

ABSTRACT

Title of dissertation: RELATIVE AND OBJECTIVE, ON BALANCE:
Detailing the Best Systems Analysis of Laws

Max Bialek, Doctor of Philosophy, 2017

Dissertation directed by: Professor Aidan Lyon
Department of Philosophy
and
Professor Jan-Willem Romeijn
University of Groningen,
Faculty of Philosophy

Variations on Lewis' Best Systems Analysis (BSA) of laws of nature have tended to emphasize the aspects of the view that allow it to accommodate the peculiarities of scientific practice. That move has allowed such views to do a lot of good work in solving old and new challenges for the BSA, but at the cost of strengthening the argument against the BSA that it is insufficiently objective. I argue that the "insufficiently objective" objection is overcome by a balance of relativity in the laws and limits to that relativity, each properly motivated by appeal to scientific practice. I then explore what relativity in the laws, and limits to it, may be required by scientific practice.

RELATIVE AND OBJECTIVE, ON BALANCE:
Detailing the Best Systems Analysis of Laws

by

Max Bialek

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2017

Advisory Committee:
Professor Aidan Lyon, Chair
Professor Jeff Bub
Professor Peter Carruthers
Professor Robert Friedel
Professor Barry Loewer
Professor Jan-Willem Romeijn

© Copyright by
Max Bialek
2017

Dedication

To Brittany and Solomon, for the perspective on what's interesting.

Acknowledgments

The intellectual journey that has led to this dissertation has not been a straightforward one. Since the start of my undergraduate degree I have changed fields at least four times, and, once I found my way to philosophy, it still took a few years to pass through as many subfields before finding myself working in philosophy of science. A consequence of all this movement, and a large part of what has enabled it, is that I have enjoyed the support of a long series of wonderful advisors, official and unofficial. Alongside them have been other teachers and fellow students whose enlightening teaching, engaging conversation, and clear passion for their subjects, have made we wish on so many occasions that I could set down whatever I was doing and work on something new. I suppose that this may be taken as a subtle dig against what I have stuck with, or, conversely, to suggest that the other subjects did not quite measure up in my mind. I do not mean for it to be taken those ways. An important theme of this dissertation is respecting the breadth of things that people find interesting. That we owe such respect is not, for me, a consequence of some theory of ethics or epistemic due diligence. Rather, it is something fundamental in my mind, and each of the people that I thank below have helped to put it there bit by bit as they have shared their interests with me.

Let me start by thanking the members of my committees: (at Groningen) Mathias Frisch, Jeanne Peijnenburg, Markus Schrenk, and Mauricio Suárez, and (at Maryland) Jeff Bub, Peter Carruthers, Robert Friedel, Barry Loewer, Aidan Lyon, and Jan-Willem Romeijn. I am enormously grateful to them for the time

they have spent reading, their comments, and their patience with the complications of organizing two defenses in the space of a week.

As an undergraduate at Rutgers, I had as de facto advisors Randy Gallistel and Barry Loewer. Weekly meetings with them, for several years with Randy, and for one year (the only one that I actually studied philosophy) with Barry, were always more interesting and (sometimes problematically) always received more of my attention than any class I was taking. Even as my interest in going into academia waned in the years before I discovered philosophy, I remained engaged with my nascent research program because checking in with Randy each week about it always ended up being so much fun. When I did move on to philosophy, a whirlwind year of independent study with Barry covered an enormous amount of intellectual ground. Only in retrospect have I appreciated how influential those meetings were, as the things we talked about in them—philosophy of mind, formal epistemology, philosophy of probability, philosophy of science, and even specifically laws of nature—have been almost perfectly retread as I moved through subjects during my graduate career.

As a graduate student, starting at Maryland and then also at Groningen, I have reaped the benefits of having three advisors. Aidan Lyon, when circumstances permitted, helped me to become lost in an enthralling world of research where a daily cycle of all-day discussions followed by furiously writing at night was the norm. (If I could live two lives, that would surely be one of them.) Jan-Willem Romeijn has been a guiding light, first for bringing me to Groningen, and then for helping me in the process of getting a great miscellany of thoughts from my head onto a coherent page. Leah Henderson has been an unrivaled reader of, and commenter on, my

work, whose every question suggested some improved formulation of my writing, if not requiring some substantial revision to my thinking. I am sure that I have more to learn from all of them.

After Rutgers and before starting graduate school, I had the pleasure of working for two years in Carlos Brody's neuroscience lab at Princeton. I will always be grateful to Carlos for taking a chance on me when I didn't know what I wanted to do and for having built up a lab with an atmosphere that supported and indulged me as it became ever more apparent that what I really wanted to do was philosophy. Contributing to that atmosphere were lab members I must also thank: Jeff Erlich, who patiently taught me (rat) neurosurgery and some delicate electrical engineering; Chuck Kopec, with whom I would always like to have lunch; and Bing Brunton, for exemplifying how I hoped to be as a graduate student.

I would like to thank a number of teachers whose lessons have had a profound impact on me. At Rutgers, Eduardo Sontag's courses on the mathematics of biology were inspiring when little else was, and Brian Loar's introductory epistemology course turned my world upside down and sent me off in search of more philosophy. Jeff Harty, Jerry Levinson, and Chip Manekin each taught multiple, deeply enjoyable, seminars at Maryland, and I like to think that in not too distant worlds I would have been writing dissertations supervised by them on logic, philosophy of art, and Jewish philosophy. Also at Maryland, Eric Pacuit should be singled out for special thanks, for classes he taught, his advice when I taught my own, reading groups he joined, and the steady stream of good advice and conversations I enjoyed while in the office next to his. While I did not have the opportunity to take courses

at Groningen, I have learned more than it seems time should have permitted from Catarina Dutilh Novaes, Barteld Kooi, Allard Tamminga, and Jeanne Peijnenburg.

My academic career has been made better as much from the advice of faculty as it has from the company of fellow students. When I first joined the philosophy department at Rutgers, Derek Anderson made me feel welcome immediately, and Zee Perry, who I now count as a dear friend, did me the great honor of thinking I was a worthy competitor. At Maryland, I enjoyed many long days and long nights, in bars and houses, in offices and seminar rooms, in the department lounge and library, with a host of fun and intelligent peers that included Heather Adair, Mike Dascal, Lane DesAutels, Lucas Dunalp, Quinn Harr, Ryan Ogilvie, Brendan Ritchie, Chris Vogel, and Evan Westra. Among them, Brock Rough stands out as a most reliable friend, interlocutor, confidant, and ally in life and fiction. At Groningen, I have benefitted immensely from the counsel and friendship of Pieter van der Kolk and Marta Sznajder.

In addition to teachers and peers, none of this would have happened without enormous administrative support. My thanks to Georges Rey, who made sure I found my way to Maryland, and Chris Morris, who made sure I found my way to Groningen and supported me stateside during my time there. Marga Hids and Sipeie Blom offered support and endless patience as I rushed at double time through four years of paper work and institutional peculiarities. And Louise Gilman, even when I was an ocean away, always knew exactly what needed to be done and was ready to do it—I cannot thank her enough.

Special thanks go to my great-aunt Paula de Waard, her children, and her

grandchildren, who have helped to make Groningen and the Netherlands feel like home since I first arrived two years ago. It is an honor to follow in the footsteps of my relatives who were also doctoral students at Rijksuniversiteit Groningen: my great-uncle Hendrik de Waard, grandfather Tjeerd van Anandel, and great-grandfather Hans Dekking.

Finally, I must thank those who were there before this academic journey began, and who each in their own ways were instrumental in helping me see it to its end. A friend for more than twenty-five years, Damien Chazelle's successes in the years that I have been a graduate student have been inspiring and a surprisingly personal source of joy for me. My wife Brittany, a teacher and an artist, has supported me in more ways that I can count over the last eleven years, and I hope that I have and can continue to do the same. I will succumb to the temptation to relate the birth just under a year ago of our son, Solomon, to the completion of this dissertation for only as long as it takes to assert that I love him and Brittany immeasurably more than I care about any of what I have written, and to thank them for affording me some time to write, for I surely would have spent all my time with them if not for their encouragement to do otherwise. My sister Fannie, who I am proud to say has just become a professor of religious studies, has been a brilliant scholar and original thinker since long before I tried to do such things. The overlapping similarities and differences between what we both now study have compelled me to think more carefully about what philosophy is, and I hope that I am a better philosopher, and person, for it. My mother Charlotte, an artist, architect, engineer, and organizer, has been a singular influence on me. While many people around

when I was growing up were celebrated for being among the best to do the thing that they did, she always did the things she did the best that they could be done, and more and more I find myself seeking out and aspiring to that sort of consistent perfection. My father Bill, a physicist, proved by example that you can find success by working on whatever interests you as long as you do it well. To be brief in describing his influence on me, and especially my intellectual life, I will say that it would not detract from the content of this dissertation if it was read as my attempt to understand the philosophical underpinnings of his work.

I say again to everyone mentioned above, and to the many that I have surely neglected (for which I am so sorry): Thank you. I hope this work pays back some of the debt I owe to you who have shared your interests with me.

Table of Contents

Dedication	ii
Acknowledgements	iii
List of Abbreviations	xii
1 Introduction	1
1.1 The Mill-Ramsey-Lewis View	5
1.2 Reactions to the BSA	11
1.2.0.1 Armstrong's Objection	12
1.2.0.2 Special Science Laws	12
1.2.0.3 The Trivial Systems Problem	13
1.2.1 Armstrong's Objection	14
1.2.2 Special Science Laws	18
1.2.3 The Trivial Systems Problem	22
1.3 The Four-Part Model	26
1.3.1 Armstrong's Objection, Expanded	29
1.3.2 Special Science Interests	32
1.3.3 Avoiding Trivial Systems	34
1.4 Dissertation Overview	37
1.5 Summary of Chapters	41
I Relativity	44
2 Interest Relativity	45
2.1 Armstrong's Objection Against Subjectivity	46
2.2 Armstrong's Objection Against Relativity	50
2.2.1 Respecting Scientific Practice	53
2.2.2 Relativity and Subjectivity	56
2.2.3 Responding to Armstrong	58
2.3 Degrees of Relativity	60
2.4 Required Relativity	67

2.5	Limiting Relativity	71
2.6	Summary	76
3	Relativizing to Kinds and Facts	78
3.1	The Better Best Systems Analysis	80
3.2	Relativity to Kinds and Facts	84
3.3	Interfield Interactions	87
3.3.1	Barlow, Between Biology and Physics	90
3.3.2	Interfield Interactions and the BBSA	92
3.3.3	Interfield Interactions and the KFRA	97
3.4	Fundamental Laws	101
3.4.1	Fundamental Laws and Fact Relativity	104
3.4.2	Fundamental Laws and Kind Relativity	107
3.5	Summary	109
II	Objectivity	112
4	Language Privileging	113
4.1	The Trivial Systems Problem	116
4.2	Lewis' BSA	117
4.3	The Package Deal Analysis	121
4.4	The Better Best Systems Analysis	124
4.5	Avoiding Trivial Systems	127
4.6	The Problem of Immanent Comparisons	134
4.7	Language-Class Relativity	143
4.7.1	Special Science Laws	144
4.7.2	Discovering Laws and Kinds Together	145
4.7.3	Limiting Language(-Class) Relativity	149
4.8	Summary	152
5	Induction Friendliness	155
5.1	Induction Friendliness and the BSA	157
5.2	Induction Friendliness and Mutual Information	166
5.3	Mutual Information and the BSA	170
5.3.1	Simplicity and Strength	171
5.3.2	Explanatory Power	174
5.3.3	Compression	177
5.3.4	Triviality	180
5.4	Summary	182
6	Conclusion	184

III	Appendix	194
A	Shannon's BSA	195
A.1	The Toy BSA	199
A.2	Degrees of Induction (Un)Friendliness	201
A.3	Toy Frequencies	206
A.4	Continuity	209
A.5	Monotonicity	213
A.6	Decomposition	215
A.7	Induction Unfriendliness	218
A.8	Induction Friendliness	219
	Bibliography	223

List of Abbreviations

AIC	Akaike Information Criterion
BBSA	Better Best Systems Analysis
BSA	Best Systems Analysis
CP	<i>ceteris paribus</i>
HM	Humean Mosaic
KFRA	Kind and Fact Relative Analysis
MI	mutual information
PDA	Package Deal Analysis
PIC	Problem of Immanent Comparisons
PTC	Problem of Transcendent Comparisons
ROBA	Relative and Objective, on Balance Analysis
RSP	An account of laws must respect scientific practice.
TSP	Trivial Systems Problem

Chapter 1: Introduction

What is a “law of nature”? We might understand this question to be asking “What is on the list of the things that are called ‘laws of nature’?”. The history of science is replete with things that might belong on that list. Some are helpfully named as laws, like Kepler’s laws of planetary motion, Newton’s laws of (general) motion, or the Laws of Thermodynamics. Other likely candidates, even if they are not standardly referred to as “laws”, are things like each of Maxwell’s Equations, or the symmetry principles that play a central role in much of modern physics. All of these candidate laws are *physical* laws, but not everything that gets called a law comes from the field of physics. There is the Weber-Fechner law in psychology that is concerned with perceptual sensitivity to changes in physical stimuli, there are the laws of supply and demand in economics, Malthus’ law of population growth in ecology, and many more. Which of these are rightly called laws of nature? I don’t know. *That* question is primarily the concern of scientists. It is the business of physicists to pick out the laws physics. Psychologists will identify the laws of psychology, economists the laws of economics, ecologists those of ecology, and so on.

But there is another way to understand the question “What is a ‘law of nature’?”. We might take it to be asking “What is it *to be* a ‘law of nature’?”. If all

the things named above are laws, what are the traits they share that make them such? If some of them *aren't* laws, what traits do those failed candidate laws possess or lack that distinguish them from the successful candidates? Some candidate laws, like Newton's, we now know to not be true universally, but they are still perfectly useful in suitably restricted circumstances. Are Newton's Laws not really laws because they aren't universally true, or might they still deserve the name as long as they are suitably qualified? *These* are questions that are amenable to philosophical analysis. Whatever the laws of nature may be, it is the business of scientists to say what they are. But it is the business of philosophers—and the overarching concern of this dissertation—to say what the laws *may be*.

So what is it to be a law of nature? One of the leading answers to that question is due to David Lewis and called the Best Systems Analysis (BSA). According to [Lewis, 1973], the laws of nature are the axioms and theorems of the true systematization of the world that is simpler and stronger than all other such systems. There are lots of ways that scientists might formally express—that is, systematize—what they know about the world. The preferred systematizations, the ones that end up in textbooks and whose first principles are dubbed the laws, tend to be simpler and stronger or more informative than their competitors. Such is the very rough sketch of the *epistemology* of the search for laws in scientific practice that inspires the BSA. What the BSA aims at is a *metaphysical* analysis of laws that parallels that epistemological picture.

Drawing so closely on the epistemology of laws for their metaphysics creates a special sort of challenge for the BSA. It is a mantra of contemporary analytic

philosophy that “you should not confuse your metaphysics and epistemology”. If the BSA is to succeed, it must walk a fine line between confusing and exploiting the epistemology of laws. Going too far in the ‘confusing’ direction will make the BSA dependent on the actual practice of science in a way that violates an expectation that the laws of nature are objective. But too far in the other direction will lead the BSA to describe laws that are ill prepared to play the right role in scientific practice.

The BSA has two significant virtues that make it worth developing despite the complications just described. Indeed, the view’s virtues are inseparable from the delicate balancing act that must be performed between metaphysics and epistemology. On the metaphysics side, the BSA is appealing for being a thoroughly realist and reductive analysis. Competing views of laws give up—too soon, I think—the search for such an analysis. Other realist views take as primitive laws (e.g. [Maudlin, 2007]), or things one might have hoped to analyze in terms of laws (e.g. [Armstrong, 1983]). Anti-realists deny that there is an account of laws to be had at all (e.g. [van Fraassen, 1989]). On the epistemology side, the BSA’s paralleling of scientific practice should make it that its laws are especially well prepared to play the role that laws play in scientific practice.

Extant formulations of the BSA have all recognized the balancing act that the view must perform, but they nevertheless direct their attention to just one side or the other. The BSA as described in [Lewis, 1983] requires that laws be expressed in the language of the metaphysically “natural properties”, and thus has been rightly criticized for rendering the laws inaccessible and potentially uninteresting to scientists.

Post-Lewis variants of the BSA (e.g. [Loewer, 2007], [Schrenk, 2008], and [Cohen and Callender, 2009]) have typically chosen to ensure accessibility and relevance to science by embracing anthropocentrism—they take “the best system” to mean the best system *for us*—and so their good work is achieved at the cost of improving the position of those who object to the BSA on the basis of its apparent subjectivity.

The goal of this dissertation is to help resolve for the BSA the tension between, on the one side, the presumption that the laws of nature should be objective, and, on the other side, the appeal of an analysis of laws that closely parallels scientific practice.

This introductory chapter will be dedicated to reviewing the challenges to the BSA in which the tension is realized and deploying a novel characterization of the BSA to illuminate the reviewed challenges and the variety of routes that may be taken to answer them. In Section 1.1, I go through the first formulations of the BSA that appear in [Mill, 1882], [Ramsey, 1928], and finally the modern standard in [Lewis, 1973, Lewis, 1983, Lewis, 1994]. In Section 1.2, I look at some responses to the BSA. Some of these are negative, like Armstrong’s objection to the BSA’s apparent subjectivity, Lewis’ realization that how good a system is will depend on the language in which it is expressed, and van Fraassen’s complaints about the laws’ accessibility and interest to scientists. I also consider positive responses in the form of variants of the BSA developed by [Loewer, 2007], [Cohen and Callender, 2009], and others, that improve upon Lewis’ original formulation. In Section 1.3, I introduce “the four-part model” of a BSA-style analysis of laws that identifies what elements all the BSA variants have in common, and makes clear how those

elements may vary to yield different versions of the view. The four-part model is then employed to expand upon the responses to the BSA discussed in Section 1.2. I conclude in Sections 1.4 and 1.5 with, respectively, a discussion of what motivates the structure of this dissertation and a summary of the coming chapters.

1.1 The Mill-Ramsey-Lewis View

At the heart of any Best Systems style analysis of the laws of nature is a nearly ubiquitous procedure: With some goal to be achieved, the path actually taken to that goal should be the best from among the candidate paths. To choose the best path is a matter of running a competition among the candidates with rules to determine each candidate's score, and then choosing the highest scoring candidate. A simple pros and cons list does this, with the candidates of "do X" and "do not do X" being evaluated by the count of considerations that may be made in favor or against the doing of X. This picture, though, is clearly not enough to provide an analysis of anything, much less laws. What matters is how we fill in the details. Saying that the goal is to identify the laws of nature is still not enough. Suppose that I enter extant bread baking recipes into a competition in which recipes are evaluated based on the quality of the bread produced, flexibility with respect to small errors on the part of the baker, and how easily they may be adjusted to make a wide variety of breads. It doesn't matter what my goal is—it may be to identify the laws of nature—what I get out of such a competition is the best recipe to recommend to an aspiring bread baker.

What must be provided in order for this sort of procedure to be the centerpiece of an analysis of the laws of nature is a description of the competitors and competition that would yield a victor that may plausibly be interpreted as (providing) the laws. Filling in these details (though not explicitly as such) is done first by Mill, who, concerned with induction and the uniformity of nature, notes that

the uniformity in question is not properly uniformity, but uniformities [...] These various uniformities, when ascertained by what is regarded as a sufficient induction, we call, in common parlance, Laws of Nature. Scientifically speaking, that title is employed in a more restricted sense, to designate the uniformities when reduced to their most simple expression.

[[Mill, 1882](#), p. 229]

We are provided with an example just after this passage in which there is a collection of seven uniformities and, since four are derivative of the other three, we are told it is just the three that should be considered laws. Mill's version of the view is more or less apparent: Candidates are sets of uniformities, and the competition evaluates the candidate sets according to the simplicity of their expression. To be a law of nature is to be an element of the most simply expressed set of uniformities of nature.

Ramsey, concerned with distinguishing the laws of nature from mere "universals of fact", presents a view in [[Ramsey, 1928](#)] of laws similar to Mill's that he later summarizes (and rejects¹) as being a view according to which laws are the

¹ Why? I suspect a better understanding than my own of the practice of philosophy in Cambridge in the late 1920s is required to properly answer this question. My best guess is that the rejection has to do with a shift from trying to identify the laws themselves to, as it is put in [[Ramsey, 1929](#)], "trying to explain the meaning of asserting the existence" of a law. [[Lewis, 1973](#), p.

consequences of those propositions which we should take as axioms if we knew everything and organized it as simply as possible in a deductive system.

[Ramsey, 1929]

As in Mill, the competition is concerned with the simplicity of the competitors. Competitors are full deductive systems from which the laws may be extracted.² This marks a slight shift from Mill in two ways. First, laws for Ramsey are the theorems of a deductive system, whereas Mill, as evidenced by the example mentioned earlier, would reject on grounds of simplicity the full set of theorems of a deductive system in favor of just the axioms. Ramsey does distinguish between the axioms as being “ultimate laws” and the non-axiom theorems as being “derivative laws”, but takes the unqualified *laws of nature* to be the full set of theorems since, while the choice of axiomatization might be somewhat arbitrary, the set of theorems “is less likely to be arbitrary if any simplicity is to be preserved” [Ramsey, 1928].

The views on laws of Mill and Ramsey may appear problematic as precursors to the modern BSA since, as presented, they seem more concerned with epistemology than metaphysics (and, I claimed, the BSA is a *metaphysical* analysis of the laws of nature). Mill writes of how “the title [of Law of Nature] is employed”, but nothing in the immediate discussion denies that the uniformities (deserving the title of law) are

73, fn. *] suggests the rejection has to do with Ramsey saying that we can’t know everything, but I think that is not enough to reject the view without also considering the change in Ramsey’s interests.

² In Mill the laws had to be “extracted” from the best competitor as well, but there it was just a matter of picking out each of the elements of the competing set. In Ramsey it is the more complicated procedure of finding all the theorems of the deductive system. In general I will say of any view that it “extracts the laws from the victor of the competition” since that is true of all of them while the method of extraction varies depending on the nature of the competitors.

present in nature independent of anyone employing the title. While the summary in [Ramsey, 1929] is concerned with, as [Lewis, 1973] puts it “a counterfactual about omniscience”, Ramsey was sensitive to this issue. He writes there that the counterfactual “is only a *spurious* one” (emphasis in original) since being a law is a matter of “the facts that form the system in virtue of internal relations, not people’s beliefs in them”; furthermore, he says of the competition’s valuing simplicity that “this is another vague formal property, not a causal one” that makes the laws a matter of our conception of simplicity [Ramsey, 1928].

The balance between metaphysics and epistemology in the BSA is something about which we must be very careful. Any presentation of a BSA-style view will benefit greatly in spirit by including epistemological analogies, but we must be careful in the proper statement of a view to make clear the distinction between the metaphysical analysis of laws and any analogous epistemology of law discovery. Lewis writes (when he is concerned with extending the BSA into an analysis of laws and chance) that

Despite appearances and the odd metaphor, this is not epistemology! You’re welcome to spot an analogy, but I insist that I am not talking about how evidence determines what’s reasonable to believe about laws and chances. Rather I am talking about how nature [...] determines what’s true about the laws and chances. Whether there are any believers living in the lawful and chancy world has nothing to do with it.

[Lewis, 1994, pp. 481–482]

The fact that Lewis worked out his analysis of laws as much as he did while always being sensitive to this issue is, I think, an important part of why he is (and should be) considered to have provided the first real statement of the BSA. So let us consider that first real statement of the BSA from Lewis.

In [Lewis, 1973] we are provided with a “restatement” of Ramsey’s view, formulated to avoid any dependence on epistemology or (since Lewis is preparing to use laws as an element of his analysis of counterfactuals) counterfactuals (all emphasis in the original):

Whatever we may or may not come to know, there exist (as abstract objects) innumerable true deductive systems: deductively closed, axiomatizable sets of true sentences. Of these true deductive systems, some can be axiomatized more *simply* than others. Also, some of them have more *strength*, or *information content*, than others. [...] a contingent generalization is a *law of nature* if and only if it appears as a theorem (or axiom) in each of the true deductive systems that achieves a best combination of simplicity and strength.

[Lewis, 1973, p. 75]

Two major additions are made by Lewis to the view in his subsequent writing on the subject. In [Lewis, 1983] a requirement is introduced that the axioms of each competing system must be expressed in a language with basic predicates for only the “perfectly natural properties” in the world. Then, in [Lewis, 1994], the BSA is expanded to provide an analysis of chances according to which, very roughly, chances

are whatever the laws of nature say they are. This expansion comes with changes to both the competition and competitors. The set of competitors is, of course, now allowed to include systems that describe chancy regularities (which will, in the victor of the competition, be named as the laws that define the chances of the world). To deal with the growth in the set of competitors, a consideration of *fit* is added to the competition—where a better fitting system is one that makes the actual history of the world more probable—such that the victor will be the system striking a best balance of simplicity, strength, and fit.

Another feature of Lewis’ BSA that does a lot to distinguish it from Mill and Ramsey is the work dedicated to describing what it is exactly that is being systematized. [Ramsey, 1929] suggests that what is being systematized is the set of propositions known to us were we omniscient. [Ramsey, 1928], when trying to deal with the problem of epistemology infecting the metaphysics of laws, speaks of “the facts” as being what is systematized, as opposed to “people’s beliefs in [the facts]”. For Lewis, what is systematized is the Humean Mosaic (HM), the fundamental facts of the world, which are concerned entirely with basic spatiotemporal relations and what properties obtain locally at each space-time point. With only the spatiotemporal relations and local properties, there are no fundamental facts about causes, or chances, or necessity, or laws. This specificity comes as a part of Lewis’ larger metaphysical project of “Humean Supervenience”, according to which anything not in the HM (like laws or chances) “supervenes on the spatiotemporal arrangement of local qualities throughout all of history, past and present and future” [Lewis, 1994].

Going forward it will be helpful to distinguish between *the* BSA and *Lewis’*

BSA. Lewis' BSA, as described above, says (roughly) that something is a law just in case it is a theorem in all of the systems which strike the best balance between being simple and strong/informative and fitting with respect to the Humean Mosaic, and whose axioms are expressed in the language of perfectly natural properties. A proper statement of *the* BSA—which I take to be the core of the view that is shared between Lewis' BSA and any (actual or potential) post-Lewis variant of Lewis' BSA—will not appear until Section 1.3, but for now let the following formulation of the BSA suffice: A regularity is a law just in case it appears in the best systematization of all the fundamental matters of fact.

1.2 Reactions to the BSA

I am concerned in this section with reviewing the literature that has developed following Lewis' modern formulation of the BSA. Dialectically, responses to Lewis fall into two broad categories: (1) objections to the view, and (2) adaptations of the view to answer questions found elsewhere in philosophy of science and metaphysics. Responses of the second variety often do double duty by both extending the BSA and answering objections to it. Topically, I have picked out three groups of responses to the BSA that are important to the tension between objective laws and paralleling practice that is the concern running through this dissertation:

1.2.0.1 Armstrong’s Objection

Armstrong’s objection³ challenges the BSA on the grounds that laws of nature should be objective, but the BSA makes laws subjective due to their dependence on what particular, and presumably subjective, answer is given to the question of what makes a system “best”. Armstrong’s objection receives only a few brief responses despite a trend in post-Lewis variants of the BSA towards making the view more dependent on the subjective interests of scientists.⁴ I discuss Armstrong’s objection and reactions to it in Sections 1.2.1 and 1.3.1 below, and at length in Chapter 2.

1.2.0.2 Special Science Laws

A distinction is often made between *fundamental laws* and *special science laws*, and a well developed account of laws should say something about this distinction. Unlike Armstrong’s objection, distinguishing between these two types of laws does not pose a unique challenge to the BSA, but a general challenge to accounts of laws for which the BSA appears to offer good answers. Precise characterizations of what makes a scientific field “fundamental” or “special” can vary, but often the matter turns on the fundamentality of the matters of interest to the field. Physics—being concerned with the fundamental nature of reality—is *the* fundamental science. Chemistry, biology, psychology, and pretty much any scientific field other

³ So named for its place as first in the list of complaints against the BSA in [Armstrong, 1983].

⁴ I use “scientists” here and throughout the dissertation loosely to refer to those who seek out the laws of nature and those who employ the laws (or their best estimates of the laws) to achieve some end (e.g. provide an explanation, engineer some device,...). I write as though the search for, and usage of, laws of nature is ubiquitous in science, but recognize that this is contentious at best.

than physics—being concerned with phenomena presumed to be not fundamental—are special sciences. The matter of accommodating special science laws in the BSA is taken up by a number of authors but its most significant treatment comes in the Better Best Systems Analysis (BBSA), a view developed independently in [Schrenk, 2008] and [Cohen and Callender, 2009]. While Cohen and Callender focus on the relevance of kinds to the special science laws, Schrenk’s concern is directed primarily at the issue of *ceteris paribus* (CP) laws, whose defining feature of having exceptions is sometimes taken to also be the defining feature of special science laws. I discuss the work of accommodating special science laws in the BSA in Sections 1.2.2 and 1.3.2 below, and at length in Chapter 3.

1.2.0.3 The Trivial Systems Problem

The Trivial Systems Problem (TSP) is concerned with the existence of certain systems that, when expressed in certain languages, are overwhelmingly the best systems, but offer only trivial or deeply problematic laws that do not seem to deserve the title. The problem is first introduced in [Lewis, 1983], who responds immediately by adding a language requirement to his version of the BSA. Lewis’ language requirement is criticized by [van Fraassen, 1989]. Answers to the TSP that are also sensitive to van Fraassen’s criticisms are provided by [Loewer, 2007] and [Cohen and Callender, 2009]. I discuss this work in Sections 1.2.3 and 1.3.3 below, and at length in Chapter 4.

1.2.1 Armstrong’s Objection

Variations on the following story often accompany summaries of the BSA. A scientist dies and appears before God, who asks if there is anything the scientist would like. The scientist replies “I would like to know how the world works”. God accepts the request and begins “At time and position $\langle t_1, p_1 \rangle$ such-and-such properties obtained, and at $\langle t_1, p_2 \rangle$ such-and-such other properties obtained, and—” at which point the scientist interrupts and says “Sorry, if that is how things *actually* work that is all well and good, but is there a *pithy* version?”. God’s reply to that question is to provide the scientist with the (BSA-style) laws of nature.

When the scientist asks for the “pithy version” of all the fundamental matters of fact, what is being asked for is something that the scientist would find more accessible but still informative, and what it means to be more accessible and informative to the scientist will of course depend on the abilities and interests of the scientist. The ability and interest relativity apparently inherent in the BSA is the source of a standard critique of the view:

“The *first* objection which may be made to the [BSA] is that an element of subjectivism remains... May there not be irresolvable conflicts about the exact point of balance [between simplicity and strength]?”

[Armstrong, 1983, p. 67]

What it is to be the “best system” looks to be a subjective matter. Some systems may be better for some than for others, and if there is disagreement about which

system is best, then there will likely be disagreement about which regularities are the laws of nature. Our intuition that the laws of nature should be objective make this look like very bad news for the BSA.

Lewis calls this “the worst problem” for the BSA, and then offers a less than satisfying response: “If nature is kind to us, the problem needn’t arise” [Lewis, 1994, p. 479]. The best system, he hopes, will be “robustly best” such that no defensible change in what *makes* a system the best will actually change which system *is* the best. If nature is unkind, then Lewis would sooner see there be no laws than non-BSA laws. Presumably this is due to his commitment to Humean Supervenience, which is incompatible with most realist accounts of laws other than the BSA. Proponents of the BSA who are not committed to Humean Supervenience—a group which includes, at least explicitly, [Loewer, 2007] and [Cohen and Callender, 2009]—should be less pessimistic about laws even if the BSA fails, and so a better response must be found.

A more common response than Lewis’ “nature is kind”, though not obviously more satisfying, is a sort of acceptance. The laws come from the system that is best *for us*, whoever we are. This attitude is apparent among a number of authors: “we understand the notion of a Best System flexibly according to our needs” [Cohen and Callender, 2009, p. 21]; the Package Deal Analysis (PDA), a variant of the BSA, “goes farther in the direction of anthropo[centrism]... in that what counts as a final theory depends on the tradition of fundamental physics” [Loewer, 2007, p. 325]; and, in a version of the BSA concerned with probabilities instead of laws, “the ‘best’ in my use of ‘Best System’ means best *for us*” [Hofer, 2007, p. 571,

emphasis in original]. These authors present the flexibility, the anthropocentricity, the subjectivity, of their views as a virtue—not, as Armstrong would have it, a failure—on the grounds that it makes their variants of the BSA better at respecting scientific practice.

The best response to Armstrong’s objection has been to attack the claim about subjectivity (as opposed to the previous two responses that were concerned with the possibility of disagreement amongst scientists) directly by pointing out that that the BSA is a strictly metaphysical view into which the epistemology of law seekers does not actually enter. This “brute metaphysics” response is already implicit in the earlier quotations from Lewis and Ramsey. It is made explicit by Cohen and Callender when they write that “relativized MRL”, which is their variant of the BSA⁵,

does not make the laws subjective in the sense that they depend for their existence on subjects. This is because the laws (relative to basic kinds K) that hold of a world w would satisfy relativized MRL’s criterion for lawhood... whether or not there are subjects in w (or any other world).

[Cohen and Callender, 2009, p. 30]

If there are disagreements amongst scientists, so what? The laws are what they are according to *some* standard of “best” and any law seekers not employing that standard will fail in their search.

This is right, but may still be unsatisfying because it misses the spirit of

⁵ In [Callender and Cohen, 2010] the same view is titled “the Better Best System analysis”, or BBSA, and I will typically refer to their view as such.

Armstrong's objection. A big part of the appeal and motivation of the BSA is the competition's ability to parallel the epistemic practices of scientists and Armstrong's objection is meant to cause trouble as a consequence of that aspect of the BSA. Suppose that one of the disagreeing scientists will be wrong. Why is that one wrong and the other right? Cohen and Callender adopt a permissive approach according to which both scientists can be right *and* what the laws are doesn't depend on the scientists themselves. The cost of this is a glut of laws. There are enough to be the target for a scientist of *any* persuasion, so there are far more than will ever plausibly be needed or wanted by any real scientist. The laws may not be subjective in the technical sense, but technically speaking Armstrong only wrote that "an element of subjectivism remains". That vagueness in the original statement of the problem suggests that there may be a different, albeit nearby, concern related to the sorts of disagreements that give rise to the problem. The existence of such an additional concern is an open question, and, if it does exist, it will be important to check if it is successfully addressed by the brute metaphysics response.

While Armstrong's objection is compelling, it is in need of a better statement. In Section [1.3.1](#) below, I will show how Armstrong's objection evolves in response to a proper characterization of the BSA. In Chapter [2](#), I will be concerned with defending the possibility of giving a response to Armstrong's objection and arguing that such a response will have certain features. Specifically, proponents of the BSA must pursue a balance of relativity and objectivity, with the former motivated by what is necessary to accommodate variety in scientific practice, and the latter by what is universal in scientific practice. The entirety of this dissertation is, more

or less explicitly, devoted to exploring the balance of relativity (in Part I) and objectivity (in Part II) that should exist in the BSA.

1.2.2 Special Science Laws

Psychology, biology, and economics are paradigmatic special sciences. Fundamental physics is decidedly *not* a special science. Perhaps the most important feature of the distinction between special and, for lack of a more general term, non-special sciences in the context of a discussion of laws is that it is assumed that non-special sciences have laws (*pace* anti-realists about laws), while special sciences may or may not have laws. If there are special science laws, they may or may not be reducible to the laws of non-special sciences. The kinds that are the concern of non-special science laws are the fundamental kinds, and the kinds that are the concern of special science laws are not (all) fundamental. Laws of non-special sciences are assumed to be exceptionless— $E = mc^2$ everywhere—whereas laws of the special sciences are typically thought to have exceptions indicated by an “all else being equal” or “*ceteris paribus*” clause—as in “all else being equal, taking acetaminophen relieves a headache”.

The above is all said in the name of intuition-pumping. Particular views coupled with particular interests will tend to prefer one way of characterizing the distinction over, or to the exclusion of, others. For our present purposes it is okay to leave the details of the distinction unspecified, but follow Cohen and Callender in loosely basing it on whether the interests of a field are directed towards

non-fundamental kinds—in which case what is sought are *special science laws*—or towards fundamental kinds exclusively—in which case what is sought are *fundamental laws*. (This question of how to properly distinguish between special science and fundamental laws in the BSA will be taken up in Chapter 3.)

One strategy for accommodating special science laws in the BSA is to make them derivative of the fundamental laws. [Loewer, 2012] and [Albert, 2000] pursue such a strategy, and offer a version of the BSA that I will refer to as “the Mentaculus view”.⁶ The defining feature of the Mentaculus view is (not surprisingly) the Mentaculus, which has three parts:

- (i) The fundamental dynamical laws.
- (ii) The claim that the initial macro state of the universe is $M(0)$ and that the entropy of $M(0)$ is very tiny.
- (iii) A law specifying a uniform probability over the micro states that realize $M(0)$.

These are not devoid of reference to non-fundamental properties or kinds, but still the Mentaculus is taken to be “the fundamental theory of the world” and special science laws are merely consequences of it in conjunction with further propositions that connect the fundamental laws to the higher level kinds of the special sciences [Loewer, 2012]. [Frisch, 2011] rightly argues that the derivative status conferred on special science laws by Albert and Loewer is at odds with the presence of non-fundamental properties/kinds in the Mentaculus and the “pragmatic dimension” of the view in which the laws are connected to the practice and interests of scientists.

⁶ The Mentaculus view is distinct from, but compatible with, the PDA of [Loewer, 2007].

A need to accommodate special science laws in the BSA becomes greater as proponents of the BSA increasingly emphasize pragmatic considerations that tie the laws to the needs and interests of practicing scientists. The need becomes greater when evaluating the success of the BSA because, insofar as we think special scientists do, can, or should, seek laws, it is harder to discount the needs and interests of special scientists over those of fundamentally inclined scientists. The need also becomes greater internal to the BSA if features of the view motivated by pragmatic considerations may make it such that distinctly special science laws are part of the best system whether we like it or not (as Frisch argued with respect to the willingness of Albert and Loewer to let the laws refer to non-fundamental kinds). Helpfully, the greater need to accommodate special science laws in the BSA comes hand in hand with greater ease in making such accommodations, as it is precisely the pragmatic considerations requiring the accommodation that provide the tools to do it.

What I take to be the best strategy for accommodating special science laws in the BSA is found in the BBSA, developed independently by [Schrenk, 2008] and [Cohen and Callender, 2009] (and endorsed by [Frisch, 2011]). The BBSA—specifically, the version found in [Cohen and Callender, 2009] and [Callender and Cohen, 2010]—works roughly as follows. There are true fundamental kinds K_{fund} and, built from that, there is the set of all supervenient and fundamental kinds \mathcal{K} (i.e. the set of kinds that supervene on the fundamental kinds in union with K_{fund}). Each subset K_i of the set of supervenient kinds corresponds to a language L_i that renders the kinds of K_i as its basic predicates.⁷ For every such L_i , a *separate* best system

⁷ Cohen and Callender tend to speak of kinds and predicates, and not languages, but use the

competition is run in which every competing system is expressed in the language L_i . To be a law is to be a law *relative* to a language. The laws relative to L_i are extracted from the victor of the competition in which all competing systems were expressed in L_i ; there are no laws of nature simpliciter. Finally, the interests of a special science field are identified with the set of kinds that the field treats as basic. So, if field X treats as basic the set of kinds K_X , and L_X is the language of the kinds K_X , then the laws determined relative to L_X are *the laws of the field X* . Assuming, as Cohen and Callender do, that subsets of the supervenient kinds are exhaustive of the interests of scientific fields, then this strategy guarantees that for any scientific field there exists a set of laws that is the field's target.

The version of the BBSA introduced in [Schrenk, 2008] is less concerned with choices of kinds distinguishing special science laws from each other and the fundamental laws than it is with the possibility of there being exceptions in the laws (typically assumed to be special science laws, but also, possibly, fundamental laws). Building on [Braddon-Mitchell, 2001], [Schrenk, 2008] rightly argues that the BSA is well suited to accommodating exception having CP-laws, because (in short) such exceptions may contribute more to the simplicity of a system than they cost in strength. To the extent that special science laws are characterized by the presence of CP exceptions, [Schrenk, 2008] clearly demonstrates how they may be accommodated in the BSA.

terms more or less interchangeably when discussing others who speak of languages; e.g. “written in the language of the [...] fundamental kinds” [Cohen and Callender, 2009, p. 10]. I will try to speak of languages as the things in which candidate systems are expressed, and reserve talk of kinds for when it is kinds as such that are at issue (e.g. when a scientist treats a set of kinds as basic, and not the language that corresponds to those kinds).

1.2.3 The Trivial Systems Problem

[Lewis, 1983] identifies a problem—which I have called the “Trivial Systems Problem” (TSP)—with his BSA. How simple a system is depends on how the system is expressed. For any system S , one could always identify as primitive a predicate F that is true of everything in just those worlds where S holds, and then axiomatize S with just the sentence $\forall xFx$. This is, presumably, about as simple as a system can get. And, since we could do this for any system, let S be the strongest system. S , then, is the best system, since all systems have equally maximal simplicity and S is the strongest. The problem is this: We expect that there should be a way to distinguish between accidental regularities and the lawful regularities. For example, it may be true of the world that nowhere is there a solid sphere of gold one mile in diameter and that nowhere is there a solid sphere of uranium-235 one mile in diameter. The absence of such a gold sphere is accidental; one could exist, it just so happens that none do. The absence of such a uranium sphere seems lawful; physics tells us that a solid uranium-235 sphere would blow up long before it could reach a mile in diameter. Moving back to the trivial system problem, since S is the strongest system, the theorems of S will include every regularity that holds in the world, every true regularity will be a law, and so we will fail to distinguish between the lawful and accidental regularities. To avoid this outcome we must find a way to rule out the possibility of there being a system like S .

Lewis’ solution is to require that the axioms of a system be expressed in the language of perfectly natural properties. Since F is presumably not a perfectly

natural property, “ $\forall xFx$ ” cannot be the axiomatization of any candidate system, and thus S , so axiomatized, could not be chosen as the best. This strategy is supposed to have the added virtue of explaining why “laws and natural properties get discovered together” [Lewis, 1983, p. 368].

[van Fraassen, 1989] identifies two problems with Lewis’ natural language requirement (I will call them “van Fraassen’s problems”, following [Loewer, 2007]) that have to do with the prospect of actually discovering the laws and natural properties together.

The first of van Fraassen’s problems is concerned with the accessibility of the laws. Suppose $S_{\text{best}}, L_{\text{nat}}$ provides us with the true BSA laws; L_{nat} is the language of the perfectly natural properties, and S_{best} is the best system when all competing systems are expressed in L_{nat} . Because S_{best} is only guaranteed to be the best when competing systems are expressed in L_{nat} , there is the possibility of a competing system, S_{vf} , that, when expressed in some language L_{vf} (that is not L_{nat}), is *better* than S_{best} expressed in L_{nat} . We already think such a system exists in the form of the trivial system, but suppose further that there is nothing so problematic with S_{vf} as there is with the trivial system. In such a situation, there is no way that scientists could ever pick out L_{nat} over L_{vf} as the proper language with which to identify S_{best} as the source of the laws. If $S_{\text{vf}}, L_{\text{vf}}$ is the overall best system-language pair (setting aside problematic pairs like those of the trivial system problem), L_{vf} will be the source of scientist’s best estimate of the perfectly natural properties, and that best estimate will be wrong. Even if nature is kind in that there is no pair better than the best system as expressed in the language of perfectly natural properties,

any sensibly skeptical scientists will know that they can never be sure about the laws.

The second of van Fraassen's problems is concerned with how deserving Lewis' laws are of their title. If $S_{\text{vf}}, L_{\text{vf}}$ really is the best (unproblematic) system-language pair, then why is it not the source of the laws? If scientists somehow know that L_{nat} is the language of the perfectly natural properties, then van Fraassen's accessibility problem goes away. But $S_{\text{vf}}, L_{\text{vf}}$ is still the better system-language pair. The only reason to stick with L_{nat} and S_{best} in these circumstances is a question-begging desire to have the laws expressed in the language of perfectly natural properties.

van Fraassen is right about all of this, and, together, these problems are a serious blow to Lewis' BSA and its requirement that the language of the competing systems be the language of the perfectly natural properties. If the perfectly natural properties are not discoverable independent of the laws, then the laws are inaccessible. And, even if the perfectly natural properties *are* discoverable independent of the laws, there may be law-like regularities (but not actual laws) available that are, strictly speaking, *better* than the actual laws. In the end, Lewis' language requirement "has produced unchartable distances between Lewis's best theories—and hence laws—and the theories we could reasonably hope for at the ideal end of science" [van Fraassen, 1989, p. 55].

The responses to the TSP after Lewis are in agreement with Lewis that the problem requires the privileging of an unproblematic language prior to the best system competition. The differences between them come in *how* they privilege a language.

We have already seen the sort of language privileging that exists in the BBSA, so let us consider that view’s treatment of the TSP first. Recall of the BBSA that, for every subset K of the set of all fundamental and supervenient kinds, there is a set of K -relative laws. Because F is a supervenient kind, there are actually many sets of laws in the world that succumb to the TSP. If $F \in K$, then $\forall xFx$ will be among the competing systems, and, as per the TSP, on track to be the problematically best system. We are assured, however, that this is not actually a problem. The laws of a scientific field are the laws determined relative to the set of kinds that are of interest to the field. As long as no field is interested in the kind F , then no sets of laws of *interest* will be troubled by the TSP. There are still sets of laws that succumb to the TSP, and fail to distinguish between lawful and accidental regularities, but such laws were already deemed to be uninteresting. “Properties like F and the ensuing threatened trivialization [...] are ruled out for lack of interest rather than any intrinsic deficiency” [Cohen and Callender, 2009, p. 23]. This also addresses van Fraassen’s problems, since it binds the laws to predetermined—and, necessarily, accessible—kinds, and makes the laws quite intimately a matter of scientific interest.

The other major variant of the BSA to address the TSP is Loewer’s PDA, the “Package Deal Analysis”. The titular “package deal” comes from the best system competition evaluating systems and languages together (as a system-language pair/package). Doing that suggests that all we need to do to block the TSP is block any languages that contain the predicate F , and let all the other languages compete, but this is not quite Loewer’s strategy. “Lewis’s argument does show that [...] the BSA requires a preferred language” [Loewer, 2007, p. 325]. “Language”, singular.

What Loewer proposes for the preferred language is that it be the language of an “optimal final theory” that would be arrived at through a succession of rational developments on the present theory of fundamental physics. Because the present theory of fundamental physics does not include F , and, presumably, no rational development of it would introduce F , the trivial system problem is overcome. Loewer has also avoided van Fraassen’s problems, since the accessibility and interest of the laws to scientists (or, at least, physicists) is ensured by the laws being what they are as a consequence of the tradition of fundamental physics being carried out to its ideal end. But these successes come with the apparent cost of strengthening Armstrong’s objection by tying the laws so closely to the actual practice of physics.

Overall, Lewis, Loewer, and Cohen and Callender each offer language privileging solutions to the TSP. *What* language(s) is (are) privileged varies greatly between the three, but the general strategy is the same. In Section 1.3.3, I will argue that a proper characterization of the BSA presents us with possible alternative strategies for answering the TSP. In Chapter 4, I argue that the TSP may be overcome without language privileging, but rather by carefully constructing the BSA’s competition.

1.3 The Four-Part Model

I suggested at the start of this chapter that the BSA may be summarized tentatively as saying that laws are drawn from the best systematization of all the fundamental matters of fact. That doesn’t say much about the facts being systematized. That doesn’t say much about the candidate systematizations. That doesn’t

say much about the language(s) in which competing systems are to be expressed. And that doesn't say much about what it means to be best. These are details that were to be filled in by any BSA variants, of which we now have three. Lewis' BSA says laws are drawn from the simplest, strongest, and most fitting (on balance) systematization of the HM expressed in the language of the perfectly natural properties. The PDA says laws are drawn from the system-language pair that is the best final theory of the world according to the practices of fundamental physics. The BBSA says that the laws of a language are drawn from the best systematization of the facts of the world where all competing systems are expressed in that language, and that there is such a set of laws for every language corresponding to a subset of the set of kinds that supervene on the true fundamental kinds.

What would be nice is a way of stating the BSA that makes these particular views clearly special cases of a general view, and which makes salient the ways that details may be added to the BSA to yield any possible variant. For this I recommend what I will call the *four-part model* according to which the BSA says that there are *the facts* (part 1) to be systematized, *the candidate systems* (part 2) from which the laws will be drawn, *the languages* (part 3) in which the candidate systems and facts may be expressed, and *the competition* (part 4) which determines the best system-language pair.

I will often use the following symbols when discussing the the BSA and its parts: There is the set \mathcal{F} of all the facts of the actual world w_{actual} ; we will sometimes be concerned with particular sets of facts $F \subseteq \mathcal{F}$. There is the set \mathcal{S} of all the systematizations S that compete to be the best. There is the set \mathcal{L} of all

languages L that may be used to express the facts of \mathcal{F} and systems of \mathcal{S} ; we will sometimes be concerned with a set of all kinds \mathcal{K} and assume there is a one-to-one correspondence between the sets of kinds $K \subseteq \mathcal{K}$ and languages $L \in \mathcal{L}$. Finally, there is the competition function $C(-)$, which (typically) takes \mathcal{F} , \mathcal{S} , and \mathcal{L} , as arguments and outputs a system-language pair $S_{\text{best}}, L_{\text{best}} \in \mathcal{S} \times \mathcal{L}$. The laws R , then, are the regularities that appear in S_{best} and are to be expressed in L_{best} .

We can summarize our three major BSA variants again with little change but for a new emphasis on the four parts. Lewis says (1) the facts are those of the HM, (2) the candidate systems must be true of the facts, (3) the language in which the systems and facts must be expressed is that of the perfectly natural properties, and (4) the competition picks a winner that maximizes a balance of simplicity, strength, and fit. The PDA keeps with Lewis in that (1) the facts are those of the HM, but (2) the candidate systems and (3) language is that of the final theory of physics, and (4) the competition picks a winner that best satisfies the demands on laws made by fundamental physicists. According to the BBSA there are, again, (1) the facts of the HM, and (2) the systematizations of those facts, but (3) the set of viable languages includes one for every subset of the set of all supervenient and fundamental kinds, and (4) the competition(s) will output a best system for every language, with the laws of each such system being the laws *relative* to the language for which the system is best.

Being able to make sense of extant BSA variants is a necessary feature of the four-part model, but whether or not it is worthwhile as a general characterization of the BSA will be decided by its ability to illuminate the challenges that face the BSA.

I begin to show how it does so below for Armstrong’s objection (Section 1.3.1), the accommodation of special science laws (Section 1.3.2), and the TSP (Section 1.3.3), in anticipation of putting the model to work throughout the dissertation.

1.3.1 Armstrong’s Objection, Expanded

Armstrong complains about the BSA’s possessing “an element of subjectivism” on account of the possibility of “irresoluble conflicts about the exact point of balance” between simplicity and strength [Armstrong, 1983, p. 67]. In the more general language of the four-part model, Armstrong’s objection is about the possibility of disagreement about the details of the competition. But, as the four-part model makes clear, there is a lot more about which there could be disagreement. Beside the competition there are the facts, the set of candidate systematizations, and the set of viable languages. Just as Armstrong imagined rationalists and empiricists arguing over the primacy of simplicity and strength, we can imagine reductionists and anti-reductionists arguing over the facts to be systematized since, for example, if you’re an anti-reductionist about psychological states, the facts that would be systematized to yield the laws of psychology can’t just be the fundamental facts appropriately expressed (as they are in the reduction friendly BBSA). Different views about the structure of scientific theories will likely yield disagreement about the set of candidate systems, and different views about natural kinds will likely yield disagreements about the set of viable languages. The weight of Armstrong’s objection presumably grows with the number of points of disagreement that could give rise to

“an element of subjectivism”.

As we saw earlier (in Section 1.2.1), the brute metaphysical response thwarts accusations of subjectivism. But what about relativism? For every variant of the BSA the laws are what they are relative to how the four parts are specified. And, more than that, a single variant of the BSA (like the BBSA) can allow for multiple ways of filling in the details such that there is a multitude of things rightly called the laws relative, explicitly, to what in the four parts was allowed to vary. A minimally relative BSA variant (like Lewis’) will say that there is exactly one acceptable way of filling in the details of the four-part model. But not all BSA variants need be (or are) minimally relative. The BBSA holds the facts, systems, and competition fixed, but says that laws may be yielded for any language of supervenient or fundamental kinds.⁸ There could similarly be variants of the BSA that hold all but the facts to be systematized fixed, yielding laws for every possible set of facts. Or which hold all but the competition fixed, yielding laws for every way of determining which system is the best. Or which hold all but the systems fixed, yielding laws for every way of characterizing the space of possible systematizations. And of course these could be combined to hold only one or two parts of the BSA fixed, yielding laws that are relative to, say, language *and* facts. Continuing in this way, there are maximally⁹

⁸ Cohen and Callender (2009) actually leave open the possibility that the rules of the competition might vary from one special science to another, but do not develop the possibility as strongly as they do their language relativity.

⁹ One very quickly gets infinitely many sets of laws even when the laws are relative to a single part. The “minimally relative” variants are supposed to yield a single set of laws, but their sense of minimality has to do with the fact that, internal to the particular variant, the actual laws are relative to nothing. The “maximally relative” are maximal in the sense that they make the laws relative to the greatest number of kinds of things (namely, all of the parts of the view that can vary). Within these grades of relativity are more or less permissive variants of the BSA in the sense that they are less or more (respectively) committed on ancillary philosophical issues. And, presumably, the more permissive variants are also, in a sense, more relative, but saying exactly

relative variants of the BSA according to which there are actual laws for *every* way of filling in *all* the details of the four-part model.

A more interesting version of Armstrong's objection is worried about relativity instead of, or in addition to, subjectivity. Considering relativity, Armstrong's point about disagreement becomes an argument against the possibility of a minimally relative variant of the BSA. If there are "irresolvable conflicts about the exact point of balance" between simplicity and strength, then, one might imagine, we must make the BSA at least competition-relative. If there is a good argument for the one acceptable way of filling in the details of a minimally relative variant of the BSA, that seems like the kind of thing that might have satisfied Armstrong. But if there is intractable disagreement among philosophers and scientists about any details in the BSA, then Armstrong cannot be wholly satisfied, and will fault the BSA proponent for either relativity or the unjustified endorsement of a minimally relative BSA.

Subjective laws are clearly problematic. Relative laws are less clearly problematic. Some relativity (so long as it can avoid descending into subjectivity) might be a good thing, as when it allows the BBSA to accommodate special science laws by making that to which the laws are relative parallel the particular interests of special sciences. However, in the extreme of a maximally relative variant of the BSA, it is not hard to sympathize with the relativity concerned Armstrong, and think that there is *too much* relativity. There may be laws suitable for any particular interests on such a view, but that isn't obviously a good thing. If anything, it's rather ob-

how (if that can be done at all) strikes me as more trouble than it's worth (certainly, at least, at this point).

viously a bad thing. Pick any true regularity and it will be a law in a system that is best relative to some choice of facts, systems, languages, and competition. As in the TSP, it has become impossible (albeit in a slightly different way) to distinguish accidentally true regularities from lawful regularities. Cohen and Callender argued that this isn't an issue in the BBSA since the problem only appears in those systems for which there is no interest, and we might accept this response since what relativity the laws have in the BBSA is clearly motivated. It's not enough, though, to motivate the relativity in a view. We must also argue for what specificity there is in a view. And the relativity and specificity should meet in the middle, with all relativity and specificity properly motivated and defended, and no part of the view left untouched after doing that.

This is, in brief, the claim argued for in Chapter 2. There I offer a more thorough development of these ideas, starting with the expansion and reconstruction of Armstrong's objection, running through the challenges posed by that new version of Armstrong's objection, and concluding with strategies for providing an answer. Importantly, the strategies discussed there are realized in part by the work done later in the dissertation.

1.3.2 Special Science Interests

It must be that what makes a set of laws the laws of a particular special science is the existence of a connection between those laws and what makes the special science the particular one that it is. What makes a special science (or any

scientific field) what it is is a variety of factors, mostly intellectual (we would hope), but also social, political, financial, and so on. Thankfully for our purposes we need not be concerned with everything that makes a special science, but rather just what matters for the purposes of determining the laws of the special science. While non-intellectual factors may be distally responsible for making some set of laws the laws of a particular special science, it seems safe to assume that the proximal causes of a special science having what laws it has are the intellectual interests of the practitioners of that special science.

The four-part model of the BSA tells us that what makes laws laws is the particular way that facts, systems, languages, and competition come together to yield a best systematization of the facts from which the laws may be extracted. So let there be the set of laws R determined relative to the quadruple (F_R, S_R, L_R, C_R) . And let there be some special science X whose practitioners have interests I_X . The laws R are *the laws of X* just in case the right relation holds between (F_R, S_R, L_R, C_R) and I_X . I take this to be the generalized form of the BBSA's strategy for accommodating special science laws in the BSA in the terms of the four-part model.

In the BBSA, the right relation has almost everything to do with the language(s) used to determine the best system. The interests of a field X are characterizable by the set of kinds K_X that are treated as basic in the field, and the laws R are the laws of X just in case L_R is the language whose basic predicates correspond to the kinds K_X . The possibility is left open in [Cohen and Callender, 2009] that different scientific fields may have different competitions, but it is not deeply explored there. Nor will it be here, except, in a fashion, in Chapters 4 and 5,

where I argue for *constraints* on any competition relativity that there might be. I am even less concerned with system relativity, except in Chapter 5, where I discuss how different constraints on the class of competing systems will influence the details of the best system competition.

I think the most interesting way to relativize laws is to facts. Implicit in the language relativity is a sort of fact relativity since a field can only systematize those facts that are expressible in the language it employs. But the fact relativity that comes with language relativity is incomplete; once a kind is admitted into our privileged language, *all* of its associated facts become subject to systematization. In Chapter 3, I argue that this incomplete fact relativity causes trouble when fields are interested in only a subset of the facts associated with a kind. Specifically, I develop in Section 3.3 an example based on the work of [Barlow, 1952] in which biology is concerned with photons, but must constrain its interests to only those photons interacting with biological systems. Fact relativity in the laws has the additional benefit, which I argue for in Section 3.4, that it makes better sense of the distinction between special and fundamental sciences that turns on universality.

1.3.3 Avoiding Trivial Systems

Recall the issue that gives rise to the TSP: There can be system-language pairs that are far and away the best according to the best system competition, but which yield laws that fail to distinguish between accidental and lawful regularities. Solutions to the problem have all focused on privileging a language (or languages,

as in the BBSA) that are assumed to not be in the group that would give rise to overwhelmingly best but problematic (i.e. *trivial*) systems. The challenge of answering the TSP was made greater by van Fraassen, who plausibly insisted that the laws of nature should be accessible and of interest to scientists, and noted that privileging the language of the perfectly natural properties, as Lewis had done to avoid trivial systems, renders the laws inaccessible and potentially uninteresting. The language privileging of the PDA and BBSA seems to succeed in answering both the TSP and van Fraassen's problems, but that success comes at the cost of subjectivity in the PDA and a glut of relative laws—and an imperfect solution to the TSP—in the BBSA.

The four-part model of the BSA suggests that language privileging might not be the only solution to the TSP (and van Fraassen's problems). Perhaps, by looking to our choice of facts, systems, or competition (or some combination of those possibly in conjunction with language privileging), an answer may be found that does not entail the less appealing consequences of Lewis' BSA, the PDA, or BBSA.

Right away we can rule out considering facts, for it is with respect to those that the distinction between accidental and lawful regularities must be made. To privilege some subset of the facts, or some altogether distinct set of facts, would just serve to change the set of true regularities that needs to be divided between the accidental and the lawful.

Considering systems on their own could be a promising endeavor. One solution to the TSP would be to simply exclude any system from competing that fails to distinguish between lawful and accidental regularities. This has a dubious ad-hoc

feel to it, though. It might be that such a distinction, being essential to laws, is permissibly built into an analysis of laws from the start. But, while it's seems true of the actual world that this distinction exists, we might think that there can be possible worlds where there is no such distinction. It may also be that this strategy makes the laws subjective insofar as its defense may depend on an appeal to our interests in the laws making the accidental-lawful distinction. Depending on one's leanings, directly constraining the set of competing systems might be less appealing than the bad making features of the extant language privileging solutions to the TSP.

Consider, then, the competition. Note that any pre-competition privileging, like that found in all the extant solutions to the trivial system problem, can be realized, often in more nuanced ways, by enriching the rules of the competition. Lewis' natural language requirement could be realized precisely by a rule in the competition that sets the score of any system-language pair in which the language is not that of the perfectly natural properties to zero. A rule that weights a system-language pair by the degree of naturalness of the language might be more interesting, at least if one's metaphysics allows for such a measure. Either way, these rules would be as troubled by van Fraassen's problems as Lewis' original approach, because they still depend on a sort of language privileging.

Central to the TSP is that it seems to have found a way of violating the need for a tradeoff between simplicity and strength through nothing more than a change in language. A purely competition based strategy, then, would be to give a more careful account of simplicity and strength, or other measures factoring into the best

system competition, that does not allow for such a violation of the tradeoff. I pursue such a strategy in Chapter 4. There I argue that taking seriously the PDA’s idea of evaluating systems and languages together as a “package” can block the TSP by replacing strict language privileging with relativity (as needed) to the best system competition.

1.4 Dissertation Overview

In this chapter I have reviewed three challenges to the BSA—Armstrong’s objection, the accommodation of special science laws, and the TSP—each of which would be worthy of extensive individual treatment. But to offer truly individual treatment would be impossible, as these challenges are intimately related. Armstrong’s objection, being concerned with a lack of objectivity in the laws of the BSA, is strongest only when proponents of the BSA try to accommodate the variety of interests found in scientific practice. The accommodation of special science laws seems sure to strengthen Armstrong’s objection by tying laws to a deeply heterogeneous array of scientific interests. And the TSP, being concerned with the languages in which competing systems are expressed, may have a straightforward solution in language privileging, but, depending on what language is privileged, that solution either breaks the connection between laws in the BSA and laws in scientific practice or strengthens that connection at the cost of strengthening Armstrong’s objection.

The consequence of these interrelations is that no one of these challenges to the BSA may be treated in full without eventually coming to the others. It is thus

largely a matter of preference that decides where one begins and how one progresses through these challenges. If your primary concern is the accommodation of special science laws, then you might begin there, solve the TSP along the way, and then upon encountering Armstrong's objection argue that resolving the two challenges is worth making the third worse. If you are taken with the TSP, then van Fraassen's problems will require you to grapple Armstrong's objection and the BBSA raises the bar for what we may expect a solution to the TSP to accomplish.

My primary concern is neither special science laws nor the TSP. Both are very interesting, but I think they are most interesting when approached in contexts *adjacent* to the one where we are concerned with the metaphysics of laws. What we say about special science laws should influence and be influenced by what we say about, for example, physicalism, reductionism, and philosophy of mind. From the TSP, analogies can and should be made to problems that appear in the epistemology of theory choice. Armstrong's objection, in contrast, appears only when we are concerned with the BSA. And in it is the basis for the tension that is exquisitely realized in the BSA between a presumption that the laws of nature should be objective and the appeal of an analysis of laws that parallels scientific practice. Whatever interests compel us to look at the BSA—from questions about consciousness to curve fitting—that tension will always be encountered, and so addressing it serves as the overarching goal of this dissertation.

Viewed in terms of the tension between objectivity and scientific practice in the BSA, the topic of this dissertation is quite open ended. We can move from one part of, or problem for, the BSA to another, and examine each in light of that tension.

Some of the parts or problems will be more interesting than others to examine in that light, and I hope that I have picked out some of the more interesting ones. There is, however, a particular interest that has led me to look at the BSA, and pick out what subjects I have, that is separate from the tension that will be my focus. That interest is in information theory, and it plays a significant, if somewhat quiet, role in determining the structure of this dissertation. It might be helpful before moving on to say a bit about how examining the relationship between the information theory and the BSA is where I hope this dissertation will lead.

To start, I will just say that information theory and the BSA seem like they were made for each other. The BSA is concerned with finding simple and strong or informative systematizations of the world. Information theory is, very roughly, a bit of mathematics that has two primary interpretations. The first takes it to be a theory of the limits to how efficiently communication can be carried out—it tells us how many sounds or symbols we need at a minimum to exchange our ideas, how much space it will take on computers or in the media of telecommunications to store and transfer images, audio, and any other data. The second interpretation takes it to be a theory of how to quantify information—relating it to the first interpretation,¹⁰ it says that carrying out more informative communication will require more space (or sounds or symbols, etc.) at a minimum, and less informative communication will require less space. These two interpretations cry out to be used in an attempt to formalize what we mean in the BSA by simple and strong/informative, respectively.

¹⁰ Strictly speaking, the ‘information’ interpretation does not depend on the ‘efficient communication’ interpretation, but relating the former to the latter is the quickest way to get at what is going on without digging in to the formal details of the theory.

But, of course, it's not that simple. How we address the issues surrounding the BSA that I consider in this dissertation will inform how we put information theory to use in the BSA. The possibility of the disagreements that give rise to Armstrong's objection makes it look like there won't just be one answer to how information theory and the BSA connect. The language privileging that does so much good work for accommodating special science laws and dealing with the TSP seems to render as irrelevant information theory's great ability to compare different ways of expressing the same content. As already discussed, these issues are a part of basic work that needs to be done in support of the BSA. Dealing with them has the added benefit of setting the stage for work on information theory's possible role in the analysis of laws.

Such is the secondary motive for considering the topics that I will in this dissertation. It will not be until Chapter 5 that information theory makes an appearance, and even then its connection to the BSA will not be definite. The primary motive of this dissertation is—as the title suggest—to pursue a balance between the need for objective laws and the relativity in the laws that might be needed to preserve a parallel with scientific practice. It may, however, be helpful to the reader who has some familiarity with information theory (and especially the various places it has been discussed in philosophy) to keep this secondary motive in mind. To the reader unfamiliar or uninterested in information theory, I say: “No worries.” Nothing of what is about to happen depends on such familiarity or interest.

1.5 Summary of Chapters

To conclude, let me offer a brief summary of what is to come.

Part I of this dissertation is concerned with the appearance of relativity in the BSA.

In Chapter 2, I take up Armstrong's objection from its first expression through the responses given by BSA proponents, and distinguish two compatible versions of the objection: one opposed to mind or subject dependence and the other opposed to relativity. The BSA, I argue, can answer both. Answering the anti-relative version of Armstrong's objection poses the more interesting challenge because it requires that the BSA be no more or less relative than is required by scientific practice. A spectrum of relativity is introduced with extremes of minimally and maximally relative variants of the BSA, and extant variants of the BSA are located on it. I then sketch what work remains to be done with respect to Armstrong's objection for BSA proponents depending on where in the spectrum of relativity they hope to locate their view.

Chapter 3 contributes to the project of motivating relativity in the BSA. Starting with the BBSA and its language relativity, I argue that the view suffers from two significant problems: (1) it will run afoul of cases of interfield interactions that blur the boundary between the basic kinds of individual fields (e.g. when photons are of interest to biology), and (2) it is unable to capture the distinction between fundamental and special science laws. I introduce an extension of the BBSA, the Kind and Fact Relative Analysis (KFRA), according to which laws are relativized

to kinds *and* the matters of fact that are being systematized. I go on to show how the KFRA can do all the same good work as the BBSA and answer the interfield interaction and fundamental/special distinction problems raised in the chapter.

Part II of this dissertation is concerned with the objectivity of the BSA, particularly how it might be achieved through constraints on the best system competition.

Chapter 4 examines the TSP and the language privileging that is the standard solution to it in [Lewis, 1983] and successive variants of the BSA. I offer a way of measuring strength that attends to both the competing system and its paired language that blocks the TSP. A related problem, the Problem of Immanent Comparisons (PIC) of [Cohen and Callender, 2009], is concerned with the commensurability between competing system-language pairs. The PIC, as it happens, has the overly strong conclusion that systems cannot be compared at all unless they are expressed in the same language. I argue that this is not right, citing the Akaike Information Criterion as a prime example (among others) of a measure that can (and plausibly might) be used to compare systems expressed in different languages. However, TSP and PIC both are based in the problems of language sensitivity that are well known throughout philosophy. These problems cannot be overcome, but, because the language sensitivity lies in the measures to be used in the best system competition, I argue that the language relativity of the BBSA can and should be replaced with competition relativity.

Chapter 5 steps away from explicitly addressing the three challenges to the BSA discussed in this chapter. Instead, the chapter focuses on the question of what, beyond simplicity and strength, might matter for making a system(-language pair)

the best. I argue that the best system(s) will be one(s) that make the actual world the one that scientists would most like to be in, and motivate this requirement by looking at the “induction friendliness” of a world given a particular system-language pair as the output of science’s best inductive practices. I propose quantifying induction friendliness with the information theoretic measure of mutual information, and then highlight a number of ways that the BSA might benefit from introducing mutual information into it. Appendix [A](#) further develops the idea of induction friendliness through the reconstruction of a classic theorem of information theory from [[Shannon, 1948](#)] in the context of a toy model of the BSA.

The dissertation concludes in Chapter [6](#). There I offer a summary of the dissertation in the course of describing the variant of the BSA—which I call the Relative and Objective, on Balance Analysis (ROBA)—that emerges from the work that has been done.

Part I

Relativity

Chapter 2: Interest Relativity

The BSA has it, roughly, that a regularity is a law just in case it appears in the best systematization of all the particular fundamental matters of fact, where “best” means the simplest and strongest on balance. A standard critique of the view—which I call *Armstrong’s objection* after its appearance as the first objection to the BSA in [Armstrong, 1983]—is that what makes a system best is only determinable relative to a subject’s conception of “best”. So, in conflict with strong intuitions that laws of nature should be objective, the laws as described by the BSA are subjective. Despite Armstrong’s objection, recent proponents of the BSA (or variations of it) have tended to embrace versions of the idea that the best system is the best *for us* (e.g.: [Loewer, 2007, Cohen and Callender, 2009, Frisch, 2011]).

I argue that Armstrong’s objection read as being concerned with the subject independence of the laws is deficient, since BSA laws need not depend on the existence of subjects. The specifics of Armstrong’s objection as it is first presented suggest an alternative, and more threatening, reading according to which the concern is with the possibility of profligate laws determined relative to the variety of ways of filling in the details of the BSA. The two readings of Armstrong’s objection are compatible and, taken together, pose a substantial challenge to the BSA with

a range of possible responses. On one end of that range is the maximally relative BSA, where there are sets of laws for every possible way of filling in the details of the view. On the other end are minimally relative variants of the BSA, where limits on permissible ways of filling in the details of the view are sufficient to pick out a unique set of laws. I advocate a compromise.

This chapter proceeds as follows. In Section 2.1, I review the standard subject (in)dependence reading of Armstrong’s objection and the case against it. In Section 2.2, I introduce the anti-relativity reading of Armstrong’s objection and argue that a proper response depends on making the laws of the BSA relative exactly as much as is required by paying due respect to scientific practice. Section 2.3 is concerned with how the relativity of the laws of the BSA can come in degrees, and how those degrees of relativity correspond to a range of responses to the combination of the subject independence and relativity versions of Armstrong’s objection. In Section 2.4, I look at how one might motivate relativity in the laws by appeal to scientific practice. In Section 2.5, I briefly develop three general strategies for identifying limits on the relativity of the BSA.

2.1 Armstrong’s Objection Against Subjectivity

Let us begin by looking at Armstrong’s full statement:

The first objection that may be made to the Systematic solution [i.e. the BSA] is that an element of subjectivism remains. We have already noticed that it has to involve our standards of simplicity, which, even

granted that they are shared by all rational mankind, may not be shared by other rational creatures. The same point seems to hold for standards of strength. Lewis also refers to ‘our way of balancing’ simplicity and strength. May there not be irresolvable conflicts about the exact point of balance?

[[Armstrong, 1983](#), p. 67]

We are told at the outset that the problem is the subjectivity of the BSA laws, so let us make that part of the argument explicit:

SBL: The laws of the BSA are subjective.

NSL: Laws cannot be subjective. [unstated premise]

Conclusion: (From SBL and NSL.) The BSA cannot be a correct analysis of laws.

The argument is certainly valid. How to read “subjective” is unclear in the reconstructed argument just as it is in the passage from Armstrong. For the moment we will assume that the relevant reading of “subjective” is the one concerned with the dependence of the laws for their content and existence on the presence of minds or subjects.

The unstated premise NSL seems plausible. Proponents of the BSA might try to reject NSL from the start. One strategy for this may be to argue that it is an intuition that does not apply to laws of the BSA, in the style of the rejection in [[Beebe, 2000](#)] of “governing” as necessary to a conception of laws. But simply rejecting NSL weakens the standing of the view in comparison with competing accounts of

laws that can accommodate the intuition. Most notably, proponents of the BSA do not reject NSL: [Lewis, 1994], [Loewer, 1996], and [Cohen and Callender, 2009] all (to be discussed at greater length below) direct their attention towards rebuffing the claim that the relation between the BSA and the standards of simplicity, strength, and their balance, makes the laws mind or subject dependent in a problematic way.

The main challenge for an employer of this objection to the BSA is to defend SBL. And, indeed, the remainder of the passage from Armstrong is dedicated to explaining just how the laws of the BSA are subjective. The matter turns on the potential for “irresoluble conflicts” between rational seekers of laws. These conflicts—call them *real disagreements*—are ones that persist through some ideal end of the progress of science, where nothing new could happen to decide the question and to suspend judgement would be to leave that corner of science incomplete. Armstrong suggests that a source of real disagreement may be found in the choice of measures for simplicity, strength, and their balance. The exact source of a real disagreement does not matter for our present purposes; it is enough if they occur at all. So let us grant Armstrong that there is such a real disagreement in the form of Scientist 1, who says that the BSA laws are R_1 , and Scientist 2, who says that the BSA laws are R_2 , where R_1 and R_2 are not equivalent.¹ From this we are supposed to extract the claim that the laws of the BSA are subjective.

S1R Scientist 1 says that the BSA laws are R_1 .

S2R Scientist 2 says that the BSA laws are R_2 .

¹ I use “ R ” here to stand for a set of (possible) laws because later on “ L ” will be used for the language(s) used to express systems, and sets of laws and merely candidate sets of laws share the quality of being sets of *Regularities*.

I12 The laws R_1 and R_2 are not equivalent.

SBL (From S1R, S2R, and I12.) The laws of the BSA are subjective.

This argument is invalid. They each *say* that the BSA laws are one way or another, but that does not make them so. The BSA is, despite the inspiration it draws very directly from the epistemology of science, a realist metaphysical analysis of laws of nature. There are no variables in the BSA (like the balance between simplicity and strength) that are metaphysically free to be set at will by some subject, and so the laws of the BSA are not subjective in the sense that they depend on subjects.

I will refer to this point about the BSA as the *brute metaphysics* response to Armstrong's objection. Employing it is to put one's foot down and insist that there just is some metaphysical fact of the matter about what the laws are according to the BSA. It does not matter what anyone says or thinks about the laws, nor does it matter if anyone exists to say or think anything about the laws. Lewis adopts this position when he writes

I am talking about how nature [...] determines what's true about the laws and chances. Whether there are any believers living in the lawful and chancy world has nothing to do with it.

[[Lewis, 1994](#), pp. 481–482]

Cohen and Callender similarly insist in their discussion of their variant of the BSA, “relativized MRL”, that

relativized MRL does not make the laws subjective in the sense that they depend for their existence on subjects. This is because the laws (relative to basic kinds K) that hold of a world w would satisfy relativized MRL's criterion for lawhood... whether or not there are subjects in w (or any other world).

[Cohen and Callender, 2009, p. 30]

Even Ramsey, whose view is summarized in [Lewis, 1973] as “a counterfactual about omniscience”, writes that what matters to the laws are “the facts that form the system in virtue of internal relations, not people's beliefs in them” [Ramsey, 1928, p. 132].

It is possible to formulate a version of the BSA that is thoroughly subjective, but it is by no means necessary. It is certainly not the case for the version of the BSA that was Armstrong's target that its laws were subjective in the sense that they depend on minds/subjects. As long as we read Armstrong's objection this way, it is less an objection than a stern reminder that we should be careful in our statements of the BSA and its variants.

2.2 Armstrong's Objection Against Relativity

Taking a second, sympathetic, look at Armstrong and Lewis tells against interpreting Armstrong's objection as being exclusively about mind/subject (in)dependence. Armstrong never accused the BSA of being subjective, but only that, on account of (real) disagreements in scientific practice, “an *element of* subjectivism remains”

(emphasis added). In light of this and what was said in the preceding section, it may be worth giving Armstrong the benefit of the doubt by not assuming he was objecting simply to a subject dependence in the laws.

Consider next how Lewis responds to Armstrong. Lewis' expression of the brute metaphysics sentiment is not presented as a response to Armstrong, but only as a general remark about the BSA. He does, however, write

The worst problem about the best-system analysis is that when we ask where the standards of simplicity and strength and balance come from, the answer may seem to be that they come from us. [...] Maybe some of the exchange rates between simplicity, etc., are a psychological matter, but not just anything goes. If nature is kind, the best system will be robustly best—so far ahead of its rivals that it will come out first under any standard of simplicity and strength and balance.

[[Lewis, 1994](#), p. 479]

This *nature is kind* response to Armstrong's objection stops just short of denying that there will ever be the problematic sort of disagreement with which Armstrong is concerned. If nature is kind, we can deny the conjunction of S1R, S2R, and I12. But what if nature is unkind? Lewis says just down the page: "in this unfortunate case there would be no very good deservers of the name of laws". This may be true if nature is *wildly* unkind, such that no account of laws can make sense of what is going on. The real threat to Lewis is that nature is unkind in just the right sort of way to make the BSA laws untenable but the laws as described by some other view

acceptable; e.g. stochastic governing laws might allow for quite misleading/unkind runs of events in nature without undermining how deserving its laws are to be so called, and similarly for the impoverished possible worlds of [Tooley, 1977] and others. Such a particular threat must be warded off with more than wishful thinking.

What is most telling against the subjectivity reading of Armstrong's objection in this response from Lewis is the suggestion that the laws of the BSA may depend in principle on law seeking subjects. No amount of kindness in nature will break the subject dependence of the laws if it is there. The same issue arises with the *rigidification* response, considered briefly in [Lewis, 1986], according to which the standards of simplicity, strength, and balance employed in the BSA are fixed to "our actual standards" [Lewis, 1986, p. 123]. Giving Lewis the benefit of the doubt, there must be something other than subject dependence addressed by these responses.

What Armstrong and Lewis share is a concern with the presence of multiple different sets of laws, not subject dependence. The intuition underwriting Armstrong's objection is less the unstated premise (NSL) that the laws cannot be subjective (i.e. dependent on subjects), and more that the laws must be *objective*, which may also be interpreted as a requirement that the BSA pick out a unique set of laws. Armstrong's objection, then, is

S1R Scientist 1 says that the BSA laws are R_1 .

S2R Scientist 2 says that the BSA laws are R_2 .

I12 The laws R_1 and R_2 are not equivalent.

DBL (From S1R, S2R, and I12.) The BSA yields multiple sets of laws.

USL There must be a unique set of laws.

Conclusion (From DBL and USL.) The BSA cannot be a correct analysis of laws.

But this is still invalid. Now, as before, we can appeal to a version of the brute metaphysics response to block any concern about extra sets of laws: The BSA will say that there is exactly one real set of laws, and so at least one of Scientist 1 and Scientist 2 is wrong.

The remainder of this section is broken up into three parts. The first is concerned with salvaging what one may from (and for) this version of Armstrong's objection by forcing the BSA to accommodate through relativity the disagreement between Scientists 1 and 2. In the second part it is shown how the BSA can be simultaneously relative and non-subjective. The third part is concerned with what would count as a response to Armstrong's objection.

2.2.1 Respecting Scientific Practice

It is easy to be unhappy with the version of the brute metaphysics response that was just employed. Proponents of the BSA should be unhappy because, in requiring that one of the two scientists be wrong, the response risks abandoning the close connection between the metaphysics of laws and the epistemology of the search for laws that is such an appealing feature of the BSA. The close connection between the BSA and scientific practice is, in fact, reliably treated as a virtue. But that virtue is in conflict with the brute metaphysics response as it has just been used: to dismiss some part of the practice of science (namely, the part associated

with at least of one Scientists 1 and 2).

Opponents of the BSA should also be unhappy with this version of the brute metaphysics response because it denies that the claims about laws made by the two scientists are the sort of thing that could count against the BSA. Suppose, for example, that a petulant employer of the brute metaphysics response just insists that the details of the BSA—e.g. the preferred balance between simplicity and strength, and what measures are used to determine the simplicity and strength of a system—are such that the laws are R_1 . An employer of Armstrong's objection would like to be able to object: Why aren't the laws R_2 ? Nothing has been said, and, by assumption, nothing can be said, in favor of Scientist 1's settings of the details BSA over Scientist 2's. If BSA proponents are as interested as it seems in a close connection between the details of the BSA and the particulars of scientific practice, on what grounds could the position of Scientist 2 be dismissed as they have been? Some other petulant proponent of the BSA may decide to favor the details that accord with Scientist 2 over Scientist 1, or endorse some third option that is incompatible with anything in scientific practice. In any of these cases, the brute metaphysics response as just presented will ward off any criticism.

So proponents of the BSA and employers of Armstrong's objection share an interest in thwarting this version of the brute metaphysics response. Its arbitrariness and apparent disinterest in scientific practice is anathema to the methods of philosophy of science. The solution is clear: Everyone should (and should be willing to) endorse the principle that the consequences of the BSA (or any account of laws) fit, parallel, conform to, answer to, (or, to just pick a word and stick with it)

respect scientific practice. Thus we add to the argument of Armstrong's objection the premise

RSP An account of laws must respect scientific practice.

What does RSP actually require? Saying exactly is difficult. van Fraassen explains that an account of laws

should make it plausible that laws of nature are the truths which science aims to discover. My phrasing should not be too strictly or prejudicially construed. If the account makes it plausible that the laws, as defined, are part of the theoretical description of the world provided by science in the long run, if all goes ideally well—that is enough.

[van Fraassen, 1989, p. 55]

For our purposes it should be emphasized that RSP requires that, whenever scientists say X about laws, the BSA is in a position to agree that X , or at least illuminate why X was said. For example, we may not have yet discovered the true BSA laws of the world, but if the BSA were run given just what we know, then it should say pretty much what we already think. Another example: The premise USL that there can only be one set of laws might be essential to the conception of laws found in scientific practice. In that case, USL follows directly from RSP, and the BSA should not generate multiple sets of laws.

Unlike in the case of the petulant employer of the brute metaphysics response, the BSA is now able to fail by its own lights if satisfying RSP leads to contradictory commitments. Of particular interest to the discussion of Armstrong's objection

is what RSP requires of the BSA in the event of real disagreements in scientific practice—recall from the previous section that those are ones that persist through some ideal end of the progress of science—like in our example of Scientists 1 and 2 and their not equivalent sets of laws R_1 and R_2 . Satisfying RSP will require the BSA to say that the laws of R_1 are real (not just possible/candidate) laws *and* the laws of R_2 are real laws; that is how the inference from S1R, S2R, and I12 to DBL is properly made. But, on pain of contradiction (since R_1 and R_2 are not equivalent), it must be qualified that each set of laws are the laws *relative to something*.

2.2.2 Relativity and Subjectivity

The first reading of Armstrong’s objection was opposed to laws that were dependent on subjects, but was overcome by way of the brute metaphysics response. But after adding RSP and requiring the laws be relative, it is possible that the first reading of the objection becomes a problem again. If the laws of the BSA are determined explicitly relative to the peculiarities of particular scientists, then it seems that they do actually depend on subjects. Before moving on with the discussion of this second reading of Armstrong’s objection, it is important to ensure that the BSA has not already succumbed to the first reading. To do that we must answer the question “To what are the BSA laws relative?”.

Scientists 1 and 2 presumably don’t just disagree about the laws for no reason. In Armstrong’s example three points of disagreement are suggested: They might disagree on the meaning of “simple”, on the meaning of “strong”, and on the

weighting between them that determines the overall goodness of a system. Since we have denied the possibility that there is a correct side in these disagreements (supposing that any one, or some combination, of them is the reason for the real disagreement about the laws between Scientists 1 and 2), I will say that the reasons for real disagreements are differences in *interests*. Scientist 1 may be interested in simple laws, while Scientist 2 is interested in strong laws. Scientist 1 may be interested in what the laws are when using one measure of simplicity, while Scientist 2 is interested in the laws when using some other measure. Since their disagreement is real, neither could ever convince the other to change their interests. But that's okay at least because each can appreciate that, were their interests to coincide, they would agree on the laws.

We thus want the laws to be interest relative. Interest relative laws guarantee satisfying RSP as long as any differences of interest may be reflected in how the details of the BSA are filled in. But the laws of the BSA cannot be explicitly interest relative since interests are had by subjects and that would make the laws of the BSA subjective in conflict with the first reading of Armstrong's objection. Suppose that the BSA run with details D_1 yields the laws R_1 . And, if run with details D_2 , the BSA yields the laws R_2 . Nothing about these relations realized by the BSA, between the details D_i and corresponding laws L_i , result in laws that are subject dependent. The R_i are out there just as the non-relative BSA laws were when we first considered how the brute metaphysics response blocks the first reading of Armstrong's objection. What has changed is this: There are not just *the laws*. There are the laws *of Scientist 1*, and the laws *of Scientist 2*, and so on if need be.

To say that the laws R_1 are the laws of Scientist 1 certainly makes it sound like the laws depend on Scientist 1, but the dependence need not be present. The laws R_1 are those relative to the details D_1 , whether there are any subjects or not. When Scientist 1 comes along with interests I_1 , those interests point to the already present laws R_1 —determined relative to details D_1 that correspond to the interests I_1 —as the laws of interest to Scientist 1. This is a generalized version of the picture I assume one must have in order to reconcile the relative laws of [Cohen and Callender, 2009] (to be discussed further in Section 2.4) with their claims to subject independence.

We can say that the laws are interest relative insofar as there are different ways of filling in the details of the BSA (to which the laws are actually relative) corresponding to the different interests of scientists. In this way the laws of the BSA can be interest relative to satisfy RSP without succumbing to subjectivity.

2.2.3 Responding to Armstrong

Now we may return to addressing the second reading of Armstrong’s objection. Can the laws of the BSA be relative and guarantee a unique set of laws (as required by USL)? Yes, in two possible ways. The first way is something like what was suggested by Lewis in the “nature is kind” response. It may be, if nature is kind, that all the sets of laws (each relative to the different ways of fixing the free variables of the BSA) turn out to be equivalent, and thus there is effectively just one set of laws yielded by the BSA. But we have already set aside this possibility as too much

wishful thinking. The second way is if, in fact, there are no real disagreements or differences of interest in scientific practice. However plausible this may or may not be, it amounts to a denial of the conjunction of S1R, S2R, and I12, and so simply asserting it is not a satisfying response to Armstrong's objection.

If there are real disagreements in scientific practice, then the BSA will fail to satisfy USL when it tries to satisfy RSP by admitting relative laws that accommodate those disagreements. But, at the same time, the justification for USL will be undermined by RSP since it is scientific practice itself that is requiring that there be multiple sets of laws. In light of asserting that there is real disagreement among scientists (premises S1R, S2R, and I12) and that the BSA must respect scientific practice (RSP), one cannot insist on there being a unique set of laws (USL). It is not that there can only be one set of laws, but rather that there should be only be as many sets of laws as is required by respecting scientific practice. Since multiple sets of laws are yielded by allowing the laws of the BSA to be relative to corresponding details, we may say that the BSA is more or less relative when it admits more or fewer sets of laws (we'll see, though, that this connection comes apart a bit in [2.3](#)). One would need to introduce (and defend) principles that supersede RSP in order to request or endorse any more or less relativity (or more or fewer sets of laws) than what is required to satisfy RSP.

Before moving on, let us pause to take stock of what has been said and see where it may take us. There are two (related) intuitions underlying Armstrong's objection, each a version of the intuition that laws should be objective. The first is that the laws cannot be subject dependent, and is satisfied by the BSA with little

issue. The second is that there should be just one set of laws. Making the second intuition problematic for the BSA requires endorsing the claim that an analysis of laws must respect scientific practice *and* guaranteeing that there are real disagreements in scientific practice that would force (on account of respecting scientific practice) the BSA to allow for the existence of more than one set of laws. If the two intuitions are to be defended, it is presumably by examination of scientific practice. Therefore, insofar as scientific practice requires more than one set of laws, the second intuition must be weakened. It is not that there must be just one set of laws. Rather, there should be only as much relativity in the laws as is required by respecting scientific practice. Responding to Armstrong's objection is thus a matter of finding just the right amount of relativity in the laws of the BSA. Determining how relative the laws of the BSA *should* be is a project beyond the scope of this chapter, but in the remaining two sections I address how the BSA *can* be more or less relative, and what that means for present and future variants of the view.

2.3 Degrees of Relativity

The BSA may be illuminated by a story in which a scientist dies and appears before God, who asks if there is anything the scientist would like. The scientist replies "I would like to know how the world works". God accepts the request and begins "At time and position $\langle t_1, p_1 \rangle$ such-and-such properties obtained, and at $\langle t_1, p_2 \rangle$ such-and-such other properties obtained, and—" at which point the scientist interrupts and says "Sorry, if that is how things *actually* work that is all well and

good, but is there a *pithy* version?”. God’s reply to that question is to provide the scientist with the (BSA-style) laws of nature. [Beebee, 2000] introduces a version of this story in order to highlight the fact that the laws of the BSA are non-governing. It also provides a nice context for thinking about how the laws of the BSA are relative. All the same questions are raised when the scientist asks for the “pithy version” as when it is said that the laws are drawn from the “best system”. And if two scientists appear before God and receive different laws, that is akin to our earlier thinking about real disagreements between scientists at some ideal end of science.

To emphasize the possibility of different responses, we can join Armstrong in worrying about whether or not the standards of simplicity, strength, and balance are the same for “other rational creatures”. Let’s call these other rational creatures Martians, and consider a version of the story in which our (human) scientist is joined by a Martian scientist. They both want a pithy version of how the world works, and for each God is happy to provide it, which is to say, the laws. Thinking of Armstrong’s objection as being concerned with relativity, the question to ask is “How different are the respective laws?”.

Maybe the Martian’s brain is much bigger than a human’s, and thus there need not be as great a trade off of strength in favor of simplicity. Such a scenario deprives the human scientist of any sort of cognitive aid without good reason. Science is conducted with many cognitive aids—computers are a stand out example—so let the human scientist possess whatever such aids they may. The Martian will probably have cognitive aids as well. And, if the aids of one are superior to those of the other,

the superior aids could be reproduced and the new user could be trained to use them (presumably they have an eternity in God’s audience, but we should also imagine that this is what would happen in the normal course of science). If one scientist can measure, calculate, or by some other means, have access to, or the ability to process, some piece of information, then just one of the following will be true: (1) Any other scientist may be trained to appreciate the measurement(/calculation/what-have-you) (given sufficient time and resources) or (2) The measurement(/...) is not worthy of being incorporated into scientific practice, precisely because it fails (1). If anyone can overcome some limitation, then the means by which they did so should be shared. In the long run, everyone should be subjected to precisely the same limitations, augmented from their start as needed with equalizing education and engineering. Everyone must reject anything failing (1). But if (1) holds, the only grounds for dismissing the relevant means and resultant information is disinterest. Despite the introduction of Martians (or any “other rational creatures”), there being differences in the laws continues to be a matter of capturing whatever differences there are in interests.

Differences in the laws are realized by different ways of filling in the details of the BSA. Earlier this was likened to setting a metaphysically free variable like the balance between simplicity and strength. Let us make that more explicit with a toy model of the BSA that is as naive and straightforward as possible: A system S is judged best (or not) according to its goodness $Good(S)$ which is the sum of its simplicity $Simp(S)$ and strength $Str(S)$ weighted by a balance factor b with a real

value from 0 to 1 such that

$$Good(S) = b \times Simp(S) + (1 - b) \times Str(S) \quad (2.1)$$

Suppose further that the measures *Simp* and *Str* are fixed so that there is only the one free variable b . The BSA relativized to b says that there are laws $R_{b=x}$ for each way of substituting a value x in for b .

We can ask again: How different are the respective laws? Without saying anything about what the laws look like in detail, it will be hard to talk of the difference between any two sets of laws $R_{b=x}$ and $R_{b=y}$ (with $x \neq y$) beyond simple equality or inequality. We can make some progress, though, in talking about the degree of relativity of the set of all law sets $\mathcal{R} = \{R_{b=x} \mid \text{for all } x \in R \text{ such that } 0 \leq x \leq 1\}$.

When Lewis introduces the “nature is kind” response it is meant to minimize relativity by having \mathcal{R} have a single member as a result of the fact that $R_{b=x} = R_{b=y}$ for all x, y . As soon as there is not one all encompassing equivalence class of law sets there is an abrupt jump from there being effectively no relativity in the laws to there being some (e.g. if there is some critical boundary value of balance v such that $R_{b=x} = R_{b=y}$ iff either $v \leq x, y$ or $x, y < v$).

In the limit where $R_{b=x} \neq R_{b=y}$ for all x, y such that $x \neq y$, there are uncountably many distinct law sets. Relativity goes down from that point in one of two ways. The first is as we’ve built up to this point in reverse: If there are any equivalent sets of laws, then it seems there is less relativity. For example, if there

is some value of balance v such that $R_{b=x} = R_{b=y}$ iff $v \leq x, y$. Then, unless $v = 0$, there will still be uncountably many nonequivalent law sets. But the claim that the laws will come out the same as long as a weight of at least v is given to simplicity seems to constitute a substantial reduction in relativity compared to every value of b yielding different laws.

The second way for relativity to go down is if there are additional constraints on the value of b . It is assumed that b has some real value ranging from 0 to 1. What if it could be shown that the extreme values are not allowed, that b must have a real value *between* 0 and 1? Again, there are still uncountably many nonequivalent law sets—only two distinct law sets have been excluded. But, again, the claim that the goodness of a system cannot be determined entirely by its simplicity or entirely by its strength seems substantial as a constraint on relativity.

The implausibility of the “nature is kind” scenario is tempered by the possibility of identifying equivalencies *and* limiting the set of permissible values b . There will be just one set of laws as long as all of the balance relative laws are equivalent for any permissible value of b . That is *minimal relativity* (at least in the limited context of relativity to b). It is, presumably, what one has in mind when they insist that the laws of the BSA be objective in the strong sense of being opposed to the possibility of any disagreement.

Is there a point at which the relativity of the BSA cannot be increased? It cannot just be that maximum relativity is achieved where there is an uncountable infinity of nonequivalent law sets. As suggested above, it seems that there can be qualitatively less relativity even when an infinite number of nonequivalent laws has

been reduced a finite amount. Going in the other direction, it seems we can *increase* relativity by adding additional variables to the BSA. Consider an extension of our toy model where the goodness of a system involves a balance of simplicity, strength, *and* fit. Now there are two free balance variables b_1 and b_2 that can take any real value from 0 to 1 as long as $(b_1 + b_2) \leq 1$, and

$$Good(S) = b_1 \times Simp(S) + b_2 \times Str(S) + (1 - b_1 - b_2) \times Fit(S) \quad (2.2)$$

The laws of this BSA are more relative than the laws of the BSA without fit, since here the laws are relative to the values of both b_1 and b_2 . But, in the limit where there are no equivalent law sets, there are still only uncountably many nonequivalent sets of laws. For another example, the same thing would be true if we removed the constraints that ensure the weights on each of simplicity, strength, and fit, have values between 0 and 1 inclusive.

We can keep increasing the relativity of the laws by letting more and more of the BSA be variable. One step in that direction would be to allow for many ways of measuring each of simplicity, strength, and fit. But the potential for variability in the details of the BSA runs much deeper than just how we measure and balance the good-making features of a system. As noted in the Chapter 1, Lewis' BSA and the variants of the view introduced since all fall under a four-part model: First, there are the matters of fact to be systematized. Second, the various candidate systems. Third, all the languages that might be used to express the facts and candidate systems. Fourth and finally, the best system competition itself, which serves as a

function from the sets of facts, systems, and languages, to a system-language pair (or set of such pairs) that is the source of the laws. When we consider laws that are relative to the exact balance of simplicity and strength, this is a worry about the rules of the best system competition. But there could also be relativity to how the set of facts to be systematized is populated. And relativity to the set of systems that will be considered as candidates in the competition. And relativity to the set of languages in which the systems and facts may be expressed.

Maximal relativity is achieved when all four parts of the BSA—the facts, systems, languages, and competition—are wholly unconstrained beyond the loose structure that brings those four parts together. On such a view, in order for some regularity r to be a law it only has to be that there exists some fact(s), some system(s), some language(s), and some competition function(s) such that the r is made a law relative to the specified fact(s), system(s), language(s), and competition(s). At this point, reductions achieved through the kindness of nature—namely, through equivalencies among law sets that have different settings of the BSA variables—do not reduce relativity away from the maximum. Suppose that despite all the allowed variability it still is the case that all the sets of laws are equivalent. Then the maximally relative BSA would also be a minimally relative BSA.

But nature is probably not as kind as that. Between the maximally relative BSA and the minimally relative BSAs there is a wide range of BSA variants that feature a middling amount of relativity. And hopefully among those, or among the extreme variants if that is how things go, is a variant that captures exactly the relativity, no more and no less, that is required by respecting scientific practice. If

there is, then Armstrong's objection as it has been reconstructed in this paper can be answered.

2.4 Required Relativity

It has been left open so far whether or not there actually is any relativity required by respecting scientific practice. A proponent of the BSA who seeks a minimally relative variant of the view has to either identify limits on free variables of the BSA, hope (or, preferably, show) that nature is kind, or do some combination of both enough to ensure that whatever relative laws there are are all equivalent to each other. There are lots of different minimally relative BSAs, and one cannot just be picked. The starting point for any proponent of the BSA should be the unique maximally relative BSA, and any limits on relativity will be accomplished by arguing down from there. If the goal is minimal relativity, then there is a lot of work to be done. But minimal relativity is not the goal if respecting scientific practice requires that there be some relativity in the laws. In this section, I consider the ways that relativity might be required by respecting scientific practice.

To start, return to the story of scientists appearing before God. Suppose that it is a physicist, an economist, and a biologist that are having an audience with God. When each asks for the pithy version of "how the world works", they are each asking different questions, since each is interested in different parts of the world. Thus (barring an incredibly kind nature, and only insofar as these fields have laws) each will receive different laws. To the physicist go the laws of physics, which are

the laws determined relative to the setting of variables of the BSA that matches the interests of the physicist. To the biologist the laws of biology, which are the laws determined relative to (...) the interests of biologists. And similarly do the laws of economics go to the economist.

Some concern should be paid to the possibility of a kind nature, or just the confidence we have in the differences between two fields, when there are substantial overlaps in interests. A neurobiologist and a cognitive psychologist, each operating under the broad heading of “cognitive science”, may come at the same problem from different directions. To the extent that it really is the same problem, it is not clear if their respective laws will be different. At that point it might be worth it to think about differences in the sets of laws beyond them being equivalent or not.

For our purposes it is enough to believe that there are different scientific fields with interests different enough that they would receive different laws from each other. Even if one doubts the autonomy of different scientific fields from fundamental science with respect to laws², we might expect that whatever principles bridge the divide between the fundamental laws and the laws of some non-fundamental field would be included in God’s reply, and thus could be counted among the BSA laws for the field in a way that distinguishes the field’s laws from the fundamental laws.

The maximally relative BSA that was presented in the previous section is a generalization of the [Cohen and Callender, 2009] variant of the BSA, the “Better Best System Analysis”, or BBSA. The BBSA is a language relative variant of the

² There are some good reasons to not have such doubts. For examples: [Lange, 2004] argues on behalf of the autonomy of functional biology, and [Callender and Cohen, 2010] argues for the compatibility of autonomy and a certain brand of reductionism in the context of the BSA.

BSA: There is the set of supervenient kinds—the set of all kinds that supervene on the true fundamental kinds—and a best system competition is run for each subset of the set of supervenient kinds, where the facts and all competing systems are expressed in the language that has the particular set of kinds as its basic predicates. The interests of a field (for the purpose of picking out the field’s laws) are identified with the kinds that are treated as basic by the field, so the laws of a field are the laws that are determined relative to the set of kinds that happens to be the set of interest to the field.

In the BBSA we have language relativity (though perhaps not *total* language relativity, since the details of a language are not exhausted by specifying the basic predicates of the language) on the grounds that it is required to accommodate the differences in interests of special sciences as those interests are directed at laws. Competition relativity is also a part of the BBSA, though less discussed than the kind relativity. It is explained that

Ecologists are not looking over their shoulders at the simplicity, strength and balance metrics of physics. They are using their own metrics tailored to their own field.

[Cohen and Callender, 2009, p. 24]

Similar reasoning probably extends to what counts as a candidate system; e.g., ecologists probably are not counting field theories among their candidate systems.

What about fact relativity, the last of the four major parts of the maximally relative BSA? If a field cares only about certain kinds of things (say, *trees*) then

there will be certain facts that will go unconsidered in determining the laws of the field simply because the language of the field is incapable of expressing all of the facts. But fact relativity will become important if there are two fields that cannot distinguish their respective domains of interest *except* by specifying which facts they will or will not try to systematize. This seems to happen when looking at the differences between the relativistic and non-relativistic varieties of classical mechanics: Classical mechanics of either variety is concerned with the same sort of basic kinds, and so the differences in the interests that lead to their respective formulations are unlikely to be captured by language/kind relativity. Non-relativistic classical mechanics is concerned with just those things whose speeds are trivially small compared to the speed of light. Relativistic classical mechanics is concerned with things moving at any speed. What distinguishes them is a fact relativity. The former is a systematization of the subset of matters of fact in which nothing is traveling at relativistic speeds, while the latter systematizes all of those facts and facts in which things are moving at relativistic speeds (both presumably ignore some facts that are outside the domain of classical mechanics generally).

If respecting scientific practice requires that there be laws for every scientific field, then it seems the BSA may be pushed quickly to maximal relativity. And that is before we consider the possibility of real disagreements within fields (though perhaps we have, if the relativistic v. non-relativistic classical mechanics example is read as involving an intra-physics disagreement). So, does the maximally relative BSA answer Armstrong's objection? If it does, then those looking for something closer to perfect agreement on the laws will be left to hope with Lewis that nature

is kind. But Armstrong’s objection against relativity, with an underlying intuition against a glut of laws, requires that the BSA capture exactly the amount of relativity that is required by respecting scientific practice, *no more and no less*. As much as relativity may have been motivated in the above, nothing said so far guarantees that there aren’t also limits on relativity to be found in scientific practice. Maximal relativity in the laws is maybe more plausible than minimal relativity, but there is still room for a better BSA to be found somewhere in the middle.

2.5 Limiting Relativity

Perhaps the easiest way to limit relativity in light of the preceding discussion is to deny that the BSA needs to accommodate laws in all fields. But to do so would be to make two errors. The first is to miss that, even if we can point to a field and say with confidence that laws are not a part of the practice of the field, laws may still become a part of the practice of the field. The second is to miss that, even if we cared only to identify laws in one field, there remains the possibility of relativity being required to respond to real disagreements within the practice of that field.

The great challenge of limiting relativity is that what limits are imposed must be *universal*; that is, they must constrain what all law seekers do whatever their interests may be. If someone claims to be seeking laws, but violates a universal limit, then we should be in a position to argue that they are mistaken in claiming that it is laws that they seek. To borrow Armstrong’s language, we must be concerned not just with “all rational mankind”, but also with all “other rational creatures”.

Looking for agreement in scientific practice will not, in general, be enough to identify a universal limit on relativity, for agreement now does not guarantee the same agreement will persist. Loewer suggests a strategy for dealing with this by setting the details of the BSA to what they would be in a “final theory” of physics that succeeds the present theory in a process of “developments that are considered within the scientific community to increase the simplicity, coherence, informativeness, explanatoriness, and other scientific virtues” that is carried out until no such increases may be achieved [Loewer, 2007, p. 325]. Setting aside any question of whether this strategy makes the laws subjective, we can say that it does not guarantee universal limits to the BSA. From a given starting theory, it may be that there are choices in the formulation of successive theories with different responses leading to different final theories. And one must consider starting theories other than our actual current theory to avoid subjectivity; then universality may be undermined by a lack of convergence.

What Loewer gets right is the idea that identifying universal limits on the relativity of the BSA will be a matter of looking at not just what is true of scientific practice now, but also at what will to be true in the future. Such limits may be grouped under three broad categories: (1) limits that emerge from the internal structure of the BSA, (2) limits imposed by desiderata for an analysis of laws, and (3) limits imposed by insurmountable practical challenges. I now consider each of these categories in turn, and give brief examples of how they may limit relativity.

Limits that emerge from the internal structure of the BSA were actually considered already (though not as such) in the previous section. The kind relativity of

the BBSA seemed to do the same work as some amount of fact relativity. Suppose, for example, that a field does not care about electrons. This could be reflected in how the details of the BSA are set by not counting electrons among the kinds the field treats as basic, or by ignoring any specifically electron related matters of fact. Similarly with the example of ecologists not considering field theoretic systems, which could be ensured by system relativity or by kind relativity when “fields” are not counted among the basic kinds of ecology. In either example one could, of course, adopt both sorts of relativity. But in some cases two sorts of relativity will be redundant because of how they interact in the BSA. If every time one sort of relativity is needed it is redundant with some other already admitted sort, then the first sort of relativity may be eliminated.

Consider now the second category: limits imposed by desiderata for an analysis of laws. Some desiderata for an analysis of laws will not directly impose limits on the relativity of the BSA. To use an example running throughout this paper, take the non-subjectivity of the laws as a desideratum. That constrains how the free variables of the BSA are set in the sense that they cannot be set directly by subjects, but limiting relativity involves saying what variables are not free, or what values of free variables are unacceptable.

A long standing desideratum for an analysis of laws is that it be able to distinguish between lawful and accidental regularities. The BSA may fail to do this if gerrymandered predicates are allowed, as in the trivial systems problem of [Lewis, 1983] where the predicate F that is true of everything in just those worlds where the system $\forall xFx$ is true. The kind that corresponds to F is included in the set

of supervenient kinds over whose subsets the kind relativity of the BBSA ranges. In that way there are sets of laws in the BBSA (any that are relative to a set of kinds that make F a predicate) that make every regularity a law, and thus fail to distinguish between lawful and accidental regularities.

The sets of laws that fail to distinguish between lawful and accidental regularities are supposed to be unproblematic for the BBSA precisely because they include the problematic predicate. The desideratum indicates that we are only interested in sets of laws that distinguish between lawful and accidental regularities. If we are uninterested in any kinds that yield a predicate like F , then the sets of laws that fail to distinguish will not be among the sets of interest to us. And if we are interested in such kinds, then the desideratum is not as strong as one might imagine precisely because we are committed to an interest in laws that violate it.

Another response to the presence of these non-distinguishing law sets is available to a BSA/BBSA proponent. BBSA laws are determined relative to any subset of the set of all kinds supervening on the true fundamental kinds. Among the set of supervenient kinds are kinds that guarantee a violation of the distinguishing desideratum. We can satisfy the desideratum (or at least take steps necessary for satisfying it) by limiting relativity: Do not let just any subset of the supervenient kinds yield a set of laws, rather only let there be laws relative to any subset of the supervenient kinds *minus* any desideratum violating kinds. This is, perhaps, only a small reduction in the relativity of the BBSA, but it is a reduction that is required by respecting scientific practice by way of the distinguishing desideratum.

Lastly, consider the third category: limits imposed by insurmountable prac-

tical challenges. For a first example, recall that a best system competition is a function from facts, systems, and languages, to the best system-language pair(s). The possibility of identifying the laws depends on the possibility of practically implementing the competition function. A best system competition may simply fail to identify a best system if there is no effective means of calculating its output. To ensure that that does not happen, we may limit competition relativity by allowing only computable functions to serve as the best system competition function.

Another such limit may come from a need to write things down. Every time we do, whether it is part of figuring out what the laws are, or using the laws to achieve some end, we incur physical space and energy costs. While in practice those costs are determined by our current technology, in principle their minima will be subject to physical limits characterized by the laws themselves. In the interest of minimizing those costs, and any other practical limitations whose details depend on the laws themselves, we may further reduce relativity by the imposition of a kind of consistency requirement that must be included with every competition function: A system S is a best system only if its laws impose the kinds of constraints on the BSA that would yield S as a best system.

These particular ideas for how relativity in the BSA may be limited are an eclectic lot. All of them, and some more than others, warrant or require more development. And, to varying degrees, they will receive it later in this dissertation. Specifically: In Chapter 3, the relationship between kind and fact relativity is explored. In Chapter 4, the distinguishing desideratum plays a prominent role in the discussion of the trivial systems problem. And, in Chapter 5, the matters of

computability and writing things down become relevant to the use of information theoretic principles in the best system competition.

2.6 Summary

My purpose in the this chapter was two-fold: First, to provide an analysis of Armstrong's objection to the BSA. Second, to sketch out the routes proponents of the BSA may follow towards a solution to the objection. Armstrong's objection had two possible readings: The first, which objects to laws being subjective, was found wanting because the laws of the BSA are not subject dependent in general. The second reading, which objects to relativity in the laws, is only able to require of the BSA that its relativity match exactly what is required by scientific practice. There is a spectrum of variously relative BSA variants, from a single maximally relative variant that lets sets of laws be identified relative to any way of filling in the details of the BSA, all the way down to a host of minimally relative BSAs that each identify only a single set of laws. First pass considerations of how to motivate relativity suggest that the maximally relative BSA is a viable view, but the range of available strategies for identifying limits to relativity in the BSA tells against that extreme response to Armstrong's objection.

If one's goal is to find a minimally relative variant of the BSA, then a lot more work has to be done to build on the suggestions given here for limiting relativity. If one's goal is to defend the maximally relative BSA, then something must be said about why the suggested means of limiting relativity, or any means like them, are

not required to respect scientific practice. But answering Armstrong's objection requires neither a deferent minimally relative variant of the BSA, or the defiant maximally relative variant of the BSA. Something in between should be just right, with relativity, and limits to it, properly motivated by appeals to scientific practice.

Chapter 3: Relativizing to Kinds and Facts

A major part of what motivates the BSA is the “Humean Supervenience” thesis, according to which “the whole truth about a world like ours”, such as what the laws are, “supervenes on the spatiotemporal distribution of local qualities” [Lewis, 1994, p. 473]. That supervenient base for “the whole truth” of a world is called the Humean Mosaic (HM), and the BSA may be summarized as saying that the laws of nature are the axioms and theorems of the best systematization—the simplest and strongest, on balance—of the (HM. Importantly, the systems that compete to be the best in the BSA are not systematizations of the HM directly. Rather, they are systematizations of the facts of the HM as expressed using a privileged language.¹ Recent variations on the BSA have paid special attention to the role of the privileged language in identifying laws (e.g.: [Schrenk, 2008, Loewer, 2007, Cohen and Callender, 2009]). The broad aim of this chapter is to (re-)direct some attention toward the facts that are being systematized and their role in identifying the laws. The more specific aim is to argue that the role for language(s) championed by [Cohen

¹ Of course it cannot be that *all facts about* the HM are being systematized since it is presumably a fact about the HM that particular regularities are laws, and we should not expect every competing system to systematize those facts (on pain of trivializing the whole analysis). What is being systematized are the *basic facts of* the HM, as determined by the predicates treated as basic by the privileged language. Throughout this dissertation I have spoken and will speak of “facts” in the sense of these basic facts.

and Callender, 2009] in their “Better Best Systems Analysis” (BBSA) view should be extended to facts.²

According to the BBSA, laws are *kind relative* in the sense that different sets of laws exist relative to every set of kinds that may be treated as basic in the language used to express the laws (and competing systems and systematized facts). The extension I am proposing, dubbed simply the “Kind and Fact Relative Analysis” view (KFRA), has laws determined relative to the set of kinds treated as basic *and* the set of facts to be systematized. The KFRA, by including the kind relativity of the BBSA, can do all the same good work as the BBSA. Furthermore, I argue that the fact relativity of the KFRA does a better job of answering two new desiderata. The first desideratum calls for the regularities that appear in *interfield interactions* to be accommodated as laws. This depends, at least in some cases, on there being a shared interest in some kind(s) but not all the facts related to the shared kind(s). The second desideratum calls for an *egalitarian distinction* between fundamental and special science laws. Such a distinction provides metaphysically equal standing to fundamental and special science laws, while also allowing for the distinction between them that exists in scientific practice.

The outline of this chapter is as follows. In Section 3.1, I introduce the BBSA and its accommodation of special science laws. In Section 3.2, I introduce the kind and fact relative KFRA, and discuss its ability to accommodate special science laws in the same way as the BBSA. Section 3.3 is concerned with developing the

² Variations of the BBSA was developed independently in [Schrenk, 2008] and [Cohen and Callender, 2009]. I generally focus on the version of the view developed by the later authors because of their explicit efforts to reconcile the BBSA with those of non-Humeans.

interfield interactions desideratum, and showing how it may be met by the KFRA but not the BBSA. Section 3.4 is concerned with developing the desideratum of having an egalitarian distinction between fundamental and special science laws, and arguing that the KFRA captures features of the distinction that the BBSA cannot. I conclude in Section 3.5 that anyone who endorses the BBSA would do better by endorsing the KFRA, and there offer a summary of the chapter.

3.1 The Better Best Systems Analysis

It is helpful to start with a closer look at Lewis' thoughts on facts and languages in the BSA. Lewis took it to be that the privileged language of the laws was the true fundamental language of "perfectly natural properties" [Lewis, 1983, p. 368]. Assuming classical physics to be more or less in the right when describing the fundamental nature of the world (i.e. the HM), the facts of the HM as expressed in the fundamental language are concerned with

spatiotemporal relations: distance relations, both spacelike and timelike, and perhaps also occupancy relations between point-sized things and spacetime points. And [...] local qualities: perfectly natural intrinsic properties of points, or of point-sized occupants of points.

[Lewis, 1994, p. 474]

The Humean Supervenience thesis and Lewis' particular characterization of the HM have largely fallen out of favor, but that has not done much harm to the general

standing of the BSA.³ Authors since Lewis have offered a number of BSA-style views of laws that depart from Lewis' version of the view in a variety of ways.

One such departure from Lewis' BSA is the BBSA of [Cohen and Callender, 2009]. According to the BBSA, a best system competition is run for every language—as determined by the set of kinds⁴ treated as basic in the language—that may be chosen to express the facts and systematizations of the world.⁵ In each competition, every competing system is judged according to its expression in the privileged language. If S is the best system when the privileged language treats the set of kinds K as basic, then the theorems of S are the laws *relative to* K . This kind relativity in the laws of the BBSA helps resolve several issues related to language choice in the BSA.⁶ It also allows for the existence of laws in the special sciences that are *egalitarian* primarily in the sense (to be expanded upon below) that no one set of laws is dependent upon any other set (e.g. the laws of biology don't depend on the laws of physics for their existence). These egalitarian special science laws are part of

³ Of particular relevance for this chapter, [Cohen and Callender, 2009] restrict their interest in supervenience to kinds. And, while they do use the term “Humean mosaic”, it is not the capital-H, capital-M, *Humean Mosaic* of Lewis, but rather the term for the various sets of language-relative facts of the world—e.g., in comparing physics and ecology, they speak of “two quite different mosaics” [Callender and Cohen, 2010, p. 18]. More generally, the BSA is compatible with a rejection of Humean Supervenience, and such a rejection has been a boon to the BSA since it frees the view from a host of extraneous metaphysical commitments. See [Loewer, 1996] for a discussion of the problems with Humean Supervenience and an explanation of how the BSA can survive despite them.

⁴ Where the possible sets of kinds are subsets of the set of the fundamental kinds *and* all kinds that supervene on the fundamental kinds.

⁵ Like with Humean Supervenience, the HM comes with substantial metaphysical baggage. So speaking of “facts of the HM” is dropped in favor of the more general “facts of the world”. It is still the *language determined basic* facts, but what those might be does not necessarily come with the commitments of the HM.

⁶ [Cohen and Callender, 2009] are explicitly motivated by three desiderata: *epistemic accessibility*, *immanent comparisons*, and *supervenient kind laws*. Suffice it to say that I think [Cohen and Callender, 2009] are right that the BBSA satisfies these three desiderata. Additionally, it should not be hard to see if one reviews the relevant arguments in that paper that they may all be run with little to no adjustment on behalf of the view introduced in this chapter.

an effort to reconcile “Non-Humean Pluralist” views like those of [Cartwright, 1999] or [Dupré, 1993] with a Best Systems view of laws [Callender and Cohen, 2010]. Advancing that effort, or at least preserving what has already been done, will be a largely backgrounded but recurring theme in this chapter.

The idea at the heart of the BBSA’s egalitarian accommodation of special science laws is this: If the law-directed interests of a field⁷ are captured by what kinds are treated as basic by the field, then the laws of the field are the laws that are relative to the kinds treated as basic by the field. For every set of kinds K_i , there are corresponding K_i -relative laws R_i . If biologists are interested in the kinds K_{bio} , then the K_{bio} -relative laws R_{bio} are the laws of biology. Similarly, if ecologists are interested in kinds K_{eco} , the laws R_{eco} are the laws of ecology. And so on, for every scientific field. There will also be lots of sets of kinds that are not, and may never be, of any interest to anyone. There are still laws for such sets of kinds, but they are just as (un)interesting as the kinds to which they are relative.

That this provides *egalitarian* special science laws stems from two related ideas in [Callender and Cohen, 2010]. The first is that the laws of special sciences should have “autonomy” from the fundamental laws (or the laws of any other field). Bridge principles (and the like) may exist incidentally because two fields are closely related;

⁷ By “law-directed interests” I mean just those interests that correspond to the feature(s) of the BSA to which the laws are relative. There are a lot of things that could be counted in the interests of a field, but not all of them have to do with identifying the laws of the field. For example, astronomers are interested in stars and planets and moons, and insofar as those are kinds of interest, the BBSA would say they are among the law-directed interests of astronomy. But astronomers might also be interested in moving from a funding model dependent on a few large government funded international collaborations to a funding model with more small projects supported by private industry. Strictly speaking, the funding stuff might be a part of the interests of astronomy, but it’s not a part of the *law-directed* interests of astronomy (at least not according to the BBSA and KFRA).

e.g., because of direct overlap in their basic kinds, or because of easy translation between one set of kinds and the other. But that the laws of one set deserve the title of “law” cannot depend on the existence of any bridge principles connecting the one set to any others [Callender and Cohen, 2010, p. 439]. The other idea relevant to the egalitarian nature of the special science laws in the BBSA is an “ontological egalitarianism”, according to which no ontological level enjoys any special status. The fundamental kinds are not, for example, more deserving of interest than any other set of kinds from an ontological perspective. Of course, we humans may care more about some kinds than others for any number of reasons—the kinds of biology and economics are of great practical importance to us, and curiosity may drive us to search for the fundamental kinds—but a truly disinterested observer should find no reason to privilege one set over any other. [Callender and Cohen, 2010] relate this to the “dappled world” picture of science in [Cartwright, 1999] that encourages a preponderance of equal and more or less independent scientific enterprises, and discourages stratification and any strained unification. However, while Cartwright is explicitly anti-fundamentalist, Callender and Cohen are not; a thoroughgoing anti-fundamentalism must be dropped on account of the fundamental kinds playing a role in determining the set of all kinds for the BBSA.

The BBSA’s egalitarian special science laws are a great addition to the BSA, but the precise implementation leaves something to be desired. The success of the BBSA depends on a conditional: *If* the law-directed interests of a field are captured by what kinds are treated as basic by the field, *then* the laws of the field are the laws that are relative to the kinds treated as basic by the field. Going forward, I

will challenge the antecedent. The law-directed interests of a field are not captured entirely by the kinds treated as basic by the field. Systems competing to be the best are systematizations of *the facts of the world* as expressed using a particular language. Once a language is chosen—as in the BBSA by a choice of kinds—there is still room for a field to be interested in only a subset of the facts of the world. What is important to the individuation of fields and laws are the kinds treated as basic *and* the set of matters of fact that are being systematized.

3.2 Relativity to Kinds and Facts

What kinds are treated as basic—by the members of a scientific field, or in the language in which a set of laws is expressed—is a matter intimately related to what facts are relevant—as being of interest to a field, or as part of what is actually being systematized by the laws. Increasing the set of available kinds increases the set of expressible facts, and limiting the set of available kinds limits the set of expressible facts. To illustrate with a very toy-like example, suppose for a moment that the fundamental level of the world is just the arrangement of atoms. The fundamental kinds include the atomic elements: hydrogen, helium, lithium, beryllium, etc. The fundamental matters of fact are concerned with the spatio-temporal arrangements of the atomic elements: something like ‘there is helium at time t and position p ’. There can also be non-fundamental, supervenient, kinds: caffeine is a certain arrangement of carbon, hydrogen, nitrogen, and oxygen. With the non-fundamental kinds comes the possibility of expressing non-fundamental facts that describe the same state of

affairs as the fundamental facts, but in different terms: ‘there is caffeine at time t occupying position(s) $P = \{p_1, p_2, \dots\}$ ’, where the p_i in P are the positions of the constituent atoms of the caffeine molecule.

In the direction of constraining the set of facts, we can see that a suitably chosen set of kinds can make some facts, even under redescription, wholly inaccessible. Consider a set of kinds that includes caffeine, but not carbon. With such a set of kinds there are no strictly carbon related matters of fact. But some carbon related matters of fact are recoverable by way of the caffeine facts and additional principles of translation having to do with the relation between caffeine and carbon. However, not all carbon related matters of fact may be recoverable; e.g., if a set of kinds doesn’t include carbon and it doesn’t include graphite (and it doesn’t include pencils or any other kinds involving graphite), then those carbon related matters of fact involved in graphite related matters of fact are both inexpressible and unrecoverable.

The fact relativity provided by kind relativity is incomplete. There are all the carbon related matters of fact, and as subsets of those there are the caffeine and graphite related carbon facts. The kind carbon may be of interest to members of the Department of Coffee Chemistry and the Department of Pencil Chemistry. But (presumably) coffee chemists only care about caffeine related carbon facts, and pencil chemists only care about graphite related carbon facts. However, once committed to having carbon as a kind, nothing short of fact relativity will allow the coffee chemists (or pencil chemists) to ignore the carbon facts related to graphite (or caffeine). This might be overcome somewhat by employing gerrymandered supervenient kinds—in

our example they might be ‘coffee carbon’ and ‘pencil carbon’—but then, strictly speaking, the respective fields are not both interested in the same kind—namely, *carbon*—and that will be a problem if the two fields try to interact over their shared interest.⁸

The BBSA must be extended with fact relativity if it is to accommodate the possibility that two fields treat some of the same kinds as basic while differing in what facts related to those shared kinds are of interest to them. That is a big if, but it will become central to the issue of accommodating interfield interactions in Section 3.3. And, as will be seen in Section 3.4’s discussion of distinguishing between fundamental and special science laws, this is not the only reason to admit fact relativity. With all that being said, let us extend the BBSA into the *Kind and Fact Relative Analysis* (KFRA).

In the BBSA, there is the set of all kinds \mathcal{K} and its subsets K_i . A best system competition is run for every K_i , and all competing systems in the K_i competition are expressed with the kinds of K_i as the basic kinds. A regularity is a law (according to the BBSA) relative to K_i if it appears in the best system of the K_i competition. The KFRA adds to this the set of all facts \mathcal{F} (that are expressible when the available kinds are the whole of \mathcal{K}) and its subsets F_j . A best system competition is run for every (K_i, F_j) , where F_j provides the set of facts to be systematized and K_i is the set of kinds treated as basic in the language in which all facts and competing systems are expressed. For every such competition the best system provides the

⁸ This issue of gerrymandered kinds will come get a fuller treatment in Section 3.3.3, where the concern is specifically with interfield interactions.

(K_i, F_j) -relative laws $R_{i,j}$.

The KFRA allows for egalitarian special science laws in a manner structurally the same as in the BBSA. If the interests of a field are captured by what kinds are treated as basic *and what facts are intended to be systematized* by the field, then the laws of the field are the laws that are relative to the kinds treated as basic *and the facts intended to be systematized* by the field. If biologists are interested in the kinds K_{bio} and systematizing the facts F_{bio} , then the $(K_{\text{bio}}, F_{\text{bio}})$ -relative laws $R_{\text{bio},\text{bio}}$ are the laws of biology. If ecologists are interested in kinds K_{eco} and systematizing the facts F_{eco} , the laws $R_{\text{eco},\text{eco}}$ are the laws of ecology. And so on.

The special science laws of the BBSA are all a special case of the KFRA where the set of facts to be systematized is always the set of all facts \mathcal{F} . Insofar as we were happy with the treatment of special science laws by the BBSA, we may assume by default that the fact set of interest is \mathcal{F} . Thus, laws of biology are the $(K_{\text{bio}}, \mathcal{F})$ -relative laws $R_{\text{bio},\mathcal{F}}$, the laws of ecology are the $(K_{\text{eco}}, \mathcal{F})$ -relative laws $R_{\text{eco},\mathcal{F}}$, etc. What the KFRA allows for beyond the BBSA is the possibility that the law-directed interests of biology, ecology, or any other field, involve being interested in only some proper subset of all the matters of fact.

3.3 Interfield Interactions

Interfield interactions, as a group and individually, are an important test case for the BBSA (and, by extension, the KFRA). The egalitarianism of Callender and Cohen should require that laws resulting from interfield interactions be located

somewhere with the same status as the laws of any one field.

Darden and Maull nicely describe interfield interactions as occurring when fields share an interest in explaining different aspects of the same phenomenon, and when questions arise about that phenomenon within a field which cannot be answered with the techniques and concepts of that field.

[Darden and Maull, 1977, p. 50]

The existence of such interfield interactions blurs the boundaries between fields. With blurred boundaries come unclear interests, and with that comes unclear laws. Something should be said to clear things up.

Callender and Cohen write:

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.

[Callender and Cohen, 2010, p. 444]

These questions are assumed to have no answer in the ecological laws because the kinds relevant to their explanantia are not ecological kinds, but rather physical and physiological kinds. But then we should also not assume that these questions have answers in physics and physiology, because the kinds relevant to their explananda (e.g., rabbits) are not physical or physiological kinds. However, these questions

must have answers *somewhere*. Some field must be interested in the relevant kinds of physics, physiology, and ecology. The field may be one of those just named, or it may be an entirely new field that arises exclusively for capturing the particular interfield interaction. Either way, interfield interactions (and their resultant laws) should be accommodated by views like the BBSA and KFRA on account of their egalitarianism. I will refer to this call for accommodation as the *interfield interaction desideratum*.

Going forward in this section, I begin by offering a concrete case of interfield interaction from [Barlow, 1952] that involves blurring the boundaries between biology and physics on account of how the behavior of photons influences the structure of insect eyes. Following that, I argue against the ability of the BBSA to make sense of Barlow's work. My argument, in brief, will be this: The best system of physics, especially the relevant fragment of it concerned exclusively with just some of the behavior of photons, is very strong and simple compared to the best system of biology. The consequence of this is that when the BBSA tries to locate a field for Barlow's work that has interests in both physics and biology, systems containing the physical laws will tend to outcompete systems containing the biological laws, and the "best" system will fail to include all the laws relevant to the interfield interaction. The KFRA, considered following the BBSA, need not suffer from such an outcompeting problem. Biology, if that is where Barlow's work is to be located, surely is interested in photons as a kind, but it is interested in only a small subset of photon related facts. The strength that comes with the relevant physical laws can be greatly reduced by disregarding, by way of the KFRA's fact relativity, the physical

facts that are irrelevant to Barlow’s work. In this way the KFRA can succeed at making sense of interfield interactions when the BBSA fails.

3.3.1 Barlow, Between Biology and Physics

[[Barlow, 1952](#)], “The Size of Ommatidia in Apposition Eyes”, published in *The Journal of Experimental Biology*, begins not so much with biology as physics: Think of an insect’s compound eye as a simple sphere of lenses each in front of a single photoreceptor.⁹ The more lenses (and associated photoreceptors) there are, the higher the resolution of the image that may be generated. The smaller the lenses are, the blurrier the image becomes as a result of diffraction and as a function of the wavelength of the incoming light.

For a given sized eye sphere, and photoreceptors sensitive to a particular wavelength of light, there is a physical limit on the resolution of the image available to an insect. This is because more lenses increases the resolution, but more lenses also results in smaller lenses (since the size of the sphere is fixed) and smaller lenses decreases resolution. Barlow calculated this limit, proceeded to the Cambridge Museum of Zoology, measured the eye sphere and lens sizes of 27 species of insect in the collection, and found that those sizes were approximately those that would be predicted if we supposed that the resolution of insect eyes was at the physical limit.

Barlow’s discussion of actual compound eyes concludes that

⁹ An ommatidium is a segment of a compound eye including the lens, photoreceptor(s), and various support structures. Actual ommatidia typically have more than one photoreceptor, but the number is at least on the order of one, as opposed to vertebrate eyes like our own that can have millions behind a single lens.

the way in which ommatidial size and inter-ommatidial angle are adjusted in eyes of different sizes suggests strongly that the wave structure of light is the limiting factor in the design of the compound eye.

[Barlow, 1952, p. 672]

Underwriting this conclusion is what I will call *Barlow's regularity*:

in [compound apposition] eyes of different sizes the number of ommatidia is adjusted so that inter-ommatidial angle is just below the limiting resolving power of the ommatidia

[Barlow, 1952, p. 674].

Because it contributes to an explanation of the conclusion of the paper, Barlow's regularity deserves the status of a law in the BBSA and KFRA.

Answering the interfield interaction desideratum will require us to say in which set of laws Barlow's regularity should be located. Barlow's regularity might be a law of physics or a law biology (to focus on the two most salient extant fields). There is also an important third possibility to consider. Regularities sourced to interfield interactions may belong to their own field and set of laws, distinct from, but of course related to, the interacting fields. So, not only could we locate Barlow's regularity among the laws of physics or biology, we could also locate it among the laws of a new field constructed, so to speak, just for Barlow's work.

Also, it is not just Barlow's regularity that has to be located in a set of laws. Barlow's concluding remark goes beyond the regularity itself and relates "the wave structure of light" to "the design of the compound eye". To do justice to Barlow's

work will require the locating of a set of laws that contains all of (1) Barlow’s regularity, (2) the physical laws relevant to “the wave structure of light”, and (3) the biological laws relevant to “the design of the compound eye”.¹⁰ So, if we try to locate Barlow’s regularity in the laws of physics, we must find some way of introducing the biological laws of (3) into the laws of physics. If we try to locate Barlow’s regularity in the laws of biology, we must introduce the physical laws of (2) into the laws of biology. And, if we try to locate Barlow’s regularity in some alternative set of laws, we must ensure that the physical laws of (2) and the biological laws of (3) can coexist in the alternative set.

Such is the sketch of what the the BBSA and KFRA must do to satisfy the desideratum of interfield interactions. With that done, we can now discuss in detail how the BBSA runs into trouble with Barlow’s work and how the KFRA succeeds.

3.3.2 Interfield Interactions and the BBSA

What varies from one set of laws to the next in the BBSA is the set of kinds that are treated as basic. The laws R_{phys} are the laws *of physics* because they are the laws of the system that is best when all competing systems are expressed using the language with basic kinds K_{phys} , and the kinds K_{phys} are the kinds treated as

¹⁰ As I will use the terms, being a physical (or biological) law (or kind, or fact, or what have you) is to be a law (kind, fact, etc.) that is *typically* a law of physics (or of biology). (I do not want to be saying anything about what it means to be “physical” in any broader sense.) A law (kind, fact, etc.) is a law (...) “of physics” (or “of biology”) just in case it is actually among the laws of the best system determined relative to kinds (and facts) of interest to physics (or biology) So, if Barlow’s regularity is located in physics, then it will be a law of physics, but it is not a physical law, since it is not a law that we typically think of as appearing among the laws of physics. If the relevant biological laws are also located in the laws of physics, then those laws may be called biological laws and laws of physics, since they are typically found in the laws of biology, but here they are found in the laws of physics.

basic by the field of physics. The laws R_{bio} and kinds K_{bio} of biology are similarly related. Locating Barlow’s regularity for the BBSA will be a matter of finding one set of kinds with the relevant physical kinds and the relevant biological kinds. And that set of kinds must yield a set of laws that includes Barlow’s regularity, the relevant physical laws, and the relevant biological laws.

There are three sets of kinds that are likely to do the work we need. The kinds of physics might be expanded to include the needed biological kinds, like ‘ommatidia’. The kinds of biology might be expanded to include the needed physical kinds, like ‘photons’. And there could be a third set of kinds that includes what is needed, and nothing more, from both the biological and physical kinds.

Consider expanding the kinds of biology first. There are the biological kinds K_{bio} . And there is the system S_{bio} that has as its axioms and theorems exactly the biological laws. According to the BBSA, it should be that S_{bio} is the best system relative to K_{bio} .¹¹ Assume that the only needed physical kind is ‘photon’, or k_{p} . Now consider the candidate systems when the kinds are $K_{\text{bio}} \cup \{k_{\text{p}}\}$. We do not want S_{bio} to persist as the best system, since then the laws of biology would not include Barlow’s regularity or the needed physical laws. And we should not expect S_{bio} to persist as the best system.

I take it to be a commonsensical feature of strength that (roughly), other things being equal between systems S_1 and S_2 , if S_2 is silent on more facts than S_1 ,

¹¹ Note that the subscript to a system refers to what is being successfully systematized, and need not match, in general, the subscript on the relevant kinds. In the competition to determine the best system relative to the kinds K_{bio} , S_{bio} is a *candidate* system (as well as the best). Another candidate system might be named $S_{\text{bio}+\text{X}}$, which is a system that has all the biological laws, plus the laws of X. If S_{bio} was not defined in a way that guarantees its victory, then it would be reasonable to consider the possibility of $S_{\text{bio}+\text{X}}$ being the best system relative to K_{bio} ,

S_1 is stronger than S_2 . In the typical biology situation with just the basic kinds K_{bio} , every system was forced to be silent on the photon related matters of fact, and so any weakness resulting from that silence was irrelevant. But now, with k_p included in the set of basic kinds, not all systems must be silent on the photon facts. A newly available system, $S_{\text{bio+p}}$, which covers all the facts that S_{bio} does *plus* the photon facts, will be much stronger than S_{bio} (since there are a lot of photon facts). $S_{\text{bio+p}}$ should only be slightly less simple than S_{bio} , since, relative to the overall complexity of biology, the complexity required to cover the photon facts should be small.

Regarding this comparison of simplicity: I take it to be a commonsensical feature of simplicity that (roughly), other things being equal between systems S_1 and S_2 , if S_1 requires employing fewer basic kinds than S_2 , then S_1 is simpler than S_2 . $S_{\text{bio+p}}$ requires one basic kind more than S_{bio} , and that would be, considering the many kinds presumably in K_{bio} , a pretty minimal loss of simplicity when all else is equal. But surely not all other things are equal between S_{bio} and $S_{\text{bio+p}}$, so that feature of simplicity only gets us so far. What I really have in mind when suggesting that “relative to the overall complexity of biology, the complexity required to cover the photon facts should be small” is something like the contrast between Maxwell’s equations and the enormous map being developed, as in [Thiele et al., 2013], just to cover the small slice of biology that is human metabolism.¹²

¹² Someone might object on the grounds that the simplicity of Maxwell’s equations is an illusion because, as Cartwright puts it, “a multitude of highly complicated unknown laws must be assumed to save Maxwell” when applying the theory [Cartwright, 1999, p. 155]. I am not worried about this is because, roughly speaking, Cartwright’s objection depends on trying to apply Maxwell’s equations in many circumstances. If someone thinks that Maxwell’s equations may be applied to any non-physics kinds (like the radiometers in Cartwright’s example), then I am inclined to join

Now, we might assume that $S_{\text{bio+p}}$ is the system we need, in that it yields as laws the needed biological laws, the needed physical laws, and Barlow’s regularity. But, even if that is true, the difficulty is we should not expect $S_{\text{bio+p}}$ to be the best system. Consider another system, S_{p} , that covers only the photon facts, and none of the biological kind related matters of fact. Certainly this system will be very simple compared to $S_{\text{bio+p}}$. And S_{p} should not be that much weaker than $S_{\text{bio+p}}$, for, even though S_{p} doesn’t cover any of the biological facts, there are a lot more photon facts than there are biological facts. Both $S_{\text{bio+p}}$ and S_{p} are a lot stronger than S_{bio} , but while $S_{\text{bio+p}}$ is a bit less simple, S_{p} is a lot simpler. On balance, it seems likely that S_{p} will be the better system relative to the set of kinds $K_{\text{bio}} \cup \{k_{\text{p}}\}$.

One might object, or at least worry, about my claim that “there are a lot more photon facts than there are biological facts”. I assuage the worried first: The stars in the observable universe release on the order 10^{67} photons per second.¹³ In contrast, there are about 10^{37} basepairs of DNA on Earth [Landenmark et al., 2015]. For the number of biological facts to get close to the number of facts related to photons originating from stars, there would need to be (assuming the DNA base-pair count is the high end for numbers of facts concerned with any particular biological kind)

Cartwright and think that the relevant (best) system will have to be a lot more complicated than the inclusion of Maxwell’s equations would appear to make it at first glance. However, in the BBSA (and KFRA), relativizing means that the complex baggage that comes with Maxwell’s equations is limited by the kinds (and facts) to which a best system competition is relative. Barlow doesn’t employ “a multitude of highly complicated unknown laws” when applying what physics he does to biology, which suggests that such a multitude is not required in this case.

¹³ The Sun has an energy output of approximately 4×10^{26} J/s, and photons in the visible spectrum (which is most of what is released by the sun) have an energy of approximately 4×10^{-19} J. So, roughly, the Sun releases 10^{45} photons per second. The number of stars estimated to be in the observable universe is on the order of 10^{22} . Assuming the Sun is an average star, that makes for $10^{45} \times 10^{22} = 10^{67}$ photons per second from all the stars in the observable universe.

10^{30} biological kinds each with a number of associated facts comparable to the count of base pairs of DNA on Earth. For example, base pairs of RNA might fit the bill, but that is just one kind of the one thousand billion billion billion that are needed. While there are surely a lot biological facts, there are *a lot more* photon facts.

Now one might object, and justifiably so, that this is only one way of individuating the relevant facts, and other ways—especially if they don't involve counting individual photons—will not result in there being vastly more photon facts than biological facts. It's absolutely right that there are other ways of individuating the relevant facts, and surely among those are ones that attribute more facts to biology than photons. But I do not think those ways of individuating will comport with scientific practice. Our understanding of photon diffraction through a slit (which is central to the physics relevant Barlow's work) involves measuring the trajectory deviations of individual photons. And, more on the side of the relevant biology, eyes are actually sensitive to individual photons [Rieke and Baylor, 1998]. There may be ways of individuating the facts that are problematic for my argument, but the burden is on the objector to identify one *and* argue that it better matches how the facts are individuated in practice.

To summarize: The introduction of the needed physical kinds into the set of biological kinds leads, because of their preponderance and relative simplicity, to the physical kinds and laws outcompeting the biological kinds and laws. The result will be a set of laws for biology that do not at all resemble the biological laws. Of particular significance, the needed biological laws will not appear in the laws of biology. Nor should we expect Barlow's regularity to appear in the laws of

biology. Because of this outcompeting, the BBSA cannot accommodate the interfield interaction exemplified by Barlow's work by introducing the relevant physical kinds into the kinds of biology.

The same problems arise for the BBSA if it tries to accommodate Barlow's work by introducing biological kinds into physics or by collecting all the relevant kinds into a new field. In the latter case the problem is exactly the same as when introducing physical kinds into biology, but with just the needed biological kinds being weighed against the needed physical kinds. With even fewer biological kinds being considered, the physical kinds (and related facts and laws) are even more likely to outcompete their biological counterparts. In the former case, the biological kinds will never get a foothold in the best system relative to the now-expanded set of physics kinds, since the gains in strength will be small compared to the losses in simplicity.

Ommatidia are rare and complicated, and while that shouldn't undermine our interest in them, it does undermine their relevance when our interests also include every photon in the world.

3.3.3 Interfield Interactions and the KFRA

Let us consider again how we might include Barlow's work in the laws of biology, but this time with the KFRA as our view of laws.

Photons as a kind are undoubtedly of interest to biologists; beyond their role in Barlow's work, they are relevant to the study of vision more generally, photosyn-

thesis, bioluminescence, camouflage, color-ation, and more. With the all-or-nothing fact relativity provided by kind relativity, every photon related fact is the concern of biology, whether it is related to anything remotely biological or not. But not all photon facts are of interest to biologists. The photon facts of interest to biologists are the ones that are near to biological facts. Biologists will care to systematize the facts related to photons that interact with eyes, leaves, skin, and scales. But the vast majority of photon related facts—whose relevant photons never interact with anything biological, or which might only be of interest to biology for their role in higher level processes like heat exchange—are ones biologists will want to ignore.

The KFRA, because it allows for there to be laws relativized to any subset of the facts of the world, can exploit biology’s particular interests in facts. In the typical biology situation, the laws of biology $R_{\text{bio},\text{bio}}$ are those taken from the best system relative to the biological facts F_{bio} and biological kinds K_{bio} . As was tried with the BBSA, the KFRA can add the relevant physics kinds—i.e., photons, k_{p} —to the set of kinds of biology. Doing that, the laws of biology are the laws relative to $(K_{\text{bio}} \cup \{k_{\text{p}}\}, F_{\text{bio}})$. If we assume, as the BBSA does, that F_{bio} just is the set of all facts, then the same outcompeting problem will arise for the KFRA as arises for the BBSA.

But we need not assume that F_{bio} is the set of all facts. F_{bio} is supposed to correspond to just those facts that are of interest to biologists. So exclude from the set of facts F_{bio} any physical facts that are irrelevant to biology, while including all those that are relevant. Since all the relevant physical facts (i.e. the photon facts) that appear among the facts of biology will be associated with some biological fact(s),

the strength gained by a system for covering the included photon facts should not be enormous relative to the strength gains associated with covering the biological facts. The system S_p , which covers all the photon facts and none of the biological facts, was problematic for the BBSA because it was overwhelmingly strong and simple. On the KFRA, since the set of facts being systematized includes only the biologically relevant photon facts, S_p 's covering "all" the photons facts just means that it covers all the biologically relevant photon facts, which are, presumably, not large in number compared to all the biological facts. The result is that S_p does not benefit from overwhelming strength, and so it is unlikely to be the best system of biology. Without S_p to outcompete it, the way is clear for $S_{\text{bio+p}}$, the system assumed to satisfy all the requirements of accommodating Barlow's work, to be the best system of biology according to the KFRA.

If anything, the risk of outcompeting runs the other way for the KFRA compared to the BBSA. Relative to the biological facts, there might not be enough included photon facts to make the gains in strength worth the losses in simplicity for a system that includes the needed physical laws and Barlow's regularity. In the face of such a worry, looking to a third field (not biology and not physics), the interests of which include only the biological and physical facts relevant to Barlow's work, seems like an attractive option. In this third field, the strength to be gained for a system by including the needed physical laws that cover the physical facts will be more in balance with the strength to be gained from including the biological laws that cover the biological facts. And the relative strength gains of including Barlow's regularity in a system go up because the regularity covers a higher fraction

of the facts being systematized. We can thus be more confident than before (when Barlow’s work was being located in biology) in saying that the KFRA allows for a set of laws—determined relative to the union of the relevant biological and physical kinds and facts—that includes Barlow’s regularity and the biological and physical laws relevant Barlow’s work. Locating Barlow’s work in its own field has the further benefit of fitting better into the Cartwright-style “dappled world” view by avoiding unification.

Given what has just been said on behalf of the KFRA, one might be inclined towards the following defense of the BBSA’s ability to accommodate Barlow’s work and interfield interactions in general. The problems with the all-or-nothing sort of fact relativity provided by kind relativity can be eliminated by moving to gerrymandered kinds. For example, consider the kind *bio-photon* that is associated with just the photon facts that are biologically interesting. If the set of basic kinds for biology, or the particular Barlow field, includes bio-photon (instead of photon), then there isn’t the glut of photon facts that made S_p overwhelmingly strong, and the BBSA will succeed at accommodating Barlow’s work as well as the KFRA. This may be correct, strictly speaking, but I will insist that Barlow and biologists are not interested in bio-photons. Barlow and biologists are interested (at least in part) in the same kinds as physicists, and that shared interest is essential to Barlow’s work being an example of interfield interaction. Furthermore, if it is bio-photon that is taken as the basic kind, then there will be no physical laws, which refer to photons, to be had in the laws of biology or the Barlow field, but only mimics that refer to bio-photons. So, while the BBSA may exploit gerrymandered kinds to accomplish

something like the fine-grained fact relativity of the KFRA, it can do so only at the cost of breaking the points of shared interest (i.e. the kinds and laws) that make interfield interactions truly *interfield*.

3.4 Fundamental Laws

A distinctive feature of the BBSA's (and by extension the KFRA's) accommodation of special science laws is that it is *egalitarian*; it confers equal metaphysical status upon the laws of all fields. But being egalitarian does not require abolishing worthwhile distinctions. For example, while Callender and Cohen are “ontological egalitarians” about kinds, they must retain a distinction between the fundamental and supervenient kinds. The egalitarianism comes from that fact that, while the fundamental and supervenient kinds may be distinguishable, the fundamental kinds (or any other set of kinds) are not treated as having any status over that of any other set of kinds.

A similar sentiment may be directed towards laws. There are the fundamental laws and the special science laws, and while they may be distinguishable, the fundamental laws should not be treated as having any privileged status. One might wonder why, if fundamental laws are not to be treated as having any special status, we should bother to even furnish some laws with the title “fundamental”. The reason to do so is because such a distinction is present in scientific practice.

Part of what motivates the development of an analysis of laws of nature is the desire for a better understanding of scientific practice. As discussed in the

preceding chapters, this is especially true among proponents of the BSA (and its variants) who—despite regular admonishment for having a view of laws that is too beholden to the peculiarities of actual human science (as in: [Armstrong, 1983, van Fraassen, 1989, Carroll, 1990])—celebrate each opportunity to have their view parallel actual human science (as in: [Lewis, 1983, Loewer, 2007, Cohen and Callender, 2009]). There clearly are scientists looking for fundamental laws (e.g., in their own ways: [Anderson, 1972, Weinberg, 1992]). Even a staunch anti-fundamentalist like Cartwright recognizes that “there is a tendency to think that all facts must belong to one grand scheme” [Cartwright, 1999, p. 25]. The existence of such a tendency, and a desire to parallel scientific practice, should be enough to make any proponent of the BSA (or its variants) interested in distinguishing fundamental laws from special science laws. To that end I propose a desideratum for an *egalitarian distinction* between fundamental and special science laws: No more or less respect should be paid to either fundamental or special science laws, but respect should be paid to their being distinct.

To illustrate the challenge that this poses, recall that part of the egalitarianism of the BBSA was that no set of laws depends for its existence on the existence of any other laws. The standard way to violate this requirement is to have special science laws be derivative of fundamental laws. It is easy to avoid this if there simply are no laws marked as fundamental for other laws to depend upon. But as soon as one makes the distinction between fundamental and special science laws, it is compelling to think that the fundamental laws are somehow more important, be that metaphysically, or in the practice of science, or both. Maybe that presumed

importance has to do with the existential dependence already mentioned. Maybe it's that an explanation that appeals to fundamental laws is better than one that appeals to non-fundamental laws. Or maybe fundamental laws are more important when it comes to thinking about scientific realism. I say that we should pay equal respect to fundamental and special science laws primarily in the sense that we should resist the compulsion to equate the fundamentality of some laws with their importance. In some contexts it may turn out that the fundamental laws are more important than special science laws—and vice versa in other contexts—but that inequality should be a consequence of considerations that go beyond our analysis of laws.

For both the BBSA and KFRA, we may approach the fundamental laws as we do the laws of any field. Just as the laws of biology are the laws determined relative to the interest of biologists, the fundamental laws are the laws relative to the interests of those scientists seeking fundamental laws. It is this similar treatment that contributes to the distinction being egalitarian.

The challenge for the BBSA is to identify the set of kinds of interest to fundamental science. Similarly, for the KFRA, it will be a matter of identifying the set of kinds and the set of facts of interest to fundamental science. I argue below (in Section 3.4.1) that the defining interest of fundamental science is an interest in systematizing all facts, while special sciences are interested in a proper subset of the facts. This is apparent from the literature on *ceteris paribus* laws, as well as from the often self-described aim of fundamental physicists being a “theory of everything”. The BBSA, without fact relativity, cannot capture this critical difference in interests between fundamental and special sciences. But kind relativity is still an

issue for the KFRA, and so I consider (in Section 3.4.2) a couple of strategies for picking out the set of kinds to which the fundamental laws are relative.

3.4.1 Fundamental Laws and Fact Relativity

The standard characterization of the distinction between fundamental and special science laws is this: Fundamental laws are universal, exceptionless, regularities. Special science laws are non-universal, exception-having, regularities, characterized by *ceteris paribus* (“all else being equal”, “CP” for short) clauses. That special science laws are so characterized is on display in a recent special issue of *Erkenntnis* on CP laws. It is noted in the introduction to that issue that

there seems to be broad agreement that the special sciences include non-universal generalizations ... The main motivation for accounts of CP laws seems to lie in the fact that these accounts should also be plausible for special science laws.

In fact, all authors of the special issue take this motivation as their starting point for CP laws.

[Reutlinger and Unterhuber, 2014, p. 1708]

While the connection between special science and CP laws is generally agreed upon, the connection between fundamental and exceptionless laws is more contentious. Led by [Cartwright, 1983], a number of authors have argued that fundamental laws do, in fact, have exceptions.¹⁴ But such exceptions need not infect

¹⁴ See also: [Cartwright, 2002, Pemberton and Cartwright, 2014, Hüttemann, 2014].

the *interests* of fundamental science. A common line in contemporary fundamental physics is that the goal is to find “the theory of everything”, a phrase that first appears in [Ellis, 1986]. And, as noted earlier, even Cartwright recognizes the fundamentalist’s “tendency to think that all facts must belong to one grand scheme” [Cartwright, 1999, p. 25]. Whether or not the fundamental laws actually are exceptionless may be set aside here. For our purposes it is enough if, insofar as there are scientists looking for fundamental laws, they are looking for a theory of everything. Cartwright can be right that the fundamentalist idea of a “grand scheme” is a bad one. There may be very good reasons for being more interested in the non-fundamental areas of science than in fundamental science. But fundamental science is still a part of science, and Cartwright’s characterization of the fundamentalist idea fits well with a “theory of everything” view of the *interests* of fundamental science. This attitude towards the fundamental laws (not Cartwright’s, but the openness to Cartwright’s) is part and parcel of the distinction between fundamental and special science laws being egalitarian. The fundamental laws do not have any special status; they are simply the laws of the system that is best relative to the interests of those scientists seeking fundamental laws.

So, as far as fact relativity is concerned, the distinction between fundamental and special science laws is clear. The fundamental laws are those that attempt, or are intended by their seekers, to be exceptionless. Special science laws are not exceptionless, do not attempt to be, and are not intended to be by their seekers. To be more precise for the KFRA: The laws of a system are fundamental only if that system is the best system relative to the set of all facts \mathcal{F} . The laws of a system are

special science laws if that system is best relative to some $F \subset \mathcal{F}$. The question of to which set of kinds the fundamental laws are relative will be considered later in this section.

The objection (from Cartwright and others, noted above) to the exceptionlessness of fundamental laws already has a response in the literature on the BSA, and it seems worthwhile to sketch how what I have just said relates to that work. Several authors [Braddon-Mitchell, 2001, Schrenk, 2008, Schrenk, 2014, Unterhuber, 2014] argue for various ways that laws (possibly fundamental) with exceptions may appear in the BSA (or its variants) when the introduction of such laws comes with gains in simplicity greater than the assured losses in strength. Such exceptions—call them *simple system exceptions*—are different from the sort of exceptions on offer by the fact relativity of the KFRA—call them *fact relative exceptions*. This difference is clear from the fact that simple system exceptions are available to any variant of the BSA, while fact relative exceptions require fact relativity. If there are exceptions in the laws sought by those seeking fundamental laws, they may be captured in the KFRA by simple system exceptions. But, to reiterate, these exceptions need not undermine the distinction between fundamental—as exceptionless—and special science—as exception-having—laws, because the fundamental laws are still the laws of the best system relative to all facts—without exception!—and special science laws are those relative to only some proper subset of the facts.

The relationship between fact relative laws and CP laws likely warrants attention beyond what may be provided here. For the moment it is enough that the all facts strategy is satisfactory as a way to capture an egalitarian distinction between

fundamental and special science laws with respect to fact relativity. With that in hand, we may turn to considering how to make the distinction with respect to kind relativity.

3.4.2 Fundamental Laws and Kind Relativity

If the above is correct—that the defining interest of fundamental science is an interest in systematizing all facts—then the BBSA cannot properly distinguish fundamental from special science laws because (according to the BBSA) every law set is determined relative to the set of all facts. But we must still consider kind relativity since it is a part of the relativity of the KFRA. At this point for the KFRA we have it that the laws are special science laws at least if they are (K, F) -relative laws for any K and any $F \subset \mathcal{F}$. This leaves as candidates for the fundamental laws the laws of systems that are best relative to (K, \mathcal{F}) for any K . These are precisely the laws of the BBSA. Answering completely the egalitarian distinguishing desideratum for the KFRA involves providing an answer to the desideratum on behalf of the BBSA (imperfect though it may be compared to the full answer given by the KFRA).

So, how might the BBSA, and the kind-relative part of the KFRA, distinguish between the kinds of fundamental science and the kinds of the special sciences? Of course, there are at least as many ways to pick out a set of kinds for the fundamental laws as there are sets of kinds. But two strategies in particular seem salient.

The first is to follow the BBSA's treatment of other fields to the letter. Let the fundamental laws be the laws of the system that is best relative to the kinds

that are treated as basic in actual fundamental science practice. A consequence of this strategy is that the fundamental laws become unstable. The kinds treated as basic in fundamental science one hundred years ago are different from those treated as basic now, and so the fundamental laws of nature (being identified with the laws relative to the kinds treated as basic by fundamental science) must have changed over the past hundred years. This change would be fine if we were talking about laws as determined directly by a scientific field. Of course the laws of fundamental science *as a field*¹⁵ have changed. But no one engaged in fundamental science would—and in the interest of paralleling scientific practice, no philosopher should—think that the fundamental laws *themselves* have changed.

The second salient strategy is to follow [Lewis, 1983] and privilege the true fundamental kinds. This suffers from the issue that, just like the laws of [Lewis, 1983], the fundamental laws are made potentially epistemically inaccessible.¹⁶ Compared to the strategy of making the fundamental laws relative to the kinds treated as basic by fundamental science, I am inclined to go this Lewisian route of privileging the true fundamental kinds, and bite the bullet with regards to concerns about inaccessibility. I suspect that it better parallels the attitude towards fundamental laws in scientific practice to prioritize (1) constancy in the fundamental laws and the potentially unsuccessful search for the true fundamental kinds, over (2) inconstancy

¹⁵ It might be more natural here to say “physics *as a field*”, but I want to avoid conflating physics and fundamental science. Certainly it is the case that fundamental science is carried out predominantly under the auspices of physics, but it is just as certain that lots of physics is explicitly not fundamental.

¹⁶ For details on this issue, see where it is first raised in Chapter 3 of [van Fraassen, 1989]. It is also worth mentioning that this is precisely the issue that motivates the epistemic accessibility desideratum of [Cohen and Callender, 2009] that is mentioned in footnote 6 of this chapter.

in the fundamental laws and a guarantee of epistemic accessibility.

In the next chapter I will argue that we can and should do without language relativity. In its place I put a combination of fact relativity—to deal with accommodating special science laws as discussed in this chapter—and any needed competition relativity—to deal with the issues surrounding language privileging/relativity discussed in the next chapter. If no competition relativity is needed, then it will be enough to distinguish fundamental and special science laws on the basis of fact relativity as discussed earlier. If competition relativity is needed, then we will be faced with answering the new question of which competition function (comprised of measures of simplicity, strength, their balance, etc.) are the ones that yield the fundamental laws and which yield special science laws. Either way, the current question of how to distinguish the fundamental from special science kinds will not need to be answered. It will not, however, be entirely uninteresting. As will be seen, using a particular competition function is tantamount to privileging a particular *class* of languages. If distinguishing between fundamental and special science laws with regards to competition relativity is needed, it may be best understood in terms of what kinds/languages different fields are willing to entertain.

3.5 Summary

The BBSA of [Cohen and Callender, 2009] makes a dramatic change to Lewis' BSA by allowing for a plurality of law sets each relative to the set of kinds treated as basic in the respective best system competitions. The egalitarian special science

laws of the BBSA depend, because of the BBSA's kind relativity, on the law-directed interests that distinguish different fields of science being captured by what kinds are treated as basic in the various fields. But the interests (in general) of fields go beyond what kinds are treated as basic. In particular, some fields may be interested in only a subset of the facts that are available for study. Accommodating such interests in a BSA-style view of laws may be done by relativizing laws to subsets of the facts of the world. While the kind relativity of the BBSA offers a certain degree of fact relativity on its own, it cannot unproblematically relativize laws to facts with great precision.

I introduced (in Section 3.2) the KFRA, which extends the BBSA with fact relativity. With both kind and fact relativity at its disposal, the KFRA can do the same good work as the BBSA and more beyond that. It was proposed (in Section 3.3) that views like the BBSA and KFRA should be able to make sense of interfield interactions. In an example of an interfield interaction between physics and biology concerning insect vision, it was found that the fact relativity of the KFRA was needed to avoid having laws of common and simple kinds outcompeting the laws of comparatively rare and complex kinds. Another desideratum was developed (in Section 3.4) according to which an egalitarian distinction between fundamental and special science laws should be available. Fact relativity is required to capture the standard picture of the fundamental laws as being exceptionless parts of a theory of everything. There does not seem to be a similarly strong strategy for satisfying the desideratum with respect to kind relativity. Because both the BBSA and KFRA feature kind relativity, neither satisfies the desideratum without issue, but the KFRA

benefits from the strength of the distinguishing allowed by fact relativity. Overall, if one is inclined towards a relative laws view like the BBSA, it seems like it would be entirely beneficial to adopt fact relativity and move to the KFRA.

Part II

Objectivity

Chapter 4: Language Privileging

In the tension between the BSA's relativity and objectivity, language privileging seems like it would be a good thing. Shouldn't having one privileged language for all competing systems reduce the need for relativity? In principle, perhaps. In practice, not so much. If we are really only going to privilege one language, then we are faced with a challenge of saying *which* language gets privileged. As will be seen, that challenge has not been met without taking on other substantial burdens. The challenge may be avoided, as in the BBSA, with the introduction of language relativity into the laws. There is language privileging in the BBSA, since, in any given best systems competition, a single privileged language is being used to express all competing systems. But the challenge of saying which language gets privileged is avoided because *every* language is privileged in its own language relativized competition.

There is a non-privileging and non-relativizing way of dealing with language in the BSA: Let system-language pairs compete in the same competition without restriction on the choice of language. This is, more or less, the default assumption about how the BSA works. But it is not how any developed version of the BSA works. Language privileging is ubiquitous, appearing in various forms in the BSA

of [Lewis, 1983], the PDA of [Loewer, 2007], and the BBSA of [Cohen and Callender, 2009].¹

Why do all of these authors adopt language privileging? The main reason is what I have called “Trivial Systems Problem” (TSP), according to which, in brief, allowing for suitably gerrymandered languages can guarantee that the “best” system will have axioms and theorems undeserving of the name “law”. Language privileging provides a quick fix to the TSP as long as the privileged language is not of the sort that gives rise to the TSP. Another reason for adopting language privileging is suggested by [Cohen and Callender, 2009]. Their “Problem of Immanent Comparisons” (PIC) is concerned with there being only “immanent” measures for simplicity, strength, and their balance—that is, measures defined for only one language. With language privileging, no two systems ever need to be compared when expressed in different languages, and so using only immanent measures is not an issue.

Relieving the tension that is the focus of this dissertation has been made a matter of properly tuning the four parts of the BSA to strike the right balance between the relativity needed to accommodate scientific practice and limits to that relativity to respect the metaphysical objectivity the view is supposed to possess. It is light of that aim to limit relativity that the overarching project of this chapter is to argue against the use of language relativity in the BSA. The TSP, it will be seen, can be solved without relying on language privileging. And the PIC is undermined by the existence of measures that, while still language sensitive, are non-

¹ As noted in previous chapters, the BBSA was introduced independently in [Schrenk, 2008] and [Cohen and Callender, 2009]. I will focus on the version of the BBSA appearing in [Cohen and Callender, 2009], since it is there that the problems discussed in this chapter receive the most attention.

immanent. The TSP and PIC are both reminiscent of, but not quite the same as, the well known problem of language sensitivity of measures in theory choice, statistical model selection, evidential favoring relations, prior selection, and so on. That well known problem does not require us to relativize to single languages, but rather *classes* of languages. Single language privileging/relativity also has a disadvantage in being unable to accommodate the idea that laws and basic kinds/predicates/etc. are discovered together. Even then, because which class of languages is relevant will be determined by the measures being used in the best system competition, language class relativity will prove to be redundant with competition relativity. In the end, language relativity of either the single or class variety will be unnecessary for the BSA.

The outline of this chapter is as follows. I begin, in Section 4.1, by laying out the TSP. I then review the various ways language privileging has been used to address the TSP by [Lewis, 1983], [Loewer, 2007], [Cohen and Callender, 2009], and discuss the additional problems that appear because of them, in Sections 4.2, 4.3, and 4.4, respectively. I mostly conclude the discussion of the TSP in Section 4.5, where I argue that the TSP can be avoided if we are careful with how we measure strength. With the TSP largely out of the way, I move on to the PIC in Section 4.6. There I review the PIC's introduction by [Cohen and Callender, 2009] and argue that the problem is weaker than its initial presentation suggests, and that it only requires us to privilege and relativize to classes of languages as opposed to single languages. I conclude in Section 4.7 by arguing for the adoption of language-class privileging over single language privileging, and then showing how language-class

privileging will be redundant with however much competition relativity is required for the BSA.

4.1 The Trivial Systems Problem

Prior to the introduction of the TSP, Lewis' BSA worked as follows. There are all the deductive systems S_i that are true of the world. Some of these are very strong, and some are very simple. In principle, strength and simplicity trade off. So, for example, a system that enumerates all the facts of the world is (presumably) very strong, but it is also (presumably) very complex. A system that only specifies the location in space and time of a single electron would be very simple, but it is also very weak. The best system should be one that strikes a balance between these extremes, with axioms and theorems—which will be named as the laws—yielding substantial strength without too much loss in simplicity.

The TSP begins with the realization that how simple a system is might vary depending on the language in which it is expressed. The simple and weak system mentioned above could be a very complex system if expressed in a language that does not have 'electron' as a basic kind, and instead requires the specification of numerous properties to pick out an electron as the object of a sentence. This also works for an incredibly strong system, as Lewis notes:

Given system S , let F be a predicate that applies to all and only things at worlds where S holds. Take F as primitive, and axiomatize S (or an equivalent thereof) by the single axiom $\forall xFx$. If utter simplicity is so

easily attained, the ideal theory may as well be as strong as possible.

[Lewis, 1983, p. 367]

With such a procedure available to make any system incredibly simple, the trivial systems problem is this: An analysis of laws should be able to distinguish between accidental regularities and the lawful regularities. But the theorems of a maximally strong system will include every regularity that holds in the world, and so every true regularity will be a law. To avoid this outcome we must find a way to rule out the possibility of there being such system-language pair. Lewis tells us that “the remedy, of course, is not to tolerate such a perverse choice of primitive vocabulary” as F , and so begins the tradition of language privileging in the BSA [Lewis, 1983, p. 367].

4.2 Lewis’ BSA

Lewis’ response to the trivial systems problem is this:

We should ask how candidate systems compare in simplicity when each is formulated in the simplest eligible way; or, if we count different formulations as different systems, we should dismiss the ineligible ones from candidacy. An appropriate standard of eligibility is not far to seek: let the primitive vocabulary that appears in the axioms refer only to perfectly natural properties.

[Lewis, 1983, pp. 367–368]

The problematic predicate F presumably does not correspond to a perfectly natural property. So, now that no competing system is allowed to be expressed with F , no competing system will have the trivializing qualities that come with F .²

Problems with Lewis’ solution to the TSP are identified by [van Fraassen, 1989]. These problems, described at length in the remainder of this section, show how Lewis’ privileging of the ‘perfectly natural properties’ yields laws that are inaccessible and possibly uninteresting to scientists. Given the commitment of BSA proponents—myself included—to drawing a connection between the BSA and scientific practice, these problems are decisive against Lewis’ version of language privileging.

To set up the problems, call the set of possible systematizations of the world \mathcal{S} , the set of languages that those systems may be expressed in \mathcal{L} , and $C(S_i, L_j)$ the score the system $S_i \in \mathcal{S}$ receives in the best system competition when it is evaluated according to its expression in the language $L_j \in \mathcal{L}$. The heart of the trivial systems problem when formulated this way is that there will be law-wise problematic systems $S_{\text{bad-}i}$ —that is, systems whose associated regularities are undeserving of the name “law” for some reason (such as failing to distinguish between accidental and lawful regularities)—for which there exist companion languages $L_{\text{bad-}i}$ that make $C(S_{\text{bad-}i}, L_{\text{bad-}i})$ arbitrarily large. Lewis solution, where $L_{\text{nat}} \in \mathcal{L}$ is the language of perfectly natural properties, is to evaluate each system according to

² The possibility remains that F does, in fact, correspond to a perfectly natural property. I suspect Lewis would say in response to this possibility the same thing he says about the possibility of very different systems tying as best: “in this unfortunate case there would be no very good deservers of the name of laws. But what of it? We haven’t the slightest reason to think the case really arises” [Lewis, 1994, p. 479].

$C(S_i, L_{\text{nat}})$, rather than allow every system-language pair (S_i, L_j) to be entered into the competition. Let S_{nat} be the best system when all systems are expressed in the language of perfectly natural properties L_{nat} .

What van Fraassen asks us to consider is possibility of there being a system-language pair $(S_{\text{vf}}, L_{\text{vf}})$ such that $C(S_{\text{vf}}, L_{\text{vf}}) > C(S_{\text{nat}}, L_{\text{nat}})$. Suppose further (and just for convenience) that, if there is more than one such system-language pair satisfying these conditions, $(S_{\text{vf}}, L_{\text{vf}})$ is the best of them. Now, it is certain that $C(S_{\text{vf}}, L_{\text{nat}})$ is less than $C(S_{\text{nat}}, L_{\text{nat}})$, for otherwise S_{nat} would not be the Lewisian best system. However, without some outside assurance that L_{nat} is the proper language to use when identifying the laws, S_{vf} seems like it may be the *scientifically* best system (barring it having any problematic qualities on par with failing to distinguish between lawful and accidental regularities). The two reasons for that seeming constitute “van Fraassen’s problems”, so named following [Loewer, 2007] (whose formulations of the problem I loosely follow).

The first of van Fraassen’s problems is concerned with the accessibility of the laws. Even if scientists are devout Lewisians with access to every matter of fact in the world, they will never discover the laws as long as they are naive to the true perfectly natural properties. And we should expect that scientists are naive to the true perfectly natural properties—the perfectly natural properties aren’t known prior to the identification of the laws, but rather, as Lewis says, “the laws and natural properties are discovered together” [Lewis, 1983, p. 368]. That the laws and properties will go undiscovered is on account of the quality of the pair $(S_{\text{vf}}, L_{\text{vf}})$. Since $C(S_{\text{vf}}, L_{\text{vf}}) > C(S_{\text{nat}}, L_{\text{nat}})$, and C is meant to capture all that may be

considered in the identification of the laws, there is no basis on which scientists could ever pick out L_{nat} over L_{vf} as the proper language with which to identify the laws. L_{vf} , since it is part of the best system-language pair, will be the source of scientist's best estimate of the perfectly natural properties, and that best estimate will be wrong. Even if nature is kind in that there is no pair $(S_{\text{vf}}, L_{\text{vf}})$ better than the best system as expressed in the language of perfectly natural properties, any sensibly skeptical scientists will know that they can never be sure about the laws.

The second of van Fraassen's problems is concerned with how deserving Lewis' laws are of their title. If $(S_{\text{vf}}, L_{\text{vf}})$ really is the best (unproblematic) system-language pair, then why is it not the source of the laws? Suppose that scientists know somehow that L_{nat} is the language of the perfectly natural properties. Then there is no issue about the accessibility of the laws. Scientists know that S_{nat} is the best of the systems when expressed in L_{nat} . They also know that there is a simpler and more informative system S_{vf} in the offing if they just abandon L_{nat} in favor of L_{vf} . The only reason to stick with L_{nat} and S_{nat} in these circumstances is a question-begging desire to have the laws expressed in the language of perfectly natural properties. One cannot even appeal to an expectation that laws and natural properties are discovered together since, by assumption, the natural properties were identified independently of the laws.

van Fraassen's problems are strongest when taken together. Take Lewis' BSA with its requirement that language of the competing systems be the language of the perfectly natural properties and the presumption that the perfectly natural properties are not discoverable independent of the laws. Then the laws are inaccessible—

since science’s best estimate of the laws will either be wrong or, at best, it will be unknowable that it is correct—and potentially uninteresting—since there may be a language other than that of the perfectly natural properties that will yield a simpler and stronger, by definition *better*, best system. As van Fraassen puts it: Lewis’ language privileging “has produced unchartable distances between Lewis’s best theories—and hence laws—and the theories we could reasonably hope for at the ideal end of science” [van Fraassen, 1989, p. 55].

4.3 The Package Deal Analysis

The quickest solution to van Fraassen’s problems is to remove the perfectly natural property language requirement from the the set of systems. But simply removing the requirement reintroduces the trivial systems problem. To compensate, the BSA must be enriched somewhere else. The strategy pursued by the the “Package Deal Analysis” (PDA)—so named after its emphasis on evaluating systems and languages together as a “package deal”—of [Loewer, 2007], is, I think, almost right. The view focuses on the choice of language for expressing the laws being constrained by the theoretical virtues that decide which system(-language pair/package) is best; I will argue for more or less the same thing in Section 4.5. But Loewer is also committed to language privileging in a way that makes the PDA dependent on subjects (namely, currently practicing physicists). The PDA and its subjectivity problem are described below.

The PDA begins by putting constraints on what can count as a candidate

system in the best system competition. The set of candidate best systems $\mathcal{S}_{\text{PDA}} \subseteq \mathcal{S}$ for a world w is decided by three conditions that must be met by each $S_i \in \mathcal{S}_{\text{PDA}}$:³

1. S_i is formulated in L .
2. S_i is true of w .
3. (S_i, L) is a final theory of w , meaning “ (S_i, L) is true and best satisfies the criteria of simplicity, informativeness, comprehensiveness, and whatever other conditions the scientific tradition places on a final theory for w ” [Loewer, 2007, p. 324].

This leaves us with a set of system-language pairs (S_i, L) . From among those the best have the greatest value of $C(S_i, L)$ and the laws of w come from the best.

The first condition admits of two interpretations. It may simply be a statement of the fact that there must be *some* language in which S_i is eventually formulated. In that case no work is done by the first condition to avoid the problem of overly strong and simple theories. It is the third condition that can then come to the rescue. In order for (S_i, L) to be a final theory, $C(S_i, L)$ must be large relative to other competing pairs, and part of achieving that is meeting “whatever other conditions the scientific tradition places on a final theory for w ”. Loewer writes in his remarks on the TSP that

from the perspective of the aims of science the obvious trouble with
‘ $(x)\text{Fx}$ ’ is not that ‘ Fx ’ doesn’t refer to a perfectly natural property but

³ I will follow Loewer’s characterization of a theory as the pair of a system and language it is expressed in, but while Loewer refers to the system sometimes as a theory, and labels it T , I will strictly use S to stand for the system and adjust quoted material accordingly for clarity.

that ‘(x)Fx’ is not a credible scientific theory.

[Loewer, 2007, p. 324]

And so, presumably, the failure of ‘(x)Fx’ to be a credible scientific theory is codified in the “other conditions” for a final theory.

The second interpretation of “ S_i is formulated in L ” is that it is saying there is some *single* privileged language L in which all the considered systems must be formulated. Loewer appears to be following this interpretation when he writes that “Lewis’s argument does show that the PDA version of the BSA requires a preferred language” [Loewer, 2007, p. 325]. And what is that preferred/privileged language? This is precisely Lewis’ strategy if it is the language of perfectly natural properties L_{nat} . But Loewer has a different idea. The privileged language for the PDA must be the language L_{final} of the final theory because (S_i, L) is a final theory by the third condition and it would fail to be if L was not L_{final} . As to the nature of L_{final} , we are provided the following working hypothesis: Let L_{present} be the present language of science. L_{final} is L_{present} or a successor to L_{present} as arrived at by “the rational development of science” such that “no development of that language and [the system S employed by science] leads to an increase in the satisfaction of the scientific virtues” [Loewer, 2007, p. 325].

The first way of interpreting “ S_i is formulated in L ”—as saying simply that there must be *some* language L in which S_i is expressed—solved the trivial systems problem by appeal to the scientific credibility, or lack thereof, of theories that have the negative qualities that come with F . This second interpretation does something

similar, but instead of appealing to problematic theories, F is ruled out directly as a predicate that is not employed in the present language of science, and would not be employed in a successor (particularly, the final) language of science.

The problem with the PDA is that it makes the laws dependent on subjects. The view features explicit references to things like “the scientific tradition” and “the present language of science”, and in so doing violates the objectivity of the laws. Loewer is open about this, writing that “what counts as a final theory depends on the tradition of fundamental physics”, but adds “I see this as an advantage of rather than an objection” [Loewer, 2007, p. 325].⁴ This is an advantage insofar as it reconnects physics and the laws in a way that answers to van Fraassen’s problems, but it will be a mark against the PDA if there are any similarly accomplished competitors that avoid subject dependence.

4.4 The Better Best Systems Analysis

Lewis went too far with the BSA being a metaphysical analysis of laws, privileged the language of “perfectly natural properties”, and in so doing broke the connection between the BSA’s metaphysics and the epistemology of the search for laws in scientific practice. Loewer went too far in the other direction, privileging the “the present language of science”, and in so doing broke the objectivity of the

⁴ Loewer has suggested in personal communication that a historical examination of the concept of laws might reveal that subject independence is not a real desideratum. Such an examination has yet to be done, but suppose that Loewer is right about what it will reveal. It seems like it could still be *advantageous*, if not strictly necessary, for the BSA or its variants to yield subject independent laws when they are being weighed against competing analyses of laws. The disinterest of subjects in subject independence will never dissuade those who are already committed to subject independence. What would be problematic for the purposes of this dissertation is if science is unequivocally committed to subject dependence, but that strikes me as highly implausible.

BSA/PDA. In ways that have already been described in the preceding chapters, the BBSA strikes an elegant balance between the metaphysics and epistemology of laws. To quickly review:⁵ The fundamental kinds (the perfectly natural properties, in Lewis' parlance) are grouped with all the kinds that supervene on the fundamental kinds to make a set \mathcal{K} of all kinds. Each subset $K_i \subseteq \mathcal{K}$ corresponds to a language L_i whose basic predicates correspond to the kinds of K_i . A best system competition is run relative to every K_i , with all competing systems expressed in L_i . The victor of each such competition will be the source of the K_i -relative laws. This gets us objectivity, because there is no explicit reference or dependence on subjects. And it gets us the connection between the BSA and scientific practice, because (supposedly) the full breadth of scientific interests (at least as far as the laws are concerned) may be captured by what variations in what kinds are treated as basic.

The BBSA, then, seems to avoid the issues that caused trouble for Lewis' BSA and the PDA. But does it solve the TSP? Well... not quite. To avoid the subject dependence that would come from there being laws just and explicitly relative to the kinds treated as basic by various scientific fields, the BBSA yields laws relative to *every* subset of the set of all kinds. Since the trivializing predicate F can be associated with a supervenient kind, this means that there will be many sets of laws that suffer from the trivial systems problem as a result of the inclusion of the F kind among the kinds to which those laws are relative. The response to this concern is that, since we aren't interested in F , we'll never have to worry about being interested

⁵ See Chapter 3 for the more detailed introduction of the BBSA.

in the problematic laws that an *F*-inclusive language yields. “Properties like *F* and the ensuing threatened trivialization of [BSA-style views] are ruled out for lack of interest rather than any intrinsic deficiency” [Cohen and Callender, 2009, p. 23]. But the laws yielded by *F* are intrinsically deficient—because they fail to distinguish between lawful and accidental regularities—and they are real laws, of interest or not, according to the BBSA. I rather like the strategy of appealing to what is of interest or not, and do not take this near miss of solving the TSP to be a decisive objection to the BBSA. That said, it is still a deficiency. If another otherwise equal capable view could avoid it, that other view should be preferred over the BBSA.

The real problem for the BBSA stems from the idea that there is an close relationship in scientific practice between the discovery of laws and the discovery of the basic kinds. Lewis claims that language privileging in the BSA

explains [...] why the scientific investigation of laws and of natural properties is a package deal; why physicists posit natural properties such as the quark colours in order to posit the laws in which those properties figure, so that laws and natural properties get discovered together.

[Lewis, 1983, p. 368]

But language privileging does just the opposite. By privileging a language prior to the best system competition, the discovery of the basic properties/kinds of that language is divorced from the discovery of the laws. The relationship should be there, though. From groupings of organisms in biology to classifying particles in fundamental physics, scientists are constantly reflecting upon and revising their sets

of basic kinds. And, insofar as these scientists are looking for laws, these revisions happen in tandem with the identification of the laws; see, for example, the BSA-friendly analysis in [Schulte, 2008] of (as the paper’s informative title puts it) “the co-discovery of conservation laws and particle families”. Privileging single languages prior to running a best system competition prevents this co-discovery feature of laws and kinds from being captured by the BBSA and any other language privileging variant of the BSA. The discussion of this issue is extended below in Section 4.7.2.

4.5 Avoiding Trivial Systems

In reassessing the TSP, it is helpful to start with its original setup from Lewis:

Given system S , let F be a predicate that applies to all and only things at worlds where S holds. Take F as primitive, and axiomatize S (or an equivalent thereof) by the single axiom $\forall xFx$. If utter simplicity is so easily attained, the ideal theory may as well be as strong as possible. Simplicity and strength needn’t be traded off. Then the ideal theory will include (its simple axiom will strictly imply) all truths, and a fortiori all regularities. Then, after all, every regularity will be a law. That must be wrong.

[Lewis, 1983, p. 367]

The crux of the argument is that $\forall xFx$ is problematic, so it should not be the best system, but its utter simplicity and strength guarantee that it is the best. The standard assumption about what makes a system best is that simplicity and

strength trade off—a simpler system will tend to be weaker, and a stronger system will tend to be more complex—and the basis of the trivial system problem is a counter example to this give and take.

All the BSA proponents discussed so far take the trivial systems problem to be a sort of *reductio* against the legitimacy of F as a predicate that may be considered when looking for laws. But there is another standout part of the argument against which we might apply the *reductio*, and that is the claim that $\forall xFx$ is utterly simple and strong.

I will focus here on the purported strength of the trivial system.⁶ The standard gloss of strength is that the stronger theory excludes more possibilities (see, for example: [Woodward, 2013, Loewer, 2012]). In some of Lewis' more substantial remarks on the subject, he writes that

the stronger theory may have fewer actual realizations or it may not;
but it must have less risk of multiple realization [...] the stronger theory
must also have more risk of nonrealization.

[Lewis, 1970, p. 434]

The strength of a system, crucially, depends on the space of possible worlds of which it may (or may not) be true. Roughly, a stronger system is one that has, given a set of possible worlds, the higher ratio of possible worlds in which it is false to possible worlds in which it is true. This ratio is surely not all that may be said about

⁶ A more careful examination of other theoretical virtues might also be able to block the TSP. For example, in Chapter 5, I show that mutual information will come out to be zero for the trivial system. When, as in that chapter, mutual information is introduced into the BSA as a measure of the goodness of the of a system, that score of zero undermines the trivial system's claim to being the obvious best.

strength. For example, it will give us nonsense when there is a countable infinity or unbounded continuum of possible worlds. But whatever the true measure of strength is (or measures are, if we want to admit a variety) it should make this rough characterization true in the appropriate circumstances; namely, when it is applied to possible worlds that are finite and discrete or bounded continua.⁷ The important thing here is getting at what is meant by “less risk of multiple realization” and “more risk of nonrealization”, which the ratio measure does at least for the the simple example that is developed below.

To illustrate, consider the space of possible ‘coin flip’ worlds with four flips. Each such world consists of four ordered positions at which either the property H or the property T obtains. One of these worlds goes HHHH, another HHHT, HHTH, HHTT, and so on up to TTTT (as if we were counting from 0 to 15 in binary). Let S_{2H} be the system that says that two points in the world are H. This is true of six of the possible sixteen worlds. Let S_{1H} be the system that says that only one point in the world is H. This system is true in four of the sixteen worlds, and so we might say, relative to this set of possible worlds, that S_{1H} is stronger than S_{2H} . Now suppose that there are only four possible worlds, HHTT, TTHH, HTTT, and TTTH. S_{2H} is true in 2 out of the four worlds and S_{1H} is also true in two of the four. So, relative to this space of possible worlds, the two systems are equally strong. And clearly we could identify a set of possible worlds in which S_{2H} is stronger.

⁷ I actually think the true circumstances are likely to be the appropriate circumstances. Science itself can only attend to bodies of data that are either finite and discrete or bounded continua. Preserving the connection between the BSA and scientific practice may require being sensitive to the finitude of that practice and so force the BSA to treat worlds as being either finite and discrete or bounded continua.

We can get up to all sorts of mischief by tweaking the space of possible worlds against which we are measuring strength. This parallels the way that we seem to be able to mess with measures of simplicity by changing the language used to express the system. The question then becomes: What determines the space of possible worlds?

The Lewisian answer is surely to appeal to the *true* space of possible worlds. This plays into both Lewis' realism about possible worlds and his "perfectly natural properties" response to the trivial systems problem. And, like with privileging the language of perfectly natural properties, versions of van Fraassen's problems will crop up with regards to the privileged space of possible worlds. Scientists do not have special access to the true space of possible worlds—just as they do not have special access to the true perfectly natural properties—and so there is a risk that the laws are epistemically inaccessible and uninteresting.

The PDA and BBSA solution to this problem, following the parallel with language, is probably to look at the space(s) of possible worlds that is (are) used in science. This cannot be, though, the space of nomologically possible worlds. If that's what we used, then either we are stumped, since this is happening prior to the identification of the laws, or strength is trivialized if we use the laws of the system in question, since then every system would be true of exactly all the (nomologically) possible worlds and have zero strength.

A better set of possible worlds to use is the set of *constructed worlds*. To say what the constructed world are, let us follow Lewis' characterization of the fundamental nature of the world as a Humean Mosaic (HM). Recall that the HM

consists in “the spatiotemporal distribution of local qualities” [Lewis, 1994, p. 473]. If our language contains the predicates for the set of basic kinds K , then a world is constructed by assigning to each space-time point in it at least one element of K . Then the constructed (HM-style) worlds correspond to all the ways of assigning non-empty⁸ sets of basic kinds to every space-time point.

Let’s revisit the example of the four-flip coin flip worlds. Our language, call it $L_{4\text{-HT}}$, gives us a four position HM and the properties H and T. The constructed worlds, then, are the sixteen that were first considered. Befitting the package-deal aspect of the PDA, and the stated dependence of strength on language choice in [Cohen and Callender, 2009], the strength of a system depends on the system and the language in which it is expressed. We thus want to compare the strengths of the pairs $(S_{2\text{H}}, L_{4\text{-HT}})$ and $(S_{1\text{H}}, L_{4\text{-HT}})$, and (as was said at first) we will get that $(S_{1\text{H}}, L_{4\text{-HT}})$ is stronger than $(S_{2\text{H}}, L_{4\text{-HT}})$ since it is true of a smaller fraction of the worlds that may be constructed following in $L_{4\text{-HT}}$.

Now we can consider the strength of the trivial system. The system is just $\forall xFx$. What is the relevant language? In the interest of simplicity, let us require that we pair systems with “minimal” language(s) that contain no more basic predicates than the ones required to express the system in that language. I do not mean by this requirement that every system should be paired with a language that has the fewest basic predicates. Requiring that would compel us to pair every system with a language containing a single F -like predicate. Rather, what I intend to prevent

⁸ It is fair to ask at this point why there can be no part of a constructed world is empty of any basic kind. I will address that question, and the objection it gives rise to, at the end of this section.

is $\forall xFx$ being paired with a language that contains F and some other arbitrary predicate G . The following, however, is perfectly acceptable: Let $Px \vee Qx$ be true just in case Fx is true. The system $\ulcorner \forall xFx \urcorner$ (with $\ulcorner \urcorner$ denoting that the system is unpaired with a language) may be expressed as $\forall xFx$ or $\forall x(Px \vee Qx)$. When expressed as $\forall xFx$, it is required that the expressing language have F as its only predicate. But the system $\ulcorner \forall xFx \urcorner$ can also be expressed as $\forall x(Px \vee Qx)$. In that case, the expressing language should have P and Q as its only predicates.

For $\forall xFx$, the minimal languages will be ones that have F as their only predicate—any language that lacks even that will make the system inexpressible. If F is the only predicate in our language, which we'll call L_F , then $\forall xFx$ will be true of all L_F -constructible worlds. That means that $(\ulcorner \forall xFx \urcorner, L_F)$ has zero risk of non-realization, and so its strength is actually zero. Since the strength of $(\ulcorner \forall xFx \urcorner, L_F)$ is zero, it's not the case that $\forall xFx$ is utterly strong (when joined by L_F).

Let $F_{P,Q}$ be the language containing just the predicates P and Q as defined above. The strength of $(\ulcorner \forall xFx \urcorner, L_{P,Q})$ is *also* zero. As long as we pair $\ulcorner \forall xFx \urcorner$ with a minimal language, its strength will always be zero. This is because its unconditional universality—precisely the quality that makes it fail to distinguish between accidental and lawful regularities—will always make the system true of every (minimal language constructible) possible world.

So it is not true, as [Lewis, 1983] thought, that “simplicity and strength needn't be traded off”. Simplicity and strength *do* trade off: Simpler languages yield more limited spaces of possible worlds. Strength depends on a system's risk of non-realization relative to the space of possible worlds yielded by the paired language.

Thus, simpler languages constraint the amount of strength that can be possessed by a system. Utter simplicity may still be easily achieved by changing our choice of language. But utter simplicity cannot be so easily achieved simultaneously with utter strength, and thus the TSP is blocked.

The objection could be made that this is just one way of measuring strength, and other ways might not be so helpful in addressing the TSP. For example, suppose that we give up the requirement that space-time points in constructible worlds be non-empty. If, say, worlds have fixed size of n space-time points, then there will be 2^n possible worlds of which $\forall xFx$ will be true of exactly one. Against this particular proposal, I might try to argue that emptiness seems like the sort of thing that might itself be a basic kind.⁹ But there is a more general counter to the object that the preceding solution to the TSP depends on the specific measure of strength: To start, *of course* the solution depends on the specific measure of strength that is used. If some other measures are to be used in the best system competition, it is quite possible that they too might block the TSP. Indeed, if some measures are proposed for the best system competition that do not block the TSP, I would think that is a mark against them when there are measures in the offing that do. The broad point is that there is at least one measure that can block the TSP, and that should give us pause when we try to infer from the existence of the TSP to a need for single language privileging in the BSA.

⁹ You could say that at each spatio-temporal point that either Fx or $\neg Fx$ is true. But maybe $\neg Fx$ is better understood as making a positive claim that $\emptyset x$ is true. I think this is plausible, but do not intend to pursue such a counterargument.

4.6 The Problem of Immanent Comparisons

The “Problem of Immanent Comparisons” (PIC) begins with an appeal in [Cohen and Callender, 2009] to a distinction in Quine between *immanent* and *transcendent* notions. Quine writes: “A notion is immanent when defined for a particular language; transcendent when directed to languages generally” [Quine, 1970, p. 19]. Measurements of simplicity, since they depend on the language in which a system is expressed, are taken by Cohen and Callender to be immanent in this Quinean sense. Strength is similarly immanent, since it is assumed to depend on the expressive power of the language in which a system is expressed. And, to finish out the set, balance is said to be immanent as well, since it will be a measure dependent on immanent measures of simplicity and strength. If two systems are competing to be the best and are expressed in different languages, then we would need transcendent measures of simplicity, strength, and balance, in order to implement the best system competition. But “there are too few (viz. no) transcendent measures” of simplicity, strength, and balance [Cohen and Callender, 2009, p. 8]. Cohen and Callender write that

Prima facie, the realization that simplicity, strength, and balance are immanent rather than transcendent—what we’ll call *the problem of immanent comparisons*—is a devastating blow to the [BSA and its variants]. 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.

[Cohen and Callender, 2009, p. 6, emphasis in original]

The only solution to the PIC, since systems can only be compared when they are expressed in the same language, is to adopt language privileging.

My objection to the PIC is that it ignores the existence of a large middle ground of measures that are neither immanent nor transcendent. At the PIC's core are the well known problems of language dependence that crop up in many guises—as gruesome predicates spoiling evidential favoring relations, reparameterizations that recommend conflicting priors or statistical models, and so on—in epistemology and philosophy of science. After illustrating the existence of the middle ground measures that undermine the force of the PIC, I will discuss what remains of the PIC as a result of the better known problems of language dependence.

To start, let us examine the central claim of the PIC: that simplicity, strength, and balance must be immanent measures. In defense of the idea that simplicity is immanent, [Cohen and Callender, 2009, p. 5] defer to [Goodman, 1954] by way of Loewer, who writes: “Simplicity, being partly syntactical, is sensitive to the language in which a theory is formulated” [Loewer, 1996, p. 109]. Loewer (and Goodman) are exactly right. Simplicity is language sensitive. For example, let us adopt a naive version of simplicity, $SimpC(-)$, that is measured by the number of characters it takes to express a sentence (including spaces and punctuation). Consider the following sentence.

This sentence is simple.

Its *SimpC*-simplicity is 24 characters. The same sentence in Dutch is

Deze zin is eenvoudig.

The sentence's *SimpC*-simplicity now is 22 characters. So the *SimpC*-simplicity of a sentence depends or is sensitive to the language in which the sentence is expressed. Does that language sensitivity mean that *SimpC* is immanent? It depends on what is meant by being "defined for a particular language".

SimpC is, in some sense, "defined for a particular language". Insofar as the measure gives conflicting results for a sentence expressed in different languages, it would be ill-defined if we took it to be directed at sentences irrespective of the language in which they are expressed. One way of dealing with this would be to think that we have a multitude of distinct simplicity measures $SimpC_{\text{English}}(-)$, $SimpC_{\text{Dutch}}(-)$, and so on. But doing that disguises an important fact; each of these measures of simplicity is *the same measure*, just relativized to particular languages. Drawing our inspiration from the "package deal" of the PDA, we could just as easily deal with the language sensitivity of *SimpC* by saying it is defined for sentence-language pairs. We don't need, then, different measures of simplicity. Just the one will do:

$$SimpC(\ulcorner \text{This sentence is simple.} \urcorner, \text{English}) = 24 \text{ char.}$$

$$SimpC(\ulcorner \text{This sentence is simple.} \urcorner, \text{Dutch}) = 22 \text{ char.}$$

In this way, *SimpC* is better understood as transcendent, and not immanent, be-

cause it is, as Quine put it, “directed to languages generally”.

Of course, *SimpC* can’t be directed to *all* languages, since it will be undefined for any languages that don’t have a written form with discrete characters. This suggests that there is an important middle ground between immanent and transcendent measures. When a measure falls in that middle, as *SimpC* seems to, I will say that it is a “non-immanent measure”.

So which conception of *SimpC* is the right one? The “devastating blow” that immanence deals to the BSA and its variants is that it “undercut[s] the possibility of intersystem comparisons” [Cohen and Callender, 2009, p. 6]. In our naive example,

$$SimpC_{\text{English}}(\ulcorner \text{This sentence is simple.} \urcorner)$$

is—if *SimpC* is immanent—incomparable to

$$SimpC_{\text{Dutch}}(\ulcorner \text{This sentence is simple.} \urcorner).$$

But obviously it’s not. $\ulcorner \text{This sentence is simple.} \urcorner$ is *SimpC*-simpler in Dutch than in English (when being *SimpC*-simpler means having a lower value of *SimpC*).

Nothing prevents a transcendent or non-immanent measure from taking a language as one of its arguments. Such a measure is transcendent (or non-immanent), but language sensitive, and, importantly, it allows for comparisons even when a variety of languages are involved. That being the case, the mere language sensitivity of simplicity, strength, and their balance is not enough to guarantee that they are

immanent, nor is it enough to guarantee the incomparability of systems expressed in different languages.

In response to the existence of a measure like *SimpC*, Cohen has suggested¹⁰ that there may well be transcendent (or non-immanent) measures plausibly named “simplicity” (etc.), but these are not the ones relevant to the BSA; the measures that *do* appear in BSA will be immanent. Cohen is absolutely right to question the plausibility of a measure as naive as *SimpC* having a role to play in the BSA. I certainly do not intend to defend *SimpC* as the right measure of simplicity for the BSA.

But I do not think it is clear why we should follow [Cohen and Callender, 2009] in assuming that the right measures are immanent. Cohen suggests that the burden falls on the defender of transcendent (or non-immanent) measures to show that the measures of the BSA will be such. Given the dearth of proposed measures for the BSA on either side, I imagine that one could argue that the burden falls the other way. But that is not what I intend to do. Rather, I think I have *just shown* in the preceding section that the BSA is likely to include non-immanent measures. Strength, I concluded, depends on a system’s risk of non-realization relative to the space of possible worlds yielded by the paired language. There is language sensitivity, but, like with *SimpC*, it is language sensitivity that can be captured by treating strength as a non-immanent measure of system-language pairs. The case for non-immanent measures will be made stronger (implicitly) in the next chapter when I argue for including mutual information in the best system competition of the BSA.

¹⁰ In personal communication.

Non-immanent measures are, if anything, the norm, and an example may be found in the selection of statistical models. Following [Forster and Sober, 1994], statistical model selection has standardly been associated in philosophy with the Akaike Information Criterion (AIC):

$$\begin{aligned} AIC(M) &= 2[\text{number of parameters of } M] \\ &\quad -2[\text{maximum log-likelihood of } M] \end{aligned}$$

The full details of AIC are not terribly important for our purposes here; it is enough to point out that that first term is concerned with the *number of parameters* of M . Forster and Sober note that the number of parameters “is not a merely linguistic feature” of models [Forster and Sober, 1994, p. 9, fn. 13]. But the number of parameters is *a* linguistic feature of a model. Since AIC can compare models with different numbers of parameters, it can—if we think of statistical models as the system-language pairs of the BSA and AIC as central to the best system competition¹¹—compare systems expressed in different languages. AIC is thus a non-immanent measure.

It is important to note, however, that AIC is also not a transcendent measure. [Kieseppä, 2001] offers a response to critics of AIC who are concerned that the measure is sensitive to changing the number of parameters of a model by changing the model’s linguistic representation. The response turns on the justification of

¹¹ To make the connection between AIC and the BSA even stronger, it is worth noting that [Forster and Sober, 1994] take the “number of parameters” term to be tracking the simplicity of a model. This is not unlike (though also not the same as) in the previous section when, in the interest of simplicity, we required systems to be paired with languages that contained a minimal number of basic predicates.

“Rule-AIC”, which says to pick the model with the smallest value of AIC, on the grounds that the predictive accuracy of model M is approximately the expected value of the maximum log-likelihood of M minus the number of parameters of M . Crucially,

the theoretical justification of using (Rule-AIC) is valid when the considered models are such that the approximation [just mentioned] is a good one.

[Kieseppä, 2001, p. 775]

Let M be parameterized to have either k or k' parameters. Then there are two claims that are relevant to the justification of Rule-AIC:

predictive accuracy of $M \approx E[(\text{maximum log-likelihood of } M) - k]$

predictive accuracy of $M \approx E[(\text{maximum log-likelihood of } M) - k']$

The predictive accuracy of M is independent of the number of parameters used to express M .¹² But the right side of the approximation in each claim *does* depend on the number of parameters. In general, both of these claims will not be true. Since Rule-AIC is only justified by the truth of these approximations, it will only be applicable to whichever parameterization of M makes the approximation true. The only time when both claims are true, and thus when AIC is applicable to both parameterizations, is when the difference between $E[(\text{maximum log-likelihood$

¹² This is, of course, intuitively true. It is also true in the formal definition of predictive accuracy given in [Kieseppä, 1997] and used in this argument from [Kieseppä, 2001].

of $M) - k]$ and $E[(\text{maximum log-likelihood of } M) - k']$ is negligible. Kieseppä concludes:

This simple argument shows once and for all that the fact that the number of the parameters of a model can be changed with a reparameterisation does not in any interesting sense make the results yielded by (Rule-AIC) dependent on the linguistic representation of the considered models.

[Kieseppä, 2001, p. 776]

From the epistemic perspective that is Kieseppä's concern, I can find room to agree that there is no "interesting sense" in which Rule-AIC is language dependent. This is because, if we are looking to employ Rule-AIC in statistical model selection, what is available to us is a procedure to check if the given parameterization is one that can support the justification of Rule-AIC. If the justification will work, then Rule-AIC applies, and if not, not. Rule-AIC isn't language dependent "in any interesting sense" insofar as it simply doesn't apply to the problematic languages/parameterizations that undermine its justification.

However, from the perspective of the BSA and the PIC, these failures of Rule-AIC *are* interesting. AIC (the measure) is not immanent, but it is also not transcendent; it is merely non-immanent. *Some* reparameterizations of considered models will lead to the inapplicability of Rule-AIC. If Rule-AIC was how we were deciding which system was best, the existence of these problematic reparameterizations would be, as Cohen and Callender put it, a *prima facie* devastating blow to the BSA.

Towards the end of their introducing the PIC, Cohen and Callender write that

What is needed to solve the problem is a *transcendent* simplicity/strength/balance comparison of each axiomatization against others. The problem is not that there are too many immanent measures and nothing to choose between them, but that there are too few (viz., no) transcendent measures.

[Cohen and Callender, 2009, p. 8, emphasis in original]

Cohen and Callender are probably right that there are “too few (viz., no) transcendent measures”. In response to this, PIC says that measuring the goodness of a system must be done with immanent measures, and so no systems expressed in different languages may be compared in the best system competition. But non-transcendence is not a guarantee of immanence. We might call the problem that remains the Problem of Transcendent Comparisons (PTC). It is the PTC that gets at the well known problems of language sensitivity found outside of the literature on the BSA. Measures like AIC are not immanent, but they are problematically language sensitive, and that may be understood as a symptom of their non-transcendence.

I do not deny that language relativity of the BBSA is good a strategy for dealing with the PIC. But we have at our disposal many non-immanent (and simultaneously non-transcendent) measures. In the face of the non-transcendence of these measures—that is, in the face of the PTC—the BBSA’s strategy of language relativity is still a good one. Our language relativity does not, however, have to involve privileging *single* languages. The alternative is to relativize to *classes* of languages constructed to ensure the applicability of the measures employed in our

best system competition.

4.7 Language-Class Relativity

We have just concluded that the PIC should be replaced with the PTC, which will not require us to relativize to single languages because of the immanence of measures in the best system competition. Rather, we can relativize to *classes* of languages constructed to include only languages that may be unproblematically compared using the measures that appear in the best system competition. At the start of this chapter I said that the goal was to rid the BSA of language relativity entirely. That goal is achieved here in three steps. First, in Section 4.7.1, I consider the role of single language privileging in accommodating special science laws in the BSA. There I revisit the discussion from Chapter 3 about the relationship between kind relativity (which we have been thinking of in this chapter as single language relativity) and fact relativity, and argue that fact relativity alone is enough for the accommodation of special science laws. Next, in Section 4.7.2, I develop what I take to be the greatest problem with single language privileging—that it fails to make sense of the idea that laws and kinds are discovered together—and argue that it may be solved by replacing single language privileging with language-class relativity. Finally, in Section 4.7.3, I propose a way of constructing classes of languages that will make language-class relativity depend on—and thus be redundant with—any competition relativity that is included in the BSA.

4.7.1 Special Science Laws

One of the most compelling features of the BBSA is its ability to accommodate special science laws in the BSA. This is accomplished through the view's single language privileging/relativity. A particular scientific field X is interested in the kinds K_X , so the laws of X come from the best system to emerge from a competition run with every system expressed in the language that treats the K_X kinds as basic. It does not seem plausible that scientific fields are interested in particular language classes,¹³ and so the value of replacing single language relativity for language-class relativity will depend in part on being able to provide an alternative account of accommodating special science laws in the BSA.

Most of what is needed to do this has already been presented in the Chapter 3. It was noted there that there is a very close relationship between what set of kinds is treated as basic and what facts are to be systematized. In one example the fields of "Coffee Chemistry" and "Pencil Chemistry" were considered and contrasted. If coffee chemists are interested in the kind 'caffeine', then they are also interested in all the caffeine related matters of fact. Coffee chemists will also be interested, at least implicitly, in all the facts that subvene the caffeine facts, such as facts relating to the various atoms that comprise caffeine molecules. Pencil chemists, in contrast, are interested in the 'graphite' kind, graphite related matters of fact, and (at least implicitly) all the facts that subvene graphite related matters of fact. Coffee and

¹³ At least not *explicitly*. Economists are explicitly interested in the set of kinds that contains things like "market" and "currency". It is not yet clear how we could characterize the class of languages that is of particular interest to economists. It is even less clear that doing so would be sufficient to distinguish economics (and its laws) from any other field.

pencil chemists are going to both be interested in the kind ‘carbon’ and carbon related matters facts. But they are not interested in *all of the same* carbon related matters of fact. Coffee chemists will only be interested in the subset of carbon facts related to carbon atoms appearing in caffeine molecules, while pencil chemists will only be interested in the subset of carbon facts related to carbon atoms appearing in graphite molecules. Fact relativity has a finer grain than kind relativity—this is precisely what made its introduction so useful.

When we give up single language relativity, fact relativity can become *the way* (and not merely additional insurance) to individuate fields. Being interested in a particular set of kinds K is the same, in practice, as being interested in the set of facts related to those kinds. We already are committed to including fact relativity in the BSA. And any two fields that may be individuated by their kinds of interest may also be individuated by the facts that they intend to systematize (i.e., the facts that correspond to the individuating kinds). Replacing single language relativity with language-class relativity will not, then, spoil our ability to accommodate special science laws in the BSA. The burden of individuating fields and their respective laws will simply fall to the fact relativity that was argued for in Chapter 3.

4.7.2 Discovering Laws and Kinds Together

One concern that we might have with letting fact relativity do the heavy lifting of individuating scientific fields is this: In conjunction with language-class relativity, we lose the guarantee offered by single language relativity that the kinds treated as

basic in the laws of a field will be the same as the kinds that are explicitly of interest to a field. The guarantee is certainly appealing. *Of course* there should be a match between the kinds of interest in a field and the kinds in the field's laws. But the *way* in which such a match is guaranteed by single language relativity fails to capture the interplay in scientific practice between the discovery of the basic kinds and the discovery of the laws.

This problem has already been discussed with respect to single language privileging (which is essential to single language relativity). Lewis claims that his variant of single language privileging has the virtue of “explaining” why “laws and natural properties get discovered together” [Lewis, 1983, p. 368].

For Loewer's PDA, the idea that laws and kinds are discovered together is central to the view. Indeed, the phrase “package deal” has its roots in Lewis, who says just before the “discovered together” remark that “the scientific investigation of laws and of natural properties is a package deal” [Lewis, 1983, p. 368]. While Loewer ultimately endorses a version of single language privileging, it is accompanied with a rough account of how a “final theory”—i.e., a candidate system-language pair—is arrived at:

a final theory is evaluated with respect to, among the other virtues, the extent to which it is informative and explanatory about truths of scientific interest as formulated in [the present language of science] SL or any language $SL+$ that may succeed SL in the rational development of the sciences. By ‘rational development’ I mean developments that are

considered within the scientific community to increase the simplicity, coherence, informativeness, explanatoriness, and other scientific virtues of a theory.

[Loewer, 2007, p. 325]

If the practice of science parallels the PDA, then the processes of discovering the laws and basic kinds are the one and the same.

And it seems Cohen and Callender are also on board with laws and kinds being discovered together when they offer this nice example of the phenomenon:

historical disputes between theorists favoring very different choices of kinds seem to us to be disputes between two different sets of laws [...] it has happened in the history of science that people have objected to particular carvings—most famously, consider the outrage inspired by Newton’s category of gravity. But given the link between laws and kinds, this outrage is probably best seen as an expression of the view that another System is Best, one without the offending category. If that other system doesn’t in fact fare so well in the best system competition—as in the case of the systems proposed by Newton’s foes—then the predictive strength and explanatory power of a putative Best System typically will win people over to the categorization employed. While it’s true that some choices of [kinds] may strike us as odd, no one would accuse science—the enterprise that gives us entropy, dark energy, and charm—as conforming to pre-theoretic intuitions about the natural kinds of the world. Yet these

odd kinds are all embedded in systematizations that would produce what we would consider laws.

[Cohen and Callender, 2009, pp. 17–18]

With everyone in agreement, what is the problem? Language privileging, essentially, happens *before* the identification (in the BSA and its variants) or discovery (in scientific practice) of the laws. As van Fraassen’s problems made clear for Lewis’ BSA, the natural properties would have to be discovered prior to the discovery of the laws in order to ensure the accessibility of the laws. Though Cohen and Callender will not “accuse science” of “conforming to pre-theoretic intuitions about the natural kinds of the world”, that is exactly what the BBSA does when it privileges sets of kinds prior to a best system competition. Furthermore, PIC makes it such that “the predictive strength and explanatory power of a putative Best System” cannot “win people over to the categorization employed” because comparing two putative Best Systems expressed in different languages (with different “categorizations”) is supposed to be impossible.¹⁴

Relativizing to classes of languages (and embracing PTC) solves this problem. Scientists are able to approach the discovery of laws and kinds with pre-theoretic intuitions about how to systematize the world, the language to use when doing that, and the best system competition. As we will see below, the intuitions regarding language and the best system competition will locate them in a particular language class. Scientists will move away from their intuitions about language (and

¹⁴ At least, it is impossible according to PIC for the BSA and its variants. If it *is* possible for scientists, then it is wholly unclear why it would be impossible for the BSA.

systematizing) when, much as Loewer describes above, there are languages in the relevant language class that may be paired with systems to yield a system-language pair that is scored better by the best system competition than the pre-theoretic system-language pair. Even without single language privileging, there will still be a guaranteed match between the kinds that are of interest to a field and the kinds that appear in the laws: The movement away from pre-theoretic intuitions about the kinds is constrained by the relevant language class and only pursued when it is accompanied by gains in the goodness of the best candidate system-language pair.¹⁵ Thus, by a field’s own lights, the new sets of kinds will be the kinds of interest to the field.

4.7.3 Limiting Language(-Class) Relativity

Let us begin addressing how language-class relativity can work by looking again at the single language relativity of the BBSA. In the BBSA, there are the fundamental kinds K_{fund} . The set of all kinds \mathcal{K} is the set including K_{fund} closed with respect to supervenience relations—that is, \mathcal{K} includes every kind that can be defined as supervening on the arrangement of the K_{fund} kinds in the actual world.¹⁶

A language L is determined by the set of kinds for which it has basic predicates,

¹⁵ This movement is only metaphorical for the BSA, where all the possibilities are considered and judged simultaneously. It is helpful, though, to think in the more methodical terms—of considering particular transitions from one system-language pair to another, the benefits that they might bring, and then adopting them or not—because that is what will happen in actual scientific practice.

¹⁶ It would be more proper to say of \mathcal{K} ’s closure that it includes every kind that can be defined as supervening on the kinds in \mathcal{K} (and their arrangement in the world), which of course includes the kinds of K_{fund} . Since supervenience is transitive—that is, if B supervenes on A , and C supervenes on B , then C supervenes on A —we can say that every kind in \mathcal{K} supervenes on K_{fund} (and the world). And it is convenient to say that because it highlights the supervenience base that is shared all kinds .

and there is a language L_i for every $K_i \subseteq \mathcal{K}$. For any two languages L and L' , the supervenience relations between the kinds of the languages and K_{fund} can be thought of as schemes for *translation* between L and L' . The set of all languages \mathcal{L}_{all} can be thought of as the set of languages that includes L_{fund} closed with respect to all translations. A class of languages \mathcal{L}_i is a set of languages including L_{fund} closed with respect to some acceptable (all, in the case of \mathcal{L}_{all}) translations.

To illustrate, let us return to the ‘coin flip’ worlds from earlier in the chapter. Such a world is a (finite) string of Hs and Ts, which we will assume are the only two fundamental kinds. Another set of kinds might be $K_{\text{ex}} = \{\mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{d}\}$, where the translation that gets us to the corresponding language L_{ex} from L_{fund} maps the pairs HH, HT, TH, and TT, to \mathbf{a} through \mathbf{d} , respectively. An example of a class of languages that includes L_{ex} could be $\mathcal{L}_{\text{n-tuple}}$: Let an acceptable translation for $\mathcal{L}_{\text{n-tuple}}$ be one that, for a given n takes the set of all n -tuples of H and T, and maps them to a set of kinds $K_n = \{k_{n,1}, k_{n,2}, \dots, k_{n,2^n}\}$. L_{fund} , then, is just L_1 . When \mathbf{a} through \mathbf{d} are $k_{2,1}$ through $k_{2,4}$, our K_{ex} and L_{ex} are precisely K_2 and L_2 . All, and only, the languages that may be formed through this procedure will be members of the class $\mathcal{L}_{\text{n-tuple}}$.

A language-class relative variant of the BSA will run a best system competition for every class of languages \mathcal{L}_i . Then \mathcal{S} is the set of all systematizations of the world, the set of all competing system-language pairs for the \mathcal{L}_i -relative best system competition is given by $\mathcal{S} \times \mathcal{L}_i$.

We can apply this conception of language-class relativity to our other running example of statistical model selection with AIC. Recall that *some* reparameteriza-

tions of statistical models would prove problematic for the use of AIC. To reparameterize a model is akin to translating it from one language to another. We can understand, then, the problem of language sensitivity for AIC as being related to some set of problematic translations. If we subtract these problematic translations from the set of all translations, then we have a set of acceptable translations which defines a class of languages that we can call \mathcal{L}_{AIC} . \mathcal{L}_{AIC} is precisely the set of all languages such that a systems expressed in any one of them will be comparable to a system expressed in any other using AIC. As long as the non-immanent measures used in the best system competition have clearly problematic and/or acceptable translations associated with them, then the class of languages that may be used to express competing systems will be determined by the measures used in the best system competition.

This will have one of two effects on the extent to which the BSA must be relativized to classes of languages: Either the BSA will be committed to competition relativity or not. Suppose that it is not. For convenience, suppose further that Rule-AIC is all that there is to the best system competition. In that case, the BSA will always be run using the \mathcal{L}_{AIC} class of languages. Language-class relativity is not required since there is only one language class that will ever be relevant to the BSA—namely \mathcal{L}_{AIC} , as determined by the best system competition. Now suppose that there is competition relativity. A different best system competition must be run for every competition function C_i in the set of all possible competition functions \mathcal{C} . In principle we will need to run best systems competitions for every pair in $\mathcal{C} \times \mathbb{L}$, where \mathbb{L} is the set of all language classes. Let \mathcal{L}_j be the class of languages

constructed according to the translations that are acceptable for the measures that comprise C_i when $i = j$. In practice, however, it will only make sense to run a competition once for each $C_i \in \mathcal{C}$, since the pairs C_i, \mathcal{L}_j will be unproblematic only when $i = j$. Language-class relativity in this situation will be redundant with competition relativity.

We also have it that, in either case, single language relativity remains unnecessary for all the same reasons that recommended language-class relativity. This means that there is no apparent need for *any* language relativity in the BSA. Its role in the accommodation special science laws will be redundant with fact relativity. Its role in solving the TSP will be redundant with more careful choices of the measures used in the best system competition. And, finally, its role in solving the PIC will be unnecessary (if a single non-immanent best system competition can be identified) or redundant with competition relativity.

4.8 Summary

Two problems have made it standard among contemporary defenders of the BSA and its variants to adopt single language privileging. The trivial systems problem (TSP), introduced in Section 4.1, is concerned with the existence of system-language pairs that are utterly strong and simple, guaranteeing their being the best, but whose axioms and theorems include all true universal generalizations, making them undeserving of the name “law”. The problem of immanent comparisons (PIC), introduced in Section 4.6, is concerned with the inexistence of transcendental mea-

asures of simplicity and strength. Both of these problems seem to be solved by privileging a single language as the one in which all systems competing to be the best are expressed. The TSP is solved as long as the privileged language does not include the predicates that give rise to the problem. The PIC is solved because the measures used in such a competition need only be defined immanently for the privileged language.

While single language privileging might solve these problems, it brings with it new problems depending on how it is implemented. The version of single language privileging adopted by [Lewis, 1983] (discussed in Section 4.2), in which the privileged language is that of the “perfectly natural properties”, made the laws potentially inaccessible and uninteresting to scientists. The PDA (discussed in Section 4.3) privileged the language of actual practicing scientists, and in doing so it appears to have made the laws dependent on subjects. The BBSA (discussed in Section 4.4) lets there be competitions run with privileged languages for every set of kinds that might be treated as basic, which means that there will still be some sets of laws that are still subject to the TSP. Lewis’ BSA and the BBSA also suffer from breaking the intimate connection in scientific practice between the discovery of the laws and the discovery of the basic kinds of nature.

One of the main things that I argued for in this chapter was that single language privileging is not required to solve the trivial systems and immanent comparisons problems. In Section 4.5, I show how the TSP can be solved by employing a suitable measure of strength. After its introduction in 4.6, I argued that the PIC is mistaken in assuming that, in the absence of transcendent measures for

the best system competition, only immanent measures may be used. There are, I argued and illustrated with the Akaike Information Criterion, non-immanent and non-transcendent measures. PIC, then, should be replaced with the slightly weaker problem of transcendent measures (PTC), which does not require single language privileging, but only privileging to classes of languages. Finally, in Section 4.7, I developed the idea of language-class relativity. There I showed how its adoption over single language privileging does not undermine the ability of the BSA to accommodate special science laws, and how it can respect the idea that laws and kinds are discovered together. With all of that done, I was able to argue that language relativity will, in general, be unnecessary for the BSA, because what it has so far accomplished will be addressed through a combination of fact and competition relativity.

Chapter 5: Induction Friendliness

What does it mean in the BSA for a system(-language pair) to be the titular *best*? The standard line—that the best system is the simplest and strongest on balance—is not without its issues. As Armstrong writes:

even granted that [our standards of simplicity] are shared by all rational mankind, [they] may not be shared by other rational creatures. The same point seems to hold for standards of strength. Lewis also refers to ‘our way of balancing’ simplicity and strength. May there not be irresolvable conflicts about the exact point of balance?

[[Armstrong, 1983](#), p. 67]

Not only might we disagree about how simplicity and strength are actually measured, we may disagree about their relative importance. Part of my aim in this chapter is to step back from simplicity and strength to ask, more or less from scratch, what makes a system-language pair the best. A result of stepping back in this way is to focus on the role of induction in scientific practice, and to argue for including in the best system competition a notion of “induction friendliness”

In our four-part model of the BSA there is the world, the systems, and the languages, and all of those are fed into the competition function, which outputs the

best system-language pair. So what should our competition function look like? The only real guide we have for saying something about the competition function is the practice of science itself (that is, after all, where simplicity and strength are supposed to come from). One strategy for dealing with the potential variety highlighted by Armstrong might be to add competition relativity into our version of the BSA by allowing for any competition function that is or might be used in scientific practice and letting there be distinct sets of laws for every choice of competition function. This is certainly what it looked like we might have to do at the end of Chapter 2. However, in the interest of objectivity (and reducing relativity), we should look closer at scientific practice to see if there are any universal features that might constrain competition relativity.

One candidate for being a universal feature of scientific practice is induction. Armstrong's "other rational creatures" might have tremendous brains or computing power at their disposal, and so be unconcerned with simplicity. But what if they were not in the business of making inductive inferences? What if they were unconcerned with generalizing from their limited data to regularities that are supposed to apply to the rest of the world? In such a case, we would say they are not doing science. At least, they are not doing any kind of science that is relevant to the BSA. Picking out regularities—namely, lawful regularities—is the whole point of the BSA. So, perhaps, induction has a role to play in constraining the relativity that might be required of our competition function(s).

What I argue in this chapter is that the competition functions of the BSA are concerned with "induction friendliness", which is, roughly, the extent to which the

world and competing system-language pairs are conducive to induction. I further argue that the induction friendliness of a system-language pair and world may be quantified by the mutual information (relative to the system-language pair) that is available between different parts of the world. Developing the relationships between “induction friendliness”, the best system competition of the BSA, and mutual information (MI), may be divided into three parts:¹ First, in Section 5.1, I argue that induction friendliness must play a role in the BSA. Second, in Section 5.2, I propose that induction friendliness may be quantified by mutual information. Third, Section 5.3, the various connections between mutual information and the BSA are sketched to highlight the potential benefits of adding induction friendliness as quantified by mutual information to the BSA. I conclude in Section 5.4 with a brief summary of the chapter.

5.1 Induction Friendliness and the BSA

Before arguing for a connection between the BSA and induction friendliness, it will be helpful to say something about what is meant by “induction friendliness”. (The BSA, at this point, needs no introduction.)

Part of what is meant by induction friendliness is best understood by saying what would make things induction *unfriendly*. Ready examples of induction unfriendliness may be found in the circumstances that give rise to the standard problems of induction. Hume’s classic worry highlights the potential for induction

¹ To explore the connection between these three further, I reconstruct in Appendix A a classic theorem of information theory from [Shannon, 1948] with the requisite assumptions being motivated as suiting a search in a toy model of the BSA for a measure of induction friendliness.

unfriendliness.

Thus not only our reason fails us in the discovery of the ultimate connexion of causes and effects, but even after experience has inform'd us of their constant conjunction, 'tis impossible for us to satisfy ourselves by our reason, why we shou'd extend that experience beyond those particular instances, which have fallen under our observation. We suppose, but are never able to prove, that there must be a resemblance betwixt those objects, of which we have had experience, and those which lie beyond the reach of our discovery.

[[Hume, 1896](#), Book 1, Part 3, Section 6]

The world around us is not guaranteed to respect the results of our inductive practices. Our circumstances are induction unfriendly if we have been led to suppose that there is some regularity in the world, but then that regularity fails.

Induction friendliness need not be concerned just with failures of our inductive practices (or a lack thereof). Successful inductive inferences may be made more or less quickly, and circumstances will be comparably more or less induction friendly. Suppose that there is a factory that produces biased coins for philosophers and statisticians who want their thought experiments to be more grounded in reality. Half of the factory's operating days are spent making coins that are balanced to come up heads 70% of the time and tails 30% of the time. The other half are spent making coins where the heads/tails split is 30%/70%. Suppose further that we have a coin that we know to be from this factory, but we do not know that it is a heads

biased coin. We start flipping the coin in order to test what kind of bias it has. Each time the coin comes up heads, we raise our degree of belief in it being a heads biased coin. Each time it comes up tails, we raise our degree of belief in a tails bias. A long run of heads at the start of our test flips would be very induction friendly because that would quickly lead us to adopt the true belief that the coin is heads biased. Our circumstances will be less induction friendly the more that tails come up during our test flips, since each tail detracts from our (right) confidence that the coin is heads biased. It is entirely possible (albeit not terribly likely), that our first 100 test flips of the coin might have tails come up 70 times and heads 30. In such an event, we would have good reason to believe that they coin is tails biased, but since that is not the truth, we, unknowingly, are in circumstances with low induction friendliness.

Like with this coin testing example, the progress of science by any measure—its rate, its accuracy, its practical implications, etc.—will be constrained (or bolstered) by the induction friendliness of our circumstances. Take the following remark from a physicist at the CERN particle accelerator:

I would consider us lucky if we discovered new phenomena or a new state of matter in two or three years [...] It would mean nature has been kind to us, but nature might be more subtle.

(Dr. Tiziano Camporesi, in [[Overbye, 2016](#)])

The kindness or subtlety of nature is precisely the sort of thing that I intend for induction friendliness to track. If “nature has been kind” to our scientific—read:

inductive—practices, then our circumstances have been induction friendly. If nature is “more subtle”, then our circumstances are less induction friendly.

A concern with induction friendliness is ever present in science, but not much can be done about it. However kind or friendly nature is to our inductive practices is out of our hands. But scientist can *hope* that their circumstances are induction friendly. More than that, absent countervailing evidence, scientist do—and, in an important sense, *must*, if they think their work is going to accomplish anything beyond merely describing their data—assume their circumstances are broadly induction friendly.

Given the importance of induction friendliness to science, it might seem obvious that it will have a role to play in the best system competition of the BSA. The standard argument for including some measure of system goodness in the BSA is something like the following argument from interest.

1. Scientists are interested in laws that are X.
2. An account of laws must respect scientific practice. (RSP)
3. (From 1 and 2:) The BSA should produce laws that are X.

We can see this argument at work in the background of Lewis’ first presentation of the BSA.

Of these true deductive systems, some can be axiomatized more *simply* than others. Also, some of them have more *strength*, or *information content*, than others. [...] What we value in a deductive system is a properly balanced combination of simplicity and strength.

[Lewis, 1973, p. 73, emphasis in original]

It is just a fact that about systems that they can be simpler or stronger. But there are also lots of other features of systems that we could track. The reason that we think simplicity and strength are the relevant ones for the BSA is that they are the ones that are of value or interest to science.²

So, since science is interested in induction friendliness, do we have it—by RSP and the argument from interest—that the BSA should produce laws that are induction friendly? It’s not quite that simple. Scientists are interested in laws (or system-language pairs) that are simple and strong. But their interest in induction friendliness is not directed at laws and system language pairs. From the perspective of scientists, it is the world itself that is more or less induction friendly. Once their best inductive practices have made an estimate of what the laws are, they are interested in the world being induction friendly. As the term has been developed here, to say that laws or some system-language pair is induction friendly is a category mistake. Worlds are more or less induction friendly with respect to a given system-language pair, and we may only say that a system-language pair is induction friendly insofar as it makes some pre-defined world (e.g., the actual world) induction friendly. Induction friendliness, it seems, is something that matters only *after* the events—namely, the estimation/discovery of the laws—that fall under the BSA’s purview.

This puts an effort to combine induction friendliness and the BSA in an awkward

² It is important to note that the argument from interest, like the argument given later in this section, is only an argument for why we *should think* that X is relevant to the BSA. The BSA has what ever best system competition(s) it has independent of the interests of any scientists—this is the “brute metaphysics” point from Chapter 2. The role of scientific interests is to guide our best estimates of how the BSA works (not to dictate how the BSA works directly).

position.

The trick to avoiding that awkwardness is to appreciate that induction friendliness is not a property of either the world alone or a given system-language pair, but of those things combined—the “circumstances” that have been mentioned already in this chapter. If S, L is the system language pair believed to be the best by science, then scientists will prefer that the actual world be the one among the set of possible worlds \mathcal{W} that maximizes the measure of induction friendliness, $IF(S, L, w)$.

Now adopt the perspective of the BSA. The BSA’s competition function \mathcal{C} takes as arguments the actual world w_{actual} , the set \mathcal{S} of all competing systems, and the set \mathcal{L} of all competing languages, and outputs a best system-language pair, $S_{\text{best}}, L_{\text{best}}$. Once $S_{\text{best}}, L_{\text{best}}$ is in hand, we can ask: Which possible world $w \in \mathcal{W}$ maximizes $IF(S_{\text{best}}, L_{\text{best}}, w)$? If the answer is not the actual world w_{actual} , then we have violated the interests/preferences of science as they pertain to induction friendliness (we will worry about considering interests in simplicity and strength later in this section).

To see that this is the case, recall the story of the scientist who dies and appears before God. When asked by God if the scientist would like anything, the scientist says “I would like to know how the world works”. God accepts the request and begins “At time and position t_1, p_1 such-and-such properties obtained, and at t_1, p_2 such-and-such other properties obtained, and—” at which point the scientist interrupts and says “Sorry, if that is how things *actually* work that is all well and good, but is there a *pithy* version?”. God’s reply to that question is to provide the scientist with the BSA-style laws of nature.

God is beholden to provide the scientist with the laws that best suit the scientist’s interests—not because God is good, but because God is a stand-in here for the BSA, and RSP binds the BSA to the interests of scientists. Prior to this chapter, we have characterized those interests in a variety of ways. We’ve noted that the scientist might have particular interests in what language(s) can be used in expressing the laws, what particular facts should contribute to determining the laws, and what kinds of systems should be considered. The scientist is interested explicitly in the laws being “pithy”—i.e., simple and strong—but that is something that we have set aside in this chapter in favor of looking at an interest in induction friendliness. The scientist’s interest in induction friendliness means that, when God recommends $S_{\text{best}}, L_{\text{best}}$ as the source of the laws, the scientist will be interested in the world being such that it maximizes $IF(S_{\text{best}}, L_{\text{best}}, w)$. Since God knows what the actual world is it must be that $IF(S_{\text{best}}, L_{\text{best}}, w_{\text{actual}})$ is maximal with respect to the set of possible worlds \mathcal{W} . Otherwise, God would be knowingly setting up the scientist for disappointment, and that goes against RSP.

We thus have the following constraint on the BSA’s best system competition:

Induction Friendliness (strong): The best system competition C must be such that,

if $C(\mathcal{S}, \mathcal{L}, w_{\text{actual}}) = S_{\text{best}}, L_{\text{best}}$, then the induction friendliness

$IF(S_{\text{best}}, L_{\text{best}}, w_{\text{actual}})$ is maximal with respect to the set of possible worlds

\mathcal{W} .

At least, that should be the case if we are ignoring (or letting induction friendliness supersede) other factors in the best system competition like the simplicity and

strength of a system-language pair.




Figure 5.1: Comparing worlds and system-language pairs: The dashed curve sitting on the w_{actual} horizontal represents the goodness of every system-language pair with respect to the actual world. The dotted curves set against the S_1, L_1 and S_2, L_2 verticals represent the induction friendliness of every world in \mathcal{W} on the assumption that the particular system-language pair is best.

Look at Figure 5.1 for an illustration of how induction friendliness might interact with other measures of the goodness of a system-language pair. The standard picture for the BSA is captured by just the dashed curve sitting on the horizontal labeled w_{actual} . There is the actual world w_{actual} , and every system-language pair in $\mathcal{S} \times \mathcal{L}$ is good—that is, strikes some balance between simplicity and strength—to some degree with respect to w_{actual} . S_2, L_2 has the highest goodness, and so it will be the best system. The induction friendliness of every world in \mathcal{W} with respect to S_2, L_2 is represented by the dotted curve set against the S_2, L_2 vertical. If scientists think that S_2, L_2 is the best system-language pair, if it is the output of their

best inductive practices, then scientists will want to be in w_1 since that is the most induction friendly world with respect to S_2, L_2 . In contrast, w_{actual} is the most induction friendly world with respect to S_1, L_1 (as may be seen from the dotted curve set against the S_1, L_1 vertical).

If we consider *just* induction friendliness, then S_1, L_1 will be preferable to S_2, L_2 since it makes the actual world the most induction friendly. But we will not, in general, be interested in just induction friendliness. S_1, L_1 is a very poor system with respect to w_{actual} . And w_{actual} isn't far from the peak of induction friendliness with respect to S_2, L_2 . It seems like we will prefer S_2, L_2 to S_1, L_1 because it strikes the better balance between the standard measures of system goodness and the induction friendliness of the actual world with respect to it.

Considering the simplicity and strength of a system-language pair may weaken the constraint above to be something like the following:

Induction Friendliness (weak): If two system language pairs S_1, L_1 and S_2, L_2 strike an equally good balance between simplicity and strength, then the best system competition should prefer S_1, L_1 over S_2, L_2 if $IF(S_1, L_1, w_{\text{actual}})$, the induction friendliness of S_1, L_1 , is closer to $\max_{w \in \mathcal{W}} IF(S_1, L_1, w)$ than $IF(S_2, L_2, w_{\text{actual}})$ is to $\max_{w \in \mathcal{W}} IF(S_2, L_2, w)$.

This is the weakest constraint that might be imposed on the competition function; it makes induction friendliness a mere tie-breaker after system-language pairs have been evaluated for their balance of simplicity and strength. Stronger constraints are possible. Simplicity and strength may be the mere tie-breakers after induction

friendliness is considered, as they would be under the strong induction friendliness constraint given at first. In general, the best system competition might look for a balance between simplicity, strength, and induction friendliness to various degrees.

5.2 Induction Friendliness and Mutual Information

One way that we might fill in the formal details of induction friendliness is with mutual information (MI). In this section, I show how that can be done. The case for associating the two will be bolstered more by a survey in the next section of some of the benefits that may be had from combining MI, via induction friendliness, with the BSA.

We have a good idea of what induction friendliness is after the work of the previous section. What about MI? Formally, MI is a measure $I(-)$ of two random variables X and Y given by

$$I(X;Y) = k \sum_{x \in X} \sum_{y \in Y} p(x,y) \log \frac{p(x,y)}{p(x)p(y)}$$

where k is an arbitrary constant that determines the units of the measure, $p(x)$ is the probability of some value x obtaining in X , $p(y)$ the probability of some value y obtaining in Y , and $p(x,y)$ the joint probability of x and y obtaining simultaneously in X and Y (respectively).

The measure has its roots in information theory. Using the most casual information theoretic terms, MI is the “amount of information” that is available on average about the values of X and Y together. When “amount of information” is

associated, as it is in the original development of information theory in [Shannon, 1948], with minimum lengths of codes for representing values of random variables, $I(X;Y)$ is a measure of the difference between (1) the average amount of space it takes to write down the value of X and (2) the average amount of space it takes to write down the value of X if the value of Y is already known. (Or, since MI is symmetric, it is the difference between the space required to write down Y 's value and the space required to write down Y 's value when X 's is known.) This “code length” interpretation of MI is, for our purposes at least, superior for a number of reasons. For one, it lets us avoid any complications associated with whether or not information theory provides a good analysis of “information”—a highly contentious issue among philosophers (see, for example, [Floridi, 2011] or, for an excellent overview of the use of “information” specifically in biology, [Godfrey-Smith and Sterelny, 2016]). The code length interpretation is of great practical importance because of its role in making information theory a theory of efficient data storage and exchange, and thus part of the theoretical foundation of modern computing and telecommunications.³

³ It is also a helpful way of understanding the relationship between MI and two other important information theoretic quantities: The *entropy* of a random variable X is given by

$$H(X) = -k \sum_{x \in X} p(x) \log p(x)$$

which may be interpreted following (1) above as the average amount of space it takes to write down the value of X . The entropy of X is the most basic measure in information theory. Its value is the minimum average length of codes for values of X that can be achieved when encoding the values of X in a language with a number of character types b when $k = 1/\log b$. If $b = 2$, then the units of H (and I) are *bits*, like in the zeros and ones that make up computer memory, with eight bits to a byte, and oh-so-many bytes—kilo-, mega-, giga-, tera-, etc.—in any digital storage you possess. The other measure is the *conditional entropy* between X and Y , given by

$$H(X|Y) = -k \sum_{x \in X} \sum_{y \in Y} p(x, y) \log p(x|y).$$

Fitting (2) above, the conditional entropy is a measure of the average amount of space it takes to write down the value of X if the value of Y is already known. It can be seen after some rearranging

Another helpful interpretation of MI takes it to be a measure of the correlation of X and Y . $I(X;Y)$ is the average over pairs of values of X and Y of $\log[p(x,y)/p(x)p(y)]$. If X and Y are wholly independent, then, for all $x, y \in X, Y$, $p(x, y) = p(x)p(y)$. Then $p(x, y)/p(x)p(y) = 1$, the logarithm of 1 is 0, and so $I(X;Y) = 0$. As the joint probability $p(x, y)$ goes up relative to $p(x)p(y)$, the value of $\log[p(x, y)/p(x)p(y)]$ gets a higher positive value, and so the pair x, y contributes ever more to the mutual information between X and Y . Less independent, more correlated, pairs x, y contribute more to $I(X;Y)$, and always in proportion to the joint probability of the pair. Thus, $I(X;Y)$ is a measure of the correlation of X and Y in that it is a measure of the average correlation between the possible values of X and Y .

An argument for the connection between induction friendliness and MI may be found in the interpretation of MI as a measure of correlation. Suppose that X and Y are the states of the world just slightly temporally offset, so that the present value of X represents the present state of the world, and the present value of Y represent the immediate future state of the world. If $I(X;Y)$ is high, that means that the present is highly correlated with the future on average. In such circumstances, our inductive practices should (on average) be more successful. If the present is more correlated with the future, then we can expect that what we know of the present will allow us to make better predictions about the future. Because MI is symmetric, knowledge of the future will also allow for more accurate predictions about the present; or, to put it in a way that is more relevant to the practice of science, knowledge of the

that $I(X;Y) = H(X) - H(X|Y)$.

present will allow for more accurate retrodictions about the the past.

This simple picture of MI being measured between temporally adjacent states of the world will not generally be how we want to measure induction friendliness. Our inductive practices do not consist entirely of using just the present to make predictions about the immediate future. Rather, extended regions of space and time, of various sizes and shapes, are used to make predictions about similarly complex regions. How we carve the world up into these regions, and how we expect them to relate and make predictions about each other, will be determined by the system-language pair with respect to which the induction friendliness of the world is being calculated. A given system-language pair will carve the world up into related parts, and the induction friendliness of the world with respect to that pair will be the MI measured between those parts.

The simple picture above can still be interesting as a toy example to explore, and I do just that in Appendix A. I illustrate there the connection between the BSA, induction friendliness, and MI, through a reconstruction of the derivation in [Shannon, 1948] of entropy as a measure of uncertainty. This is done in the context of a toy model of the BSA, with the assumptions of Shannon's theorem being newly motivated as a part of a search for a measure of induction unfriendliness that is concerned with predictive success. From that we are able to derive a measure of induction friendliness that is approximately the MI between different parts of the world.

Importantly, MI does not tell us *how* to make predictions. Rather, whatever predictions we might be led to make based on our inductive practices, *our predictive*

accuracy faces an upper bound measurable by MI. It is in this way that we can see MI as a measure of induction friendliness: The average success of our best possible inductive practices will be greater or lesser in concert with higher and lower values of the MI between different parts of the world. Speaking loosely in terms of “information”, a world with higher MI with respect to some system-language pair will be one whose various parts carry more information about each other on average.

5.3 Mutual Information and the BSA

It has just been argued that MI is *a* way of dealing with induction friendliness. My interest in this section is to sketch how that inclusion may benefit the BSA in a number of ways that do not have direct connections to induction friendliness. In each case, MI is shown to be *a* way of dealing with the relevant idea. We could choose to deal with each one, induction friendliness included, with the use of other measures. But I think MI’s applicability for each should give us pause. Maybe some unificationist leanings should make us think that in MI we are getting at *the* measure to use in the BSA. Absent such leanings, it must be left to future work to make a decisive argument in favor of MI’s role in the BSA.

So what are the various ways (independent of induction friendliness) for MI to play a role in the BSA? In Section [5.3.1](#), I discuss the connections between MI and the virtues of simplicity and strength that are standardly associated with the BSA. In Section [5.3.2](#), I highlight extant discussion in philosophy that imply a connection between MI and the BSA would improve the explanatory power of the laws yielded

by the BSA. In Section 5.3.3, I draw on the code length interpretation of MI and information theory to argue that MI reflects some of the “insurmountable practical challenges” mentioned briefly in Chapter 2. Finally, in Section 5.3.4, I show how the incredibly strong and simple systems-language pairs of the trivial systems problem discussed in the preceding chapter score minimally on MI, a fact which will help block the trivial systems problem.

5.3.1 Simplicity and Strength

This chapter began with the idea that we should set aside our thinking about simplicity and strength in favor of starting from scratch in a search for measures to be used in the best system competition. We have done that, and identified induction friendliness, plausibly quantified by MI, as having a role to play in the BSA. But simplicity and strength haven’t just been abandoned. We have every reason to think that they still have a role to play in the BSA alongside induction friendliness, and that means we should still be asking “What does it mean to be ‘simple’ and ‘strong’?”. MI may play a role in addressing that question.

MI’s relationship with predictive accuracy is one way that we might connect it with simplicity. Predictive accuracy is the ultimate justification given in [Forster and Sober, 1994] for the Akaike Information Criterion (AIC) method of statistical model selection and, in the context of that method, it is also the basis for defending the pursuit of simplicity and unification in science. [Kieseppä, 2001] notes that predictive accuracy in the AIC formally corresponds to the Kullback-Leibler divergence

between two probability distributions. And Kullback-Leibler divergence is equal to MI when it is measuring the divergence between the distributions given by $p(x, y)$ and $p(x)p(y)$ for the random variables X and Y that are the concern of MI.

MI has further connection to simplicity by way of the two broad classes of simplicity—*simplicity in formulation* and *simplicity in derivation*—that [Frigg and Hofer, 2010] distinguish. With respect to simplicity in formulation, they write that

a linear relation between two variables is simpler than a polynomial of order 325, a homogenous first order differential equation is simpler than a non-linear integro-differential equation, etc. It is not easy to pin down what general rule drives these judgments, but this does not represent a serious obstacle to us because nothing in what follows depends on simplicity judgments of this kind. Another component of simplicity in formulation is how many distinct probability rules a system contains. Ceteris paribus, the fewer rules a system has in it, the simpler it is.

[Frigg and Hofer, 2010, p. 359]

One general rule that might drive these judgments is that all the given examples of something being simpler involve something that standardly takes less space to write down. In this way, simplicity in formulation may be related to MI through the code length interpretation of information theoretic measures.

[Frigg and Hofer, 2010] describe simplicity in derivation as having to do with “the computational costs incurred in deriving a desired result”, and recommend quantifying this with Kolmogorov complexity [Frigg and Hofer, 2010, p. 359]. Im-

portantly, Kolmogorov complexity is not a computable function. This means that, if it were to play a role in the BSA, the laws of nature may turn out to be epistemically inaccessible just because the simplicity of competing systems, as measured by Kolmogorov complexity, cannot be calculated. Epistemic accessibility is typically a desideratum for an analysis of laws, so this incalculability is not something that we want in the BSA. There is, however, a definite class of situations where Kolmogorov complexity is guaranteed to be computable, and, in all of them, measuring Kolmogorov complexity is equivalent to measuring entropy as in [Shannon, 1948]. Thus, MI, through its connections to related information theoretic measures, has connections to simplicity in derivation as well as simplicity in formulation.

MI might also have close ties to strength. Many authors— [Lewis, 1973], [Loewer, 2007], [Cohen and Callender, 2009], and [Woodward, 2013], among others— refer to strength as “informativeness”. Though I have so far preferred the code length interpretation of information theoretic measures over the “information” interpretation, it cannot be denied that information theory is the most prominent *quantification* (as opposed to qualitative analysis) of “information”. Thus, if “informativeness” is how we are to understand strength, MI is well placed to be a quantification of strength. Circumstances with higher MI are ones where the world is more informative about itself (as mediated by the system-language pair of the given circumstances).

5.3.2 Explanatory Power

Explanatory power has been related to MI in the work of [Hanna, 1969] and [Greeno, 1970]. I will focus in the discussion below on [Greeno, 1970], which has garnered greater attention as part of the introduction of the Statistical-Relevance model of explanation in [Salmon, 1971]. Specifically, I will show here how the role for MI in Greeno’s analysis of explanatory power fits with the way MI has been deployed above on behalf of induction friendliness as being measured between different parts of the world.

To begin, Greeno is concerned with the relevance of possible explanantia that are statistically related to some explanandum. The motivating example involves trying to explain why Albert stole a car. Part of the explanation might be that Albert lives in San Francisco, where there is a high delinquency rate. But Albert’s father earns more than \$40,000 per year, and the children of such high earners have a low delinquency rate. It seems clear that a good explanation of Albert’s theft should not exclude Albert’s father’s high earnings, but also that an explanation will be worse for its inclusion since it makes the explanandum less likely. Greeno’s solution to this paradox involves looking at “general explanatory systems, rather than single explanations” and, in the context of those systems, using MI to identify those potential explanantia that are most relevant [Greeno, 1970, p. 280].

Greeno’s application of information theoretic ideas begins by carving the domain of a theory into two partitions M and E^4 , with M meant to be interpreted as

⁴ Greeno uses S instead of E . I introduce the substitution to avoid using S in a way other than as standing for a system. For clarity and consistency, I will make other minor adjustments

the set of variables whose values are to be explained, and E the set of variables that may contribute to any such explanations. MI measured between M and E can be thought of as the difference between the entropy of M and the conditional entropy of M given E . This gives us (when the arbitrary constant k is set to 1):

$$\begin{aligned}
 I(E; M) &= H(M) - H(M|E) \\
 &= - \sum_{m \in M} p(m) \log p(m) \\
 &\quad + \sum_{m \in M} \sum_{e \in E} p(e) p(m|e) \log p(m|e)
 \end{aligned}$$

Greeno notes that $I(E; M)$ has the nice feature of giving “an indication of the extent to which the theory is ‘close’ to the goal of nomological deductive explanatory power” [Greeno, 1970, p. 283]. The reasoning behind this claim is that concerns about statistical explanation will give way to the deterministic explanations of the deductive-nomological model in the special case where $p(m|e)$ is either 0 or 1 for all $m, e \in M, E$. $I(E; M)$ ranges in value from 0, when $H(M) = H(M|E)$, to $H(M)$, when $H(M|E) = 0$. The maximum value will be achieved in the deductive-nomological model’s special case since either $p(m|e) = 0$ or $\log p(m|e) = 0$, zeroing out every value in the sums of $H(M|E)$.

This is just one of a number ways Greeno relates MI to explanatory power. The paper goes on to show that MI approximates “the usefulness of the theory in improving the accuracy of predictions about the explananda”, that theories with higher MI will be more testable under certain conditions, and how gains (or lack

to notation in any quotations from Greeno.

thereof) in the MI of a theory when new variables are added to the theory can illuminate the development of new hypotheses in scientific practice [Greeno, 1970, p. 286]. I will not go into any of the details of these here.

What is important to see is that Greeno's partitions M and E are of a kind with the way I have been talking about MI being measured between parts of the world. Extending the example of Albert's auto theft to a general theory of delinquency, Greeno suggests that

the domain might be American males between 12 and 18 years old. The variables of the theory might be degrees of delinquency D , the kind of neighborhood a boy lives in B , and the income of the boy's family C .

[Greeno, 1970, p. 280]

These variables that can make up M and E are like the basic kinds or quantities picked out by the language of a system-language pair. How the statistics associated with these variables and the domain are like the relationship between a system and a world. Following these associations, Greeno's calculation of MI as a measure of explanatory power is carried out in the same way as mine for MI as a measure of induction friendliness. Insofar as we are right to use MI in these ways (or, rather, this way), then it may be said that circumstances—the triple w, S, L of a world, system, and language—will present with a more induction friendly world precisely when the world makes the given system-language pair more explanatorily powerful.

5.3.3 Compression

Back in Chapter 2, it was suggested that the relativity of the BSA may be limited by “insurmountable practical challenges”. Two such challenges were mentioned. The first was that the best system competition must be computable, on pain of leaving the laws incalculable—and this already came up earlier in this section in the context of using Kolmogorov or Shannon-style approaches to information theory in measuring simplicity. The second challenge had to do with writing things down in the practice of science, and how, every time we do, “we incur physical space and energy costs”.

To develop this second challenge more, think again of the story of a scientist appearing before God and receiving the laws of nature. It does a lot to illuminate the nature of BSA laws, but it gets an important feature of actual scientific practice wrong: *Science is done in the world*. The results of experiments must be written down in the same space that is (broadly speaking) the subject of the experiments. The systematic calculations of practicing scientists must be done in a way that comports with the very collection of facts that are to be systematized. As with a concern for induction friendliness, any limits to relativity imposed by the fact that science is done in the world are going to be universal—nobody has an eternity outside of space and time in the audience of a boundless intellect.

When, in our story, God starts listing all the facts, the scientist asks for the pithy version. But what if we set aside the desire for the pithy version of the facts? Suppose that scientists were content to collect every fact they can without a care

for simplicity (or distinguishing lawful from accidental regularities, or any other consideration that might speak against this list as the aim of science). We may ask naively: How much space would it take to write down all the facts? How big, in the actual world, is “God’s big book of facts”, as [Beebee, 2000] calls it?

Consider the question first with our current technology. Suppose we tried to write down just 1 byte worth of data for each of the estimated 10^{80} atoms in the known universe, and that we can store 1 TB (10^7 bytes) on a drive 5×10^{-5} m³ in volume and 100g in mass (typical for a commercially available external hard drive). To do that we would need a storage device whose size is comparable to a stack of 100 Milky Way galaxies (which in the grand scheme of things isn’t that big) and whose mass is 10^{14} times *greater the mass of the known universe* (a galaxy is mostly empty space, but this would be a very dense cube). And that impossible to achieve feat is just trying to get 1 byte’s worth of data on each atom at one point in time.

If we were to try and record every fact of the universe with perfect accuracy throughout all of space and time we would need, in essence, a perfect duplicate universe to serve as our big book of facts. Since the big book must be located in the actual universe, the best we could ever hope to do is get one half of the universe (the big book region) to look exactly like the other half (the to-be-written-down fact region). In getting the book region to match the fact region, we effectively prevent ourselves from writing down the facts of the book region of the universe. Still more of the universe would have to be abandoned for inclusion in the book to provide the energy required to arrange the book region of the universe to match the arrangement in the facts region (with each region now only being sized to half of

what is left in the world after extracting the needed energy for the rearrangement). Writing anything down on the scale of God's big book will incur enormous costs in the form of abandoning as never-to-be-written-down some fraction of all the facts in the world

But this isn't quite right. Suppose we do want to be able to record every fact of the universe with perfect accuracy throughout all of space and time. Doing so would require a perfect duplicate universe to serve as our big book of facts, but only if we deny ourselves any means of *compression*—in some sense, only if we are uninterested in the pithy version of the big book. But, of course, we *are* interested in the pithy version of the big book. The whole point of the BSA is to pick out the regularities of the universe, and in doing so it frees us up from having to write everything down. If we know, or have a law saying, for example, that *As* always precede *Bs*, then we don't need to write down every instance of an *A* in the world. It will be enough to write down the regularity and let it be implicit whenever we write down an instance of a *B* that it was preceded by an *A*.

Including MI in the BSA puts us on track to address all of this by way of the code length interpretation of information theory. Indeed, this problem of efficiently writing things down was precisely what led to the development of information theory in [Shannon, 1948]. Shannon was working for Bell Labs, still owned by AT&T at the time, and they were very much concerned with the practical problem of efficiently writing things down (in the form of electrical impulses propagating across telecommunications wires) as a part of the telephone and telegraph business that is the company's namesake. Preferring system-language pairs that have higher MI will

lead to sets of laws that allow for greater compression of God’s big book of facts, thus avoiding more of the physical costs associated with writing down every fact.

5.3.4 Triviality

The last reason that I will give for adding MI to the BSA is that it seems to offer a solution a longstanding problem for the BSA—our much discussed *trivial systems problem*.

To quickly review: [Lewis, 1983] notices that how good a system is depends on the language in which it is expressed. The dependency between the goodness of a system and the language in which it is expressed leads to the trivial systems problem for the BSA: Let F be a predicate true of all and only things in the actual world. Then the system “ $\forall xFx$ ” seems to be both incredibly strong and incredibly simple, making it the best system. But $\forall xFx$ would make every true regularity of the world a law (since the laws are the axioms and theorems of the best system), and that violates the desideratum that there be a distinction between lawful and accidental regularities. The standard solutions to this problem—offered in [Lewis, 1983], [Loewer, 2007], and [Cohen and Callender, 2009]—all involve blocking F in various ways from being a legitimate predicate in whatever language is paired with a system. In Chapter 4 it was argued that an alternative solution to the problem may be found by properly defining the best system competition, and specific attention was paid to how strength is measured.

Adding MI to the best system competition does similar work. Even if $\forall xFx$

is incredibly strong and simple, its MI is exactly zero. Consider our simple example of systems that treat the world as Markovian with degree one (though, in this case, the degree doesn't matter). The world w , described in the language of F , is just a long string of F s. On its face, one might think that such a world has high MI. But remember that MI is concerned with *how much less* has to be written down about the random variables of interest. This is most clear when MI, $I(X; Y)$, is measured as the difference between the entropy $H(X)$ and conditional entropy $H(X|Y)$ (see fn. 3 of this chapter). In the circumstances of the trivial system, there is zero conditional entropy; that is, if you know the present state of the world (Y), then nothing more needs to be written down in order for the immediate future state of the world (X) to be communicated to you. And the non-conditional entropy is *also* zero; it doesn't take any space to write down the immediate future state of the world because it is guaranteed to be F , and we don't need to waste space writing things down about outcomes that are guaranteed.

In both cases, this can be seen formally. For the entropy

$$H(X) = -k \sum_{x \in X} p(x) \log p(x)$$

when $x = F$, $p(F) = 1$ —since F always occurs—then $\log p(F) = 0$, and so 0 is contributed to the value of $H(X)$. And, since F always occurs, $H(X) = 0$. Similarly for the conditional entropy

$$H(X|Y) = -k \sum_{x \in X} \sum_{y \in Y} p(x, y) \log p(x|y).$$

Since F always occurs, the only relevant term is when $x, y = F, F$. And so it goes: $p(F_{(x)}|F_{(y)}) = 1$, $\log 1 = 0$, and so $H(X|Y) = 0$.

MI is the difference between these two measures, and the difference between zeros is zero. So, even if $\forall xFx$ is overwhelmingly simple and strong, it has minimal MI. The TSP depends on the trivial system being overwhelmingly the best, and that is unlikely to be the case when MI is a part of the best system competition. Notably, unlike the strength measure in Chapter 4, the success of MI in undermining the TSP does not depend on how we construct the space of possible worlds.

5.4 Summary

I have argued in this chapter that the role of induction in scientific practice should be taken very seriously in the BSA—specifically, *induction friendliness* should be a part of the measure of the goodness of a system-language pair. Furthermore, I recommended that induction friendliness be quantified by mutual information, and pointed to a number of additional benefits that that might have for the BSA.

I began, in Section 5.1, by arguing that the BSA must be concerned with induction friendliness. After saying a bit about what is meant by ‘induction friendliness’, I pointed out that the usual argument from interest that supports the inclusion of simplicity and strength in the BSA does not easily apply to induction friendliness. Despite that, science’s interest in induction friendliness does place a requirement on the BSA according to which the best system system-language pair must make

the the actual world the most induction friendly among all the possible worlds. In Section 5.2, I recommended that induction friendliness be quantified by mutual information. This was argued for on the basis of being able to construe mutual information as a measure of the correlation between past, present, and future states of the world. Finally, in Section 5.3, I argued that mutual information might bring a number additional of benefits to the BSA that are independent from its connection to induction friendliness. Specifically, I noted that mutual information has extant connections to measures of simplicity, strength, and explanatory power, that it will help to minimize the physical costs of doing science through its connection to data compression, and that it may have a role to play in avoiding the trivial systems problem.

Chapter 6: Conclusion

Let us return to the question that prompts all the work that has just been done: What is it to be a law of nature? The rough answer offered by the BSA is that the laws are the regularities that appear in the best systematization of the world. As we quickly saw, this rough answer raises many new questions. What counts as a system? What is being systematized? What makes a system best? The rough answer to these questions is to say that what counts as a system, what makes a system best, etc., will be determined by scientific practice. Laws of nature have a role to play in scientific practice, and a good account of laws will respect that.

The quickest way to respect scientific practice is to build it into your theory. For example, if scientists prefer simple and strong theories, then the BSA should say that the best systems are the simplest and strongest. But it better not be that *within* the BSA what makes it true that the best theory is the simplest and strongest is the preferences of scientists. If that were the case, then what it is to be a law of nature would depend on subjects (namely, scientists), in conflict with a widely held expectation that the laws of nature should be objective. This issue, dubbed ‘Armstrong’s objection’ and discussed at length in Chapter 2, admits (at first) of a rather straightforward solution. The ‘brute metaphysics’ response to Armstrong’s

objection says that there is just a metaphysical fact of the matter about what makes a system the best. Insofar as the laws have a role to play in scientific practice, we can and should look to that practice to make our best philosophical guess as to what “best” means. But that does not make scientific practice the determiner of the meaning of “best” in the BSA, and so the laws (as described by the BSA) are not in violation of the expectation that they be objective.

If that is the end of the story, then there is a lot of work to do for the BSA in the form of going through scientific practice and finding precise answers to the questions like “What makes a system best?”. But things are not quite that simple. As soon as we set out to search for answers to such questions in scientific practice, what we find are a multitude of incompatible answers. The sorts of systems considered in physics just aren’t the same as those considered in biology (at least not in general). Maybe physics and economics consider similar systems (insofar as they employ similar formal tools), but the basic kinds of those systems, and generally the languages use to express them, are wildly different. We likewise don’t find physics, biology, or economics to be interested in systematizing the same things. Nothing could happen in the realm of economics that would count as evidence for or against the Standard Model of particle physics. *Maybe* we could learn something about the inner workings of black holes that will be relevant to biology, but barring (say) some revelation about the prevalence of micro black holes in the primordial soup in which life on Earth started, biologists can safely ignore the latest advances in the physics of black holes.

So what does this do to our response to Armstrong’s objection? The practice

of science taken as a whole disagrees with itself on the answers to those questions which are important to the details of the BSA. RSP, the principle that an account of laws must respect scientific practice, tells against privileging one set of answers available in scientific practice over others. The alternative is to *relativize* the laws, letting there be different sets of laws for every way you could fill in the details of the BSA corresponding to every way scientific practice could answer the relevant questions. This could be done on a very large scale, with every detail of the BSA being allowed to vary without constraint. But, I argued in Section 2.2, Armstrong's objection has a good alternative reading as being opposed to the profligate sets of laws relativity and accommodating scientific diversity would bring. This anti-relativity reading of Armstrong's objection is compatible with the anti-subjectivity reading, and, I claimed, the objection is best understood as the combination of both readings. Once it was sorted out how there can be relativity in the laws without succumbing to subjectivity (in Section 2.2.1), responding to Armstrong's objection became a matter of identifying just what relativity is required to accommodate differences (actual and potential) in the interests of scientists. If you don't include enough relativity, then you will break (at least some of) the BSA's laws' connection to scientific practice in violation of RSP. If you include more relativity than what is needed to satisfy RSP, then you do harm to the expectation that the laws are objective in the sense of being absolute (as opposed to non-subjective).

Properly answering Armstrong's objection requires striking a balance between relativity and objectivity in the BSA. *That* has been (give or take some additional concerns) the project of this dissertation following the general framework of the BSA

offered in Chapter 1 and the setup of the challenge posed by Armstrong's objection in Chapter 2. The project of the remainder of this concluding chapter is to continue reviewing the results of this dissertation and spell out how they come together into a coherent variant of the BSA. In reference to the tension that is the overarching theme of this dissertation, I will call my BSA variant the *Relative and Objective on Balance Analysis*, or ROBA for short.

Before getting into the details of the ROBA, it will be helpful to review all the moving parts of the BSA whose details need to be filled in. In Chapter 1, the four-part model of the BSA was introduced. The four parts are:

1. The facts of the world to be systematized.
2. The candidate systems from which the laws will be drawn.
3. The languages in which the facts and systems may be expressed.
4. The competition function that determines the best system language pair.

A *relativized* variant of the BSA will have it that there are multiple sets of laws each relative to different ways of filling in the details of these four parts. Such a relativized BSA variant will be more or less relative depending on it having fewer or more constraints on the ways to fill in the details.

Much of what has happened in this dissertation played on relationships between these parts. In Chapter 3, the relationship between facts and languages was explored in some detail. The languages of the BSA, following [Cohen and Callender, 2009], are individuated by the kinds that they treat as basic, with the set of all kinds

\mathcal{K} being the set of fundamental kinds K_{fund} closed under supervenience relations. The most straightforward way kinds and facts are related is that some facts will be unable to be systematized if we deny ourselves access to the languages that are able to express those facts. Going in the other direction, we can make it unlikely that certain kinds will be part of the best system if we do not attempt to systematize the relevant facts. What came out in Chapter 3 is that fact relativity has a finer grain than kind relativity. In the example of insect vision from Section 3.3.1, it became apparent that biologists cared about photons, but not *all* photons. Capturing the particular interests of biologists depended on being able to sort some (namely, the biologically interesting) photon related facts from others, and not just on picking out the biological kinds. Properly accommodating special science laws, the second (after Armstrong's objection) of the three major topics highlighted in Chapter 1, thus turned out to depend on the relationship between facts and language in the four parts of the BSA.

The third major topic, discussed at length in Chapter 4, was the trivial systems problem (TSP), the central concern of which is the relationship between language and the competition function. The TSP points out that a system can be made overwhelmingly simple by choosing a suitably gerrymandered language, and so you can find very strong and problematic systems and seemingly guarantee that they will be the best despite their problems. There is also the related problem of transcendent comparisons (PTC), which tells us that there are no measures of system goodness that are 'transcendent' in the sense that they can compare two systems expressed in any two different languages. There are, however, 'non-immanent' measures of

system goodness—contra the problem of immanent comparisons (PIC)—that can compare systems expressed in different languages (just not *any* different languages). A competition function defined by a non-immanent measure of system goodness has an associated class of languages. A non-immanent competition function C_1 may be associated with a class of languages \mathcal{L}_1 that includes L_{fund} (the language of the fundamental kinds) and every language L_i such that there exists a language $L_j \in \mathcal{L}_1$ and C_1 can compare two systems where one is expressed in L_i and the other in L_j . That is, C_1 is associated with a language class \mathcal{L}_1 comprised of all the languages that allow systems expressed in them to be compared using C_1 . Pairing C_1 with any other language class would make it possible to have two system-language pairs competing to be the best that are incomparable.

What is most interesting about these relationships between the parts of the BSA is that relativizing laws to one part can eliminate the need to relativize to other parts. The language relativity that is the hallmark of the BBSA must be supplemented by fact relativity (as per the arguments of Chapter 3). Indeed, language relativity, insofar as it is being used to ensure the existence of sets of laws for each special science, may be wholly replaced by fact relativity. It was argued in Chapter 4 that language relativity (or the related language privileging) is also not required to address the TSP, PIC, or PTC. In place of language relativity, we can have a combination of fact relativity and, if needed, competition relativity. Eliminating language relativity in this way has two advantages. First, it gives the BSA a way to make sense of the idea, developed in Section 4.7.2, that laws and kinds are discovered together. Second, it should make the BSA no more relative than when

language relativity is included, and, insofar as we can constrain the set of acceptable competition functions, it should make the BSA less relative (which serves to assuage the concern with relativity in Armstrong’s objection).

The question we might now ask is “Can the set of acceptable competition functions be constrained?”. The answer, I think, is yes. I argued in Chapter 5 that, while scientists may be more or less interested in simplicity and strength, they will always care, albeit to possibly greater or lesser degrees, about ‘induction friendliness’. However significant the role of induction friendliness in a best system competition will be, the ability to compare systems with respect to their induction friendliness is an all-or-nothing affair. If, as I recommend Section 5.2, we quantify induction friendliness with mutual information (MI), then, in terms of language relativity, the only language classes to which we can (effectively) relativize are those that are subsets of the class associated with MI.

So, with all that done, what does the ROBA say? How does it fill in the details of the four-part model of the BSA? To start, I will be bold with respect the more speculative parts of Chapter 5, and say that MI is all that is needed to capture what is meant by “best”—induction friendliness of course, but also the other components discussed in Section 5.3 like simplicity and strength—without need for competition relativity. This means that the set of all languages to be considered, $\mathcal{L}_{\text{ROBA}}$, is the set of all and only the languages that allow for comparisons between systems using MI. With $\mathcal{L}_{\text{ROBA}}$ in hand, we can define the set of all facts $\mathcal{F}_{\text{ROBA}}$ as being all the atomic sentences true of the actual world expressible by any language $L \in \mathcal{L}_{\text{ROBA}}$.

It is now helpful to note that the ROBA is fact relative, and runs a separate

best system competition for each $F \subseteq \mathcal{F}_{\text{ROBA}}$. Entering into each such competition are the facts F , the possible systematizations of those facts \mathcal{S}_F , and the languages $\mathcal{L}_{\text{ROBA}}$. (Note that \mathcal{S}_F may not include every system true of the world, and may include some systems that, while they must be true of F , are not true of the world as a whole.) The F -relative laws, then, are the regularities drawn from the system deemed best by the competition $C(F, \mathcal{S}_F, \mathcal{L}_{\text{ROBA}})$.

A complication crops up at this point with respect to induction friendliness and fact relativity. Induction friendliness, it was said in Section 5.1, is a feature of a world with respect to some system language pair. A preference for induction friendliness was understood there as a desire for the world to be maximally induction friendly (or close to it) with respect to a space of possible worlds. The complication is that, when we relativize to a subset of the facts in the actual world, it is not clear what subset of facts are to be considered when we move to some possible world. The most straightforward way of dealing with this may be to note that the actual world (or some subset of it) will be more or less induction friendly depending on the considered system-language pair. It seems plausible that scientists will prefer, absent an ability to ensure that they are in the most induction friendly of the possible worlds, that they describe the world they are stuck with (or some subset of it) in a way that makes it as induction friendly as they can. Thus, $C(F, \mathcal{S}_F, \mathcal{L}_{\text{ROBA}})$ will pick a system-language pair $S_{\text{best}}, L_{\text{best}}$ as the best if it maximizes (with respect to all $S, L \in \mathcal{S}_F \times \mathcal{L}_{\text{ROBA}}$) MI measured between the parts of F (as they are carved up by the respective pairs S, L).

In total, we have the following.

ROBA. A regularity is an F -law of nature just in case it appears in, and is expressed

according to, any system-language pair $S_{\text{best}}, L_{\text{best}}$ such that:

- (1) $F \subseteq \mathcal{F}_{\text{ROBA}}$, where $\mathcal{F}_{\text{ROBA}}$ is the set of atomic propositions expressible by any language in $\mathcal{L}_{\text{ROBA}}$.
- (2) $S_{\text{best}} \in \mathcal{S}_F$, where \mathcal{S}_F is the set of systems that are true of F .
- (3) $L_{\text{best}} \in \mathcal{L}_{\text{ROBA}}$, where $\mathcal{L}_{\text{ROBA}}$ is the set of languages that allow for inter-system-language pair comparisons using MI.
- (4) The MI of F is maximized with respect to all $S, L \in \mathcal{S}_F \times \mathcal{L}_{\text{ROBA}}$ when F is divided into related parts according to $S_{\text{best}}, L_{\text{best}}$.

There is, of course, a lot of work left to be done. The ROBA is by no means a complete analysis of laws of nature. Much more needs to be said about dividing sets of facts into parts following a given system-language pair, about just what the structure of $\mathcal{L}_{\text{ROBA}}$ is, in defense of using MI exclusively, etc. But in these respects, it is a no less complete analysis than any other BSA variant that casually appeals to, for example, “the best balance of simplicity and strength”.

It is also important to note that I do not take the ROBA to be the ultimate accomplishment of this dissertation. It is just one way of applying the lessons learned in each chapter. Those lessons—about the general structure of the BSA (introduced in Chapter 1), constraints on that structure, and how they relate to subjects such as (in Chapter 2) the variable interests of scientists, (in Chapter 3) interfield interactions and distinguishing special and fundamental laws, (in Chapter 4) the relationship between the languages used in science and the measures used to

choose best systems, and (in Chapter 5) the significance of the problems of induction to system choice—are where the value of this dissertation lies. My hope is that those lessons will be fruitfully applied in future work.

Part III

Appendix

Chapter A: Shannon's BSA

In this appendix, I provide a reconstruction of a classic theorem from [Shannon, 1948], but newly motivated to be about induction friendliness and the BSA. The theorem and its setup is worthy of being included here in full.¹

We have represented a discrete information source as a Markoff process. Can we define a quantity which will measure, in some sense, how much information is “produced” by such a process, or better, at what rate information is produced?

Suppose we have a set of possible events whose probabilities of occurrence are p_1, p_2, \dots, p_n . These probabilities are known but that is all we know concerning which event will occur. Can we find a measure of how much “choice” is involved in the selection of the event or of how uncertain we are of the outcome?

If there is such a measure, say $H(p_1, p_2, \dots, p_n)$, it is reasonable to require of it the following properties:

1. H should be continuous in the p_i .

¹ The proof, however, I leave for the most committed of readers to find in the second appendix to [Shannon, 1948].

2. If all the p_i are equal, $p_i = 1/n$, then H should be a monotonic increasing function of n . With equally likely events there is more choice, or uncertainty, when there are more possible events.
3. If a choice be broken down into two successive choices, the original H should be the weighted sum of the individual values of H . The meaning of this is illustrated in Fig. [A.1].




Figure A.1: Decomposition of a choice from three possibilities.

At the left we have three possibilities $p_1 = \frac{1}{2}$, $p_2 = \frac{1}{3}$, $p_3 = \frac{1}{6}$. On the right we first choose between two possibilities each with probability $\frac{1}{2}$, and if the second occurs make another choice with probabilities $\frac{2}{3}$, $\frac{1}{3}$. The final results have the same probabilities as before. We require, in this special case, that

$$H\left(\frac{1}{2}, \frac{1}{3}, \frac{1}{6}\right) = H\left(\frac{1}{2}, \frac{1}{2}\right) + \frac{1}{2}H\left(\frac{2}{3}, \frac{1}{3}\right)$$

The coefficient $\frac{1}{2}$ is the weighting factor introduced because this second choice only occurs half the time.

In Appendix 2, the following result is established:

Theorem 2: The only H satisfying the three above assumptions is of the

form:

$$H = -K \sum_{i=1}^n p_i \log p_i$$

where K is a positive constant.

[[Shannon, 1948](#), p. 10]

This is what gives rise to the “uncertainty” interpretation of entropy. Shannon himself notes that this theorem is unnecessary to the minimum average code length interpretation of entropy that is the primary concern of information theory (at least as far as Shannon was concerned in that paper). But it is illuminating.

And it is that illuminating quality that prompts the remaining work of this appendix. I reconstruct this theorem in the context of the BSA, leading to an interpretation of entropy as, roughly, “induction unfriendliness when the present state of the world is known”. More importantly, from there we can derive (approximately) MI as a measure of induction friendliness.

I begin by introducing a toy model of the BSA (Section [A.1](#)) that parallels Shannon’s toy model of a communication system. In the context of that toy model, I qualitatively develop the idea that different worlds may be more or less induction friendly (Section [A.2](#)). Instead of directly looking for the induction friendliness of the world, we will start by asking about the induction *unfriendliness* of a world, and doing that requires quite a bit of setup. First, in Section [A.3](#), I introduce toy frequencies that are like, but just not quite, the probabilities of individual (or combinations of) states in our toy worlds. The next, and crucial, step is to begin by

considering the induction unfriendliness of a world in just the simple situation of focusing on just the regularities in the given system have the same antecedent kind. It is this situation that most closely parallels Shannon’s theorem, and I re-motivate Shannon’s three assumptions for the purpose of quantifying induction unfriendliness in Sections [A.4](#), [A.5](#), and [A.6](#), respectively. In Section [A.7](#), I derive the induction unfriendliness of a world when the given system is allowed to include regularities with differing antecedent kinds as being the conditional entropy of consequent states in the world on antecedent states. Finally, in Section [A.8](#), I show how the induction friendliness of a world is approximately the MI between the antecedent and consequent states of the world.

As in Shannon, the success of the work in this section is not required for successfully arguing that induction friendliness, MI, and the BSA, should be brought together. The preceding sections should stand on their own. But I do think this exercise in reconstructing one of the original results of information theory is worthwhile. For one, it illustrates in significant detail *how* induction friendliness, MI, and the BSA, can be brought together. I also think it is emblematic of how the sorting out of the details of the best system competition should be done—namely, by carefully going through the needs and interests of scientific practice, and seeing what emerges from that.

A.1 The Toy BSA

To begin, there are the local qualities, the basic kinds, that are distributed throughout the world. The toy-ness of our model appears most strongly in the assumption that the world is a discrete and finite string of the basic kinds. Any toy world may be described as a ‘coin flip’ world where there are only two kinds 0 and 1 (or **heads** and **tails** or H and T or whatever one prefers). Here is an example of a coin flip world:

$$w_e = 10100101001010101111100101$$

By default I assume that a toy world is a coin flip world. I will refer to any position in the world where a kind obtains a “state”, as in “the first two states of w_e are 1 and 0, respectively”. When we are interested in more than two kinds I will say that the world is what it is “as expressed in L ”, where L consists in a set of kinds K and a function $T(-)$ from binary sequences to K . So, for an example language L_e , if $K_e = \{\mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{d}\}$, and T_e maps 10 to **a**, 01 to **b**, 00 to **c**, and 11 to **d**, we may say of our example world from above that

$$w_e =_{L_e} \mathbf{aabbcaaaddabb}.$$

That is, the world w_e as expressed in L_e —or, with kinds K_e according to the translation T_e —is “**aabbcaaaddabb**”.

To accompany our toy worlds is a toy conception of the BSA and induction. A system S , always to be paired with a language L , is characterized by a set of regularities of the form “If antecedent kind a obtains at any time t in the world, then the consequent kind c obtains at $t + 1$ ” for any $a, c \in K$. I will just write “ $a \rightarrow c$ ” when there will be no loss of clarity. A best system competition for a world w scores pairs $S_i, L_i \in \mathcal{S} \times \mathcal{L}$ according to (usually) the balance of simplicity and strength S_i has with respect to w when both are expressed in L_i . The laws of world w are the regularities that appear in S_{best} , where $S_{\text{best}}, L_{\text{best}}$ is the highest scored of the competing system-language pairs.

Epistemologically, the talk of simplicity and strength is providing a particular, if rather vague, story about the inductive practices of scientists. To say that S, L is the (epistemologically) best system-language pair—whatever regularities appear in our best estimates of the laws and however we generalize from what limited data we have to the whole world—is to say that it is the product of our best inductive practices. Part of the aim of Chapter 5 was to step back from the standard picture of the best system being the simplest and strongest on balance. In the place of simplicity and strength, we would like to install induction friendliness, but that cannot be done without the arguments to come. For now, we will suppose for any S, L that it has been given to science from on high as the conclusion of some ideal inductive practice.

In the context of the BSA, we normally think of S, L as being better or worse for a particular world w . But for the moment we are being neutral on what makes S, L better or worse; all we know is that it is the best. It is helpful, then, to flip our

usual thinking around to consider, for a given S, L , “How good is the world w ?”. Or, since S, L is the product of an ideal inductive practice: “How induction friendly is w ?”

A.2 Degrees of Induction (Un)Friendliness

We can start looking for an answer to the question of “How induction friendly is w ?” by illustrating Hume’s worry about the failure of constant conjunction in our toy worlds. The worry, at least on the surface, is that we see constant conjunction until some time at which there is a failure of the pattern. Such a possible world would be like w_{h} below.²

$$w_{\text{h}} = \overline{01}|_t 00\dots$$

In w_{h} we see a repetition of 0 followed by 1 until the step from t to $t + 1$ when the constant conjunction fails and a 0 is followed by another 0 instead of a 1. Hume is not really concerned about what happens after that (hence the ellipsis)—what matters is just that the constant conjunction did fail, and nothing that happened prior to t assures us that we aren’t in a world with such a failure.

We can motivate the possibility of there being more or less inductively friendly worlds by going beyond Hume and considering what the world looks like after that

² The subscript of vertical bar names the following point in time. An ellipsis indicates an indefinite but finite arbitrary string. An over-line indicates indefinite but finite repetitions of the over-lined sequence, or, when followed by a superscript, that superscript’s number of repetitions.

initial failure of the 01 pattern, as in the worlds described below.

$$\begin{aligned}
 w_{1 \times 00} &= \overline{01}|_t 00\overline{01} \\
 w_{2 \times 00} &= \overline{01}|_t 0000\overline{01} \\
 &\vdots \\
 w_{i \times 00} &= \overline{01}|_t \overline{00}^i \overline{01}
 \end{aligned}$$

Each of the $w_{i \times 00}$ worlds is consistent with w_h .³ In $w_{1 \times 00}$ the violation of the 01 pattern occurs only once. As failures of induction go, this presumably is not a bad one—yes there was a failure, but it is just the one and everywhere else our adoption of the $0 \rightarrow 1$ regularity will work so well. Similarly for $w_{2 \times 00}$ —clearly it is a 01 world with just a small wrinkle in the middle—but things are bit worse than they were in $w_{1 \times 00}$ since the dominant regularity isn't quite as dominant. The inductive inference that we presumably would make from observing the world prior to t to thinking that 0 is always followed by 1 seems less and less successful as the period in which that regularity fails to apply gets larger.

So it seems as though the world looks less induction friendly as i becomes larger. But things shouldn't be quite so simple. As i grows to make the 00 pattern a significant fraction of the world, we might stop thinking that we're looking at a 01 patterned world with some failure in the middle, but rather a world with a 01 regime, followed by a 00 regime, and then another 01 regime, and perhaps there is

³ Note that w_h actually describes a set of possible worlds on account of the openness of the ellipsis and number of initial repeated 01 pairs, and included in that set are the $w_{i \times 00}$ worlds. That is what I mean here by “consistent”.

some threshold where that makes for a more induction friendly world than the one in the series of worlds before it. Further still, as i becomes large relative to the total size of the world, what we will see is not a 01 world with some mess in the middle, or a 01 to 00 to 01 world, but rather a 00 world with some mess on the ends, and that may present another inflection point in the plot (against i) of the induction friendliness of the $w_{i \times 00}$ worlds.

We may also consider the problem of gruesome predicates, the titular “new riddle of induction”, raised in [Goodman, 1954]: “Grue” is a predicate that “applies to all things examined before t just in case they are green [in appearance] but to other things just in case they are blue” in appearance [Goodman, 1954, p. 74]. Similarly, “bleen” applies to a thing just in case it appears blue when examined prior to t and green after that. Notably, anything that we can say about green and blue can also be said in terms of grue and bleen, and vice versa. Prior to time t , all the emeralds we have examined have *appeared* green, but, prior to t , appearing green is indistinguishable from appearing grue. While we might like to generalize from our examinations of emeralds to the conclusion that all emeralds are green, the evidence is just as good in favor of the incompatible alternative conclusion that all emeralds are grue.

The obvious appeal of treating this problem in the context of the (toy) BSA is that either (1) t is at or after the end of our finite world, and so green and grue are effectively indistinguishable, or (2) we have access to the facts after t , and so can say with confidence whether emeralds are actually green or grue (depending on whether they appear green or blue after t , respectively). In the case of (1), generalizing

to all emeralds are green is not incompatible with generalizing to all emeralds are grue—the two regularities will either both be true or both be false in every possible world that ends at or before t —and so the problem goes away. In the case of (2): If emeralds do, in fact, persist in appearing green, then the employer of grue (and bleen) will be confronted with a world in which there is an abrupt shift from emeralds appearing grue to emeralds appearing bleen. When described using the language of grue and bleen, such a world seems less friendly to induction, but when described using the language of green and blue it seems perfectly induction friendly (with respect to emeralds) since emeralds persist in their greenness. Similarly, if the world is such that emeralds *do* change appearance from green to blue, then the world will seem more induction friendly if we describe it in terms of grue and bleen as opposed to green and blue.

The toy worlds that I have described cannot quite capture what is happening in Goodman's grue emerald thought experiment because the time-stamped translations from green and blue to grue and bleen (or vice versa) cannot straightforwardly be done with the time invariant translation functions $T(-)$. We can, though, appreciate that there should be a difference in induction friendliness between a world in which the same predicate always applies (as when the emerald always appears green) and a world with an initial segment of one kind occurring followed by some other kind occurring in the rest of the world (as when we find the emerald switches appearance from green to blue). We can likewise appreciate that the difference in induction friendliness should be perfectly reversed if we were to redescribe each world in such a way as to make the first world switch the prevailing kind (as when the emerald

switches from grue to bleen) and the second world have a single prevailing kind (as when the emerald is always grue).

The time-stamping involved in the translations from green and blue to grue and bleen is not necessary for the induction friendliness of a world to change under redescription. Consider the following worlds (with some spacing added for clarity):

$$w_a = 10\ 01\ 10\ 01\ 10\ 01\ 10\ 01\ 10\ 01$$

$$w_b = 110\ 011\ 110\ 001\ 110\ 001\ 100\ 011\ 100\ 001$$

The world w_a is clearly quite regular, but not in a way that can be captured by our restricted conception of what regularities may appear in a system because we can use just those concerned with what single state follows another single state. To correct this issue we could expand the sort of regularities that we allow for to include what pairs of states follow other pairs, but that is unnecessary. All we need to do is redescribe the world. Using the same language L_e introduced earlier—according to which each 10 goes to **a** and 01 goes to **b**—we get that

$$w_a =_{L_e} \mathbf{a\ b\ a\ b\ a\ b\ a\ b\ a\ b.}$$

Now we can say that w_a , redescribed using L , is perfectly regular; **a** is always followed by **b**, and **b** is always followed by **a** (at least if we ignore the end of the world not having any consequent state).

In contrast to w_a , w_b does not seem particularly regular. But what if we

care about only the end points of triplets of states? Let L'_e be like L_e except that whenever the T_e of L_e takes a pair of states to some element of K_e , T'_e takes two triplets to the same element of K_e where the first and last states of the triplet and pair are the same, and one triplet as has a center state of 0 and the other a center state of 1. For example, while T_e maps 10 to a, T'_e maps both 100 and 110 to a. Expressing w_b in L'_e gets us that

$$w_b =_{L'_e} \text{ a b a b a b a b a b }$$

So w_b is regular, at least when described in the right sort of way. This should come as no surprise. How induction friendly a world is depends on the language we use to describe the world. We could try to privilege as the “true” induction friendliness of the world whatever value is yielded by the “true” language for expressing the world, but in our toy worlds there is no such true language.

A.3 Toy Frequencies

Let us return to our first example world

$$w_e =_{L_e} \text{ aabbcaaaddabb.}$$

There are, straightforwardly, some actual relative frequencies for this world. Namely, there are the frequencies with which kinds obtain in the world. In w_e , expressed

according to L_e , those frequencies are

$$f_{w_e, L_e}(\mathbf{a}) = 6/13$$

$$f_{w_e, L_e}(\mathbf{b}) = 4/13$$

$$f_{w_e, L_e}(\mathbf{c}) = 1/13$$

$$f_{w_e, L_e}(\mathbf{d}) = 2/13.$$

For convenience, I will drop the subscripts on f when they are clear in context.

There are also frequencies associated with seeing pairs of kinds obtaining.

Continuing with our example,

$$f(\mathbf{aa}) = 1/4$$

$$f(\mathbf{ab}) = 1/6$$

$$f(\mathbf{ac}) = 0$$

$$f(\mathbf{ad}) = 1/12$$

and so on for all 16 possible pairs. These frequencies are interesting in our toy worlds since they measure the frequency with which a corresponding regularity is the correct one to apply. For example, $f(\mathbf{ab}) = 1/6$ indicates that the regularity $\mathbf{a} \rightarrow \mathbf{b}$ correctly characterizes a sixth of all transitions from one state to the next. With this point in mind, I will call these “significance” frequencies.

In addition to the significance frequencies, it will be helpful to consider the frequency with which a regularity will be deployed successfully by someone with

knowledge of the antecedent state but not the consequent state. For some $a, c \in K$, the “success” frequency $f_{w,L}(c|a)$ is the frequency with which c is the second kind in all state pairs that begin with a . In our example

$$f(\mathbf{a}|\mathbf{a}) = 1/2$$

$$f(\mathbf{b}|\mathbf{a}) = 1/3$$

$$f(\mathbf{c}|\mathbf{a}) = 0$$

$$f(\mathbf{d}|\mathbf{a}) = 1/6$$

and so on.

These frequencies behave almost like regular probabilities. The sum over all $k \in K$ of $f(k)$ will equal one, as a regular probability would. We also have it that

$$f(c|a) \approx \frac{f(ac)}{f(a)}$$

where the failure to achieve equality of the two sides is due to the first and last states in a world not having antecedent or consequent states (respectively). However, as worlds become large, the difference from equality goes to zero. This near equality is reminiscent of the standard definition of conditional probability—the conditional probability $p(x|y)$ is the joint probability $p(x, y)$ divided by the probability $p(y)$ —suggesting that success frequencies are analogous to conditional probabilities and significance frequencies to joint probabilities. The biggest failure of this analogy will be due to the fact that $f(ac) \neq f(ca)$ but $p(x, y) = p(y, x)$, and the rest will be due

to the end effects already mentioned that become smaller as worlds become larger.

A.4 Continuity

Suppose we have observed that some kind a ($\in K$) obtains at t . We would like to know what happens at $t + 1$. If our given system S contains just the single regularity $a \rightarrow c$, then we at least know what to expect. It of course could fail to be the case that in this particular instance c is what obtains at $t + 1$; S may be the best system, but that doesn't guarantee that the regularity will hold in every situation in which it is applicable (where a regularity is "applicable" at a time t just in case its antecedent kind obtains at t).⁴ What is at issue is the success frequency associated with the regularity: If $f(c|a) = 1$, then every a is followed by a c and the regularity $a \rightarrow c$ holds everywhere it is applicable. If $f(c|a) < 1$, then the regularity fails with a frequency of $1 - f(c|a)$ among the points in time when it is applicable. In general (for all $a, c \in K$), the use (when it is applicable) of a regularity $a \rightarrow c$ will be successful with frequency $f(c|a)$ —this is precisely the conditional-probability-like "success" frequency introduced above.

Situated in this world where our best inductive practices recommend the adoption of the $a \rightarrow c$ regularity, what would we like to be true of $f(c|a)$? Clearly we would do best if $f(c|a) = 1$, as then every time we see an a we would correctly

⁴ This is in conflict with the idea that laws should be universal. We can assume that the non-universal nature of these regularities is just part of the toy-ness of the toy model. It might also be that laws are not—and should not be expected to be—universal. This is the position of [Braddon-Mitchell, 2001] and [Schrenk, 2008, Schrenk, 2014], and, if correct, the non-universality of these regularities is not a problem at all. In Chapter 3, I discuss these violations of universality and refer to them as "simple system exceptions".

predict the appearance of a c . As the value of $f(c|a)$ falls away from 1, its successful use becomes less frequent, and, presumably, the induction friendliness of a world where $a \rightarrow c$ is truly the single best regularity to adopt goes down proportionally.

Of course, $a \rightarrow c$ need not be the only regularity that has been recommended by our best inductive practices. Continuing to restrict ourselves to the case where we have observed a , the set of applicable regularities (adopted on the recommendation of our best inductive practices or not) is comprised of the regularities $a \rightarrow c_i$ for all $c_i \in K$. Compare two worlds using the same language. If $f(c_i|a)$ —which, since a is fixed for the time being, I will write as $f_a(c_i)$ for convenience and clarity—is lower in the first world than the second, then necessarily the success frequency of at least one of the other applicable regularities will be lower in the second world than in the first. In other words, if $f_{a,w,L}(c_i) < f_{a,w',L}(c_i)$ there exists a $c_j \in K$ such that $j \neq i$ and $f_{a,w,L}(c_j) > f_{a,w',L}(c_j)$. It is thus not immediately obvious how strongly, when considering the full set of applicable regularities, deviation from a success frequency of 1 impacts the induction friendliness of a world, since being further from 1 is less friendly, but ensures that there is at least one other regularity that is more successful.

Let H be the function that is meant to quantify the induction unfriendliness of a world with respect to the best system whose regularities all feature a as their antecedent kind. It seems clear, at the very least, that H is a function of the $f_a(c)$ for all $c \in K$. I will assume something slightly stronger (and make all the dependencies explicit in subscripts):

Continuity. $H_{a,w,L}$ is a continuous function of just the frequencies $f_{a,w,L}(c)$ for all $c \in K$.

CONTINUITY is stronger than what is immediately apparent in light of the preceding discussion on account of insisting that (1) H is continuous, and that (2) H is a function of *just* the success frequencies.

Defending (2) is a matter of reemphasizing the simplifying assumptions surrounding our toy model and what we are trying to quantify. To say that H is a function of *just* the success frequencies is a bit disingenuous since the success frequencies themselves depend on the world and choice of language. The only information about the world that is not contained in the success frequencies has to do with the ordering and number of states (and even then there is some information since success frequencies are concerned with pairs). We should not worry about H not exploiting that information because such large scale ordering information is not generally available to someone trying to exploit regularities arrived at by induction.

In defense of (1), the continuity of H : Suppose we have two worlds w and w' . Without loss of generality, suppose further that we have left them as being described in the basic coin-flip form—this allows us the convenient fact that in each world $f_a(1) = 1 - f_a(0)$, so there is no question when there is a difference in the success frequencies between the two worlds about how the difference is distributed across the relevant regularities. Let $f_{a,w}(1) = f_{a,w'}(1) + \Delta$. We should take H to be continuous if, for every possible value of $f_{a,w}(1)$, in the limit as Δ goes to 0, the difference between H_w and $H_{w'}$ also goes to zero.

Now consider an arbitrary value of $f_{w,a}(1)$. Say, $f_{a,w}(1) = .75$. Since the relevant success frequency is close to, but not, 1, this world w isn't ideal, but it's also not terrible. Similarly, for some small positive value of Δ , the world w' won't be too bad. With respect to just the $a \rightarrow 1$ regularity, w' is definitely more induction unfriendly than w (since $f_{a,w}(1) > f_{a,w'}(1)$)—note though that CONTINUITY makes no demands on the direction of the difference. There will be some mitigation of the difference between the induction unfriendliness of the two worlds when we consider all the applicable regularities because $f_{a,w}(0) < f_{a,w'}(0)$, and so, with respect to just the $a \rightarrow 0$ regularity, w' is better than w . As $f_{a,w'}(1)$ gets closer to .75 (i.e., as Δ goes to zero), the induction unfriendliness of the two worlds should also converge as they become more and more alike. And so it should be in general for any value of $f_{a,w}(1)$.

There are two salient values of $f_{a,w}(1)$ where one might reasonably have intuitions against the continuity of H . At $f_{a,w}(1) = .5$, any difference in the success frequencies would make w' better. At $f_{a,w}(1) = 1$, any difference in the success frequencies would make w' worse. We might take the salience of these points to indicate their possessing *very*—i.e., discontinuously—high and low values of H , respectively. This is question begging, but perhaps not more so than the above appeals to intuition in favor of the continuity of H . Here then is an indirect argument for continuity via an attack on discontinuity. Discontinuity has the prima facie disadvantage of leaving us with a measure H that is not mathematically well behaved. Its advantage is that, as long as we are using H to quantify a cardinal ranking of induction unfriendliness, the discontinuity might better reflect the ranking we are

trying to quantify. But this is only an apparent advantage. If there is discontinuity to be found in relation to induction unfriendliness, it will be ambiguous whether it is discontinuity in H or discontinuity in our response to H . As a parallel example, think of a full belief threshold in our degrees of belief. Our degrees of belief can take values in the continuum from 0 to 1, but the move from partial or no belief to full belief, if thresholded, will be discontinuous (as in, once my degree of belief in P passes .95 I will say that I have a full belief in P). Things may be similar for induction unfriendliness. H itself can be continuous (with all the mathematical advantages that brings), and any apparent discontinuity may be pushed from H itself to how we respond to H .

A.5 Monotonicity

To motivate the second assumption that we will make about H , consider the following scenario: You can choose to be a scientist in either world w or w' . L is the language of the best system-language pair in w , and L' the best in w' . The antecedent kind a that is the concern of H appears in and is treated the same way by both L and L' . No kind in K or K' fails to obtain in the respective worlds, and $|K| < |K'|$. Lastly, the L and L' are such that

- (1) for all $c_i, c_j \in K$, $f_{a,w,L}(c_i) = f_{a,w,L}(c_j)$, and
- (2) for all $c'_i, c'_j \in K'$, $f_{a,w',L'}(c'_i) = f_{a,w',L'}(c'_j)$.

Stated less formally: In two worlds w and w' where you have just observed a , the consequent kinds are all equally likely, and there are more possible consequent

kinds in w' (as described by L'), than w (as described by L). In which world would you rather be? We've assumed that, relative to a particular regularity $a \rightarrow c$, a world is more friendly, or less induction unfriendly, the closer $f_a(c)$ is to 1. When we move to considering all the applicable regularities the issue is confused because, as the success frequency of one regularity goes down, the success frequency of at least one other must go up. In this case we know exactly how all the success frequencies compare. Since the $f_{a,w,L}(c)$ are all equal, each is equal to $1/|K|$. Similarly in w' we have it that, for all $c' \in K'$, $f_{a,w',L'}(c') = 1/|K'|$. Thus the success frequency of every applicable regularity in w is strictly closer to 1 than the success frequency of any applicable regularity in w' . So w is preferable to—that is, more induction friendly than— w' . When the success frequencies of the applicable regularities are all equal, a world is less preferable—that is, more induction unfriendly—when there are more applicable regularities to consider.

We make the above explicit in our second assumption about H :

Monotonicity. $H_{a,w,L}$ is a monotonically increasing function of $|K|$ when, for all

$$c_i, c_j \in K, f_{a,w,L}(c_i) = f_{a,w,S,T}(c_j).$$

Of the assumptions that we have made (and will make) about H , MONOTONICITY is the only one that influences the direction of changes in H relative to changes in the relevant success frequencies. It is also worth noting that MONOTONICITY makes H a tracker of a kind of ontological simplicity⁵ (or lack thereof) by counting a world and

⁵ [Schulte, 2008] employs “ontological simplicity” in an analysis of the discovery of particle families. He characterizes ontological simplicity as encouraging the use of “as few ontological categories as possible”—for our purposes, this means having fewer kinds—and “as many particle families as possible that are disjoint, that is, categories whose boundaries do not overlap” [Schulte, 2008, p. 307]

associated system-language pair as being less friendly to induction precisely when more kinds are being used.

A.6 Decomposition

Consider a particular relation between two languages for a world that I will call *decomposition*. L' is a decomposition of L in w iff

(D1) $K = K'$ except that there are $n \geq 2$ kinds $c_1, c_2, \dots, c_n \in K$ that do not appear in K' , and $n + 1$ kinds $c'_0, c'_1, c'_2, \dots, c'_n \in K'$ that do not appear in K , and

(D2) T and T' are such that $w_L = w_{L'}$ except, for each of the c_i among the $c_1, c_2, \dots, c_n \in K$, wherever c_i obtains in w_L , the ordered pair (c'_0, c'_i) obtains in $w_{L'}$.

In our running example with language L_e and world w_e : The language L'_e is a decomposition of L_e in w_e when $K_e = \{a, b, c, d\}$, $K'_e = \{c, d, x, y, z\}$, and, with spacing for clarity,


$$w_{e=L_e} \quad a \quad a \quad b \quad b \quad c \quad a \quad a \quad a \quad d \quad d \quad a \quad b \quad b$$

$$w_{e=L'_e} \quad xy \quad xy \quad xz \quad xz \quad c \quad xy \quad xy \quad xy \quad d \quad d \quad xy \quad xz \quad xz.$$

It will be convenient to be able to characterize a decomposition relation by saying what states decompose into what pairs of states. In our example we can characterize the decomposition relation by saying that “a decomposes to xy and b decomposes to xz” (in w , from L_e to L'_e).

Changing the language being used also changes how our given system S is expressed. If S, L contained the regularities $a \rightarrow a$ and $a \rightarrow b$, S, L'_e should contain

the regularities $a \rightarrow x$, $x \rightarrow y$, and $x \rightarrow z$. This is illustrated in the picture below.



We can quickly say some things about the frequencies involving \mathbf{a} , \mathbf{b} , \mathbf{x} , \mathbf{y} , and \mathbf{z} :

$$f_a(\mathbf{z}) = f_a(\mathbf{a}) + f_a(\mathbf{b})$$

$$f_a(\mathbf{a}) = f_a(\mathbf{y}|\mathbf{x})f_a(\mathbf{x})$$

$$f_a(\mathbf{b}) = f_a(\mathbf{z}|\mathbf{x})f_a(\mathbf{x}).$$

The success frequencies associated with the L_e kinds \mathbf{a} and \mathbf{b} are no different from chained success frequencies in L'_e of going from a to \mathbf{x} and then from \mathbf{x} to \mathbf{y} or \mathbf{z} . Effectively all that has happened is that we changed the names of the \mathbf{a} and \mathbf{b} kinds to \mathbf{xy} and \mathbf{xz} , respectively. It thus seems fair to assume that our move from L_e to L'_e makes no difference to the inductive unfriendliness of the world. In particular, the contribution to induction friendliness made by having the kinds \mathbf{a} , \mathbf{b} , and their associated regularities, in S, L_e should be the same as the contributions made by having \mathbf{x} , \mathbf{y} , \mathbf{z} , and *their* associated regularities in S, L'_e . Similarly, the contributions made by the kinds \mathbf{c} and \mathbf{d} should be the same because they were unaffected by the decomposition.

The tricky part with this is that the measure of unfriendliness H that we have been building up is only defined for collections of regularities that have the same antecedent state, which is a rule we want to break when considering the unfriendliness of S, L' . To get around this, we can decompose the regularities in S, L'_e : First we look at the H of the regularities taking us from a to \mathbf{x} , \mathbf{c} , and \mathbf{d} . Then to that we add the H associated with the regularities taking us from \mathbf{x} to \mathbf{y} and \mathbf{z} , but weighted by the frequency with which those regularities are applicable (conditional on a having preceded the \mathbf{x}). This is our third assumption about H :

Decomposition. If L' is a decomposition of L in w where the kinds $k_1, \dots, k_n \in K$ are decomposed into pairs of K' kinds $(k'_0, k'_1), \dots, (k'_0, k'_n)$ and the kinds a and k_{n+1}, \dots, k_{n+m} appear in both K and K' . Then

$$H_{w,S,L}(w_{(c|a)}) = H_{w,S,L'}(w_{(c|a)}) + f_{a,w,L'}(k'_0) \times H_{w,S,L'}(w_{(c|a,k'_0)})$$

In the above, $w_{(c|a)}$ stands for the sequence of states in the world that are consequent to a and, similarly, $w_{(c|a,k'_0)}$ stands for the states consequent to the ordered pair of states a, k'_0 .

Expanding on this to make the relevant (non-zero valued) frequencies explicit, we get

$$\begin{aligned} H_{w,S,L}(f_a(k_1), \dots, f_a(k_{n+m})) &= H_{w,S,L'}(f_a(c'_0), f_a(k_1), \dots, f_a(k_m)) \\ &+ f_{a,w,L'}(k'_0) \times H_{a,w,L'}(f_a(k'_1|k'_0), \dots, f_a(k'_n|k'_0)). \end{aligned}$$

And, spelled out in our running example, we get that

$$\begin{aligned}
 H_{w,S,L_e}(f_a(\mathbf{a}), \dots, f_a(\mathbf{d})) &= H_{w,S,L'_e}(f_a(\mathbf{x}), f_a(\mathbf{c}), f_a(\mathbf{d})) \\
 &\quad + f_{a,w,L'_e}(\mathbf{x}) \times H_{w,S,L'_e}(f_a(\mathbf{y}|\mathbf{x}), f_a(\mathbf{z}|\mathbf{x})).
 \end{aligned}$$

A.7 Induction Unfriendliness

Our three assumptions about H —CONTINUITY, MONOTONICITY, and DECOMPOSITION—are shown in [Shannon, 1948] to be satisfied uniquely by the function

$$H_{w,S,L}(w_{(c|a)}) = -k \sum_{c \in K} f_a(c) \log f_a(c)$$

where k is an arbitrary constant that determines the units of H (e.g. $k = 1/\log(2)$ gives units of bits). For a variety of reasons—most notably, its formal resemblance to the concept from statistical physics— H is sometimes called the *entropy* of $w_{(c|a)}$.

The entropy of $w_{(c|a)}$ is the amount of induction unfriendliness for a world associated with the regularities in the best system that have a as their antecedent kind. In our running example, a could be any of \mathbf{a} , \mathbf{b} , \mathbf{c} , or \mathbf{d} . The best system S of world w expressed in L presumably contains regularities that feature all four kinds in the antecedent state position, and so we have all of $H(w_{(c|\mathbf{a})})$, $H(w_{(c|\mathbf{b})})$, $H(w_{(c|\mathbf{c})})$, and $H(w_{(c|\mathbf{d})})$, as relevant to the total unfriendliness of the world. Let each contribute in proportion to how frequently their respective regularities are applicable; that is, let the total unfriendliness of the world be the average over

possible values of a of the unfriendliness measures $H(w_{(c|a)})$:

$$\begin{aligned} E_a[H(w_{(c|a)})] &= \sum_{a \in K} f(a) \left[-k \sum_{c \in K} f_a(c) \log f_a(c) \right] \\ &= -k \sum_{a \in K} \sum_{c \in K} f(a) f_a(c) \log f_a(c). \end{aligned}$$

Noting that $f_a(c) = f(c|a)$, and that $f(ac) = f(c|a)f(a)$, we can rewrite the above as what is known as the *conditional entropy* of consequent states given antecedent states

$$H_{w,S,L}(w_{(c)}|w_{(a)}) = -k \sum_{a \in K} \sum_{c \in K} f(ac) \log f(c|a).$$

If one is inclined to adopt some of the terminology of information theory, then this result has a very sensible reading: The induction unfriendliness of a world is equivalent to our uncertainty about the immediate future when we know the immediate past. If you are not so inclined, then the preceding argument can stand on its own, and perhaps offer a more palatable alternative reading for some other uses of information theoretic language.

A.8 Induction Friendliness

Let us pause to take quick stock of what has been done so far. We are assured that S, L is the best system-language pair of world w , and have set out to answer the question “How induction friendly is w ?”. We have gotten as far as saying that the induction *unfriendliness* of w is $H(w_{(c)}|w_{(a)})$, the conditional entropy of w 's

consequent states given the antecedent states.

One way of thinking about our measure of induction unfriendliness is that it tells us how bad things *still* are even after we are given the best system and the current (antecedent) state of the world. This suggests that there is a sort of induction unfriendliness to the world before we are given the best system. I say “sort of” here because there is no induction involved at that point—without the best system in hand, there are no generalizations around that have been made based on past experience—just raw guessing about what kind is about to obtain. Conveniently, we already know how to measure this. Recall that the induction unfriendliness of the world for a particular given antecedent state a is

$$H_{w,S,L}(w_{(c|a)}) = -k \sum_{c \in K} f_a(c) \log f_a(c).$$

When we were looking for the total induction unfriendliness of the world above, we were going to know the antecedent state, we just didn’t know which one it would be and so we took the average for all $a \in K$. In the current situation, anything we know about the antecedent state is irrelevant since we do not have a system on hand to exploit that knowledge. Thus the only applicable frequencies are the raw unconditional frequencies of the consequent states, and we get an “initial unfriendliness” of a world w and language L as the entropy simply of the consequent states:

$$H_{w,S,L}(w_{(c)}) = -k \sum_{c \in K} f(c) \log f(c).$$

Now what of induction friendliness? The system language pair S, L is assumed to be the best because it is the output of best inductive practices. We have worked out that $H(w_{(c)}|w_{(a)})$ is a measure of how unfriendly the world is *after* the best inductive practices recommend adopting S, L . And we have just determined that $H_{w,S,L}(w_{(c)})$ is a measure of how unfriendly the world is *before* the best inductive practices recommend adopting S, L . Let the induction friendliness of the world be how much less the world is unfriendly after implementing the best inductive practices. That is, let the induction friendliness for a world w given the best system-language pair S, L , be

$$I_{w,S,L}(w_{(a)}, w_{(c)}) = H_{w,S,L}(w_{(c)}) - H(w_{(c)}|w_{(a)}).$$

This is *almost* the MI between the antecedent and consequent states of w because, while MI is symmetric—i.e. it would be that $I(w_{(a)}, w_{(c)}) = I(w_{(c)}, w_{(a)})$ —this will not generally be true for us because of the asymmetry of $f(ac)$. However, we can see that it is formally similar by exploiting the fact that $f(ac) = f(c|a)f(a)$ and $f(c) = \sum_a f(c|a)f(a)$ are approximately true (with only edge effects violating equality):

$$\begin{aligned} I_{w,S,L}(w_{(a)}, w_{(c)}) &= H_{w,S,L}(w_{(c)}) - H(w_{(c)}|w_{(a)}) \\ &= k \sum_{c \in K} \sum_{a \in K} f(ac) \log f(c|a) - k \sum_{c \in K} f(c) \log f(c) \end{aligned}$$

$$\begin{aligned}
&= k \sum_{c \in K} \sum_{a \in K} f(ac) \log f(c|a) \\
&\quad - k \sum_{c \in K} \sum_{a \in K} f(c|a) f(a) \log f(c) \\
&= k \sum_{c \in K} \sum_{a \in K} f(ac) [\log f(c|a) - \log f(c)] \\
&= k \sum_{c \in K} \sum_{a \in K} f(ac) \log \frac{f(ac)}{f(a)f(c)}
\end{aligned}$$

while the MI of two random variables X and Y is

$$I(X; Y) = k \sum_{x \in X} \sum_{y \in Y} p(x, y) \log \frac{p(x, y)}{p(x)p(y)}.$$

Why should we think that there is more than mere formal similarity between these two measures? The kinds of systems that we allow for has an effect on what our measure of induction friendliness will look like. If we restrict ourselves to asymmetric regularities in our systems—as we have—then we should expect an asymmetric measure of induction friendliness. If, for example, we allow for symmetric systems that can treat $w_{(a)}$ and $w_{(c)}$ as random variables running alongside each other, and not just as slightly offset fragments of the same sequence, then it makes perfect sense to equate $f(ac)$ and $p(a, c)$ since there will no longer be a privileged ordering between the states of $w_{(a)}$ and $w_{(c)}$. In general, something like the following should be true:

Induction Friendliness. The induction friendliness of a world w and given system-language pair S, L is the mutual information between w 's parts, where what “ w 's parts” are is determined by S, L .

Bibliography

- [Albert, 2000] Albert, D. Z. (2000). *Time and Chance*. Cambridge, MA: Harvard University Press.
- [Anderson, 1972] Anderson, P. W. (1972). More is different. *Science*, 177(4047):393–396.
- [Armstrong, 1983] Armstrong, D. M. (1983). *What is a Law of Nature?* Cambridge: Cambridge University Press.
- [Barlow, 1952] Barlow, H. B. (1952). The size of ommatidia in apposition eyes. *Journal of Experimental Biology*, 29(4):667–674.
- [Beebe, 2000] Beebe, H. (2000). The non-governing conception of laws of nature. *Philosophical and Phenomenological Research*, pages 571–594.
- [Braddon-Mitchell, 2001] Braddon-Mitchell, D. (2001). Lossy laws. *Noûs*, 35(2):260–277.
- [Callender and Cohen, 2010] Callender, C. and Cohen, J. (2010). Special sciences, conspiracy and the better best system account of lawhood. *Erkenntnis*, 73(3):427–447.
- [Carroll, 1990] Carroll, J. (1990). The humean tradition. *The Philosophical Review*, 99(2):185–219.
- [Cartwright, 1983] Cartwright, N. (1983). *How the Laws of Physics Lie*. Oxford University Press.
- [Cartwright, 1999] Cartwright, N. (1999). *The Dappled World: A Study of the Boundaries of Science*. Cambridge University Press.
- [Cartwright, 2002] Cartwright, N. (2002). In favor of laws that are not ceteris paribus after all. *Erkenntnis*, 57(3):425–439.

- [Cohen and Callender, 2009] Cohen, J. and Callender, C. (2009). A better best system account of lawhood. *Philosophical Studies*, 145(1):1–34.
- [Darden and Maull, 1977] Darden, L. and Maull, N. (1977). Interfield theories. *Philosophy of science*, pages 43–64.
- [Dupré, 1993] Dupré, J. (1993). *The Disorder of Things: Metaphysical Foundations of the Disunity of Science*. Harvard University Press.
- [Ellis, 1986] Ellis, J. (1986). The superstring-theory of everything, or of nothing? *Nature*, 323:595–598.
- [Floridi, 2011] Floridi, L. (2011). *The philosophy of information*. Oxford University Press.
- [Forster and Sober, 1994] Forster, M. and Sober, E. (1994). How to tell when simpler, more unified, or less ad hoc theories will provide more accurate predictions. *The British Journal for the Philosophy of Science*, 45(1):1–35.
- [Frigg and Hoefer, 2010] Frigg, R. and Hoefer, C. (2010). Determinism and chance from a humean perspective. In Dieks, D., Gonzalez, W., Hartmann, S., Weber, M., Stadler, F., and Uebel, T., editors, *The Present Situation in the Philosophy of Science*, pages 351–372. Berlin and New York: Springer.
- [Frisch, 2011] Frisch, M. (2011). From arbutnot to boltzmann: The past hypothesis, the best system, and the special sciences. *Philosophy of Science*, 78(5):1001–1011.
- [Godfrey-Smith and Sterelny, 2016] Godfrey-Smith, P. and Sterelny, K. (2016). Biological information. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, summer 2016 edition.
- [Goodman, 1954] Goodman, N. (1954). *Fact, Fiction, and Forecast*. Cambridge, MA: Harvard University Press.
- [Greeno, 1970] Greeno, J. G. (1970). Evaluation of statistical hypotheses using information transmitted. *Philosophy of Science*, 37(2):279–294.
- [Hanna, 1969] Hanna, J. F. (1969). Explanation, prediction, description, and information theory. *Synthese*, 20(3):308–334.
- [Hoefer, 2007] Hoefer, C. (2007). The third way on objective probability: A sceptic’s guide to objective chance. *Mind*, 116(463):549–596.
- [Hume, 1896] Hume, D. (1896). A treatise of human nature by david hume, reprinted from the original edition in three volumes and edited, with an analytical index, by l.a. selby-bigge, m.a.
- [Hüttemann, 2014] Hüttemann, A. (2014). Ceteris paribus laws in physics. *Erkenntnis*, 79(10):1715–1728.

- [Kieseppä, 1997] Kieseppä, I. (1997). Akaike information criterion, curve-fitting, and the philosophical problem of simplicity. *The British journal for the philosophy of science*, 48(1):21–48.
- [Kieseppä, 2001] Kieseppä, I. (2001). Statistical model selection criteria and the philosophical problem of underdetermination. *The British journal for the philosophy of science*, 52(4):761–794.
- [Landenmark et al., 2015] Landenmark, H. K., Forgan, D. H., and Cockell, C. S. (2015). An estimate of the total dna in the biosphere. *PLoS Biol*, 13(6):e1002168.
- [Lange, 2004] Lange, M. (2004). The autonomy of functional biology: a reply to rosenberg. *Biology and Philosophy*, 19(1):93–109.
- [Lewis, 1970] Lewis, D. (1970). How to define theoretical terms. *The Journal of Philosophy*, 67(13):427–446.
- [Lewis, 1973] Lewis, D. (1973). *Counterfactuals*. Malden, MA: Blackwell Publishers.
- [Lewis, 1983] Lewis, D. (1983). New work for a theory of universals. *Australasian Journal of Philosophy*, 61(4):343–377.
- [Lewis, 1986] Lewis, D. (1986). *Philosophical Papers Vol. II*. Oxford University Press.
- [Lewis, 1994] Lewis, D. (1994). Humean supervenience debugged. *Mind*, 103(412):473–490.
- [Loewer, 1996] Loewer, B. (1996). Humean supervenience. *Philosophical Topics*, 24(1):101–127.
- [Loewer, 2007] Loewer, B. (2007). Laws and natural properties. *Philosophical Topics*, 35(1/2):313–328.
- [Loewer, 2012] Loewer, B. (2012). Two accounts of laws and time. *Philosophical Studies*, 160(1):115–137.
- [Maudlin, 2007] Maudlin, T. (2007). *The metaphysics within physics*. Oxford University Press on Demand.
- [Mill, 1882] Mill, J. S. (1882). *A System of Logic Ratiocinative and Inductive*. Harper & Brothers.
- [Overbye, 2016] Overbye, D. (2016). The particle that wasn't. *The New York Times*, August 5.
- [Pemberton and Cartwright, 2014] Pemberton, J. and Cartwright, N. (2014). Ceteris paribus laws need machines to generate them. *Erkenntnis*, 79(10):1745–1758.

- [Quine, 1970] Quine, W. V. O. (1970). *Philosophy of logic*. Harvard University Press.
- [Ramsey, 1928] Ramsey, F. P. (1928). Universals of law and of fact. In Mellor, D. H., editor, *Foundations: Essays in Philosophy, Logic, Mathematics and Economics*, pages 128–132. London: Routledge & Kegan Paul.
- [Ramsey, 1929] Ramsey, F. P. (1929). General propositions and causality. In Mellor, D. H., editor, *Foundations: Essays in Philosophy, Logic, Mathematics and Economics*, pages 133–151. London: Routledge & Kegan Paul, London.
- [Reutlinger and Unterhuber, 2014] Reutlinger, A. and Unterhuber, M. (2014). Thinking about non-universal laws. *Erkenntnis*, 79(10):1703–1713.
- [Rieke and Baylor, 1998] Rieke, F. and Baylor, D. (1998). Single-photon detection by rod cells of the retina. *Reviews of Modern Physics*, 70(3):1027.
- [Salmon, 1971] Salmon, W. C. (1971). *Statistical explanation and statistical relevance*, volume 69. University of Pittsburgh Pre.
- [Schrenk, 2008] Schrenk, M. (2008). A theory for special science laws. In H. Bohse, K. Dreimann, S. W., editor, *Selected Papers Contributed to the Sections of GAP.6*, pages 121–131. Paderborn: Mentis.
- [Schrenk, 2014] Schrenk, M. (2014). Better best systems and the issue of cp-laws. *Erkenntnis*, 79(10):1787–1799.
- [Schulte, 2008] Schulte, O. (2008). The co-discovery of conservation laws and particle families. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 39(2):288–314.
- [Shannon, 1948] Shannon, C. E. (1948). A mathematical theory of communication. In Sloane, N. J. A. and Wyner, A. D., editors, *Claude Elwood Shannon Collected Papers*, pages 5–83. Piscataway, NJ: IEEE Press.
- [Thiele et al., 2013] Thiele, I., Swainston, N., Fleming, R. M., Hoppe, A., Sahoo, S., Aurich, M. K., Haraldsdottir, H., Mo, M. L., Rolfsson, O., Stobbe, M. D., et al. (2013). A community-driven global reconstruction of human metabolism. *Nature biotechnology*, 31(5):419–425.
- [Tooley, 1977] Tooley, M. (1977). The nature of laws. *Canadian Journal of Philosophy*, 7(4):667–698.
- [Unterhuber, 2014] Unterhuber, M. (2014). Do ceteris paribus laws exist? a regularity-based best system analysis. *Erkenntnis*, 79(10):1833–1847.
- [van Fraassen, 1989] van Fraassen, B. C. (1989). *Laws and symmetry*. Oxford: Oxford University Press.

- [Weinberg, 1992] Weinberg, S. (1992). Dream of a final theory, the scientist's search for the ultimate laws of nature.
- [Woodward, 2013] Woodward, J. (2013). Laws, causes, and invariance. In S. Mumford, M. T., editor, *Metaphysics and Science*, pages 48–72. Oxford: Oxford University Press.