

Title	Scientific metaphysics and social science
Authors	Ross, Don
Publication date	2023-10-27
Original Citation	Ross, D (2023) 'Scientific metaphysics and social science', Synthese, 202, 146 (34pp). doi: 10.1007/s11229-023-04358-1
Type of publication	Article (peer-reviewed)
Link to publisher's version	10.1007/s11229-023-04358-1
Rights	© 2023, the Author. Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit http://creativecommons.org/licenses/by/4.0/. - https://creativecommons.org/licenses/by/4.0/
Download date	2026-09-23 07:04:16
Item downloaded from	https://hdl.handle.net/10468/15179



University College Cork, Ireland
Coláiste na hOllscoile Corcaigh



Scientific metaphysics and social science

Don Ross^{1,2,3} 

Received: 9 June 2022 / Accepted: 18 September 2023
© The Author(s) 2023

Abstract

Recently, philosophers have developed an extensive literature on social ontology that applies methods and concepts from analytic metaphysics. Much of this is entirely abstracted from, and unconcerned with, social science. However, Epstein (2015) argues explicitly that analytic social metaphysics, provided its account of ontological ‘grounding’ is repaired in specific ways, can rescue social science from explanatory impasses into which he thinks it has fallen. This version of analytic social ontology thus directly competes with radically naturalistic alternatives, in a way that helps to clarify what makes some metaphysics genuinely scientific (that is, part of the scientific enterprise and worldview). I consider this competition, marshal considerations against the value to social science of analytic metaphysics, and sketch a contrasting scientific metaphysics for understanding the implications of revisionist social ontology in unified scientific ontology.

Keywords Philosophy of social science · Social ontology · Scientific metaphysics · Analytic metaphysics · Supervenience

1 Introduction

My aim in this essay is to describe a program for extending radically naturalistic metaphysics to the philosophy of social science. By a ‘naturalistic metaphysics’ I refer to metaphysics that applies no transcendental concepts or categories that do

✉ Don Ross
don.ross931@gmail.com

¹ School of Society, Politics, and Ethics, University College Cork, Cork, Ireland

² School of Economics & Research Unit in Behavioural Economics and Neuroeconomics, University of Cape Town, Cape Town, South Africa

³ Center for the Economic Analysis of Risk, Robinson College of Business, Georgia State University, Atlanta, USA

not feature explicitly in first-order scientific theories, models, or explanations.¹ By a ‘radically naturalistic metaphysics’ I refer to a metaphysics that is motivated by service it can perform to the delivery of scientific knowledge. This need not involve ‘helping’ any particular science along, but might involve unifying or synthesising scientific discoveries.

There are two relevant contrasting views that help to frame the intended point of radically naturalistic metaphysics. One view, common among some naturalists, especially those who emphasise a strong empiricist stance, is that metaphysics is irrelevant to the pursuit of scientific knowledge or, if taken seriously as a source of constraints on hypotheses and models, should be expected to impede this pursuit. The other view is that useful metaphysics requires some special metaphysical concepts that do not derive from applications in science or are not answerable to their uses in science where refinement is concerned.

The three views just described are, as formulated, both sweeping and inexact. There is often room for argument about which of them is closest to the stance of a particular philosopher or specific application, particularly if the philosopher in question is only tangentially concerned with meta-philosophy. But most explicit defenses and criticisms of the sweeping positions have taken ‘natural’ science, especially physics, as the relevant test-bed where relationships between metaphysics and science are concerned. The object of critical attention here is the much smaller literature that addresses the relevance and value of metaphysics applied to or integrated with social science.

I will not be concerned except in passing with the view that metaphysics is irrelevant to social science. Since I will defend a positive view, the reader will be able to see directly how it contrasts with anti-metaphysical empiricism, at least in general. But scepticism about the value of metaphysics to social science cannot simply be bypassed altogether, because it is entangled with varying opinions about what metaphysics *is* and *is for*, which are central to what follows. If a philosopher has a very narrow conception of metaphysics, less is implied or at stake if they then say that it doesn’t matter to science.

It will be crucial to what follows that I will not suppose that a theorist is engaged in metaphysics whenever they take a view on which things and processes exist. As several philosophers (e.g. Kincaid, 1996; Guala, 2016; Lohse, 2017; Lauer, 2019) have argued, social scientists, like natural scientists, often pose and investigate ontological questions of this sort. But questions such as ‘Are there such things as business cycles?’ or ‘Are multiple personalities real?’ are generally framed within the assumptions of discipline-specific models. An ontological question is metaphysical, I will suppose, just in case it partly depends on a view of what *kinds* of things and process exist, in general, that is intended to transcend specific models and contexts of disciplinary investigation.

¹ I do not think there are decisive considerations either for or against the view that mathematical knowledge transcends empirical inquiry. But in any case many or most mathematical ideas feature in first-order science.

For my main purposes here, the foil of radically naturalistic – or, as I will henceforth say ‘scientific’ – metaphysics is *analytic* metaphysics.² By this I refer to a tradition that, in its contemporary form, finds early exemplars in Strawson (1959) and the papers gathered in Lewis (1983, 1999). The tradition generally supposes that all people who aim at systematic and general knowledge, including scientists, presuppose some system of categories or modes or forms of existence that transcend specific observations, measurements, model-relative constructs and paradigm-relevant classifications. Often this framework is taken to be culturally inherited ‘folk’ metaphysics. But analytic metaphysicians generally appeal to a specifically philosophical system of ideas. Such systems are usually viewed as descended from historical schools of thought in the accepted philosophical canon, but usually heavily refined by applications of technical ideas from modal logic and semantic theories. Often they appeal to distinctions between grades of possibility and necessity that are unknown in first-order science. Analytic metaphysicians appeal to special varieties of dependence relations between elements of reality, going beyond causal dependence, in which scientists do not traffic.

By contrast, a naturalistic or scientific metaphysic as I will understand it here is restricted to use of concepts and distinctions that feature in first-order science. The motivation for this restriction, explored at length in Ladyman and Ross (2007), is the naturalist’s belief that science is the only reliable institutional setting for filtering objective, general, systematic (Hoynigen-Huene 2013) knowledge (Bird, 2022) from the penumbra of human beliefs, including individuals’ subjective convictions (ibid.), and philosophical beliefs, either folk or professional. Naturalists doubt that argument and reflection, divorced from controlled empirical experimentation and measurement, is a reliable knowledge filtering process, and they deny that there is any basis other than conventions of scientific practice for pronouncing general facts about the ‘correct’ or ‘best’ uses of concepts. Naturalists do not deny that conceptual analysis can sometimes be useful for clarity of communication, but deny that it should serve a regulatory role with respect to the custodianship of scientific knowledge. The naturalist can acknowledge that philosophers may be experts on how concepts have historically been used – though in the case of scientific concepts, practicing scientists often know better – but denies that it is possible for anyone to be an expert on the *normative* use of concepts. From time to time, what has seemed like a substantive question of fact to scientists dissolves into a merely semantic issue, and philosophers might assist in identifying some such dissolutions. But then further study of the issue in question is for anthropologists or linguists studying ways in which some communities of scientists express themselves.

² Simply for practical reasons, I will entirely bracket questions about the status of various ‘continental’ flavours of metaphysics with respect to naturalism and science. The kind of naturalism I defend has much in common with that of Rorty (1999), except that it is not anti-metaphysical. Rorty’s anti-metaphysical rhetoric is descended from Heidegger’s, and that is why I mention it in connection with ‘continental’ attitudes. But these are complicated issues that merit full treatment in themselves on another occasion. My view is that Heidegger and Rorty declare opposition to metaphysics on the basis of taking historically narrow views of what metaphysics must be, privileging the conceptions of Plato and Kant at the expense of Aristotle’s.

It is relatively straightforward to distinguish naturalistic from analytic conceptual frameworks in physics because the ontologies explicitly found in physical theories are obviously not folk ontologies, and require significant technical effort to be interpreted in terms of professional metaphysical frameworks. Thus one can ask with at least presumptive clarity how the scientific ontology might be interpreted in the frame of the externally articulated metaphysics; or, following Ladyman and Ross (2007, 2013), one can criticise the motivations and value of such interpretation. One can then further inquire as to whether resistance to such interpretation leads to a denial of a metaphysics of physics, or to a new style of metaphysics that is based on physics – a radically naturalistic or scientific metaphysics. The stakes in this assessment are also relatively clear, because if one were to deny the value of a metaphysics of physics, one would thereby automatically put the burden of argument onto anyone who wanted to still defend the prospects of *any* metaphysics. This is because metaphysics, whatever else it is supposed to do, should provide insight into *general* structures of reality. And physics is the part of science that applies most generally. So it would at least be *prima facie* odd to claim that there is a defensible – i.e., more or less most factually adequate – metaphysics, but that the generalisations of physics stand outside of it.

These clarifying ground rules do not apply when we consider social science. The ontologies of social sciences interact multifariously with the elements of the manifest social image. Some philosophers (Ruben, 1989; Thomassen, 2003) maintain that they are necessarily directly answerable to the terms of that image. Furthermore, there is important work in philosophy of social science that depends on the interanimation of everyday and social scientific ontologies. Philosophers such as Ásta (2012) and Haslanger (2018) seek to critically disrupt social ontologies that trap people in oppressive and restrictive categories or binary options, and they do so, in part, by appealing to well confirmed social science. Many philosophers conceptualise this kind of work as useful social metaphysics.

That the reformist work is useful, indeed important, is not in doubt. And clearly it is about ontology. But it does not require appeal to analytic metaphysics, even if its practitioners occasionally deploy language from that tradition.³ Guala (2016, pp. 194–205) provides a persuasive basis for understanding reform of social ontology in naturalistic terms. I interpret work such as Ásta's and Haslanger's as urging displacement of folk ontologies by ontologies that better accord with work in social science, but motivating such replacement partly on moral and political grounds, rather

³ Sarkia and Kaidesoja (2023) criticise the broad philosophical 'social ontology' movement as being (i) too detached from social-scientific practice; (ii) too tolerant of a priori reasoning, and (iii) aimed at generalisations too broad in scope to respect the rich heterogeneity and methodological flexibility of social science. My criticism of the use of analytic metaphysics in social ontology will echo points (i) and (ii). My target is analytic metaphysics, not social ontology in general. But my point about "overly relaxed" naturalists indicates why I think that SK's criticisms gain wider traction. As will be clear, analytic metaphysics inherits its fatal flaws from the application of folk metaphysics to science. And much of the social ontology literature combats oppressive folk metaphysics by promoting normatively preferable *alternative* folk metaphysics, instead of either avoiding metaphysics or being stringently naturalistic about it. *Any* folk metaphysics treated as prior to science, even if it aligns with currently progressive conceptions, is likely to ultimately become a conservative barrier to empirical knowledge. And conservatism *always* leads to oppression of people, ideas, and practices deemed eccentric.

than exclusively on epistemic ones. I suggest that the tendency in this mainstream philosophical literature to treat ‘ontology’ and ‘metaphysics’ as loosely synonymous stems from the fact that most contemporary philosophers ecumenically believe that ontology is discoverable (though often difficult to discover), that it is constrained by both logic and (philosophically refined) folk insights, and that empirical science is essential to discovering interesting ontological novelties. That is, most philosophers are naturalistic in a relaxed way, but not radically so. They blend appeals to science and to science-transcending conceptual logic in flexible and unsystematic ways and thereby stay focused on their main prizes: accounts of human ontologies that are in empirically good shape, that escape the dead (and deadly) hand of culturally inherited alleged essences, *and* that are readily accommodated by the informal logic of conceptual criticism that we teach to young philosophers.

My view is that we best achieve the practical ambitions of the critical social ontologists by being more consistently scientistic about ontology. Logicist conceptual arguments for and against ontological restrictions *or* innovations do not generally carry persuasive weight with non-philosophers. My contention here is that they should not persuade philosophers either.

Epstein (2015) *also* rejects the ecumenical and relaxed view. His objective is to synthesise a reformist analytic metaphysics for social science. By ‘reformist’ I now refer not to the reformist social project of Ásta and Haslanger, but to Epstein’s contention that social scientific practice is significantly hampered by its lack of attention to metaphysics. The style of metaphysics he has in mind is analytic metaphysics, as characterised above. Specifically, he claims that model selection in social science is undermined by specific confusions that conceptual tools of analytic metaphysics can be used to repair.

In defending the value of radically naturalistic metaphysics to social science, my method will be to use Epstein’s reformist project as a critical foil. This helps to illuminate the approach I prefer, for two reasons.

First, Epstein rhetorically commits himself to concern for social *science* rather than merely, as in the large and well-known literature to which he mainly responds, the status of ‘the social’ in metaphysics itself. Thus his book allows the contest between analytic and naturalistic metaphysics to be joined on shared ground with respect to a main objective: facilitating the production of good social science.⁴

Second, Epstein shares with a radical naturalist the view that the vast literature on supervenience in philosophical social ontology has been delivering diminishing returns with respect to elucidating social science. The bodies of reasoning that generate our agreement in this opinion, which I will review in Sect. 2, are different. But they both advert to a shared view of the point and ambitions of metaphysics as not simply equivalent to specifying ontologies.

The core premise for Epstein’s general argument is that social sciences as practiced fail to attend to the relative *fundamentality* of different facts and properties. Radical naturalists, by exact contrast, deny that there is any sound basis for applying

⁴ I say ‘rhetorically’ because, as Guala [2022] argues, the conception of a ‘social property’ on which Epstein’s arguments rely is much more general than the conception of a social *scientific* property to which attention should be restricted by his reformist ambition.

any concept of fundamentality in science. Fundamentality is a concept derived from a priori philosophical metaphysics, not from science. By Epstein's lights, then, it would seem that the naturalist must deny the possibility of metaphysics. It would be uncharitable to Epstein to suppose that he is unaware that some philosophers seek to naturalise metaphysics. In his book he never considers the issue because, almost certainly, he regards it as uninteresting for *social* science. Here lies a key part of my motivation for the present project of extending the radically naturalist metaphysical agenda in the context of social science. Many, perhaps most, of the relaxed naturalists in the philosophy of social science, as I characterised them above, doubt that the extension can work.

Lohse (2017, p. 4) states the challenge well. Some philosophers who are interested in metaphysical issues in social science, Lohse says,

... might be aiming at naturalizing metaphysics, that is, pursuing some kind of scientifically informed metaphysics that attempts to paint an accurate overall picture of the world – in our case, the social world – that is compatible with, or constrained by, or based on current state of the art of our best social sciences. This seems to be a sensible project in the philosophy of the natural sciences (Ladyman et al. 2007) [Ladyman & Ross, 2007 in references here], and it would be a legitimate – though at the current state of the social sciences, hardly achievable – project in POSS [philosophy of the social sciences].

This project, Lohse goes on, is not “*meant*, at least not in a straightforward sense, to be relevant for the social sciences in the first place”. Likewise, Lauer (2019, footnote 1) describes the project to radically naturalise metaphysics as “running parallel to” the issues that concern Epstein. With an important qualification related to Lohse's remark about “achievability”, my aim here is to indirectly contest this complacency about the radical implications of scientific metaphysics. I will do this through identifying the *form of value* that scientific metaphysics could bring to the social sciences, by contrasting it with Epstein's analytic approach. To stress: I do not think there is a current third alternative on the table, other than denying that metaphysics is relevant to social science, because I agree with Epstein – and Lauer and Guala – that refining supervenience concepts is not a path to a general metaphysical view of social science.

The paper is structured as follows. In Sect. 2, I review debates over supervenience, and explain why the radical naturalist joins with Epstein in regarding these as red herrings where genuine metaphysics is concerned. In Sect. 3, I criticise Epstein's programme for reform of social science based on analytic metaphysics. In Sect. 4, I provide a case study, in which I criticise Epstein's view of the relationship between microeconomic and macroeconomic ontologies from the perspective of practice in economics. In Sect. 5, I indicate how naturalistic (or ‘scientific’) metaphysics departs from the analytic variety, and discuss the implications of such metaphysics for social ontology. Section 6 briefly concludes the essay.

2 Beyond supervenience

Before getting to the main argument of the paper, I turn to a point of agreement between the ‘unrelaxed’ analytic metaphysician, as represented by Epstein, and the scientific metaphysician. This concerns the limited returns available from the long-standing preoccupation with technical concepts of supervenience in the literature that applies metaphysics to science. The issue is important here because if a philosopher thinks that social-scientific ontologies generally supervene on uncontroversial folk ontology (individual people and their individual actions) plus ontologies that are entangled with science only at the levels of biology, chemistry, and physics, then she can happily set aside the issues that arise in the dialectic between Epstein’s reformist programme and the radical naturalist one – which, as we will go on to see, are complex and challenging. Supervenience, that is, promises the philosopher a quiet life where metaphysics is concerned. But the promise can’t be delivered on.

The primary spark for the emergence of supervenience as a central idea in philosophy of science was Fodor’s classic (1974) paper that generalised issues that had arisen from the displacement of mind-brain identity theory by functionalism as the dominant position in the philosophy of psychology. Fodor’s topic of concern was not metaphysics *per se*. What he aimed at was showing philosophers that identification of abstract scientific types with sets of (in principle) relatively directly observable tokens was not a general explanatory or modeling strategy in the sciences, and that abandonment of this strategy by cognitive scientists represented no *ad hoc* innovation. The exact title of Fodor’s paper, “Special sciences, or the disunity of science as a working hypothesis”, is clearly indicative of its objectives. “Special sciences” denotes every science outside of physics. The rest of the title signals that Fodor’s foil is another classic paper, Oppenheim and Putnam (1958), which had aimed to empirically verify the logical empiricists’ reductive account of the unity of scientific theory by offering evidence that the world is accurately described as a monotonic *assembly* of ‘fundamental’ physical constituents. Of course this fact, if true, would be highly favourable to the traditional metaphysical doctrines of atomism and physicalism. But the phrase “as a working hypothesis” in the titles of both papers tells us with utter clarity that what is propounded is not a transcendental doctrine, or its denial in Fodor’s case, but an empirical proposition about the basis for successful scientific theory and model development. Fodor’s paper was that rarest of philosophical achievements, an argument that ultimately convinced almost the entire community of intended readers.

The wide scope of influence of Fodor’s paper in the philosophy of science is explained not merely because its conclusion reaches across most of the sciences, but because it also offered deep illumination on central topics beyond reduction and unity, specifically the relationship between statements of laws of nature and causal generalisations. These are not topics on which science institutionally aims to establish empirical consensus. They are topics of metaphysics (including naturalistic metaphysics, since they arise within science). But Fodor did not attempt in his paper to *analyse* them. Analytic metaphysicians, whose enterprise was just then starting to become respectable after decades of suppression by logical positivism and empiri-

cism, took up the challenge energetically.⁵ As in most post-Lewis analytic metaphysics, the primary concept on which they relied was the logically possible world.⁶

A next milestone in the literature was Kim (1998), which argued that under this style of analysis supervenience is *unstable*, tending to collapse into either classic reductionism or eliminativism with respect to presumptively supervenient types. Though Kim returned the focus of attention to the literature's original home in the philosophy of mind, his argument is based entirely on *a priori* analytic metaphysics – ultimately, issues in applied modal logic – and not at all in empirical psychology or any other science. Ross and Spurrett (2004) argue that Kim's analysis indeed resists any attempt at useful application even to interesting ontological problems that actually arise in cognitive science, a conclusion that is taken up and expanded upon in Ladyman and Ross (2007, Chap. 4). While philosophers are naturally dissatisfied about a loose end until and unless they can agree on a solution to Kim's cluster of problems, no scientists are waiting on the solution.

But supervenience need not be regarded as a metaphysical relationship that requires any reference to possible worlds. Observation might establish an empirical relationship between a specific type and a set of tokens at a less abstract level of analysis such that any inference about the type licenses, with potential for error and pending empirical surprises, an analogous inference about the tokens of another type from a different 'level'. It is not obvious that one can defend physicalism, as the belief that every existing empirical structure supervenes on the world as described by physics, without entertaining a metaphysical belief. However, many philosophers and scientists still think that actual known minds supervene on actual brains, without therefore having to suppose that minds supervene on brains across all or even any *a priori* delimitable sets of possible but non-actual worlds. Similarly, that chemical bonds supervene on sub-atomic structures is surely the majority view. It is routinely claimed by philosophers of economics (Guala, 2022), and at least one economist (Hoover, 2009), that macroeconomics supervenes on microeconomics. So belief in supervenience is compatible with naturalistic metaphysics, as long as the concept is not used as a trojan horse for dragging in analytic metaphysics.

That said, a naturalist philosopher is typically sceptical that there is a single *general* concept of supervenience that applies in a scientifically informative way across all of these examples. Shorn of the analytic apparatus for defining it, each real supervenience relation is approximate and specific in its supporting mechanisms and value for inference. For my part, I am sceptical about all three of the above examples, and am unconvinced that there are *any* real cases of supervenience at the scales of whole disciplines or major sub-disciplines. There are very likely some local instances within sciences, and around the fuzzy boundaries between basic physics and basic chemistry.

Epstein, while holding supervenience to be a general metaphysical relationship specifiable in terms of possible worlds, also doubts that it is a general basis for the unity of science. This is because he thinks that relatively abstract ontologies tend

⁵ The first two decades of the resulting technical literature is represented in Savellos and Yalçın (1995).

⁶ Strictly speaking, they often distinguished logically and metaphysically possible worlds. But the analyses are always given in terms of propositional or deontic or modal logic.

to be ‘open’ in the less abstract properties that their principles of observation track. Both Epstein and radical naturalists think that many philosophers are complacent in supposing that they can resist the conclusion of Dupré (1993) and Cartwright (1999) that science presents us with a disunified or ‘dappled’ general picture of the world merely by appealing to a general supervenience concept. And it is not clear that Epstein’s reasons for scepticism about the philosophical value of supervenience claims depends on interpreting it metaphysically. His argument that macroeconomics does not supervene on microeconomics (Epstein 2014) *does* rely on a metaphysical interpretation of the relationship and on toy examples rather than attention to actual working economic models, an issue to which I will return in detail later. So in this instance I agree with his conclusion but not his reasoning. However, he elsewhere attacks another putative supervenience relationship, of social network relationships on the structures modelled by agent-based simulations (Epstein, 2011), on the basis of considerations that are empirical and derived from the real scientific ambitions of the relevant models. For my part, as a naturalist, I am persuaded by the conclusions of that paper for the reasons Epstein gives; so it does not depend on our divergent opinions about metaphysics.

It is not as surprising as it might seem that a strong proponent of the relevance of analytic metaphysics to science, and a promoter of scientific metaphysics, are in general agreement about the significance of supervenience. Both are motivated to resist the hypothesis of the dappled world. In the case of the analytic metaphysician the reasons for this are obvious, and then the scepticism about supervenience can be motivated by the failure of philosophers to agree on a general solution to Kim’s cluster of technical problems. For the scientific metaphysician, the problem is that the claim that the world described by science is dappled is equivalent to following van Fraassen (2002) in saying that science leaves no place for metaphysics (Ladyman & Ross, 2007, Chap. 2), scientific or otherwise.⁷

The aim of Epstein (2015) is to construct stronger and much more demanding general foundations for social science than supervenience, based on analytic metaphysical concepts. Core to his enterprise is the conviction that some very general kinds of facts are more fundamental than others, and that judgments about relative fundamentality transcend empirical discoveries that are assessed according to methods and conceptual schemes specific to scientific disciplines. Philosophers are held to have special expertise where knowledge of fundamentality is concerned, and can use this expertise to help scientists, and social scientists in particular, achieve more stable ontologies that can foster improved epistemic progress. By exact contrast, naturalism of the kind I defend denies that there are general relations of fundamentality, and that there is expertise about what should be explained by reference to what that transcends the specific histories and practices of scientists. On the other hand, this form of naturalism allows that metaphysics is possible and potentially useful not just to ‘natural’ science but to social sciences as well, without depending on general fundamentality judgments. The core source for such metaphysics is Ladyman and Ross (2007), and

⁷ Dupré (1993) sometimes says that disunity is itself a metaphysical hypothesis. Ladyman and Ross (2007, Chap. 4) argue that this is not a coherent view.

the substantial subsequent literature that it inspired; other key sources are Ross et al. (2013) and French (2014).

3 Epstein on analytic metaphysics and social science

Epstein's book is full of clear structural signposts, making the core features of his enterprise *for my purposes* easy to summarise. The point of the emphasised phrase is that I will bypass intricate details that an analytic metaphysician might highlight for *her* purposes. That is, I will neglect fights in which a naturalist has no dog. The focus is exclusively on what metaphysics can allegedly do for science.

Epstein's book has two main objectives. One is to alert social scientists to the idea that explaining facts about social kinds requires attention not only to causes but to *constituting grounds* that (he says) make a kind the kind that it is. The second is to refute what he claims to be the "consensus view" that human collective agencies, both formal institutions and informally coordinated groups, are constituted exclusively by individual people. This is the thesis of ontological individualism (OI). As Epstein says, OI is distinct from methodological individualism (MI), the view that group behaviour is best modelled by modelling the behaviours of its members. Neither sort of individualism implies the other. Epstein is satisfied that social scientists generally understand that MI is sometimes a useful policy and sometimes isn't. But he claims that social science is hobbled by a mistaken general commitment to OI.

OI is compatible with either reduction of social elements to individual ones, or of the supervenience of the social on the individual. If reductionism is denied, then OI is the specific application of a supervenience claim to the relationship between social and individual scales of ontology. Epstein insists on a metaphysical interpretation of OI. Guala (2022) resists this insistence, but in doing so makes OI equivalent to contingent supervenience of the kind discussed in Sect. 2. Epstein's view that commitment to OI hobbles social science in general depends on his metaphysical analysis of that alleged commitment. Otherwise he would be claiming, implausibly, that social scientists as a community have failed to spot some empirical facts that he has noticed, or, at least, have estimated unrepresentative samples.

Epstein's direct targets of criticism are in fact philosophers who have failed to give best advice to social scientists because they have in various ways misunderstood the general nature of the constitutive grounding relationship between facts about individuals and facts about groups and institutions, in particular confusing grounds – the "metaphysical" reasons that a group (or, more generally, social fact of a certain kind) exists in the world – with 'anchors' – the typically diffuse, processes that establish frame principles for grounds. These frame principles are specifications of sets of possible worlds that determine which elements of descriptions of grounding conditions can and can't vary. It is partly because of these modal restrictions that frame principles and grounding relationships are regarded by Epstein as metaphysical. The other reason, which also applies to anchors, is that they rest on constitutive relationships that are matters of fact, but not derived from causal generalisations. He therefore devotes much of his book to producing a theoretical specification of the basis of social facts that strictly separates anchors from grounds, and to contrasting and

defending this theory against alternatives from the (exclusively) philosophical literature. I have no opinion on the extent to which he succeeds in this intermural contest with other analytic metaphysicians. I take Epstein's account as the naturalist's foil because of his objectives with respect to reforming social science.

My main critical targets in engaging with Epstein are the constitution/causation distinction and the idea of metaphysical grounding. Naturalistic metaphysics does not admit that there is any basis for regarding these ideas, which do not emerge from scientific observation or inference, as tracking any facts about objective reality. Of course, it is a fact that some people, when thinking about the world, apply these ideas. But for the naturalist this is an anthropological fact, and not metaphysically significant. Epstein's metaphysical interpretation of constitutive grounding rests on a general intuition that some structural facts are *more fundamental* than others. I do not think it is possible to explicate this notion of fundamentality from outside of a non-naturalistic metaphysical framework. To the naturalist, the idea that the world could be objectively structured in this way makes no sense. So I cannot argue with Epstein by promoting an alternative account of fundamentality. What I will do instead is consider the most closely analogous idea that is compatible with naturalism. This is a variable that turns up as a conditioning variable in relatively many successful structural models.

With respect to Epstein's effort to displace OI, the issues between the analytic and the naturalistic metaphysician are subtle. (That is exactly why using Epstein as a foil is productive strategy for illuminating naturalism.) The naturalist need not deny that OI has exerted influence on social science *as a metaphysical idea*. The reformist social ontologists mentioned in Sect. 1 take themselves to be resisting pernicious influences on Western societies of OI as an element of traditional folk metaphysics.⁸ That is at least part of the reason that they characterise their project as partly metaphysical. OI is a specific manifestation, in the social domain, of atomism, the idea that structures reduce without residue to interactions of basic constitutive elements that carry their properties, including their causal powers, across contexts.

Western culture has often interpreted science atomistically, as it has everything else. However, atomism as a metaphysical program for science never had genuine empirical support, even during the period in which classical mechanics was widely taken as the template for successful science. Newton's own mechanics relied on pervasive gravitational fields that resisted atomistic interpretation. But opponents of atomism have often failed to gain persuasive traction by offering holism – that is, general scepticism about stable and observer-independent system boundaries – as the alternative. Holism can seem particularly attractive in application to highly complex systems, such as human societies. But it directly undermines the division of scien-

⁸ Examples of such pernicious influences are many. Folk individualism undermines efforts to mitigate structural oppression and wasted welfare in most if not all European and European-settler countries, and it feeds and is amplified by ideologues and materially comfortable people who prefer to assign responsibility for poverty to poor individuals. It fuels the widespread and tenacious fallacy that reform of institutions should prioritise, and indeed can be fulfilled by, identifying and weeding out 'bad apples'. It encourages the widespread recent policy focus on 'nudging' individuals into more welfare-promoting behaviour, which is stringently criticised by Chater and Loewenstein (2023) for distracting governments from responsibility for structural interventions.

tific labour and isolation of mechanisms and relatively closed causal networks on which both engineering control and scientific explanation depend. This is the sort of circumstance in which analytic metaphysicians are inspired to come to the rescue: folk metaphysics generates a dichotomy of crude alternatives that are both unsatisfactory. The philosopher then seeks a refined synthesis. In the case of the conflict between crude atomism and crude holism, the idea of supervenience as metaphysical construction seems to promise the best of both worlds. Epstein rejects this dialectic, rightly in my view. Supervenience still allows OI, and OI reflects the atomistic intuition that Epstein joins the reformist social ontologists in resisting.

Much of Epstein's book is preoccupied with trying to convince the reader that his theory of grounds and anchors provides a successful technical alternative to OI, which does not slide into holism, where previous attempts have failed. The details of argument follow the standard method of analytic philosophy: generating counterexamples to previous theories and parrying apparent counterexamples to the newly proposed one. Again, this is not a debate I aim to enter into; the aim is to generally characterise and ultimately reject an analytic metaphysics for social science, not to identify a preferred analysis. The naturalist's ambition is not to find a preferred refinement of folk metaphysics; it is to promote the entire displacement of all *a priori* metaphysics, folk and analytic versions alike, by generalising results from science.

Therefore, instead of criticising Epstein's account using his own tools of analysis, I will consider it from the 'grander' perspective of the history of ideas. Let us imagine for the sake of argument that Epstein's account succeeded in the narrow sense of winning all the technical debates decisively. That is, suppose that every analytic philosopher who carefully examined Epstein's account concluded that it was superior to any alternative any of them could think of. How might we then summarise the extended intellectual adventure that had ended with Epstein's triumph? First, we might note that a major share of the success of his solution would reside in his recognition that anchoring conditions for social kinds are (effectively) limitlessly heterogeneous and contingent: to be able to state *a priori* constraints on how human practice and thought can anchor new dependence relations between social facts and non-social facts we would need to be able to identify a complete set of limits on collective human capabilities. This is simply to acknowledge the complexity of social processes, the consideration that makes holism a philosophical temptation – but a threat to the possibility of successful social science – in the first place. If Epstein's account works, this is in considerable part because he does not try to deny social complexity. But then the *rest* of the credit for his hypothetical success, the analytically *hard* part, would have to lie in the account of grounding, for this is what is supposed to achieve specific, orderly, stable bases for identification of isolable regularities. This diagnosis is consistent with the burden of effort in Epstein's book: most of it is about developing and defending his technical account of social grounding and anchoring. These relations are held to be metaphysical: facts about structures of possible worlds that (contingently) apply to specific entities in the actual world. That is why Epstein supposes that he can do more than provide the folk with a less politically damaging metaphysical story, but

can in addition help social scientists achieve firmer and clearer results.⁹ In particular, the metaphysician's toolkit is viewed by Epstein as essential technology for showing how social scientists can arrive at more robust identification of regularities underlying the social facts that interest them.

This brings us to the core critical question for my present purpose: *why* should we be convinced in the first place that practical problems of social scientists seeking regularities they can project out of sample should have a *metaphysical* solution? Recall how the path to that view is set up. We start with a metaphysical interpretation of OI: every set of structurally connected social facts constitutively depends (exhaustively) on facts about individual people. This dependence relation is not causal (since if it were it would be empirically contingent). Nor does dependence refer to contingent structural composition. 'Depends' here means nothing independently of an intuition to the effect that some facts are 'more fundamental' than other facts, so some important *scientific* generalisations are taken to state constitutive relations rather than causal regularities or mechanisms.

Intuitions about hierarchies of inferential dependence are arguably inevitable for beings that need to continuously identify links between actions they could take that might be successful, and varying conditions in which these actions could be taken. Such beings are bound to (at least behaviourally) treat the first domain of variation as depending on the second domain. If the beings have very limited flexibility with respect to actions – suppose they are starfish, lacking central nervous systems – then natural selection must have pre-solved this problem for them, in the sense of hard-wiring the practical dependencies that govern their choices. A starfish need not learn to treat sizes of encountered objects as more 'fundamental' facts than facts about whether objects have chemical composition that allow them to be ingested; the starfish automatically behaves in accordance with that practical dependence. But if an animal has enough capacity to run internal simulations of counterfactual situations – that is, if it can partially decouple behavioural control from direct response to sensory affordances – it will need to *model* actual and possible experiences into practical equivalence classes, and facts about some equivalence classes as depending on facts about others. In that sense, some domains of perception and potential action will be modelled as more 'fundamental' than others.

So far there is no basis for thinking that these models incorporating relations of inferential dependence involve any metaphysical dimension. They are merely instruments for optimising expected utility – 'as if' fundamentality, so to speak. But now consider a being who aims to do science, by which we mean: she aims to discover propositions that are true independently of the instrumental value to her of modelling them. In principle – that is, abstracting away from limitations that natural selection might have built into her information-processing hardware – she has (at least) two options here. She could continue to try to identify some domains as more 'fundamental' than others *in general*, but now with respect to truth-conduciveness rather than utility-conduciveness. Alternatively, she could operate on the assumption that where

⁹ These two goals are complementary, on Epstein's picture: the folk can refine their picture of society both by deploying better metaphysics *and* by enjoying the epistemic fruits of a social science that is aligned with the metaphysics in question, rather than working against its grain.

truth is concerned, no domains are generally more fundamental than others (though estimations of the probability of any given fact statistically depends on estimates of the probabilities of various other facts). If she goes the second way, she is on the road to the radical naturalist's picture of reality. So I will postpone consideration of this option until the next section, when I consider naturalistic metaphysics for social science. For now we'll follow the implications of the first approach, which set the agent on the road to analytic metaphysics.

The being who identifies some domains of fact as more fundamental *in general* than others will build models with what amount, effectively, to axioms that reflect this hierarchy. And she'll need a logic that, if followed, will generate truth-preserving inferences from sets of premises that include the axioms. Of course the logic must do more than that if she's under time pressure to derive conclusions, or is more interested in some truths than in others. What she in fact needs is both a logic and a program, and these can't be selected independently. We thus find ourselves in the well explored terrain of classical (i.e., pre-connectionist) computer science and general-purpose AI. This tells us that the scientist under development has a lot of options.

This agent has a metaphysical model, which is equivalent to her axioms about relative fundamentality of domains of facts plus her chosen logic. As Smith (1996) has demonstrated in rich detail, if she is not to arbitrarily foreclose her scientific modelling capacities she will need to construct objects – a catalogue of co-occurring properties that share common causes and effects and depend on the same facts from more fundamental domains – posterior to having selected her logic. Some of these objects can be tagged as types and others as tokens of those types, where such tagging will depend on observations plus restrictions derived from her metaphysics. If she discovers increasing complexity in the network of all these relations as she goes along, then she will need to be able to revise this ontological catalogue, potentially without limit. This is because she could not anticipate in advance all the co-occurring properties meeting the metaphysical criteria for object-hood that she might observe. But she will be constrained in such revisions by her metaphysics – unless her basic operating program allows axioms or rules of inference to be revised, perhaps when she receives a signal that others are consistently doing better science than she is. A program could be built to allow for such deep resets.

Epstein's social scientists, prior to their being cajoled by him into shared metaphysical awareness, are networked truth-seekers who, having encountered in their relationships with one another a new domain of high complexity, have evolved varying ontological catalogues, and this variation interferes with their capacities to settle on common revisions. Epstein – remarkably – never explains what he thinks actually breaks down *in modelling* when social scientists are in ontological misalignment. In his most extended examples from outside of his book (Epstein, 2008, 2011, 2014), they make incorrect predictions when they project the wrong reference classes, but we aren't shown how their lack of metaphysical attentiveness leads to structural identification problems that fail to allow for the risk of the mistakes.¹⁰ So I will

¹⁰ For example, in the extended example in Epstein (2014), his imagined fiscal policy makers cause a crisis when they are surprised by the retrospective reclassification of a tropical storm as a hurricane. But what was their economic model? Why did it not include the possibility of issuing bonds to offset the reclassifi-

conjecture an account that seems consistent with his project. Social scientists working on closely related problems need to try to reach agreement on both the ontological structures they should model, and on measurement protocols for identifying these models with observable data. If their working ontologies aren't clearly and consensually grounded, then they have too many degrees of freedom when designing and assessing empirical research programs. They can't distinguish between cases in which their models incorporate inconsistent ontologies from cases in which they're measuring inconsistent grounding conditions as proxies. But they are presumed to be able to agree on a common metaphysics. Evidently there is noise in their reliability in applying this metaphysics, since they have failed to keep anchors and grounds well distinguished. But they can consult experts, philosophers, to help them shake off these confusions. *Given* the right partitioning between anchors and grounds, they will be able to identify their points of disagreement. Sometimes they might turn out to disagree about anchoring conditions, which will generate inconsistent ontologies. But anchoring conditions can be empirically assessed – this is the job for social theory or economic theory. (If bodies of social and economic theory use disjoint or incommensurable ontologies, disciplines will agree to study different phenomena. If they study elements identified by the same anchoring conditions, they can achieve interdisciplinarity or disciplinary complementarity, depending on whether their methodologies converge.) On other occasions their disagreement will be limited to the causal structure among the items in their shared social ontology. Then their problem is just the main business of everyday social science: identifying models through statistical tests, running experiments and surveys, estimating correlations and regression coefficients, conjecturing and testing directed acyclic graphs, and so on. This hardly guarantees ultimate convergence on uniquely best models – even disputes that are only methodological can be intractable. But at least everyone can agree on what everyone is trying to measure and model.

Thus we have a picture of a scientific method that requires metaphysics and involves identification of metaphysical grounding relationships. It is the picture we would implement if we set out to design a scientist, and a networked community of scientists, using the principles of classical AI. This should not be surprising, on reflection. The crucial metaphysical relationships are specified logically, and classical AI systems are essentially logical inference machines. If metaphysics can really guide scientists in the way that Epstein supposes, then we should be able to imagine how this would work in an engineered inference engine that shares the basic technology of the analytic philosopher.

But this should be worrying. Classical AI systems *only* work in tightly constrained worlds. That is why they did not deliver general-purpose intelligence, are useful only as narrow expert systems (e.g., operating an assembly line, diagnosing breast cancer), and have been displaced by statistical deep learning architectures (neural networks or neural network simulations) for 'open' problem spaces like language processing and ... axiom-free scientific discovery. There are several ways in which classical AI

cation risk, given that the imagined law on damage compensation that gets them into trouble is imagined as having been known to them? In general: how do ontological uncertainties translate into parametric structural restrictions? Epstein never addresses this.

systems don't work in 'wild' inference environments, but one of these is that they either get trapped in narrow ranges of discovery space *or* (if their programs allow for controlled revision of axioms), they fall into endless loops of frame revision when confronted with novel kinds of data. That is, they are crippled by frame problems (Pylyshyn, 1987; Ford & Pylyshyn, 1996).

Frame problems loosely defined are partly attributable to the permissiveness of logic: too many logically consistent models are compatible with any finite body of data. But these can often be tempered through artful design hacks (such as cuts when logic programs are undone by the louche inferential tolerance that comes from negation operators; see Lloyd, 1984, pp. 56–60). The frame problem 'proper', the deadly one for a system trying to solve a problem as unconstrained as 'discover scientific truth', arises because fundamentality assumptions are both powerfully restrictive and arbitrary unless one can base them on extensive knowledge of the structure of the domain the computer is supposed to model. There is only one way to acquire the extensive knowledge in question: by scientifically studying the domain. If we want to successfully model the kinds of dependence judgments that grounding of Epstein's kind requires – where by 'successfully' we mean 'find the fundamentality assumptions consistent with optimising discovery of truths that license out-of-sample predictions' – then such modelling would have to *follow* the science. It cannot serve as the foundation for populating the ontologies we use for fixing, structurally relating, and measuring variables and parameters.

I have approached this point in the way I have, via the roundabout strategy of imagining an effort to design a science machine, in order to avoid the pure standoff approach of refusing to grant Epstein's unargued assumption that it makes sense to regard some bodies of fact as more fundamental with respect to the constitution of ontological elements than others. Empiricists sometimes write as if the idea of metaphysical grounding on which Epstein depends is *unintelligible* unless we imagine access to truth achieved by unaided intuition and logic. In fact we could, as I have argued, operationalise a metaphysical model, which could be distinguished from first-order scientific models and play an independent role in constructing ontology, in a working scientific discovery program. But we could have no basis for believing that the metaphysics in question describes the actual structure of reality unless *something* prior to science gives us confidence in the relative fundamentality judgments that determine which sorts of facts ground which sorts of other facts.

Epstein says nothing in his book about possible sources of fundamentality judgments. I infer from this that he thinks they reflect natural, default beliefs. Does he suppose that these beliefs constitute knowledge – that is, that metaphysical facts include facts about which strata of existence are more and less fundamental? That is one possibility, in which case we have an instance of a view with very old and deep roots in Western philosophy, according to which knowledge of very general metaphysical structure precedes scientific knowledge of empirical contingencies and is indeed a condition of its possibility. Alternatively, someone following Epstein's program might take a Kantian or quasi-Kantian stance and hold that some widely shared *conceptions* of relative fundamentality are necessary foundations for orderly empirical inquiry and therefore cannot themselves be usefully interrogated. Kant thought of this as denying the possibility of metaphysics. But metaphysical *theory*, in the

form of an analytic program, could survive the demise of belief in reliable intuitive insight into structures of reality. ‘Metaphysics’ could be identified with ‘metaphysical method’. Science could be understood as a methodologically well-ordered *bet* on the epistemic value of (shared) brute fundamentality judgments.

We must be explicit that if fundamentality judgments were treated as unproblematic, this would by no means imply a view that accurate judgments about particular grounding relationships are easy to come by. Here Epstein puts his cards clearly on the table, and we need not speculate. His reason why we need to do fresh metaphysical work to improve social science is that grounding relationships are generally much more heterogenous, complicated, and hard to identify in social science – and, perhaps, more generally in the ‘sciences of the artificial’ (Simon, 1969) – than in ‘natural’ science. “If there is any moral to this book,” Epstein says, “this may be it: facts in the social sciences are grounded differently than those in the natural sciences. Compared to the social sciences, the ontology of natural science is a walk in the park” (Epstein, 2015, pp. 163–164). This is followed by a remarkable excursion in how and why natural scientists can supposedly get away with being casual about identifying grounds in a way that social scientists cannot.

I will return to this discussion, which is illuminating with respect to the deep rupture between the assumptions of analytic and scientific metaphysicians, later. For the moment, however, I want to emphasise a basic point. Epstein believes that a general, albeit inexact, metaphysical structure of natural reality is common currency. Perhaps it’s taken as just obvious that facts about solid, everyday natural objects are more fundamental than facts about fields and forces, or facts about biological kinds such as cells and synapses. Then we can ground facts about cells in facts about chemicals, and facts about chemicals in facts about molecules, and facts about molecules in facts about observed stability in solid, everyday natural objects. (Epstein’s example of an easily grounded natural object is “a rock”.) But when we set out to identify grounds for facts about a social entity such as Starbucks (one of Epstein’s examples) we find that we must attend to facts about people (e.g., Starbucks employees) but also facts about coffee beans and espresso machines. (This reflects, according to Epstein, the failure of OI.¹¹) So multiple, different fundamentality judgments are required to be able to know which kinds of facts can ground the facts about Starbucks.

Many of Epstein’s extended examples of social entities are of formal institutions – the US Supreme Court, a parliament, an intramural college basketball team, a corporation (e.g. Starbucks) – in which anchoring conditions are explicit rules. These rules almost inevitably encode relevant fundamentality judgments. The rules for constituting the Supreme Court make clear that it will partly be grounded in members (judges) who are individual people. So facts about individual people can be inferred to be regarded as more fundamental than facts about the Court. The anchoring conditions for the basketball team include rules requiring a sufficient number of basketball players, again individuals, to show up for scheduled games. So again we can infer that

¹¹ As Guala (2022) argues, the easy-going naturalist who understands supervenience non-metaphysically need not accept this. She can grant that the material beans and machines are important to Starbucks accountants, but deny that economists, who are only interested in relatively abstract features of Starbucks’ business model and market interactions, need to pay any attention to these material tokens and their properties.

facts about players are treated as more fundamental than facts about teams. Another grounding element for the basketball team, according to Epstein, is the occurrence of a specified “initiating event”, completion of a form by a manager of the (then) prospective team. The anchoring conditions specify this event, so we can infer that the event of submitting a form is a more fundamental kind of event than the extended event in the history of the league that is the existence of a specific team. As Epstein makes explicit, laws, and by obvious implication regulations and explicit rules generally, can be frame principles.

Most social science is not about social entities or events that are constructed by explicit rules. That so many of Epstein’s main examples feature such social types may obscure the extent to which fundamentality judgments often lack any persuasive justification. Note that laws and regulations function in a way that would effectively create closed worlds for the classical AI inference engine we imagined as operationalising Epstein’s conception of science. There are AI programs, expert systems, that can substitute for legal clerks in identifying precedent cases. Thanks to the explicitness of legal definitions *plus* the rigidity of their applications, they set frame principles that have (local) modal force. But the source of this is the nature of law, not the metaphysical structure of reality. If we are thought to enter the domain of metaphysics merely by virtue of using reference to possible worlds to spell out tolerance conditions on variation within a rule people wrote down and promulgated, then this ‘metaphysics’ is merely a methodological residue of what medieval Western philosophers historically took metaphysics to be about, when they supposed that all of reality was governed by a lawmaker.

Epstein recognises that anchoring conditions for informal human groups are typically more complicated. In a footnote, in the context of discussing how group actions are grounded, Epstein says, concerning groups that aren’t anchored by explicit rules, “constraints on action are anchored in more complex ways. Family structures, for instance, involve membership conditions and hierarchies of power. These are anchored by historical tokens, practices, environmental facts, and more” (Epstein, 2015, p. 235). In these cases, reliance on intuitions about relative fundamentality looks much more questionable. How do we know that facts about historical practices are more fundamental than facts about families, instead of the other way around? Epstein might answer that we simply see the relevant fundamentality judgments expressed in social scientists’ explanatory projects: the anthropologist explains the characteristics of a particular family by reference to historical practice. But this doesn’t necessarily reflect any judgment about relative fundamentality; it more plausibly just reflects the fact that we generally explain present conditions by reference to past ones rather than the other way around.

Sugden (2016) complains that, in his book, Epstein never works through a case of a model from actual social science that he shows to suffer from a problem that his conceptual treatment would fix. He reviews Virchow’s 19th -century effort to reduce the entire biology of the organism to cytology, and argues that the obvious ultimate failure of this project is analogous to prevailing programmes in social sciences. But the question of whether contemporary social scientific models are infected with such atomism must be addressed by direct reference to such models. All of Epstein’s other detailed examples concern either explicitly rule-governed entities or quotidian social

objects and everyday conceptions of them. As noted in Sect. 2, Epstein in other work has considered examples that are at least generically drawn from scientific literature, but only his (2011) discussion of agent-based simulation engages with *details* of actual models, and his argument there, as I pointed out, does not depend on metaphysics.¹² Indeed, a general point on which Epstein is surprisingly *not* clear is what he means by ‘model’.

This disengagement from social science cases might charitably be explained as reflecting a deep bias that separates the analytic metaphysician from the naturalist. Perhaps the former are more likely to view science as an extension of everyday epistemic practices, where the naturalist conceives science as rendered special and sui generis by its institutional structure. Outside the typically tight constraints imposed by a specific family or sequence of scientific models, it is easy to generate lists of social kinds in which judgments about plausible explanations of constitution aren’t likely to be controversial, especially in the absence of any situational details. But in such settings it is equally easy to generate cases that flip the underlying fundamentality judgments, just by adding special context. Suppose I want to explain why there are no conservatives on a particular intermural basketball team, and I have among my data that the manager who registered the team called it the Social Justice Warriors? Now it is facts about the nature of teams, including sports teams, as anchors of generalised solidaristic identification, which explain facts about some individuals (those that want to play basketball but won’t play on *this* team). The point isn’t that the case would be difficult for a social scientist; the first hypothesis to be tested would be obvious. The point is rather that the issue appears to have nothing to do with any general metaphysical principle to the effect that facts about individuals are more fundamental than facts about teams when it comes to deciding which kinds of considerations should ground which others in explanations.

In summary, Epstein provides us with cases in which people make fundamentality judgments that could provide a basis for analysis in terms of grounding, which plausibly reflect folk metaphysical conceptions. But these are relevant to improving social science only if it is agreed that scientists should include folk metaphysical ideas in their inferential priors; and the naturalist does not agree with this. Epstein also provides examples where it seems reasonable to specify grounds and anchors for social entities that are governed by explicit rules. In these cases, however, it is unclear why any appeal to metaphysics is warranted. He does not provide an instance where a model, explanation, or hypothesis from social science is improved by specifying grounds and anchors for types. In some work that preceded his book, however, he has dealt in more detail with actual social science. The next section considers one of those cases.

¹² Epstein (2014), on microfoundations for macroeconomics, obviously addresses a real scientific issue. But even there, Epstein does not consider relationships between actual models, which are the working objects of economists. In any event, I will revisit this argument of Epstein’s specifically below. In other work where Epstein *does* discuss real cases, most clearly in Epstein (2011), the elaborate machinery we are asked to understand in *The Ant Trap* is not used, and the argument can be reinterpreted naturalistically without much strain.

4 An example: economic microfoundations

In the opening pages of his book, Epstein claims that there is a “paradox” afflicting social sciences. Social scientists, he says, lately have been blessed with enormously more data than they used to have, thanks to digital tracking of people and their thoughts and actions. Yet, he alleges “the social sciences are hardly budging” with respect to their record of success in explaining poverty, educational success, or relative financial-system stability.¹³ The example of vivid failure he provides is the current favourite among critics who are disgruntled about economists: the fact that macroeconomic modelers in 2008 didn’t predict the financial crisis that erupted that year. But Epstein never uses the example to illustrate metaphysical confusion of the kind he claims to provide resources for fixing. The case is *prima facie* unpromising, because we can fully understand the apparent failure by reference to the policy contexts in which the relevant macroeconomic models had been developed. Macroeconomists did not include the aggregate health of balance sheets of financial institutions in their models because their main implicit policy clients, central banks, had no direct influence on these balance sheets. Two facts speak directly against any diagnosis of what went wrong in terms of mistaken ontological assumptions about ‘the macroeconomic domain’. First, economists in general were *not* ignorant in 2008 about the transmission mechanism, from unsustainable home mortgages through short-term corporate credit markets to the real economy, that brought about the crisis. It had been identified, years before the crisis, in a well-circulated and widely cited paper by famous and highly prestigious economists (Holmström & Tirole, 1997).¹⁴ Second, when central banks *did* assume influence over the values of corporate financial assets by adopting the innovation of quantitative easing, macroeconomists duly added these values as elements of standard models.

The relationship between macroeconomics and microeconomics is a rare instance from social science on which there is extensive literature on what looks like the classic kind of scenario for application of ideas about ontological supervenience. The majority of macroeconomic theorists avow that their models should have ‘microfoundations’ (Blanchard, 2016). But there is significant expert dissent about this (Hoover, 1988; Janssen, 1993; Duarte & Lima, 2012; King, 2012). Furthermore, there are multiple specific formulations of microfoundations, and varying preferences among them have significant implications for economists’ opinions about monetary and fiscal policy options (Ragot, 2012). So this looks at first glance like a scene of ontological disagreement that crucially affects model selection and policy applications. Perhaps economists should be calling in analytic metaphysicians for help?

Epstein (2014) frames the problem in part by asking what distinguishes “micro-economic properties” from “macroeconomic properties”. He notes that this is challenging because economists often use the same families of models at what intuitively

¹³ That an expert on applied logic thinks that ‘paradox’ is the right word to use about this alleged state of affairs is surprising.

¹⁴ What economists, along with almost everyone else, were not alert to before 2008 was the extent to which regulation of the financial services industries in both the US and Europe had become dysfunctional. Complicating this was that regulatory failures on the two sides of the Atlantic were importantly different in both cause and effect, and they compounded one another.

look like *both* micro and macro scales. Economists treat any entity that adjusts its behaviour in response to incentives as an agent (Ross, 2014). Thus individual people are modelled as agents, but so for many purposes are firms, which might operate on a global scale, and national governments. “This creates a conundrum,” Epstein says, “as to whether the properties of nations ought to be included in the macroeconomic property set ... or [the] microeconomic property set” (Epstein, 2014, p. 9). He goes on to dissolve the “conundrum” in just the way a typical economist would: “If all it takes for an entity to be microeconomic is that somewhere it is modelled as an economic agent, then microeconomic and macroeconomic properties are not likely to be distinct at all ... Equally, if we understand microeconomics simply as a set of methods, which apply to objects at various levels depending on our explanatory interests, then any supervenience claim about microeconomics is empty” (*Ibid.*, pp. 9–10). This diagnosis feeds into Epstein’s main conclusion that indeed macroeconomics does not supervene on microeconomics. That is as far as his explicit treatment of the issue goes in his paper on that subject. But in the wake of his book, we might see how we could still make sense of the widespread commitment to microfoundations by wheeling in the machinery of frames, anchors, and grounding of less metaphysically fundamental properties in more fundamental properties.

Though economists are at least as inclined to atomism when they make casual remarks as any other inheritors of Western folk metaphysical culture,¹⁵ leading defenders of the importance of microfoundations for macroeconomics almost never appeal to general ontological principles such as reduction or supervenience.¹⁶ The view that macroeconomics should have microfoundations arises from the concern that many macroeconomic policy proposals require some individuals, firms, and households to remain ignorant of the intended effects of the policies in question; otherwise they are incentivised to choose actions that would undermine or entirely undo the intended policy outcomes. There are many subtleties to these debates. For example, some policies might not be able to shift long-run macroeconomic equilibria due to microfoundational strategy adaptation, but can redistribute *timing* of investments or other corporate strategy elements in such a way to avoid destabilising coordination of decisions. This is the basic motivation for the much-discussed central bank innovation of quantitative easing, which is generally thought to have prevented the 2008 crisis from triggering a global depression, and to have saved the Euro from collapse in 2012.

The problem with integrating microeconomic and macroeconomic models arises not from the fact that relationships between their ontological frames are obscure, but because they model processes on radically different timescales that causally interact. This kind of problem, which is very common in sciences that study dynamic processes, poses a direct challenge to Epstein’s reliance on shared fundamentality judgments.

¹⁵ It remains true that the majority of prominent non-Western economists were educated at Western universities, but this is changing. Will rhetorical adversions to OI, which have never aligned very well with actual economic theory (Ross, 2014), begin to fade away among a coming generation of economists trained at Chinese universities?

¹⁶ Hoover (2009) is among the few economists who has explicitly maintained that macroeconomics supervenes on microeconomics. But he argues that the microfoundations project fails despite this.

Some proposed solutions to scale integration rely on the claim that stability in economies must be generated from the micro scale, because that is the scale at which expectations can be at least approximately rational, in the sense that utility and production functions can be assigned to agents (Lucas, 1976; Kydland & Prescott, 1977; Long & Plosser, 1983). This is the domain where microfoundations are most strongly emphasised. A standard theoretical manoeuvre is not to study actual agents at the micro-scale, but to model the whole economy as a single idealised (infinitely lived) agent with fully rational expectations. Many economists, including me, find this approach deeply unsatisfying. The problem here truly is ontological, but not metaphysical. It is that the representative agent is not even an idealisation of anything that has a counterpart in reality outside the models. In consequence, such models tend to be untestable (Romer, 2016). Increasingly many macroeconomic models with microfoundations consequently use heterogeneous agents who represent common expected response patterns – for example, in a model of impacts of international trade flows, different representatives for producers of goods that expect terms of trade to be altered in their favour by some exogenous shock and producers that expect a deterioration in terms of trade. But other theorists – so-called ‘post-Keynesians’ – defend the opposite perspective, according to which only expert management at the policy level can coordinate aggregate-scale responses so as to dampen random and in-principle unpredictable fluctuations (Cencini, 2005; Taylor, 2010). It is utterly unlikely that a day will come when one of these modelling methodologies drives out the other – each approach sheds light on interesting questions, and occasionally each sheds *different* light on the *same* question. In practice economists typically adopt a pragmatic approach of assuming a ‘mild’ rational expectations constraint on all macroeconomic policy proposals regardless of the preferred style of modelling them: “If the economist builds a theory or model in which [some of] the agents fail to do something that it is in their interests to do, then the economist must justify why they did not do it” (Ragot, 2012, pp. 187–188). Satisfactory justifications sometimes invoke market structures that institutionally constrain choices, and sometimes special information constraints faced by individuals, households, or firms. Thus some models implicitly treat microeconomic relationships as more fundamental than macroeconomic ones, while others do the opposite (e.g. Taylor, 2010).

In discussing the microfoundations issue, I have deliberately avoided following Epstein’s lead in referring to ‘microeconomic properties’ and ‘macroeconomic properties’. As a practicing economist, I can make little sense of these ideas. If asked to produce lists of properties, I would not know how to begin or what principles should be used. The distinction between macroeconomic models and microeconomic models is often clear, but sometimes not. We might say as a first approximation that microeconomic models feature designated agents responding to changes in incentives, while macroeconomic models do not. As Epstein (2014) worries, this seems to have the consequence that all models involving representative agents are microeconomic models. But that is not the view of the matter that economists take. A key difference is buried in the word ‘designated’. This does *not* mean, as someone framing the issue in Epstein’s terms might think, that the agents in the microeconomic models must be anchored in actual, named, people or firms or governments (though from time to time they are). Rather, it means that they can be assigned *structural* response

functions that reflect information that is available specifically to them (or available at costs specific to them). By contrast, in models with heterogeneous representative agents all representatives typically know everything that the modeler does. This basis for distinction is admittedly strained for models with single representative agents. Here there is some justice in Epstein's thinking that the models amount to replacing macroeconomics by microeconomics. As I noted, however, there is a growing turn against models of this kind, and they are steadily decreasing in frequency of use. I predict that future historians of economics will find it necessary to devote energy to explaining why Laureate-class theorists thought that such a bizarre modelling practice might make sense.

It is ontology – the causal structure of the world – that explains the need for both microeconomic and macroeconomic models, and it is a scale feature of that world, the existence of policy-relevant contexts in which no choice by any specified agent or agents can make any marginal difference to outcomes, that determines when macroeconomic models should be used. But economic practice reflects no *general* judgment about what kinds of facts in economics should be expected to serve as grounds for their special concepts, because they maintain no general doctrine about what is 'fundamental' across all economic model designs. This does not imply that concepts used in models are unstable across applications. Economic models are mathematical, so specifications are precise, and so are identification restrictions that tell economists what they need to measure to estimate a given model in a specific set of data. The point, rather, is that what plays the role Epstein seeks from the metaphysical concept of grounding is specific to modelling assumptions – there are no *general* frame principles that apply across all models, but models can be cleanly cross-evaluated by reference to their formal structures. Economists avoid general fundamentality judgments, because they doubt that the economic world is structured that way. In effect, good advice to young economists learning to build models could be glossed as: avoid metaphysical assumptions.¹⁷

Naturalist metaphysicians think that the absence of general foundational relations in economics is characteristic of the world in general.

5 Scientific metaphysics and social science

Consider again Lohse's (2017) remark that radical naturalisation of metaphysics is "a sensible project" in application to 'natural' sciences, but "hardly achievable" where social sciences and their ontological foundations are concerned. In one sharply restricted sense, Epstein has a similar view. He does not think that analytic metaphysics fails to apply to physical objects. However, as cited previously he says that a "main moral" of his 2015 book is that "facts in the social sciences are grounded differently than those in the natural sciences. Compared to the social sciences, the ontology of natural science is a walk in the park" (Epstein, 2015, pp. 163–164). The reason

¹⁷ I wouldn't put it that way to an economics graduate student, however, because they likely would have no idea what I meant. I would say: avoid supposing that you have a grand model of how a whole economy works, or of how any economic mechanism *always* works.

he believes this is that “objects treated in the natural sciences, like rocks and planets and cells” are typically “nearly intrinsically individuated” (p. 166). What he means here is indicated by analysis of his example of an “ordinary object we might treat in natural science ... : a rock” (p. 164). What makes the rock the rock it is, he argues, is constitution by “particles” that are contiguous and relatively strongly bonded to one another, plus absence of essential parts (one can chip off any piece of the rock without changing its individuation).¹⁸

My burden of argument here is to show how scientific metaphysics applies to social sciences by showing why the asymmetry claimed by Lohse and Epstein does not hold. I will do this by chiselling, as it were, on both ends of the stick, explaining why Epstein is wrong about the metaphysics of physics and physical objects, *and* arguing that he is wrong about the kind of metaphysics that supports *successful* social science.

5.1 Physics

Epstein’s contention that objects in the ‘natural’ sciences are “nearly intrinsically individuated” is an expression of the atomistic tradition from which his philosophy of *social* science is aimed at escaping. That formerly prevailing grand metaphysic of the natural world was a core element of the first analytic philosophy (Russell, 1911, 1918, 1927), of which contemporary analytic metaphysics is in many ways a true descendant, following the interlude of logical empiricist scepticism about all metaphysics. The atomistic vision is that self-subsistent objects shift through fields of relations, carrying their essential properties – in Epstein’s terms, the facts that ground them – around with them as they shift relative positions. Change is driven by their collisions.

The core sources of contemporary scientific metaphysics (Ladyman & Ross, 2007; French, 2014; see also Lewis, 2016, and Ney, 2021) reject atomism as an interpretation of physical theory, and argue that it must consequently fail as a general metaphysic.

Insofar as they are object-like, the particles of quantum theory are not basic elements in physical ontology. In being subject to entanglement, they are shorn of the most basic characteristic of a metaphysical atom, individuality across measurement contexts (French, 1989, 1998, 2014; French & Redhead, 1988; Ladyman & Ross, 2007, pp. 132–145). In contemporary physical theory, particles are excitations of fields; so under the analytic framework followed by Epstein, fields are fundamental and particles are reference points for theoretically and operationally crucial

¹⁸ Epstein says that the rock is only *nearly* intrinsically individuated because he recognises that boundaries of natural objects are typically fuzzy. Thus we could change the temperature of the rock, without heating or cooling anything, by deciding on theoretical grounds that we should adjust its boundary specification. Here again the stringently analytic metaphysician and the scientific metaphysician agree with one another, as against the relaxed contemporary majority who temper their atomism by invoking supervenience. As argued in Sect. 1, Epstein and the scientific metaphysician agree on things like this *because* they both take metaphysics seriously. I will add, however, that the mathematics involved in physically characterising an object like a naturally occurring rock are very complex – more difficult, in my view, than any part of quantum theory.

measurement coordinations. In traditional metaphysical terms, quantum fields are much more like processes than objects. Furthermore, they characterize the *large-scale* structure of reality; the idea that only micro-scale physical phenomena are ontologically bizarre to Western folk metaphysics is not tenable.

Ladyman and Ross (2007) refer to quantum theory, along with general relativity, as ‘fundamental’ physics. In the present context this invites confusion.¹⁹ They do not mean ‘fundamental’ in the sense of grounding all ontology. They mean that these are the parts of physical theory believed to apply everywhere in the universe. I will therefore refer henceforth to ‘universal’ physics. Other branches of physics – for example, the physics of solid states – are parts of the special sciences. Universal physical theory does not ascribe properties to objects that persist through time; it applies statistical distributions to measurement values of observable processes. This is exactly what economic models, both microeconomic and macroeconomic, are up to also.

The failure of atomism in universal physics does not imply the silly conclusion that there are no objects in the world. Special sciences have much to say about them. Nor does the absence of individuals in universal physics imply that there are no individuals in the quotidian sense, as long as the analytic metaphysical characterization of individuality is dropped in favour of acknowledging unique clusters of networked events that resist entropy for long enough to be tracked, and are worth tracking for various purposes even if, *qua* individuals, they aren’t important for law-like generalisations.²⁰ What the failure of atomism in universal physics *does* portend is the failure of *metaphysical* atomism, the idea that facts about objects and individuals are ontologically fundamental.

The core commitment of naturalism is that empirical science is the sole source of objective knowledge, that is, knowledge that is not a record of historical human subjectivity. Naturalists deny that there is any reason to believe any principles of ‘first philosophy’ that can be known prior to science (Maddy, 2007). Consistently with this, they assign a privileged epistemic role to universal physics *because* scientists do. It is a fact about the organisation of science that no hypotheses produced by special scientists that contradict currently accepted generalisations of universal physics are taken seriously. No symmetric restriction governs universal physics, developers of which thus need pay no attention to results of special sciences, except in so far as they may generate interesting cases for specific applications. In this limited sense, universal physics does a crucial *part* of the job historically performed by metaphysics, that of providing a model of *general* reality.

If there were reason to expect special sciences to reduce, even in principle, to universal physics, then that physics *would* be ontologically fundamental in the philosophers’ sense. Unification of sciences would then be best pursued simply by waiting for work in individual special sciences to discover each one’s particular physical foundations. In effect, no distinctive metaphysical project would remain. However,

¹⁹ Indeed, it has invited confusion in *all* the contexts that matter most to scientific metaphysicians. The alternative terminology adopted here should be used by all naturalistic metaphysicians.

²⁰ Scientists are often very interested in individuals – biological species are often taken to be individuals – and scientific research is by no means always conducted for the sake of contributing to law-like generalisations.

Fodor (1974) convinced most philosophers that trends across the sciences point away from general reduction of special sciences (including non-universal parts of physics) to universal physics. Thus naturalistic philosophers face a choice between seeking to provide a general account of reality by showing how, specifically, universal physics *constrains* special sciences, or accepting that reality is dappled and philosophers should follow van Fraassen (2002) in abandoning metaphysics.²¹ It is not possible to *refute* scepticism about the unity of nature. But nor can such scepticism be regarded as established in the absence of a reason to suppose that it can be shown in detail how universal physics *constrains* the special sciences, in specific ways, by relationships other than reduction or grounding.²² There are various potential mathematical and statistical technologies for exploring such relationships. Different clusters of special sciences may be knitted together by different such technologies. This enterprise does not fall within the remit of any special science, though the most promising work currently tends to be performed by applied mathematicians. This is the space for naturalised, or scientific, metaphysics.

Ladyman and Ross (2007) (LR) motivate several general discovery principles for naturalised metaphysics. They defend a version of non-anthropomorphic perspectival realism, grounded in physical restrictions on information flow (Barwise & Seligman, 1997), according to which ontology is scale-relative, and how many scales of measurement are required to specify the range of domains of scientific generalisation without redundancies is a purely empirical question (see also Thalos, 2013). Relative to these scales sciences discover *real patterns* (Dennett, 1991; Ross, 2000; Ladyman & Ross, 2007, pp. 220–238; Wallace, 2014), structural data models in the absence of which some in-principle observable processes couldn't be explained or predicted even by the most computationally powerful physically possible computer. So LR's naturalised metaphysics incorporates a variety of scientific realism according to which special sciences (along with universal physics) *discover* ontology.

The computational processes that discover real patterns should not be imagined as classical AI systems testing logical relationships between qualitative propositions. They should be understood as statistical search systems constrained by theories of causal or pseudo-causal²³ inference. This is not a philosopher's external interpretation of what social scientists do, as is *all* analytic metaphysics; it is how economists who think explicitly about the 'identification' of real structures in the world by elements of models understand their ontological obligations (see Leamer, 1978, a classic in methodology of economic inference).

²¹ Abandonment of metaphysics would not necessarily force a philosopher to also join van Fraassen in rejecting scientific realism. Fine (1986) and Wilson (2006) defend (in very different ways) a picture according to which scientists take themselves to be studying a mind-independent world by modelling the results of experiments, but treat the ontological details reflected in these models merely as provisional approximations for comparatively evaluating measurements. That is, they take a realist attitude to specific scientific discoveries but share the empiricist's disinterest in metaphysics.

²² As per the argument of Sect. 2, supervenience doesn't do the job, as Epstein agrees.

²³ I here have in mind so-called 'Granger causality' (Granger, 1969; Leamer, 1985).

5.2 Social sciences

The similarity noted above in how ontology is stabilised in physics and economics without any need for metaphysical ‘grounding’ is reflected in the fact that in both disciplines ‘theory’ is understood as indicating the appropriate mathematics for representing various target problems and phenomena. However, in some social sciences, particularly experimental psychology, the practitioners typically speak of ‘theories’ rather than ‘theory’. This difference is not a reflection of greater pluralism in psychology.²⁴ I suggest that psychology *is* a case of a (partly) social science that is conducted as if it were answerable to Epstein’s analytic-metaphysical approach to ontology management. But far from being stabilised by this implicit philosophy, such science comparatively suffers from it with respect to successful knowledge accumulation.

Individual psychological ‘theories’ typically denote hypotheses about specific causal dependencies, of the kind we might ‘test’ by associating each such dependency with a qualitative relationship between a null hypothesis and observations we should not expect unless the null hypothesis is false. Because these relationships between models and individual experiments are qualitative, some criterion external to the theory must be used to decide what constitutes inconsistency with the null hypothesis. The standard such criterion is statistical significance. However, use of such low-powered inference as anything other than an exploratory indicator involves an entirely unjustified general assumption about uniformity of distribution of statistics, particularly standard errors, across observation contexts (Ziliak & McCloskey, 2008). There are of course many branches of science in which such *F*-testing and *t*-testing are routine procedure, but these are precisely the branches, such as experimental social psychology, that are suffering from replication crises and consequent collapse of confidence in their allegedly confirmed hypotheses (Yarkoni, 2020). This is not an external philosophical complaint. Psychologists were warned to adopt more properly quantitative methods by their own best mid-20th-century theorists, Edwards (1961)²⁵ and Meehl (2006). The failure is institutional. But it has a philosophical moral, directly relevant to the present context: confusing governance by quantitative theory with qualitative hypothesis testing, one experiment at a time, relies on the idea that recurrent processes can have context-independent grounds that can be qualitatively fixed prior to quantitative measurement – that we know relevant ontologies before they emerge from data. Social psychologists have come too close to acting in accord with Epstein’s advice. Reliable social science should do the opposite.

Effective theoretical constraints best take the form of priors in hierarchical Bayesian search architectures (Gelman et al., 2013; Kruschke, 2014). These priors concern dependencies of observable measurements on parameters and distributions of variables in structural models. Identification conditions specify magnitudes, not qualitatively interpreted outcomes of individual experiments. Inferences concern the whole

²⁴ Having published in both economics and psychology, my experience is that methodological templates are much more restrictively applied by editors and referees in psychology than in economics, and the former are considerably less tolerant of theoretical speculation.

²⁵ Edwards said “Unless someone ... leads psychological statistics out of the wilderness of *t* and *F* tests, psychological experiments may find themselves still seeking the magic 0.05 level when to statisticians and other applied scientists the notion of significance level is only a historical curiosity” (1961, p. 474).

structural model, to which effect sizes of potentially causally influential variables are relative. Posteriors used as priors for the next round of inference could involve deleting or adding variables, or adjusting parameters, or both. At this level of specification there is no place for anything that resembles Epstein's frame principles. But there is as yet no explicit identification of causal relationships, as opposed to autocorrelational ones, either. The point of insisting on structural models is that the scientist is ultimately interested in discovering quantitative causal effects. Hierarchical Bayesian inference can be used to identify structures represented by directed acyclic graphs (DAGs).²⁶ Foundational work on this approach to causal identification, in particular Pearl (2009), has consistently focused on social science applications. Kincaid (2021) provides detailed and illuminating diagnosis and examples of complementarities between DAGs and associated structural equation modeling (SEM), on the one hand, and structural econometrics, on the other hand, for cases of hypothesized sufficient causes of measurable effects. Much more methodological work lies ahead in this area. As Kincaid recognizes, economists and other social scientists are often interested in *enabling* conditions that are necessary but not sufficient causes of outcomes, and in factors that *constrain* causal influences beyond thresholds. Analysis based on DAGs and SEMs waits to be extended to these varieties of causal influence.

So far these are methods used to identify statistical patterns, including patterns that allow for control interventions of the kind that Woodward (2005) and other philosophers of science identify as the indicator of causation. But nothing has yet been said about ontology. Convergence of Bayesian inference provides evidence of real patterns, not only 'mere' patterns, because if patterns were redundant this would imply that unobserved causal influences were still hiding in the model's error term; and part of the point of the methodology is to flush these out. According to Ladyman and Ross (2007), to be is to be a real pattern. Though that is a metaphysical proposition, discovery of real patterns is not in itself metaphysics according to the naturalist; it is just first-order science. It is important to note, however, that causal interpretation of results of hierarchical Bayesian inference depends technically on information theory and maximum entropy calculation. Applied information theory and maximum entropy are motivated not by formal axioms but by universal physical constraints on possible information flow. Thus preference for the method, and interest in unifying science, mutually support one another. It is necessary that no a priori domain restrictions be placed on measured independent variables, so the method is not consistent with a hypothesis to the effect that disciplinary boundaries reflect a dappled world.²⁷

Thus modelling the world, and in particular the social world, using Bayesian statistical analysis, SEMs, and DAGs is not divorced from natural metaphysics. However, it doesn't in itself produce a metaphysical typology that is more informative than the network of discovered real patterns that is incrementally built up. The pos-

²⁶ McElreath (2020) is one of several manuals on carrying out such inference.

²⁷ If some of the variables that make a causal difference are incentives, then that furnishes a basis for interpreting the model as involving economic response (or just 'choice', as an economist who understands choice as equivalent to revealed responsiveness to incentives [see Ross, 2005, 2014] would put it). Alternatively, if values of the dependent variable are mainly driven by adjustments in variables associated with indices of relative social power, we might say that the real pattern is primarily a sociological phenomenon. But these assignments of disciplinary 'ownership' are not metaphysically significant.

sibility of such a typology is not precluded. However, it could not be generated by mapping propositions onto the models that identify real patterns, and then logically analysing the propositions in question. These propositions, no matter how carefully considered, would at best be pragmatic approximations. The practical purposes in question would often be important to many people, and to policy. But this amounts to science popularisation or journalism. Performing intricate logical reconstructions of this journalism would amount to imposing a degree of semantic precision on the popular translations that would have no systematic relationship to anything in the inferential machinery that generated the original source material. Translating a statistical structural model into a natural-language description is *not* relevantly like translating a statement from one natural language into a statement in another natural language; it is an exercise in what Quine (1960) called ‘radical translation’. Radical translation is precisely translation in which stability of ontology across mappings cannot be assumed. Specifically, we must not commit ourselves to preserving the intuitive fundamentality judgments, or in Epstein’s terms the grounding facts, that are woven into the manifest image as folk metaphysics.

There are contexts in which more is at stake than scientific accuracy and inferential power, where people need to create very stable social ontologies for the sake of consistent and fair practice and policy. Here we rely on special language – roughly, legal language – that is (tellingly) often difficult for non-professionals to produce and understand. Co-opting the evolved technology of natural language to this purpose is at best a hack; laws invariably turn out to have loopholes. But these are the cases to which Epstein’s analysis might usefully apply. There is no evident reason, however, why we should regard the analysis in these kinds of cases as metaphysical; it is *stipulative*. As I suggested earlier, many of the examples in Epstein’s book are of this kind. Though they might be regarded as pedagogically motivated analogies, they do not support his programme for reforming social science.

Quantitative social science, then, gives rise to new ontologies, and to the extent that radical translation into quotidian contexts is motivated and culturally successful, these may shift the manifest image. People will go on assuming fundamentality judgments for practical purposes regardless of whether social scientists approve, but specific such judgments may change under scientific influence. *No one* should interpret them as *metaphysically* significant.

New scientifically generated ontologies are not only *different* from folk ontologies; they can appeal to principles unknown outside of science that make them systematically more inferentially powerful. The increasing emphasis in social science (and in other special sciences) on modelling complex systems using deep learning algorithms running on connectionist processing architectures makes such profound ontological adjustment more likely. Here is where the scientific metaphysician finds her call to practically useful action of the kind to which Epstein’s programme aspires. A specific ontology that emerges from the patterns constructed in a deep learning application should no more be taken as metaphysically indicative than the patterns constructed by human cultures.²⁸ However, our knowledge of the histories of AI sys-

²⁸ Each human culture can be regarded as a kind of deep learning system in itself. They work similarly to connectionist architectures at an abstract level. Cultural learning does *not* interestingly resemble theorem

tems we build may give rise to a tractable project of trying to develop a meta-theory of ontological construction – the project taken up by Smith (1996) in a classical logicist AI setting, but in the context of a less restrictive understanding of computation. The crucial theory for such an enterprise will derive from the physical theory of information flow, as Ladyman and Ross (2007) conjecture. Social sciences, the disciplines that study the most complex systems we know of, can be expected to furnish the most powerful data for such an enterprise. Note that this is not, like Epstein's, a vision in which metaphysics is used to boost the success of social science; according to my picture, it is richer social science, by grappling fully with causal and structural complexity, that offers promise for improved metaphysics.

This activity I have described will be metaphysics in Aristotle's sense – it will amount to effective (not simply declarative) unification of the scientific world picture. A sign of this is that it will almost certainly require mathematical frameworks for ontology construction that are more sensitive to structural modelling than set theory and its associated logics. This is not mere conjecture, since quantum physicists have already had to face this challenge (Bub, 1974). There, group theory generally does the job (French, 2014), but its adequacy likely depends on scale uniformity. The scale-relativity of ontology implied by failure of special sciences to reduce to (or supervene on, I add with Epstein!) universal physics (Ladyman & Ross, 2007; Thalos, 2013) indicates that, in the social (and biological) sciences, mathematics that is still more powerful in the discriminations it forces than group theory will be required to specify ontological principles that emerge from nature's array of generated statistics. Homotopy type theory (Univalent Foundations Program, 2013) is a promising prospect for such work. But as far as I am aware, researchers remain far from applying it to anything like the networks of multiple data generating processes that econometricians routinely identify in economic data sets (including even data sets from single experiments). But this is an idea for a naturalistic metaphysical research programme. If it succeeds, there is no reason to expect that revealed homotopy types in social science should correspond to intuitive kinds derived from a priori ontological hunches. More radically, we will not focus on restrictions about possible worlds framed in terms of set-theoretic logic. The enterprise will not resemble analytic philosophy except in its broad ambition to make sense of one whole world that includes quasars, quarks, animals, populations, power hierarchies, economic agents, and monetary economies.

6 Conclusion

I have argued that we should not suppose that metaphysics is usefully brought to bear on science, including social science, merely whenever ontology is critically considered. Application of metaphysics must involve, as Aristotle thought, addressing questions that first-order sciences do not regard as their proper business. This is a point of agreement between Epstein, as a proponent of serious analytic metaphysics, and a promoter of radically naturalistic metaphysics. But the latter is intended as a *com-*

proving or logical deduction.

plete replacement of analytic metaphysics. In particular, it rejects the logicist conception of science that is the essential starting point of the distinctly analytic tradition in philosophy, and it eschews all use of special concepts derived from that tradition. Its tools are mathematical and statistical.

Radical naturalism has attracted some support from philosophers of science in application to physics. But the extant literature indicates that its relevance to social (and behavioural) sciences is less well appreciated. A plausible reason for this asymmetry is that because physics is formulated in mathematical terms to begin with, and pays increasingly less attention to folk ontology over the course of its history, it does not seem so challenging to extend those practices into the philosophy of physics. Then the naturalist can simply be read as developing metaphysical implications of that move. But it has evidently been far from obvious how to go from there to a metaphysics of sciences, such as social sciences, that seem to be grounded in folk ontologies. Unless this part of the naturalist's programme is understood, it may be thought that a serious metaphysics of social science must use the tools of philosophical analysis, because there is no other game in town. 'Relaxed' naturalists have been comfortable with this, on grounds that the tools in question can be substantially drained of transcendental interpretation. I argued, however, that this relaxed attitude amounts to not taking metaphysics seriously.

I have sketched a naturalist metaphysical programme, to ontologically unify the sciences, that anticipates radical ontological revisions in social science based on applications of mathematics and statistics that are not anchored in set theory and its associated logics. This does not require speculation, because AI and formalistic cognitive science are already passing through that revolution. To understand what the revolution portends where social science is concerned, we must be explicit about both what we should expect to do on the other side of the transformation, and what we should *stop* trying to do. The foil, Epstein's analytic metaphysical programme for social science, furnishes a guide to the second part: we should have done with context-independent fundamentality judgments, the constitutive / causal distinction as applied to structural models, and grounding relations specified by reference to possible worlds. Refining concepts of supervenience is a red herring. Future social science may be as ontologically surprising as twentieth-century physics was. We need a metaphysical perspective that will not struggle against the scientific tide and obscure the depth of change.

Funding Open access funding provided by University of Cape Town.

Statement and declaration The author has no competing interests to declare that are relevant to the content of this article.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/>

[licenses/by/4.0/](#).

References

- Ásta. (2012). *Categories we live by*. Oxford University Press.
- Barwise, J., & Seligman, J. (1997). *Information flow*. Cambridge University Press.
- Bird, A. (2022). *Knowing science*. Oxford University Press.
- Blanchard, O. (2016). *Macroeconomics*, 7th edition. Pearson.
- Bub, J. (1974). *The interpretation of quantum mechanics*. Reidel.
- Cartwright, N. (1999). *The dappled world*. Cambridge University Press.
- Cencini, A. (2005). *Macroeconomic foundations of macroeconomics*. Routledge.
- Chater, N., & Loewenstein, G. (2023). The i-frame and the s-frame: how focusing on individual-level solutions has led behavioral public policy astray. *Behavioral and Brain Sciences*, forthcoming.
- Dennett, D. (1991). Real patterns. *Journal of Philosophy*, 88, 27–51.
- Duarte, P., & Lima, G. (Eds.). (2012). *Microfoundations reconsidered*. Edward Elgar.
- Dupré, J. (1993). *The disorder of things*. Harvard University Press.
- Edwards, W. (1961). Behavioral decision theory. *Annual Review of Psychology*, 12, 473–498.
- Epstein, B. (2008). When local models fail. *Philosophy of the Social Sciences*, 38, 3–24.
- Epstein, B. (2011). Agent-based modeling and the fallacies of individualism. In P. Humphreys, & C. Imbert (Eds.), *Models, simulations, and representations* (pp. 115–144). Routledge.
- Epstein, B. (2014). Why macroeconomics does not supervene on microeconomics. *Journal of Economic Methodology*, 21, 3–18.
- Epstein, B. (2015). *The ant trap*. Oxford University Press.
- Fine, A. (1986). *The natural ontological attitude*. University of Chicago Press.
- Fodor, J. (1974). Special sciences, or the disunity of science as a working hypothesis. *Synthese*, 28, 97–115.
- Ford, K. (1996). In Z. Pylyshyn (Ed.), *The robot's dilemma revisited*. Praeger.
- French, S. (1989). Identity and individuality in classical and quantum physics. *Australasian Journal of Philosophy*, 67, 432–446.
- French, S. (1998). On the withering away of physical objects. In E. Castellani (Ed.), *Interpreting bodies: Classical and Quantum Objects in Modern Physics* (pp. 93–113). Princeton University Press.
- French, S. (2014). *The structure of the world*. Oxford University Press.
- French, S., & Redhead, M. (1988). Quantum physics and the identity of indiscernibles. *British Journal for the Philosophy of Science*, 39, 233–246.
- Gelman, A., Carlin, J., Stern, H., Dunson, D., Vehtari, A., & Rubin, D. (2013). *Bayesian data analysis* (3rd edition). Chapman and Hall.
- Granger, C. (1969). Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, 37, 424–438.
- Guala, F. (2016). *Understanding institutions*. Princeton University Press.
- Guala, F. (2022). Rescuing ontological individualism. *Philosophy of Science*, 89, 471–485.
- Haslanger, S. (2018). *Resisting reality*. Oxford University Press.
- Holmström, B., & Tirole, J. (1997). Financial intermediation, loanable funds, and the real sector. *Quarterly Journal of Econometrics*, 112, 663–691.
- Hoover, K. (1988). *The new classical macroeconomics*. Blackwell.
- Hoover, K. (2009). Microfoundations and the ontology of macroeconomics. In H. Kincaid, & D. Ross (Eds.), *The Oxford handbook of philosophy of economics* (pp. 386–409). Oxford University Press.
- Hoyningen-Huene, P. (2013). *Systematicity*. Oxford University Press.
- Janssen, M. (1993). *Microfoundations*. Routledge.
- Kim, J. (1998). *Mind in a physical world*. MIT Press.
- Kincaid, H. (1996). *Philosophical foundations of the social sciences*. Cambridge University Press.
- Kincaid, H. (2021). Making progress on causal inference in economics. In H. Kincaid, & D. Ross (Eds.), *Modern guide to philosophy of economics* (pp. 28–65). Edward Elgar.
- King, J. E. (2012). *The microfoundations delusion*. Edward Elgar.
- Kruschke, J. (2014). *Doing bayesian data analysis* (2nd edition.). Academic Press.
- Kydland, F., & Prescott, E. (1977). Rules rather than discretion: The inconsistency of optimal plans. *Journal of Political Economy*, 85, 473–491.

- Ladyman, J., & Ross, D. (2007). *Every thing must go*. Oxford University Press.
- Ladyman, J., & Ross, D. (2013). The world in the data. In D. Ross, J. Ladyman, & H. Kincaid (Eds.), *Scientific metaphysics* (pp. 108–150). Oxford University Press.
- Lauer, R. (2019). Is social ontology prior to social scientific methodology? *Philosophy of the Social Sciences*, 49, 171–189.
- Leamer, E. (1978). *Specification searches* Wiley.
- Leamer, E. (1985). Vector autoregressions for causal inference? *Carnegie-Rochester Conference Series on Public Policy*, 22, 255–304.
- Lewis, D. (1983). *Philosophical papers, volume 1*. Oxford University Press.
- Lewis, D. (1999). *Papers in metaphysics and epistemology*. Cambridge University Press.
- Lewis, P. (2016). *Quantum ontology*. Oxford University Press.
- Lloyd, J. W. (1984). *Foundations of logic programming* Springer.
- Lohse, S. (2017). Pragmatism, ontology, and philosophy of the social sciences in practice. *Philosophy of the Social Sciences*, 47, 3–27.
- Long, J., & Plosser, C. (1983). Real business cycles. *Journal of Political Economy*, 91, 39–69.
- Lucas, R. (1976). Econometric policy evaluation: A critique. In K. Brunner, & A. Meltzer (Eds.), *The Phillips curve and labor markets* (pp. 19–46). Elsevier.
- Maddy, P. (2007). *Second philosophy*. Oxford University Press.
- McElreath, R. (2020). *Statistical rethinking* 2nd edition. Chapman and Hall.
- Meehl, P. (2006). *A Paul Meehl reader*. Edited by N. Walter, L. Yonce, W. Grove, D. Faust, & M. Lenzenweger. Routledge.
- Ney, A. (2021). *The world in the wave function*. Oxford University Press.
- Oppenheim, P., & Putnam, H. (1958). Unity of science as a working hypothesis. In H. Feigl, G. Maxwell, & M. Scriven (Eds.), *Minnesota studies in the philosophy of science*, v. 2 (pp. 3–36). University of Minnesota Press.
- Pearl, J. (2009). *Causality*, 2nd Edition. Cambridge University Press.
- Pylyshyn, Z. (Ed.). (1987). *The robot's dilemma: The frame problem in artificial intelligence* Ablex.
- Quine, W. V. O. (1960). *Word and object*. MIT Press.
- Ragot, X. (2012). The economics of the laboratory mouse: Where do we go from here? In R. Solow, & J. P. Touffut (Eds.), *What's right with macroeconomics?* (pp. 181–194). Edward Elgar.
- Romer, P. (2016). The trouble with macroeconomics. Working paper: https://www.law.yale.edu/system/files/area/workshop/leo/leo16_romer.pdf.
- Rorty, R. (1999). A world without substances or essences. In R. Rorty (Ed.), *Philosophy and social hope* (pp. 47–71). Penguin.
- Ross, D. (2000). Rainforest realism. In D. Ross, A. Brook, & D. Thompson (Eds.), *Dennett's philosophy: A comprehensive assessment* (pp. 147–168). MIT Press.
- Ross, D. (2005). *Economic theory and cognitive science: Microexplanation*. MIT Press.
- Ross, D. (2014). *Philosophy of economics* Palgrave Macmillan.
- Ross, D., & Spurrett, D. (2004). What to say to a skeptical metaphysician: A defence manual for cognitive and behavioral scientists. *Behavioral and Brain Sciences*, 27, 603–647.
- Ross, D., Ladyman, J., & Kincaid, H. (Eds.). (2013). *Scientific metaphysics*. Oxford University Press.
- Ruben, D. (1989). Realism in the social sciences. In H. Lawson, & L. Appingnanesi (Eds.), *Dismantling truth* (pp. 58–75). Weidenfeld and Nicolson.
- Russell, B. (1911). Le réalisme analytique. *Bulletin de la Société Française de Philosophie*, 11, 53–61.
- Russell, B. (1918). The philosophy of logical atomism. *The Monist*, 28 Reprinted in B. Russell, *Logic and knowledge*, (pp. 177–281), George Allen & Unwin, 1956.
- Russell, B. (1927). *The analysis of matter*. Routledge and Kegan Paul.
- Sarkia, M., & Kaidesoja, T. (2023). Two approaches to naturalistic social ontology. *Synthese*, 201, 1–28.
- Savellos, E., & Yalçın, Ü. (Eds.). (1995). *Supervenience: New essays*. Cambridge University Press.
- Simon, H. (1969). *The sciences of the artificial*. MIT Press.
- Smith, B. C. (1996). *On the origin of objects*. MIT Press.
- Strawson, P. (1959). *Individuals*. Routledge.
- Sugden, R. (2016). Ontology, methodological individualism, and the foundations of the social sciences. *Journal of Economic Literature*, 54, 1377–1389.
- Taylor, L. (2010). *Maynard's revenge*. Harvard University Press.
- Thalos, M. (2013). *Without hierarchy*. Oxford University Press.
- Thomassen, A. (2003). Realism and human kinds. *Philosophy and Phenomenological Research*, 68, 580–609.

- Univalent Foundations Program (2013). *Homotopy type theory* <https://homotopytypetheory.org/book>.
- van Fraassen, B. (2002). *The empirical stance*. Yale University Press.
- Wallace, D. (2014). *The emergent multiverse*. Oxford University Press.
- Wilson, M. (2006). *Wandering significance*. Oxford University Press.
- Woodward, J. (2005). *Making things happen*. Oxford University Press.
- Yarkoni, T. (2020). The generalizability crisis. *Behavioral and Brain Sciences*. <https://pubmed.ncbi.nlm.nih.gov/33342451/>. <https://doi.org/10.1017/S0140525X20001685>.
- Ziliak, S., & McCloskey, D. (2008). *The cult of statistical significance*. University of Michigan Press.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.