

Constraint Accounts of Laws

Christopher J. G. Meacham

Forthcoming in Ergo

Abstract

In recent work, Adlam (2022b), Chen & Goldstein (2022), and Meacham (2023) have defended accounts of laws that take laws to be primitive global constraints. A major advantage of these accounts is that they're able to accommodate the many different kinds of laws that appear in physical theories. In this paper I'll present these three accounts, highlight their distinguishing features, and note some key differences that might lead one to favor one of these accounts over the others. I'll conclude by briefly discussing a version of a "constraint" account that I think is especially attractive.

1 Introduction

A number of different accounts of laws have been offered in the literature. The most popular accounts include Humean accounts (e.g., Lewis (1994)), Necessitation Relation accounts (e.g., Armstrong (1983)), Dispositional Essentialist accounts (e.g., Bird (2007)), and Primitive Local Dynamics accounts (e.g., Maudlin (2007)). Although these accounts have different pros and cons, they have a shared weakness. All of these accounts have difficulty accommodating the full variety of nomic possibilities and laws that physicists have taken seriously.

Adlam (2022b), Chen & Goldstein (2022), and Meacham (2023) have recently proposed novel accounts of laws which can accommodate these possibilities. Although these accounts differ in detail, all three suggest a similar picture of laws. First, all three accounts posit primitive nomic features of the world. Second, all three accounts take these nomic features to be global – i.e., features of the world as a whole. Third, all three accounts suggest a picture of laws in which they're best thought of as things that constrain the world, instead of things that (say) produce or generate the world.

This paper aims to examine these primitive global constraint accounts. The goal of the paper is to consolidate a recent trend of research on constraint accounts, and lay out a framework for debate. In section 2 I'll describe the features that these three accounts have in common, and contrast them with some popular accounts of laws. In

section 3 I'll provide a more detailed description of these three accounts. In section 4 I'll explore some of the key differences between these accounts, and assess the pros and cons of these differences. In light of this discussion, I'll briefly present a Primitive Global Constraint account that I think looks particularly appealing in section 5. I'll conclude in section 6.

2 Primitive Global Constraint Accounts

Adlam (2022b), Chen & Goldstein (2022), and Meacham (2023) have recently advanced accounts of laws that resemble each other in a number of respects. I'll call them:

Primitive Global Constraint Accounts: Primitive Global Constraint accounts of laws maintain that (i) the world has primitive nomic features, (ii) these nomic features are "global"; they work at the level of the world as a whole, and (iii) while these nomic features constrain what the world must be like, they don't do anything more than that – e.g., they don't engage in substantive time asymmetric relations, like production or generation.

Thus all Primitive Global Constraint accounts (or "constraint accounts", for short) take there to be fundamental irreducible nomic features of the world. And they take these nomic features to do something like *constrain* what the world is like; where "constraining" is like "governing", but without the temporal connotations.¹ Thus these nomic features do something more metaphysically heavyweight than merely

¹Ott (2022) offers a critical discussion of the notion of "constraining" employed by Chen & Goldstein (2022), and his commentary applies equally well to the other constraint accounts. Ott raises a dilemma regarding the Chen & Goldstein's (2022) notion of constraining: either constraining is a form of causation or it's some weaker relation. If we take it to be a form of causation, we face worries about laws being the right kinds of things to cause anything, and worries regarding causal overdetermination. If we take it to be a weaker relation, and take "constraints" to just be, say, primitive modal constraints, then Ott worries that the notion is so weak as to be compatible with Humeanism about laws. Although proponents of constraint accounts haven't said much about the notion of "constraining", here's how I think they should respond to Ott's worry. They should agree that constraining is not a form of causation, and take the second horn of Ott's dilemma, but resist the claim that their theory collapses into a form of Humeanism. The more modest way to do this is to hold that constraint accounts aren't Humean because these modal constraints violate the Humean denial of necessary connections – e.g., these constraints won't supervene on the Humean mosaic. (And they can distinguish "constraining" from "production" by simply holding that the latter is committed to some kind of fundamental time asymmetric relation, which "constraints" do not commit you to.) A more ambitious response is to add that Ott is right that more should be said to flesh out the notion of "constraining" they accept (and the notion of "production" they reject), and to provide these details in a way that further distinguishes constraint accounts from Humean accounts. There are various ways in which one might do this; see footnote 2 for two possibilities. (I thank an anonymous referee for bringing this worry to my attention.)

describing what the world is like, but not as metaphysically heavyweight as “generating” or “producing” future states of the world from past ones. Though none of the proponents of constraint accounts provide a reductive definition of the notion of constraining, one might try to spell it out in terms of essences or grounding relations. The question of how to best spell it out in such terms would lead us too far afield, however, so I’ll limit my discussion of these options to a footnote.²

Constraint accounts take these nomic features to be “global”, in the following sense. Let one region be a *full copy* of another if they contain the same number of things, those things have the same fundamental properties and relations, those properties and relations have the same higher-order fundamental properties and relations, and so on.³ Let’s say that an account of laws is *local* if any regions that are full copies of one another will instantiate the same laws and chances, and *global* otherwise. Primitive global constraint accounts are global in this sense – intuitively, you can know everything fundamental that there is to know about a region, its qualities, its quality’s qualities, and so on, without knowing what laws and chances are instantiated in that

²Spelling out what “constraining”/“governing” and “producing”/“generating” amount to is a non-trivial task. One possibility is to cash them out using real definitions and essences, perhaps along the lines suggested by Wilsch (2021). Another possibility is to cash out these notions in terms of grounding. (See Coates (manuscript) for a discussion of some different ways of linking grounding and laws.) For example, one might understand the claim that laws *constrain* or govern the world (in this sense) as the claim that laws ground their worldly instances. And one might understand the latter claim along the lines suggested by Emery (2019), where the fact that it’s a law that L grounds instances of L , where instances of L are understood as collections of events such that the law and one of these events entails the others. (E.g., one fact – that the laws are L – grounds the fact that E_1 at t_1 will be followed by E_2 at t_2 .) By contrast, one might understand the claim that the laws *produce* future states of the world from past ones as the claim that laws and past states ground future states. (E.g., *two* facts – the fact that the laws are L and the fact that E_1 at t_1 – grounds the fact that E_2 at t_2 .) (Note that constraining and production, so understood, are compatible – they can both obtain with respect to the same collection of events. Note further that there aren’t any obvious overdetermination issues, since different facts are being grounded in cases of constraining and cases of production.)

³This notion is more demanding than Lewis’s (1983) characterization of duplication. We can spell it out more precisely as follows. For simplicity, I assume substantivalism about spacetime, that there are no extended simples, and that mereological and location relations are fundamental (see Eddon (2017) for an argument for the former; the latter claim is relatively uncontroversial). Let the *stuff in region* R , S_R , consist of: (a) all of the spatiotemporal parts of R ; (b) all of the material objects that bear the exact-location relation to some part of R ; (c) all of the first-order fundamental properties and relations that hold of or between the things in (a) and (b); and (d) all of the n th-order fundamental properties and relations that hold of or between the $(n - 1)$ st-order properties and relations in (c) and (d). Then we can say that a region R is a *full copy* of a region R^* iff there’s a one-to-one and onto mapping f between S_R and S_{R^*} such that for any fundamental properties and relations F : if $x_1, x_2, \dots \in S_R$, and $F(x_1, x_2, \dots)$, then $F(f(x_1), f(x_2), \dots)$. Intuitively, the stuff in a region is just all of the things that exist in that region, and all of the fundamental properties and relations that hold of those things, and all of the fundamental properties and relations that hold of those fundamental properties and relations, and so on. And one region is a full copy of another if all of the stuff in one region can be paired with the stuff in the other in a way that preserves fundamental properties and relations.

region.

To help us get a clearer picture of this view, let's look at how constraint accounts differ from some standard accounts of laws.

On Humean accounts, such as Lewis's (1994), the laws are just helpful descriptions of regularities found in the distribution of local occurrent facts. Humean accounts join constraint accounts in taking the laws to be global features of the world. But they don't take the laws to be doing something as metaphysically substantive as *constraining* the world.⁴ And they deny that the world has primitive nomic features – any nomic features the world has are reducible to facts about the distribution of local occurrent properties.

On Necessitation Relation accounts, such as Armstrong's (1983), the laws are determined by a fundamental "necessitation relation" N that holds between fundamental properties, where if $N(F, G)$ holds at a world, then anything at that world that has property F will also have property G .⁵ Like constraint accounts, Necessitation Relation accounts take worlds to have irreducible nomic features (certain instantiations of the necessitation relation). And they take these nomic features to do something like constrain what the world is like. However, unlike constraint accounts, Necessitation Relation accounts are local. For on Necessitation Relation accounts, regions that are full copies of each other will instantiate the same laws and chances, since they'll instantiate the same fundamental properties and the same necessitation relations between them.

On Dispositional Essentialist accounts, such as Bird's (2007), the laws are determined by the essential dispositions of fundamental properties. Dispositional Essentialist accounts join constraint accounts in taking there to be irreducible nomic features (the dispositions of fundamental properties). Likewise, they take these nomic features to do something like constrain what the world is like. But unlike constraint accounts, Dispositional Essentialist accounts are local. Since the dispositions of fundamental properties are essential – the same at all worlds – regions that are full copies of each other will instantiate the same laws and chances.

⁴Of course, most proponents of Humean accounts don't make claims about grounding, real definitions, and the like. Likewise, most proponents of the other accounts discussed below (Necessitation Relation, Dispositional Essentialist, and Primitive Local Dynamics accounts) don't make claims about grounding or real definitions. And the proponents of Primitive Global Constraint accounts follow suit. Thus if we spell out the notions of "constraining" and "production" in one of the ways suggested in footnote 2, then none of these accounts are clearly committed to any particular stance on whether laws constrain or produce. That said, some ways of adding claims about grounding and the like to these accounts are clearly in the spirit of such accounts, while some other ways of adding such claims are not. For the purposes of this discussion I assume we're adding such claims in a manner which accords with the spirit of these accounts.

⁵Some versions of Necessitation Relation accounts, such as Tooley's (1987), allow for a much broader range of necessitation relations; relations that can hold between any number of fundamental properties, and necessitate any kind of nomic relation between them (not just entailment). As we'll see, this more flexible view has pros and cons relative to more austere accounts like Armstrong's. In the text, I'll focus on accounts like Armstrong's, and flag how things differ for Tooley's account in footnotes.

On Primitive Local Dynamics accounts, like Maudlin's (2007), laws are primitive features of regions that determine how future regions are generated or produced from prior ones. (Thus on these accounts, all fundamental laws are laws of temporal evolution.) Like constraint accounts, Primitive Local Dynamics accounts take worlds to have primitive nomic features. But they differ from constraint accounts in being local; since the laws of temporal evolution that govern a region are features of that region, any full copy of a region will instantiate the same laws and chances.⁶ Likewise, they differ from constraint accounts in taking the laws to do more than merely constrain what the world is like. For the laws not only constrain the future states of the world, they also *produce* or *generate* these states.

2.1 An Advantage of Constraint Accounts

In section 2 we saw how constraint accounts differ from some standard accounts of laws. Now let's look at why one might prefer a constraint account over these more established alternatives. The most obvious advantage of constraint accounts, relative to these alternatives, is that they have an easier time accommodating the variety of nomic possibilities and fundamental laws considered by the scientific community. To see this, let's briefly rehearse some difficulties facing the standard accounts here, and how the constraint accounts do better.⁷

Let's start with Humean accounts. Humean accounts can't recognize laws and nomic possibilities that don't supervene on the distribution of local occurrent facts. Thus, for example, Humean accounts are unable to recognize otherwise identical worlds where there's a single chancy coin flip, but the chance of heads is 0.5 at one world and 0.6 at the other. Likewise, Humean accounts have difficulty recognizing worlds with uninstantiated laws, such as laws governing charge at a world where no charges are instantiated.⁸

Necessitation Relation accounts run into difficulties because most of the fundamental laws scientists have entertained are hard to characterize in terms of necessitation relations between fundamental properties. For example, it's unclear how to characterize laws concerning particular times or places, like laws regarding initial conditions (e.g., the Past Hypothesis), since such laws generally aren't reducible to

⁶Strictly speaking, it's hard to assess whether Maudlin's account is local or global, because Maudlin suggests that laws don't fall into any of the usual ontological categories – they're not individuals, properties/relations, etc. And my characterization of being a full copy (upon which the local/global distinction relies) is like Lewis's characterization of duplication in that it presupposes that everything falls into one of the usual ontological categories. Without further details regarding this new ontological category, it's hard to know how to provide a satisfying extension of the characterization of full copies that fits Maudlin's account. (Just as it's hard, without further details, to know how to provide a satisfactory characterization of duplication that fits Maudlin's account – e.g., are two worlds with different Maudlin laws duplicates?)

⁷See Adlam (2022b) and Chen & Goldstein (2022) for a more in-depth discussion of these issues.

⁸For some classic criticisms along these lines, see Armstrong (1983) and Van Fraassen (1989).

claims about entailments between fundamental properties.⁹ Likewise, it's unclear how to characterize laws that concern holistic rather than particular facts about the instantiations of properties, like conservation laws (e.g., charge conservation), since global claims about the total amount of charge don't seem reducible to claims about the co-occurrence of different fundamental properties.¹⁰ And it's unclear how to provide a satisfactory characterization of functional laws (e.g., Coulomb's law), since it's difficult to see how to characterize laws relating different magnitudes of properties in terms of claims about the co-occurrence of particular properties.¹¹

⁹A classic case in favor of adopting a lawful Past Hypothesis is offered by Albert (2000); for a discussion of why it's crucial that it be a *law* and not merely a contingent boundary condition, see Meacham (2010). This worry regarding Necessitation Relation accounts is discussed by Tooley (1977) and Armstrong (1983). Both Armstrong and Tooley have suggested ways of modifying their account to get around this worry, with Armstrong suggesting that we allow necessitation relations to hold of "quasi-universals", non-fundamental relatives of fundamental properties that make reference to particulars, and Tooley suggesting that we allow necessitation relations (broadly construed) that involve ineliminable reference to particulars. But these moves have struck many as *ad hoc* maneuvers.

¹⁰See Bigelow et al. (1992) for a discussion of this worry for Armstrong's account. Note that Tooley's account, which imposes virtually no constraints on the kinds of necessitation relations one can employ, has the tools to handle this problem easily. Consider the (possibly infinite) description, in the language of fundamental properties, of a model of (say) Newtonian mechanics. And consider the disjunction of all of these descriptions of models. Tooley can posit a necessitation relation which takes all of the relevant properties as arguments, and necessitates that this disjunction of descriptions must obtain – i.e., that one of the models of Newtonian mechanics must obtain. One can perform a similar trick for any set of laws.

Two comments. First, this move makes the account global, and so is a kind of Primitive Global Constraint account. Second, for fans of Necessitation Relation accounts, this "brute force" approach has a number of drawbacks. It seems *ad hoc*, provides a messy and complicated picture of laws, and gives up much of the intuitive appeal of the account, which is to appeal to relatively simple connections between local fundamental properties to explain the patterns we see in nature.

¹¹Necessitation Relation accounts can, of course, posit necessitation relations between the determinates for any particular instance of these functional laws. For example, with respect to the gravitational force law, they can posit that particular magnitudes of mass and distance necessitate particular magnitudes of force. But this move raises some questions. First, is there anything metaphysically substantive that bundles these different instances of a functional law together, or is the fact that we can express them all using one simple function a fortuitous coincidence that needs to be explained away? Second, these different nomic relations between magnitudes seem to smoothly co-vary in a way that tracks certain ordering and distance relations between these magnitudes; indeed, these relations seem to smoothly co-vary in a similar way in *different* functional laws. What explains this, on the Necessitation Relation account?

Armstrong (1983) attempts to address these questions by positing a further higher-order law that bundles instances of functional laws together. But it's hard to see how this move provides a satisfactory answer to these questions, especially the second one. At bottom the problem is that standard formulations of Necessitation Relation accounts yield laws that don't substantively engage with the ordering and distance structure of the magnitudes that appear in functional laws. And without such engagement, it seems one is going to run into problems, whether they're problems of trying to explain away the apparent fact that there are objective features of magnitudes of this kind, trying to explain away the apparent fact that these structural features of magnitudes play a substantive role in the laws (and the same role in different laws), or trying to show how these structural features do play a substantive role in the laws, despite the postulated

Dispositional Essentialist accounts run into problems because many of the laws physicists have considered aren't easily derivable from or consequences of the dispositions of fundamental properties. In fact, Dispositional Essentialist accounts and Necessitation Relation accounts have difficulties with many of the same kinds of laws. For example, it's unclear how to use the dispositions of fundamental properties to derive laws concerning initial conditions, since it's hard to find properties that plausibly (i) are fundamental, (ii) yield dispositions that impose the right constraints on initial conditions, and (iii) apply to the kinds of things which these would need to have those dispositions.¹² Likewise, it's hard to see how to use the dispositions of fundamental properties to derive holistic laws like charge conservation, since constraints on global facts about the total amount of charge don't seem reducible to facts about the dispositions of particular charged objects.¹³ And it's unclear how to use the dispositions of fundamental properties to yield the functional laws posited by physics, since it's hard to provide a derivation of such laws that (i) is appropriately unified, (ii)

form of these laws. See Eddon (2007) and Forge (1999) for discussions of some of these issues.

Tooley (1987) addresses these questions by (a) adopting a Field (1980)-style approach to quantitative laws, but applied to first-order fundamental properties instead of objects, and (b) adopting a functional understanding of the names we assign to properties, according to which the label "mass" applies to all and only the properties that play a certain role in the laws (the mass role), though (c) there's still an underlying layer of fundamental categorical properties. This approach can answer the two questions given above, but has a number of counterintuitive consequences. For example, it gives up the intuition that the force between masses could have been governed by an inverse cube law instead of an inverse square law (since that would no longer be "force"), it gives up the intuition that a perfect duplicate of me will have the same mass and shape (since duplicating the fundamental categorical properties doesn't entail these properties will play the same roles), it gives up the intuition that there's a sense in which different magnitudes of mass are objectively similar regardless of what the laws are like (since the laws could have unrelated fundamental properties playing the role of different mass magnitudes), it gives up the intuition that there are objective (non-nomic) structural facts about quantitative properties (since quantitative structure is determined by the laws, not the properties), and so on.

¹²What we want is for something to be disposed to be such that it will have the appropriate initial state. So it looks like the kind of thing we want to have this disposition is something big – perhaps time slices of the universe, or the universe as a whole. But it's hard to see what plausible fundamental property would both apply to such a thing and yield the relevant constraint. Perhaps the most straightforward option, following Bigelow et al. (1992), is to take the relevant object to be the entire world, and take the relevant property to be a very fine-grained fundamental property, like that of being such-and-such a kind of world. But note two things. First, this move makes the account global, and so is a kind of Primitive Global Constraint account. Second, for fans of Dispositional Essentialist accounts, this move has a number of drawbacks. It seems *ad hoc*, to provide poor explanations, and to give up on the much of the intuitive appeal of the account, which is to appeal to the common dispositions of local fundamental properties to explain the patterns we see in nature (see Ioannidis et al. (2021)).

¹³See Bigelow et al. (1992) and Bird (2007) for discussions of this worry. As with the worry concerning initial condition laws, one can escape this worry by positing fine-grained fundamental properties of the world as a whole, like being such-and-such a world, and taking these dispositions to yield conservation laws. But as noted earlier, there are several reasons why fans of Dispositional Essentialist accounts might worry about this move (cf. footnote 12).

assigns plausible dispositions to properties, and (iii) appeals only to the usual kinds of fundamental properties.¹⁴

Primitive Local Dynamics accounts run into difficulty because a number of the laws physicists have taken seriously are either hard to understand dynamically, hard to understand locally, or conflict with the idea of temporally asymmetric production. For example, it's hard for Primitive Local Dynamics accounts to accommodate laws regarding initial conditions, since such laws are not dynamical. Likewise, it's hard for Primitive Local Dynamics accounts to accommodate laws whose solutions are complete histories (such as Einstein's Equation), since such laws are neither dynamical nor local. And it's hard to square Primitive Local Dynamics accounts with retrocausal laws (such as the Wheeler-Feynman absorber theory, or the two-state vector accounts of quantum mechanics), or possibilities involving closed time-like curves (such as Gödel solutions in General Relativity), since such laws seem incompatible with a picture of temporally asymmetric production.¹⁵

Primitive Global Constraint accounts, by contrast, don't face any of these difficulties. Constraint accounts avoid the worries facing Humean accounts because they don't take laws to supervene on the distribution of local occurrent facts. Thus constraint accounts have no problems recognizing possibilities that are identical with respect to their occurrent facts, but differ with respect to their laws. Constraint accounts also avoid the worries facing Necessitation Relation and Dispositional Essentialist accounts, because they don't need the laws to be something you can derive from entailments between or dispositions of fundamental properties. Thus they have no problem accommodating, e.g., initial condition laws, conservation laws, or functional laws. And constraint accounts avoid the worries facing Primitive Local Dynamics accounts because they don't require laws to be dynamical and local, and they don't presuppose some temporally asymmetric notion of production. Thus they have no problems recognizing laws that aren't easily formulated in terms of local dynamical laws, or recognizing retrocausal laws or laws which admit the possibility of closed

¹⁴Some versions of Dispositional Essentialism, like Bird's (2007), are naturally understood as assigning dispositions to determinates, raising questions about whether there's anything metaphysically substantive that bundles these different instances of functional laws together (see Vetter (2012)). An alternative approach is to assign something like functional dispositions to determinables. But even if we can capture functional laws like Coulomb's law and the gravitational force law in such a manner, further problems arise here for Dispositional Essentialists. For functional laws generally engage with each other in various ways – e.g., Coulomb's law and the gravitational force law both yield forces that bear on the behavior of objects – and the ways in which they engage with each other seem lawfully governed. But if we take all laws to fall out of the dispositions of fundamental properties, we need to attribute these further “combination” laws to some fundamental property or other, and it's not clear there are any plausible candidates to attribute these dispositions to. (See Ioannidis et al. (2021) for discussion.) One could attempt to address this worry by positing fine-grained global fundamental properties – like being such-and-such a kind of world – and assigning the relevant dispositions to these properties. But there are several reasons why fans of Dispositional Essentialism might be uncomfortable with this move (cf. footnote 12).

¹⁵See Wheeler & Feynman (1945), Aharonov et al. (1964), and Gödel (1949).

time-like curves.

3 Three Constraint Accounts

Adlam (2022b), Chen & Goldstein (2022), and Meacham (2023) have all advocated adopting Primitive Global Constraint accounts of laws. But while these accounts are similar in a number of respects, they're distinct proposals. Let's take a more careful look at these views.

3.1 Adlam's Account

On Adlam's (2022b) proposal, the laws of nature are parts of the objective modal structure of the world. Adlam characterizes the modal structure associated with the laws in terms of constraints induced by these laws. And these constraints metaphysically necessitate that the world be certain ways.

More formally, let a "Humean mosaic" be a distribution of local categorical properties across spacetime.¹⁶ Adlam proposes to associate laws with functions that take sets of Humean mosaics as inputs and yield real numbers as outputs, in a manner that satisfies the probability axioms. For non-probabilistic laws, these values are all 0s or 1s, and the function serves to identify the set of Humean mosaics that this law requires the actual world to "belong" to.¹⁷ For probabilistic laws, these values can take any value in the $[0, 1]$ interval, and the function serves to determine the likelihood of different sets of Humean mosaics being the ones the actual world has to belong to.¹⁸ Note that these laws do not need to be the complete laws of a world – many different laws can obtain at the same world. When there are multiple laws that obtain at a world, the combined effect of the laws is to require that the world belong to the intersection of the set of Humean mosaics picked out (deterministically or probabilistically) by each law.

Two notes. First, Adlam spells out the account in terms of functions over Humean mosaics instead of worlds because worlds contain more information than we want, such as information about what the laws and chances are. And characterizing laws

¹⁶Note that on Adlam's understanding, neither these categorical properties nor the spatiotemporal relations specifying their locations have to be fundamental properties. This allows for physical theories according to which spacetime is emergent. It's unclear whether Adlam's claim that these are "local" categorical properties means these categorical properties have to be instantiated at points, or whether they just have to be instantiated in some spatiotemporally restricted manner. I take the latter position to be more attractive, so I'll assume that here.

¹⁷Strictly speaking, since Humean mosaics aren't worlds, sets of Humean mosaics aren't something a world can belong to. When I talk of a world having to "belong" to a set of Humean mosaics, I mean the Humean mosaic of that world must be identical to one of the members of that set.

¹⁸See Adlam (2022a) for a discussion of one way one might develop an account of chances on her account in more detail (though she doesn't commit herself to this picture).

and chances in terms of functions over worlds which are partially individuated by what laws and chances obtain raises worries about circularity/triviality. Second, Adlam characterizes these functions in terms of sets of Humean mosaics instead of individual Humean mosaics because laws generally don't tell you precisely what Humean mosaic the actual world must belong to. For example, the gravitational force law will pick out a set of Humean mosaics that satisfy that law, but doesn't say anything further about which of those worlds is the actual one.

Adlam takes one of the advantages of her account to be its neutrality with respect to many of the other questions one might ask about laws and chances. For example, the account leaves open what, if anything, the laws of nature add to our ontology. The account leaves open whether laws of nature are qualitative intrinsic features of the world, and whether they'd be preserved by duplication. The account leaves open how we should think about the modal status of the laws of nature.¹⁹ And so on.

So, to sum up, Adlam's account posits constraints that metaphysically necessitate that the Humean mosaic have certain features; constraints whose content can be characterized using probability functions over sets of Humean mosaics. And Adlam's account associates/identifies laws with these constraints.

3.2 Chen and Goldstein's Account

Chen & Goldstein (2022) propose to take the laws of nature to be primitive facts about the world. Their discussion suggests that these laws have something like sentential form – for example, they hold that the laws refer to particular properties.²⁰ As a consequence, they take the laws to be relatively fine-grained. E.g., a law that refers to the properties green and blue can be distinct from an intensionally equivalent law that refers to the properties grue and bleen. That said, Chen and Goldstein are explicitly agnostic about how fine-grained the laws are, and don't provide a precise characterization of the form of the laws or how they're individuated.²¹ For concreteness, I'll

¹⁹It's tricky on Adlam's account to specify which Humean mosaics are nomically possible. A natural thought is to take the union of all of the sets of Humean mosaics that are assigned a value of 0 by any of the laws that obtain, and take those to be the nomically impossible Humean mosaics. But that's not what we want, since for laws that assign continuous probability distributions, every Humean mosaic will belong to some set that gets assigned a 0 value, and thus every Humean mosaic will be nomically impossible. (One might try a variant of this proposal that appeals to non-zero probability densities. But this requires positing more structure than you get from a probability function over sets of mosaics. For example, a probability density needs to be a probability with respect to something, which requires a viable reference measure over sets of Humean mosaics. And the resulting probability densities need to be combinable, which requires (among other things) that the account specify how the probability functions associated with different laws are related (e.g., are they independent, or do they bear some other relation? Cf. section 4.6).)

²⁰"Most of the fundamental laws we discover refer only to fundamental properties. But it is reasonable to consider candidate fundamental laws that refer to non-fundamental properties." (Chen & Goldstein (2022), p22.)

²¹"We should not think that, in every case, a law is equivalent to the set of possibilities it generates. The

assume from here on that they take laws to have a form similar to that of structured propositions, and are individuated in a similar manner.²²

On Chen and Goldstein's account, non-probabilistic laws are taken to pick out a set of models, which the world that instantiates that law must be a member of. Probabilistic or chancy laws are given a more ambiguous treatment; Chen and Goldstein discuss five options for how to incorporate chances into their account. Since this is relevant to the discussion to come, it's worth briefly reviewing what each of these options are.

The first option they consider is pairing their non-Humean account of laws with a Humean account of chances.²³ The second option they consider is to take chances to be distinct primitive facts about the world, facts that aren't directly related to the primitive constraints on nomic possibility imposed by non-probabilistic laws. The third option they consider is to take chances to be primitive facts about the world, and relate these facts to the constraints on nomic possibility imposed by non-probabilistic laws by adopting a gradable notion of constraining. In particular, they suggest understanding both probabilistic and non-probabilistic laws as imposing gradable constraints on nomic possibility – constraints which can take on values between 0 and 1 – and taking the non-probabilistic laws to be the ones that only impose extreme values (0 or 1).

The fourth option they consider is to adopt a typicality account of chances. There's a literature on how to do this for non-dynamical chances, but one might try to extend this interpretation to dynamical chances by positing primitive facts about which dynamical histories are typical. And one can relate these facts to the constraints on nomic possibility imposed by non-probabilistic laws by taking typicality facts to constrain nomic possibility too – i.e., by requiring nomic possibilities to be typical. Unfortunately, because there are atypical models of laws, this option entails that there are models of laws – *prima facie* nomic possibilities – that aren't possible.

The fifth option they consider is similar to the fourth, positing primitive facts about which histories are typical. But this option takes typicality and nomic possibility to be nested modal notions – every typical world is a nomic possibility, but not all nomic possibilities are typical. And both modal notions intuitively constrain the world – the actual world must be both nomically possible and typical. Unlike the fourth option,

two can be different. For example, there are many principles and equations that can give rise to the same set of possibilities denoted by Ω^H . But we expect laws to be simple. One way to pick out the set Ω^H is by giving a complete (and infinitely) long list of possible histories contained in Ω^H . Another is by writing down simple equations... which express simple laws. Hence, the equivalence of physical laws is not just the equivalence of their classes of models. For two laws to be equivalent, it will require something more... It is an interesting question, on MinP, what more is required and how to understand the equivalence of physical laws... we do not provide such an account as it is orthogonal to our main concerns in the paper." (Chen & Goldstein (2022), p20-21.)

²²See King (2019).

²³Although this might seem like a strange option to consider, there is precedent for adopting this combination; for example, see Hoefer (2018).

this option doesn't immediately rule out atypical nomic possibilities. But it does rule out atypical actualities, and one might be unhappy with this result. For example, if we follow Lewis (1973) in taking actuality to be an indexical property, then every world will be actual with respect to itself. But if so, then to forbid atypical actual possibilities is to forbid atypical possibilities full stop.

One of the main advantages of constraint accounts is that they can accommodate the nomic possibilities taken seriously by the scientific community. In particular, unlike Humean accounts, they can allow that every model of a set of laws corresponds to a nomic possibility. The first, fourth and fifth approaches considered here all give up this advantage, for they require us to deny that every model of a set of laws corresponds to a nomic possibility.²⁴ Thus if we want to keep this advantage intact, we need to adopt either the second or third approach. So while Chen and Goldstein don't rule out any of these options, I'll restrict my attention to the second and third options in what follows.²⁵

Chen and Goldstein's account is neutral with respect to a number of other questions one might ask about laws and chances, and, like Adlam, they take this to be one of the strengths of their account. For example, the account is agnostic about whether there are objective facts about how to divide the complete laws at a world into individual laws. The account is neutral about what kind of ontological category the laws belong to.²⁶ The account leaves open whether the laws are qualitative intrinsic features of the world, which would be preserved by duplication. And so on.

So, to sum up, Chen and Goldstein's Account posits primitive facts about the world that impose modal constraints on what can be actual, and it identifies laws with these primitive facts.

3.3 Meacham's Account

Meacham's (2023) proposal offers an account of laws and chances that mirrors certain accounts of quantitative properties suggested by measurement theory. Meacham posits a primitive nomic likelihood relation that tells us, roughly, when one state of affairs is more nomically likely than another. Meacham then characterizes the laws and chances in terms of these nomic likelihood relations.

More formally, let $\langle A, C, w \rangle$ be an ordered triple consisting of a pair of propositions (understood as sets of possible worlds) A and C , which aren't themselves about nomic facts, and a world w . Meacham's account begins by positing a fundamental nomic likelihood relation between such triples that satisfies certain constraints. Intuitively,

²⁴This is not to say that the other options Chen and Goldstein consider are not worthy of consideration. But due to lack of space, I'll confine my attention to these two options here.

²⁵Barrett & Chen (manuscript) offer an account of chance compatible with this framework that develops the fourth option described above.

²⁶Chen and Goldstein present the view in terms of "facts", but this is a stylistic flourish; they don't intend to commit themselves to an ontology of facts, or to the claim that the laws are such facts.

this relation holds when *A given C at w is at least as nomically likely as A' given C' at w'*. Meacham then provides a representation and uniqueness theorem showing, roughly, that there's a unique way of assigning to these triples numbers between 0 and 1, and the status of being nomically required or nomically forbidden, in a manner that lines up with these nomic likelihood relations.²⁷

The complete laws of a world w are then identified with the property of being a world that bears the same nomic likelihood relations as w . (The account doesn't take there to be any objective facts about how to divide the complete laws at a world into individual laws.) What's nomically required/forbidden at w is determined by the triples that are assigned the status of being nomically required/forbidden. And the chances of w are determined by the numbers assigned to those triples. It's worth noting that on this account the status of being nomically required and the status of having a chance of 1 are distinct. While every triple that is nomically required has a chance of 1, there can be triples with a chance of 1 that aren't nomically required. Likewise, the status of being nomically forbidden and the status of having a chance of 0 are distinct. While every triple that is nomically forbidden has a chance of 0, there can be triples with a chance of 0 that aren't nomically forbidden.

So, to sum up, Meacham's Account posits a fundamental nomic likelihood relation, and identifies laws with the property of being a world at which certain nomic likelihood relations hold.

4 Some Key Differences

Although the three constraint accounts described in section 3 are similar in many ways, there are also some important differences. In this section I'll look at some of the differences between them, and discuss why one might take these differences to tell for or against these accounts.²⁸

²⁷What does it mean to require that these assignments "line up" with the nomic likelihood relations? Roughly, that one triple is assigned a greater number than another *iff* it has a greater nomic likelihood, and that a triple is assigned the status of being nomically required (forbidden) *iff* it's nomically on a par with a tautology (anti-tautology).

²⁸There are two further worries one might raise for these accounts that are independent of the differences discussed below. Since these worries aren't central to the accounts, and relatively easy to fix, I'll confine my discussion of them to this footnote.

First, one might worry that Adlam's account can't accommodate the possibility of lawful worlds without spacetime. This is because Adlam pairs laws with functions over sets of Humean mosaics, and characterizes "Humean mosaics" in terms of distributions of properties over spacetime. But this problem can be fixed by characterizing "Humean mosaics" in a less restrictive way.

Second, one might worry that Chen and Goldstein's account is incomplete in some important ways. In particular, Chen and Goldstein don't specify how fine-grained laws are, or how to individuate laws. This is an intentional choice on their part, since they take these questions to be orthogonal to their main concerns. But it nonetheless makes it hard to answer key questions about the account, like what kinds of laws it can

Before we proceed, let me introduce some terminology. I'll say that two laws have the same *content* if they're compatible with the same worlds (or in the case of probabilistic laws, they assign the same probabilities to the same sets of worlds). I'll say that two laws have the same *non-nomic content* if they're compatible with the same Humean mosaics (or in the case of probabilistic laws, assign the same probabilities to the same sets of Humean mosaics). And I'll say that two laws are *otherwise identical* if they have the same non-nomic content.

4.1 Is There an Objective Way to Carve the Complete Laws into Individual Laws?

Consider the complete laws that obtain at a world w . Here are some questions one might ask about w . How many different fundamental laws obtain at w ? What way of dividing the complete laws at w into individual laws cuts nature at the joints? Suppose that at w the only laws that obtain are two fundamental non-probabilistic laws whose contents are A and B , respectively. Is w distinct from a world w' with only one fundamental non-probabilistic law whose content is $A \wedge B$?

Adlam's account starts by positing individual fundamental laws, and takes the complete laws to just be the bundle of the individual laws. Thus on Adlam's account there are non-trivial answers to each of the above questions: there's a particular number of fundamental laws that obtain at w , a correct way of dividing the complete laws at w into individual laws, and a difference between a world in which there are two laws A and B and a world in which there's only one law $A \wedge B$.

By contrast, Meacham's account starts with an account of the complete laws of a world, and takes the question of how to divide the complete laws into individual laws to be a purely pragmatic question. Thus Meacham's account provides deflationary answers to the above questions: metaphysically speaking, there is no right number of fundamental laws that obtain at a world, there is no correct way to carve the complete laws into individual laws, and there is no difference between a world with two laws A and B and a world with one law $A \wedge B$.

Finally, Chen and Goldstein's account sits in the middle – it's agnostic about whether there are substantive answers to these kinds of questions.²⁹

Why might one favor one of these approaches over the other? Those with certain inegalitarian intuitions about laws will be attracted to non-deflationary approaches like Adlam's. For example, some, such as Lange (2009), have argued that some laws (e.g., conservation laws) are more modally robust than others (e.g., force laws). This

recognize. But again, I take this to be a problem that is straightforward to fix, for one can add these details to the account. And indeed, as I noted in my discussion of Chen and Goldstein's view, for the purposes of this discussion I've filled in these details in what I think is a charitable way – by assuming that the laws have a form similar to that of structured propositions, and are individuated in a similar manner.

²⁹This issue doesn't arise in Chen and Goldstein's discussion, and in correspondence they've confirmed that they're agnostic between these two approaches.

position is hard to maintain on deflationary approaches, which don't take there to be metaphysical facts supporting any one way of dividing the complete laws into individual laws over any other.

By contrast, those with certain skeptical intuitions about the structure of laws will be attracted to deflationary approaches like Meacham's. For example, consider the classical formulation of Maxwell's Equations in terms of electric and magnetic fields. This yields four equations, and suggests that there are four fundamental laws governing electromagnetic fields. But if we consider the more modern formulation of Maxwell's Equations in terms of electric and magnetic potentials, we end up with two equations, which suggests that there are two fundamental laws governing electromagnetic fields. And one might be skeptical that there's a fact of the matter as to which of these ways of counting laws is correct.

4.2 How Fine-Grained Are The Laws?

We can use a sentence to express the content of a law. But do the laws themselves have something like sentential form? In particular, are laws like sentences in being more fine-grained than the propositions (sets of worlds) expressing their content? This raises a number of interesting questions regarding how fine-grained laws are. For example, are laws fine-grained enough to distinguish between otherwise identical laws which appeal to different properties? Are laws fine-grained enough to distinguish between otherwise identical laws which have different logical forms?

Chen and Goldstein adopt a fine-grained account according to which the laws have something like sentential form. On their account, there are facts about which properties the laws refer to. So otherwise identical laws which appeal to different properties are distinct. Likewise, on their account there are facts about the logical form of the laws – a law corresponding to a simple characterization of the states the laws allow is distinct from a law corresponding to an infinite disjunction of the individual states the laws allow. So otherwise identical laws which have different logical forms are distinct.

By contrast, both Adlam's and Meacham's accounts are relatively coarse-grained. On Adlam's account the laws correspond to functions that assign values to sets of Humean mosaics. Since such functions aren't sensitive to the properties used to pick out these sets, or the logical form of the claims used to pick out these sets, otherwise identical laws which appeal to different properties or have different logical forms are the same law. Likewise, on Meacham's account two worlds can only differ with respect to their laws if they differ with respect to their nomic likelihood relations. And since the relata of nomic likelihood relations are propositions which aren't themselves about nomic facts, and such propositions aren't sensitive to properties or logical form, nomic likelihood relations aren't sensitive to such factors either.

Fine-grained accounts are attractive to those who want to capture hyper-intensional distinctions between laws. For example, Adlam (2022b) has us consider the difference

between retrocausal laws that have temporally forwards and backwards dynamics, and retrocausal laws that impose an “all at once” constraint on possible world histories.³⁰ On coarse-grained approaches like Adlam’s and Meacham’s, there is no straightforward way to distinguish between such laws – e.g., nothing about a set of Humean mosaics tells us whether these states have been picked out by laws that impose forwards and backwards dynamics instead of laws that impose an all-at-once constraint. By contrast, on fine-grained approaches like Chen and Goldstein’s, we can distinguish between laws that are formulated in terms of dynamical equations of motion (e.g., Newtonian formulations of classical mechanics) and laws which are not (e.g., Lagrangian formulations of classical mechanics).

Coarse-grained accounts are attractive to those who are skeptical about such fine-grained laws. For example, consider a law that says “All emeralds are green and all sapphires are blue”. And consider a law that says “All emeralds first observed before t are *grue*, and all other emeralds are *bleen*, and all sapphires first observed before t are *bleen*, and all other sapphires are *grue*”.³¹ If one feels that there isn’t a substantive distinction between these two laws (i.e., that they’re just two different ways of expressing the same law) then one will be unhappy with hyper-intensional accounts that take these to be different laws.

4.3 Are Laws and Chances of a Kind?

At first glance, it seems like probabilistic and non-probabilistic laws are of a kind. Non-probabilistic laws tell us that if one state obtains, then certain other states are required to obtain. Probabilistic laws tell us that if one state obtains, then certain other states have a certain likelihood of obtaining. Both probabilistic and non-probabilistic laws determine how likely certain states are; non-probabilistic laws are just what you get when you turn that likelihood “all the way up”.

How well do the three constraint accounts do at vindicating this thought? On Adlam’s account, laws and chances are the same kind of thing – probability functions over sets of mosaics. Laws and chances differ only in the values they assign – non-probabilistic laws correspond to the functions that only assign values of 0 or 1. Likewise, on Meacham’s account, both probabilistic and non-probabilistic laws represent the nomic likelihoods of some states given others. Non-probabilistic laws just correspond to the special case where these likelihoods are maximized.

On Chen and Goldstein’s account, matters are more complicated – it depends on how we add chances to their account. As noted earlier (section 3.2), I’ll restrict my attention here to options 2 and 3. Option 2 takes probabilistic laws to be primitive facts that are distinct from the non-probabilistic laws – facts that have no direct connection

³⁰Adlam (2022b), section 5.5.

³¹Following Goodman (1955), the predicate *grue* applies to something *iff* it’s green and observed before t , or blue and not observed before t , and the predicate *bleen* applies to something *iff* it’s blue and observed before t , or green and not observed after t .

to the primitive constraints on nomic possibility imposed by the non-probabilistic laws. So on this option, probabilistic and non-probabilistic laws are not of a kind. Rather, probabilistic laws are distinct primitive facts that are added independently of the primitive constraints that yield non-probabilistic laws.

By contrast, option 3 adopts a gradable notion of constraining that can take values between 0 and 1, and takes probabilistic and non-probabilistic laws to just impose different grades of nomic constraint. On this approach, non-probabilistic laws are the ones which only impose extreme values of nomic constraint (0 or 1), while probabilistic laws are ones which impose some non-extremal values. So on this approach, probabilistic and non-probabilistic laws are of a kind. They're both nomic constraints, just ones that take different values.

4.4 Is Being Nomically Required/Forbidden Distinct From Having a Chance of 1/0?

There's a difference between possibilities that are nomically required and possibilities that have a chance of 1. For while it's plausible that every possibility that's nomically required has a chance of 1, not every possibility that has a chance of 1 is nomically required. For example, the chance of at least one of infinitely many fair coin tosses landing heads is 1, but this possibility is not nomically required – it's nomically possible that all these coins land tails. Likewise, *prima facie* there's a difference between possibilities that are nomically forbidden and possibilities that have a chance of 0 – for while every possibility that's nomically forbidden might have a chance of 0, there are possibilities with a chance of 0 that aren't nomically forbidden (e.g., every one of infinitely many fair coin tosses landing tails).

How well do the constraint accounts do at capturing this distinction? Adlam's account doesn't appear able to distinguish between such possibilities. For on this account, the only thing that distinguishes between probabilistic and non-probabilistic laws is whether they assign non-extremal values. And nothing about this account provides us with the tools to distinguish between those chance 1 possibilities that are nomically required and those that are not. Likewise, nothing about the account allows us to distinguish between those chance 0 possibilities that are nomically forbidden and those that are not.

Meacham's account, on the other hand, can distinguish between such possibilities. The representation and uniqueness theorem Meacham provides, coupled with his account of chances and nomic requirements, entails that while every possibility that's nomically required has a chance of 1, there can be possibilities with a chance of 1 that aren't nomically required. Likewise, while everything that's nomically forbidden has a chance of 0, there can be possibilities with a chance of 0 that aren't nomically forbidden.

On Chen and Goldstein's account, it will depend on how we decide to incorporate chances into the account. As before, I'll restrict my attention to option 2 (adding prim-

itive probabilistic facts that aren't directly related to the non-probabilistic law facts) and option 3 (adopting a degreed notion of constraining, where non-extremal values yield non-trivial chances, and extremal values (1 and 0) yield nomic requirements and forbiddings). Option 2 has the tools to distinguish between having a chance of 1 and being nomically required, since the facts that yield chances and the facts that yield nomic requirements are distinct posits that aren't reducible to one another. Likewise, this option has tools to distinguish between having a chance of 0 and being nomically forbidden. Option 3, however, seems unable to distinguish between such possibilities. Since having a chance of 1 and being nomically required are identified with the same primitive fact – a constraint with the maximum value – we can't distinguish between them. Likewise, we can't distinguish between having a chance of 0 and being nomically forbidden.

This consideration, together with the consideration discussed in section 4.3, suggest that proponents of Chen and Goldstein's account are going to face a choice about which bullet they want to bite. If they adopt option 2 for incorporating chances into the account, then they'll be able to distinguish between having a chance of 1/0 and being nomically required/forbidden, but won't be able to vindicate the thought that chances and nomic requirements/forbiddings are of a kind. Whereas if they adopt option 3 for incorporating chances into the account, then they'll be able to maintain that chances and nomic requirements/forbiddings are of a kind, but won't be able to distinguish between having a chance of 1/0 and being nomically required/forbidden.

4.5 What Vindicates Our Numerical Assignments to Chances?

We typically assign numbers to the different outcomes of chance events. For example, given a fair coin toss, we assign 0.5/0.5 to the coin landing heads/tails. And it seems like these assignments are *correct*; if we were to assign 0.6/0.4 to a fair coin landing heads/tails, we would be doing something wrong. But what makes these first assignments correct and these alternative assignments wrong? How does the structure of chance events vindicate assigning numbers to outcomes in the way that we do?³²

The most straightforward way to answer this question would be to posit fundamental relations between the outcomes of chance events and numbers – e.g., between a fair coin landing heads and the number 0.5. But this answer is implausible. It's a merely conventional choice to assign numbers between 0 and 1; we could assign numbers between 0 and 2 just as well. So what vindicates these numerical assignments can't be a relation between chance events and *particular numbers*. What's more plausible is that chance events have a kind of quantitative structure which is perspicuously represented by numbers. If that's right, then a satisfactory account of probabilistic laws and chances should tell us what that underlying quantitative structure is, and how it vindicates our assigning numbers to chances in the way that we do.

³²These types of questions are generally explored in the literature on measurement theory. For an overview of the relevant literature, see Eddon (2013).

How well do the different constraint accounts do at satisfying this desideratum? Meacham's account satisfies this desideratum. The account's nomic likelihood relations provide the quantitative structure underlying chances. The account's representation and uniqueness theorem shows that there is a way of assigning numbers in the $[0, 1]$ interval to triples that lines up with the nomic likelihood relation, that this assignment is unique, and that these assignments will satisfy the probability axioms. And since the account identifies chances with these assignments, it vindicates our assigning numbers to chances in the way that we do.

Adlam's account does not satisfy this desideratum. Adlam's account takes laws to correspond to probability functions over sets of Humean mosaics, and probabilistic laws to correspond to the subset of such functions that assign non-extremal values. In some places Adlam's discussion suggests that we might *identify* the laws with such functions, in other places Adlam merely suggests that the laws "induce" such functions. Let's consider each of these possibilities in turn. If we identify the laws with such functions, then it's a primitive fact that these numbers are assigned to the corresponding sets. As we've seen, this kind of view is implausible. Alternatively, we might take the laws to merely induce such functions, and take our primitive metaphysical posits to consist of something like fundamental non-numerical relations that encode the quantitative structure of chances. But then the account needs to tell us what these metaphysical posits are, and how they vindicate our numerical assignments to chances. Adlam's account, as presented, does neither.

For similar reasons, Chen and Goldstein's account fails to satisfy this desideratum. If we incorporate chances into the account via option 3 – adopting a degreed notion of constraining that takes values between 0 and 1 – then their account is in a similar position to the first version of Adlam's account, appearing to ground the chance facts in something like primitive relations to numbers. If we incorporate chances into the account via option 2 – adding probabilistic laws as independent primitive facts – then their account can avoid positing implausible relations to numbers by taking these primitive facts to consist of non-numerical relations that ground the quantitative structure of chances. But then their account is in a similar position to the second version of Adlam's account, needing to tell us what this metaphysical structure is, and how it vindicates our numerical assignments to chances.

It's natural to wonder whether one could add something like Meacham's approach to Adlam and Chen and Goldstein's accounts in order to satisfy this desideratum. As it turns out, it is difficult for these accounts to incorporate something like the representation theorem framework employed by Meacham without giving up some of the distinctive ways in which these accounts differ from Meacham's account: taking there to be objective facts about how to carve up the complete laws, and taking laws to be fine-grained. Let's see why.

First, consider the issue of whether there's an objective way to carve the complete laws into individual laws. Representation and uniqueness theorems work holistically. These theorems take all of the instances of a certain kind of relational fact (in this case,

instantiations of the nomic likelihood relation) and tell us that they can be uniquely represented in a certain way. But these theorems generally won't work if we only input part of the relevant relational facts. Likewise, they won't output parts of a representation – they'll output a single unified representation. So it's hard to see how a representation theorem approach could yield the division into individual laws that Adlam's account requires.

Second, consider the issue of whether there are fine-grained laws. The representation theorem Meacham employs, like the variants one can find in the literature on measurement theory, assigns a representation to an algebra defined over a space of elements. So the relevant representational statuses (having a certain chance, being nomically required/forbidden) are assigned to sets of these elements. The straightforward way to interpret this framework is to think of the elements as possible worlds, and take these sets of them to correspond to (coarse-grained) propositions. But this yields objects of representation that are too coarse-grained to capture the kind of fine-grained distinctions that, e.g., Chen and Goldstein want to capture. To get these theorems to apply to fine-grained propositions, the objects of representation would need to correspond to sets of something more fine-grained than possible worlds: impossible worlds. But this replacement renders the theorem uninteresting, since the axioms of these theorems and the representations they yield will no longer mean what we want them to mean.

To see why, consider the most straightforward way of making this replacement. Take some appropriately expressive (interpreted) language and let each set of sentences in this language correspond to an impossible world; intuitively, this set will correspond to all the sentences that are true at that (impossible) world. So while a sentence S_1 expressed in terms of green and blue will be true at the same set of possible worlds as the corresponding sentence S_2 expressed in terms of grue and bleen, those two sentences will pick out different sets of *impossible* worlds – there will be some impossible worlds with S_1 but not S_2 , and some impossible worlds with S_2 but not S_1 . There aren't any formal obstacles to applying the axioms to these new elements (impossible worlds), and obtaining a representation regarding sets of these elements. After all, the only thing that's changed formally is that we've replaced one set of elements with another, and nothing about the theorem cares about what the elements we're working with are.

But this replacement undercuts the interest of the theorem, because the axioms and results will no longer mean what we want them to mean. Consider, for example, the first axiom of Meacham's approach. Let S stand for the set of impossible worlds containing S . The first axiom of Meacham's approach entails (among other things) that if S is assigned a likelihood, then \bar{S} is as well. Meacham motivates this axiom by noting that, intuitively, if there's some nomic likelihood of it raining, then there should also be some nomic likelihood of it not raining. And if S corresponds to a set of possible worlds where S is true, then the axiom vindicates this intuition, since the set of worlds at which it's raining (S is true) is the complement of the set of worlds

where it's not raining ($\neg S$ is true). But if S corresponds to the set of *impossible* worlds where it's true, this motivation no longer works. Because the set of worlds at which it's raining (S is true) is not the complement of the set of worlds where it's not raining ($\neg S$ is true); there are impossible worlds containing both S and $\neg S$, and impossible worlds containing neither. Thus the complement of the set of worlds where it's raining doesn't correspond to anything in particular, and the axiom's insistence that this complement be assigned a nomic likelihood is unmotivated.

For another example, consider one of the results of the representation theorem, which says that the relevant sets of elements will be assigned values which satisfy the probability axioms, such as additivity. Additivity tells us that if S_1 and S_2 are disjoint, then the probability of $S_1 \cup S_2$ is equal to the sum of the individual probabilities assigned to S_1 and S_2 . So if these elements are possible worlds, then additivity entails that the probability of it raining or it snowing is equal to the probability of it raining plus the probability of it snowing. But if S_1 and S_2 correspond to sets of *impossible* worlds, this is no longer true. For $S_1 \cup S_2$ won't pick out the set of worlds where $S_1 \vee S_2$ is true – there are impossible worlds where S_1 or S_2 is true, but $S_1 \vee S_2$ is not. So additivity *won't* entail that that probability of it raining or it snowing is equal to the probability of it raining plus the probability of it snowing.

4.6 What Constraints Are There on Laws and Chances?

We typically take there to be constraints that require laws to line up with other laws, chances to line up with other chances, and laws and chances to line up with each other. For example, we think that laws that hold at the same world must be compatible with each other – if one law entails that all masses have charge, there can't be another law that entails that there are charge-less masses. Likewise, we think that chance distributions that hold at the same world must line up with each other in certain ways – if the chance at t_1 of the next coin toss landing heads is $1/2$, and the chance of the next two coin tosses landing heads is $1/4$, and at t_2 the first coin toss does land heads, then the chance at t_2 of the second coin toss landing heads must be $1/2$. And we think that laws and chances must line up with each other as well – an event with a chance of $1/2$ can't be nomically required.

This suggests another desideratum for a satisfactory account of laws and chances: a satisfactory account should yield the kinds of constraints on laws and chances that we typically take to obtain. Since “the constraints that we typically take to obtain” is somewhat vague, this desideratum is somewhat vague as well. But even without a precise list of constraints in mind, we can prefer theories that provide us with plausible restrictions of this kind over those that do not. So how do the constraints views we've looked at do at providing us with such restrictions?

Let's start with Adlam's account. On Adlam's account, both probabilistic and non-probabilistic laws correspond to probability functions that assign values to sets of Humean mosaics. In the case of non-probabilistic laws, Adlam's account requires

the actual world to be a member of the intersection of the Humean mosaics the non-probabilistic laws allow. This requires non-probabilistic laws to be consistent with each other.

Turning to probabilistic laws, Adlam tells us to think of the probabilities assigned to sets of mosaics by probabilistic laws as representing the likelihood that each set of mosaics will be “chosen”. Adlam’s account requires the actual world to be a member of the set “chosen” for each law. Thus each set that might be “chosen” for one probabilistic law must be compatible with every set that might be “chosen” for every other probabilistic law; i.e., the probabilistic laws that obtain at a world must have outcomes that partition each other. Likewise, every outcome of these probabilistic laws must be compatible with the set of mosaics picked out by every non-probabilistic law that holds at that world. So Adlam’s account also imposes some important consistency constraints on probabilistic laws, and on which probabilistic and non-probabilistic laws can obtain at the same world.

That said, Adlam’s account doesn’t tell us as much about probabilistic laws as one might like. One issue is that the account doesn’t tell us how the probabilities of different probabilistic laws are related to one another. Suppose one probabilistic law assigns a chance of 0.5 to A and 0.5 to $\neg A$, while another assigns a chance of 0.3 to B and 0.7 to $\neg B$. What is the chance of $A \wedge B$? Since Adlam’s account doesn’t tell us how to combine the chances of different laws, it doesn’t provide us with an answer to this question.³³

Another issue arises when we try to make precise the requirement that each set that might be “chosen” for one probabilistic law must be compatible with each set that might be “chosen” for every other probabilistic law. What does this requirement amount to? A natural thought is that the each set must be assigned a non-zero probability by their law (so that they might be “chosen”), and must have a non-empty intersection (so they’re compatible). But this thought runs into two problems. First, unless we know how the probabilities of each law are related to one another, the fact that each set is assigned a non-zero probability by their respective laws doesn’t entail that their intersection has a non-zero probability; that would require some further assumption, such as that different laws are probabilistically independent. Second, it’s unable to account for probabilistic laws that assign a chance of 0 to each of a continuum of outcomes (e.g., the chance of a dart landing at any particular point on a dart board). For the above thought would entail that no outcome of such a law could be “chosen”.³⁴

³³There are a couple natural answers available here. For example, we might assume that these probability functions are independent, and take the chance of $A \wedge B$ to be the product of the values assigned by each probabilistic law. Or we might take the chance of $A \wedge B$ to be undefined if it’s not directly provided by one of the probabilistic laws of the world. (Thus if we think $A \wedge B$ has a well-defined chance, we have reason to think that the two probabilistic laws described above can’t be the only probabilistic laws that obtain.) But this is something that we would like an account to tell us.

³⁴A natural thought is to appeal to probability densities here. But this requires more structure than the

A third issue that arises has to do with how Adlam's account incorporates probabilistic laws. On Adlam's account, probabilistic laws correspond to probability functions that assign numbers directly to sets of Humean mosaics. This effectively makes them chance distributions with one argument – that argument being the possibility a chance value is assigned to. But this makes it hard to make sense of the kinds of probabilistic laws posited by our physical theories, which posit chance distributions which take two arguments – the possibility the chance value is assigned to, and the background state that picks out the chance distribution.

For example, consider a toy time-dependent chance theory that assigns chances to coin tosses. Suppose that the chance of a coin toss landing heads at t_1 is 0.5, and that after that coin lands heads, the chance of that coin toss landing heads (at t_2) is 1. Since these are two different probability distributions, they need to correspond to different laws on Adlam's account. But these two probability functions seem to assign values that are incompatible, and so seem to correspond to probabilistic laws that can't obtain at the same world. In a similar vein, statistical mechanical probabilities vary depending on the macrostate they're assessed with respect to, and the corresponding probability functions will generally assign incompatible values. But if they assign incompatible values, they can't correspond to probabilistic laws that obtain at the same world.³⁵

Adlam's account prides itself in being neutral with respect to a number of issues. In some ways this lack of detailed commitments is a strength – it broadens the range of potential buyers. But it's also a weakness – it leaves the account unable to provide us with the constraints on laws and chances that we typically take to obtain.

Turning to Chen and Goldstein's account, they require the world that instantiates a (non-probabilistic) law to belong to the set of models generated by that law. This imposes a consistency constraint on the non-probabilistic laws instantiated at a world. But because their account is agnostic about so many other details, it says little else about what kinds of laws and chances there can be. For example, we haven't been told anything that ensures that different probabilistic laws will make consistent

account currently provides. E.g., densities require well-defined reference measures over sets of Humean mosaics (the probability density over what?), and these densities need to be combinable, which requires (among other things) telling us something about how these different probabilities are related.

³⁵One might try to fix this problem by taking only one of these probability functions as "real" (perhaps the initial one, or the one corresponding to the maximal macrostate, or what have you), and then take the rest of these probabilities to be just what you get when you conditionalize these "ur-chance" functions on some further facts. This won't work in all cases, however; there is no probabilistic "ur-chance" function to appeal to in the case of statistical mechanics, for example (see Meacham (2005)).

A satisfactory treatment of this problem requires modifying Adlam's account in some way, such as having it posit probability functions with two arguments, an "object" argument that it assigns values to and a "background" argument that picks out the relevant facts that the probability function is assigned to. Of course, such moves highlight the need for the account to impose further constraints, such as constraints requiring chance distributions with different background arguments to line up with each other in certain ways.

assignments. We haven't been told how to combine the chance assignments from different probabilistic laws. We haven't been told anything that requires different chance distributions at a world to line up with each other. And so on. So, like Adlam's account, Chen and Goldstein's account does little by way of imposing constraints on laws and chances.

What about Meacham's account? Because Meacham's account commits itself to more details than Adlam's or Chen and Goldstein's accounts, it's better placed to provide us with the kinds of constraints on laws and chances that we'd like an account to provide. In particular, from the constraints the account imposes on the nomic likelihood relation, one can derive a number of the constraints on laws and chances that we typically take to obtain. For example, one can derive that if C is nomically required then it will have a chance of 1, and if it's nomically forbidden it will have a chance of 0. One can derive that if C is nomically required if some state A obtains, and A does obtain, then C will be true. One can derive that, given certain assumptions, different chance distributions at a world will be related by conditionalization. And so on.³⁶

5 An Attractive Constraint Account

I think an adequate account of laws and chances should be able to accommodate the kinds of laws and possibilities that physicists have taken seriously. And since all three constraint accounts can do so, and their mainstream competitors cannot, I take these three constraint accounts to be among the most attractive options on the market.

Of these constraint accounts, I find Meacham's account to be the most appealing. Meacham's account is the only one that vindicates our assigning numbers to chances in the way that we do, the only one that yields the kinds of constraints on laws and chances that we typically take to obtain, and the only one that can both treat laws and chances as a kind and distinguish between nomic requirements/forbiddings and having a chance of 1/0. Of course, Meacham's account also holds that there's no objective way to carve the complete laws into individual laws, and holds that laws are relatively coarse-grained. Whether these features of Meacham's account are pros or cons will depend on one's philosophical sensibilities, but I'm inclined to think these are plausible verdicts.

That said, the default version of Meacham's account has some unattractive features. First, the account entails that the laws aren't intrinsic features of the worlds that instantiate them. This is because the laws of a world are determined by the nomic likelihood relations involving that world, and these relations hold between the world

³⁶I've argued that Adlam's and Chen and Goldstein's accounts don't yield the kinds of constraints on laws and chances that we typically take to obtain. But these are largely sins of omission; difficulties that arise from the lack of detail these accounts provide. Thus both accounts could address this issue by adding further details to the account. How satisfactory the resulting accounts would be will, of course, depend on how these details are added.

and things external to the world, such as other worlds. In a similar vein, this account entails that a perfect duplicate of a world needn't have the same laws, and that the laws aren't qualitative features of the world.³⁷ These are implausible consequences.³⁸ Second, by postulating a fundamental relation that holds between worlds and propositions, it adopts an ontological commitment to such things. It would be preferable to have an account of laws without such commitments.

In light of this, I suggest we adopt a different version of Meacham's account.³⁹ The default version of Meacham's account focuses on triples $\langle A, C, w \rangle$ consisting of a world and a pair of propositions. It posits a fundamental nomic likelihood relation that holds between such triples (intuitively, the relation that holds when the likelihood of A given C at world w is at least as great as the likelihood of A' given C' at world w'). And it identifies laws with the property of being a world bearing certain nomic likelihood relations.

This alternative version starts by positing a number of mutually exclusive fundamental properties of worlds, L_i . It replaces the triples $\langle A, C, w \rangle$ with triples of properties $\langle \mathbf{A}, \mathbf{C}, L_i \rangle$, where \mathbf{A} is a property that holds of a world *iff* that world makes A true, \mathbf{C} is a property that holds of a world *iff* that world makes C true, and L_i is one of the fundamental properties mentioned above. It posits a fundamental nomic likelihood relation that holds between such triples. And it identifies the different complete laws with the different L_i properties.⁴⁰ (So on this view, intuitively, the nomic likelihood relation holds when the likelihood of \mathbf{A} given \mathbf{C} given laws L_i is at least as great as the likelihood of \mathbf{A}' given \mathbf{C}' given laws L_j .)

This alternative version of Meacham's account inherits all of the advantages of the default version of Meacham's account described above.⁴¹ But it also allows us to

³⁷This is assuming we adopt something like Lewis's (1983) accounts of intrinsic properties, duplication, and qualitative properties. In section 7, Meacham (2023) considers some alternative characterizations that would yield different verdicts.

³⁸What about Adlam's and Chen and Goldstein's accounts? On their accounts will the laws be intrinsic, qualitative features of the world that are preserved by duplication? Because their accounts are agnostic about many of the relevant details, it's unclear.

³⁹This alternative is the marriage of two potential variants that Meacham (2023) discusses in section 7: a variant which adopts a two-layer view of fundamental nomic properties, and a variant which takes properties to be the relata of the nomic likelihood relation.

⁴⁰As Meacham (2023) notes in section 7, one might complain that a view like this is similar to a primitivist account of laws. (I thank an anonymous referee for encouraging me to address this concern.) But this is not a special worry for this account; all the proponents of "constraint" views are sympathetic to primitivism. (Chen & Goldstein (2022) even call their view "Minimal Primitivism"). The question of interest here is which view of this kind is most attractive.

⁴¹This alternative is also subject to several complaints one might raise about the default version of Meacham's account. (I thank an anonymous referee for encouraging me to address these concerns.) In particular, one might complain that these accounts (1) don't restrict the content of laws to things that might be easily expressed in terms of perfectly natural properties, (2) don't tie the chances to frequencies, and (3) don't provide a justification for the Principal Principle. (1) is true of all of the constraint accounts discussed in this paper, but this is a feature, not a bug. This is because they want to allow for laws like the Past Hy-

maintain that laws are intrinsic and qualitative features of worlds that are preserved by duplication. And it only commits us to a single ontological category – properties.

The main demerit of this alternative version of Meacham’s account is that it needs to posit a collection of fundamental L_i properties. I suspect people will differ on whether they take this alternative version’s benefits to be worth this cost. But I’m inclined to think they are.

6 Conclusion

The most popular accounts of laws in the literature share a common problem – they’re unable to accommodate many of laws and possibilities seriously considered by the scientific community. Recently, a number of new accounts of laws – what I’ve called Global Primitive Constraint accounts – have been proposed that don’t have this problem. I’ve presented a version of a constraint account that I think is particularly attractive. Others may find some other constraint account more to their liking. Either way, these constraint accounts are a welcome addition to the literature on laws.⁴²

References

- Adlam, E. (2022a). Determinism beyond time evolution. *European Journal for Philosophy of Science*, 12(4), 1–36.
- Adlam, E. (2022b). Laws of nature as constraints. *Foundations of Physics*, 52(1), 1–41.
- Aharonov, Y., Bergmann, P. G., & Lebowitz, J. L. (1964). Time symmetry in the quantum process of measurement. *Phys. Rev.*, 134, B1410–B1416.
URL <https://link.aps.org/doi/10.1103/PhysRev.134.B1410>
- Albert, D. Z. (2000). *Time and Chance*. Harvard University Press.
- Armstrong, D. (1983). *What is a Law of Nature?*. Cambridge, United Kingdom: Cambridge University Press.
- Barrett, J. A., & Chen, E. K. (manuscript). Algorithmic randomness and probabilistic laws. ArXiv preprint:2303.01411.

pothesis whose content can’t easily be expressed in terms of perfectly natural properties. Likewise, (2) is a feature, not a bug. Since we want to accommodate the non-Humean intuition that every outcome that gets a positive chance is a possible outcome, we don’t want to directly tie chances to frequencies. Finally, (3) is true, but this isn’t something that one generally expects an account of laws and chances to provide. E.g., none of the popular accounts discussed in section 2.1 provide a justification for the Principal Principle.

⁴²I’d like to thank Emily Adlam, Eddy Chen, Maya Eddon, three anonymous referees, attendees of the 2022 UMass Brown Bag, and attendees of the 2023 Society for the Metaphysics of Science conference, for helpful comments and discussion.

- Bigelow, J., Ellis, B., & Lierse, C. (1992). The world as one of a kind: Natural necessity and laws of nature. *British Journal for the Philosophy of Science*, 43(3), 371–388.
- Bird, A. (2007). *Nature's Metaphysics: Laws and Properties*. Oxford University Press.
- Chen, E. K., & Goldstein, S. (2022). Governing without a fundamental direction of time: Minimal primitivism about laws of nature. In Y. Ben-Menahem (Ed.) *Rethinking the Concept of Law of Nature*, (pp. 21–64). Cham: Springer.
- Coates, A. (manuscript). The grounding conception of governance.
- Eddon, M. (2007). Armstrong on quantities and resemblance. *Philosophical Studies*, 136(3), 385–404.
- Eddon, M. (2013). Quantitative properties. *Philosophy Compass*, 8(7), 633–645.
- Eddon, M. (2017). Parthood and naturalness. *Philosophical Studies*, 174(12), 3163–3180.
- Emery, N. (2019). Laws and their instances. *Philosophical Studies*, 176(6), 1535–1561.
- Field, H. H. (1980). *Science Without Numbers: A Defence of Nominalism*. Princeton, NJ, USA: Princeton University Press.
- Forge, J. (1999). Laws of nature as relations between quantities? In H. Sankey (Ed.) *Causation and Laws of Nature*, (pp. 111–124). Kluwer Academic Publishers.
- Gödel, K. (1949). An example of a new type of cosmological solutions of einstein's field equations of gravitation. *Rev. Mod. Phys.*, 21, 447–450.
URL <https://link.aps.org/doi/10.1103/RevModPhys.21.447>
- Goodman, N. (1955). Fact, fiction, and forecast. *Philosophy*, 31(118), 268–269.
- Hoefer, C. (2018). *Chance in the World: A Humean Guide to Objective Chance*. New York, NY, USA: Oup Usa.
- Ioannidis, S., Livianos, V., & Psillos, S. (2021). No laws and (thin) powers in, no (governing) laws out. *European Journal for Philosophy of Science*, 11(1), 1–26.
- King, J. C. (2019). Structured Propositions. In E. N. Zalta (Ed.) *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Summer 2019 ed.
- Lange, M. (2009). *Laws and Lawmakers Science, Metaphysics, and the Laws of Nature*. Oxford University Press.
- Lewis, D. (1983). New work for a theory of universals. In D. H. Mellor, & A. Oliver (Eds.) *Properties*. Oxford University Press.

- Lewis, D. (1994). Humean supervenience debugged. *Mind*, 103(412), 473–490.
- Lewis, D. K. (1973). *Counterfactuals*. Cambridge, MA, USA: Blackwell.
- Maudlin, T. (2007). *The Metaphysics Within Physics*. Oxford University Press.
- Meacham, C. J. G. (2005). Three proposals regarding a theory of chance. *Philosophical Perspectives*, 19(1), 281–307.
- Meacham, C. J. G. (2010). Contemporary approaches to statistical mechanical probabilities: A critical commentary - part i: The indifference approach. *Philosophy Compass*, 5(12), 1116–1126.
- Meacham, C. J. G. (2023). The nomic likelihood account of laws. *Ergo: An Open Access Journal of Philosophy*, 9(9), 230–284.
- Ott, W. (2022). *The Metaphysics of Laws of Nature: The Rules of the Game*. Oxford, UK: Oxford University Press.
- Tooley, M. (1977). The nature of laws. *Canadian Journal of Philosophy*, 7(4), 667–98.
- Tooley, M. (1987). *Causation: A Realist Approach*. Oxford University Press.
- Van Fraassen, B. C. (1989). *Laws and Symmetry*. Oxford, England: Oxford University Press.
- Vetter, B. (2012). Dispositional essentialism and the laws of nature. In A. Bird, B. Ellis, & H. Sankey (Eds.) *Properties, Powers, and Structures: Issues in the Metaphysics of Realism*. Routledge.
- Wheeler, J., & Feynman, R. (1945). Absorber theory and the radiation arrow of time. *Review of Modern Physics*, 17, 157–181.
- Wilsch, T. (2021). The governance of laws of nature: Guidance and production. *Philosophical Studies*, 178(3), 909–933.