowledge economy". (Recently, we were said to have moved toward a "service economy".) That is

⁹Wittgenstein, 1998, §§143–46, 151–57, Wittgenstein, 1983, §§I-148–152, Cavell, 1999, pp. 168–190 and esp. 183. Arvan, 2024, adapts the Kripke/Goodman discussions of rule-following and projectability to argue that we can't tell how large language models and other reinforcement-trained AI will 'go on,' but on grounds that are perhaps so general that it's hard to tell how much practical bite they have.

¹⁰For fascinating nuts-and-bolts documentation of that observation in a entirely non-Marxist accent, see White, 1964.

misleading: the economically central product, that occupies the role formerly held first by agriculture and then by factory-based mass production, is self-executing procedures: algorithms, realized in the first place as apps or as web pages.¹¹

So we should expect to see emerging new concepts of assets to be controlled: perhaps IP (intellectual property), reconceived as ownership of code; perhaps semi-monopolies supported by network effects. We should expect to see new ideologies emerge, supplying justifications for the activities and privileges of the new logical upper class. And if our argument has been on track, we should expect a new focus on the means of production, a critical aspect of which is the phenomenon we have been discussing here: the ability to treat inferences as defeasible, and to be able to call down exemptions from those who design and implement the endlessly many procedures that, intangible as they are, now form the primary structural material out of which our lived world is constituted.

References

- Adams, E., 1966. Probability and the logic of conditionals. In Hintikka, J. and Suppes, P., editors, *Aspects of Inductive Logic*, pages 265–316, North-Holland, Amsterdam.
- Arvan, M., 2024. ‘Interpretability’ and ‘Alignment’ are fool’s errands. *AI & Society*, *TBA(TBA)*, 1–16.
- Cavell, S., 1999. *The Claim of Reason*. Oxford University Press, Oxford.
- Hlobil, U., 2018. Choosing your nonmonotonic logic: A shopper’s guide. In Arazim, P. and Lávička, T., editors, *The Logica Yearbook 2017*, pages 109–123, College Publications, London.
- Horty, J., 2012. *Reasons as Defaults*. Oxford University Press, Oxford.

¹¹Here’s a way to convince yourself: It’s no longer possible to protect information; companies’ networks are regularly broken into, and their data is distributed to all the world, or sold online. This doesn’t seem to interrupt their business activities, however. As long as their procedures keep running, the business stays viable, even when they no longer control the information they have accumulated.

- Mill, J. S., 1967–1991. *Collected Works of John Stuart Mill*. University of Toronto Press/Routledge and Kegan Paul, Toronto/London.
- Millgram, E., 2005. *Ethics Done Right: Practical Reasoning as a Foundation for Moral Theory*. Cambridge University Press, Cambridge.
- Millgram, E., 2015. *The Great Endarkenment*. Oxford University Press, Oxford.
- Millgram, E., 2019. *John Stuart Mill and the Meaning of Life*. Oxford University Press, New York.
- Reutlinger, A., Schurz, G., and Hüttemann, A., 2017. Ceteris paribus laws. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*, Metaphysics Research Lab, Stanford University.
- Schneier, B., 2023. *A Hacker's Mind*. W. W. Norton, New York.
- Velleman, J. D., 2009. *How We Get Along*. Cambridge University Press, Cambridge.
- White, L., 1964. *Medieval Technology and Social Change*. Oxford University Press, Oxford.
- Wittgenstein, L., 1983. *Remarks on the Foundations of Mathematics (revised edition)*. MIT Press, Cambridge. Edited by G. H. von Wright, Rush Rhees and G. E. M. Anscombe; translated by G. E. M. Anscombe.
- Wittgenstein, L., 1998. *Philosophical Investigations (third edition)*. Blackwell, Oxford. Translated by G. E. M. Anscombe.