Humean Laws of Nature: The End of the Good Old Days

November 16, 2020

I show how the two great Humean ways of understanding laws of nature, projectivism and systems theory, have unwittingly reprised developments in metaethics over the past century. This demonstration helps us explain and understand trends in both literatures. It also allows work on laws to "leap-frog" over the birth of many new positions, the nomic counterparts of new theories in metaethics. However, like leap-frogging from (say) agriculture to the internet age, it's hardly clear that we've landed in a good place. My reactionary advice is to return to Hume and work on the central insights that motivated Humeanism about modality in the first place. When updated with contemporary insights, there we will find an attractive naturalistic theory of laws, or so I'll argue, and along the way we'll see how projectivism and systems theory both get something right about this overall theory.

In the Good Old Days philosophers knew where they stood regarding moral realism. Irrealists and realists each came in two varieties. Irrealists could pick between emotivism and error theory. Realists meanwhile had a choice between non-naturalism and a naturalistic ideal observer theory. If someone claimed that cheating is bad, you would ask yourself whether that statement is truth evaluable, and if so, determine what makes it true or false. Your answers located you in conceptual space. Life was simple. Turning to the topic of laws of nature, we are still - naively - living in the Good Old Days. The sepia toned geography is the same. Irrealists can choose between projectivism and a "no laws" counterpart to error theory, and realists divide between non-Humean governing views and Humean systems theories that are counterparts to ideal observer theory. For philosophers of science, life is still simple, the days warm and easy.[1]

[1]The phrase, metaphor and conceptual geography in ethics are all due to the excellent Dreier 2004. Ayer is an example of projectivist/expressivist style accounts of both morality and laws of nature (Ayer 1956). Mackie is an example of a moral error theorist and van Fraassen 1989 might be considered a nomic counterpart. Firth 1952 defends ideal observer theory in ethics and Mill 1843, Ramsey 1990 and Lewis 1973 do in laws. Moore 1903 is a non-reductionist realist about moral properties and Armstrong 1983 is one about nomic properties.

Unfortunately, those days are over in metaethics and I will show that they are also finished for theories of laws of nature. Focusing on theories that find inspiration in David Hume's thought (systems theory and projectivism) I'll begin with a problem for systems theory. I'll demonstrate that the natural resolution of that problem explicitly parallels moves made in metaethics, moves that led to the end of the Good Old Days. Just as metaethics now faces an uncertain future - one where the difference between realism and irrealism is unclear - Humean theories of laws face this same predicament. Life is now complicated. How do we progress? I tentatively suggest that we move forward by looking back, back to the Really Old Days, namely, Hume and what I call the Humean Core.

1 The Best System Theory of Laws of Nature

The best system theory of laws (Mill 1843, Ramsey 1929, Lewis 1973) holds that the laws of nature are a kind of elegant summary of the non-nomic facts of the world. It is Humean in the sense that, like Hume, it denies that there are necessary connections in the world. On this theory, some true generalizations qualify as laws not due to metaphysical facts that these generalizations represent but rather because they express especially powerful compact summaries of the world. David Lewis imagined the set of non-nomic facts as a great mosaic of perfectly natural fundamental properties such as mass and charge distributed across spacetime. The laws, he said, are the axioms that systematize this mosaic while optimally balancing the two virtues of theoretical simplicity and predictive strength.

I suggest that we regard Lewis' formula–a trade-off between simplicity and strength–as merely a first pass at characterizing how science discovers projectable generalizations, not the final word. Science cares about scores of theoretical virtues. Simplicity alone can be understood in dozens of distinct ways that often compete against one another. General relativity contains many more equations than Newtonian gravitation, yet it posits one fewer force - which is simpler?[2] Strength is just as complicated. We should understand the best system as that theory that optimizes whatever metric science actually employs when judging theoretical goodness and not get too bogged down in Lewis' gesture at this metric. Like Hall 2015, I agree that the "central, nonnegotiable idea" behind systems theory is that science's "implicit standards for judging lawhood are in fact constitutive of lawhood."

Is there such a metric? Assuming so is a substantial assumption, one that Feyerabend 1975, for instance, may have denied. But the assumption is defensible. When we look at physics we see this metric hard at work. Quark models were proposed in the early 1970's. Many were empirically adequate. Some were ruled out for not constraining the data enough and others for constraining it too tightly. A delicate balance was sought. We see this balancing act, which is especially clear in curve-fitting, operate throughout science. It is behind the criticism that the Ptolemaic model was too complex and today in the complaint that superstring theory is too unconstrained. Science uses a kind of rough implicit standard in picking theories and the generalizations central to these theories.

[2]See, e.g., Woodward 2014.

Once we accept that there is such a metric function, Humeans have essentially solved the original problem of lawhood, the one bequeathed to us from Hempel 1965. That problem asked us to distinguish intuitively "accidental" true universal generalizations from intuitively "necessary" true universal generations. Syntactically, nothing distinguishes 'all gold spheres are less than 1 mile in diameter' from 'all uranium spheres are less than 1 mile in diameter', even though the first is intuitively accidental and the second intuitively necessary (because uranium is unstable). For positivists with few tools besides syntactic structure, this is a puzzle. But for the Humean who says that the metric actually used is constitutive of lawhood there is a solution. Although the example is somewhat contrived and we must solve it indirectly, we can appreciate that the statement about gold spheres - albeit true and simple - is not a simple corollary of a statement that plays a central role in the theoretical system describing our world whereas the statement about uranium is a corollary of such a statement (ultimately, quantum theory).

Two toy theories may help us present this account as an ideal observer theory.

First, consider the silly theory that declares that the laws of nature are whatever Steven Weinberg, the Nobel Prize-winning physicist, says are the laws. Weinberg uses some implicit metric in his law judgements. He would probably not declare either the gold or the uranium statements as laws, but he will distinguish some true statements as laws and others as not. Hempel's problem is therefore answered. However, as smart as Weinberg is, this theory isn't remotely satisfactory. Weinberg simply doesn't have all the data, nor will he.

Second, consider actual computer programs that "discover" the laws of nature (e.g., Schmidt and Libson 2009). These are programs trained on data from simple systems, e.g., double pendulums. Accurate predictions and compact summaries are rewarded. A genetic algorithm is employed, choosing the best of the failures, modifying, and trying again. Using this method, the (overhyped[3]) Eureka Machine in 2009 produced a conservation law and Newton's Second Law. In the Eureka Machine we have our metric explicitly coded into the program and one can understand the genetic algorithm as an actual best system competition played out in real time. Here we have a timeless program that can handle an indefinite amount of data. Of course, the Eureka program and subsequent ones don't come close to encoding the creativity and insight of Copernicus, Darwin, and Einstein. It's doubtful that we'll ever have an actual program that does because we have no idea how to program the metric science uses.

To make systems theory more vivid, Hall anthropomorphizes the systematizer. Imagine someone who shares the criteria our scientists prize. How would she describe the mosaic if she had perfect understanding of these standards, unlike the Eureka Machine, but also full information about the mosaic, unlike Weinberg? Hall insists that this anthropomorphizing heuristic is entirely dispensable, that it is only a narrative device. That's right. It's just another way of expressing systems theory. The point I want to emphasize is that the "person" behind this heuristic must be hypothetical, an Ideal Observer. The

[3] *The Guardian* reported "'Eureka machine' puts scientists in the shade by working out laws of nature:; see https://www.theguardian.com/science/2009/apr/02/eureka-laws-nature-artificial-intelligence-ai, but the program is essentially one that performs regression analysis.

laws are the axioms the Ideal Observer uses when she systematizes the mosaic. Seen this way, we can appreciate how the laws are objective — they are not relative to Weinberg or Eureka or anyone else – and how they do not share the limitations of an actual person or program. What decides the laws is the best system. That system is the hypothetical system we employ at the ideal limit of science. The Ideal Observer anthropomorphizes this hypothetical element.

2 The Problem of Alien Laws

Framing systems theory via an Ideal Observer helps make obvious specific connections to ethics. Before we get there, however, we must confront a set of problems that have grown over time for systems theory. Putting the theory in terms of an Ideal Observer helps here too, as it better allows us to see that all these individual problems are particular instances of a more general problem, one I'll dub *the problem of alien laws*.

Here is the general problem: wouldn't an Ideal Observer declare as laws propositions that no human scientist would ever find acceptable? Indeed, shouldn't we *expect* an Ideal Observer to declare as laws propositions that are alien to us? By "alien" I mean not merely that these winning propositions might be surprising or counter-intuitive, but worse, that they would not play the roles laws play in actual science, e.g., supporting counterfactuals, prediction and explanation.

To have something to work with, let's stick with the original best system. The Ideal Observer cares only about balancing Lewis' strength and simplicity in a theoretical system. Her goal is to recover as much of the Humean mosaic as possible while optimizing these two virtues. Scrutinizing the properties of our best current candidates for laws, we are led to wonder whether she would produce propositions with these properties. Or would she instead produce something alien to us, something lacking some of the crucial properties we favor in our best current candidates? Summarizing and condensing a vast literature, we can ask, why would the Ideal Observer find laws that...[4]

1. apply to both systems and sub-systems?
2. yield results even when supplied with only initial data?
3. are Markovian?
4. contain a division between initial conditions and dynamics?
5. permit various types of error tolerance?
6. can be approximately solved by tractable mathematics?
7. are different from $(x)Fx$, where F is the predicate true of all and only those non-nomic events that exist?

[4]See Jaag and Loew 2019, Eddon and Meachem 2015, Hicks 2018, Dorst 2019, Ismael 2015, Hall 2015 and more. In these articles one can find some of the questions below and others in the same spirit. Jaag and Loew, Dorst and Hicks all identify the same kind of pattern I do here.

The Ideal Observer's goal is to recover as much as it can about the mosaic while balancing the two virtues of simplicity and strength. Why would it "care" about these other features? None obviously increase simplicity or strength individually or the optimization of these virtues. Let me expand on this point

Regarding 1, the Ideal Observer works with only one system, the grand mosaic. Most propositions we've ever contemplated as fundamental laws - e.g., Hamilton's equation, Schrödinger's equation, Maxwell's equations, Einstein's field equations - have the notable feature of working for systems *and* sub-systems. Classical mechanics describes projectile motion of cannon balls, but also small bits of cannon balls. Quantum mechanics works for water but also hydrogen. Relativity works for the solar system but also the earth-moon system. There is simply no reason for this to occur from a systems perspective. So it's surprising, even suspicious, that what we take to be candidate laws so often work for sub-systems.

Question 2 reminds us that the Ideal Observer is not limited in space or time and yet we are. One way to appreciate the huge difference this makes is by noticing that most candidate fundamental laws are described by hyperbolic partial differential equations (Callender 2017). These equations have the property that we can get something useful out of them (i.e., a non-trivial domain of dependence) by plugging information at a moment into them. We don't need to input temporally non-local data into them. Most types of equations do not have this feature. Instead of what's called a Cauchy Problem, where one puts data across a spacelike hypersurface and marches it forward or backward in time, why shouldn't the Ideal Observer produce a Dirichilet Problem, where one puts the system in a "box" and solves for the inside? That could be very simple and powerful, yet a four-dimensional version of such a problem - as opposed to an actual box - would be utterly unusable for human beings.

Relatedly, question 3 asks why are most of the laws Markovian? A time-dependent system is Markovian if and only if the distribution of future outcomes depends only on the present state of the system. In such systems the present screens off the past, so we do not now need input from the year 1837 to correctly predict the future path of a (say) satellite. This property of the laws is very convenient for creatures like us who lack detailed knowledge of the past (and must pay to store it when we have it). And it's a puzzle why the Ideal Observer should produce such laws given that there are indefinitely many intuitively "simple" equations that are non-Markovian. Thanks to the thermodynamic arrow of time human beings are "stuck in time". Conveniently, the laws tend to accommodate this limitation by being hyperbolic and Markovian.

Question 4 is raised by Hall 2015. Our laws tend to allow free or nearly free initial conditions and only constrain the dynamics (probabilistically or deterministically). We can use the same laws for cannonball trajectories no matter the mass of the ball, angle or direction of the cannon, and so on. Not only does there seem no reason for the Ideal Observer to allow free initial conditions, that seems downright *against the spirit of her job*, which is to increase strength.

Question 5 points to a family of questions. The Ideal Observer needs to recover the mosaic, nothing more. Yet our laws have the curious and again convenient feature of not being too finicky. When sending the Rover to Mars, do we need to input into the

laws the full decimal expansion of its weight? Suppose we only go down to the nearest femptogram (10^{-15}g) after 533 kilograms. Does the Rover go to the Moon rather than Mars? No. Plenty of possible theories, however, are such that small errors in initial data lead to disastrously large differences in solution. Not most of our theories. To put the point sharply, many of our fundamental equations have well-posed initial value formulations. A well-posed problem is one where the mathematical model for the system has three features: the existence of a solution, uniqueness of solution, and the solution varies continuously with the initial data. The first two features make perfect sense for an Ideal Observer because she prizes strength. The third feature is a mystery. For an Ideal Observer, existence and uniqueness should be enough.

Question 6 reminds us that the Ideal Observer is very good at math. Our laws are often hard to work with. Analytic solutions rarely exist. One can barely do any quantum mechanics with the naked Schrödinger equation. Fortunately, they all seem to have the nice property of admitting good approximations by more tractable math. One can plug in, say, the WKB ansatz into the Schrödinger equation and get sensible empirically verified predictions across a large domain, e.g., in quantum tunneling. To have that feature, and others like it, the laws have to have very specific forms, forms that again seem hard to connect to simplicity and strength.

Question 7 is familiar from the philosophical literature and never raised with respect to the other problems, but we can usefully see it as expressing this same problem. Lewis 1986 notices that if we allow *any* type of predicate into a system then $(x)Fx$ will be as strong as possible and as simple as can be. Lewis uses this worry as part of an argument for restricting the lawful predicates to natural kinds. The worry is the same as ours, however, for the concern is that $(x)Fx$ is an alien law. Unlike our laws, it lacks any modal latitude. As a result, if $(x)Fx$ is a law, it's hard to see it playing any of the law roles (supporting counterfactuals, playing a role in explanation, connecting to causation, being used in science by human beings with finite resources, and so on.). As this and many of the examples show, we dearly prize modal latitude but it's hard to understand why an Ideal Observer would.[5]

In sum, it's suspiciously convenient for we human beings that the candidate law statements are so nice to us and have the above features.

In reply one may try to link simplicity or strength to one or more of the above properties. Descartes said that the way God preserves straight line motion is simple because "He always preserves the motion in the precise form in which it is occurring at the very moment when he preserves it, without taking into account of how it was moving [a moment before]" (1644, II, 39). In this passage it sounds like Descartes is linking simplicity to the Markov property. Even if this were plausible, as one goes through the list this type of move becomes more and more a reach. There is no reason why the Ideal Observer, as presently characterized, should be expected to produce such marvelously practicable features when it devises laws.

[5]Real-life examples similar to this can happen when machine learning tools are used to discover laws. Sometimes the laws produced are ones no human being would ever use or want. Presumably these laws are often the result of overfitting the data. That is perhaps another way to put some of the above worries: shouldn't the Ideal Observer overfit the data?

The supercomputer Deep Thought is tasked with computing the Ultimate Question of Life, the Universe and Everything in Douglas Adams' *Hitchhiker's Guide to the Galaxy.* Famously, after 7.5 million years of computation, the computer reveals the answer, 42. What makes this funny is that the answer is so alien to us. We have no idea how 42 could answer the question of the meaning of life. The problem of alien laws is similar. We worry–indeed, we expect–that the Ideal Observer will spit out the counterpart for laws of 42.

3 Metaethical Interlude 1

While these worries about the best system come up here and there, putting them all together in terms of the Ideal Observer helps us see the general pattern. And once one sees the pattern, one quickly realizes that we've previously encountered this worry. It arises in Hume scholarship (see Beebee 2016, Radcliffe 1994, Sayre-McCord 1994) and prominently in metaethics. According to ideal observer theory in ethics (Sidwick 1907, Firth 1952), ethical expressions are claims about the attitudes of a hypothetical observer who is fully informed and rational. Like ideal observer theory for laws of nature, it is a cognitivist (law claims are truth-evaluable) and realist theory.

With the problem of alien laws fresh in mind, the reader can quickly appreciate the comparable problem for ideal observer theory in ethics. In Firth's theory, the ideal observer is omniscient, disinterested, dispassionate, immune to subconscious effects, and perfectly consistent - but otherwise normal!

These features of ideal observers raise epistemological and motivational problems. Epistemologically, the idealization is so drastic that we can have little faith that we'll approximately know what the ideal observer thinks. Real human beings are riddled with inconsistencies, subconscious biases, and passions right to the core. Removing these and making me omniscient, I have little idea of what preferences the ideal observer will want satisfied. We also expect that the good for us will be motivating. Here too we have trouble. Even if we solved the epistemological hurdle and learned what the ideal observer wants, why should what moves this god-like creature also move us? Often I'm inclined to find more information, but clearly the Ideal Observer isn't. The internalist link to motivation seems threatened if the advice is no different than what I might learn from reading a book. As with laws of nature, the idealization process leaves us expecting that the pronouncements of ethical ideal observers will be alien to us.

As a result of these well-known and widely accepted criticisms, Railton 1986 proposed that we replace the Ideal Observer with what is now sometimes called an Ideal Advisor. On this theory, the hypothetical entity's reactions don't determine the good for you, but rather she recommends what is best for non-ideal you. If you imagine the hypothetical being as a kind of guardian angel sitting on your shoulder, the idea is that your good is not determined by what the guardian angel wants but rather by *what she wants for you.* Now the hypothetical entity considers not merely what information you lack and acts rationally, but she also takes account of your psychological traits, motivational system, and even the way you have lived your life. Suppose it is best for you that you believe in

climate science. Studies of climate change deniers show that "more information" – more climate science – isn't always what is needed for someone to change their mind. What climate change deniers sometimes require is having the facts presented in a certain way. Appreciating that need may change the advice to you, which is a feature that the Ideal Advisor but not Ideal Observer may exploit. For Railton, we hold the non-moral features of a person "as nearly constant as possible when asking what someone like him would come to desire." That means that if the Ideal Advisor is to be a normative authority for you, she had better take account of *you*, warts and all.[6]

Before leaving, let me mention another relevant worry raised about ideal observer theory in ethics. Brandt famously pressed another question:

> The facts of ethnology and psychological theory suggests that there could (causally) be two persons, both "ideal observers" in Firth's sense, who would have different or even opposed reactions... with respect to the same act, say on account of past conditioning, as different system of desires, etc. (1955, 26)

Need ideal observers agree? Brandt claims no, arguing that any assumed convergence will beg the question. Theorists appealing to a hypothetical thinker need to decide whether to embrace this relativism to evaluative perspective or not. This question will come up with laws of nature too.

4 Ideal Advisor Theory

The challenges faced by the best system theory, we now see, are essentially counterparts to those faced by Firth's ideal observer theory. The idealization process left Firth's Ideal Observer god-like and not human enough. We expect its recommendations to seem alien to us, and therefore, not motivating. The same is true of the Ideal Observer for laws. An ideal observer or even just an algorithm seeking to balance simplicity and strength would declare as laws propositions that no human scientist would find acceptable. Worse, these laws wouldn't play the law-role that motivated the project in the first place. Just as "42" is not an action guiding answer to the question of the meaning of life, neither are the propositions likely delivered by the Ideal Observer action-guiding in science.

How to answer this worry seems clear. In an amusing parable, David Albert imagines an audience with God where God agrees to tell you about the world. God starts listing all the facts, whereupon

> you explain to God that you're actually a bit pressed for time, that this is not all you have to do today, that you are not going to be in a position to hear out the whole story. And you ask if maybe there's something meaty and pithy and helpful and informative and short that He might be able to tell you about the world which (you understand) would not amount to everything, or nearly everything, but would nonetheless still somehow amount to a lot. Something that will serve you well, or reasonably well, or as well as possible, in making your way about in the world. (Albert 2015, 23)

[6] Whether this move is successful is debatable; see Rosati 1995 and Sobel 1994 for criticism.

God then replies (presumably) with something like the Schrödinger equation, an algorithm with which we finite fallible creatures can make successful predictions. The Ideal Observer has now turned into Ideal Advisor for someone in our predicament. The laws aren't what God Them-self would use, but they are what They would recommend to creatures such as us.

To a Humean, the answer to all our questions 1-7 above is staring at us from the mirror: the laws are partly about us. Laws are useful, and they're laws *because* they're useful.

Once we view the laws as designed for us - given by an Ideal Advisor who knows us warts and all - then the reason why the laws have all of the features mentioned is crystal clear. We always deal with sub-systems in science. These are what matter to us. This chip, that chemical solution, those electrons... In cosmology we sometimes aspire to describe the entire system, but even there we never know that we're dealing with the system as opposed to a sub-system in a larger cosmos. We gather information in spatiotemporally limited regions. If our predictions varied significantly by requiring data from distant spatial locations or from distant times past, we wouldn't be able to use our theories. If they didn't permit various types of error tolerance and allow for mathematically tractable approximations, then again laws would be impractical. We need approximate predictions from equations that are solvable in polynomial time—the shorter, the better. Due to the arrow of time, we don't know the future but want to predict it from whatever state we find ourselves in. That's why we need great latitude in initial conditions and want little latitude in dynamics. We don't know what situation we'll be in, but once we do we want to narrow down what will happen. (The Ideal Observer, by contrast, doesn't suffer from a temporal knowledge asymmetry.) And we care about predicates that we can measure, intervene on, and so forth, not useless gruesome predicates like F.

The key to all of our puzzles is that we see laws as "partially prepared solutions to frequently encountered problems" (Ismael 2015, 197) or Albert's "meaty and pithy and helpful and informative and short" statements that "serve you well, or reasonably well, or as well as possible, in making your way about in the world." We move from an Ideal Observer theory to an Ideal Advisor theory. And that is in fact the way the literature has reacted. Long ago, Earman 1984 advocated thinking of strength as strength *for us*. Elsewhere Cohen and Callender 2009 eliminate predicates like F in Question 7 not for reasons of metaphysics but because such predicates aren't useful to us. Hicks 2019 modifies the best system competition so that it rewards the ability to be confirmed by experiment, thereby helping explain why we have laws that work for subsystems and that divide initial conditions from dynamics. Dorst 2019 changes the desiderata so as to emphasize prediction, producing "principles that are predictively useful to creatures like us." Along the way he shows how this alteration explains a host of features besides 1-7 above (e.g., why we prize symmetries) that together explain why scientists would care about system laws. Jaag and Loew 2019, focusing on these issues but also modal latitude, argue that the criteria for laws must maximize their "cognitive usefulness for creature like us." In each case, the best system competition is modified by adding to strength and simplicity pragmatic criteria, resulting in laws that we might care about.

These moves to make the best system best for us are the counterparts (for essentially

the same reasons) of the changes transforming an Ideal Observer into an Ideal Advisor in ethics. Not all the resulting system theories are the same, of course; for instance, Hicks' ideal observer is perfectly rational but not omniscient. Zooming out, however, they are each moves in a pragmatic direction, moves that bring the Ideal Observer down to earth. As in ethics, I think they're all moves in the right direction.

I want to highlight one further development that also has a counterpart in ethics. As mentioned, Brandt asked: why are we confident that ideal observers would converge on the same propositions? In the case of laws we can similarly ask why we think they would converge on the same set of laws? This question is especially pressing once ideal observers become advisors to the limited creatures that we are. The best advice depends on the audience. If the audience changes, so then should the advice. For laws of nature, the question is sometimes broached by asking whether all scientific communities need discover the same laws we do.

This question gets buried in the details of the systems approach. For example, the best system in its Lewisan formulation demands a preferred language. Lewis insists on perfectly natural properties. Since what is perfectly natural is in principle impossible to know, Cohen and Callender and Loewer 2015 loosen things up and allow an indefinite number of languages. What chooses the language is us and our theorizing, not the world. Since simplicity, strength and their balance are all "immanent" notions - that is, they depend on the predicates being used - it's plausible that different systems formulated in different languages will yield different laws. A system using green and blue may have different laws than one using grue and bleen. If there are no right or wrong languages but only more or less useful ones, then the laws are hostage to what language is pragmatically best. If what is pragmatically best doesn't converge on one language, then one admits that the laws are relative to language and system. Different communities of scientists, using the language that is best for them, might arrive at different laws.

Laws can also become relative to system if one believes that the standards of simplicity, strength, balance, or any other additional theoretical virtue varies with need. The history of science arguably displays change in our standards for a good theory (Doppelt 1978). Maybe this change can happen at a time too? Maybe it depends on the science involved — biology, chemistry, or economics?

Advocates of the best system such as Taylor 1993, Halpin 2003, Callender and Cohen 2009 all advocate making the laws relative to either language or metric or both (for discussion, see Eddon and Meacham 2015). One possible benefit of this move is that it arguably makes understanding the special sciences easier from a systems perspective (Callender and Cohen 2010). In any case, the more pragmatic the theory goes, i.e., the more the features of the audience being advised matters, the more pressure there will be to allow relativity.

In sum, for good reasons, the best system theory has progressed from an ideal observer theory to an ideal advisor theory. In so doing, it also moves closer to its historical rival, projectivism, as we'll now see.

5 Metaethics Interlude 2

Shifting to an ideal advisor theory of laws of nature is a natural transition for the Humean. It more or less solves the problem of alien laws. Once we've made that transition, however, a new parallel with metaethics appears and this time the best response is less clear.

To appreciate this new parallel, I need to sketch a quick potted history of metaethics. Back in the Good Old Days, Hume's thought was developed in two ways. There were non-cognitivist accounts of moral discourse which are the heirs of Hume's internalist tendencies. Think here of Ayer's emotivism, Stevenson's expressivism, and R.M. Hare's prescriptivism. And there were cognitivist accounts that accommodate some of Hume's more externalist claims. One might have naturalistic theories like Firth's ideal observer theory in mind as examples.

These theories were each subjected to many well-known criticisms. Ideal observer theory came under attack by challenges coming from the more "subjective" aspects of our moral language. Meanwhile the emotivist/prescriptivist/expressivist/projectivist strand faced challenges from the more "objective" aspects of our moral language. For instance, we speak as if there is a fact of the matter when two people disagree about a moral claim, e.g., eugenics is bad. We do not shrug such disagreement away as we do when someone likes pickles and someone else doesn't. Also, there are problems in developing the semantics for a non-cognitivist position, e.g., the Frege-Geach problem.

Due to these challenges, each strand of thought developed in sophistication. We saw that ideal observers became ideal advisors. And on the (let's abbreviate this strand to) expressivist side, Gibbard 1993's norm expressivism and Blackburn 1993's quasi-realism were developed. Both expressivists propose semantics that better handle the objectivist functions of moral language than previous versions of the theory. In fact, expressivists felt that their non-cognitivism freed them to do better than cognitivists like Firth. They were able to mimic what a realist about moral properties like Moore would say.

That brought the two historical rivals — Moorean non-naturalism and Humeanism — closer to each other. But if we focus on the Humean views, to the degree that contemporary ideal advisor theories can reproduce what non-naturalists can say, then this result also brings the two Humean views closer together. Railton's ideal advisor takes you and all your features into account when determining the good for you. But Gibbard points out that the rules for expressing yourself morally are much more constrained than had been appreciated. These twin moves left the two Humean theories, expressivism and ideal advisor theory, with little distance between them.

However, it was always still possible to distinguish the positions thanks to the cognitivism versus non-cognitivism divide. For the non-cognitivist, moral statements were not truth-evaluable, did not represent moral properties, and so on. For cognitivists they did, for moral statements expressed truths about what the hypothetical observer would want for you. Then one day expressivists adopted minimalism about truth, representation, propositions and properties and this convenient distinction potentially vanished. *Real* trouble in distinguishing all these historical rivals began, including the two Humean views considered here. This position, expressivism + minimalism, allows expressivists to mimic the language of moral realism. Minimalism entails that anything with content

found in a 'that' clause can be said to be truth apt, represent, and so on. The sentence 'keeping a promise is good' represents that act as good and is true iff that act is good. Minimalism even allows us to say that the property goodness exists. Coupled to minimalism, contemporary expressivists can accept that ethical claims are beliefs that represent mind-independent facts. Blackburn 2015 famously accepts all three defining tenets of Richard Boyd 1988's moral realism.

Dreier 2004 labels this the Problem of Creeping Minimalism. Coming from the other side, Price 2015 dubs it the Problem of Creeping Cognitivism. Who it's a problem for depends on one's default perspective. From outside the debate, it seems a problem for everyone because one now wonders what all the fuss was about in metaethics regarding moral realism. Gibbard, the arch expressivist, announces at the beginning of his book (Gibbard 2003, x [preface]) that he is ambivalent about whether there is an issue at stake or not.

Of course there are replies to the Problem of Creeping, both by Dreier and others. But it's fair to say that there is no widely accepted answer to the Problem of Creeping. Whether the moral realism debate is best described as in a state of *ennui* or *detente* is not clear; what's clear is only that one must resort to French to describe it.

The two great traditions emanating from Hume have been developed to the point where few if anyone can tell them apart. Hybrid views abound: cognitive expressivism, cognitive and non-cognitive sentimentalism, and more. All these views agree that there are no Moorean non-natural properties in the world but that nonetheless contract cheating on university essay assignments is truly very evil. After that they fragment into dozens of views differing mostly over questions about the meaning and function of moral language.

6 Nomic Projectivism

For the final piece of our story, we begin with Ramsey's about-face on laws. Ramsey developed Mill's system theory in 1928, explicitly connecting it to an ideal future scientific theory where we know everything non-nomic. Interestingly, under a year later, Ramsey switches to a form of projectivism about laws of nature. Hints of projectivism can be found in both Hume and Pierce. In Ramsey 1929 this idea is developed. A law is no longer viewed as a summary of events but instead a recommendation about one's confidence in a way of inferring future events. Laws are "not judgments but rules for judging 'If I meet a φ, I shall regard it as a ψ.'" This cannot be negated but it can be disagreed with by one who does not adopt it" (Ramsey 1990, p. 149). The crucial insight we incorporated into ideal observer theories is front and center: laws are guides to the future. Like the warning "prepare yourself, winter is coming," they are not truth evaluable, even if rules exist for their use.

Despite Ramsey's switch, few followed. Projectivism about the nomic is discussed by Blackburn 1986 but is still mostly associated with Ayer's theory. That theory, if it makes an appearance anywhere, is typically found only in undergraduate courses where it is "counter-exampled" to death. Sophisticated developments of projectivist theories of laws are rarely discussed.

But they exist. Inspired by Ayer, Ramsey and Blackburn, Barry Ward develops projectivism in detail through an impressive series of articles (e.g., 2002, 2003). What is interesting about Ward's theory is that he explicitly models it on the most detailed form of ethical expressivism available at the time he was writing, namely, Gibbard's norm expressivism. In Gibbard's theory, when one makes an evaluative judgement, one is accepting a norm that permits or forbids the relevant action. It is an endorsement. Because these norms play social roles, there are rules and logic behind how they function. One goal of such norms is social cooperation. Gibbard modifies possible world semantics to provide a semantics for normative judgements. Using this semantics, he is able to recover the logic underlying most normative thought and language. Ward takes over Gibbard's apparatus wholesale. For him, as for Ramsey, a law of nature is an endorsement, but for Ward the goal of law discourse is not only prediction but also explanation. Saying a generalization is a law is a recommendation that using it will be fruitful to both. The theory is non-cognitivist about language involving laws of nature; but like Gibbard's theory, by focusing on the function of such language and using a modified form of Gibbard's semantics, one can again recover the logic underlying modal discourse.

Note the parallels. For the new "pragmatic" systems theorists, one demands that the Ideal Advisor produce laws that are useful for actual scientists to use in explanation, prediction and experiment. Dorst, Hicks, and Jaag and Loew explicitly build this kind of criterion into the systems view. For the new projectivist, the laws are what you would advise someone to use if they care about explanation, prediction and experiment. There are differences between the two views, but they are not great.[7]

7 Leap-frogging to Creepiness

Sometimes in political science writers speak of countries "leap-frogging" over the Industrial Age. A country, due to its peculiar history and circumstances, might jump more or less straight from an agriculturally dominated society to the internet age, going from cows to Facebook without smoke stacks in between. We can in a similar vein spare the philosophy of laws some toil by leap-frogging to the present situation in metaethics.

In ethics, ideal observer theories and projectivism/expressivism inched closer and closer

[7] Apart from the difference I'm about to discuss (cognitivism), another is that the projectivist doesn't expect Humean supervenience to hold for the *content* of law assertions. Humean laws are summaries of the actual world but describe worlds that aren't actual. We need this modal latitude to use laws. However, many of those lawful possibilities are such that, if they were actual, the most elegant summary of that world wouldn't be what the system theory dubs as lawful. Newton's equations have solutions — say a single particle always traveling inertially — whose simplest and strongest summary are not Newton's equations. The projectivist has no such problem. Suppose, for comparison, that ghosts played an important role in navigating through life. Ghosts almost by definition don't supervene on the natural world. A "projectivist" treatment in which the overly credulous paint the world with such spirits doesn't demand that we identify anything actual with ghosts. It allows a naturalistic understanding of ghosts without an implausible supervenience thesis, e.g., identifying ghosts with creaky noises in the dark. Same with laws. If I endorse using $F = ma$ in making predictions, that doesn't entail that its content supervenes on the actual. See Ward 2002 and Ismael 2015. In what follows we'll assume that the nomic ideal advisor view can overcome this challenge.

to each other for decades. Pushed by the demands of making sense of the more "objective" aspects of our moral language, the former adopted rules of use that mirrored the language of moral realism. Meanwhile, pushed by internalist aspects of morality, ideal observer theory became subjectivist and in some cases even relativist (in response to Brandt), becoming about Ideal Advisors. Throw in minimalism about truth, representation, facts and properties and now few can tell the difference between realists and irrealists anymore.

This same story is unfolding with laws of nature. Suppose that you are in a physics lab doing a spin measurement on some neutrons. The projectivist makes a recommendation: for you, given your goals (prediction, explanation) and circumstances, I recommend using quantum mechanics. The systems theorist likewise says: the ideal advisor, who takes into account your goals (prediction, explanation) and circumstances, advises you to use quantum mechanics. The only substantial difference between the two theories lay in their semantic properties, and in particular, one being non-cognitivist and the other cognitivist. Now add minimalism about the relevant metaphysical and semantic notions to something like Ward's projectivism and mix. The projectivist can then say that it's a fact that Schrödinger's equation is a law, that Schrödinger's equation is a law is true, that the proposition possesses the property of lawfulness, and so on. Just as Blackburn can agree with Boyd on the central tenets of moral realism, so can Ward agree with Lewis and Armstrong on features of laws. The Problem of Creeping equally affects laws of nature as it does moral language. Concentrating on Humean views, it is now very hard to tell the difference between a major form of Humean nomic realism and a major form of Humean nomic irrealism.

We didn't see this problem coming in the case of laws, I suspect, because the internalist, motivational aspects of morality were always front and center whereas these aspects of laws of nature were systematically downplayed. The challenges to systems theory force us to face these features. We just don't care about alien laws. A good recommendation needs to take the audience into account and alien laws don't do this. We expect laws to help us navigate through life. Alien laws don't. As a result, they don't motivate us in any way. In retrospect, we in the metaphysics of science ought to have paid more attention to this aspect of laws. The laws have always been suspiciously kind to us.

8 The New Landscape

We need to sharply distinguish the two issues I've raised.

The first issue is that in meta-ethics, for good reasons, expressivism and ideal observer theory each made moves that brought them very close to one another. I've pointed out that, for similiar good reasons, projectivism and systems theory about laws of nature have also made moves that leave them almost indistinguishable. This point is like Parfit's famous claim that the apparently deep disagreements amongst Kantians, Contractualists, and Consequentialists are not in fact so profound. He views them as "climbing the same mountain on different sides" (2011, 385). The first point is similiar, that projectivists and systems theorists are climbing the same mountain, differing only in certain semantic features.

The second issue, the Problem of Creep, is that minimalism about truth and other semantic properties allows this remaining difference to vanish. This realization happened in meta-ethics and it is a problem for more than only expressivism and ideal observer theory. It is also a challenge one meets in distinguishing a Moorean realism from expressivism, for instance. I've argued that the same Problem of Creep affects the conceptual landscape in laws of nature as much as in morality. This news will be received as no news to global expressivists such as Blackburn and Price, but it should be news to Humeans and others about laws of nature. The problem of alien properties suggested shifts that make it apparent that systems theorists and expressivists are climbing the same mountain from different sides. Yet one could still distinguish the views due to their different semantic properties. The Problem of Creep removes that ability. Keeping the mountain metaphor, the Problem of Creep places us at the summit where the climbers merge into one.

How should we respond to these two issues?

Taking the second first, note that I have *not* argued that we *should* adopt minimalism about truth. I am not saying that we are at the summit just described. I have only pointed out that *if* one is a minimalist, then philosophy of laws must face the same problem that meta-ethics currently does.

In meta-ethics it's not so clear how to respond to the Problem of Creep. Naturally, one reaction has been for metaethicists to find new ways to draw the line between realism and irrealism. Even if moral realists and irrealists can say the same things, moral realists may demand types of explanation unacceptable to moral irrealists. Or perhaps inference patterns will reveal a difference. The literature contains a proliferation of explanationist and inferentialist ways of distinguishing the two views (see Dreier 2018 for a survey and references). Unfortunately, the different ways of drawing a line between realism and irrealism tend not to agree with one another. None seems obviously best. Dreier 2018 argues that there is no One True Distinction, and he adopts an "irenic and pragmatist perspective, allowing that different ways of drawing the line are best for different purposes."

If I'm right that the analogy is strong, then we can expect philosophers of science keen to maintain a difference between projectivism and systems theory to also find new lines of differentiation between the two. If the parallels hold up, as I think they will, however, we should not expect to find the One True Distinction in laws of nature either.

Another response is to "move on" past cognitivism and non-cognitivism, representationalism and non-representationalism. I don't know precisely how to characterize this path due to its diversity, but the result seems to be articles in metaethics with titles that sound like self-help guides to bad break-ups. One other reaction in this neighborhood would be to embrace a global expressivism such as that held by Blackburn and Price. None of these moves will deliver a clear difference between realism and irrealism in ethics, nor can they be expected to do so between Humean realism and irrealism.

Of course one can deny minimalism about truth and avoid the Problem of Creep. As I said, I have not argued that we ought to embrace mininalism so this move is perfectly fine for all I have said. However, it does leave us with the first problem. And it means that the principal difference between realism and irrealism hangs on the correct theory of truth and representation, not anything specifically about laws. That is

a very unwelcome result for someone who felt that projectivism or systems theory was correct due to reflection on science and its laws of nature. The moves toward pragmatism among systems theorists and the moves toward norms of law-talk among projectivists took the views so close that only semantic differences like truth-evaluability remained, differences that Creep potentially obliterates. Rejecting minimalism about truth doesn't resolve the first problem raised by this paper, namely, learning that projectivists and systems theorists are climbing the same mountain.

Like leap-frogging to the internet age, it's not clear that we've landed in a good place. We can take solace in skipping the painful birth of many sophisticated new positions and directions, yet it's not clear how Humeans should now understand laws of nature.

9 Laws as Negotiated Settlements

Philosophy papers are often convincing when they deliver bad news and go awry when they try something positive. I don't know how to respond to the Problem of Creep – so this paper will end safely negative on this point. Risking the fate of going positive, in the space remaining I want to suggest a way of thinking about the "two sides of the same mountain" problem. I think we can gain insight on the "two sides" problem by re-examining some issues that arise in Hume's work–by going back to the Really Good Old Days. When we do, we'll see that drawing lines between Humean theories of laws and declaring a winner seems less interesting than it was before. We can also see some moves made that will help us in our current predicament.

Hume's theory of natural necessity is very rich. He is famous for his skeptical attack on necessity. Necessity is found, he says - or most commentators say he says - in us and not the world. Hume does not rest with skepticism. He also provides a great origin story of why creatures like us would manufacture (natural) modality in a world lacking naked modal facts. Hume's insight is that modal reasoning is highly adaptive or "fitness" enhancing given the predicament we're in. We live in an uncertain and risky world. We're cognitively and perceptually limited creatures who receive no information from the future. We don't know what will happen next, but if we're to survive and thrive we'd better have some guidance. We dearly need to predict what will happen next, prepare for it, and possibly intervene. For this the actual world is not enough. Our epistemic limitations and practical deliberative contexts require us to have theories about what is non-actual, just in case what we think is non-actual turns out to be actual (for attempts to begin to spell this out, see Ismael 2015, Strevens 2007).

Hume famously provides a theory of how all this works with causation. The source of necessity is a psychological faculty, the feeling of expectation. When we meet an event of type F, we come to expect one of type G because we've witnessed events like G follow events like F many times before. The feeling of expectation plays a key role: it gives us the ability to predict, prepare for, and possibly intervene on G by making F obtain. The expectation enhances our "fitness." And he also tells us how this works, claiming that the mind employs certain "rules" for judging cause and effect, e.g., that like causes regularly precede spatiotemporally contiguous like effects.

Notice that this Humean theory can be updated to contemporary times. In outline, he proposes that we have a psychological faculty that follows some rules in taking as input some non-nomic facts (i.e., regularities) and yields as output a type of mental state (i.e., expectations) that is justified by it usefulness. Recent work in casual learning theory and developmental psychology offer psychological faculties quite different than Hume's state of expectation. One might, for example, replace the feeling of expectation with a special type of cognitive map, namely, a causal map, which is a representational system that maintains an updated representation of the causal structure of one's environment (Carey 1985; Gopnik et al 2004; Gopnik and Meltzoff 1997; Wellman 1990). Work in causal methodology might substitute Hume's "rules" with more sophisticated theories such as structural causal models (Pearl 2000). But the original rationale is still very much the same: these rules and faculties help us get by. The general picture is one wherein the brain is a prediction engine. To play its role, given the knowledge asymmetry (that we know "more" about the past than future), it creates various models of the future and runs them forward using rules it has found useful in the past.

This core Humean picture is very compelling. There is a beautiful consilience between theory and empirical work in it, and it has the advantage of parsimony because it needn't attribute necessary connections to the world itself. But where do laws of nature fit into this?

Hume doesn't really focus much on laws as opposed to causes (except when discussing miracles).[8] For Hume they are empirical generalizations formed from the patterns of covariation we encounter, but not a lot is said. I want to suggest a way of thinking of laws in Hume's picture that falls out from a problem his theory of causation faces, a problem that is basically the same as the one we've just encountered. The problem is one of mismatches between the deliverances of our psychological system – understood as feelings of expectation or as modern causal maps – and our considered judgements about causation.

To succeed in the world, we need to be able to change our causal maps. When we gather new data or come to better understand the rules by which causation is derived, we need to be able to learn and update our maps. Alignment with others is also important. In a social environment, it's crucial that we share expectations about what will happen. It becomes important to bring my causal map in line with yours. Early hunters could only bring down prey if they shared similar causal maps. But not all causal maps are equal. There are many pressures to change one's causal map. Some are better than others. What does "better" mean for the Humean?

Hume faces this question throughout his corpus, and especially in his theories of causation, aesthetics and morality. Causation is an expectation formed by seeing constant

[8]Hume does recognize that

> the utmost effort of human reason is, to reduce the principles, productive of natural phaenomena, to a greater simplicity, and to resolve the many particular effects into a few general causes, by means of reasonings from analogy, experience, and observation. (E 4.12)

So he sees value, like modern Humeans, in devising generalizations from which one can get a lot from a little.

conjunction for Hume. But no one person is likely to observe all instances of F's and G's. Maybe they're constantly conjoined only in my set of observables. Maybe they're not even conjoined there but due to perceptual error it seems like they are. Nonetheless the psychological feeling of expectation may arise. We agree that some can be experts in taste. Nonetheless, it might be that some low art pleases me, e.g., zombie films. And my moral sentiments vary with whether people are near or far, family or not, even though when we make moral judgments we agree that "virtue in rags is still virtue" (T 584). In all of these cases our psychological faculties may deliver responses that depart from some standard.

To deal with the mismatches between the deliverances of our psychological faculty (e.g., expectation, approval, pleasure) and our judgements, Hume often appeals to a "general point of view." In resolving the mismatch between our steady moral discriminations and our less fixed actual sentiments, Hume writes:

> Our situation, with regard both to persons and things, is in continual fluctuation; and a man, that lies at a distance from us, may, in a little time, become a familiar acquaintance. Besides, every particular man has a peculiar position with regard to others; and 'tis impossible we cou'd ever converse together on any reasonable terms, were each of us to consider characters and persons, only as they appear from his peculiar point of view. In order, therefore, to prevent those continual contradictions, and arrive at a more stable judgment of things, we fix on some steady and general points of view... (T 581-82).

A general point of view is a "method of correcting our sentiments, or at least of correcting our language, where the sentiments are more stubborn and inalterable" (T 582). Elsewhere he writes that we need "to correct these inequalities [the above fluctuations] by reflection, and retain a general standard of vice and virtue, founded chiefly on general usefulness (E 229, n. 1).

The "general point of view" motivates some commentators (e.g., Rawls 1971, 183-92) to interpret Hume as espousing a kind of ideal observer theory (Radcliffe 1994). Moral and aesthetic judgements are based on the sentiments of someone occupying the general point of view. For causation, Hume defines causation as

> an object precedent and contiguous to another, and so united with it, that the idea of the one determines the mind to form the idea of the other, and the impression of the one to form a more lively idea of the other (T 1.3.14.31/170)

Garrett 1997 points out that we can give "the mind" in the above a subjective or an "idealized" spectator interpretation. If the latter, we would be identifying "true" causation with the expectation that occurs in a hypothetical mind, a mind that accurately perceived all the relevant constant conjunctions and perfectly followed the rules. Put in terms of our contemporary example, we can make sense of "better" in terms of an objectively best causal map. This map portrays the links between variables that one would draw if one perfectly followed the rules (say, Pearl's structural causal models) and had all the facts.

Making sense of a 'best causal map' leads us to speak of hypothetical beings who are fully informed and rational. The best causal map is the map that such a being has in her head.

Note that this question — whether there is in the limit a best causal map — is a very esoteric one. It is essentially the question of whether there is a limit to science. I don't know of good arguments for thinking there is or isn't such a limit. In any case we've seen where this route will take us. These hypothetical "in-the-limit" causal maps will have little or nothing to do with psychological faculties in individuals and or the modifications made to them through learning. And since both are used in our predictions, preparations, and interventions, the hypothetical causal maps will have little to do with helping us navigate through life. They will be alien.

Sayre-McCord 1994 urges us *not* to interpret Hume's general point of view in morality and aesthetics as the perspective of an idealized observer. He understands Hume as envisioning a kind of Hobbesian jungle of diverse sentiments and the general point of view as a kind of negotiated settlement that smoothes away inconsistencies amongst them. For this to work, the standard provided by the general point of view must be salient, mutually accessible, and tend toward stable consensus (217). By salient, he means that the standard must engage the sentiments. I can bring myself to see that a person from outside my narrow circle is virtuous by imagining their act done in my circle. When I do, my sentiments are triggered. Mutual accessibility means that you can do the same. And if together we can form a consensus that selfless sacrifice for one in need is virtuous abstracted away from one's circle, then we will have a standard that irons out the bumps in the peculiarities of our varying sentiments.

Transferring this picture to natural modality, we can agree with Gopnik 2000 when she writes that "Science simply puts these universal and natural capacities [to detect causation] to work in a socially organized and institutionalized way." Science becomes a kind of Hobbesian civil society (minus dictator) that avoids a war of all against all. It tries to reconcile the many conflicting and changing causal maps we all have, settling on ones that serve our goals — of the lab, of the field, of the public, and so on. On this picture there is no final correct causal map, no guarantee that science settles on the "right" one. Yet there are ongoing negotiated standards for how to reconcile conflict — look to experiment, and if experiment can't decide, turn to theoretical virtues like simplicity, fruitfulness, and unification.

If we adopt this picture — which I wish I had space to develop — then what do we say about laws of nature? Laws of nature are the product of this social organization and institutionalization, particularly compact and powerful ways of saying a lot about the consensus-best causal maps in the relevant fields. They are the projectible generalizations that emerge from this Hobbesian jungle of vying causal maps. They play an important role, for they are the recommendations on which our sciences have achieved consensus. But metaphysically laws are not so interesting: they are the somewhat imprecise results of negotiation among our individual causal maps. As with the products of any messy social negotiation, they are bound to be somewhat loose and contested. This is the reason why science doesn't fuss too much about annointing one package as nomic when there are competing systems–for instance, choosing between Schrodinger's wave mechanics and

Heisenberg's matrix mechanics, or figuring out *precisely* which propositions are the laws of evolutionary biology or general relativity. As with legal laws, consensus is difficult to achieve, so further negotiation is only warranted when it is matters a great deal to achieving our goals.

We can now appreciate what both Humean theories got right. Focusing on our psychological capacities, e.g., expectations, causal maps, we can understand why projectivism is tempting. The source of Humean modality lay in our psychological faculties and the laws are what science has achieved consensus on as its best recommendation. Projectivism nicely captures these aspects of laws. But focusing on the later improvements, organization and consensus building regarding our causal maps, we can see why a systems theory is appealing. The best system competition represents the Hobbesian battle to determine what causal map is best. Both theories get something right. Does the metaphysician or philosopher of science really need to decide between the two? No, for neither tells the full story. Each is only climbing half of the mountain.

Stepping back, in metaethics whether moral judgments are literally true or false might seem (or have seemed) like a life and death matter. In the laws debate the situation has never been like that. What's important is the origin and purpose of modal discourse: why do we engage in modal discourse if the world is fundamentally amodal, just one thing after another? Given this focus, it seems that philosophers of science should be less worried than meta-ethicists about creepiness and about from which side to climb the mountain. The Problem of Creep perhaps demonstrates that it is even less worth having this debate than it initially seemed. It might not matter whether realism or irrealism about (Humean) laws is true. What matters is developing the origin and purpose stories for modal discourse. Laws of nature are then a kind of almost optional late-stage wrapping up of this developoment, one connecting the story to practice in scientific. Here it turns out that the two camps (realist and iirealist Humeans) have so far approached things in slightly different ways, but the two approaches in fact complement each other nicely.[9]

References

[1] Albert, David. 2015. *After Physics*. Cambridge: Harvard University Press.

[2] Ayer, A. J. 1956. "What is a Law of Nature?" *Revue Internationale de Philosophie* 10, 144–65.

[3] Armstrong, David. 1983, *What Is a Law of Nature?*, Cambridge: Cambridge University Press.

[9] My longtime colleague, Nancy Cartwright, reflecting on my work on the best system, said to me (I paraphrase), "Craig, for someone who doesn't really believe in them, you care a lot about laws." In some ways this essay is me finally working through this tension. Many thanks for comments to Nancy Cartwright, Eddy Chen, Jonathan Cohen, Jamie Dreier, Michael Hicks, Carl Hoefer, Siegried Jaag, Barry Loewer, Elizabeth Miller, Markus Schrenk, Elliott Sober, the philosophers at Ranch Metaphysics 2020, and especially Elanor Cranor and Christian Löw.

[4] Beebee, Helen. 2016. "Hume and the Problem of Causation", in P. Russell (ed.), *The Oxford Handbook of Hume*, NY: Oxford University Press.

[5] Brandt, Richard, 1955. "The Definition of an "Ideal Observer" Theory in Ethics", *Philosophy and Phenomenological Research* 15, 408.

[6] Blackburn, Simon. 1993. *Essays in Quasi-Realism*. Oxford: Oxford University Press.

[7] Blackburn, Simon. 1986, "Morals and Modals," in *Fact, Science and Morality*, G. Macdonald and C. Wright (eds.), Oxford: Basil Blackwell.

[8] Blackburn, Simon. 2015. "Blessed are the Peacemakers." *Philosophical Studies* 172, 843–853.

[9] Boyd, Richard. 1988. "How to Be a Moral Realist." In *Essays on Moral Realism*, edited by Geoffrey Sayre-McCord, 181–228. Ithaca, NY: Cornell University Press.

[10] Callender, Craig. 2017. *What Makes Time Special?* Oxford: Oxford University Press.

[11] Callender, Craig. and Cohen, J. 2010. "Special Sciences, Conspiracy and the Better Best System Account of Lawhood". *Erkenntnis* 73, 427–447.

[12] Carey, Susan. 1985. *Conceptual Change in Childhood*. Cambridge, Mass: MIT Press.

[13] Cohen, Jonathan and Callender, Craig. 2009. "A Better Best System Account of Lawhood", *Philosophical Studies* 145, 1–34.

[14] Descartes, Rene. 1644. *Principles of Philosophy*. In J. Cottingham , R. Stoothoff, and D. Murdoch (eds. and trans.), *The Philosophical Writings of Descartes* (vol. 1). Cambridge University Press, 1985.

[15] Doppelt, Gerald. 1978. "Kuhn's Epistemological Relativism: An Interpretation and Defense", *Inquiry* 21, 33-86.

[16] Dorst, Christopher. 2019. "Toward a Best Predictive System Account of Laws of Nature", *The British Journal for the Philosophy of Science* 70, 877–900.

[17] Dreier, Jamie. 2004. "Meta-Ethics and The Problem of Creeping Minimalism." *Philosophical Perspectives* 18, 23–44.

[18] Dreier, Jamie. 2018. "The Real and the Quasi-real: Problems of Distinction", *Canadian Journal of Philosophy* 48, 532-547

[19] Earman, John. 1984, "Laws of Nature: The Empiricist Challenge," in D. M. Armstrong, R. Bogdan (ed.), Dordrecht: D. Reidel Publishing Company.

[20] Eddon, Maya and Meacham, C.J.G. 2015. "No Work for a Theory of Universals", In Jonathan Schaffer & Barry Loewer (eds.), *A Companion to David Lewis*. Wiley-Blackwell, 116-137.

[21] Feyerabend, Paul. 1975. *Against Method: Outline of an Anarchistic Theory of Knowledge*.

[22] Firth, Roderick. 1952. "Ethical Absolutism and the Ideal Observer", *Philosophy and Phenomenological Research* 12, 317–345.

[23] Garrett, Donald. 1997. *Cognition and Commitment in Hume's Philosophy*. Oxford: Oxford University Press.

[24] Gibbard, Alan. 1990. *Wise Choices, Apt Feelings: A Theory of Normative Judgment*, Harvard University Press & Oxford University Press.

[25] Gibbard, Alan. 2003. *Thinking How to Live*. London: Harvard University Press.

[26] Gopnik, Alison, Clark Glymour, David Sobel, Laura Schulz, Tamar Kushnir, and D. Danks. 2004. "A Theory of Causal Learning in Children: Causal Maps and Bayes Nets", *Psychological Review* 111, 3-32.

[27] Gopnik, Alison. 2000. "Explanation as Orgasm and the Drive for Causal Understanding: The Evolution, Function and Phenomenology of the Theory-formation System", in F. Keil & R. Wilson (Eds.) *Cognition and Explanation*. Cambridge, MA: MIT Press, 299-323.

[28] Hall, Ned. 2015. "Humean Reductionism about Laws of Nature", in B. Loewer & J. Schaffer (eds.), *The Blackwell Companion to David Lewis*, 262–277. Oxford: Blackwell.

[29] Halpin, John. 2003. "Scientific Law: A Perspectival Account." *Erkenntnis* 58, 137-168.

[30] Hare, R. M. 1952. *The Language of Morals*. Oxford: Clarendon Press.

[31] Hempel, Carl. 1965. *Aspects of Scientific Explanation and Other Essays in the Philosophy of Science*. New York: The Free Press.

[32] Hicks, Michael. T. 2018. "Dynamic Humeanism", *The British Journal for the Philosophy of Science* 69, 983–1007.

[33] Horgan, Terry. and Timmons, Mark. 2006. "Cognitivist Expressivism," in Terry Horgan and Mark Timmons (eds.), *Metaethics after Moore*, Oxford: Oxford University Press, 255- 298.

[34] Hume, David. 1975. *An Enquiry Concerning the Principles of Morals*, edited by L. A Selby-Bigge, 3rd ed. Oxford: Clarendon Press.

[35] Hume, David. 1978. *A Treatise of Human Nature*, edited by L. A. Selby-Bigge, 2nd ed. Oxford: Clarendon Press.

[36] Ismael, Jenann. 2015. "How to be Humean," In B. Loewer & J. Schaffer (Eds.), *The Blackwell Companion to David Lewis*, 188–205. Oxford: Blackwell.

[37] Ismael, Jenann. 2017. "An Empiricist's Guide to Objective Modality", in M. Slater, and Z. Yudell (eds), *Metaphysics and the Philosophy of Science: New Essays*, New York: Oxford University Press, 109–126.

[38] Jaag, Siegfried and Loew, Christian. 2018. "Making Best Systems *Best for Us*", *Synthese*, 1-26.

[39] Lewis, David K. 1973. *Counterfactuals*. Oxford: Blackwell.

[40] Lewis, David K. 1983. "New Work for a Theory of Universals", *Australasian Journal of Philosophy* 61, 343–377.

[41] Lewis, David K. 1986. *Philosophical Papers Vol. II*. Oxford University Press.

[42] Loewer, Barry. 1996. "Humean Supervenience", *Philosophical Topics* 24, 101–127.

[43] Loewer, Barry. 2007. "Laws and Natural Properties", *Philosophical Topics* 35, 313–328.

[44] Mill, John Stuart. 1843. *A System of Logic, Ratiocinative and Inductive*. Cambridge University Press.

[45] Michotte, Albert. 1946. *La Perception de la Causalité*. Paris: Vrin.

[46] Moore, G.E., 1903. *Principia Ethica*, NY: Cambridge University Press.

[47] Parfit, Derek. 2011. *On What Matters*. Oxford University Press.

[48] Pearl, Judea. 2000. *Causality: Models, Reasoning, and Inference*. Cambridge University Press.

[49] Price, Huw. 2015. "From Quasi-Realism to Global Expressivism – And Back Again?" In *Passions and Projections*, edited by Robert N. Johnson and Michael Smith, 134–152. Oxford: Oxford University Press.

[50] Radcliffe, Elizabeth. 1994. "Hume on Motivating Sentiments, the General Point of View, and the Inculcation of "Morality"", *Hume Studies* 20, 37-58.

[51] Railton, Peter. 1986. "Facts and Values." *Philosophical Topics* 14, 5–31.

[52] Ramsey, Frank P. 1990. *Philosophical Papers*. Edited by D.H. Mellor. Cambridge University Press. [Includes the 1928 note "Universals of Law and of Fact" and the 1929 "General Propositions and Causality"]

[53] Rawls, John. 1971. *A Theory of Justice*. Harvard University Press.

[54] Roberts, John. 1999. "'Laws of Nature' as an Indexical Term: A Reinterpretation of Lewis's Best-System Analysis." *Philosophy of Science* 66 (3): S502-S511.

[55] Rosati, Connie. 1995. "Persons, Perspectives, and Full Information Accounts of the Good." *Ethics* 105, 296–325.

[56] Sayre-McCord, Geoffrey. 1994. "On Why Hume's "General Point of View" Isn't Ideal– and Shouldn't Be", *Social Philosophy and Policy* 11, 202-228.

[57] Schmidt, Michael and Lipson, Hod. 2009. "Distilling Free-Form Natural Laws from Experimental Data", *Science* 324, 81-85.

[58] Sidgwick, Henry. 1907. *Methods of Ethics*. 7th ed. Macmillan.

[59] Sobel, David. 1994. "Full Information Accounts of Well-Being", *Ethics* 104, 784–810.

[60] Stevenson, Charles. 1944. *Ethics and Language*. Yale University Press.

[61] Strevens, Michael. 2007. Why Represent Causal Relations? In A. Gopnik and L. Schulz (eds.), Causal Learning: Psychology, Philosophy, Computation, Oxford University Press, New York.

[62] Taylor, Barry. 1993. "On Natural Properties in Metaphysics." *Mind* 102, 81– 100.

[63] Van Fraassen, Bas. 1989. *Laws and Symmetry*. Oxford: Oxford University Press.

[64] Ward, Barry. 2002, "Humeanism without Humean Supervenience: A Projectivist Account of Laws and Possibilities," *Philosophical Studies* 107, 191–218.

[65] Ward, Barry. 2003. "Sometimes the World is not Enough: The Pursuit of Explanatory Laws in a Humean World" *Pacific Philosophical Quarterly* 84, 175-197.

[66] Wellman, Henry. 1990. *The Child's Theory of Mind*. Cambridge, Mass.: MIT Press.

[67] Woodward, James. 2014. "Simplicity in the Best Systems Account of Laws of Nature", *British Journal for the Philosophy of Science* 65, 91–123.