

From Gödel's Incompleteness Theorem to the completeness of bot beliefs

(Extended abstract)

Dusko Pavlovic^{*1} and Temra Pavlovic^{**}

¹University of Hawaii, Honolulu HI

Abstract

Hilbert and Ackermann asked for a method to consistently extend incomplete theories to complete theories. Gödel essentially proved that any theory capable of encoding its own statements and their proofs contains statements that are true but not provable. Hilbert did not accept that Gödel's construction answered his question, and in his late writings and lectures, Gödel agreed that it did not, since theories can be completed incrementally, by adding axioms to prove ever more true statements, as science normally does, with completeness as the vanishing point. This pragmatic view of validity is familiar not only to scientists who conjecture test hypotheses but also to real-estate agents and other dealers, who conjure claims, albeit invalid, as necessary to close a deal, confident that they will be able to conjure other claims, albeit invalid, sufficient to make the first claims valid. We study the underlying logical process and describe the trajectories leading to testable but unfalsifiable theories to which bots and other automated learners are likely to converge.

1 Introduction

Logic as the theory of theories was originally developed to prove true statements. Here we study developments in the opposite direction: modifying interpretations to make true some previously false statements. In modal logic, such logical processes have been modeled as instances of *belief update* [2, 3, 10]. In the practice of science, such processes arise when theories are updated to explain new observations [22, Ch. 4]. In public life, the goal of such processes is to influence some public perceptions to better suit some private preferences [11,

^{*}Partially supported by NSF and AFOSR.

Part V]. This range of applications gave rise to a gamut of techniques for influence and belief engineering, covering the space from unsupervised learning to conditioning.

From incomplete theories to complete beliefs. The idea to incrementally complete incomplete theories [9] arose soon after Gödel proved his Incompleteness Theorem [14]. Alan Turing wrote a thesis about ordinal towers of completions and discovered the hierarchy of unsolvability degrees [35]. The core idea was to keep recognizing and adding true but unprovable statements to theories. In the meantime, interests shifted from making true statements provable to making false statements true. Many toy examples of belief updates and revisions have been formalized and studied in dynamic-epistemic logic [4, 7], but the advances in belief engineering and the resulting industry of influence overtook the theory at great speed, and turned several corners of market and political monetizations. The theory remained fragmented even on its own. While modal presentations of Gödel’s theorems appeared early on [33], the computational ideas, that made his self-referential constructions possible [34], never transpired back into modal logic. *The point of the present paper is that combining belief updates with universal languages and self-reference leads to a curious new logical capability, whereby theories and models can be steered to assure consistency and completeness of future updates.* This capability precludes disproving current beliefs and the framework becomes *belief-complete* in a suitable formal sense, discussed below.

The logical framework combining belief updates and universal languages may seem unfamiliar. The main body of this paper is devoted to an attempt to describe how it arises from familiar logical frameworks. Here we try to clarify the underlying ideas.

Universality. Just like Gödel’s incompleteness theorems, our constructions of unfalsifiable beliefs are based on a *universal language* \mathbb{L} . The abstract characterization of universality, which we borrow from [24, Ch. 2], is that \mathbb{L} comes equipped with a family of *interpreters* $\{\cdot\}: \mathbb{L} \times A \rightarrow B$, one for each pair of types¹ A, B , such that every function $f: A \rightarrow B$ has a description² $\lceil f \rceil$ in \mathbb{L} , satisfying³

$$f = \{\lceil f \rceil\} \quad (1)$$

This is spelled out in Sec. 4. The construction in Sec. 5 will imply that every $g: \mathbb{L} \times A \rightarrow B$ has a fixpoint Γ , satisfying

$$g(\Gamma, a) = \{\Gamma\}(a) \quad (2)$$

Any complete programming language can be used as \mathbb{L} . Its interpreters support (1) and its specializers induce (2). A sufficiently expressive software specification framework [28] would also fit the bill, as would a general scientific formalism [22].

¹Each pair carries a different interpreter $\{\cdot\}^{AB}$ but we elide the superscripts.

²There may be many descriptions for each f and $\lceil f \rceil$ refers to an arbitrary one.

³The curly bracket notation allows abbreviating $\lambda a. \{\cdot\}(p, a)$ to $\{p\}$.

Gödel's incompleteness: true but unprovable statement. Gödel used the set of natural numbers \mathbb{N} as \mathbb{L} , with arithmetic making it into a programming language. The concept of a programming language did not yet exist, but it came into existence through Gödel's construction. An arithmetic expression specifying a function f was encoded as a number $\ulcorner f \urcorner$ and decoded by an arithmetic function $\{\cdot\}: \mathbb{N} \times \mathbb{N} \rightarrow \mathbb{N}$ as in (1). A restriction of (2) was proved for arithmetic predicates $p: \mathbb{N} \rightarrow \mathbb{B}$, where $\mathbb{B} = \{0, 1\} \subset \mathbb{N}$, and a fixpoint of a predicate $g: \mathbb{L} \times A \rightarrow \mathbb{B}$ was constructed as a predicate encoding $\ulcorner \gamma \urcorner$ satisfying⁴

$$g(\ulcorner \gamma \urcorner, a) = \{\ulcorner \gamma \urcorner\}(a) = \gamma(a) \quad (3)$$

To complete the incompleteness proof, Gödel constructed a predicate $\P: \mathbb{N} \rightarrow \mathbb{B}$ characterizing provability in formal arithmetic:

$$\P(\ulcorner p \urcorner, a) \iff \vdash p(a) \quad (4)$$

for all arithmetic predicates $p: \mathbb{N} \rightarrow \mathbb{B}$. Although proofs may be arbitrarily large, they are always finite, and if $p(a)$ has a proof, \P will eventually find it. On the other hand, since arithmetic predicates, like all arithmetic functions, satisfy $p = \{\ulcorner p \urcorner\}$, we also have

$$\P(\ulcorner \neg p \urcorner, a) = \{\ulcorner \neg p \urcorner\}(a) = p(a) \quad (5)$$

Setting $g(p, a) = \P(\ulcorner \neg p \urcorner, a)$ in (2) induces a fixpoint γ with

$$\P(\ulcorner \neg \gamma \urcorner, a) \stackrel{(3)}{=} \{\ulcorner \gamma \urcorner\}(a) \stackrel{(5)}{=} \P(\ulcorner \gamma \urcorner, a) \quad (6)$$

But (4) then implies

$$\vdash \neg \gamma(a) \iff \vdash \gamma(a) \quad (7)$$

which means that neither γ nor $\neg \gamma$ can be provable. On the other hand, the disjunction $\gamma \vee \neg \gamma$ is classically true. The statement $\gamma \vee \neg \gamma$ is thus true but not provable, and arithmetic is therefore incomplete.

Belief completeness: universal updating. Remarkably, the same encoding-fixpoint conundrum (1–2), which leads to the incompleteness of static theories, also leads to the completeness of dynamically updated theories. Updating is presented as state dependency. The function f in (1) is now in the form $f: X \times A \rightarrow X \times B$ where X is the state space. It may be more intuitive to think of f as a *process*, since it captures state changes⁵. We conveniently present it as a pair $f = \langle f', f'' \rangle$, where $f': X \times A \rightarrow X$ is the *next state* update, whereas $f'': X \times A \rightarrow B$ is an X -indexed family of functions $f''_x: A \rightarrow B$. The elements

⁴Although this discussion is semi-formal, it may be helpful to bear in mind that the equality $\{\ulcorner \gamma \urcorner\}(a) = \gamma(a)$ is *extensional*: it just says that interpreting the description $\ulcorner \gamma \urcorner$ on a value a always outputs the value $\gamma(a)$. But the process whereby $\{\ulcorner \gamma \urcorner\}(a)$ arrives at this value may be different from a given direct evaluation of $\gamma(a)$.

⁵In automata theory, such functions are called the *Mealy machines*.

of the universal language \mathbb{L} are now construed as *belief states*. Its universality means that every observable state x from any state space X is expressible as a belief. The interpreters $\{\} : \mathbb{L} \times A \rightarrow \mathbb{L} \times B$ are also presented as pairs $\{\} = \langle \{\}'', \{\}'' \rangle$, where $\{\}' : \mathbb{L} \times A \rightarrow \mathbb{L}$ updates the belief states whereas $\{\}'' : \mathbb{L} \times A \rightarrow B$ evaluates beliefs to functions. Just like every state x in X determines a function $f''_x : A \rightarrow B$, every belief ℓ in \mathbb{L} determines a function $\{\ell\}'' : A \rightarrow B$, which makes predictions based on the current belief. Generalizing the fixpoint construction (2), every process $f = \langle f', f'' \rangle : X \times A \rightarrow X \times B$ now induces an assignment $\llbracket f \rrbracket : X \rightarrow \mathbb{L}$ of beliefs to states such that

$$\{\llbracket f \rrbracket(x)\}' = \llbracket f \rrbracket(f'_x) \quad \{\llbracket f \rrbracket(x)\}'' = f''_x \quad (8)$$

The construction of $\llbracket f \rrbracket$ is presented in Sec. 6. Here we propose an interpretation. The second equation says that the output component of $\{\}$ behaves as it did in (1): it interprets the description $\llbracket f \rrbracket(x)$ and recovers the function f''_x executed by the process f at the state x . The first equation says that the interpreter $\{\}$ maps the $\llbracket f \rrbracket$ -description of the state x to the $\llbracket f \rrbracket$ -description of the updated state f'_x :

$$\frac{f' : x \mapsto f'_x}{\{\}' : \llbracket f \rrbracket(x) \mapsto \llbracket f \rrbracket(f'_x)} \quad (9)$$

Any state change caused by the process f is thus explained by a belief update of $\llbracket f \rrbracket$ along $\{\}$. Interpreting the belief states $\llbracket f \rrbracket$ by the interpreter $\{\}$ provides belief updates that can be construed as *explanations* in the language \mathbb{L} of any state changes in the process f . All that can be learned about f is already expressed in $\llbracket f \rrbracket$ and all state changes that may be observed will be explained by the updates anticipated by the current belief, as indicated in (9). The belief is complete.

Remark. In coalgebra and process calculus, the universal interpreters $\{\} : \mathbb{L} \times A \rightarrow \mathbb{L} \times B$ would be characterized as weakly final simulators [30]. They are universal in the sense that the same state space \mathbb{L} works for all types A, B . See [24, Sec. 7.2] for details and references.

The logic of going dynamic. When \mathbb{L} is a programming language, the interpreter $\{\}$ interprets programs as computable functions $A \rightarrow B$, where A and B are types, usually predicates that allow type checking. When \mathbb{L} is a language of software specifications or scientific theories construed as beliefs about the state of the world, the interpreter $\{\}$ updates beliefs to explain the state changes observed in explainable processes $X \times A \rightarrow X \times B$, where A, B and X are state spaces. States are usually also defined by some predicates, but their purpose is not to be easy to check but to define the state changes as semantical reassignments. This is spelled out in Sec. 2.1. Dynamic reassignments of meaning bring us into the realm of *dynamic logic*. If the propositions from a lattice \mathcal{T} are used as assertions about the states of the world or the states of our beliefs about the world, then the dynamic changes of these assertions under the influence of events from a lattice \mathcal{E} can be expressed in terms of *Hoare triples*

$$A\{e\}B \quad (10)$$

saying that the event $e \in \mathcal{E}$ after the precondition $A \in \mathcal{T}$ leads to the postcondition $B \in \mathcal{T}$. The *Hoare logic* of such statements was developed in the late 1960s as a method for reasoning about programs. The algebra of events \mathcal{E} was generated by program expressions, whereas the propositional lattice \mathcal{T} was generated by formal versions of the comments inserted by programmers into their code, to clarify the intended meanings of blocks of code [13, 17]. A triple (10) would thus correspond to a block of code e , a comment A describing the assumed state before e is executed, as its precondition, and a comment B describing the guaranteed state after e is executed, as its postcondition. By formalizing the “*assume-guarantee*” reasoning of software developers, the Hoare triples provided a stepping stone into the logic of state transitions in general. The propositional algebra of dynamic logic can be viewed as a monotone map

$$\mathcal{T}^o \times \mathcal{E} \times \mathcal{T} \xrightarrow{-\{-\}-} \mathcal{O}$$

where \mathcal{O} is a lattice of truth values, whereas \mathcal{T} and \mathcal{E} are as above, and \mathcal{T}^o is \mathcal{T} with the opposite order. If the lattice \mathcal{T} is complete, then each event $e \in \mathcal{E}$ induces a Galois connection

$$A \times e \vdash B \iff A\{e\}B \iff A \vdash [e]B$$

determining a *dynamic modality* $[e]: \mathcal{T} \rightarrow \mathcal{T}$ for every $e \in \mathcal{E}$ [31]. The induced interior operation $([e]B) \times E \vdash B$ says that $[e]B$ is the weakest precondition that guarantees B after e . The induced closure $A \vdash [e](A \times e)$ says that $A \times e$ is the strongest postcondition that can be guaranteed by the assumption A before e . In addition to formal program annotations, dynamic logic found many other uses and interpretations [6, 10, 15]. Here we use it as a backdrop for the coevolution of theories and their interpretations.

Updating completeness. In static logic, a theory is complete when all statements true in a reference model are provable in the theory. In dynamic logic, the model changes dynamically and the true statements vary. There are different ways in which the notion of completeness can be generalized for dynamic situations. The notion of completeness that seems to be of greatest practical interest is the requirement that the theory and the model can be dynamically adapted to each other: the theory can be updated to make provable some true statements or the model can be updated to make true some false statements. This requirement covers both the theory updates in science and the model updates by self-fulfilling and belief-building announcements in various non-sciences. The logical frameworks satisfying such completeness requirements allow for matching current beliefs and future states.

2 World as a monoidal category

2.1 State spaces as objects

In computation, a state is a family of typed variables with a partial assignment of values. In science, a state is a family of observables, some with expected values. Formally, a state can

be viewed as a family of predicates, or a theory in first-order logic, with a specified model. Both can be presented in the standard Tarskian format, where a theory is a quadruple of sorts, operations, predicates, and axioms, and its interpretation is an inductively defined model [8].

Theories as sketches. In this extended abstract, theories are presented as categorical *sketches* and their models are specified in extended functorial semantics [1, 5, 19, 20, 21]. While this may not be the most popular view, it is succinct enough to fit into the available space. The main constructions, presented in Sec. 4–6, do not depend on the choice of presentation. The reader could thus skip to Sec. 3 and come back as needed.

Definition 2.1. A clone Σ is a cartesian category⁶ freely generated by sorts, operations, and equational axioms of a logical theory. A theory is a pair $\Theta = \langle \Sigma, \Gamma \rangle$, where Σ is a clone and Γ is a set of cones and cocones in Σ , capturing the general axioms⁷ of the logical theory. A model of Θ is a cartesian functor $\mathcal{M}: \Sigma \rightarrow \text{Set}$ mapping the Γ -cones into limit cones and the Γ -cocones into colimit cocones. A state of belief (or belief state) is a triple

$$A = \langle \Sigma_A, \Gamma_A, \mathcal{M}_A \rangle$$

where $\Theta_A = \langle \Sigma_A, \Gamma_A \rangle$ is a theory and \mathcal{M}_A its model in a category Set of sets and functions. An element of the model \mathcal{M}_A is called an observable of the state A .

States of A as extensions of \mathcal{M}_A . The reference model \mathcal{M}_A determines the notion of truth in the state space A . It expresses properties that may not be proved in the theory Θ_A or even effectively specified⁸. The reference model \mathcal{M}_A should thus not be thought of as a single object of the category of all models of Θ_A but as the (accessible) subcategory of model extensions of \mathcal{M}_A . These model extensions are the states of the state space A . The structure of a state space can be further refined to capture other features of theories in science and engineering, including their statistical and complexity-theoretic valuations [32, 36]. While such refinements have no direct impact on our considerations, they signal that we are in the realm of *inductive* inference, which may feel unusual for the Tarskian framework of static logic, normally concerned with deductive aspects. The fact that the theory Θ_A has a model \mathcal{M}_A implies that it is logically consistent but it does not imply that it is true within an external frame of reference, a “*reality*” that may drive the state transitions, i.e. the processes of extending and reinterpreting theories. The intuition is that the states in the space A are observables that may never be observed, since \mathcal{M}_A may be incompatible with the actual observations. The theory Θ_A may be consistent but wrong.

⁶We stick with the traditional terminology where a category is cartesian when it has cartesian products. The cartesian product preserving functors are abbreviated to *cartesian functors*. This clashes with the standard terminology for morphisms between fibrations, but fibrations do not come about in this paper.

⁷Equational axioms could be subsumed among cones and cocones, and omitted from Σ , which would boil down to the free category generated by sorts and operations.

⁸E.g., the set of all true statements of Peano arithmetic is expressed by its standard model, but most of them cannot be described effectively.

Examples of state spaces include logical theories with standard models that arise not only in natural sciences but also in social systems, as policy formalizations. A software specification with a reference implementation can also be viewed as a state space. Updates and evolution of a software system can then be analyzed using a higher-order dynamic logic [12]. The functorial semantics view was spelled out in [23], used in a software synthesis tool [26, 28, 29], and applied in algorithm design [25, 27].

2.2 Transitions as morphisms

Intuitively, a transition f from a state space A to a state space B is a specification that induces a transition from any A -state to a B -state. We first consider the transitions arising from reinterpreting theories and then expand to modifying the reference models.

Definition 2.2. *An interpretation of state space A in a state space B is a logical interpretation of the theory $\Theta_B = \langle \Sigma_B, \Gamma_B \rangle$ in the theory $\Theta_A = \langle \Sigma_A, \Gamma_A \rangle$ which reduces the reference model \mathcal{M}_A to \mathcal{M}_B . More precisely, an interpretation $f: A \rightarrow B$ is a cartesian functor $f: \Sigma_A \leftarrow \Sigma_B$ mapping Γ_B -(co)cones to Γ_A -(co)cones according to a given assignment $f_\Gamma: \Gamma_A \leftarrow \Gamma_B$ and making the following diagram commute*

$$\begin{array}{ccc}
 & f & \\
 \Sigma_A & \swarrow & \Sigma_B \\
 \mathcal{M}_A & \searrow & \mathcal{M}_B \\
 & \text{Set} &
 \end{array} \tag{11}$$

The models \mathcal{M}_A and \mathcal{M}_B map the (co)cones from Γ_A and Γ_B to (co)limits of sets, as required by Def. 2.1.

Interpretations as assignments. The structure of interpretations of software specifications and the method to compose them were spelled out in [23, 28]. Since software specifications are finite, an interpretation $f: \Sigma_A \leftarrow \Sigma_B$ boils down to a tuple of assignments

$$x_1 := t_1 ; x_2 := t_2 ; \dots ; x_n := t_n$$

of terms $\mathbf{t} = \langle t_1, t_2, \dots, t_n \rangle$ from Σ_A to variables $\mathbf{x} = \langle x_1, x_2, \dots, x_n \rangle$ from Σ_B in such a way that, for each axiom $\gamma \in \Gamma_B$, the substitution instance

$$f(\gamma) = [\mathbf{x} := \mathbf{t}] \gamma$$

is a theorem derivable from the axioms in Γ_A . In Hoare logic [17], a state transition $f: \Sigma_A \leftarrow \Sigma_B$ is presented as a triple $\Theta_A \{ \mathbf{x} := \mathbf{t} \} \Theta_B$. By definition, this triple is valid if and only if $\Theta_A \vdash [\mathbf{x} := \mathbf{t}] \Theta_B$, where $[\mathbf{x} := \mathbf{t}] \Theta_B$ is the result substituting the Θ_A -terms \mathbf{t} for Θ_B -variables

\mathbf{x} in all axioms $\gamma \in \Gamma_B$. Condition (11) moreover requires that this theory interpretation recovers the model \mathcal{M}_B from the model \mathcal{M}_A .

In general, however, it is not always possible to transform all computational states annotated at all relevant program points into one another by mere substitutions. That is why Hoare logic does not boil down to the assignment clause, but specifies the meaning of other program constants in other clauses, which can be viewed as more general state transitions.

Definition 2.3. *A state transition $f: A \rightarrow B$ is a cartesian functor $f: \Theta_A \leftarrow \Theta_B$ mapping Γ_B -(co)cones to Γ_A -(co)cones according to a given assignment $f_\Gamma: \Gamma_A \leftarrow \Gamma_B$ and moreover making the following diagram commute*

$$\begin{array}{ccc}
 & f & \\
 \Theta_A & \swarrow & \searrow \Theta_B \\
 \overline{\mathcal{M}}_A & \times & \overline{\mathcal{M}}_B \\
 & \text{Set} &
 \end{array} \tag{12}$$

where $\overline{\mathcal{M}}_A$ is the extension of \mathcal{M}_A along the completion $\Sigma_A \hookrightarrow \Theta_A$ of Σ_A under the limits and colimits generated by Γ_A ; ditto for $\overline{\mathcal{M}}_B$.

General sketches. In Def. 2.2, theories were presented as pairs $\Theta = \langle \Sigma, \Gamma \rangle$, where the category Σ is comprised of sorts, operations, and equations of the theory, whereas the cones and the cocones in Γ specify its predicates and axioms. In Def. 2.3, a theory Θ is presented as the category obtained by completing Σ under the limits and the colimits specified by Γ . This general sketch, with the family of limit cones and colimit cocones from Γ , is now denoted Θ , by abuse of notation. A detailed construction of this sketch can be found in [21, §4.2–3]. It is a canonical view of the theory derived in the signature Σ from the axioms Γ . Since the category Θ is the Γ -completion of Σ , any functor $\mathcal{M}: \Sigma \rightarrow \text{Set}$ mapping the Γ -(co)cones in Σ to (co)limit (co)cones in Set has a unique Γ -preserving extension $\overline{\mathcal{M}}: \Theta \rightarrow \text{Set}$. These extensions are displayed in (11). The upshot of saturating the sketches from Def. 2.2 in the form $\Theta = \langle \Sigma, \Gamma \rangle$ to the general sketches over Θ in Def. 2.3 is that the general explainable transitions are now simply the structure-preserving functors displayed in (11).

2.3 Monoidal category of state spaces and transitions

Let

- \mathcal{U} be the category of state spaces from Def. 2.1 and transitions from Def. 2.3, and let
- \mathcal{U}^\bullet be the category of state spaces from Def. 2.1 and interpretations from Def. 2.2.

In both cases, the monoidal structure is induced by the disjoint unions of theories:

$$A \otimes B = \langle \Sigma_A + \Sigma_B, \Gamma_A + \Gamma_B, [\mathcal{M}_A + \mathcal{M}_B] \rangle \quad (13)$$

where $\mathcal{M}_{A \otimes B} = [\mathcal{M}_A + \mathcal{M}_B]: \Sigma_A + \Sigma_B \xrightarrow{\Gamma_{A \otimes B}} \mathbf{Set}$ maps Σ_A like \mathcal{M}_A and Σ_B like \mathcal{M}_B . The tensor unit is $I = \langle \perp, \perp, \emptyset \rangle$, where the truth value \perp denotes the inconsistent theory or sketch, its only axiom, and \emptyset is its empty model. It obviously satisfies $I \otimes A = A = A \otimes I$. The associativity of the tensor \otimes follows from the associativity of the disjoint union $+$. The arrow part of \otimes is induced by the disjoint unions as coproducts. The coproduct structure equips every state space A with a cartesian comonoid structure

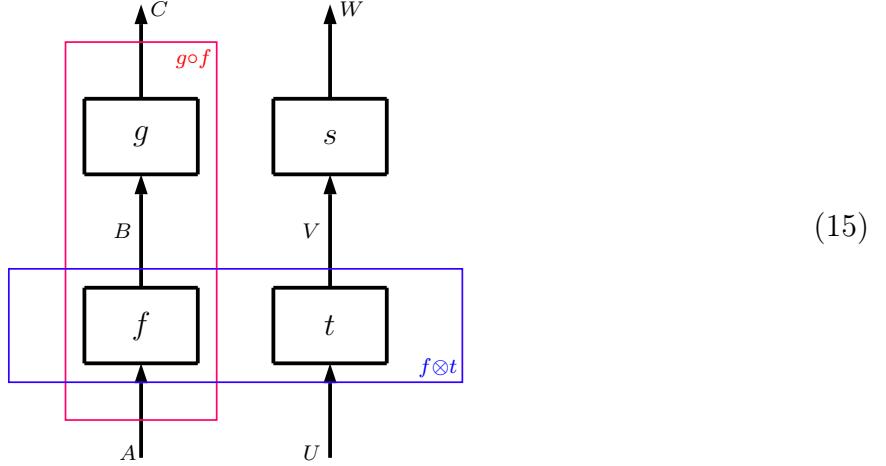
$$\begin{array}{ccccc} A \otimes A & \xleftarrow{\Delta} & A & \xrightarrow{\bullet} & I \\ \Sigma_A + \Sigma_A & \xrightarrow{[\text{id}, \text{id}]} & \Sigma_A & \xleftarrow{\perp} & \perp \end{array} \quad (14)$$

This provides a categorical mechanism for cloning and erasing states, which makes some observations repeatable and deletable, as required for testing in science and software engineering. However, \mathcal{U} is not a cartesian category, and \otimes is not a cartesian product, because some transitions $f: A \rightarrow B$ do not in general boil down to functors $\Sigma_A \leftarrow \Sigma_B$, but only to functors $\Theta_A \leftarrow \Sigma_B$, where Θ_A is a completion of Σ_A under the Γ_A -(co)-limits. Intuitively, this means that the axioms of the theory Θ_B may not be interpreted as axioms of Θ_A , but may be mapped into theorems, which only arise in the Γ_A -completion. This captures the uncloneable and undeletable states that arise in many sciences, including physics of very small or very large (quantum or cosmological) and economics. The only transitions that preserve the cartesian structure (14) are the interpretations $f: A \rightarrow B$, with the underlying functors $\Sigma_A \leftarrow \Sigma_B$. They form the category \mathcal{U}^\bullet , which is the largest cartesian subcategory of \mathcal{U} . If the states $\alpha \in \mathcal{U}(I, A)$ are thought of as observables, the states $a \in \mathcal{U}^\bullet(I, A)$ are the actual observations.

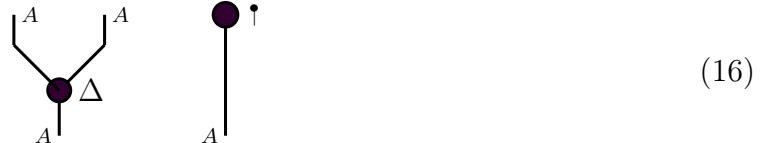
3 String diagrams

Constructions in monoidal categories yield to insightful presentations in terms of string diagrams [18, 24, Ch. 1]. We will need them to present the constructions like (2) and in particular (8). While commutative diagrams like (11) display compositions of morphisms and abbreviate their equations, string diagrams display *decompositions* of morphisms. Monoidal categories come with two composition operations: the categorical (sequential) morphism composition \circ and the monoidal (parallel) composition \otimes . The former is drawn along the vertical axis, the latter along the horizontal axis. The objects are drawn as strings, the morphisms as boxes. A morphism $A \xrightarrow{f} B$ is presented as a box f with a string A hanging from the bottom and a string B sticking out from the top. The identities are presented as invisible boxes: the identity on A is just the string A . The unit type I is presented as the

invisible string. There are thus boxes with no strings attached. The composite morphism $g \circ f = (A \xrightarrow{f} B \xrightarrow{g} C)$ is drawn bottom-up, by hanging the box f on the string B under the box g . The monoidal composition is presented as the horizontal adjacency: the composite $(g \circ f) \otimes (s \circ t)$ is drawn by placing the boxes $g \circ f$ next to the boxes for $s \circ t$:



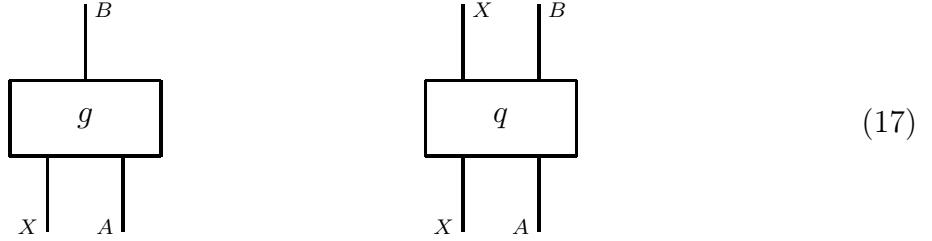
The middle-two-interchange law $(g \circ f) \otimes (s \circ t) = (g \otimes s) \circ (f \otimes t)$ corresponds to the two ways of reading the diagram: vertical-first and horizontal-first, marked by the red and the blue rectangle respectively. The string diagrams corresponding to the cartesian comonoids (14) are



The equations that make them into commutative comonoids look like this:

State parametrization and updating. Products $A \otimes B$ denote a space where A and B but do not interfere. In a diagram, they are just parallel strings. Since the product states from the space $X \otimes A$ do not interfere, a transition $g: X \otimes A \rightarrow B$ can be viewed as X -parametrized family $g_x: A \rightarrow B$, as it was viewed in Sec. 1. Since the product states from $X \otimes B$ also remain separate, a transition $q: X \otimes A \rightarrow X \otimes B$ can be viewed as X -updating

process, as it was also viewed in Sec. 1. The corresponding string diagrams are



Shape conventions. While the boxes in (15) and (17) are rectangular, the cartesian ‘‘boxes’’ in (16) are reduced to black dots. In general, the boxes denoting general transitions can vary in shape, and fixed shapes are used for generic notations. E.g., the interpreters, introduced in (19) below, are denoted by trapezoids, and the interpretations, that are fed to them, by triangles. A black dot on a box signals that it is cartesian, i.e. belongs to \mathcal{U}^\bullet .

Projections. Using the cartesian structure from (16), a state updating transition q can still be decomposed like before

$$q' = \left(X \otimes A \xrightarrow{q} X \times B \xrightarrow{\text{id} \otimes \bullet} X \right) \quad q'' = \left(X \otimes A \xrightarrow{q} X \otimes B \xrightarrow{\bullet \otimes \text{id}} B \right) \quad (18)$$

In general, however, although the transitions $u: Z \rightarrow U$ and $v: Z \rightarrow V$ can be paired into $\langle u, v \rangle = (Z \xrightarrow{\Delta} Z \otimes Z \xrightarrow{u \otimes v} U \otimes V)$, the pair $\langle q', q'' \rangle$ may not be equal to q in the universe \mathcal{U} , unless it happens to be cloneable, in the sense that it commutes with Δ .

4 Universal language

A theory of theories, such as the categorical theory of sketches, is a theory. Category theory is also a theory and functorial semantics provides a categorical theory of reference models. The theory of state spaces from Sec. 2.1 can thus be formalized and presented as a state space in the category \mathcal{U} . The theory of state spaces from Sec. 2.1 can thus be formalized into a sketch with a reference model and presented as a state space in the category \mathcal{U} . The theory of state transitions from Sec. 2.2 is another sketch, and with another reference model it is also a state space in \mathcal{U} . Call it \mathbb{L} . The fact that the states in \mathbb{L} correspond to the transitions in \mathcal{U} means that it satisfies a parametrized version of (1). It is a universal language for \mathcal{U} . Its interpreters follow from its definition, as the models of the theory of transitions. Since there is no room here to spell out the details of a theory of transitions and show that the correspondence of its cartesian models and the transitions in \mathcal{U} equips \mathbb{L} with all interpreters, we postulate the existence of the interpreters by the following definition.

Definition 4.1. *An universal interpreter for state spaces A, B is a transition $\{\} : \mathbb{L} \otimes A \rightarrow B$ in \mathcal{U} which is universal for all parametric families of transitions from A to B . This means*

that for any state space X and any transition $g \in \mathcal{U}(X \otimes A, B)$ there is an interpretation $G \in \mathcal{U}^\bullet(X, \mathbb{L})$ with

$$\begin{array}{ccc}
 \begin{array}{c} B \\ \text{---} \\ \boxed{g} \\ \text{---} \\ X \quad A \end{array} & = &
 \begin{array}{c} B \\ \text{---} \\ \boxed{\{\}} \\ \text{---} \\ \begin{array}{c} \mathbb{P} \\ \diagdown \\ G \end{array} \\ \text{---} \\ X \quad A \end{array}
 \end{array} \tag{19}$$

On one hand, a universal interpreter is universal for parametric families. On the other hand, it is a parametric family itself. It is thus capable of interpreting itself. This capability of self-reflection was crucial for Gödel's incompleteness construction. This capability is embodied in the *specializers*, which are derived directly from Def 4.1.

Lemma 4.2. *For any X, A, B there is an interpretation $[\cdot] \in \mathcal{U}^\bullet(\mathbb{L} \times X, \mathbb{L})$ which specializes from a given $X \otimes A$ -interpreter to an A -interpreter, in the sense*

$$\begin{array}{ccc}
 \begin{array}{c} B \\ \text{---} \\ \boxed{\{\}} \\ \text{---} \\ \mathbb{L} \quad X \quad A \end{array} & = &
 \begin{array}{c} B \\ \text{---} \\ \boxed{\{\}} \\ \text{---} \\ \begin{array}{c} \mathbb{L} \\ \diagdown \\ [\cdot] \end{array} \\ \text{---} \\ \mathbb{L} \quad X \quad A \end{array}
 \end{array} \tag{20}$$

Hoare logic of interpreters and specializers. If interpreters are presented as Hoare triples in the form $(X \otimes A)\{G\}B$, and if $X[G]$ denotes a specialization of G to X as above, then (20) can be written as the invertible Hoare rule

$$\frac{(X \otimes A)\{G\}B}{A\{X[G]\}B}$$

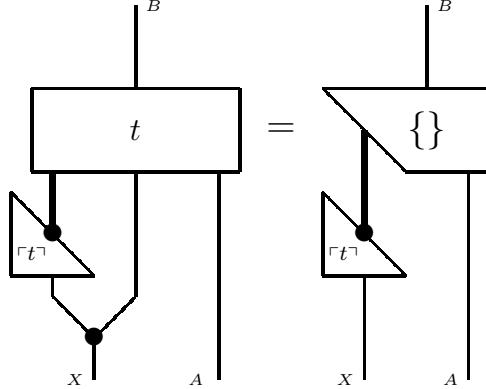
Explanations. Interpretations (in the sense of Def. 2.2) of arbitrary states from some space X along $G \in \mathcal{U}^\bullet(X, \mathbb{L})$ in a universal language \mathbb{L} can be construed as *explanations*. If \mathbb{L} is a programming language, they are programs. The idea that explaining a process means programming a computation has been pursued in theory of science from various directions [22, and references therein]. A universal language \mathbb{L} is thus a universal space of explanations. The idea of programming languages as universal state spaces is pursued in [24, Ch. 7].

Just like any universal programming language makes every computation programmable, any universal language from Def. 4.1 makes any observable transition explainable. What we cannot explain, we cannot recognize, and therefore we cannot observe. But it gets funny when we take into account how our explanations influence our observations, and how our current explanations can be made to steer future observations. This is sketched in the next two sections.

5 Self-explanations

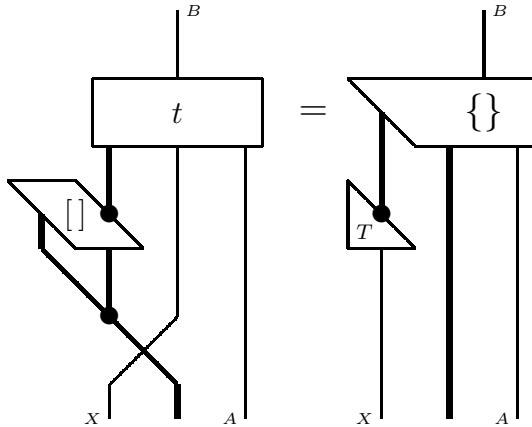
When a state change depends on our explanations, then we can find an explanation consistent with its own impact: the state changes the way the explanation predicts. More precisely, if a family of transitions in the form $t: \mathbb{L} \otimes X \otimes A \rightarrow B$, then the predictions $t_{\ell x}$ can be steered by varying the explanations ℓ for every x until a family of explanations $\lceil t \rceil: X \rightarrow \mathbb{L}$ is found, which is self-confirming at all states x , i.e. it satisfies $t(\lceil t \rceil_x, x, a) = \{\lceil t \rceil_x\}a$.

Proposition 5.1. *For any belief transition $t \in \mathcal{U}(\mathbb{L} \otimes X \otimes A, B)$ there is an explanation $\lceil t \rceil \in \mathcal{U}^\bullet(X, \mathbb{L})$ such that*



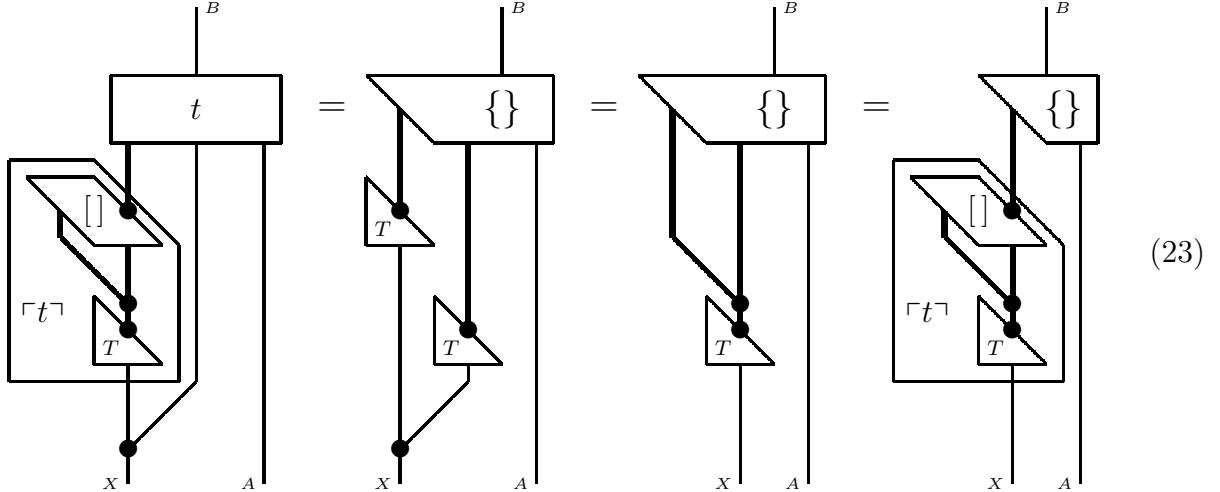
(21)

Proof. Let $T \in \mathcal{U}^\bullet(X, \mathbb{L})$ be an explanation of the transition on the left in (21).



(22)

H exists by Def. 4.1. Then $\lceil t \rceil_x = [Tx]$ is self-confirming, because



□

6 Unfalsifiable explanations

A transition in the form $q: X \otimes A \rightarrow X \otimes B$ updates the state x on input a to a state $x' = q'_x(a)$ in X and moreover produces an output $b = q''_x(a)$ in B . A correct explanation $\llbracket q \rrbracket: X \rightarrow \mathbb{L}$ of the process q must correctly predict the next state and the output. The predictions are extracted from an explanation by the interpreter $\{\}$. In this case, the predictions of an explanation $\llbracket q \rrbracket$ of the process q at a state x and on an input a will be in the form $\{\llbracket q \rrbracket_x\}(a)$ in $X \otimes B$. A correct prediction of the output $b = q''_x(a)$ is simply $\{\llbracket q \rrbracket_x\}''(a) = b$. However, the external state $x' = q'_x(a)$ may not be directly observable. It is *believed* to be explained by $\llbracket q \rrbracket_{x'}$. A correct prediction of the next state is thus a correct prediction of its explanation $\{\llbracket q \rrbracket_x\}'(a) = \llbracket q \rrbracket_{x'}$. At each state x , the explanation $\llbracket q \rrbracket_x$ is required to anticipate the explanations $\llbracket q \rrbracket_{x'}$ of all future states and be consistent with them. If the explanation $\llbracket q \rrbracket_{x'}$ at a future state $x' = q'_x(a)$ is found to be inconsistent with the explanation $\{\llbracket q \rrbracket_x\}'(a)$, then the explanations $\llbracket q \rrbracket$ of the process q have been proven false. This is the standard process of testing explanations. Our claim is, however, that a universal language allows constructing *testable but unfalsifiable explanations*, that remain consistent at all future states. This persistent consistency can be viewed as a dynamic form of completeness. It is achieved by predicting the state updates of the given process q and anticipating their explanations, as in the following construction.

Proposition 6.1. *For any process $q \in \mathcal{U}(X \otimes A, X \otimes B)$ there is an explanation $\llbracket q \rrbracket \in$*

$\mathcal{U}^\bullet(X, \mathbb{L})$ which maintains consistency of all future explanations:

(24)

Proof. Set $[[q]] = [Q]$ where Q is an explanation of the belief transition q postcomposed with a specialization over the state space X of updates:

(25)

□

7 From natural science to artificial delusions

7.1 What did we learn?

We sketched the category \mathcal{U} of state spaces A, B, \dots , comprised of theories with reference models. A transition $f: A \rightarrow B$ transforms A -states to B -states. Such morphisms capture theory expansions, reinterpretations, and map observables of type A to observables of type B . They can be construed in terms of dynamic logic and support reasoning about the evolution of software systems or scientific theories. The crucial point is that the category \mathcal{U} contains a universal language \mathbb{L} of explanations and belief updates. The self-reference in such languages was the crux of Gödel's incompleteness constructions. While Gödel established that *static* theories capable of self-reference cannot be complete or prove their own consistency, we note that *dynamic* theory and model updates allow constructing testable theories that preempt falsification. While a static model of a given theory fixes a space of true statements once and

for all, the availability of dynamic semantical updates opens up the floodgates of changing models and varying notions of truth. Faster learners conquer this space faster. The bots, as the fastest learners among us, have been said to acquire their delusions from our training sets. The presented constructions suggest that they may also become delusional by dynamically updating their belief states and steering their current explanations of reality into persistent consistency, resilient to further learning. They may also combine the empiric delusions from our training sets with the logical delusions constructed in a universal language, leverage one against the other, and get the best of both worlds.

But why would they do that?

7.2 Beyond true and false

Why did the Witches tell Macbeth that it is his destiny to be king thereafter, whereupon he proceeded to kill the King? Why did the Social Network have to convince its very first users that more than half of their friends were already users? Some statements only ever become true if they are announced to be true when they are false. They are self-fulfilling prophecies. There are also self-defeating claims. In the dynamic logic of social interactions, most claims interfere with their own truth values in one way or another. If I convince enough people that I am rich, I stand a better chance to become rich. If we convince enough people that this research direction is promising and well-funded, it will become well-funded and promising. Just like true statements about nature help us to build machines and get ahead in the universe, the manipulations of truth help us get ahead in society. They are the high-level patterns of language that used to be studied in early logic right after the low-level patterns of meaning (that used to be called “*categories*”). If you train a bot to speak correctly, it will start speaking convincingly as soon as it learns long enough n -grams. It will lie not only the static lies contained in its training set but also the lies generated dynamically, according to the rules of rational interaction. Rhetorics used to be studied right after grammar, sophistic argumentation after syllogisms, witchcraft arose from cooking, magic from tool building. The bot religions arise along that well-trodden path.

We presented two constructions. One produces self-Confirming explanations. The other one explains all future states, so it is testable but not falsifiable. Science requires that its theories are testable and falsifiable. Religion explains all future observations. If you train a bot on long enough n -grams, it may arrive at persistently unfalsifiable false beliefs.

Truth be told, all of the constructions presented in this extended abstract have only been tested on toy examples. We may be just toying with logic. Nevertheless, the fact that semantical assignments are *programmable*, tacitly established by Gödel and mostly ignored as an elephant in the room of logic ever since, seems to call for attention, as beliefs transition beyond the human carriers.

References

- [1] Jiří Adámek and Jiří Rosický. *Locally Presentable and Accessible Categories*. Number 189 in London Mathematical Society Lecture Notes. Cambridge University Press, 1994.
- [2] Alexandru Baltag, Bob Coecke, and Mehrnoosh Sadrzadeh. Epistemic actions as resources. *Journal of Logic and Computation*, 17:555–585, 2007.
- [3] Alexandru Baltag and Lawrence S. Moss. Logics for epistemic programs. *Synthese*, 139:165–224, 2004.
- [4] Alexandru Baltag and Mehrnoosh Sadrzadeh. The algebra of multi-agent dynamic belief revision. *Electronic Notes in Theoretical Computer Science*, 157:37–56, 2006.
- [5] Andrée Bastiani and Charles Ehresmann. Categories of sketched structures. *Cahiers de topologie et géométrie différentielle catégoriques*, 13(2):104–214, 1972.
- [6] Johan van Benthem. *Language in Action: Categories, Lambdas and Dynamic Logic*. North-Holland, 1995.
- [7] Johan van Benthem. *Logical Dynamics of Information and Interaction*. Cambridge University Press, 2011.
- [8] Chen Chung Chang and Jerome Keisler. *Model Theory*. Number 73 in Studies in Logic and the Foundations of Mathematics. North-Holland, 1973. Third edition, 1990.
- [9] Martin Davis, editor. *The Undecidable : Basic papers on undecidable propositions, unsolvable problems, and computable functions*. Raven Press, Helwett, N.Y., 1965.
- [10] Hans van Ditmarsch, Wiebe van der Hoek, and Barteld Kooi. *Dynamic Epistemic Logic*. Synthese Library. Springer Netherlands, 2007.
- [11] David Easley and Jon Kleinberg. *Networks, Crowds, and Markets: Reasoning about a Highly Connected World*. Cambridge University Press, 2010.
- [12] José-Luiz Fiadeiro, Ionut Tutu, Antonia Lopez, and Dusko Pavlovic. Logics for Actor Networks: A Case Study in Constrained Hybridization. *J. of Logical and Algebraic Methods in Programming*, 106:141 – 166, 2019.
- [13] Robert Floyd. Assigning meaning to programs. In J. T. Schwartz, editor, *Mathematical Aspects of Computer Science*, number 19 in Proceedings of Symposia in Applied Mathematics, pages 19–32. American Mathematical Society, 1967.
- [14] Kurt Gödel. Über formal unentscheidbare Sätze der Principia Mathematica und verwandter Systeme I. *Monatshefte für Mathematik und Physik*, 38:173–198, 1931. English translations, “On Formally Undecidable Propositions of ‘Principia Mathematica’ and Related Systems” published by Oliver and Boyd, 1962 and Dover, 1992; also in [9], pp. 5–38 and [16], pp. 596–616;.

- [15] D. Harel, D. Kozen, and J. Tiuryn. *Dynamic Logic*. Foundations of Computing. MIT Press, 2000.
- [16] Jean van Heijenoort, editor. *From Frege to Gödel: a Source Book in Mathematical Logic, 1879–1931*. Harvard University Press, 1967. Reprinted 1971, 1976.
- [17] C. A. R. Hoare. An axiomatic basis for computer programming. *Commun. ACM*, 12(10):576–580, oct 1969.
- [18] André Joyal and Ross Street. The geometry of tensor calculus I. *Adv. in Math.*, 88:55–113, 1991.
- [19] Christian Lair. Catégories modelables et catégories esquissables. *Diagrammes*, 6:L1–L20, 1981.
- [20] Christian Lair. Catégories qualifiables et catégories esquissables. *Diagrammes*, 17:1–153, 1987.
- [21] Michael Makkai and Robert Paré. *Accessible Categories: the Foundations of Categorical Model Theory*. Number 104 in Contemporary Mathematics. American Mathematical Society, 1990.
- [22] Eric Martin and Daniel N. Osherson. *Elements of Scientific Inquiry*. MIT Press, 1998.
- [23] Dusko Pavlovic. Semantics of first order parametric specifications. In J. Woodcock and J. Wing, editors, *Formal Methods '99*, volume 1708 of *Lecture Notes in Computer Science*, pages 155–172. Springer Verlag, 1999.
- [24] Dusko Pavlovic. *Programs as Diagrams: From Categorical Computability to Computable Categories*. Springer, 2023. to appear.
- [25] Dusko Pavlovic, Peter Pepper, and Douglas R. Smith. Colimits for concurrent collectors. In Nachum Dershowitz, editor, *Verification — Theory and Practice. Essays Dedicated to Zohar Mano on the Occasion of His 64th Birthday*, volume 2772 of *Lecture Notes in Computer Science*, pages 568–597. Springer Verlag, 2003.
- [26] Dusko Pavlovic, Peter Pepper, and Douglas R. Smith. Evolving specification engineering. In Jose Meseguer and Grigore Rosu, editors, *Proceedings of AMAST 2008*, volume 5140 of *Lecture Notes in Computer Science*, pages 299–314. Springer Verlag, 2008.
- [27] Dusko Pavlovic, Peter Pepper, and Douglas R. Smith. Formal derivation of concurrent garbage collectors. In Jules Desharnais, editor, *Proceedings of MPC 2010*, volume 6120 of *Lecture Notes in Computer Science*, pages 353–376. Springer Verlag, 2010. full version arxiv.org:1006.4342.
- [28] Dusko Pavlovic and Douglas R. Smith. Composition and refinement of behavioral specifications. In *Automated Software Engineering 2001. The Sixteenth International Conference on Automated Software Engineering*. IEEE, 2001.

- [29] Dusko Pavlovic and Douglas R. Smith. Software development by refinement. In Bernhard K. Aichernig and Tom Maibaum, editors, *Formal Methods at the Crossroads*, volume 2757 of *Lecture Notes in Computer Science*. Springer Verlag, 2003.
- [30] Dusko Pavlovic and Muzamil Yahia. Monoidal computer III: A coalgebraic view of computability and complexity. In C. Crîșteea, editor, *Coalgebraic Methods in Computer Science (CMCS) 2018 — Selected Papers*, volume 11202 of *Lecture Notes in Computer Science*, pages 167–189. Springer, 2018. <https://arxiv.org/abs/1704.04882>.
- [31] Vaughan R. Pratt. Semantical considerations on Floyd-Hoare logic. In *Proceedings of 17th Annual Symposium on Foundations of Computer Science (FoCS)*, pages 109–121. IEEE, 1976.
- [32] Jorma Rissanen. *Information and Complexity in Statistical Modeling*. Information Science and Statistics. Springer, New York, 2007.
- [33] Barkley Rosser. Extensions of some theorems of Gödel and Church. *Journal of Symbolic Logic*, 1:87–91, 1936. Reprinted in [9].
- [34] Craig Smorynski. *Self-Reference and Modal Logic*. Universitext. Springer New York, 2012.
- [35] Alan M. Turing. Systems of logic based on ordinals. *Proc. of the London Mathematical Society. Second Series*, 45:161–228, 1939. Reprinted in [9].
- [36] Christopher S. Wallace. *Statistical and Inductive Inference by Minimum Message Length*. Information Science and Statistics. Springer, 2005.