

Tractability and Laws

Isaac Wilhelm

(penultimate draft; please cite final version in *Synthese*)

Abstract

According to the Best System Account (BSA) of lawhood, laws of nature are theorems of the deductive systems that best balance simplicity and strength. In this paper, I advocate a different account of lawhood which is related, in spirit, to the BSA: according to my account, laws are theorems of deductive systems that best balance simplicity, strength, and also calculational tractability. I discuss two problems that the BSA faces, and I show that my account solves them. I also use my account to illuminate the nomological character of special science laws.

1 Introduction

According to the Best System Account (BSA) of lawhood, L is a *law of nature* if and only if L is a theorem of each true deductive system of science which strikes the best balance between simplicity and strength (Lewis, 1973, p. 73).¹ Simplicity is a syntactic feature of deductive systems: a system with redundant axioms, for instance, is less simple than a system whose axioms are logically independent. Strength is measured in terms of possibilities: roughly, if deductive system A rules out more possibilities than deductive system B —that is, if A classifies more possible worlds as nomologically impossible—then A is stronger. These two theoretical virtues pull against each other: simpler systems tend to be less strong, and stronger systems tend to be less simple. Consider the systems which strike the best balance

¹For brevity, I will not discuss probabilistic laws here. For an account that does, see (Lewis, 1994).

between these competing virtues, and take the theorems which they have in common. Those, according to the BSA, are the laws.

There are many reasons to like the BSA. It promises to distinguish lawful regularities from non-lawful regularities: the former are implied by the axioms of the best systems, while the latter are not. It also explains an intuition—call it the ‘Humean intuition’—which many scientists and philosophers share: the laws are nothing more than concise codifications of all the non-modal facts. And the BSA accounts for why laws are important to science: to discover the non-modal facts, one must discover the laws (Loewer, 2004, pp. 188-189).

But the BSA faces at least two problems. Absent certain further assumptions, it implies that all regularities are laws, including those which are clearly not (Lewis, 1983, p. 367). And it implies that many intuitively non-nomological statements about initial conditions are laws (Hall, 2015, pp. 269-271).

To address these problems, I propose a different account of lawhood, related in spirit to the BSA. According to my account, laws are theorems of true deductive systems which best balance simplicity, strength, and also calculational tractability. The best systems, on this view, are not just simple and strong. The best systems also facilitate the sorts of calculations that scientists seek to perform.

In Section 2, I present my account of lawhood. In Section 3, I motivate my account by discussing the two problems that the BSA faces, and by showing that my account solves them. Finally, in Section 4, I discuss a bonus of my account: it illuminates the nomological character of special science laws.

2 A New Account of Lawhood

2.1 The BUSA

My account of lawhood invokes a new theoretical virtue: calculational tractability. The basic idea is this: to be a law is, in part, to facilitate the calculations that scientists seek to perform. And that is pretty intuitive. Laws play a crucial role in all calculations of scientific phenomena. Arguably, calculation is the primary activity for which laws are used. So it makes sense to think that in order for a statement to be lawful, it ought to support tractable calculations.

Here are the details. Let X and Y be deductive systems. X is more ‘calculationally tractable’ than Y just in case X is, overall, more computationally useful than Y when it comes to performing numerical integrations, estimating infinite series expansions, constructing idealized models of phenomena, approximating exact solutions to equations of motion, and so on. The overall computational utility of a system is determined by many factors: for instance, the speed with which Turing machines perform computations using that system, and the number of parameters which that system allows to vary freely. In the sections to come, I discuss a couple of these factors in some detail. But for now, the important point is simply that comparisons of computational utility hinge on the many different ways that systems can help facilitate computations.

My proposed account of lawhood—call it the ‘Best-Utility System Account’, or BUSA—is as follows. Rank deductive systems according to the balance they strike between strength, calculational tractability, and simplicity. According to the BUSA, L is a *law of nature* if and only if L is a theorem of each deductive system that is best on this ranking.

For example, let W be the system whose axioms are the three laws of Newtonian mechanics and Newton’s law of gravity. When supplemented with sentences about initial conditions, the axioms of W entail sentences like \mathfrak{s} : “The gravitational effect of the Sun

on the Earth is such-and-such at time t .” Let W' be the deductive system whose axioms are the axioms of W plus an additional axiom which expresses the fact—derivable from the other axioms—that in a two-body system, the acceleration of one body is proportional to the mass of the other body. Let Y be the deductive system whose axioms are the true sentences about all possible gravitational interactions among two particles. So the axioms of Y are sentences of the form “If one particle is here and another particle is there, then the gravitational effect of the former on the latter is thus-and-so.” When supplemented with sentences about initial conditions, suitably large collections of these axioms entail sentences about macroscopic phenomena, such as \mathfrak{s} .

The BUSA correctly implies that W is the best deductive system. To see why, note that W , W' , and Y are equally strong: they all imply the same sentences, and so they all rule out the same possible worlds. Both W and W' , however, are more computationally tractable than Y : Newton’s equations are far more useful than a list of particular facts about gravitational attraction among particles when integrating trajectories, constructing models of intergalactic interaction, and so on. Moreover, both W and W' are simpler than Y too, since they contain fewer predicates and fewer basic axioms. So overall, both W and W' strike a better balance between strength, calculational tractability, and simplicity, than Y strikes. In addition, though W and W' are more-or-less equally computationally tractable,² W is simpler than W' : for W' contains an additional, redundant axiom.³ So overall, W strikes a better balance between strength, calculational tractability, and simplicity, than W' strikes. Therefore, according to the BUSA, W is the best deductive system. So the BUSA gets the right result.

Calculational tractability is an extremely important theoretical virtue. Strong deductive systems rule out lots of possible worlds. But strong deductive systems may not allow us

²I say ‘more-or-less’ because strictly speaking, W' is marginally more computationally tractable than W . The greater tractability of W' comes from the fact that its additional axiom occasionally allows for slightly shorter computations. But as will become clear, the small increase in tractability allotted to systems which include that additional, redundant axiom is outweighed by the simplicity allotted to systems which leave it out.

³The additional axiom in W' is redundant because it can be derived from Newton’s equations.

to calculate, in any detail, which worlds are ruled out. Hence the importance of calculational tractability. If a system is both strong and also calculationally tractable, then it does more than just eliminate various possible worlds. In virtue of its tractability, it also gives us the tools to determine which worlds are eliminated.

One might think that calculational tractability is best understood as a kind of simplicity: one way for a theory to be simple, perhaps, is to be calculationally tractable. I have reservations about understanding simplicity in this way. For intuitively simple theories can, I think, be utterly intractable. As an example, suppose all particle motions can be encoded in the digits of an absolutely massive number: the binary expansion of $(7 \uparrow\uparrow\uparrow\uparrow 9) - (3 \uparrow\uparrow\uparrow\uparrow 8)$, say.⁴ Then there may well be a simple way to state the axioms for a candidate best system: just say that the particle motions are encoded in $(7 \uparrow\uparrow\uparrow\uparrow 9) - (3 \uparrow\uparrow\uparrow\uparrow 8)$ in thus-and-so ways.⁵ But this system would be utterly intractable, since it would take eons to actually extract the motions of particles from facts about this massive number.

But ultimately, the BUSA is compatible with the view that calculational tractability is a kind of simplicity. Proponents of this view should read the BUSA as saying that when ranking theories, one should balance strength against *one* measure of simplicity—namely, calculational tractability—as well as against other measures (the number of axioms, the number of predicates, and so on). So understood, the BUSA is more like a precisification of the BSA than an utterly new account of lawhood.⁶ The reader is welcome to adopt this way of understanding the BUSA.

⁴The ‘ \uparrow ’, which is Knuth’s up-arrow notation, represents a generalization of exponentiation. For example, $2 \uparrow 4 = 2^4 = 16$, $2 \uparrow\uparrow 4 = 2 \uparrow (2 \uparrow (2 \uparrow 2)) = 2^{2^{2^2}} = 2^{2^4} = 2^{16} = 65536$, $2 \uparrow\uparrow\uparrow 4 = 2 \uparrow\uparrow (2 \uparrow\uparrow (2 \uparrow\uparrow 2)) = 2 \uparrow\uparrow (2 \uparrow\uparrow 4) = 2 \uparrow\uparrow 65536 = 2 \uparrow (2 \uparrow \dots)$, where the ‘ \dots ’ contains a total of 65534 instances of ‘2’. So a number like $(7 \uparrow\uparrow\uparrow\uparrow 9) - (3 \uparrow\uparrow\uparrow\uparrow 8)$ is unimaginably massive. For more discussion of this notation, see (Knuth, 1976).

⁵The encoding consists of a series of metasemantic rules which map (i) digits in various places of the binary expansion of $(7 \uparrow\uparrow\uparrow\uparrow 9) - (3 \uparrow\uparrow\uparrow\uparrow 8)$, to (ii) the positions of particles at various times. These metasemantic rules may be highly gerrymandered and complex, of course. But the complexity of these metasemantic rules does not increase the complexity of the corresponding deductive system, just as the complexity of the interpretation function used to construct a set-theoretic model of a theory in first-order logic does not increase the complexity of that theory.

⁶And so understood, the BUSA goes some way towards addressing the concerns raised by Woodward about the definition—or really, lack of a definition—of simplicity (2014). For the BUSA contains a more precise characterization of simplicity than the BSA does.

The BUSA is based on a particular procedure for ranking deductive systems: in a single step, compare each system's (i) degree of strength, (ii) degree of tractability, and (iii) degree of simplicity. There are other procedures for ranking systems on the basis of their strength, tractability, and simplicity, however. Here is one alternative: in a first step, compare each system's degree of strength with its degree of tractability; in a second step, compare each system's score in the first step with its degree of simplicity. And here is another alternative: in a first step, compare each system's degree of strength with its degree of simplicity; in a second step, compare each system's score in the first step with its degree of tractability.⁷

I prefer the BUSA's one-step procedure over these alternatives. It is attractively symmetrical to compare strength, tractability, and simplicity, all at once. It is unattractively asymmetrical to compare two of those virtues against each other, and only then bring in the third.

For my purposes here, however, none of this matters much. In this paper, my goal is to develop an attractive account of lawhood in which calculational tractability plays a central role. The BUSA is just such an account. Other accounts, which differ over the details of the relevant ranking procedure—and which differ from the BUSA in other ways too—should be explored as well. That is a good topic for future work.⁸

The BUSA is a complete, stand-alone account of lawhood. But it can also be used to supplement other accounts of lawhood which are similar, in spirit, to the BSA. For instance, the BUSA can be combined with the view that laws are defined relative to classes of basic kinds and predicates (Cohen & Callender, 2009, p. 21). On that version of the BUSA, calculational tractability is 'immanent', where a theoretical virtue of a system is 'immanent' just in case it is measured relative to the system's language.⁹

⁷Thanks to two anonymous reviewers for pressing me on which, of these many different procedures, is best for the purposes of formulating the BUSA.

⁸It would be interesting to explore, for instance, whether these different ranking procedures tend to produce more-or-less the same rankings of deductive systems – and therefore, tend to agree on which statements count as laws. I suspect that they might, since I see no reason to think otherwise; and they all seem to imply that W is a better deductive system than both W' and Y .

⁹According to another version of the BUSA, calculational tractability is 'transcendent': like the accounts of simplicity and strength offered by Lewis (1973) and Loewer (2004), the calculational tractability of a

The BUSA is similar to the accounts of lawhood proposed by Hicks (2017) and Dorst (2018).¹⁰ Those accounts, and my own, agree that lawhood is a matter of balancing collections of theoretical virtues; just more virtues, and balanced in more complicated ways, than in the traditional BSA. So the BUSA is very friendly to those other accounts.

In at least one respect, however, the BUSA and those other accounts are quite different: as far as I can tell, the BUSA allows deductive systems to be computationally tractable in more ways than those other accounts allow.¹¹ For example, Hicks defines strength and modularity using the notion of a quasi-isolated system, where a system is quasi-isolated just in case it is relatively unaffected by outside influences (2017, p. 997). Some characterizations of computational utility invoke similar notions, such as the characterization discussed in Section 3.2. But other characterizations do not, such as the characterization discussed in Section 2.2. Each of these characterizations of computational utility is partially constitutive of the theoretical virtue of calculational tractability – and therefore, of lawhood. Because of that, the account of lawhood provided by the BUSA may be somewhat more general than those other accounts, at least when it comes to tractability.

2.2 The Objectivity of the BUSA

One might worry that lawhood, when understood in terms of the BUSA, is far too subjective. What it is to be computationally useful—to admit of tractable numerical integrations, estimations of infinite series, the construction of idealized models, approximating

system is evaluated independently—or at least, somewhat independently—of that system’s language. I take theoretical virtues to be transcendent, because immanent approaches tend to imply an overabundance of (relativized) laws: they often imply that for every regularity r , there is a choice of basic kinds and predicates such that relative to those kinds and predicates, r is nomological. But the choice between immanence and transcendence is not forced by the BUSA. Fans of immanent accounts of theoretical virtues can still subscribe to the BUSA, since the BUSA is neutral with respect to whether theoretical virtues are immanent or transcendent.

¹⁰There are also similarities between the BUSA and the approach to lawhood discussed by Urbaniak and Leuridan (2018). I discuss those similarities—and some differences as well—in Section 3.1.

¹¹Similarly, the BUSA allows deductive systems to be computationally tractable in more ways than Braddon-Mitchell’s account (2001) allows.

solutions, and so on—may seem to depend on human psychology.¹²

But that is not so. There are plenty of ways to understand computational utility—and so, calculational tractability—that do not depend on the capacities of human minds. For instance, computational utility could be characterized, to a first approximation, using the resources of computational complexity theory. Say that deductive system X is *more computationally useful* than deductive system Y if and only if the worst case running time of a maximally efficient Turing machine that carries out computations using the axioms of X is less than the worst case running time of a maximally efficient Turing machine that carries out computations using the axioms of Y .¹³ Call this the ‘Worst Case Account’ of computational utility.

A toy example will illustrate how the Worst Case Account may be used to compare deductive systems. Suppose that the universe consists of a single sphere traveling at constant velocity. Furthermore, suppose that the universe is a one-dimensional lattice, with the sphere jumping from point to point. And suppose it is a law that objects always travel at a particular constant velocity, jumping two spatial units to the right for every one unit of time. Finally, to keep things simple, suppose that this lattice has been assigned a coordinate system. That will help us calculate the sphere’s later states based on its earlier states.¹⁴

One systematization of the facts about the sphere’s location—call it system A —consists of all triples of the form $\langle t, x, p \rangle$ where t and x are natural numbers represented in binary notation, p is either 0 or 1, and the following two conditions obtain.

¹²A related but distinct worry: exactly which calculations are relevant for determining the degree to which a deductive system is tractable? Calculations of observables, or calculations of all quantities posited by a particular theoretical system, or calculations of something else entirely? The answer: for each collection of relevant calculations, there is a corresponding version of the BUSA. My preferred version of the BUSA takes all calculations—of all quantities, properties, constants of nature, and so on—to be relevant. But it is beyond the scope of this paper to investigate whether this version of the BUSA is better than those others.

¹³The ‘worst case running time’ of a Turing machine T is the maximum number of steps (that is, the time) required for T to halt on an input string of length n . So if T never requires more than n^2 steps to halt on an input of length n , but always requires a number of steps asymptotically close to n^2 , then the worst case running time of T is n^2 . A ‘maximally efficient Turing machine’ is a machine whose worst case running time exhibits the slowest asymptotic growth in n (for a given problem). For further discussion, see (Cormen et al., 2009, pp. 27-29).

¹⁴It is possible to drop this assumption by doing everything in terms of relative positions in space and time. But then the example becomes needlessly complicated.

- (i) Triple $\langle t, x, 1 \rangle$ is in A if and only if the sphere is at location x at time t .
- (ii) Triple $\langle t, x, 0 \rangle$ is in A if and only if the sphere is not at location x at time t .

In virtue of these two conditions, A encodes all facts about the position of the sphere at all times.¹⁵

A different systematization of the facts about the sphere’s location—call it system B —contains just one axiom. That axiom is a purely formal version of the following sentence: the sphere moves two spatial units to the right per temporal unit. In other words, the sole axiom of B says that for every unit of time, the sphere jumps two points rightward.

According to the Worst Case Account of computational utility, B is extremely computationally useful. To see why, let T_B be a Turing machine which, for any input of time elapsed t_0 , outputs the string which is just like t_0 except with an extra 0 placed to the right. This is the machine equivalent of the sole axiom of B .¹⁶ T_B only requires $n + 1$ steps to complete its code, where n is the length of the input string: T_B begins on the leftmost numeral of the input string, moves to the right n times, writes a 0 in the first blank space, and then halts. Therefore, T_B efficiently uses the machine-theoretic equivalent of the axioms of B to compute the relative location of the sphere.

A is not nearly so computationally useful, however. Let T_A be any Turing machine which, given an input t_0 of length n , outputs whichever x is the second coordinate in a triple that has t_0 as its first coordinate and 1 as its third. To do this, T_A must search through A for the relevant triple. That search will always take more than $n + 1$ steps. For it takes n steps just to read the first coordinate of any given triple, it takes at least one step to move past the second coordinate, and it takes another step to read the third coordinate. So T_A will always require at least $n + 2$ steps to output the right location. And depending on how the list comprising A is organized, that search could require a number of steps that is exponential—or factorial, or worse—in n .

¹⁵ A also includes an “and that’s all” clause – that is, a clause specifying that there are no positions or times other than those described by the triples.

¹⁶Putting a 0 to the right of a string is equivalent, in binary notation, to multiplying that string by two.

In brief, since A does not encode the fact that the sphere moves two spatial units for every temporal unit, T_A cannot take advantage of the shortcuts that T_B can. A does not ‘know’ that there is a tractable relationship between the position of the sphere and the time. So T_A can only scan a massive—indeed, possibly infinite—list.

The BUSA correctly implies that it is a law of nature that objects travel at a constant velocity of two spatial units to the right for every temporal unit. To see why, note that as was just shown, the Worst Case Account implies that B is far more computationally useful—and so, far more calculational tractable—than A . Of course, A is far stronger than B : A rules out all possible sphere worlds except the actual one, while B only rules out worlds in which the sphere jumps at a different velocity. In addition, however, A is far less simple than B : the axioms of A form a massive—possibly infinite—list of every fact about the sphere’s location; B contains just one axiom. So overall, B achieves the better balance between strength, calculational tractability, and simplicity. Assuming that no other systems are better than B , it follows that B is best.¹⁷ Therefore, the theorems of B —like the axiom that objects travel at a constant velocity of two spatial units to the right for every temporal unit—are laws.

This toy example does more than just illustrate the Worst Case Account of computational utility. It also provides a reason to prefer the BUSA to the BSA. For the BSA does not seem to imply that the axiom of B is a law. A is stronger than B , and B is simpler than A . So it is plausible that on the BSA’s method of ranking deductive systems, A and B come out equally good. Only B , of course, implies that objects travel at a constant velocity

¹⁷The situation is somewhat more complicated than I have suggested. To see why, let B^* be the deductive system which results from supplementing B with a specific triple from A : whichever triple describes the position of the sphere at the initial time. Clearly, B^* and A are equally strong, but B^* is far simpler. In addition, B^* is much stronger than—but only slightly less simple than— B . And if computational utility is only determined by worst case running time—that is, if computational utility is measured as in the Worst Case Account only—then (i) B and B^* are equally calculational tractable, and (ii) both are much more calculational tractable than A . So on balance, if computational utility is only determined by worst case running time, then B^* —rather than B —seems like the best deductive system. To avoid this somewhat unintuitive result, I suggest measuring computational utility using more than just the notions in the Worst Case Account. In particular, the notions invoked in an account of computational utility described in Section 3.2—called the ‘Variational Account’—should be used to measure the computational utilities of the systems A , B , and B^* . Then it follows that B and B^* are not equally calculational tractable: B is far more calculational tractable than B^* . And so plausibly, B is indeed the best deductive system (thanks to an anonymous reviewer for discussion).

of two spatial units to the right for every temporal unit. And the BSA requires laws to be deductive consequences of *each* of the best deductive systems. Therefore, according to the BSA, it may not be a law that objects travel at that particular velocity.

Many different characterizations of computational utility are partially constitutive of tractability, because scientists have many ways of performing computations. The Worst Case Account is one of them, since the Worst Case Account can be used to quantify the efficiency of the computer programs that scientists actually employ. But other characterizations are constitutive of calculational tractability too. The computational utility of any particular system is determined by an amalgam of many different notions: worst case running time, perhaps data compressibility, a characterization discussed in Section 3.2, and undoubtedly many others.

To put the point a bit more carefully, the computational utility of a deductive system is given by balancing each of these notions against one another. Taken on their own, each such notion partially—but not fully—determines the overall computational utility of the system in question. Each contributes to, but does not completely settle, the given system’s degree of computational utility. So a given system’s worst case running time partially, but not fully, determines that system’s computational utility. Other features of that deductive system, like a feature discussed in Section 3.2, contribute to that system’s computational utility as well. And altogether, the most computationally useful deductive system is the system which best balances all these different computationally-relevant features: worst case running time, compressibility, the feature in Section 3.2, and so on.¹⁸

Note that all these characterizations of computational utility are perfectly objective. So the BUSA provides a perfectly objective analysis of lawhood. And the BUSA uses that objective analysis to get many cases right, such as the toy case of the sphere world.

¹⁸For this reason, when I formulated the Worst Case Account using a biconditional, I was making a simplifying assumption. To keep the example concise, I assumed that no other features of deductive systems *A* and *B* contribute, in a substantial way, to those systems’ computational utility. As explained in footnote 17, however, those systems’ computational utilities are indeed affected by other features, such as the features described by the Variational Account in Section 3.2.

Before moving on, it is worth making a general point about the notions invoked in the BUSA. Follow standard practice in the literature on best system accounts of lawhood, I will not provide fully rigorous characterizations of each and every notion that the BUSA invokes: simplicity, strength, tractability, various kinds of computational utility, and balance. For the most part, in this paper, I stick to somewhat coarse-grained characterizations of these notions.

And for at least three reasons, that is perfectly fine. First, the notions invoked in the BUSA can indeed be precisely characterized. For lack of space, I will not explore that here: it is an interesting topic for future work. Second, even without fully rigorous characterizations of simplicity, strength, tractability, certain kinds of computational utility, and balance, the BUSA still provides an illuminating account of lawhood. It still tells us a great deal about what makes some statements nomological and other statements not. Third, the BUSA does lots of useful philosophical work, even given somewhat coarse-grained characterizations of the notions that it invokes. As shown in Section 3, it can be used to solve two problems facing the BSA; as shown in Section 4, it can be used to account for the nomological character of special science laws.

So there is certainly much more to say about simplicity, strength, measures of computational utility, calculational tractability, and how all these theoretical virtues balance against each other. But I need not do so here. For present purposes, it is enough to provide reasonably illuminating characterizations of the notions that the BUSA invokes, and to show how those notions help us make progress on some philosophical problems.

3 Two Solutions

In this section, I discuss two problems that the BSA faces, and I summarize the costs of some solutions that have been proposed. Then I show that the BUSA avoids the problems without incurring the costs.

3.1 The Trivialization Problem

Any deductive system may be formalized in a trivially simple way without sacrificing strength. To see how, and to see why this is a problem, let S be the deductive system for our world whose axioms express all and only the propositions which actually obtain. Let F be a primitive predicate that is true of all and only things at worlds where S holds (Lewis, 1983, pp. 367-368). Consider the deductive system S_F whose lone axiom is $\forall xFx$. S_F is just as strong as S : both systems rule out exactly the same possible worlds. But S_F is simpler: it contains only one axiom. In addition, plausibly, there are no other comparably strong and simple systems. So S_F is best. And so according to the BSA, every contingent regularity—which S_F implies—is a law.¹⁹

This is a problem for the BSA—call it the ‘trivialization problem’—for two reasons. First, in S_F , every contingent regularity in the actual world is a law because every such contingent regularity is a true proposition. And that seems wrong: plenty of contingent regularities are not nomological. Second, this system looks nothing like the laws that scientists actually propose. No physicist would accept S_F as the fundamental physical theory of the world.

To avoid this problem, Lewis revised the BSA as follows (1983, p. 368). Consider only those deductive systems whose predicates refer to perfectly natural properties; call this the ‘naturalness constraint’ on deductive systems. Rank those systems according to how well they balance simplicity and strength. Laws are theorems of the systems which score best on this ranking. So for S_F to even be eligible to enter the ranking in the first place, F must express a perfectly natural property. Since it does not, S_F is not eligible to be a best system.

Lewis’s solution has at least three costs. First, it invokes a metaphysically loaded notion: the notion of perfect naturalness. One of the attractions of the BSA was its almost

¹⁹Following Lewis, in this formulation of the problem, I take propositions to be classes of worlds, and I take implication to be the subclass relation between propositions (1983, p. 367). So the proposition expressed by one sentence implies the proposition expressed by another just in case the first proposition is a subclass of the second.

exclusive reliance on clear, explicit principles of scientific inquiry. But naturalness constraints are not part of scientific theorizing, at least not explicitly.

Second, Lewis's solution assumes that there are maximally natural—that is, perfectly natural—properties. But perhaps naturalness is not well-founded: perhaps properties get more and more natural, without any bottom layer of perfect naturalness. If so, then Lewis's solution to the trivialization problem cannot be used.

Third, the naturalness constraint on deductive systems seems overly strong. Fundamental physical theories may sometimes invoke properties which are not perfectly natural (Hicks & Schaffer, 2017). It would be better to allow for the possibility of non-perfectly natural properties in fundamental laws; properties which, perhaps, are merely almost perfectly natural.²⁰

The BUSA avoids the trivialization problem while also avoiding these costs. The solution is straightforward: S_F is computationally useless. It is nearly impossible to use S_F to perform numerical integrations, estimate infinite series expansions, construct idealized models, approximate exact solutions, and so on. So plenty of deductive systems are better than S_F . Therefore, plausibly, the BUSA implies that $\forall xFx$ is not a law.

It is worth comparing this solution, to the trivialization problem, with the solution discussed in (Urbaniak & Leuridan, 2018). According to Urbaniak and Leuridan, in order for two theories to be equally strong—or as they put it, in order for two theories to have the same content—those theories must have exactly the same consequences (2018, p. 1651). This condition, along with Urbaniak and Leuridan's characterization of the consequence relation, implies that S and S_F are not equal in strength. Therefore, Urbaniak and Leuridan conclude, S_F is not – or at least not obviously – the best system, and $\forall xFx$ is not – or at least not obviously – a law.

There are some similarities between (i) Urbaniak and Leuridan's solution to the trivialization problem, and (ii) my solution based on the BUSA. Both solutions endorse something

²⁰For more criticisms of Lewis's use of naturalness in the BSA, see (Massimi, 2017).

like a derivability condition for the best deductive system. For Urbaniak and Leuridan, the derivability condition is described in terms of a particular sort of consequence relation. For the BUSA, the derivability condition is described in terms of the many different measures of computational utility.

The main difference, between my solution and the solution proposed by Urbaniak and Leuridan, concerns the associated accounts of strength. The BUSA is compatible with many different measures of the strength of a deductive system. For instance, the BUSA is compatible with the measure of strength mentioned in Section 1: roughly put, according to this measure, the strength of a deductive system is given by the class of possible worlds which that system excludes. The solution endorsed by Urbaniak and Leuridan commits to a particular measure of strength: one based on their conception of the consequence relation. That measure is different from the measure, just mentioned, based on *possibilia*.²¹

3.2 The Initial Condition Problem

Hall has observed that sometimes, the strength of a deductive system can be dramatically increased—with only a negligible decrease in simplicity—by supplementing its axioms with a sentence that describes the world’s initial conditions (2015, pp. 269-270). Consequently, the best deductive systems may include constraints on the initial conditions of the world. If so, then those constraints are nomological. Call this the ‘initial condition problem’.

An example will illustrate why this is problematic. Suppose that the world is Newtonian, and suppose it contains exactly 2^{200} particles. Let W be the same deductive system as before: its axioms are Newton’s three laws of mechanics and Newton’s law of gravity. Let I be the system whose axioms include all the axioms of W , plus an additional axiom α which states that exactly 2^{200} particles exist. I is far stronger than W : its axioms rule out all the

²¹For this reason, I prefer the BUSA’s solution to the trivialization problem over Urbaniak and Leuridan’s solution. There is much to like about the measure of strength which invokes possible worlds. And the BUSA is attractive, insofar as it is compatible with measuring strength in that way.

worlds that W 's axioms rule out, but in virtue of \mathbf{a} , its axioms also rule out worlds with a different total number of particles. Since I is only slightly less simple than W , it achieves the better overall balance of simplicity and strength. It is therefore the better deductive system. If we suppose that I is better than all other deductive systems too, then I 's theorems are laws, and so it is a law of nature that there are exactly 2^{200} particles. But that does not seem like a nomological fact about our world. That seems nomologically contingent.

This is problematic for at least two more reasons. First, laws of nature seem like they should be consistent with a fairly wide range of initial conditions. But most initial conditions are incompatible with the actual initial conditions of the world; with \mathbf{a} , say. So if the actual initial conditions are nomological, then most initial conditions are incompatible with the laws. According to I , for instance, it is nomologically impossible that there be more or less than 2^{200} particles.

Second, if the best system specifies what the initial conditions of the world are, then the laws cannot support any counterfactual whose antecedent specifies initial conditions that are different from the actual initial conditions. For instance, I cannot support any counterfactual of the form "If the universe contained exactly 10 particles, arranged thus-and-so, then..." because according to I , a universe with exactly 10 particles is nomologically impossible. If initial condition statements are nomological, then those counterfactuals are akin to counterfactuals like "If the laws of the world were different, then..." This is problematic because evidence for laws comes from experiments on restricted subsystems of the world. Those subsystems are used to formulate counterfactual claims, which are then used to test hypotheses. If a law cannot support such counterfactuals because it fails to hold in the subsystems on which we perform experiments—as is the case with \mathbf{a} —then it cannot be tested in the usual way.

To solve these problems, Hall advocates a new way of measuring the strength of a deductive system (2015, p. 271). For any candidate best system X , the axioms of X may be sorted into two groups: axioms about what initial conditions are possible – call these

‘Initial Condition Hypotheses’, or ICHs – and axioms about how those possible conditions evolve forward in time – call these ‘Dynamical Hypotheses’, or DHs. The strength of X is measured by the number of possibilities ruled out by its DHs, and the number of possibilities *not* ruled out by its ICHs. So for example, on Hall’s measure of strength, W is stronger than I : their DHs are the same, and so rule out the same possibilities; but I has one ICH—namely, α —while W has none, and so trivially, W ’s ICHs rule out fewer possibilities than do I ’s. Therefore, W is better than I , and so α is not a law.

There are two problems with Hall’s solution. First, the distinction between DHs and ICHs is not very robust. If the dynamics of a theory assumes that all particles are small, impenetrable spheres of finite but nonzero radius, for example, then no two particles can overlap in physical space. That restricts the range of initial conditions that could obtain: situations in which two particles overlap are ruled out. So it is not clear that the distinction between DHs and ICHs is as robust as Hall’s measure of strength requires. Some dynamical hypotheses are also, implicitly, initial condition hypotheses.

Second, as Hall points out, his solution conflicts with the conception of scientific inquiry embodied by the BSA (2015, p. 272). One of the principal motivations for the BSA is that it explains the Humean intuition that the laws are concise codifications of the world’s non-modal facts. So if there is a concise way to articulate—or severely restrict—the world’s initial conditions, then according to the Humean intuition, those initial conditions should be nomological.

The BUSA provides a solution to the initial condition problem while avoiding the disadvantages of Hall’s proposal. The solution relies on a particular account of computational utility: the ‘Variational Account’. According to the Variational Account, deductive system X is *more computationally useful* than deductive system Y if and only if X allows more parameters of a specific type to vary freely than Y allows.²² In other words, more computationally

²²For the sorts of reasons mentioned in footnote 18, in order to keep the following discussion simple, I formulated the Variational Account using a biconditional. Strictly speaking, free variation among parameters only partially contributes to the overall computational utility of any given deductive system.

useful systems allow for more flexibility in the values that parameters can take. For by allowing more variation in the values of parameters, those systems allow more computations to be performed.

There are many different versions of this account, one for each type of parameter. I will focus on just one here: parameters that define subsystems.²³ Suppose, to keep things simple, that our universe is Newtonian. Consider all subsystems of our universe. Several different parameters must be provided, in order to completely characterize any given subsystem: for example, the number of particles comprising the subsystem, their masses, their positions, their velocities, their accelerations, and so on. These parameters vary from one subsystem to another, so in order to uniquely define any given subsystem, these parameters must be specified.

The resulting version of the Variational Account is as follows: *X* is *more computationally useful* than *Y* if and only if *X* allows for more free variation among particle number, mass, position, velocity, and acceleration, than *Y* allows. In other words, *X* is more computationally useful than *Y* just in case *Y* fixes more values of those parameters than *X* fixes. That is, the more of those parameters a system does not pin down, the more computationally useful that system is.

The Variational Account is quite attractive. For it captures an imprecise, yet extremely plausible, view of what computational utility is. Basically, according to that view, computational utility is largely a matter of computational flexibility. If one system of axioms facilitates more calculations of the properties of more subsystems than another system of axioms facilitates, then the former is more computationally useful than the latter. So the former system of axioms is more calculationally tractable, in the sense that more calculations—for more subsystems—can be conducted using it. The Variational Account captures all this, because it takes computational utility to be a matter of allowing for free variation among the parameters with which those calculations are performed: namely, the parameters that completely

²³This version of the Variational Account is inspired by the ideas in (Hicks, 2017).

characterize any given subsystem.

For an example application of the Variational Account, consider W and I once more. Recall that while I specifies the number of particles; W does not. Neither I nor W specify those particles' masses, positions, velocities, or accelerations. So according to this version of the Variational Account, W is more computationally useful than I . And that is clearly the right result. Computations carried out using I cannot assume that there are more or fewer than 2^{200} particles: any such assumptions would immediately yield a contradiction. So W allows for more computations. And because of that, W is more computationally useful than I .

One might object that this is a bad consequence of the Variational Account. For one might claim that intuitively, I is just as computationally useful as W . After all, I can indeed be used to perform computations which assume that there are fewer than 2^{200} particles, so long as those assumptions are formulated carefully. In particular, consider the following sentence \mathfrak{t} : there are 10 particles in thus-and-so subsystem. This sentence is perfectly consistent with I . So in that sense, computations carried out using I can indeed assume that there are fewer than 2^{200} particles. Those assumptions just need to specify that the particles in question form a subsystem of the universe. As a result, I allows for just as many computations as W . In other words, intuitively, I and W are equally computationally useful. So the Variational Account is problematic, insofar as it does not respect this.

This objection, however, does not succeed. Sentences like \mathfrak{t} are consistent with both I and W ; that is certainly true. For when assumptions like \mathfrak{t} are combined with either system, the resulting collection of postulates can be used to perform useful computations. Nevertheless, other sentences—like the sentence “There are 10 particles”—are consistent with W only; they cannot be consistently combined with I . So while W and I are equally computationally useful when it comes to computations based on sentences like \mathfrak{t} , W is far more computationally useful than I when it comes to computations based on sentences like “There are 10 particles.” Therefore, overall, W is more computationally useful than I ; it is not the

case that intuitively, they are equally computationally useful. The Variational Account is actually quite attractive, insofar as it respects this.

Unlike the BSA, the BUSA correctly implies that \mathfrak{a} —the sentence which states that the number of particles in the universe is exactly 2^{200} —is not a law. Though I is stronger than W , W is more computationally tractable than I , because W is more computationally useful on the Variational Account. In addition, W is simpler, since it has fewer axioms. Therefore, W is the better deductive system. And so according to the BUSA, \mathfrak{a} is not nomological.

All this suggests another attraction of the BUSA: it can be used to explain why some restrictions on initial conditions are nomological. The nomological ones are the ones that do not detract too much from the calculational tractability of the system to which they belong. For example, take Gauss’s law of magnetism $\nabla \cdot B = 0$, where B is the magnetic field. Gauss’s law expresses a constraint on initial conditions: it says, roughly, that there are no magnetic monopoles. In a universe that obeys the laws of classical electromagnetism, Gauss’s law is a law because it does not vary across subsystems (and it is simple and strong). No subsystem contains a magnetic monopole, and so ‘the number of magnetic monopoles’ is not a parameter that varies from one subsystem to another. So according to the Variational Account, $\nabla \cdot B = 0$ does not decrease the calculational tractability of the deductive systems in which it appears. Of course, $\nabla \cdot B = 0$ dramatically increases the strength of any system containing it, and only marginally reduces any such system’s simplicity. Therefore, $\nabla \cdot B = 0$ yields a better system. That is why it belongs to the best system for classical electromagnetism, even though it expresses an initial condition constraint.

4 The BUSA and Special Science Laws

The BUSA accounts for special science laws. So far, I have only discussed the BUSA in the context of physics, for that is where the nomological status of certain generalizations is most clear-cut. But one of the strengths of the BUSA is that it counts certain special science

generalizations as nomological. Let us see how.

Many generalizations in the special sciences have the marks of lawhood. They are used to support counterfactuals, construct predictions, and offer explanations of the phenomena that the special sciences study (Fenton-Glynn, 2016, p. 274; Loewer, 2008, p. 150). But these generalizations almost always have exceptions and *ceteris paribus* conditions, which makes them seem different from the laws of physics (Fodor, 1974, pp. 108-109; Rupert, 2008, p. 349). This raises two questions. Are the generalizations typically taken to be special science laws indeed nomological? And if so, what is it that makes them laws, despite the fact that they have exceptions?

Take Malthus's law, a principle of population dynamics. According to Malthus's law, if P_0 is the size of an initial population and r is its growth rate, then *ceteris paribus*, the size of the population at time t is $P(t) = P_0 \cdot e^{rt}$ (Ginzburg & Colyvan, 2004). This law supports many counterfactuals about how populations increase, it is used to predict population growth, and it explains many salient phenomena. Of course, it has exceptions. Populations do not always grow exponentially. So is Malthus's law really a law? And if so, what makes it nomological?

The BUSA answers both questions in a satisfying way. Malthus's law really is a law because, in part, it is extremely computationally tractable. Since Malthus's law has exceptions, a deductive system that includes it may be somewhat less strong and somewhat less accurate than a system which includes, for example, a list of particular facts about population growth instead. But Malthus's law greatly increases calculational tractability, because it facilitates many different calculations of population phenomena. Equipped with only a list of facts about how specific populations in the past have changed, or with only the laws of fundamental physics, it would be extremely difficult—if not impossible—to predict how present populations will grow into the future. Malthus's law is a calculational rubric for making predictions like that, and many other predictions too. Therefore, Malthus's law provides a net increase in the tractability of any deductive system in which it appears. So it is reasonable

to suppose that Malthus’s law is a theorem of the best deductive system, and therefore, is a law.²⁴

There are many other examples of special science generalizations which count as nomological because, in large part, of how well they facilitate calculations.²⁵ For instance, take the generalizations described by the standard model of particle physics. For a variety of reasons related to the details of certain integration procedures, those generalizations do not hold exactly.²⁶ But those generalizations still count as laws, because they facilitate so many accurate computations. In other words, the nomological character of our best scientific theories derives from their calculational tractability.

As long as other special science generalizations are extremely computationally tractable—and simple and strong too—they too will qualify as laws in the BUSA. It does not matter that they admit of exceptions or *ceteris paribus* conditions, if they facilitate enough accurate calculations. Therefore, the BUSA allows us to take seriously the claim that special science generalizations have a nomological character akin to the generalizations of fundamental physics.²⁷

²⁴Of course, Malthus’s law is only accurate in circumscribed domains. It only provides a reasonably accurate description of population growth over limited periods of time, in sufficiently isolated ecosystems, and so on. So strictly speaking, there is another theoretical virtue which contributes to the nomological character of Malthus’s law: accuracy. Following Braddon-Mitchell (2001, pp. 266-267), accuracy can be analyzed in terms of data compression. Altogether, then, Malthus’s law really is a law because it helps deductive systems strike a good balance between four theoretical virtues—strength, tractability, simplicity, and also accuracy—rather than three. So to accommodate the fact that accuracy often contributes to the nomological character of special science laws, amend the formulation of the BUSA so that it balances those four virtues rather than the three virtues from earlier: in other words, the best system is the one which strikes the best balance between strength, tractability, simplicity, and accuracy.

²⁵Tractability even contributes to the nomological character of certain generalizations about chaotic dynamics. That might be surprising: for in chaotic dynamical systems, many analytic generalizations—about subtle variations in the orbital parameters of the planets, say—are utterly intractable. Nevertheless, in chaotic systems, there are tractable generalizations which seem nomological. For examples of generalizations like those, and how they are used in astronomy and in the design of spacecraft trajectories, see (Wilhelm, 2019).

²⁶For accessible summaries of those procedures, along with explanations of why quantum field theories can be regarded as special sciences, see (Fraser, 2020; Ruetsche, 2011; Williams, 2019).

²⁷This complements other accounts of special science laws. For instance, it complements the view that special science laws derive from the Mentaculus, a fundamental physical theory of the world that assigns a probability to every physically possible event (Loewer, 2008, pp. 159-160). According to the BUSA, the postulates of the Mentaculus—like the Past Hypothesis or the Statistical Postulate (Albert, 2000, p. 96)—are laws: they increase the calculational tractability of the Mentaculus by implying that special science generalizations like Malthus’s law, which are extremely computationally tractable, are highly probable.

Loewer (2009) asks: why are there any sciences apart from physics? The BUSA answers: because, in large part, of calculation. Scientists seek to discover the laws, lawhood is partially determined by calculational tractability, and the vast majority of macroscopic phenomena cannot be calculated from fundamental physics alone. Special science laws increase the tractability of the deductive systems in which they appear. For special science laws allow scientists to calculate higher-level phenomena that, otherwise, would be utterly intractable; those calculations simply could not be performed using only the laws of fundamental physics. So to discover the laws, we must discover the calculationally tractable, strong, simple sentences which describe our macroscopic world. We must pursue the special sciences.

5 Conclusion

The BUSA is quite plausible. It provides a unified response to two different problems facing the best system account of lawhood: the trivialization problem and the initial condition problem. And it illuminates special science laws.

All this success can be explained, in part, by the connection that the BUSA draws between lawhood and the principles of scientific inquiry. Simplicity and strength are just two of the theoretical virtues which scientific theories are designed to have. Calculational tractability is another, perhaps one of singular importance. For it is through calculation that theories are put to nature. It is through calculation that the world provides answers to empirical questions.

Acknowledgements

Thanks to Craig Callender, Eddy Chen, Jonathan Cohen, Chris Dorst, Véronica Gómez, Eugene Ho, Barry Loewer, Mike Hicks, Jill North, Eric Winsberg, Yip Sze Kay, audiences at the 2016 Western Canadian Philosophical Association and at the 2017 Pacific APA, two anonymous referees, and especially Jonathan Schaffer, for much helpful feedback and discussion.

References

- Albert, D. (2000). *Time and Chance*. Cambridge, MA: Harvard University Press.
- Braddon-Mitchell, D. (2001). Lossy Laws. *Noûs*, 35(2), 260–277.
- Cohen, J., & Callender, C. (2009). A better best system account of lawhood. *Philosophical Studies*, 145, 1–34.
- Cormen, T. H., Leiserson, C. E., Rivest, R. L., & Stein, C. (2009). *Introduction to Algorithms* (3rd ed.). Cambridge, MA: MIT Press.
- Dorst, C. (2019). Towards a Best Predictive System Account of Laws of Nature. *The British Journal for the Philosophy of Science*, 70, 877–900.
- Fenton-Glynn, L. (2016). *Ceteris Paribus* Laws and *Minutis Rectis* Laws. *Philosophy and Phenomenological Research*, 93(2), 274–305.
- Fodor, J. A. (1974). Special Sciences (Or: The Disunity of Science as a Working Hypothesis). *Synthese*, 28(2), 97–115.
- Fraser, J. D. (2020). Towards a Realist View of Quantum Field Theory. In S. French & J. Saatsi (Eds.), *Scientific Realism and the Quantum* (pp. 276–292). New York, NY: Oxford University Press.
- Ginzburg, L., & Colyvan, M. (2004). *Ecological Orbits*. New York, NY: Oxford University Press.
- Hall, N. (2015). Humean Reductionism about Laws of Nature. In B. Loewer & J. Schaffer (Eds.), *A Companion to David Lewis* (pp. 262–277). Malden, MA: Wiley.
- Hicks, M. T. (2018). Dynamic Humeanism. *The British Journal for the Philosophy of Science*, 69, 983–1007.
- Hicks, M. T., & Schaffer, J. (2017). Derivative Properties in Fundamental Laws. *The British Journal for the Philosophy of Science*, 68, 411–450.
- Knuth, D. E. (1976). Mathematics and Computer Science. *Science*, 194(4271), 1235–1242.
- Lewis, D. (1973). *Counterfactuals*. Malden, MA: Blackwell Publishers.

- Lewis, D. (1983). New Work for a Theory of Universals. *Australasian Journal of Philosophy*, 61(4), 343–377.
- Lewis, D. (1994). Humean Supervenience Debugged. *Mind*, 103(412), 473–490.
- Loewer, B. (2004). Humean Supervenience. In J. Carroll (Ed.), *Readings on Laws of Nature* (pp. 176–206). Pittsburgh, PA: University of Pittsburgh Press.
- Loewer, B. (2008). Why There *Is* Anything Except Physics. In J. Hohwy & H. Kallestrup (Eds.), *Being Reduced* (pp. 149–163). New York, NY: Oxford University Press.
- Loewer, B. (2009). Why Is There Anything Except Physics? *Synthese*, 170(2), 217–233.
- Loewer, B. (2012). Two Accounts of Laws and Time. *Philosophical Studies*, 160(1), 115–137.
- Massimi, M. (2017). Laws of Nature, Natural Properties, and the *Robustly* Best System. *The Monist*, 100, 406–421.
- Ruetsche, L. (2011). *Interpreting Quantum Theories*. New York, NY: Oxford University Press.
- Rupert, R. D. (2008). Ceteris Paribus Laws, Component Forces, and the Nature of Special-Science Properties. *Noûs*, 42(3), 349–380.
- Urbaniak, R., & Leuridan, B. (2018). Challenging Lewis’s challenge to the best system account of lawhood. *Synthese*, 195, 1649–1666.
- Wilhelm, I. (2019). Celestial chaos. *Studies in History and Philosophy of Modern Physics*, 65, 97–102.
- Williams, P. (2019). Scientific Realism Made Effective. *The British Journal for the Philosophy of Science*, 70, 209–237.
- Woodward, J. (2014). Simplicity in the Best Systems Account of Laws of Nature. *The British Journal for the Philosophy of Science*, 65, 91–123.