Algorithmic Randomness and Probabilistic Laws

Jeffrey A. Barrett* Eddy Keming Chen†

*Department of Logic and Philosophy of Science, University of California, Irvine

†Department of Philosophy, University of California, San Diego

Preprint: arXiv 2303.01411

Introduction

The concept of probability

- has many aspects;
- related to beliefs and credences;
- related to the world.

Introduction

Our focus:

- Probabilistic laws of nature.
- Those laws of nature that involve probabilities in their statements.
- Examples: probabilistic laws in quantum mechanics and statistical mechanics.

Introduction

We develop a new conception of probability, and thereby develop a new approach to probabilistic laws.

Introduction

We develop a new conception of probability, and thereby develop a new approach to probabilistic laws.

Key: algorithmic randomness.

Introduction

We develop a new conception of probability, and thereby develop a new approach to probabilistic laws.

Key: algorithmic randomness.

Relations to conceptual issues about probability:

- scientific explanation,
- randomness,
- independence,
- typicality,
- objective probability,
- Principal Principle,
- perhaps Cournot's Principle.

Some background:

- Jeff: foundations of quantum mechanics
- Eddy: laws of nature; typicality; Cournot's principle.

Probabilistic laws, as they are usually understood, involve a variety of underdetermination.

Probabilistic laws, as they are usually understood, involve a variety of underdetermination.

A simple example:

Introduction

Probabilistic laws, as they are usually understood, involve a variety of underdetermination.

A simple example:

- Repeated tosses of a coin
- It produces: an infinite ω-sequence of results $\langle r_1, r_2, \ldots \rangle$.
- Each possible sequence describes a possible world.
- Let Ω^L be the set of all such worlds that accord with a law L.

Consider the probabilistic law L:

Probabilistic Law L

Each element in the ω-sequence of coin tosses $\langle r_1, r_2, \ldots \rangle$ is determined independently and with an unbiased probability of heads and tails.

Consider the probabilistic law L:

Probabilistic Law L

Each element in the ω-sequence of coin tosses $\langle r_1, r_2, \ldots \rangle$ is determined independently and with an unbiased probability of heads and tails.

One might think of L as descriptive of a fundamentally random process, something like starting with a sequence of spin-1/2 particles each in a eigenstate of z-spin, then measuring their x-spins in turn.

- As probabilistic laws are typically understood, Ω^L is the set of all ω-sequences.

Introduction

- As probabilistic laws are typically understood, Ω^L is the set of all ω-sequences.
- L does not rule out any world.

- As probabilistic laws are typically understood, Ω^L is the set of all ω-sequences.
- L does not rule out any world.
- A world compatible with L might exhibit any limiting relative frequency, or no limiting relative frequency at all.

Introduction

- As probabilistic laws are typically understood, Ω^L is the set of all ω-sequences.
- L does not rule out any world.
- A world compatible with L might exhibit any limiting relative frequency, or no limiting relative frequency at all.
- Examples: the all-heads world; a world with $1/3$ heads.

Introduction

- As probabilistic laws are typically understood, Ω^L is the set of all ω-sequences.
- L does not rule out any world.
- A world compatible with L might exhibit any limiting relative frequency, or no limiting relative frequency at all.
- Examples: the all-heads world; a world with $1/3$ heads.
- Even the full history of a world will fail to determine L, in a continuous cardinality of cases.

Introduction

- As probabilistic laws are typically understood, Ω^L is the set of all ω-sequences.
- L does not rule out any world.
- A world compatible with L might exhibit <u>any</u> limiting relative frequency, or no limiting relative frequency at all.
- Examples: the all-heads world; a world with $1/3$ heads.
- Even the full history of a world will fail to determine L, in a continuous cardinality of cases.
- Even Ω^L, the full set of worlds compatible with L, does nothing to determine L over any other probabilistic law.

The underdetermination is closely related to empirical coherence.

Introduction

The underdetermination is closely related to empirical coherence.

Empirical Coherence

A physical law is empirically coherent only if it is always in principle possible for one to have empirical support for the law if the law is in fact true.

cf: Barrett (1996), (1999) and (2020)

If a law is empirically incoherent....

Introduction

If a law is empirically incoherent....

- It may be impossible to learn that the law is true with even complete evidence.

If a law is empirically incoherent....

- It may be impossible to learn that the law is true with even complete evidence.
- There are ω-sequences that might occur but would provide no empirical evidence whatsoever for accepting L.

Introduction

If a law is empirically incoherent....

- It may be impossible to learn that the law is true with even complete evidence.
- There are ω-sequences that might occur but would provide no empirical evidence whatsoever for accepting L.
- Example: the all-heads sequence.

Introduction

If a law is empirically incoherent....

- It may be impossible to learn that the law is true with even complete evidence.
- There are ω-sequences that might occur but would provide no empirical evidence whatsoever for accepting L.
- Example: the all-heads sequence.
- In such worlds one would never have any empirical support for accepting the correct probabilistic law even with full evidence.

If a law is empirically incoherent....

- It may be impossible to learn that the law is true with even complete evidence.
- There are ω-sequences that might occur but would provide no empirical evidence whatsoever for accepting L.
- Example: the all-heads sequence.
- In such worlds one would never have any empirical support for accepting the correct probabilistic law even with full evidence.
- There is a continuous cardinality of such worlds.

We suggest: one can get a tighter fit between probabilistic laws and empirical evidence.

Introduction

We suggest: one can get a tighter fit between probabilistic laws and empirical evidence.

How? By appealing to a stronger conception of probability and a correspondingly stronger variety of probabilistic laws.

We suggest: one can get a tighter fit between probabilistic laws and empirical evidence.

How? By appealing to a stronger conception of probability and a correspondingly stronger variety of probabilistic laws.

Probabilistic Law $L^\star$

The ω-sequence of coin tosses $\langle r_1, r_2, \ldots \rangle$ is random with unbiased relative frequencies of heads and tails.

Introduction

We suggest: one can get a tighter fit between probabilistic laws and empirical evidence.

How? By appealing to a stronger conception of probability and a correspondingly stronger variety of probabilistic laws.

Probabilistic Law $L^\star$

The ω-sequence of coin tosses $\langle r_1, r_2, \ldots \rangle$ is random with unbiased relative frequencies of heads and tails.

- Here being random is a property of the ω-sequence. (Outcome randomness)
- The notion should be defined.

Introduction

- The notions of randomness we consider here are algorithmic.

- The notions of randomness we consider here are algorithmic.
- ... defined in terms of statistical tests that determine whether an ω-sequence exhibits any specifiable pattern.

Introduction

- The notions of randomness we consider here are algorithmic.
- ... defined in terms of statistical tests that determine whether an ω-sequence exhibits any specifiable pattern.
- Each sequence will either pass or fail a particular test for being random.

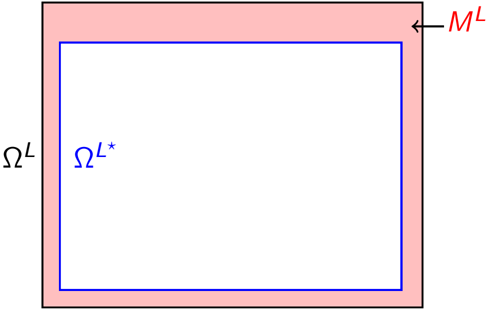

- While L is compatible with all ω-sequences of results, $L^\star$ is not.
- Let $\Omega^{L^\star}$ be the set of all worlds that accord with the law $L^\star$.
- All worlds in $\Omega^{L^\star}$ exhibit the random unbiased sequences stipulated by $L^\star$.

Introduction

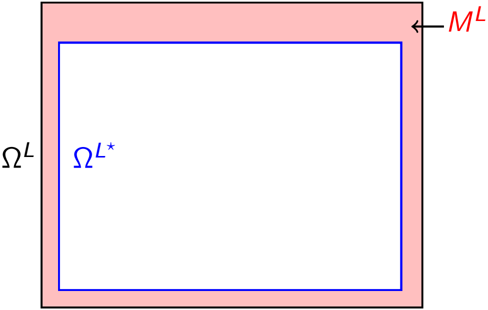

- $\Omega^{L^\star}$ contains no maverick worlds
- Maverick worlds are those that exhibit a specifiable pattern or fail to exhibit the right relative frequencies or fail to exhibit any relative frequencies at all.

Introduction

- If $L^{\star}$ is true, then any physically possible world fully determines $L^{\star}$.

Introduction

- If $L^{\star}$ is true, then any physically possible world fully determines $L^{\star}$.
- A $L^{\star}$ law is empirically coherent in the relevant sense, while a L law is not.

A nice implication:

A nice implication:

- No special probabilistic background assumptions or priors regarding what world one inhabits are required for successful inquiry.

A nice implication:

- No special probabilistic background assumptions or priors regarding what world one inhabits are required for successful inquiry.
- (The Principal Principle and Cournot's Principle are sometimes used for this purpose.)

A nice implication:

- No special probabilistic background assumptions or priors regarding what world one inhabits are required for successful inquiry.
- (The Principal Principle and Cournot's Principle are sometimes used for this purpose.)
- An inquirer in any physically possible world might determine the truth of $L^{\star}$ by considering the results of coin tosses.

A nice implication:

- No special probabilistic background assumptions or priors regarding what world one inhabits are required for successful inquiry.
- (The Principal Principle and Cournot's Principle are sometimes used for this purpose.)
- An inquirer in any physically possible world might determine the truth of $L^\star$ by considering the results of coin tosses.
- With complete evidence, one will surely learn it up to an equivalence class of computationally indistinguishable laws.

A nice implication:

- No special probabilistic background assumptions or priors regarding what world one inhabits are required for successful inquiry.
- (The Principal Principle and Cournot's Principle are sometimes used for this purpose.)
- An inquirer in any physically possible world might determine the truth of $L^\star$ by considering the results of coin tosses.
- With complete evidence, one will surely learn it up to an equivalence class of computationally indistinguishable laws.
- A probabilistic law like $L^\star$ is much like a deterministic law in this way.

Plan for this talk:

Introduction

Plan for this talk:

1. We define the randomness constraint that every $L^\star$ world must satisfy.

Introduction

Plan for this talk:

1. We define the randomness constraint that every $L^\star$ world must satisfy.
2. We consider two understandings of a $L^\star$ law.

Introduction

Plan for this talk:

1. We define the randomness constraint that every $L^{\star}$ world must satisfy.
2. We consider two understandings of a $L^{\star}$ law.
3. We discuss the costs and benefits.

Plan for this talk:

1. We define the randomness constraint that every $L^\star$ world must satisfy.
2. We consider two understandings of a $L^\star$ law.
3. We discuss the costs and benefits.
4. We argue that the second understanding has salient virtues.

Randomness Constraints

Let's define the randomness constraint.

Randomness Constraints

Let's define the randomness constraint.

- We need a test of randomness for ω-sequences.

Randomness Constraints

Let's define the randomness constraint.

- We need a test of randomness for ω-sequences.
- A random sequence of tosses with an unbiased coin should exhibit an even relative frequency of heads and tails in the limit.

Randomness Constraints

Let's define the randomness constraint.

- We need a test of randomness for ω-sequences.
- A random sequence of tosses with an unbiased coin should exhibit an even relative frequency of heads and tails in the limit.
- But this is **not sufficient**!

Randomness Constraints

Let's define the randomness constraint.

- We need a test of randomness for ω-sequences.
- A random sequence of tosses with an unbiased coin should exhibit an even relative frequency of heads and tails in the limit.
- But this is **not sufficient**!
- Example of a non-random sequence: an alternating sequence of heads and tails.

There are three further conditions for an unbiased random sequence:

Randomness Constraints

There are three further conditions for an unbiased random sequence:

1. patternless;

Randomness Constraints

There are three further conditions for an unbiased random sequence:

1. patternless;
2. generic / typical;

Randomness Constraints

There are three further conditions for an unbiased random sequence:

1. patternless;
2. generic / typical;
3. not allowing for a successful fair betting strategy.

Randomness Constraints

There are three further conditions for an unbiased random sequence:

1. patternless;
2. generic / typical;
3. not allowing for a successful fair betting strategy.

cf: Li and Vitányi (2008), A. Dasgupta (2011), Barrett and Huttegger (2021), Eagle (2021).

Let's start with the first condition: an ω-sequence is random only if it exhibits no finitely specifiable regularity.

Randomness Constraints

Let's start with the first condition: an ω-sequence is random only if it exhibits no finitely specifiable regularity.

- Algorithmic tests are helpful for characterizing "patternless."

Randomness Constraints

Let's start with the first condition: an ω-sequence is random only if it exhibits no finitely specifiable regularity.

- Algorithmic tests are helpful for characterizing "patternless."
- First pass: an ω-sequence is patternless, and hence random, if and only if there is no finite-length algorithm that produces the sequence.

Randomness Constraints

Let's start with the first condition: an ω-sequence is random only if it exhibits no finitely specifiable regularity.

- Algorithmic tests are helpful for characterizing "patternless."
- First pass: an ω-sequence is patternless, and hence random, if and only if there is no finite-length algorithm that produces the sequence.
- A finite-length algorithm would express a regularity, something that one might even think of as a deterministic law.

Randomness Constraints

But the first-pass criterion is insufficient:

Randomness Constraints

But the first-pass criterion is insufficient:

- Example: a repeated three-block pattern of one thousand heads followed by one thousand tails followed by one thousand random and unbiased heads and tails.

Randomness Constraints

But the first-pass criterion is insufficient:

- Example: a repeated three-block pattern of one thousand heads followed by one thousand tails followed by one thousand random and unbiased heads and tails.
- The relative frequency of heads and tails in the full sequence is unbiased.

Randomness Constraints

But the first-pass criterion is insufficient:

- Example: a repeated three-block pattern of one thousand heads followed by one thousand tails followed by one thousand random and unbiased heads and tails.
- The relative frequency of heads and tails in the full sequence is unbiased.
- The sequence cannot be represented by a finite-length algorithm.

Randomness Constraints

But the first-pass criterion is insufficient:

- Example: a repeated three-block pattern of one thousand heads followed by one thousand tails followed by one thousand random and unbiased heads and tails.
- The relative frequency of heads and tails in the full sequence is unbiased.
- The sequence cannot be represented by a finite-length algorithm.
- But the sequence is not random!

Randomness Constraints

But the first-pass criterion is insufficient:

- Example: a repeated three-block pattern of one thousand heads followed by one thousand tails followed by one thousand random and unbiased heads and tails.
- The relative frequency of heads and tails in the full sequence is unbiased.
- The sequence cannot be represented by a finite-length algorithm.
- But the sequence is not random!
- We can design a gambling strategy and enjoy unbounded wealth.

- Problem: there is no bound on the amount that a finite initial segment of this sequence might be compressed.

Randomness Constraints

- Problem: there is no bound on the amount that a finite initial segment of this sequence might be compressed.
- One might write a very short program that takes advantage of the regularity.

- Problem: there is no bound on the amount that a finite initial segment of this sequence might be compressed.
- One might write a very short program that takes advantage of the regularity.
- One can eventually shorten the algorithmic representation of finite initial segments of the sequence by more than any constant c.

Randomness Constraints

- Problem: there is no bound on the amount that a finite initial segment of this sequence might be compressed.
- One might write a very short program that takes advantage of the regularity.
- One can eventually shorten the algorithmic representation of finite initial segments of the sequence by more than any constant c.
- This observation provides the key idea behind Kolmogorov-Chaitin randomness.

Randomness Constraints

Kolmogorov-Chaitin Randomness

An ω-sequence is Kolmogorov-Chaitin random if and only if there is a constant c such that all finite initial segments are c-incompressible (by a prefix-free Turing machine).

Randomness Constraints

Kolmogorov-Chaitin Randomness

An ω-sequence is Kolmogorov-Chaitin random if and only if there is a constant c such that all finite initial segments are c-incompressible (by a prefix-free Turing machine).

- An initial segment is c-incompressible if and only if it is not representable by an algorithm that is c shorter than the initial segment.
- A prefix-free Turing machine is a universal Turing machine that is self-delimiting and hence can read its input in one direction without knowing what, if anything, comes next.

Randomness Constraints

Let's turn to the second condition regarding "generic/typical."

Randomness Constraints

Let's turn to the second condition regarding "generic/typical."

First-pass: an ω-sequence x is random only if x satisfies all typicality properties.

Randomness Constraints

Let's turn to the second condition regarding "generic/typical."

First-pass: an ω-sequence x is random only if x satisfies all typicality properties.

Too strong!

- Being different from x is a typicality property.

Randomness Constraints

Let's turn to the second condition regarding "generic/typical."

First-pass: an ω-sequence x is random only if x satisfies all typicality properties.

Too strong!

- Being different from x is a typicality property.
- The first-pass criterion implies that no ω-sequence is random!

Randomness Constraints

Let's turn to the second condition regarding "generic/typical."

First-pass: an ω-sequence x is random only if x satisfies all typicality properties.

Too strong!

- Being different from x is a typicality property.
- The first-pass criterion implies that no ω-sequence is random!
- We need to restrict the class of typicality properties.

Randomness Constraints

Let's turn to the second condition regarding "generic/typical."

First-pass: an ω-sequence x is random only if x satisfies all typicality properties.

Too strong!

- Being different from x is a typicality property.
- The first-pass criterion implies that no ω-sequence is random!
- We need to restrict the class of typicality properties.

Martin-Löf (1966) offers a remarkably simple and satisfactory solution.

Martin-Löf Randomness

An ω-sequence is Martin-Löf random if and only if it belongs to every effective full-measure set, i.e. it belongs to no effective measure-zero set.

Randomness Constraints

Martin-Löf Randomness

An ω-sequence is Martin-Löf random if and only if it belongs to every effective full-measure set, i.e. it belongs to no effective measure-zero set.

- Martin-Löf: a set $E \subset 2^\omega$ is effective measure-zero iff there is a uniformly effective sequence of open sets, $G_1, G_2, \ldots$ such that, for all n, $E \subset G_n$ and $\mu(G_n) < 1/n$.
- Uniformly effective open: the entire sequence of open sets is determined by a single program.

cf: A. Dasgupta 2011.

Surprisingly, the two seemingly different definitions are equivalent.

Surprisingly, the two seemingly different definitions are equivalent.

Schnorr's Theorem

An ω-sequence is Martin-Löf random if and only if it is Kolmogorov-Chaitin random.

Randomness Constraints

Surprisingly, the two seemingly different definitions are equivalent.

Schnorr's Theorem

An ω-sequence is Martin-Löf random if and only if it is Kolmogorov-Chaitin random.

Moreover, they are also equivalent to the betting-criterion, i.e. no computably enumerable martingale succeeds on the ω-sequence.

Randomness Constraints

Surprisingly, the two seemingly different definitions are equivalent.

Schnorr's Theorem

An ω-sequence is Martin-Löf random if and only if it is Kolmogorov-Chaitin random.

Moreover, they are also equivalent to the betting-criterion, i.e. no computably enumerable martingale succeeds on the ω-sequence.

Upshot: Martin-Löf randomness is simple and natural, and very surprisingly it is equivalent to seemingly different characterizations of randomness.

Randomness Constraints

Surprisingly, the two seemingly different definitions are equivalent.

Schnorr's Theorem

An ω-sequence is Martin-Löf random if and only if it is Kolmogorov-Chaitin random.

Moreover, they are also equivalent to the betting-criterion, i.e. no computably enumerable martingale succeeds on the ω-sequence.

Upshot: Martin-Löf randomness is simple and natural, and very surprisingly it is equivalent to seemingly different characterizations of randomness.

cf: Church-Turing thesis; Martin-Löf-Chaitin thesis.

Randomness Constraints

We use Martin-Löf randomness to specify the law $L^{\star}$ as a constraint on the set of physically possible worlds:

Randomness Constraints

We use Martin-Löf randomness to specify the law $L^\star$ as a constraint on the set of physically possible worlds:

Probabilistic Law $L^\star_{ML}$

The ω-sequence of coin tosses $\langle r_1, r_2, \ldots \rangle$ is Martin-Löf random with unbiased relative frequencies of heads and tails.

Randomness Constraints

We use Martin-Löf randomness to specify the law $L^\star$ as a constraint on the set of physically possible worlds:

Probabilistic Law $L^\star_{ML}$

The ω-sequence of coin tosses $\langle r_1, r_2, \ldots \rangle$ is Martin-Löf random with unbiased relative frequencies of heads and tails.

Here all of the worlds in $\Omega^{L^\star_{ML}}$ are random with well-defined unbiased relative frequencies.

Hence,

Hence,

- a non-dogmatic inquirer will surely infer unbiased relative frequencies in the limit;

Randomness Constraints

Hence,

- a non-dogmatic inquirer will surely infer unbiased relative frequencies in the limit;
- inasmuch as all initial segments of her data will be c-incompressible, she will have as good of evidence as possible that the data are patternless and hence randomly distributed.

Randomness Constraints

- Martin-Löf randomness is not the only way that one might characterize a probabilistic coin-toss law.

Randomness Constraints

- Martin-Löf randomness is not the only way that one might characterize a probabilistic coin-toss law.
- There are other algorithmic notions of randomness to choose from.

Randomness Constraints

- Martin-Löf randomness is not the only way that one might characterize a probabilistic coin-toss law.
- There are other algorithmic notions of randomness to choose from.
- Schnorr randomness is a closely-related notion with many of the same virtues.

Randomness Constraints

A Schnorr test is a special kind of Martin-Löf test, requiring a more stringent definition of "effective measure-zero."

A Schnorr test is a special kind of Martin-Löf test, requiring a more stringent definition of "effective measure-zero."

- A Schnorr test is a Martin-Löf test where the measures $\mu(U_n)$ are themselves uniformly computable. A class $C \subset 2^\omega$ is Schnorr null if there is a Schnorr test $\{U_n\}_{n\in\omega}$ such that $C \subseteq \bigcap_n U_n$. And a sequence $S \in 2^\omega$ is Schnorr random if and only if $\{S\}$ is not Schnorr null.

Randomness Constraints

As with Martin-Löf randomness, we might use the notion of Schnorr randomness to specify a probabilistic constraining law:

As with Martin-Löf randomness, we might use the notion of Schnorr randomness to specify a probabilistic constraining law:

Probabilistic Law $L_S^\star$

The ω-sequence of coin tosses $\langle r_1, r_2, \ldots \rangle$ is Schnorr random with unbiased relative frequencies of heads and tails.

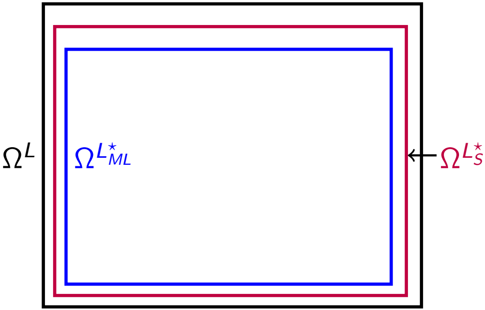

Since there are sequences that are Schnorr random but not Martin-Löf random, $L^\star_{ML}$ and $L^\star_{S}$ are different laws.

Randomness Constraints

- $L^{\star}_{ML}$ and $L^{\star}_{S}$ are different laws.

Randomness Constraints

- $L^{\star}_{ML}$ and $L^{\star}_{S}$ are different laws.
- But they are in a strong sense empirically equivalent.

Randomness Constraints

- $L^{\star}_{ML}$ and $L^{\star}_{S}$ are different laws.
- But they are in a strong sense empirically equivalent.
- No effective procedure would determine whether a particular sequence is Martin-Löf random or Schnorr random but not Martin-Löf random. (Barrett and Huttegger 2021)

Randomness Constraints

- $L^{\star}_{ML}$ and $L^{\star}_{S}$ are different laws.
- But they are in a strong sense empirically equivalent.
- No effective procedure would determine whether a particular sequence is Martin-Löf random or Schnorr random but not Martin-Löf random. (Barrett and Huttegger 2021)
- If one is limited to Turing-strength computation, one would never be able to distinguish between $L^{\star}_{ML}$ and $L^{\star}_{S}$ no matter what empirical evidence one had.

Randomness Constraints

Moving from a standard probabilistic law to a $L^{\star}$ law eliminates one variety of empirical underdetermination, but it reveals two others.

Randomness Constraints

Moving from a standard probabilistic law to a $L^\star$ law eliminates one variety of empirical underdetermination, but it reveals two others.

1. One should already expect law $L^\star$ to be empirically indistinguishable from L, insofar as one expects a sequence of coin tosses governed by a traditional probabilistic law L to be such that one can detect no discernible pattern.

Randomness Constraints

Moving from a standard probabilistic law to a $L^\star$ law eliminates one variety of empirical underdetermination, but it reveals two others.

1. One should already expect law $L^\star$ to be empirically indistinguishable from L, insofar as one expects a sequence of coin tosses governed by a traditional probabilistic law L to be such that one can detect no discernible pattern.
2. One will be unable to distinguish between different versions of $L^\star$ like $L^\star_{ML}$ and $L^\star_S$, insofar as one is limited to Turing-strength computations.

Randomness Constraints

Moving from a standard probabilistic law to a $L^\star$ law eliminates one variety of empirical underdetermination, but it reveals two others.

1. One should already expect law $L^\star$ to be empirically indistinguishable from L, insofar as one expects a sequence of coin tosses governed by a traditional probabilistic law L to be such that one can detect no discernible pattern.
2. One will be unable to distinguish between different versions of $L^\star$ like $L^\star_{ML}$ and $L^\star_S$, insofar as one is limited to Turing-strength computations.

Since Martin-Löf randomness has the sort of properties we want and as it is arguably the standard algorithmic notion (Dasgupta 2011), we shall understand $L^\star$ as $L^\star_{ML}$.

Recap

We've defined what it means to be a $L^{\star}$ law.

Recap

We've defined what it means to be a $L^\star$ law.

- Tighter fit between the law and its corresponding set of possible worlds.
- Quite different from a L law.

The Nature of Laws

What kind of physical law is $L^{\star}$?

The Nature of Laws

What kind of physical law is $L^{\star}$?

And how does it govern the world?

One might think of $L^\star$ in one of these two ways:

The Nature of Laws

One might think of $L^{\star}$ in one of these two ways:

- a generative chance* law;

One might think of $L^\star$ in one of these two ways:

- a generative chance* law;
- a probabilistic* constraining law.

One might think of $L^\star$ in one of these two ways:

- a generative chance* law;
- a probabilistic* constraining law.

These are meant to be non-Humean laws. Later, we will discuss the implications for Humeanism.

The Nature of Laws

As a generative chance* law:

As a generative chance* law:

- $L^\star$ tells us that each toss is generated by unbiased chances*,

As a generative chance* law:

- $L^\star$ tells us that each toss is generated by unbiased chances*,
- where a chance* process behaves just like an ordinary chance process...

As a generative chance* law:

- $L^\star$ tells us that each toss is generated by unbiased chances*,
- where a chance* process behaves just like an ordinary chance process…
- …except that it can never produce an infinite sequence that fails to be Martin-Löf random or fails to exhibit well-defined relative frequencies.

- As a result, a chance* process involves a subtle violation of independence.

- As a result, a chance* process involves a subtle violation of independence.
- The sequence of tosses will pass every finitely specifiable test for statistical independence.

The Nature of Laws

- As a result, a chance* process involves a subtle violation of independence.
- The sequence of tosses will pass every finitely specifiable test for statistical independence.
- But since the full sequence must satisfy the constraint imposed by the law, a chance* process is <u>holistically</u> constrained.

- The constraint is not felt on any finite set of tosses, nor is it discoverable by effective means.

The Nature of Laws

- The constraint is not felt on any finite set of tosses, nor is it discoverable by effective means.
- But it does require that a relationship hold between the full sequence of tosses that is generated by the process in the limit.

- The constraint is not felt on any finite set of tosses, nor is it discoverable by effective means.
- But it does require that a relationship hold between the full sequence of tosses that is generated by the process in the limit.
- This interdependence between outcomes may be incompatible with the usual intuitions behind wanting a generative law.

The Nature of Laws

- The constraint is not felt on any finite set of tosses, nor is it discoverable by effective means.
- But it does require that a relationship hold between the full sequence of tosses that is generated by the process in the limit.
- This interdependence between outcomes may be incompatible with the usual intuitions behind wanting a generative law.
- It may also be incompatible with how causal explanation works more generally.

The Nature of Laws

A more natural interpretation:

A more natural interpretation:

- $L^\star$ can be regarded as a law that governs by constraining the entire history of the world—in this case, the full ω-sequence of outcomes.

The Nature of Laws

A more natural interpretation:

- $L^\star$ can be regarded as a law that governs by constraining the entire history of the world—in this case, the full ω-sequence of outcomes.
- $L^\star$ tells us which sequences of outcomes are physically possible, namely those that satisfy the frequency constraint and the randomness constraint imposed by the law.

A more natural interpretation:

- $L^\star$ can be regarded as a law that governs by constraining the entire history of the world—in this case, the full ω-sequence of outcomes.
- $L^\star$ tells us which sequences of outcomes are physically possible, namely those that satisfy the frequency constraint and the randomness constraint imposed by the law.
- It meshes well with Chen and Goldstein's (2022) minimal primitivism account (MinP),

A more natural interpretation:

- $L^\star$ can be regarded as a law that governs by constraining the entire history of the world—in this case, the full ω-sequence of outcomes.
- $L^\star$ tells us which sequences of outcomes are physically possible, namely those that satisfy the frequency constraint and the randomness constraint imposed by the law.
- It meshes well with Chen and Goldstein's (2022) minimal primitivism account (MinP), according to which laws are certain primitive facts that govern the world by constraining the physical possibilities of the entire spacetime and its contents.

The Nature of Laws

Understood this way, $L^{\star}$ addresses problems encountered by both non-Humean and Humean accounts of laws.

Understood this way, $L^{\star}$ addresses problems encountered by both non-Humean and Humean accounts of laws.

We will start with the former.

- On non-Humean governing accounts of laws, there is a puzzle:

The Nature of Laws

- On non-Humean governing accounts of laws, there is a puzzle:
- **Precisely how do probabilistic laws govern?**

The Nature of Laws

- On non-Humean governing accounts of laws, there is a puzzle:
- **Precisely how do probabilistic laws govern?**
- According to the standard view, probabilistic laws do not rule out any world.

- On non-Humean governing accounts of laws, there is a puzzle:
- **Precisely how do probabilistic laws govern?**
- According to the standard view, probabilistic laws do not rule out any world.
- Instead, they merely assign some numbers between zero and one to (measurable) subsets in the space of all ω-sequences.

The Nature of Laws

- On non-Humean governing accounts of laws, there is a puzzle:
- **Precisely how do probabilistic laws govern?**
- According to the standard view, probabilistic laws do not rule out any world.
- Instead, they merely assign some numbers between zero and one to (measurable) subsets in the space of all ω-sequences.
- This raises a puzzle: what do these numbers between zero and one represent in physical reality?

- Some non-Humeans appeal to gradable notions...

The Nature of Laws

- Some non-Humeans appeal to gradable notions...
- such as "propensities" (Maudlin 2007, p.20) or "probabilities of necessitation" (Armstrong 1983, p.172).

The Nature of Laws

- Some non-Humeans appeal to gradable notions...
- such as "propensities" (Maudlin 2007, p.20) or "probabilities of necessitation" (Armstrong 1983, p.172).
- Suppose a probabilistic law assigns a 0.2 probability to the next outcome being heads.

The Nature of Laws

- Some non-Humeans appeal to gradable notions...
- such as "propensities" (Maudlin 2007, p.20) or "probabilities of necessitation" (Armstrong 1983, p.172).
- Suppose a probabilistic law assigns a 0.2 probability to the next outcome being heads.
- "The chance setup has a 0.2 propensity to bring about a heads-outcome in the next toss."

The Nature of Laws

- Some non-Humeans appeal to gradable notions...
- such as "propensities" (Maudlin 2007, p.20) or "probabilities of necessitation" (Armstrong 1983, p.172).
- Suppose a probabilistic law assigns a 0.2 probability to the next outcome being heads.
- "The chance setup has a 0.2 propensity to bring about a heads-outcome in the next toss."
- "The current state of affairs necessitates the state of affairs of a heads-outcome to 0.2 probability."

The Nature of Laws

- Some non-Humeans appeal to gradable notions...
- such as "propensities" (Maudlin 2007, p.20) or "probabilities of necessitation" (Armstrong 1983, p.172).
- Suppose a probabilistic law assigns a 0.2 probability to the next outcome being heads.
- "The chance setup has a 0.2 propensity to bring about a heads-outcome in the next toss."
- "The current state of affairs necessitates the state of affairs of a heads-outcome to 0.2 probability."

One might make sense of non-gradable notions of physical possibility and impossibility.

The Nature of Laws

- Some non-Humeans appeal to gradable notions...
- such as "propensities" (Maudlin 2007, p.20) or "probabilities of necessitation" (Armstrong 1983, p.172).
- Suppose a probabilistic law assigns a 0.2 probability to the next outcome being heads.
- "The chance setup has a 0.2 propensity to bring about a heads-outcome in the next toss."
- "The current state of affairs necessitates the state of affairs of a heads-outcome to 0.2 probability."

One might make sense of non-gradable notions of physical possibility and impossibility. But gradable notions such as propensities and degrees of necessitation are much less clear.

The Nature of Laws

- In contrast, probabilistic laws, such as $L^{\star}$, can be viewed as a special class of constraining laws.

The Nature of Laws

- In contrast, probabilistic laws, such as $L^\star$, can be viewed as a special class of constraining laws.
- They constrain what is physically possible by ruling out certain sequences of outcomes.

- In contrast, probabilistic laws, such as $L^{\star}$, can be viewed as a special class of constraining laws.
- They constrain what is physically possible by ruling out certain sequences of outcomes.
- Which worlds are ruled out?

The Nature of Laws

- In contrast, probabilistic laws, such as $L^\star$, can be viewed as a special class of constraining laws.
- They constrain what is physically possible by ruling out certain sequences of outcomes.
- Which worlds are ruled out?
- Those that fail either the frequency constraint or the randomness constraint.

The Nature of Laws

- In contrast, probabilistic laws, such as $L^{\star}$, can be viewed as a special class of constraining laws.
- They constrain what is physically possible by ruling out certain sequences of outcomes.
- Which worlds are ruled out?
- Those that fail either the frequency constraint or the randomness constraint.
- We can do away with gradable notions such as propensities or probabilities of necessitation altogether.

The Nature of Laws

- In contrast, probabilistic laws, such as $L^\star$, can be viewed as a special class of constraining laws.
- They constrain what is physically possible by ruling out certain sequences of outcomes.
- Which worlds are ruled out?
- Those that fail either the frequency constraint or the randomness constraint.
- We can do away with gradable notions such as propensities or probabilities of necessitation altogether.
- In their place, we require only non-gradable notions of physical possibilities and impossibilities.

The Nature of Laws

Consider MinP as a non-Humean example.

Consider MinP as a non-Humean example.

- We can now employ a single primitive relation, namely <u>constraining</u>.

The Nature of Laws

Consider MinP as a non-Humean example.

- We can now employ a single primitive relation, namely <u>constraining</u>.
- Both probabilistic and non-probabilistic types of laws govern by constraining what is physically possible.

The Nature of Laws

Consider MinP as a non-Humean example.

- We can now employ a single primitive relation, namely constraining.
- Both probabilistic and non-probabilistic types of laws govern by constraining what is physically possible.
- The way that $L^\star$ constrains the world is not so different from that of $F = ma$.

The Nature of Laws

Consider MinP as a non-Humean example.

- We can now employ a single primitive relation, namely constraining.
- Both probabilistic and non-probabilistic types of laws govern by constraining what is physically possible.
- The way that $L^\star$ constrains the world is not so different from that of $F = ma$.
- $L^\star$ constrains the physical possibilities to be all and only the non-maverick worlds.

The Nature of Laws

Consider MinP as a non-Humean example.

- We can now employ a single primitive relation, namely constraining.
- Both probabilistic and non-probabilistic types of laws govern by constraining what is physically possible.
- The way that $L^\star$ constrains the world is not so different from that of $F = ma$.
- $L^\star$ constrains the physical possibilities to be all and only the non-maverick worlds.
- $F = ma$ constrains the physical possibilities to be all and only the solutions of $F = ma$.

Consider MinP as a non-Humean example.

- We can now employ a single primitive relation, namely constraining.
- Both probabilistic and non-probabilistic types of laws govern by constraining what is physically possible.
- The way that $L^\star$ constrains the world is not so different from that of $F = ma$.
- $L^\star$ constrains the physical possibilities to be all and only the non-maverick worlds.
- $F = ma$ constrains the physical possibilities to be all and only the solutions of $F = ma$.

In this way, $L^\star$ removes a major obstacle to a unified understanding of probabilistic and non-probabilistic laws.

Let's turn to Humeanism.

Let's turn to Humeanism.

Humeans may also find it useful to adopt $L^{\star}$, for two reasons.

(1) It is relevant to the issue of the Big Bad Bug (Lewis 1986, pp.xiv-xvi).

The Nature of Laws

(1) It is relevant to the issue of the Big Bad Bug (Lewis 1986, pp.xiv-xvi).

- Lewis: the Principal Principle and Humean supervenience lead to a contradiction.

The Nature of Laws

(1) It is relevant to the issue of the Big Bad Bug (Lewis 1986, pp.xiv-xvi).

- Lewis: the Principal Principle and Humean supervenience lead to a contradiction.
- Certain histories of the Humean mosaic, called <u>undermining histories</u>, are assigned, according to the Principal Principle, non-zero probability, conditionalized on some probabilistic theory T being the best system.

(1) It is relevant to the issue of the Big Bad Bug (Lewis 1986, pp.xiv-xvi).

- Lewis: the Principal Principle and Humean supervenience lead to a contradiction.
- Certain histories of the Humean mosaic, called undermining histories, are assigned, according to the Principal Principle, non-zero probability, conditionalized on some probabilistic theory T being the best system.
- However, they are assigned, according to Humean supervenience, zero probability, because T would not be the best system had any of its undermining histories been actual.

What sort of history would count as undermining?

The Nature of Laws

What sort of history would count as undermining?

- wrong limiting frequencies

The Nature of Laws

What sort of history would count as undermining?

- wrong limiting frequencies
- no limiting frequencies

The Nature of Laws

What sort of history would count as undermining?

- wrong limiting frequencies
- no limiting frequencies
- containing patterns that can be summarized by a simpler system, such as a deterministic law in the case of the alternating heads-tails sequence.

What sort of history would count as undermining?

- wrong limiting frequencies
- no limiting frequencies
- containing patterns that can be summarized by a simpler system, such as a deterministic law in the case of the alternating heads-tails sequence.

Undermining histories, then, are exactly the histories of maverick worlds, as they lack the frequency or randomness patterns exemplified by typical sequences of the standard probabilistic law.

- $\star$-laws rules out maverick worlds.

- $\star$-laws rules out maverick worlds.
- If maverick worlds are physically impossible, then there are no physically possible undermining histories that can be used to derive the contradiction.

- $\star$-laws rules out maverick worlds.
- If maverick worlds are physically impossible, then there are no physically possible undermining histories that can be used to derive the contradiction.
- The Big Bad Bug is eliminated!

The Nature of Laws

- $\star$-laws rules out maverick worlds.
- If maverick worlds are physically impossible, then there are no physically possible undermining histories that can be used to derive the contradiction.
- The Big Bad Bug is eliminated!
- Inasmuch as restricting to $\star$-laws is also motivated by considerations of underdetermination and empirical coherence, a Humean may find this solution particularly natural.

(2) Adopting $\star$-laws can avoid appealing to fit as a criterion in the best-system analysis of probabilistic laws, which allows Humeans to bypass difficulties with how to characterize this notion. (cf. Elga 2004)

The Nature of Laws

- Given an ω-sequence, there is much underdetermination among standard probabilistic laws such as L,

The Nature of Laws

- Given an ω-sequence, there is much underdetermination among standard probabilistic laws such as L,
- one needs something like fit to choose the winning best system.

The Nature of Laws

- Given an ω-sequence, there is much underdetermination among standard probabilistic laws such as L,
- one needs something like fit to choose the winning best system.
- Informativeness as the quantity of worlds being excluded does not distinguish among probabilistic laws like L.

The Nature of Laws

- Given an ω-sequence, there is much underdetermination among standard probabilistic laws such as L,
- one needs something like fit to choose the winning best system.
- Informativeness as the quantity of worlds being excluded does not distinguish among probabilistic laws like L.
- Different candidates compatible with the same set of worlds.

In contrast, there is much less underdetermination among probabilistic laws like $L^\star$.

In contrast, there is much less underdetermination among probabilistic laws like $L^\star$.

- If we consider a spectrum of different probabilistic statements like $L^\star$ that differ, say, in their specifications of the relative frequencies in the ω-sequence,

In contrast, there is much less underdetermination among probabilistic laws like $L^\star$.

- If we consider a spectrum of different probabilistic statements like $L^\star$ that differ, say, in their specifications of the relative frequencies in the ω-sequence,
- then at most one of them is compatible with the ω-sequence,

In contrast, there is much less underdetermination among probabilistic laws like $L^\star$.

- If we consider a spectrum of different probabilistic statements like $L^\star$ that differ, say, in their specifications of the relative frequencies in the ω-sequence,
- then at most one of them is compatible with the ω-sequence,
- thus at most one of them is an axiom in the best system of that ω-sequence.

In contrast, there is much less underdetermination among probabilistic laws like $L^\star$.

- If we consider a spectrum of different probabilistic statements like $L^\star$ that differ, say, in their specifications of the relative frequencies in the ω-sequence,
- then at most one of them is compatible with the ω-sequence,
- thus at most one of them is an axiom in the best system of that ω-sequence.
- The best system analysis of $L^\star$ is much like that of $F = G\frac{m_1 m_2}{r^2}$.

- Given any Humean mosaic, one needs criteria such as simplicity and informativeness, but one does not need the statistical criterion of fit, to determine the best system.

The Nature of Laws

- Given any Humean mosaic, one needs criteria such as simplicity and informativeness, but one does not need the statistical criterion of fit, to determine the best system.
- The usual problems associated with fit would not arise for such Humeans!

Discussion

A $L^\star$ law may be thought of as either a generative chance* law or a probabilistic* constraining law.

Discussion

A $L^\star$ law may be thought of as either a generative chance* law or a probabilistic* constraining law.

The notions of chance* and probability* are subtly different from traditional chance or probability.

Discussion: Independence

- The results of coin tosses on $L^{\star}$ satisfy every computable test for independence and will hence appear to be statistically independent.

Discussion: Independence

- The results of coin tosses on $L^\star$ satisfy every computable test for independence and will hence appear to be statistically independent.
- One might say that the results are probabilistically* independent.

Discussion: Independence

- The results of coin tosses on $L^\star$ satisfy every computable test for independence and will hence appear to be statistically independent.
- One might say that the results are probabilistically* independent.
- But inasmuch as some sequences are impossible, there is also a sense in which the results of tosses in this full ω-sequence are interdependent.

Discussion: Independence

- To understand $L^{\star}$ as a generative chance* law, one would need to allow for a holistic causal structure that guarantees random sequences with unbiased relative frequencies in the limit.

Discussion: Independence

- To understand $L^{\star}$ as a generative chance* law, one would need to allow for a holistic causal structure that guarantees random sequences with unbiased relative frequencies in the limit.
- Depending on one's commitments regarding causal explanation, this may lead one to favor understanding $L^{\star}$ as a probabilistic* constraining law.

Discussion: Independence

- To understand $L^\star$ as a generative chance* law, one would need to allow for a holistic causal structure that guarantees random sequences with unbiased relative frequencies in the limit.
- Depending on one's commitments regarding causal explanation, this may lead one to favor understanding $L^\star$ as a probabilistic* constraining law.
- If one gives up on a generative chance* law, one is left with a useful option for both proponents of governing-law accounts and Humeans.

Discussion: Empirical Underdetermination

- chance* and probability* depend on a choice of a particular standard of algorithmic randomness.

Discussion: Empirical Underdetermination

- chance* and probability* depend on a choice of a particular standard of algorithmic randomness.
- Choice between Martin-Löf and Schnorr randomness.

Discussion: Empirical Underdetermination

- chance* and probability* depend on a choice of a particular standard of algorithmic randomness.
- Choice between Martin-Löf and Schnorr randomness.
- A computational sort of empirical underdetermination.

Still, $\star$-laws also help to eliminate some forms of empirical underdetermination.

Discussion: Empirical Underdetermination

Still, $\star$-laws also help to eliminate some forms of empirical underdetermination.

- Unlike traditional probabilistic laws but very much like deterministic laws like $F = ma$ and $F = G\frac{m_1 m_2}{r^2}$, one will surely learn $L^\star$ on complete evidence in <u>every</u> physically possible world.

Discussion: Empirical Underdetermination

- In contrast, if L is the true law, there will be a continuous cardinality of maverick worlds such that, if one were to inhabit any of them, one could never learn L from the results of the coin tosses.

Discussion: Empirical Underdetermination

- In contrast, if L is the true law, there will be a continuous cardinality of maverick worlds such that, if one were to inhabit any of them, one could never learn L from the results of the coin tosses.
- On the usual approach to thinking about laws, one needs special background assumptions to overcome this difficulty.

Discussion: Empirical Underdetermination

- In contrast, if L is the true law, there will be a continuous cardinality of maverick worlds such that, if one were to inhabit any of them, one could never learn L from the results of the coin tosses.
- On the usual approach to thinking about laws, one needs special background assumptions to overcome this difficulty.
- One needs to argue that inhabiting a maverick world of the true law is sufficiently unlikely or atypical that one has rational justification for simply ignoring the possibility.

Discussion: Empirical Underdetermination

- In contrast, if L is the true law, there will be a continuous cardinality of maverick worlds such that, if one were to inhabit any of them, one could never learn L from the results of the coin tosses.
- On the usual approach to thinking about laws, one needs special background assumptions to overcome this difficulty.
- One needs to argue that inhabiting a maverick world of the true law is sufficiently unlikely or atypical that one has rational justification for simply ignoring the possibility.
- One needs to know that the world one inhabits, and hence the statistical nature of the sequence of records that one has in fact recorded, is typical or very likely given the true law.

Discussion: Empirical Underdetermination

- If an agent is simply presented with a sequence of coin-toss records and knows nothing more than that she inhabits a physically possible world...

Discussion: Empirical Underdetermination

- If an agent is simply presented with a sequence of coin-toss records and knows nothing more than that she inhabits a physically possible world...
- ...she cannot infer that the sequence exhibits typical statistical properties given the true law.

Discussion: Empirical Underdetermination

- If an agent is simply presented with a sequence of coin-toss records and knows nothing more than that she inhabits a physically possible world...
- ...she cannot infer that the sequence exhibits typical statistical properties given the true law.
- Such an inference requires her to appeal to a background assumption like the Principal Principle or Cournot's Principle.

Discussion: Empirical Underdetermination

- If an agent is simply presented with a sequence of coin-toss records and knows nothing more than that she inhabits a physically possible world...
- ...she cannot infer that the sequence exhibits typical statistical properties given the true law.
- Such an inference requires her to appeal to a background assumption like the Principal Principle or Cournot's Principle.
- While such assumptions may be warranted given one's other commitments, they are not required for empirical coherence if one restricts one's hypotheses to $\star$-laws.

Conclusion

- We have used algorithmic randomness to characterize two types of probabilistic laws:

Conclusion

- We have used algorithmic randomness to characterize two types of probabilistic laws:
    - a generative chance* law;

Conclusion

- We have used algorithmic randomness to characterize two types of probabilistic laws:
  - a <u>generative chance* law</u>;
  - a <u>probabilistic* constraining law</u>.

Conclusion

- We have used algorithmic randomness to characterize two types of probabilistic laws:
  - a generative chance* law;
  - a probabilistic* constraining law.
- We have argued that $\star$-laws provide a novel way of understanding probabilities and chances, and help to address one variety of empirical underdetermination, but they also reveal other varieties that have been underappreciated.

Conclusion

- We have used algorithmic randomness to characterize two types of probabilistic laws:
  - a generative chance* law;
  - a probabilistic* constraining law.
- We have argued that $\star$-laws provide a novel way of understanding probabilities and chances, and help to address one variety of empirical underdetermination, but they also reveal other varieties that have been underappreciated.
- For all we know, our world might be characterized by a traditional probabilistic law or a $\star$-law.

Conclusion

- The notion of a probabilistic* constraining law has advantages over that of a generative chance* law.

Conclusion

- The notion of a probabilistic* constraining law has advantages over that of a generative chance* law.
- It meshes well with the holistic character of the randomness and relative frequency constraints,

Conclusion

- The notion of a probabilistic* constraining law has advantages over that of a generative chance* law.
- It meshes well with the holistic character of the randomness and relative frequency constraints, directly supports a unified governing account of non-Humean laws,

Conclusion

- The notion of a probabilistic* constraining law has advantages over that of a generative chance* law.
- It meshes well with the holistic character of the randomness and relative frequency constraints, directly supports a unified governing account of non-Humean laws, and provides independently motivated solutions to issues in the Humean best-system account.

Conclusion

- The notion of a probabilistic* constraining law has advantages over that of a generative chance* law.
- It meshes well with the holistic character of the randomness and relative frequency constraints, directly supports a unified governing account of non-Humean laws, and provides independently motivated solutions to issues in the Humean best-system account.
- Both notions are worthy of study and may lead to new ideas concerning the nature of laws.

Thank you for your attention!