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

12,506 words

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]

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 can be traced to J.S. Mill in 1843 and was then developed by F.P. Ramsey in 1929 (Ramsey 1990). Ramsey quickly discarded it. So it was left mostly untouched until David Lewis absorbed it into his grand metaphysical system in 1973. The theory remained unchanged again for decades. Recently, however, it has been the subject of an unprecedented amount of attention. This work illuminates new options for the theory, solves old problems, and provides a sophistication lacking in the theory's first century and a half of life. Let me briefly review how I understand it. For what comes later, it will be helpful for me to put it in its less familiar ideal observer theory format.

The best system theory of laws 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. Famously Lewis imagined the set of non-nomic facts as a great Humean mosaic of perfectly natural fundamental properties such as mass and charge distributed across spacetime. The laws, he said, are the axioms of the deductive system that best balance simplicity and strength while systematizing this mosaic.

Lewis' best system balances the trade-off between simplicity and informativeness. I regard this as merely a first pass at characterizing how science discovers projectable

[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.

generalizations, not the final word. In scientific theorizing, we care about scores of theoretical virtues. Simplicity alone can be understood in dozens of distinct ways that often compete against one another. General relativity posits 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 didn't constrain the data enough and were rejected; some constrained too much and were rejected. A delicate balance was sought. This balancing act goes right back to the beginnings of the field where we find criticism the Ptolemaic model's complexity and the balancing continues today with the complaint that superstring theory is too unconstrained. The same is true across all science, e.g., in curve-fitting. Science uses a kind of rough implicit standard in picking theories and the generalizations that are central in these theories.

Once we admit that there is such a metric function, Humeans have essentially dissolved 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. The answer is solved by the best system theory because the metric we use actually does distinguish the two types of claims. Therefore the result of the competition will see a distinction and output different verdicts. Although the example is somewhat contrived, we can appreciate that the statement about gold spheres - albeit true and simple - may not play as central a role in a theoretical system as the statement about uranium.

How do we get to an ideal observer account from this theory? Two toy theories may help.

First, consider the silly theory that declares that the laws of nature are whatever Steven

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

Weinberg, the Nobel Prize-winning physicist, declares 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 doesn't have all the data, lacks a perfect understanding of science's standards, and we don't want the search for the laws to die with him.

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. Here one sees a 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. But of course, the Eureka program and subsequent ones don't come close to encoding the creativity and insight of Copernicus, Darwin, and Einstein, and it's doubtful that we'll ever have an actual program that does.

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 but also full information about the mosaic, unlike either Weinberg or the Eureka Machine? 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 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.

[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.

2 The Problem of Alien Laws

Framing systems theory via an Ideal Observer will help 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? 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, they would not play the roles laws play in actual science, e.g., supporting counterfactuals, playing a role in prediction and explanation.

To have something to work with, let's suppose 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 like in our best current candidates? Summarizing and condensing a vast literature, we can ask, why would the Ideal Observer find laws that...[4]

1. work not only for systems but also sub-systems?
2. yield results even with a small amount of local input?
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 nonnomic 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, Schrodinger'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.[5] So it's surprising, even suspicious, that what we take to be candidate laws so often work for sub-systems. There is simply no reason for this to occur from a systems perspective.

Question 2 is really a large family of questions. The Ideal Observer is not limited in space or time, but our laws succeed with modest input. Most of the fundamental ones are described by hyperbolic partial differential equations. Plug in a small amount of initial data and they make predictions about a huge ever expanding cone-like swathe of the mosaic later on. We don't need information across all of space at an instant for the laws to give us back something very useful. 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.

For that matter, why would the Ideal Observer make a distinction between initial conditions and dynamics? This is Question 4 and one raised by Hall 2015. Our laws tend

[5]Interestingly, this doesn't seem to be the case for most special science laws.

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 her to allow free initial conditions, that seems downright against the spirit of her job.

Question 5 is again a family of questions. The Ideal Observer needs to recover the mosaic, nothing more. Yet our laws have the feature that they're not too finicky. When sending the Rover to Mars, do we need to input 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, note that for many of our fundamental equations, we have well-posed initial value problems. 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 - she gets all A's in every exam. Our laws are often hard to work with. Analytic solutions rarely exist. One can barely do any quantum mechanics with the naked Schrodinger 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 Schrodinger 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 many of the examples show, e.g., the point about error tolerance, we dearly prize modal latitude but it's hard to understand why an Ideal Observer would.

In Douglas Adams' *Hitchhiker's Guide to the Galaxy* the supercomputer Deep Thought is tasked with computing the Ultimate Question of Life, the Universe and Everything.

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 be an answer to the meaning of life. The problem of alien laws is similar. We worry — indeed, in some cases we ought to expect — that the Ideal Observer will spit out the law-version of 42. If $(x)Fx$ is the law, it's hard to see it playing any of the law roles (supporting counterfactuals, playing a role in explanation, connecting to causation, actually being used in science by human beings with finite resources, and so on.). Even if we rule out the $(x)Fx$ example for some reason, still there are plenty of on balance simple and strong equations that are just as alien to us.[6]

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. True, one could push back and 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). Perhaps that provides a link between simplicity and the Markov property. As one goes through the list, however, that move becomes more and more a reach. One might also learn to live with these features, regarding them as brute coincidences, elements of fine-tuning. In any case, there is no reason why an ideal observer, as presently characterized, should be expected to produce such marvelously practicable features when it devises laws.

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 seen this worry before. It comes up in Hume scholarship (see Beebee 2016, Radcliffe 1994, Sayre-McCord 1994) but it also arises more 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 and realist theory.

With the problem of alien laws fresh in mind, the reader can quickly appreciate the

[6] Real-life examples of 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?

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 no faith that we'll even approximately know what the ideal observer thinks. The idealization process leaves us expecting that the ideal observer's pronouncements will be alien to us. Motivationally, 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. In addition, the internalist link to motivation seems threatened if the advice is no different than what I might learn from reading a book.

As a result of these criticisms, Railton 1986 proposed that we replace the Ideal Observer with what is now sometimes called an Ideal Advisor. On this theory, the ideal 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. Climate change deniers don't always need more climate science. Some have a greater need to have the facts presented in a way to which they are receptive. Appreciating that need may change the advice to you and therefore alter what you should do. 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.[7]

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.

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

4 Ideal Advisor Theory

The challenges faced by the best system theory, we can 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)

God then replies (presumably) with something like the Schrodinger 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 would Himself or Herself use, but they are what He or She 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, then, 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 Callender and Cohen 2009 eliminate predicates like F by pointing out that they 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 laws' 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.

Clearly, these moves to make the best system best for us are essentially the same as the move from ideal observer theories to ideal advisor theories in ethics. Not all the 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, one that brings the ideal observer down to earth.[8]

I want to highlight one further development in laws that also has a counterpart in

[8]Note, incidentally, that Hall's Ideal Observer is more Ideal Advisor than Ideal Observer.

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 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 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 come up with different laws.

Laws can also become relative to system if one believes the standards of simplicity, strength, balance, and any other theoretical virtue you care about vary with need. The history of science arguably displays change in our standards for a good theory (Doppelt 1978). Maybe this variety 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 2009). 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, the best system theory has largely progressed from using an ideal observer to an ideal advisor, and in so doing, the theory 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. So why aren't we wrapping up the paper now? Unfortunately the parallel with metaethics continues and this time results in a more confusing situation.

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 make sense of some of Hume's more externalist claims. One might think of various instances of naturalism, e.g., Firth's ideal observer theory, as examples.

These theories were each subjected to many well-known criticisms. We've just witnessed how 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. Gibbard 1993's norm expressivism and Blackburn 1993's quasi-realism are examples. Both propose semantics that better handle the objectivist functions of moral language than previous versions of the theory. That brought the two historical rivals — ideal observer theory and (let's abbreviate this strand to) expressivism — closer to each other in spirit. Railton's ideal advisor really 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. The two theories approached one another, but it was still possible to distinguish the positions thanks to the cognitivism versus non-cognitivism divide.

Then one day expressivists adopted minimalism about truth, representation, propositions and properties. *Real* trouble in distinguishing these historical rivals began. This position, expressivism + minimalism, allows expressivists to mimic perfectly 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, is true iff that act is good, and 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 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. Here is Gibbard, the arch expressivist, announcing at the beginning of his book that he is ambivalent about whether there is an issue at stake or not:

> Does this mean that there are no facts of what I ought to do, no truths and falsehoods? Previously I thought so, but other philosophers challenged me to say what this denial could mean. In this book, I withdraw the denial and turn non-committal. In one sense there clearly are "facts" of what a person ought to do, and in a sense of the word 'true' there is a truth of the matter. That's a minimalist sense, in which "It's true that pain is to be avoided" just amounts to saying that pain is to be avoided - and likewise for "It's a fact that". Perhaps, as I used to think, there are senses too in which we can sensibly debate whether ought conclusions are true or false. Nothing in this book, though, depends on whether there is any such sense. (Gibbard 2003, x [preface])

Of course there are replies to the Problem of Creeping, both by Dreier and others. But the meta-metaethical worry is now so widespread that it merits a section of an entry of the *Stanford Encyclopedia of Philosophy*. My sense is that 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 out there in the world and 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 can 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 for winter," they are not truth evaluable, even if rules exist for their use.

Despite Ramsey's switch, few others 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 if ever 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. 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 significant differences between the two views[9], but they are not great.

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 much toil by leap-frogging to the present situation in metaethics.

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

[9]One attractive feature about nomic projectivism is that one 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. There are many responses, but this situation is awkward. The projectivist has no such problem. Suppose, for comparison, that ghosts played an important role in navigating through life. Ghosts are almost by definition supernatural and don't supervene on the natural world. A "projectivist" treatment in which we 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.

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 Schrodinger's equation is a law, that Schrodinger'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 on features of laws. The Problem of Creeping equally affects laws of nature as it does moral language. It is now very hard to tell the difference between a major form of nomic realism and a major form of 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. What we want from laws is that they 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 Where Now?

If I'm correct, the Problem of Creep affects the landscape in laws of nature as much as morality. This news will be received as no news to global expressivists such as Blackburn and Price, but it should be news to Humeans about laws of nature. Humean systematizers believe that they are realists about laws, but I've just shown that they are "*this close* [finger and thumb barely apart]" to irrealists about laws. I should also point out that the Problem of Creep affects the differences amongst many other views, including non-Humean ones; my focus, however, is on the two great Humean traditions. How should Humeans about laws respond?

That's not so clear. I can imagine three classes of reactions.

Naturally one reaction in the post-Creep world has been for metaethicists to find new ways to draw the line between realism and irrealism. Perhaps, for instance, even if moral realists and irrealists can now 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 and no one of them stands out as 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." With some study in philosophy of language and the meta-metaethics, philosophers of science can eventually "catch up" to the situation in ethics. But since we can already see where that will lead, I'm not convinced that reprising what happens in ethics will be worth the effort.

Another response is to embrace a global expressivism such as that held by Blackburn and Price. If one is concerned to see a difference between realist and irrealist Humean laws, then that path won't deliver a clear difference between the two. In a similar vein one sees calls to "move on" beyond cognitivism and non-cognitivism, representationalism and non-representationalism. I don't know how to precisely characterize this path, but the result seems to be articles in metaethics with titles that sound like self-help guides to bad break-ups. This more radical path will shift the laws debate more or less entirely into far more general issues.

A third response is of course to deny the Problem of Creep by rejecting minimalism. Here I want to emphasize that I'm not arguing that we *should* embrace minimalism, so this option is completely open for all I've said. Yet it seems to be a case of winning a battle without advancing much on the war. First, it means that the principal difference between realism and irrealism hangs on the correct theory of truth and representation, not anything specifically about ethics. Second, even before the Problem of Creep, we saw that the two theories of laws were moving closer together for decades. Parfit famously argues that the deep disagreements amongst Kantians, Contractualists, and Consequentialists are not so deep, that they are "climbing the same mountain on different sides" (2011, 385). I think a similar picture holds here. The Problem of Creep takes one to the summit, where one can see the other views are one's own. Before reaching the top, however, the general trends already showed that we were on different sides of the same mountain. At these lower altitudes there are differences, but they aren't so deep.

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. But it's not clear how Humeans should now understand laws of nature.

Philosophy papers are often convincing when they deliver bad news and go awry when they try something positive. Risking this fate, in the space remaining I want to suggest

that what matters is not the difference between systems theory and projectivism but instead the full Humean story about laws. This response may strike some readers as "putting one's head in the sand." Yet drawing lines between Humean theories of laws and declaring a winner now seems less interesting than it was before. My reactionary advice is to leave the Good Old Days and look back to the Really Old Days. That is, return to Hume and work on the central insights that motivated Humeanism about modality in the first place. There we will find an attractive naturalistic theory of laws that answers the main metaphysical questions we had about them, 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.

9 The Humean Core

Hume 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 and a theory of when it makes sense to attribute modalities.

Only the actual world exists, so why isn't it enough? Why bother with what could or couldn't or must happen? The best answer to this question deserves development in a book of its own (for a start, see Ismael 2015, Strevens 2007), not here. Yet Hume points us in the right direction. Hume's insight, in short, 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 I think is non-actual turns out to be actual.

I walk into a room and wonder whether I should push the button I find. I want the light on. If I knew the exact state of the system and a conditional that states that if I push the button a light will come on, then I don't need anything non-actual. Causes and counterfactuals don't add anything to my beliefs about what actually is. Yet I don't l know what actually is. I don't know the future nor do I know the exact state of the present. So I must employ hypothetical reasoning. Should I press it? Well, what might

happen if I did? To answer this I devise theories based on similar encounters of not only what actually happens but what might happen in different scenarios. These theories glue together variables across possible worlds, thereby giving us modal connections. I press the button; but I wouldn't have unless I thought that not pressing the button wouldn't have turned the light on.

Perhaps a being who knows everything actual or has unlimited time and computational abilities doesn't need modality. For those in our predicament modality is essential if we are to survive and thrive. We must reason hypothetically, and some of these hypotheticals are counterfactual. There is obviously a lot more to say, but the core idea is that reasoning this way pays off for creatures like us as we navigate through life.

Not only does Hume provide a rationale for modality, but in addition he offers a theory of when and how we make these modal attributions and further speculates about the source of this necessity. Regarding the former, our mind applies certain "rules" for judging cause and effect, e.g., that like causes regularly precede spatiotemporally contiguous like effects. Regarding the latter, he suggests that a psychological faculty is responsible for necessity, namely, a state 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."

Stepping back, Hume offers an explanation and justification for modality. 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.

Fast forward to contemporary times. While all the particular pieces of Hume's theory have been modified, the overall schema has held up. Recent work in casual learning theory and developmental psychology offer psychological faculties quite different than Hume's state of expectation. Work in causal methodology replaces 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, weights them for their probabilities of occurring, and runs them forward using rules it has found useful in the past. These are weighted future indicative conditionals, e.g., if F_1, then G_1 will happen, if F_2, then G_2 will happen, and so on. Because the antecedents are typically unknown in decision contexts, it's easy to

take these as counterfactuals. In English 'will' has both an indicative and counterfactual usage. These models can then be regarded as different possible worlds and the rules we use as the causal laws governing them.

Take, for example, the "theory theory" (Carey 1985; Gopnik and Meltzoff 1997; Wellman 1990). This is an empirically and theoretically powerful way of understanding causal learning. It's very easy to "Humeanize" such a theory: simply take it to be replacing both Hume's rules and also the psychological faculty. The rules become something much more sophisticated than Hume's eight rules, something closer to a contemporary theory of causation using Bayes nets (Gopnik et al 2004) The psychological faculty replaces a 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 a child's environment. Because it is innate such a map gives the theory a more Kantian than Humean cast, but since the map changes in response to new evidence - much as a scientist revises her scientific theories - it is in another sense an empiricist theory, and in any case the nativism is a detachable part of the theory. Cognitive maps of space are well-supported by behavioral and neurobiological evidence. The present idea is that we also employ causal maps. Just as rats explore undesirable parts of a maze in order to form a spatial map for future use, we as children play to help discover the modal relationships amongst the variables we encounter. The result is very different from Hume, of course, but we still have an origin story of a psychological faculty responsible for modality and a rationale for it via servicing future action.

The above, or something like it, is the core Humean picture. I find it 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.

What of laws? Hume doesn't really focus much on laws as opposed to causes (except when discussing miracles).[10] For Hume they are empirical generalizations formed from the patterns of covariation we encounter. For Gopnik 2000, who posits capacities for detecting causation, "Science simply puts these universal and natural capacities to work in a socially organized and institutionalized way." Presumably laws of nature are then

[10]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.

the product of this social organization and institutionalization, particularly compact and powerful ways of saying a lot about the best causal maps in the relevant fields.

That brings us to a question, one that runs throughout Hume's work and connects with our previous discussion. Consider Gopnik's child scientist. When she grows up and becomes a real scientist, she discovers that she needs a lot of help and has much to learn. Confront an individual with all the billions of events during a run at the Large Hadron Collider. Ask her what variables cause other variables. She will be at a loss. She needs computers, teams of statisticians, and knowledgable fellow scientists. She is going to need to learn about the world. Agreement 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. Contemporary particle hunters only find the Higgs boson if they also bring their causal maps into alignment.

Not all causal maps are equal. There are many pressures to change one's causal map. Some are better than others. But 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 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., I'm not proud of it and I wish it weren't so, but I enjoy 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 followed the rules (say, Pearl's structural causal models) perfectly 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. 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, 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.

Consider the famous Michotte 1945 experiments on causation. Michotte shows that some very simple non-causal visual sequences will be perceived as causal if they occur at the right positions and times. They will create in me a feeling of expectation that I ordinarily attribute to causation. The community standard for causation, however, states that there is no causation present. In my judgement, should I ignore my feeling of expectation and agree with the standard? Certainly. The community standard – say, Pearl's theory of causation – typically matches my expectations elsewhere. The standard can engage the relevant psychological faculty. It also has a great track record of success. And it helps me bring my judgement in line with that of others so that we can coordinate our actions. I should therefore treat this feeling of expectation in a Michotte case as an idiosyncratic one. I dismiss it just like I dismiss illusions of motion as cases of genuine motion. This causal standard is mutually accessible and leads to consensus, consensus that itself leads to useful progress.

If we adopt this picture — which I wish I had space to develop — then what do we say about laws of nature? Science has standards for what gets dubbed a law of nature. The systems picture is right about that. But given that the negotiation and organization and institutionalization in this Hobbesian jungle of vying causal maps is bound to be pretty messy, like any other type of social consensus making, we're unlikely to find a clean timeless unique criterion of what makes something a law of nature. We can gesture

like the systems theory does at the arguments of a function implicitly used by science. Perhaps as AI is increasingly employed in science, we might even find algorithms that generate a lot of stable consensus on the laws. Yet we shouldn't think that any one of them is uniquely correct unless we go to the ideal limit of science. Laws of nature, broadly understood, are massively important, as they are the recommendations on which our sciences have achieved consensus. Philosophers and sociologists of science have and will continue to study this process. But metaphysically laws are less interesting: they are the somewhat imprecise results of negotiation among our individual causal maps.

Stepping back, 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; focusing on the later improvements, organization and consensus building, we can see why a systems theory is appealing. They both get something right. Does the metaphysician or philosopher of science really need to decide between the two? No, neither tells the full story, and we know where picking a side will lead - we just watched it happen in metaethics.

Moreover, Humeans have a great theory. It is an empirically and theoretically rich one that provides an explanation and justification for modality. The theory is still being actively developed, but little in it is metaphysically puzzling. Trouble arises from trying to stuff it into established boxes. It is often pointed out, sometimes in frustration, that Hume never comes right out and tells us how to box his theory in contemporary categories. Does he subscribe to projectivism or an ideal observer theory or something else? Maybe in his imprecision he was on to something.... [11]

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.

[11] 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, Elanor Cranor, Michael Hicks, Barry Loewer, Elliott Sober, and especially the philosophers at Ranch Metaphysics 2020.

[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.