

Special Sciences, Conspiracy and the Better Best System Account of Lawhood*

Craig Callender and Jonathan Cohen[†]

If we are on the outside, we assume a conspiracy is the perfect working of a scheme. Silent nameless men with unadorned hearts. A conspiracy is everything that ordinary life is not. It's the inside game, cold, sure, undistracted, forever closed off to us. We are the flawed ones, the innocents, trying to make some rough sense of the daily jostle. Conspirators have a logic and a daring beyond our reach. All conspiracies are the same taut story of men who find coherence in some criminal act.

— Don DeLillo, "In Dallas," pt. 2, *Libra* (1988).

1 Introduction: Special Science Laws and the Conspiracy Problem

Consider the so-called first principle of population dynamics, Malthus's exponential law:

$$P(t) = P_0 e^{rt}$$

where P_0 is the initial population (say, of rabbits), r the growth rate, and t the time. This ecological generalization is very powerful. It supports counterfactuals and crucially enters ecological predictions and explanations. It has an undeniably central role in most presentations of the science of ecology. Indeed, it arguably has the very same central role in ecology that Newton's first law does in classical mechanics: both express a kind of ideal default behavior — exceptions to which are to be explained with further laws (Ginzburg and Colyvan, 2004). In short, pending some good reason for distinguishing Malthus's Law over Newton's first law, there is good reason for taking seriously the idea that the former should count as a bona fide law of nature. (If you don't like "law of nature", substitute "widely applicable projectible generalization" in its stead.)

However, famously, philosophers have suggested a number of reasons for doubting that there can be laws in ecology and other special sciences. In

*This work is fully collaborative; the authors are listed alphabetically.

[†]Department of Philosophy, University of California, San Diego, 9500 Gilman Drive, La Jolla, CA 92093-0119, ccallender@ucsd.edu, joncohen@aardvark.ucsd.edu

particular, recent discussion has focused on the worry that recognizing special science laws would require positing a metaphysically extravagant conspiracy among fundamental particles to explain why the physical objects constituting rabbits (etc.) “know” to move, projectibly, in ways consistent with Malthus’s Law. To be sure, the facts needed to harmonize regularities at different levels can be quite puzzling. For since Malthus’s Law supervenes on a combination of what is projectible (the fundamental physical laws) and what is not projectible (the initial conditions), it can begin to look suspiciously as if the bunnies (or their constituent particles) are conspiring to bring things off just right. But surely that can’t be right? Here are a couple of writers expressing wonder about this matter:

The very existence of the special sciences testifies to reliable macro-level regularities that are realized by mechanisms whose physical substance is quite typically heterogeneous.... Damn near everything we know about the world suggests that unimaginably complicated to-ings and fro-ings of bits and pieces at the extreme micro-level manage somehow to converge on stable macro-level properties (Fodor, 1997, 160).

One can easily join Fodor in the mood of feeling that it is “molto misterioso,” that the motions of particles (and fields) to-ing and fro-ing in accordance with $F = ma$ (or whatever the fundamental dynamical laws prove to be) lawfully end up converging on the regularities of the special sciences. How do the particles that e.g. constitute an economy “know” that their trajectories are required (ceteris paribus) to enforce Greshams law? (Loewer, 2009).

This is, indeed, curious, and, we think, demands some kind of explanation.¹

In this paper we want to face the conspiracy problem squarely. Our question will be how best to formulate a theory of lawhood that allows for laws in the special sciences without resorting to positing outlandish conspiracies. In addition, we want to insist on respecting two further constraints. First, our theory of lawhood should provide for Malthus’s Law a kind of metaphysical autonomy from what is going on at more fundamental levels; while this constraint will need some fleshing out below, the rough motivation is that ecologists shouldn’t need to know or care about the standards for

¹The so-called “problem of explanatory exclusion” is a related, and well-worn, problem in philosophy of mind (Kim, 1989). In our terminology, the problem is to explain how the realizers of (functionally individuated) mental state types “know” to act in ways compatible with higher-level laws of psychology. The standard view in philosophy of mind effectively answers this question by assuming that both psychological states and their subvening physical realizers are both causally efficacious — thus obviating conspiracy on the part of the physical realizers. But, Kim and others have objected, accepting this view threatens causal overdetermination of the outcome. And, just as in the discussion about laws, this situation has led some to abandon projectibility at the higher level (i.e., to give up the hope for a bona fide science of psychology), to make the projectibility at the high level a determinate result of goings on at the lower level (i.e., to give up on the autonomy of psychology), or to abandon the idea that psychological kinds supervene on physical kinds.

formulating physical laws in order to give Malthus's Law its central role in their science. And second, we assume that ecological (and other higher-level) kinds supervene on lower-level physical kinds — there is no *élan lapin* whose exemplification fails to be fixed by the distribution of fundamental physical kinds. Whether we can have all of these things at once is, we think, central to the metaphysical concerns many have had about reduction. Importantly, our list of desiderata is not a mere shopping list, but rather collects the crucial ingredients in what we see as a fundamental question: How can the special science laws be projectible when they supervene upon a basis that is partly projectible (lower-level laws) and partly not projectible (initial and boundary conditions)?²

In what follows we'll review a number of strategies for solving the conspiracy problem — i.e., for allowing for the projectibility of special science generalizations without positing outlandish conspiracies. Some of these strategies turn on accepting accounts of laws (e.g., non-Humean pluralist and classical MRL theories, §§2–3) on which the problem of conspiracy cannot arise in the first place. Alas, we'll argue, these accounts do less than we should want a theory of laws to do, and so should be rejected. Next (§4) we'll take on a recent strategy, due to Albert (2001) and Loewer (2009), that aims to solve the conspiracy problem by treating special science regularities as probabilistic corollaries of statistical postulates over low-level initial conditions. We'll argue that, while this view deals admirably with the supervenience constraint and may also obviate conspiracy, it rests on unrealistic and insufficiently general assumptions about the relation between special and fundamental sciences. Finally (§5), we'll present our own favored view of laws, a revised version MRL theory that we have called the Better Best System (BBS) (Cohen and Callender, 2009). BBS offers a picture on which, although all events supervene on a fundamental level, there is no one unique locus of projectibility; rather there are a large number of loci corresponding to the different areas (ecology, economics, solid-state chemistry, etc.) in which there are simple and strong generalizations to be made. While we expect that some amount of conspiracy-fear-inducing special science projectibility is inevitable given BBS, we'll argue that this is unobjectionable. It follows from BBS that the laws of any particular special or fundamental science amount to a proper subset of the laws. From this vantage point, the existence of projectible special science generalizations not guaranteed by the fundamental laws is not an occasion for conspiracy fantasies, but a predictable fact of life in a complex world.

²A further, and much-discussed, set of problems about reduction concerns scientific methodology and scientific explanation. The methodological issue is whether and how the project of formulating special science laws is appropriately constrained by the goal of integrating those laws with those at lower levels. The main explanatory issue is whether higher level generalizations such as Malthus's Law (perhaps together with other higher level generalizations) can supply genuine explanations without enlisting lower-level (physical) kinds or laws. We'll largely ignore these issues in what follows.

2 Non-Humean Pluralism

The view we're dubbing Non-Humean Pluralism is frequently motivated by noticing that special science generalizations have two features that laws are typically assumed not to allow: they are (by definition) formulated in terms of non-fundamental kinds, and they tolerate exceptions (the so-called problem of the provisos). And yet, many have thought, such generalizations are not only true but projectible: this is why we are inclined to apply Malthus's Law next time questions about bunny distribution arise. Such observations have led some to favor a picture of the special sciences as involving (not laws, but) apparatus found outside the so-called Humean mosaic of events — apparatus that includes causal principles, mechanisms, capacities, invariances, powers, dispositions, and the like. Thus, for example, Cartwright (1989) suggests that the regularity that aspirins relieve headaches must be explained by adding "capacities" or "singular causes" of particular aspirins — features not found in the Humean mosaic, but which secure for the generalization about aspirins a strength it could not otherwise obtain. Namely, these elements ensure, projectibly, that the higher-level cause bring about the higher-level effect in certain circumstances.

Although logically detachable, the view we have in mind couples this non-Humeanism with a kind of ontological egalitarianism among levels. On Cartwright's view, the world is "dappled": objects are endowed with capacities that produce regularities only in certain highly limited scenarios. These scenarios might be in the economic, biological or physical realms (among others), where no one realm is distinguished as fundamental. Although there are significant differences among them, positions similar to Cartwright's are found in the work of Dupré (1993), Giere (2006), and Teller (2009). On all of these views, science provides windows on to the world, but these windows are always partial and usually somewhat opaque. Nor can they be assembled together to form one true consistent picture. Crucially, the disunified picture of the world found in scientific methodology is mirrored for these authors by a genuine ontological disunification. As a result, there is no more reason to assert that the physical determines the ecological than the other way around.

The good news about Non-Humean Pluralist theories of the kind under consideration is that they answer to several of our desiderata. First, they are designed specifically to accommodate the projectibility of higher level generalizations, even if the latter are not officially counted as laws. Either way, the intention is to allow that these generalizations can do the things we want laws to do — figure in explanations, support counterfactuals, and so on. So we can grant that they make room for robust special science generalizations (albeit not laws). Second, the view accepts the metaphysical autonomy of the special sciences in positing genuine higher-level elements (the capacities, singular causes, and whatnot) absent at the lower levels: generalizations involving the latter entities can't even be stated in lower-level vocabulary, since these entities don't occur at the lower level. Third, there's no mystery about how lower-level entities know to conform to higher-level regularities; indeed, on this pluralistic

picture that question can't arise. Each non-fundamental level has non-Humean elements that ensure that the regularities at that level (and only at that level) obtain. But these non-Humean elements and the entities themselves aren't stitched together into a unified picture that allows the worry about how these regularities are harmonized (more on this immediately below).

For all their virtues, however, such views blithely give up one of the constraints we want to enforce — the supervenience of bunnies and other high-level entities on the lower-level kinds. According to the disunified picture, supervenience of macro-kinds/properties on micro-kinds/properties is willingly eschewed. Within each theory one can see this different ways. Cartwright (1999, 32–33), for instance, gives up supervenience in part because of her view of all science being *ceteris paribus*.³ Dupré (2001, 161) reaches the same conclusion because of his belief in “downward causation”, i.e., causation by the high-level properties on the low-level properties. And Teller (2009) (and presumably Giere, though Giere (2006) doesn't discuss supervenience itself) thinks the different windows on to the world don't have to add up to a single coherent picture of one world. In all cases, what is seen through the ecology window needn't mesh with what is seen through that of physics, and supervenience is abandoned. We can now better see how the conspiracy worry is avoided on such views. Supervenience is a crucial ingredient of the worry. Without it there is no mystery, for we don't assume that the regular behaviors of entities at different levels have to be harmonized.

Yet for us the price of abandoning supervenience is too great. If we have reason to believe anything in science, it's that macroscopic entities are constituted by microscopic ones and their relations. If we insist on this, as we do, Non-Humean Pluralism won't do all of the things we want our metaphysics of science to do.

Moreover, to the extent that the theory relies on non-Humean additions to scientific ontology — additions that contribute to the denial of supervenience — there is a question whether they earn their keep. Of course, there are many large issues surrounding Humean versus non-Humean views of the world, and we will not tackle them here. Suffice to say, however, that there exist many powerful concerns about capacities and the like — what they are, how we know them, etc. (Psillos, 2002, §7.4.2) — and whether they help us understand special science generalizations (Psillos (2002, §7.4); Schrenk (2007)). Given

³To be fair, Cartwright (1999) may be read as endorsing supervenience in certain limited cases. She writes, for instance, that “We do not want colour patches to appear in regions from which the laws of physics have carried away all matter and energy” (33). When supervenience holds, then, she does seem to face the worry about harmonizing the two levels. She sees this worry and answers with a kind of theological parable (33). In one, dubbed the reductionist picture, God creates the universe by specifying the laws of physics and the precise initial conditions; God then leaves to Saint Peter the computationally tricky task of deriving the higher-level regularities. In the other, dubbed the pluralist picture, God instead cares about regularities, and writes down all the micro- and macro-regularities the world should have; Saint Peter is then charged with finding initial conditions such that these regularities all hold. While this story is evocative, it doesn't yet answer our worry: intuitively “conspiratorial” initial conditions will still need to be specified, whether by God or by Saint Peter.

these concerns about non-Humean treatments of projectible generalizations, we'd prefer, if we can, to find a position that carries the advantages of such views without requiring supervenience-violating additions to the Humean mosaic.

3 Classical MRL

Given that we want our theory of lawhood to respect supervenience in a way that non-Humean accounts do not, we now want to turn to what is perhaps the most important of the (broadly Humean) views that make supervenience a central virtue. The account we have in mind is the "Best System" or (classical) MRL view associated with Mill, Ramsey, and Lewis.

Classical MRL starts by conceiving of our knowledge of the world as a deductive system containing axioms and results derived from those axioms.⁴ Of course, there will be different ways of arranging our knowledge of the world into such a deductive system. Some such systems will be stronger than others in the sense that they contain (among their axioms and derived consequences) more truths about the world than others — some such systematizations carry more information about the world, or, equivalently, rule out more possible ways the world might have been, than others. Other true deductive systems will be simpler than others in the senses that they include a smaller set of independent axioms or are syntactically less complex.⁵ Significantly, strength and simplicity seem to be competing virtues: adding more axioms to a system increases a system's strength at the cost of simplicity, while taking away axioms increases simplicity at the cost of strength.

The heart of the MRL approach to lawhood is to say that a true generalization is a law if and only if it is an axiom of all the "Best Systems" — axiomatic systematizations that best balance strength and simplicity. A generalization like Schrödinger's equation might plausibly count as a law on this account because any deductive system lacking it would be inferior in strength (by leaving out many truths about the world) or simplicity (by recapturing the otherwise left-out truths only by including a huge list of underived axioms about the locations of particles at each moment). In contrast, the true generalization that all your children have brown hair is plausibly something that could be captured in a Best System without making it a law (viz., it wouldn't be too costly in terms of simplicity to include axioms listing all your children and their hair color, from which we could derive the needed result as a consequence).

⁴For canonical statements, see Lewis (1973, 73ff), Lewis (1983, 365–368), Earman (1986); Lewis (1994, 478ff) extends the view to incorporate objective chances.

⁵Understanding the simplicity of a system is a difficult problem that we won't attempt to resolve here. For a non-syntactic alternative, it is possible to regard simplicity in a manner more connected to scientific practice. For instance, we might consider that a system merits the simplicity score it gets based on how well it embodies various theoretical virtues, e.g., unification, consilience.

There is a lot to like about the classical MRL account of laws. In a systematic assessment, Loewer (1996) points out that MRL captures many of the most important features of laws, including ensuring that laws are explanatory, are confirmed by their instances, and support counterfactuals. Moreover we are unconvinced by most of the criticism of MRL in the literature, which consists largely of imagined counter-examples — scenarios in which we are invited to share the intuition that there are laws other than those occurring in the Best System (see Carroll (1994) and references therein). We find these criticisms unconvincing because, as Earman (1986) and Loewer (1996) point out, these alleged counterexamples beg the crucial question in taking the regularities at issue to be laws. Finally, classical MRL respects the supervenience requirement: on this view, what makes a generalization a law is simply its occurrence as an axiom of all the Best Systems — not its relation to special entities from outside the Humean mosaic. Therefore, the view does not require the introduction of new, non-supervenient entities to figure in laws.⁶

Despite its many attractions, classical MRL also has its problems.

First, the view has difficulty making room for laws in the special sciences. This is because classical MRL requires (as do many other traditional theories of lawhood) that laws must relate fundamental natural properties; but, of course, special science generalizations (by definition) don't. One might attempt to overcome this problem by "translating" higher level predicates into low-level disjunctive predicates and then joining the disjunctive predicates into would-be laws. But this won't help proponents of MRL, since the results of such translations will be non-simple. Moreover, special science generalizations tolerate exceptions and are generally incomplete descriptions of reality; therefore, candidate special science generalizations will plausibly lose in terms of strength to candidate fundamental laws. Consequently, it is plausible that the generalizations of the special sciences are neither simple nor strong, so couldn't be MRL laws.

What about the autonomy constraint? Strictly speaking, it appears, trivially, that classical MRL cannot secure the autonomy of special science laws since (as we have just noted) it cannot recognize special science laws at all. Slightly less strictly speaking, however, classical MRL appears to abide by the spirit, if not the letter, of the autonomy constraint. While classical MRL may not recognize special science generalizations as laws, it can still recognize them — namely, as (a subset of) the conclusions derivable from the MRL laws together with a set of initial conditions. And it appears that the special science generalizations, considered as derived consequences of the MRL laws together with a set of initial conditions, *are* indeed autonomous from the fundamental laws: for the proposal is that they are derivable not just from the laws of physics, but from

⁶This is not enough to guarantee supervenience; it only shows that supervenience won't be violated in the way in which non-Humeans violate it. But in fact proponents of MRL typically go on to endorse the generalization that the distribution of high-level kinds is fixed by the distribution of fundamental kinds. That said, it is also worth noting that classical MRL is independent of the further thesis (sometimes called "Humean supervenience") that everything supervenes upon local matters of particular fact (Lewis, 1986b, ix).

those laws together with initial conditions. Alas, making the higher level autonomous in this way leads directly to the worry about conspiracy. For it means that the laws of the lower level fail to determine behavior at the higher level.

Aside from the worries we have raised concerning the desiderata of §1, there are serious further internal problems with classical MRL. In our view, the most serious of these concerns the theory's use of inter-system comparisons of strength, simplicity, and balance.⁷ This problem arises because assessments of simplicity are inherently relative to an inventory of basic kinds or basic predicates. Loewer brings out this point by reference to a famous example from Goodman (1954):

There is a problem concerning the languages in which the best systems are formulated. Simplicity, being partly syntactical, is sensitive to the language in which a theory is formulated, and so different choices of simple predicates can lead to different verdicts concerning simplicity. A language that contains 'grue' and 'bleen' as *simple* predicates but not 'green' will count 'All emeralds are green' as more complex than will a language that contains 'green' as a simple predicate (Loewer, 1996, 109).

Adapting terminology from Quine (1970), simplicity is an immanent rather than than transcendent notion: it is defined relative to a set of basic kinds or basic predicates. An immediate consequence of the immanence of simplicity is that there is no way of making relative simplicity assessments between systems that disagree with respect to the basic kinds (systems that differ in respect of the terms they make available for basic kinds). This is a problem for MRL because, as set out above, that account demands that we engage in inter-system comparisons of relative simplicity.

It has been less frequently noted (but not unnoticed) that the strength of a deductive system is, likewise, immanent rather than transcendent, as is the notion of balance among systems (see Cohen and Callender, 2009).

Prima facie, the realization that simplicity, strength, and balance are immanent — what we'll call *the problem of immanent comparisons* — is a devastating blow to MRL. For what counts as a law according to that view depends on what is a Best System; but the immanence of simplicity and strength undercut the possibility of intersystem comparisons, and therefore the very idea of something's being a Best System.

Another way the problem manifests itself is via a trivialization worry raised by Lewis (1983, 42); (cf. Loewer, 1996, 109). This difficulty arises from consideration of the kind F that holds of all and only things in the world where an arbitrary system S holds. If allowed to compete, it's hard to see why the trivial generalization $\ulcorner(x)Fx\urcorner$ wouldn't be the Best System for all such worlds: it is very simple, and also very strong (it strictly implies all truths,

⁷The worry as it applies to simplicity has been noted by Lewis (1983, 366–368); Lewis (1986a, 123–124); Earman (1986); van Fraassen (1989, 41–43, 51–55) Taylor (1993, 82); Loewer (1996, 109).

and so, if individuals are world-bound, singles out a unique possible world). Consequently, this trivial generalization would be the one law for any world.

Both versions of the problem cry out for solution by distinguishing some special set of basic predicates or kinds. If we insist that all the putative Best Systems be written in a privileged vocabulary (refer to some distinguished set of kinds), then we can make inter-system comparisons of simplicity, strength and balance. And we can also rule out noxious predicates like the above $\ulcorner F \urcorner$ for failing to occur in our privileged vocabulary/inventory of kinds. Thus, Lewis (1986a) holds that, as a matter of fact, certain “perfectly natural” fundamental kinds are basic and others are not, and that only axiomatizations expressed in terms of these perfectly natural kinds are candidates for being Best Systems. However, as we have argued elsewhere (Cohen and Callender, 2009), this solution has the defect that, because it fails what Earman (1986, 85) calls the “empiricist loyalty test,” it makes lawhood epistemically inaccessible.⁸ Briefly, we can bring out the problem by considering two possible worlds, w_1 and w_2 , such that the generalization $\ulcorner (x)(Fx \supset Gx) \urcorner$ is part of the on balance simplest and strongest system in both, but where $\ulcorner F \urcorner$ and $\ulcorner G \urcorner$ refer to perfectly natural kinds in w_1 but not w_2 . Unfortunately, there seems to be no way of telling what world we’re in, so, given Lewis’s solution, no way of telling whether the generalization at issue is a law. Normally we can’t tell whether what we think is the best system given some kinds really is the Best System, but perfect naturalness adds an extra layer of skeptical possibility. Even if we saw all of world history, possessed unlimited computing power, and so on, we still couldn’t determine if the kinds picked out by $\ulcorner F \urcorner$ and $\ulcorner G \urcorner$ were perfectly natural, and so, whether $\ulcorner (x)(Fx \supset Gx) \urcorner$ is really a law of nature.

In all, then, though we are sympathetic to the classical MRL position, it seems that we need to look elsewhere to meet our needs. In the next two sections, therefore, we’ll explore two different ways of modifying classical MRL, and consider the consequences of these modifications for the conspiracy problem.

4 Special Sciences as Probabilistic Corollaries

One of the more important developments to emerge from recent work on laws is an elaboration of classical MRL that, although formulated in terms of the fundamental lower level-kinds, is intended to capture higher-level probabilistic generalizations such as those of thermodynamics (Lewis, 1994; Albert, 2001; Loewer, 2004, 2009). In outline, the idea here is that simply adding to candidate Best Systems a statistical postulate over the universe’s initial conditions will — at very little cost to simplicity — significantly add to their strength by putting thermodynamic generalizations within their grasp. Thus Albert (2001) suggests that, to explain high-level thermal generalizations

⁸As Cohen and Callender (2009) emphasize, this criticism goes through on even a very general formulation of the empiricist loyalty test that should be acceptable to empiricists and non-empiricists alike.

from the perspective of microphysics, physics (via statistical mechanics) posits a probability distribution over the precise initial conditions of the universe. The initial conditions with greatest probabilistic weight then trace histories that we would call thermodynamic. Since adding this posit to physical candidate best systems increases their overall balance of strength and simplicity, the resulting Systems rather than those without the posit are truly Best, hence truly law-constitutive. If so, then the thermodynamic generalizations really would be laws, hence projectible.

This treatment of thermal generalizations is very satisfying. While fully embracing the idea that macroscopic thermodynamic systems supervene upon more basic physical entities, it explains how generalizations in the special sciences can be projectible — viz., by tracing out objectively likely histories in state space. Moreover, it rids the world of conspiracy by making the truth of such generalizations objectively likely: air is spontaneously spreading through its available volume because that is the most likely thing for it to do, given the probability distribution assigned to the microstates underlying the air.

This idea can be found, of course, in the work of Gibbs, Maxwell and especially Boltzmann. However, as Albert (2001) and Loewer (2009) develop it, this understanding of statistical mechanics turns out to have an astoundingly broad explanatory scope. Once one places a probability distribution over the world's precise initial microstates, then the theory implies more than probability assignments for thermodynamic macrostates and their generalizations. The theory will also assign probabilities to any generalization whatsoever. Thus Albert (2001) speaks of this fundamental chance plus the laws entailing a probability assignment for spatulas being found in kitchens versus other rooms of apartments. And Loewer (2009) extends this to all the special science generalizations in an effort to deal with the conspiracy problem. The idea is that the fundamental chance assigns probabilities not only to thermal generalizations, but also to bunnies reproducing according to Malthus's Law, economic generalizations, psychological generalizations, and even regularities not encoded in any particular science, e.g., the claim about spatulas. These regularities will either be heavily weighted, probabilistically, by this primordial chance, or not. If they are heavily weighted, then, of course, there is no conspiracy. Bunnies aren't conspiring any more than particles are when they spread out through a room. We explain why the particles are to-ing and fro-ing appropriately — we explain how they "know" to orchestrate kosher histories — by saying that most evolutions of the state space according to the privileged measure are those in which they do so. Furthermore, we can turn this reasoning around: since these special science regularities keep happening, and seem to be projectible, it must be that they are indeed heavily weighted by the fundamental chance, for otherwise they are, literally, a lucky run. If you keep rolling a die and it predictably lands on 3 one-sixth the time, one expects a probabilistic mechanism to be responsible for this. After a certain time, it can't just be luck. What Albert and Loewer do is offer the fundamental chance and laws as the mechanism responsible for why the special science regularities keep materializing. With this fundamental chance we have no conspiracies

but a “unified science” program that is, in some senses, stronger even than the aspirations of Oppenheim and Putnam (1958).

Although we have sketched it briefly, we hope we have made clear the outlines of Albert’s and Loewer’s view and shown that it offers a non-conspiratorial explanation of some higher-level generalizations. On this solution, we explain why the formerly accidental occurrence of special science patterns instead becomes probabilistically necessary, thereby making the relevant special science generalizations projectible. Moreover, because it is such an ontologically conservative extension of the classical MRL picture, it inherits from its predecessor the advantage of respecting supervenience: the Albert-Loewer view agrees with classical MRL in the distribution of fundamental and non-fundamental properties, so respects supervenience exactly to the extent classical MRL does. Although we might quibble with their take on the statistical mechanical account of the thermodynamic generalizations, overall, we find this account impressive. So far, then, so good.

Unfortunately, we are unconvinced that the view under consideration merits acceptance.

For one thing, although we accept for the sake of this paper that the view makes sense of the non-conspiratorial projectibility of the generalizations of thermodynamics, we are deeply skeptical that it can be extended to all other genuine higher-level generalizations.⁹ For this extension to succeed, it must be that all genuine and truly regular generalizations — e.g., presumably, Malthus’s Law — are likely according to the chance posited by physics. We doubt that this requirement can be sustained.

There is not a shred of evidence that the chances used in ecology are the ones used in statistical mechanics. A chance is relative to a particular measure over a particular state space. The statistical mechanical chance is based on Lebesgue measure over phase space ($3n$ -positional dimensions and $3n$ -momenta dimensions, where n is the number of particles). Ecological systems are sometimes modeled via state spaces with measures on them — sometimes even Lebesgue measure. However, the physical and ecological chances are unrelated, to our knowledge, because the state spaces are different. The (classical) physical one is parametrized with respect to position and conjugate momentum, the ecological ones are parametrized with respect to ecological variables, such as ‘number of age 1 females’, ‘number of age 2 females’, and so on. Are the generalizations that are highly probable in the one space highly probable in the other? We have no idea, and neither does anyone else. The solution in question requires that all of this work out, but we don’t see any reason for such confidence.

To press the point a little, let’s step back and note that, logically speaking, it is easy to construct the following. Take a trajectory through one state space that is typical according to the probability measure adapted to that space. Now “translate” this trajectory into another state space. Then it’s certainly possible that, with respect to a probability measure on that second state space, the

⁹See Callender (2010) for alternative views and a more thorough discussion of thermodynamics.

resulting trajectory is atypical. Just think of a sequence of coin flips of a fair coin as described by the standard measure in a ‘heads/tails’ state space. Now translate this sequence into a state space based on a different partitioning of this sequence of events (say, a ‘flipped-by-a-female/flipped-by-a-male’ state space, or an ‘initial-state-heads-up/initial-state-tails-up’ state space). Depending on the partition, there is no special reason to think that that sequence of events is typical with respect to the new measure, even if it’s an intuitively natural one.

Given that this is possible, it is difficult to believe that Albert’s and Loewer’s confidence is warranted, if for no other reason than that there are so many cases to cover. The class of so far empirically adequate generalizations — call it G — is enormous. G includes all of what’s successful in ecology, biology, economics, as well as more mundane general facts — e.g., that in households pillows tend to be on the beds. Surely it is implausible that every single generalization g in G is weighted heavily by Lebesgue measure when g is translated into phase space (whatever that would look like). Human beings are marvelous pattern detectors and find all sorts of pairs of patterns/partitions in nature. To think that everything we discover will turn out to be typical patterns in phase space according to Lebesgue measure is extremely optimistic. Surely, that is, some of these g s are due to special (non-projectible) initial conditions or lucky runs according to the Lebesgue measure chance. We have no idea how many generalizations there are like this. They might be some of our very best generalizations in some of our best theories. In the face of this ignorance, it seems fantastic to hope that all, or even most, of the g s in G turn out probabilistically likely according to Lebesgue measure on phase space. This suggests that the solution at issue non-conspiratorially explains the projectibility of only a proper subset G^* of the generalizations in G . Conspiracy is still afoot.

Moreover, the worry just developed can be turned around. Call the set of generalizations that are typical with respect to Lebesgue measure on phase space T . Then we have argued that there is no reason to think that all of G is a subset of T , and consequently, that the conspiracy problem lives. But there is also no reason to believe that every generalization t in T is a member of G . That possibility is just as awkward as the other one for this view. In this case, the view would be telling us that t is a special science law even though no scientist would have reason to think it is. This point leads us to our next criticism, namely, that from the point of view of the constraints mentioned in §1: it seems to threaten the autonomy of the special sciences to which it does extend.

First off, autonomy seems threatened from the beginning since to the extent that the view has anything to say about a generalization g , it does so precisely by treating g as a probabilistic corollary of fundamental physics.¹⁰ A proponent might respond that, although the truth of the generalizations in G is probabilistically fixed by the fundamental physics (including its

¹⁰This allows that there might be physically accessible worlds where g fails; such worlds would be treated by the view as objectively unlikely.

chance), there is still a methodological sense in which special sciences are autonomous. Ecologists can discover, formulate, and use Malthus's Law without knowing any physics at all, and without caring whether this and the other generalizations they work with are derivable from the laws of physics, or the laws of physics plus bridge laws, or the laws of physics plus statistical postulates, or anything else. Of course, the Albert-Loewer view holds that the special science laws *are* so derivable; but it needn't require that special scientists know or care about this fact. On the other hand, the position under discussion does seem to amount to a long-run constraint on the acceptability of laws in the special sciences: whether or not any particular special scientist knows or cares about what happens at the end of inquiry, the view still has it that those generalizations that fail to be probabilistic corollaries of fundamental physics are *ipso facto* not laws.

There is another sense in which this view fails to respect the autonomy of the special sciences. The Albert-Loewer view holds that a true special science generalization g not derivable from the fundamental physical laws plus statistical postulates is objectively unlikely. The continuing satisfaction of g so far is analogous to a long run of rollings 7s whose continuation cannot be assumed for even the very next case. But from the perspective of the relevant special science, g might enjoy the very same virtues as a distinct generalization g' that is deemed objectively likely by the Albert-Loewer view. Both g and g' might, as far as we can see, support counterfactuals, play a role in explanations, and even play a central role in the theoretical core of the relevant science. To impugn g for reasons entirely external to the science in which it plays a role strikes us as very much against the spirit of autonomy.

And finally, as mentioned, plausibly there will be generalizations t that are not in G . If anything is an affront to autonomy, this is. For t , in virtue of being in T , is the generalization that truly supports counterfactuals, plays a role in explanation, and so on because it is likely according to the primordial chance. Special scientists forming the best system over their predicates might even consider and reject t . Yet on this view they would be making a mistake. Special scientists, even ideal ones, ultimately don't have any input into what is projectible in their domain.

Thus, we claim, while the Albert-Loewer view offers a non-conspiratorial and supervenience-respecting account of the projectibility of some special science generalizations, it plausibly fails to extend to many cases that we care about, and in addition, extends to some cases that we don't care about. The view requires its own delicate contrivance: the sets G and T must more or less precisely overlap. Where they do not, we have trouble. For gs not in T , conspiracy lives; for ts not in G , autonomy dies.

5 The Better Best System

We have urged elsewhere (Cohen and Callender, 2009) that a different way of elaborating MRL — what we call the Better Best System theory, or BBS —

solves the internal problems facing classical MRL. We also believe that the same view can answer to the needs of §1. In this section we'll present BBS, briefly catalog some of its benefits, and then consider the light it sheds on the conspiracy problem.

5.1 BBS: A (Very) Short Synopsis

Recall that one of the most significant problems for MRL was its reliance on simplicity, strength, and balance comparisons across systems that differ in their basic kinds. It is possible to view the Albert-Loewer view as a way of circumventing that problem by requiring a single microphysical language for the expression of all true generalizations; unfortunately, for reasons already discussed, we strongly doubt that this strategy can succeed. In contrast, BBS embraces the diversity of basic kinds used by different sciences. The guiding idea is that even if there is no transcendentally Best System, nothing stops us from assessing the immanently strongest, simplest, and best balanced axiomatizations relative to a specific choice of basic kinds K (/to a specific choice of basic predicates P_K). Given such an assessment relative to a choice of basic kinds K (/predicates P_K), we can say that a true generalization is a law relative to K (/ P_K) just in case it appears in all the immanently Best Systems relative to the basic kinds K (/basic predicates P_K).¹¹

Although BBS has many attractions, we'll content ourselves with listing just three here.

First, since BBS is relativized to kinds in a way that MRL was not, it makes do without transcendent comparisons, and thereby sidesteps the problem of immanent comparisons.

Second, it renders unproblematic Lewis's odd predicate $\ulcorner F \urcorner$ that holds of all and only things in the world where an arbitrary system S holds. The proponent of BBS can accept that, if allowed to compete, the generalization $\ulcorner (x)Fx \urcorner$ is a Best System for all such worlds. Yet she can remark that $\ulcorner (x)Fx \urcorner$ is *not* a Best System relative to the kinds we care about. If we care about mass, charge, spin, etc., then relative to these kinds, Schrödinger's equation (for example) might result, not 'all events are F '. Properties like that picked out by $\ulcorner F \urcorner$, and the ensuing threatened trivialization of MRL are ruled out for lack of interest rather than any intrinsic deficiency. BBS solves the problem by appeal to scientific interests rather than by appeal to metaphysics, a solution much more in keeping, we believe, with the spirit of MRL.

Third, so long as the kinds in terms of which it is formulated are epistemically accessible, BBS will also pass the empiricist loyalty test and so will make lawhood epistemically accessible.

¹¹We are not the first to consider relativizing MRL. Earlier proponents of some such strategy include Halpin (2003); Taylor (1993, 97); Roberts (1998). Indeed, it is possible to read Lewis (1983, 368) as endorsing a relativized version of MRL in claiming that laws and kinds are chosen as a package deal. BBS was developed prior to our discovery of a similar account of Schrenk (2008).

5.2 BBS on Special Science Laws

For the reasons set out above (among others), we think BBS holds significant advantages over rival accounts. We now turn to the question of how well BBS answers to the desiderata of §1. After arguing that the view nicely reconciles lawhood with supervenience and autonomy in the special sciences, we'll use it to motivate a metaphysically deflationary treatment of the worry about conspiracy. We'll argue that BBS respects our desiderata as well as we have reason to expect that they should be respected, and additionally brings to light ways in which they should not be respected.

We begin, then, by pointing out that BBS allows for the projectibility of special science generalizations. Given BBS, all that is required for lawhood in a domain is that a generalization figure in the on balance and immanently simplest and strongest axiomatizations of that domain. We can axiomatize whatever domains we favor. Every conceivable carving up of the world will, assuming we can make sense of simplicity, strength and balance with respect to these kinds, result in a competition, and where there is a winner, a Best System. That this should be possible in ecology is no more puzzling for BBS than that it should be possible in physics. Significantly, BBS doesn't posit some threshold of strength or simplicity that must be had for lawhood; consequently, BBS does not require that aspirants to the title of lawhood should be universal and exceptionless. An incomplete generalization may be sufficiently simpler than any complete one, and thereby gain entry into the Best System. Nor does it require that they be translatable (simply or otherwise) into basic vocabulary. Thus, given BBS, the existence of exceptions to special science generalizations is no barrier to their lawhood.¹²

Consider supervenience next. Recall that the problem about supervenience we raised against non-Humean Pluralism was that that view ensured special science causal powers (and thereby sidestepped the conspiracy worry) only by adding extra entities to its ontology, thus only by committing to the denial of supervenience. By contrast, BBS is fully consistent with supervenience. While BBS has it that higher levels will contain kinds (say, rabbits) not present at lower levels, it does not require in order for the projectibility of generalizations about such higher level kinds the addition of non-supervenient entities or events. Rather, on this approach, the laws and kinds are what result from

¹²Claims to the contrary have often motivated philosophers to give up on lawhood in the special sciences. BBS not only shows that law-alternatives are not necessary to account for the special sciences, but in fact also naturally fills holes in some proposals. For example, Woodward (2000) develops a popular invariance-based account of explanation to answer problems that arise in understanding special science generalizations. According to this theory, "whether or not a generalization can be used to explain has to do with whether it is invariant rather than with whether it is lawful. A generalization is invariant if it is stable or robust in the sense that it would continue to hold under a relevant class of changes" (197). But which changes? And what provides the truth conditions for these counterfactuals? Woodward's non-reductive account is silent on these questions, and as a result, it may be viewed as circular if we insist on an answer (i.e., the invariant generalizations are those that don't change under the interventions that leave them the same). BBS, by contrast, gives non-circular answers to these questions via the Best System for that special science.

best systematizing (parts of) the Humean mosaic. Significantly, and unlike, say, the internal realism of (Putnam, 1987), the basic ontological framework associated with BBS is a single world filled with events. What it adds to this ontological framework is a conception of laws as certain kinds of sophisticated summaries of the events. Now, because interests vary among summarizers, some of these summaries will be best told with vocabularies that carve up these events in different ways; and, to be sure, projectibility will fail to transfer between different of these summaries. But, crucially, that doesn't mean that the rabbits — or anything else — fail to supervene upon the basic events.¹³

Third, there is a clear sense in which the BBS special science generalizations (/laws) are autonomous in a way they aren't according to the Albert-Loewer view, since BBS does not demand that the former are derivable from the truths (laws and initial conditions) of fundamental physics. Consequently, BBS allows that the correctness of special science generalizations does not depend on their eventual vindication by the metrics of physics. In this, BBS is clearly on the side of scientific practice: ecologists just don't look over their shoulders to see what physicists are doing before deciding whether an ecological generalization is on balance strong and simple. When considering Malthus's Law, ecologists are weighing the pros and cons of this as a generalization of rabbits, and so on, not quarks. Economists, ecologists, and so on use their own standards of balance, simplicity and strength, based on organizing knowledge in their own domains, not fundamental physics. BBS respects this fact, making the metaphysics of laws much closer to the methodological practice.

Finally, we now want to argue that BBS reveals the conspiracy "problem" to be overstated in certain ways, and therefore apt for dissolution rather than solution. We'll contend that BBS provides the resources to say just as much (but no more) than needs to be said about the avoidance of conspiracy.

Recall that, while classical MRL did not address conspiracy directly, the two other accounts we have considered so far attempt to solve (or dissolve) the worry about conspiracy by adding entities to the world that ensure higher level generalizations hold. For Non-Humean Pluralists, the solution comes from

¹³In fairness, the issue of supervenience is somewhat more complicated than this lets on, partly because BBS can be developed in so many different ways. One could, we suppose, combine a BBS-like pluralistic-MRL story about laws with an internal-realist-like metaphysics according to which different explanatory domains come with their own ontologies between which no supervenience relations hold. That sort of a view would, of course, amount to a violation of supervenience. But the point in the main text stands. Namely, BBS (unlike non-Humean pluralism) secures projectibility for the special sciences without requiring the rejection of supervenience.

It is also worth remarking here that the supervenience of (say) rabbits upon (say) quantum fields will not be visible (or, for that matter, stateable) in a Best System for ecology framed in terms of ecological kinds; for an ecological Best System will lack the kind terms necessary to say anything at all about quantum fields. On the other hand, the supervenience of rabbits on quantum fields will also be invisible from the point of view of the Best System formulated in terms of physics, which (of course) lacks the requisite higher level vocabulary. To actually see bridge principles linking the kinds of one level with the kinds of another, one needs a Best System formulated with both sets of vocabulary available. While, of course, this can happen (as it does with statistical mechanics, biochemistry, and more), there is no reason to insist on the inclusion of such bridge principles as an acceptability condition for any particular Best System.

the top: non-Humeans populate the higher-level with “genuine” movers and shakers — entities not reducible to lower-level generalizations — that make the higher-level generalizations hold, irrespective of what happens below at subvenient levels. For Albert and Loewer, the answer comes from below: proponents of this view build in to the physics the necessary ingredients to get rid of conspiratorial higher-level behavior, so that it becomes likely that the universe unfolds in ecological, economic, etc., ways. The problems with these views naturally raise the question: is there instead a way to rid the world of higher level conspiracies without adding entities?

We doubt it. We believe a certain amount of ‘conspiratorial’ behavior is virtually inevitable if we demand that the higher-level is supervenient upon and yet autonomous from the higher level but refuse to add anything designed to nomicallly ensure cooperation between levels. That said, we believe the costs of accepting conspiracy are not as high as have been thought and that BBS offers a way of looking at the issue that significantly blunts the worry. We have also argued that the costs of solving the problem by the non-Humean and Albert-Loewer methods are higher than their proponents think. A cost-benefit analysis is called for, therefore, in which we evaluate how bad the original problem is.

To begin, we note that observers are tempted to posit conspiracies when events (especially long runs of events) occur that are improbable relative to a chance. This means that the conspiracy doesn’t get off the ground in the absence of a chance. Of course, all manner of occurrences may be judged subjectively unlikely. But unless a chance makes it into the Best System for that field, and the events are unlikely with respect to that chance, then there is nothing to explain apart from the reason why certain people have a psychological state of surprise at some events and not others.

An initial observation, then, is that the Best System for fundamental physics may not posit a physical chance, contrary to what Albert and Loewer assume. If so, then, strictly speaking, there is no conspiracy problem. No chance, no conspiracy. There is only the undeniable feeling of surprise that the particles should behave that way. But surprise that the initial conditions of the universe were such as to lead to ecological patterns is like surprise at any other brute fact. One may equally be surprised that the speed of light is roughly 180,000 miles/sec. But in the absence of chance, the initial condition, like the speed of light, is not objectively unlikely. In the absence of a better theory, it’s not at all clear that one can legitimately judge different brute facts as more or less likely to be true (cf. Callender, 2004).

One might reply that statistical mechanics posits a chance over every physical initial microstate, and we’re assuming supervenience, so it automatically assigns chances to ecological generalizations, like it or not. This objection, in effect, re-asserts the Albert-Loewer position of probabilistic completeness, and it serves to remind us that there may be a cost incurred by rejecting it. The cost is that it may turn out that a certain set of positions in the foundations of statistical mechanics cannot be held consistently without adding a fundamental physical chance. Obviously we cannot assess this worry

here. However, we can point out that some (Leeds, 2003; Earman, 2006) are skeptical about the relevant positions defended by Albert (2001) for reasons independent of those we raise here. The foundations of statistical mechanics are fraught with controversy (Sklar, 1993), and there are indeed many positions — whether desirable or not, we would have to see in a prolonged discussion — that are at least consistent with the rejection of the ‘imperialist’ view of chance defended by Albert and Loewer. These remarks are no substitute for argument, of course, but they do indicate that many premises and arguments must be supplied before we get probabilistic completeness from statistical mechanics.

Suppose we put these qualms aside for the sake of argument, and assume, with Albert and Loewer, that fundamental physics does posit a chance, either over initial conditions or through a stochastic dynamics. Then via supervenience every ecological event gets a physical chance assigned to it. A conspiracy looms. However, we want to suggest that this doesn’t look so terrible from the perspective of BBS. According to BBS, and contrary to classical MRL and Albert and Loewer, ecology also has its own Best System. So we need to imagine a scenario in which there are two Best Systems in play, the ecological and the physical.

Typically these two Best Systems will disagree in the chance assignments given to any particular event. Why? The physical Best System is weighing the strengths and weakness of a vast number of events compared to the ecological Best System. Even in the unlikely event that physicists and ecologists held exactly the same standards of balance, simplicity and strength, given the two quite different mosaics, it would be amazing if they yielded the same chances for a given event. The physical Best System might have to cut corners on the chances of a few rabbits reproducing if it’s to simply describe (say) why so many galaxies are elliptical. In contrast, the ecological Best System may miss things too by not seeing how physical fluctuations might percolate up to the ecological level. Hence, given BBS, we should expect disagreement all over the place among all the different Best Systems over the chances of events. There will be some pressure toward convergence: the frequencies of rabbits will be a small part of what physics wants to get right given that rabbits supervene upon physical particles and fields. But in general we should expect disagreement between the two Best Systems.

We can now make two (related) points.

First, assume the physical and ecological chance, found in their respective Best Systems, differ with respect to a regularity. Ecology obviously rates the chances of Malthus’s Law to be very high, so from its perspective, there is no conspiracy. Let’s assume, to get a conspiracy, that the physical chance assigns patterns in accord with Malthus’s Law a low chance. The first point, then, is that the problem is symmetric. If physics deems the ecological generalization unlikely, then so does the ecological generalization deem the physical motions unlikely from its perspective. If we insist on talk of conspiracy, there’s no good

way of saying whether the particles are conspiring to behave like rabbits or the rabbits are conspiring to obey physics.¹⁴

Second, notice something already suggested above: ecology's Best System will see other conspiracies from its perspective. Why do the rabbits fall down rather than up? Why do the rabbits' speeds attain a maximum value where they do? These are questions with answers in physics and physiology, but (we assume) no answer in the Best System crafted from ecological kinds. How do the rabbits know to obey physics and biology, the ecologist might wonder.

The lesson here is that there is a symmetry among Best Systems with respect to conspiracies. The patterns of one Best System relative to a set of kinds K look conspiratorial to the Best System relative to the set of kinds K' when they are talking about the same objects. There is no reason to privilege the physical system and its worry about conspiracy over the ecological one. On a broadly Humean theory like BBS, there is no distinguished level where the 'real' laws and causes operate. The ecological events cause other ecological events just as much as the physical events cause one another. Although there is an ontological asymmetry (rabbits supervene on physical objects and not vice versa), the Best System for ecology is no less true for that, nor its laws or causes less genuine. For these reasons, there is no way to step back and say which conspiracy is the real one. Ecologists must 'solve' the conspiracy 'problem' the same way physicists do — by accepting that some of the total behavior will be a brute fact in the sense that it fails to be likely according to their useful but proprietary laws and chances.

It might appear that we have multiplied conspiracy problems instead of solving one. However, we are tempted by the opposite conclusion: if every Best System has a conspiracy problem, then no Best System does. It seems that we are always going to be stuck with some unaddressed projectible regularities at some levels that, in the right frame of mind, feel conspiratorial. While we cannot be sure without knowing more about the details of the relevant systems and events than anyone now does, we strongly doubt that it is possible to rid every Best System of every putative conspiracy by building in chances that make every conceivable pattern likely. For any regularity highly weighted by some system, there is surely another, even if framed in terms of gruesome kinds, that does not weigh that generalization (when translated into its own kinds) highly.

Even if it were possible, abolishing "conspiracy" (i.e., abolishing all projectible regularities not addressed in some particular Best System) in the way contemplated is at odds with BBS's recognition of scientists' appeal to local standards when developing Best Systems for their fields (cf. Kitcher, 2003). Given that standards are local, it is unsurprising that there are certain behaviors not made likely according to the Best System for a fixed set of kinds K . To insist that this not happen, to have all the Systems give high probabilistic weight to each other's patterns in a delicately balanced harmony, is to ask,

¹⁴Do the different values of chance of an event by different Best Systems signal that the theory is inconsistent? No: as in the account of chance of Hoefer (2007), where chances are relative to set-up, and the set-ups differ, here chances are relative to different Best Systems.

implausibly, that every system be sensitive to all interests outside its scope. We doubt strongly that this can be made (or should be expected) to work, and so see “conspiracy” as virtually inevitable.

On the other hand, from our point of view, the existence of projectible (e.g., special science) regularities not addressed by particular (e.g., fundamental physical) Best Systems seems less like a problem and more like a predictable fact of life in a complex world. This situation would indeed look problematic if the discovery of higher-level projectibility required the postulation of extra-Humean elements or probabilistic vindication at the lower-level. But BBS dispenses with these requirements. Adopting BBS means, among other things, refusing to single out a Best System as the unique locus of projectibility. Having done that, it will be no surprise to discover projectibility outside the scope of any one System. In any case, such a discovery should not motivate fantasies of conspiracy.

6 Conclusion

The existence of (projectible) special science laws not guaranteed by the fundamental laws can seem puzzling from certain points of view. However, we have argued that much of the puzzle can be rendered innocuous, especially once we have adopted an adequate framework for thinking about laws. Indeed, given our conception of laws, it would be *more* puzzling if there were *not* high level laws not guaranteed by the fundamental laws. Therefore, we suggest that special science laws should not be taken as occasions for the promulgation of conspiracy theories.¹⁵

References

- Albert, D. Z. (2001). *Time and Chance*. Harvard University Press, Cambridge, Massachusetts.
- Callender, C. (2004). Measures, explanation and the past: Should ‘special’ initial conditions be explained? *British Journal for the Philosophy of Science*, 2004–55(2), 195–217.
- Callender, C. (2010). The past histories of molecules. In C. Beisbart and S. Hartmann, editors, *Probabilities in Physics*. Oxford University Press, Oxford.
- Carroll, J. (1994). *Laws of Nature*. Cambridge University Press, Cambridge.
- Cartwright, N. (1989). *Nature’s Capacities and Their Measurement*. Clarendon Press, Oxford.

¹⁵We are grateful to an audience at the TilPS Reduction and the Special Sciences conference, and to three referees for this journal, for helpful discussion that has considerably improved this paper.

- Cartwright, N. (1999). *The Dappled World: A Study of the Boundaries of Science*. Cambridge University Press, Cambridge.
- Cohen, J. and Callender, C. (2009). A better best system account of lawhood. *Philosophical Studies*, **145**(1), 1–34.
- Dupré, J. (1993). *The Disorder of Things: Metaphysical Foundations of the Disunity of Science*. Harvard University Press, Cambridge, Massachusetts.
- Dupré, J. (2001). *Human Nature and the Limits of Science*. Oxford University Press, New York.
- Earman, J. (1986). *A Primer On Determinism*. Dordrecht Reidel, Dordrecht, Holland.
- Earman, J. (2006). The ‘past hypothesis’: Not even false. *Studies in History and Philosophy of Science*, **37**, 399–430.
- Fodor, J. A. (1997). Special sciences: Still autonomous after all these years (A reply to Jaegwon Kim’s “Multiple Realization and the Metaphysics of Reduction”). In J. E. Tomberlin, editor, *Philosophical Perspectives*, **11: Mind, Causation, and World**, pages 149–163. Ridgeview Press, Atascadero, California.
- van Fraassen, B. C. (1989). *Laws and Symmetry*. Clarendon, Oxford.
- Giere, R. (2006). *Scientific Perspectivism*. University of Chicago Press, Chicago.
- Ginzburg, L. R. and Colyvan, M. (2004). *Ecological Orbits: How Planets Move and Populations Grow*. Oxford University Press, New York.
- Goodman, N. (1954). *Fact, Fiction, and Forecast*. Harvard University Press, Cambridge, Massachusetts.
- Halpin, J. F. (2003). Scientific law: A perspectival account. *Erkenntnis*, **58**.
- Hoefer, C. (2007). The third way on objective probability: A sceptic’s guide to objective chance. *Mind*, **116**(463), 549–596.
- Kim, J. (1989). Mechanism, purpose, and explanatory exclusion. In J. E. Tomberlin, editor, *Philosophical Perspectives 3, Philosophy of Mind and Action Theory*, pages 77–108. Ridgeview Publishing, Atascadero, California. Reprinted in Kim (1993), 237–264.
- Kim, J. (1993). *Supervenience and Mind: Selected Philosophical Essays*. Cambridge University Press, New York.
- Kitcher, P. (2003). *Science, Truth, and Democracy*. Oxford University Press, New York.
- Leeds, S. (2003). Foundations of statistical mechanics: Two approaches. *Philosophy of Science*, **70**(11), 126–144.

- Lewis, D. (1973). *Counterfactuals*. Basil Blackwell, Oxford.
- Lewis, D. (1983). New work for a theory of universals. *Australasian Journal of Philosophy*, **61**, 343–377.
- Lewis, D. (1986a). *On the Plurality of Worlds*. Basil Blackwell, Oxford.
- Lewis, D. (1986b). *Philosophical Papers, Volume II*. Oxford University Press, New York.
- Lewis, D. (1994). Humean supervenience debugged. *Mind*, **103**(412), 473–490.
- Loewer, B. (1996). Humean supervenience. *Philosophical Topics*, **24**(1), 101–127.
- Loewer, B. (2004). David Lewis' Humean theory of objective chance. *Philosophy of Science*, **71**, 1115–1125.
- Loewer, B. (2009). Why there *Is* anything except physics. In J. Hohwy and J. Kallestrup, editors, *Being Reduced*. Oxford University Press, Oxford.
- Oppenheim, P. and Putnam, H. (1958). Unity of science as a working hypothesis. In H. Feigl, M. Scriven, and G. Maxwell, editors, *Minnesota Studies in the Philosophy of Science*, volume II. University of Minnesota Press, Minneapolis.
- Psillos, S. (2002). *Causation and Explanation*. Acumen, Chesham, UK.
- Putnam, H. (1987). *The Many Faces of Realism*. Open Court, LaSalle, Illinois.
- Quine, W. V. O. (1970). *Philosophy of Logic*. Harvard University Press, Cambridge, Massachusetts.
- Roberts, J. (1998). 'Laws of Nature' as an indexical term: A reinterpretation of Lewis's best-system analysis. *Philosophy of Science*, **66**, **Supplement**, S502–S511.
- Schrenk, M. (2007). Can capacities rescue us from ceteris paribus laws? In M. Kistler and B. Gnassounou., editors, *Dispositions in Philosophy and Science*. Ashgate Publishing.
- Schrenk, M. (2008). A theory for special sciences laws. In H. Bohse, K. Dreimann, and S. Walter, editors, *Selected Papers Contributed to the Sections of GAP.6, 6th International Congress of the Society for Analytical Philosophy*. Mentis, Paderborn.
- Sklar, L. (1993). *Physics and Chance*. Cambridge University Press, Cambridge.
- Taylor, B. (1993). On natural properties in metaphysics. *Mind*, **102**, 81–100. Reprinted in *The Philosopher's Annual XVI* (1993) 185–204.
- Teller, P. (2009). Reduction, emergence, and truth. *Erkenntnis*. this volume.
- Woodward, J. (2000). Explanation and invariance in the special sciences. *British Journal for the Philosophy of Science*, **51**, 197–254.