



Semantic Foundations of Reductive Reasoning

Alexander V. Gheorghiu^{1,2} · David J. Pym^{1,3,4}

Accepted: 17 April 2025
© The Author(s) 2025

Abstract

The development of logic has largely been through the *deductive* paradigm: conclusions are inferred from established premisses. However, the use of logic in the context of both human and machine reasoning is typically through the dual *reductive* perspective: collections of sufficient premisses are generated from putative conclusions. We call this paradigm, *reductive logic*. This expression of logic encompass as diverse reasoning activities as proving a formula in a formal system to seeking to meet a friend before noon on Saturday. This paper is a semantical analysis of reductive logic. In particular, we provide mathematical foundations for representing and reasoning about *reduction operators*. Heuristically, reduction operators may be thought of as ‘backwards’ inference rules. In this paper, we address their mathematical representation, how they are used in the context of reductive reasoning, and, crucially, what makes them *valid*.

Keywords Logical systems · Reductive logic · Proof-search · Theorem proving · Tactics · Tactical proof · Proof-theoretic semantics

1 Introduction

The definition of a system of logic may be given *proof-theoretically* as a collection of rules of inference that, when composed in specified ways, determine *proofs*; that is, formal constructions that establish that a conclusion is a consequence of some assumptions or axioms. In other words, proofs are objects regulated by rules from a given formal system that determine that an inference of a conclusion from a collection of premisses has been established:

$$\frac{\text{Established Premiss}_1 \quad \dots \quad \text{Established Premiss}_k}{\text{Conclusion}} \Downarrow$$

We call this proof-theoretic formulation *deductive logic*.

✉ Alexander V. Gheorghiu
a.v.gheorghiu@soton.ac.uk

David J. Pym
d.pym@ucl.ac.uk

- ¹ Department of Computer Science, University College London, Gower Street, London WC1E 6EA, UK
- ² School of Electronics and Computer Science, University of Southampton, University Street, Southampton SO17 1BJ, UK
- ³ Department of Philosophy, University College London, Gordon Square, London WC1H 0AW, UK
- ⁴ Institute of Philosophy, University of London, Malet Street, London WC1E 7HU, UK

Deductive logic is useful as a way of defining what proofs are, but it does not reflect either how logic is typically used in practical reasoning problems or the method by which proofs are found. In practice, proofs are typically constructed by starting with a desired, or putative, conclusion and applying the rules of inference *backwards*. Read from conclusion to premisses, the rules are sometimes called *reduction operators*, and denoted

$$\frac{\text{Sufficient Premiss}_1 \quad \dots \quad \text{Sufficient Premiss}_k}{\text{Putative Conclusion}} \Uparrow$$

We call the constructions in a system of reduction operators *reductions*. This proof-theoretic formulation has been dubbed *reductive logic* (Pym and Ritter 2004).

Note that reductive logic is not *a* logic in the sense of, for example, classical, intuitionistic, or modal logic, but a conception of logic in the sense of *deductive*, *inductive*, or *abductive* logic.

To clarify, reductive logic is defined with respect to a formal specification of a logic and does not encompass semantic reasoning in general. Specifically, the figures above illustrate particular inferences within formal systems, rather than arbitrary steps of problem-solving. As such, reductive logic represents a relatively narrow domain within the broader field of analysis (Beaney and Raysmith 2024). More precisely, it

pertains to the aspects of analysis employed in computer-assisted formal reasoning, as we discuss below.

Also to avoid confusion, note that reductive logic should not be conflated with the superficially related concept of reverse mathematics, as developed by Friedman (1975) and others. Reverse mathematics investigates which axioms, within a given formalization of mathematics, are necessary for specific theorems to hold. By contrast, reductive logic focuses on the proof theory and semantics underlying the process of finding proofs—via reduction and search—within fixed systems of logic.

Focusing on areas with formal, well-defined notions of inference allows us to develop a more concrete semantic theory of reductive reasoning. Specifically, our aim is to create a semantics of reductions that not only identifies complete proofs but also clarifies the meaning of unfinished or incomplete proof-searches. This approach places our investigation squarely within the domain of proof-theoretic semantics (P-tS), which seeks to explain the validity of proofs in formal logic. So far, P-tS has largely been developed in the context of deductive logic. We have argued that it is a natural and metaphysically well-motivated semantic framework that is explanatory for both reductive and deductive logic. We therefore argue that this suggests that these two views of formal logic are as valid as each other as ways of modelling reasoning.

For a comprehensive treatment of P-tS, we refer to Schroeder-Heister (2006, 2008), though we provide the key highlights here. While P-tS is typically framed in the context of deductive logic, its core principles are directly applicable to reductive reasoning. In the Dummett-Prawitz tradition of P-tS—perhaps the most extensively studied—direct proofs (closed, normalized proofs) are considered canonically valid, whereas indirect proofs rely on transformations to achieve directness. This distinction introduces a technical separation between proof structures, which represent potential proofs, and justifications, which define the transformations that validate them. Similarly, in reductive reasoning, several steps of reduction may yield a potential proof or contribute proof-theoretic content, even if additional manipulation and transformation are required to achieve the final proof of the intended goal—that is, the putative conclusion one desires to prove.

Naïvely, reductive logic may appear simply as deductive logic read backwards. This heuristic suffices in many cases and is useful for understanding deductive logic; for example, proof of cut-elimination for sequent calculi (see Szabo 1969) can be understood from this perspective. However, the space of reductions studied in reductive logic is larger than the space of proofs, as studied in deductive logic: while the reductive system can simply reconstruct all the correct proofs in the underlying system, it also allows choices to be made that yield reductions that do not correspond to proofs. For

example, when no further reduction operators can be applied to a premiss and that premiss is not an axiom of the logic. Reductive logic is thus a paradigm of logic in parallel to deductive logic and cannot be defined purely in terms of it.

While the restriction to proof systems may seem limiting, the conception of reductive logic in this paper is sufficiently general to encompass reasoning activities as diverse as proving a formula in a formal system and seeking to meet a friend before noon on Saturday. Importantly, it more closely resembles what mathematicians do when proving theorems and, more generally, how people solve problems using formal representations. It is also the paradigm of logic used for diverse applications in informatics and other systems-oriented sciences, including, but not limited to, areas such as program and system verification, natural language processing, and knowledge representation. It is the reductive perspective that underpins the use of logic in proof-assistants—for example, LCF (Gordon et al. 1979), HOL (Gordon and Melham 1993), Isabelle (Paulson 1994), Coq (Coq Team 2024; Bertot and Castéran 2004), Twelf (Frank Pfenning et al. 2024; Pfenning and Schürmann 1999), and more (xxxx xxxx). More generally, it is a restricted model of human reasoning, which is discussed extensively, for example, in the work of Kowalski (1986) and Bundy (1985). We discuss this further in Sect. 3.

We seek a semantic foundation for reduction operators. Our motivation is both philosophical and pragmatic. Philosophically, we recognize that reductive logic is a paradigm through which much practical problem solving occurs, as discussed above (examples below). Pragmatically, the increasing significance of technologies reliant on reductive logic suggests the pressing need for a more comprehensive and mathematically more developed semantic theory than currently available—see, for example, Pym and Ritter (2004). Hence, in particular, we seek a representation of reduction operators as functions encapsulating the action suggested by the picture above,

$$\begin{array}{l} \text{Putative Conclusion} \quad \longrightarrow \quad \text{Sufficient Premiss}_1, \dots, \\ \text{Sufficient Premiss}_n \end{array}$$

Having done this, we can then analyze them and formally study reductive reasoning. It is not only the representation of reduction operators as mathematical objects that is important: we seek an account of what makes a reduction operator and its deployment *valid*.

Historically, semantics in logic has been dominated by the paradigm of *denotationalism*, where the meaning of logical structures is given in terms of abstract interpretations. For instance, in model-theoretic semantics, propositions are interpreted relative to abstract algebraic structures called *models* (see Schroeder-Heister 2008). This paradigm is suitable for *deductive* logic. However, other paradigms of logic may benefit from different approaches.

For *reductive* logic, denotationalist accounts have been useful and thoroughly explored (see Pym and Ritter 2004, Komendantskaya et al. 2011; Komendantskaya and Power 2011; Komendantskaya et al. 2016, and Gheorghiu et al. (2023a)). These studies are discussed below. In this paper, we generalize these results by removing auxiliary and domain-specific aspects to concentrate on the semantics of reductive reasoning itself.

We also examine semantics with a *proof-theoretic* focus, this is the connection to P-tS. That is, we give validity conditions not in terms of abstract algebraic structures, but in terms of what ‘proofs’ are being witnessed during reductive reasoning. Pym and Wallen (1993) provided such an account in terms of *proof-valued computations*.

To this end, we draw substantively upon Milner’s theory of *tactical proof* (Milner 1984) in Sect. 4. The historical and technical significance of tactical proof in the context of computer-assisted formal reasoning is profound; it forms the basis for almost all proof assistants (see Gordon 2015). Ultimately, this inspires the P-tS account of reductive logic in this paper.

We begin, in Sect. 2, by explicating the notion of reduction and its current semantics in comparison to well-established techniques in deductive logic. This helps expose the subtlety and challenges involved in reductive logic. In Sect. 3, we provide a denotational semantics of the reduction operators and soundness and completeness conditions that determine the validity of a set of reduction operators. In Sect. 4, we give an operational account of reduction operators in terms of structural operational semantics that enables a more refined account of their use and validity. In Sect. 5, we discuss future problems for a satisfactory foundation of reductive reasoning; viz., the expression and semantics of strategies for proof-search. This moves us from a semantic foundation of reduction operators, which we may think of as the *local* perspective on reduction reasoning, to a full semantic foundation of reductive logic, the corresponding *global* perspective on reductive logic. The paper ends in Sect. 6 with a summary of the results obtained.

This paper should be seen as a companion to Gheorghiu and Pym (2023a), in which P-tS is used to provide a semantics for tactical proof, mentioned above. By contrast, this paper provides a general foundation for reductive reasoning in terms of P-tS and obtains tactical proof as a significant instance. Happily, the two analyses are mutually coherent.

Technical Background. We assume familiarity with formal logic (see, for example, van Dalen 2012) and basic proof theory (see, for example, Troelstra and Schwichtenberg 2000 and Negri (2005)). We use $\varphi, \psi, \chi, \dots$ to denote logical formulae and Γ, Δ, \dots to denote collections thereof. We use \triangleright as the sequent symbol and we distinguish it from the consequence judgement \vdash of a logic; that is, a sequent is an

expression of an assertion in a logic, which may be valid or invalid, but a consequence is an assertion that is valid.

For the technical portions of this paper, we will assume some basic concepts and terminology from category theory—see, for example, Cheng (2022)—in particular, categories, functors, adjunctions, and natural transformations. However, these technical details are not essential for the thesis, they simply give a systematic mathematical treatment of the ideas presented. As we are always working in the category of sets and (partial) functions, relatively little background is required beyond typical mathematical training.

We use $\mathcal{P}_{<\omega}$ to denote the *finite* powerset functor—that is, given a set X , we write $\mathcal{P}_{<\omega}(X)$ to denote the set of all finite subsets of X —and LIST to denote the list functor—that is, given a set X , we write $\text{LIST}(X)$ to denote the set of all finite lists of X . We will use the concept of *F-algebra* and *F-coalgebra* (for a functor F) in various parts of this paper, but they are just functions.¹ The nomenclature of category theory enables us to apply well-established techniques and constructions. For a general introduction to coalgebra, see, for example, Rutten (2000) and Jacobs (2017).

We use basic co-induction in some definitions and proofs and, accordingly, give a brief explanation of how it works presently; for a more detailed discussion, see, for example, Barwise and Moss (1996) and Rutten (2000). We have to use co-induction (as opposed to induction) because reduction spaces are co-recursively generated and so may be infinite objects. We follow Reitzig (2012) for a succinct explanation of co-induction as the dual of induction.

In an *inductive* definition of a set, one begins with some *base* elements and applies *constructors* that say how one builds new elements of the set from existing ones. For example, a set of finite sequences, or ‘strings’, \mathbb{S} may be defined by the following:

$$\begin{cases} \varepsilon \in \mathbb{S} & \text{Base Case} \\ w \in \mathbb{S} \implies aw \in \mathbb{S} & \text{Condition 1} \\ aw \in \mathbb{S} \implies baw \in \mathbb{S} & \text{Condition 2} \end{cases}$$

The set \mathbb{S} inductively defined in this way is the *smallest* set satisfying these conditions. Hence, as defined above, \mathbb{S} is the set of strings of a and b with no two subsequent bs . This is a perfectly adequate account of induction, but one can give a more subtle treatment that explains precisely how it works in terms of fixed points.

¹ *F*-algebras generalize the notion of algebraic structure. If \mathcal{C} is a category and $F : \mathcal{C} \rightarrow \mathcal{C}$ is an endofunctor, then an *F*-algebra is a pair (A, α) , where A is an object of \mathcal{C} and $\alpha : F(A) \rightarrow A$ is a morphism in \mathcal{C} . A is called the carrier of the algebra. If we take $\alpha : A \rightarrow F(A)$, then we have an *F*-coalgebra. For endofunctors F on the category of sets, *F*-algebras are functions on the carrier set and *F*-coalgebras are countably infinite streams (sequences) over the carrier set. Often, we refer to the maps α as (co)algebras.

The definition above determines the following map f on sets of strings,

$$f(X) := X \cup \{\varepsilon\} \cup \{aw | w \in X\} \cup \{baw | aw \in X\}$$

Observe that f is monotone and the sets of strings form a complete lattice under set inclusion. Hence, the Knaster-Tarski Theorem (Knaster and Tarski 1928; Tarski 1955) tells us that the set of fixed points for f forms a complete lattice. The defined set \mathbb{S} is the *least* fixed point; that is, the smallest set X such that $f(X) = X$. Since f is continuous, \mathbb{S} is the colimit of the following ω -chain,

$$\emptyset \longrightarrow \underbrace{\{\varepsilon\} \cup \{aw\}}_{\emptyset \cup f(\emptyset)} \longrightarrow \underbrace{\{\varepsilon, a\} \cup f(\{\varepsilon, a\})}_{\emptyset \cup f(\emptyset) \cup f(f(\emptyset))} \longrightarrow \dots \longrightarrow \underbrace{\mathbb{S}}_{\bigcup_n f^n(\emptyset)}$$

Here each arrow denotes an application of f and \mathbb{S} denotes the colimit of the chain.

For a *co-inductive* definition, we use *destructors* rather than constructors. Intuitively, we reverse the implications in the inductive definition; observe that in doing so the ‘base case’ becomes vacuous. That is, rather than say ‘if w is included, then so is aw ’, we say ‘if aw is included, then so was w ’. This determines a set \mathbb{S}' with the following property: we can take arbitrarily long prefixes away from any word in \mathbb{S}' and remain in \mathbb{S}' . Observe that this property does not hold for finite strings, hence the co-inductive definition requires us to consider *infinite* strings.

While induction determines the least fixed point satisfying f , the co-inductive definition determines the (unique) *greatest* fixed point—that is, the *largest* set X such that $f(X) = X$. Thus \mathbb{S}' contains all *infinite* strings of a and b with no consecutive occurrences of bs . In other words, it is the limit of any ω^{op} -chain,

$$1 \longleftarrow f(1) \longleftarrow f(f(1)) \longleftarrow f(f(f(1))) \longleftarrow \dots$$

in which 1 is a singleton (i.e., a terminal object in the category), and the arrows denote the duals of the arrows in the least fixed point case (i.e., the unique morphisms that remove a prefix from a string). This is quite intuitive: we cannot construct \mathbb{S}' by an iterative process on elements as its elements are infinite, so instead we determine it by starting with all strings and removing those that do not satisfy the conditions after an arbitrary amount of time.

This concludes the background on co-induction.

2 The Semantics of Reductive Logic

In this section, we motivate the idea of reductive logic and explicate its current semantics. We set this discussion in the

historical development of *intuitionistic logic* (IL) (van Atten 2022) as it is here that the ideas are well known and are the most developed. Subsequently, in Sect. 3 and Sect. 4, we address the semantics of reduction operators, the things that generate reductions.

In logic, we have structures called *arguments* that represent evidence for a consequence of a logic—for example, natural deduction arguments after Szabo (1969), after Fitch (1952), after Lemmon (1978), and so on. When these arguments are *valid* they represent a *proof* of some statement. The semantics of proofs for particular logics have been very substantially developed for particular logics. Perhaps the best examples is the BHK interpretation of *intuitionism*.

Intuitionism, as defined by Brouwer (1913), is the view that an argument is valid when it provides sufficient evidence for its conclusion. This is IL. Famously, as a consequence, IL rejects *the law of the excluded middle*—that is, the meta-theoretic statement that either a statement or its negation is valid. This law is equivalent to the principle that, in order to prove a proposition, it suffices to show that its negation is contradictory. In IL, such an argument does not constitute sufficient evidence for its conclusion.

Heyting (1989) and Kolmogorov (1932) provided a semantics for intuitionistic proof that captures the evidential character of intuitionism, called the Brouwer-Heyting-Kolmogorov (BHK) interpretation of IL. It is now the standard explanation of the logic—see, for example, van Atten (2022).

The *propositions-as-types* correspondence—see Howard, Barendregt, and others (Howard 1980; Barendregt et al. 2013; Barendregt 1991, 1993)—gives a standard way of instantiating the denotations of proofs in the BHK interpretation of *intuitionistic propositional logic* (IPL) as terms in the simply-typed λ -calculus. Technically, the set-up can be sketched as follows: a judgement that \mathcal{D} is an NJ-proof of the sequent $\varphi_1, \dots, \varphi_k \triangleright \varphi$ corresponds to a typing judgement

$$x_1 : A_1, \dots, x_k : A_k \vdash M(x_1, \dots, x_k) : A$$

where the A_i s are types corresponding to the φ_i s, the x_i s correspond to placeholders for proofs of the φ_i s, the λ -term $M(x_1, \dots, x_k)$ corresponds to \mathcal{D} , and the type A corresponds to φ .

Lambek (1980) gave a more abstract account by showing that the simply-typed λ -calculus is the internal language of *cartesian closed categories* (CCCs), thereby giving a categorical semantics of proofs for IPL. In this set-up, a morphism

$$[\varphi_1] \times \dots \times [\varphi_k] \xrightarrow{[\mathcal{D}]} [\varphi]$$

in a CCC, where \times denotes cartesian product, that interprets the NJ-proof \mathcal{D} of $\varphi_1, \dots, \varphi_k \triangleright \varphi$ also interprets the term M ,

Fig. 1 Curry–Howard–Lambek correspondence

$$\begin{array}{ccc}
 x : \Gamma \vdash M(x) : \varphi & \longleftrightarrow & \llbracket \Gamma \rrbracket \xrightarrow{\llbracket \mathcal{D} \rrbracket} \llbracket \varphi \rrbracket \\
 & \searrow \quad \swarrow & \\
 & \mathcal{D} \Rightarrow \Gamma \triangleright \varphi &
 \end{array}$$

Fig. 2 Constructions-as-realizers-as-arrows correspondence

$$\begin{array}{ccc}
 \llbracket \Gamma \rrbracket \vdash \llbracket \Phi \rrbracket : \llbracket \varphi \rrbracket & \longleftrightarrow & \llbracket \Gamma \rrbracket \xrightarrow{\llbracket \Phi \rrbracket} \llbracket \varphi \rrbracket \\
 & \searrow \quad \swarrow & \\
 & \Phi \Rightarrow \Gamma \triangleright \varphi &
 \end{array}$$

where the $\llbracket \varphi_i \rrbracket$ s interpret also the A_i s and $\llbracket \varphi \rrbracket$ also interprets A .

Altogether, this describes the *Curry–Howard–Lambek* correspondence for IPL. It may be summarized by Fig. 1 in which:

- $\mathcal{D} \Rightarrow \Gamma \triangleright \varphi$ denotes that \mathcal{D} is an NJ-derivation of φ from Γ ;
- $x : \Gamma \vdash M(x) : A_\varphi$ denotes a typing judgment, as described above, corresponding to \mathcal{D} ; and,
- $\llbracket \Gamma \rrbracket \xrightarrow{\llbracket \mathcal{D} \rrbracket} \llbracket \varphi \rrbracket$ denotes that $\llbracket \mathcal{D} \rrbracket$ is a morphism from $\llbracket \Gamma \rrbracket$ to $\llbracket \varphi \rrbracket$ in a CCC.

To generalize to full IL (and beyond), Seely (1983) modified this categorical set-up and introduced *hyperdoctrines*—indexed categories of CCCs with coproducts over a base with finite products. Martin-Löf (1975) gave a formulae-as-types correspondence for predicate logic using dependent type theory. Barendregt (1991) gave a systematic treatment of type systems and the propositions-as-types correspondence. A categorical treatment of dependent types came with Cartmell (1978)—see also, for examples among many, work by Streicher (1988), Pavlović (1990), Jacobs (1991), and Hofmann (1997). In total, this gives a semantic account of *proof* for first- and higher-order predicate intuitionistic logic based on the BHK interpretation.

This is all very well for explaining what is a proof in IL. In reductive logic, however, the space of objects considered contains also things that are not proofs and cannot be continued to form proofs. Pym and Ritter (2004) have provided a general semantics of reductive logic in the context of classical and intuitionistic logic through polynomial categories; that is, by extending the categories in which arrows denote proofs for a logic by additional arrows that supply ‘proofs’ for propositions that do not have proofs but appear during reduction. This generalization of the BHK interpretation is known as the *constructions-as-realizers-as-arrows correspondence* in Fig. 2:

- $\Phi \Rightarrow \Gamma \triangleright \varphi$ denotes that Φ is a sequence of reductions for the sequent $\Gamma \triangleright \varphi$;
- $\llbracket \Gamma \rrbracket \vdash \llbracket \Phi \rrbracket : \llbracket \varphi \rrbracket$ denotes that $\llbracket \Phi \rrbracket$ is a *realizer* of $\llbracket \varphi \rrbracket$ with respect to the assumptions $\llbracket \Gamma \rrbracket$; and,
- $\llbracket \Gamma \rrbracket \xrightarrow{\llbracket \Phi \rrbracket} \llbracket \varphi \rrbracket$ denotes that $\llbracket \Phi \rrbracket$ is a morphism from $\llbracket \Gamma \rrbracket$ to $\llbracket \varphi \rrbracket$ in the appropriate polynomial category.

They also defined a judgement $w \Vdash_\Theta (\Phi : \varphi) \Gamma$ which says that w is a world witnessing that Φ is a reduction of φ to Γ , relative to the indeterminates of Θ .

What has just been presented gives a semantics of *reductions* that explicates how they relate to proofs and certify the validity of sequents. In this paper, we turn to the semantics of the reduction operators themselves; that is, the things which generate the reductions. Unlike rules, which specify particular transformations, reduction operators contain a certain degree of non-determinism rendering the subject more subtle.

3 Reduction Operators I—Denotational Semantics

In this section, we give a *denotational* semantics of reduction operators beginning from first principles. We then use this semantics to express the space of reductions.

3.1 Denotational Semantics

As described in Sect. 1, reductive logic is concerned with ‘backwards’ inference,

$$\frac{\text{Sufficient Premiss}_1 \quad \dots \quad \text{Sufficient Premiss}_k}{\text{Putative Conclusion}} \Uparrow \quad (\dagger)$$

This framework not only describes reduction in a formal system but also applies to reductive reasoning in more general settings.

Example 1 (see Milner 1984) Two friends, Alice and Bob, desire to meet before noon on Saturday. We shall model how they may reason about accomplishing this goal.

Let G be the statement that Alice and Bob did meet before noon on Saturday. For G to hold, Alice and Bob reason that they need to arrive at Waterloo Station before noon on Saturday. That is, it suffices for G_1 and G_2 to hold:

- G_1 : Alice arrives under the clock at Waterloo Station before noon on Saturday
- G_2 : Bob arrives under the clock at Waterloo Station before noon on Saturday

This represents a reduction of the putative conclusion G to the sufficient premisses G_1 and G_2 .

The precise meaning of (\dagger) has been formalized by Pym and Ritter (2004), though the earliest account is perhaps due to Kleene (1961). Fixing a logic \mathcal{L} , the ideas are as follows:

- R1** We assume that our logic \mathcal{L} comes with a proof system. Each inference rule denotes a *reduction operator* mapping assertions in \mathcal{L} to lists of assertions in \mathcal{L} .
- R2** The putative conclusion is an assertion in our logic \mathcal{L} , called the *goal*.
- R3** The assertions that must be proved in order to have a proof system of the initial assertion, called *subgoals*, are given by the corresponding instance of the sufficient premisses of a reduction operator.

While **R2** and **R3** merely specify the relationship between the terms involved, **R1** does more than needed. In practice, it is clearest to begin with a proof system and read the inference rules backwards to generate an adequate set of reduction operators, but this is inessential and limiting for defining reductive logic for a given logic. Pym and Ritter (2004) acknowledge that much of their analysis can be recovered purely semantically. Indeed, Gheorghiu and Pym (2023b, c) have shown that reductive logic bridges semantics and proof theory. For example, in Gheorghiu and Pym (2023c), it is shown for the bunched logic BI (and so, implicitly, for intuitionistic propositional logic also) that a bisimulation between the transition system generated by reductive proof and the one generated by the satisfaction relation of its relational model-theoretic semantics can be used to establish completeness.

The point of **R1** is twofold: it ensures the *validity* of the reduction operators, and it ensures their correctness—that is, if the putative conclusion is indeed valid, then it will eventually be reduced to trivial premisses. Thus, it says more than what (\dagger) intends. Presently, let us drop **R1** in place of the following:

R1' We have a finite set of reduction operators **RED** for assertions in the language of L , which map assertions to sets of assertions

The assumption that the set is finite enables the mathematical treatment and is quite representative as proof systems typically have finitely many rule schemas.

Having made this simplification in the overhead, we no longer need to fix a logic \mathcal{L} but only an assertion language over which reduction is taking place. The validity of reductions in this setup may then be determined by their correctness relative to a logical system, as in the standard practice of semantics.

With heuristics **R1'**, **R2**, and **R3** in place, we can give a precise, mathematical account of (\dagger) .

Let **GOALS** denote the set of all assertions in the language \mathcal{L} . The heuristics above suggest that a reduction operator is a partial function from goals to finite sets of goals:

$$\rho : \text{GOALS} \rightarrow \mathcal{P}_{<\omega}(\text{GOALS})$$

Example 2 (Example 1 cont'd) Alice's and Bob's reasoning is modelled by the following reduction,

$$\rho : G \mapsto \{G_1, G_2\}$$

Such reduction operators describe one step of reduction. Reasoning reductively means chaining together such steps, reducing the subgoals to further subgoals. This introduces the idea of *state* as the set of goals to be reduced.

Initially, the state is a singleton consisting of a single goal, the putative conclusion. Subsequently, it consists of a finite set of (sub)goals generated so far. Mathematically, a reduction operator ρ moves a state S to a next state S' by replacing one of the goals with a list of subgoals,

$$S = \{G_1, \dots, G_n\} \xrightarrow{\rho} (\{G_1, \dots, G_n\} \setminus \{G_i\}) \cup \rho(G_i) = S'$$

Observe that a sequence of reductions may go on infinitely or terminate. When it terminates, one of the following two things can happen:

- T1** There are no further subgoals—that is, one has reduced all the goals to the empty set
- T2** The remaining subgoals are irreducible—that is, no reduction operators are applicable to the present set of goals

Example 3 (Infinite Reduction) Consider the following reduction operator:

$$\rho : (\varphi \supset \psi, \Gamma \triangleright \chi) \mapsto \{(\varphi \supset \psi, \Gamma \triangleright \varphi), (\psi, \varphi \supset \psi, \Gamma \triangleright \chi)\}$$

Observe that it essentially corresponds to \supset -elimination rule in LJ (Szabo 1969). Applying it to the goal $(p \supset p \triangleright p)$ yields the subgoals $\{(p \supset p \triangleright p), (p, p \supset p \triangleright p)\}$. Applying ρ to either subgoal does not change the list of subgoals. Hence, there is an infinite sequence of reductions,

$$\begin{aligned} \{(p \supset p \triangleright p)\} &\xrightarrow{\rho} \{(p \supset p \triangleright p), (p, p \supset p \triangleright p)\} \\ &\xrightarrow{\rho} \{(p \supset p \triangleright p), (p, p \supset p \triangleright p)\} \\ &\xrightarrow{\rho} \dots \end{aligned}$$

This is a typical example of the kind of computations encountered in co-inductive logic programming—see, for example, Gupta et al. (2007), Simon et al. (2007). Note that one can construct a more pathological example than the one given by not only having the computation never terminate, but also having the set of goals increase with each reduction.

Example 4 (Termination **T1**) Let ρ_1 be ρ in Example 3 and let ρ_2 be as follows for any formula φ :

$$\rho_2 : (\varphi, \Gamma \triangleright \varphi) \mapsto \emptyset$$

Beginning with the goal $(\varphi, \varphi \supset \psi \triangleright \psi)$, we have the following sequence:

$$\begin{aligned} \{(\varphi, \varphi \supset \psi \triangleright \psi)\} &\xrightarrow{\rho_1} \{(\varphi, \varphi \supset \psi \triangleright \varphi), (\psi, \varphi, \varphi \supset \psi \triangleright \psi)\} \\ &\xrightarrow{\rho_2} \{(\psi, \varphi, \varphi \supset \psi \triangleright \psi)\} \\ &\xrightarrow{\rho_2} \emptyset \end{aligned}$$

This shows a sequence terminating with **T1**.

Example 5 (Termination **T2**) Let ρ_1 and ρ_2 be as in Example 4. When $\varphi \neq \psi$, the goal $(\varphi \triangleright \psi)$ is irreducible using the reduction operators $\{\rho_1, \rho_2\}$. Therefore, it represents a terminating sequence (of one step) of type **T2** (not **T1**).

Intuitively, **T1** means that the putative goal is indeed valid. This determines an important soundness and completeness condition for the semantics of reductive logic. Fix a logic \mathcal{L} :

(Soundness) \mathbb{RED} is *sound* if, for any goal $G \in \mathbf{GOALS}$, if the state $\{G\}$ has a sequence of reductions ending in the empty set, then G is valid in \mathcal{L} .

(Completeness) \mathbb{RED} is *complete* if, for any goal $G \in \mathbf{GOALS}$, if G is valid in \mathcal{L} , then the state $\{G\}$ has a sequence of reductions ending in the empty set.

Observe that beginning with a goal G and witnessing an infinite sequence of reduction or a finite one terminating with **T2** does not necessarily mean G is *not* valid, only that

we have not been able to witness that it is. For these cases to show that the putative goal is invalid, we should need a meta-theorem that says that the sequence of reductions used is itself complete, which is stronger than what is prescribed above. Of course, this relates to the celebrated *negation-as-failure* protocol—see, for example, Clark (1977).

We regard these conditions as criteria for a semantics of reductive logic. Observe that **R1** essentially guarantees that these conditions hold. However, in doing so, it narrows the scope of reduction operators to the point where asking about their validity becomes nonsensical. We return to this in Sect. 4.

3.2 The Space of Reductions

While we have modelled ‘reduction’ through the use of operators and said something about how they compose to form reductive reasoning, we can be much more precise about the space of exploration during reductive reasoning. This space, the space of reductions, is larger than the space of deductions—that is, the space of objects studied in deductive logic; for example, the class of proofs in a fixed proof-system—as it also contains ‘failed’ or ‘incomplete’ reductions—that is, reductions failing to satisfy condition **T1**.

To model the space of reductions for a given goal $G \in \mathbf{GOALS}$, we move from considering individual reduction operators to considering *all* possible reductions as different sets of sets of sufficient premisses. Let \mathbb{RED} denote the set of reduction operators given in **R1'**. We define a single *destructor* for \mathbb{RED} representing the collected action of all the reduction operators that apply at any point,

$$\partial : \mathbf{GOALS} \rightarrow \mathcal{P}_{<\omega}(\mathcal{P}_{<\omega}(\mathbf{GOALS}))$$

with the action

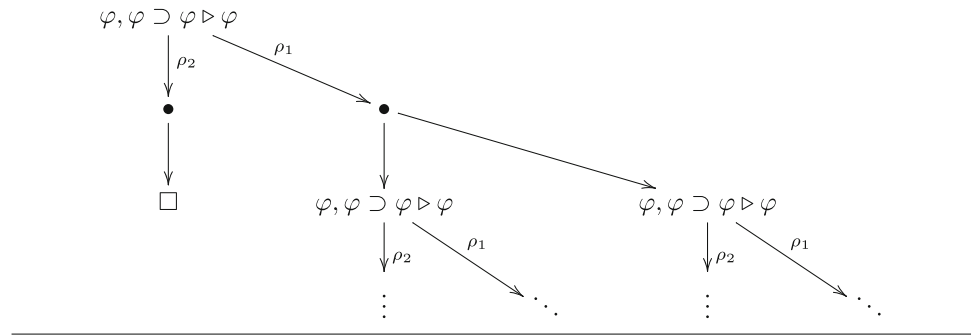
$$\partial : G \mapsto \{\rho(G) \mid \rho \in \mathbb{RED}\}$$

— recall that \mathbb{RED} is finite and each $\rho \in \mathbb{RED}$ is a $\mathcal{P}_{<\omega}$ -coalgebra. Using ∂ , one can define the space of reductions for a goal G as a limit of ∂ applied to a given goal:

Definition 6 (Reduction Space) The reduction space $\mathfrak{R}(G)$ for a goal $G \in \mathbf{GOALS}$ is the the greatest tree satisfying the following:

- the root of the tree is a node labelled by G ,
- the root has $|\partial(G)|$ children labelled \bullet ,
- for each \bullet , there exists a unique set $\{G_0, \dots, G_n\} \in \partial(G)$
- if $\{G_0, \dots, G_n\} \in \partial(G)$ corresponds to \bullet , then the children of \bullet are $\mathfrak{R}(G_1), \dots, \mathfrak{R}(G_n)$.

Fig. 3 Reduction space for $\varphi, \varphi \supset \varphi \triangleright \varphi$



Example 7 In Fig. 3, we show the reduction space for $\varphi, \varphi \supset \varphi \triangleright \varphi$. For readability, we have labelled the branches using the generating reduction operators. Here, ρ_1 corresponds to the reduction operator in Example 3 and ρ_2 corresponds to the reduction operator in Example 4. Observe that it is an infinite tree along the rightmost branch.

This is the co-inductive version of an *and/or*-tree as used in *logic programming* (LP)—see, for example, Komendantskaya et al. (2011), Komendantskaya and Power (2011), Komendantskaya et al. (2016), Bonchi and Zanasi (2015). This nomenclature arises from the fact that the \bullet denotes a choice, representing the disjunction of sets of sufficient premisses, and the children of any given or-node must all be redexes, representing a conjunction of sufficient premisses. This suggests the following representation of a reduction sequence:

Definition 8 (Reduction Tree) A reduction tree for G is any tree \mathcal{R} such that G is the root of \mathcal{A} and the immediate subtrees $\mathcal{A}_1 \dots \mathcal{A}_n$ are the reduction trees for $G_1, \dots, G_n \in \text{GOALS}$, respectively, where there is $\rho \in \text{RED}$ and $\rho(G) = \{G_1, \dots, G_n\}$.

Example 9 In Fig. 4, we present the reduction sequence in Example 4 as a reduction tree in which \square denotes a node with no sequent.

One may view the reduction space as the glueing together of all the possible reduction trees using \bullet to denote different choices. Conversely, one may extract them by skipping over \bullet -nodes: beginning from the goal, choose *one* \bullet and connect the goal to all the children of that node, then repeat the process for these subgoals—see Gheorghiu et al. (2023a).

Following work by Komendantskaya et al. (Komendantskaya et al. 2011; Komendantskaya and Power 2011; Komendantskaya et al. 2016) (see also Bonchi and Zanasi 2015 and Gheorghiu et al. (2023a)), we can model the space of reductions as the cofree comonad \mathcal{C} of $\mathcal{P}_{<\omega}(\mathcal{P}_{<\omega}(-))$ applied to GOALS .² That is, for a given X , consider the

following sequence of sets, where α is any ordinal:

$$\begin{cases} Y_0 := X \\ Y_{\alpha+1} := \text{GOALS} \times \mathcal{P}_{<\omega}(\mathcal{P}_{<\omega}(Y_\alpha)) \end{cases}$$

Each stage of the construction yields a function $\partial_\alpha : X \rightarrow Y_\alpha$ defined inductively as follows, where I is the identify function:

$$\begin{cases} \partial_0 := I \\ \partial_{\alpha+1} := I \times (\mathcal{P}_{<\omega}(\mathcal{P}_{<\omega}(\partial_\alpha)) \circ \partial) \end{cases}$$

This construction converges at the limit ordinal 2ω —see Worrell (2005). We denote the resulting map $\nabla : X \rightarrow \mathcal{C}(X)$. This is a standard construction for a cofree comonad. For $X = \text{GOALS}$, this precisely maps goal $G \in \text{GOALS}$ to its space of reductions:

Theorem 10 For any $G \in \text{GOALS}$,

$$\nabla(G) = \mathfrak{R}(G)$$

Proof We proceed by co-induction on the definition of $\mathfrak{R}(G)$ (Definition 6). For a proof by co-induction, we use the uniqueness of the greatest fixed point (see Sect. 1). That is, we need only show that applying the destructor defining $\mathfrak{R}(G)$ is

$1_{\mathcal{C}}$ denotes the identity functor on \mathcal{C} , and $\mu : T^2 \rightarrow T$ ('multiplication'), where T^2 is the functor $T \circ T : \mathcal{C} \rightarrow \mathcal{C}$ subject to the following laws:

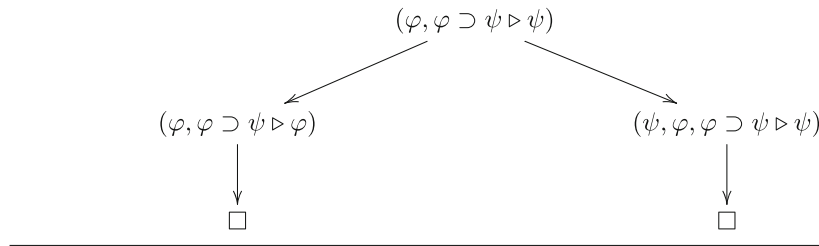
- $\mu \circ T\mu = \mu \circ \mu T : T^3 \rightarrow T$, where $T\mu$ and μT are formed by the induced 'horizontal composition' of natural transformations (see Saunders MacLane 1971)
- $\mu \circ T\eta = \mu \circ \eta T = 1_T : T \rightarrow T$, where 1_T denotes the identity natural transformation from T to T .

A *comonad* is categorical dual on the category \mathcal{C}^{op} , with axioms for counit and comultiplication, obtained by reversing the arrows in the definition above.

Generally speaking, 'free' means an instance of a structure that satisfies just the axioms of its definition, with no additional structure imposed. A *free functor* arises as a left adjoint to a forgetful functor and a *cofree functor* arises as a right adjoint to a forgetful functor. For example, a left adjoint to the forgetful functor from a given category \mathcal{C} of algebraic structures to the category of sets takes sets X to their corresponding free objects in \mathcal{C} . See nLab (2021) for an explanation and examples.

² A *monad* on a category \mathcal{C} is given by an endofunctor $T : \mathcal{C} \rightarrow \mathcal{C}$ together with two natural transformations $\eta : 1_{\mathcal{C}} \rightarrow T$ ('unit'), where

Fig. 4 A Reduction tree for $(\varphi, \varphi \supset \psi \triangleright \psi)$



equivalent to applying the destructor defining $\nabla(G)$, assuming that the results of such applications are equivalent. This is made clear presently.

Following Definition 6, we think of each \bullet as a set of possible subgoals; that is, the elements of ∂ when it has been applied to a goal. In this reading, encoding trees a lists, Definition 6 is expressed as follows:

$$\begin{aligned} \mathfrak{R}(G) &:= \langle G, \{\{\mathfrak{R}(G_1), \dots, \mathfrak{R}(G_n)\} \mid \exists \rho \in \text{RED} : \rho(G) = \{G_1, \dots, G_n\}\} \rangle \\ &= \{G_1, \dots, G_n\} \end{aligned} \quad (*)$$

This expresses \mathfrak{R} as a *destructor* for the reduction spaces.

Let $\text{SUB}(G) := \cup_{\rho \in \text{RED}} \rho(G)$. The co-induction hypothesis is the following:

$$\text{If } G' \in \text{SUB}(G), \text{ then } \nabla(G') = \mathfrak{R}(G') \quad (\text{co-IH})$$

The result follows immediately:

$$\begin{aligned} \mathfrak{R}(G) &= \langle G, \{\{\mathfrak{R}(G_1), \dots, \mathfrak{R}(G_n)\} \mid \exists \rho \in \text{RED} : \rho(G) = \{G_1, \dots, G_n\}\} \rangle \quad (*) \\ &= \{G_1, \dots, G_n\} \\ &= \langle G, \{\{\nabla(G_1), \dots, \nabla(G_n)\} \mid \exists \rho \in \text{RED} : \rho(G) = \{G_1, \dots, G_n\}\} \rangle \quad (\text{co-IH}) \\ &= \{G_1, \dots, G_n\} \\ &= \nabla(G) \end{aligned}$$

This completes the proof. \square

This completes the construction of the space of reductions.

3.3 Soundness and Completeness

Let L be a deductive system over GOALS . We elide details of L as they do not matter; concrete examples include natural deduction systems, sequent calculi, tableaux systems, and so on. Fix a logic \mathcal{L} . We have the following soundness and completeness conditions:

(Soundness) L is *sound* for \mathcal{L} if, for any goal $G \in \text{GOALS}$, if G admits a proof in L , then G is valid in \mathcal{L} .

(Completeness) L is *complete* for \mathcal{L} if, for any goal $G \in \text{GOALS}$, if G is valid in \mathcal{L} , then G admits a proof in L .

For the remainder of this section, we take a fixed L that is sound and complete for a fixed logic \mathcal{L} .

We can define the correctness of RED relative to a calculus L as follows:

(Faithfulness) RED is *faithful* for L if, for any goal $G \in \text{GOALS}$ and any $\rho \in \text{RED}$, if $\rho(G) = \{G_1, \dots, G_n\}$, then there is a rule in L from G_1, \dots, G_n to G .

(Adequacy) RED is *adequate* for L if, for any goal $G \in \text{GOALS}$, if there is a rule from G_1, \dots, G_n to G in L , then there is a reduction operator $\rho \in \text{RED}$ such that $\rho(G) = \{G_1, \dots, G_n\}$.

Relative to this set-up, the validity of reduction we automatically have soundness and completeness of RED with respect to \mathcal{L} :

Proposition 11 (Soundness & Completeness) *The following both hold:*

- if RED is faithful for L , then RED is sound for \mathcal{L}
- If RED is adequate for L , then RED is complete for \mathcal{L} .

Proof Immediate by induction since reduction operators in RED are instances of the inverses of rules in L . \square

This broadly represents the account of reductive logic used in computer-assisted formal reasoning. A central example is *Automated Theorem Proving* (ATP), which leverages algorithms to emulate the logical reasoning process that human mathematicians use, representing reasoning in many formal systems such as sequent calculi. This includes, for example, proof by induction, case analysis, and so on. It is the heart of automated reasoning tools (see, for example, Portoraro (2024)), including *logic programming* (see, for example, Miller 1989; Miller et al. 1991), and symbolic AI (see, for example, Bundy 1985 and Gillies 1996). Of course, ATP applies a *range* of techniques including, *inter alia*, analytic tableaux (Smullyan 1968), resolution (Kowalski 1979), the matrix method, Wallen (1987), and uniform/focused proof-search (Miller 1989; Miller et al. 1991). In particular, the latter has received much interest because it enables a uniform platform for proof construction as well as other reasoning methodologies (e.g., SAT/SMT solving algorithms such as DPLL Farooque et al. 2013).

Despite its efficacy in the above problem domain, this account of reductive logic is somewhat limited in that, essentially, it demands that reduction operators be exactly the duals of admissible rules, recovering R1. What is more, it speaks of the correctness of the collection of reduction operators, but not of the validity of reduction operators themselves.

In the next section, we give an alternative semantics for reduction operators in which the assertion language of goals need not even be the assertion language of the logic, as long as one has some sense of correspondence between them.

4 Reduction Operators II—Operational Semantics

In Sect. 3, we provided a denotational semantics for reduction operators from first principles. This enables the representation of reductive reasoning. In this section, we present a corresponding *operational* semantics, allowing for a more refined analysis that provides a subtle and expressive account of reduction-operator usage.

4.1 Operational Semantics

In this section, we introduce *structural operational semantics* (SOS) for reduction and reduction operators. Originally introduced by Plotkin (1981) to define the behaviour of programming constructs, SOS uses formal rules to describe how each part of a program manipulates a state to yield a new state. It adopts a meaning-as-use approach, where ‘use’ implies ‘behaviour’.

For our purposes, we focus on one of the most extensively studied rule formats for SOS: the format of *general structural operational semantics* (GSOS) as defined by Bloom et al. (1995).

We represent a state as a (finite) list of current goals, where $G \in \text{GOALS}$,

$$S ::= \square | G :: S$$

Intuitively, \square denote the empty string. We may write $[G_1, \dots, G_n]$ to denote the list $G_1 :: (\dots :: G_n)$.

We have a countable collection of action labels A and a reduction operator ρ is the transformation of a goal G into a list of subgoals $[G_1, \dots, G_n]$,

$$G \xrightarrow{\rho} [G_1, \dots, G_n] \quad (\uparrow \rho)$$

The set of reduction operators RED from Definition R1’ can be seen as comprising all such rules. Notably, while Sect. 3

used finite sets, we now work with finite strings due to syntactic considerations.

We model the free choice of goal reduction by adding the following for all $\rho \in A$:

$$\frac{G \xrightarrow{\rho} [G_1, \dots, G_n]}{G \longrightarrow [G_1, \dots, G_n]} \quad (\uparrow \rho)$$

This behaviour extends to states with more than one element via the following rules:

$$\frac{S_1 \longrightarrow S'_1 \quad S_2 \longrightarrow S'_2}{S_1 :: S_2 \longrightarrow S'_1 :: S'_2} \quad (\uparrow)$$

The operational semantics of reduction comprises the collection of all such rules,

$$\mathcal{O} := \text{RED} \cup \{\uparrow\} \cup \{\uparrow \rho \mid \rho \in A\}$$

This system enables us to describe the process of reduction as a ‘computation’ in \mathcal{O} .

Example 12 The reduction sequence in Example 4 is presented in \mathcal{O} as follows:

$$\frac{(\varphi, \varphi \supset \psi \triangleright \psi) \xrightarrow{\rho_1} (\varphi, \varphi \supset \psi \triangleright \varphi) :: (\psi, \varphi, \varphi \supset \psi \triangleright \psi)}{(\varphi, \varphi \supset \psi \triangleright \psi) \longrightarrow (\varphi, \varphi \supset \psi \triangleright \varphi) :: (\psi, \varphi, \varphi \supset \psi \triangleright \psi)}$$

and

$$\frac{(\varphi, \varphi \supset \psi \triangleright \varphi) \xrightarrow{\rho_2} \square \quad (\psi, \varphi, \varphi \supset \psi \triangleright \psi) \xrightarrow{\rho_2} \square}{(\varphi, \varphi \supset \psi \triangleright \varphi) \longrightarrow \square \quad (\psi, \varphi, \varphi \supset \psi \triangleright \psi) \longrightarrow \square} \quad (\varphi, \varphi \supset \psi \triangleright \varphi) :: (\psi, \varphi, \varphi \supset \psi \triangleright \psi) \longrightarrow \square$$

We discuss combining these to represent a sequence of reductions below.

Following Turi and Plotkin (1997), this GSOS specification can be modelled as a *bialgebra*. Given two functors S and T , and a *distributive law* $\lambda : TS \Rightarrow ST$, a bialgebra is a pair of an T -algebra $\alpha : TX \rightarrow X$ and an S -coalgebra $\beta : X \rightarrow SX$ over a common carrier (i.e., X) satisfying the following *pentagonal law*:

$$\beta \circ \alpha = S\alpha \circ \lambda \circ T\beta$$

Turi and Plotkin (1997) observe that one can construct bialgebras from a GSOS such that the algebra is a denotational model, the coalgebra is an operation model, and the distributive law says that the combination of the two models obeys the rules of the GSOS.

This approach can be applied immediately to the account of reduction above. Our state space is modelled by $X = \text{LIST}(\text{GOALS})$. The heuristics are as follows:

- the internal structure of states is modelled by an algebra $\alpha : \text{LIST}(X) \rightarrow X$ —thus, by (\uparrow) , a state $S_1::S_2$ is regarded as the concatenation of the lists of states S_1 and S_2 ,
- the behaviour (i.e., reduction) is modelled by a coalgebra $\beta : X \rightarrow \mathcal{P}_{<\omega}(\text{LIST}(X))$; that is, by \uparrow_ρ , a state is taken to the set of next possible states,

$$S \mapsto \{S' \in \text{LIST}(\text{GOALS}) \mid \exists \rho \in A : S \xrightarrow{\rho} S'\}$$

- there is a distributive law $\lambda : \text{LIST}(\mathcal{P}_{<\omega}) \Rightarrow \mathcal{P}_{<\omega}(\text{LIST})$ ensuring coherence,

$$[X_1, \dots, X_n] \mapsto \{[x_1, \dots, x_n] \mid x_i \in X_i \text{ for } i = 1, \dots, n\}$$

In the behaviour functor for β —that is, $\mathcal{P}_{<\omega}(\text{LIST}(-))$ —the internal functor LIST represents the structure of states and is conjunctive, because all the goals are potential, while the external functor $\mathcal{P}_{<\omega}$ is disjunctive, because only one collection of sufficient premisses needs to be checked. This reading justifies the particular distributive law used. For more technical details, see Bonchi and Zanasi (2015).

This operational semantics can be easily extended to express *large* steps of reduction (rather than just small ones). For example, we take the reflexive and transitive closure:

$$S \longrightarrow S \quad \frac{S \longrightarrow S' \quad S' \longrightarrow S''}{S \longrightarrow S''}$$

Now, $G \longrightarrow \square$ means that there is some sequence of reduction terminating with an empty list of remaining subgoals (T1).

Example 13 In the large-step operational semantics, $(\varphi, \varphi \supset \psi \supset \psi) \longrightarrow \square$. This is witnessed by the computation

$$\frac{(\varphi, \varphi \supset \psi \supset \psi) \longrightarrow (\varphi, \varphi \supset \psi \supset \varphi) :: (\psi, \varphi, \varphi \supset \psi \supset \psi) \quad (\varphi, \varphi \supset \psi \supset \varphi) :: (\psi, \varphi, \varphi \supset \psi \supset \psi) \longrightarrow \square}{(\varphi, \varphi \supset \psi \supset \psi) \longrightarrow \square}$$

in which both the left and right branches continue exactly as in Example 12.

A ‘composition’ can be seen as any (partial) mapping that from a collection of reduction operators maps a new reduction operator. Formally, let \mathbb{R} denote the set of all reduction mappings: $\rho \in \mathbb{R}$ if $\rho : \text{GOALS} \rightarrow \text{LIST}(\text{GOALS})$. A composition is a (partial) mapping $\circ : \mathbb{R}^n \rightarrow \mathbb{R}$. Intuitively, \circ is valid if it preserves validity: for any $\rho_1, \dots, \rho_n \in \mathbb{R}$, if ρ_1, \dots, ρ_n are all valid, then $\circ(\rho_1, \dots, \rho_n)$ is valid.

This approach, as articulated by Milner (1984), allows us to demonstrate the validity of a proposed goal in terms of individual steps. As Milner (1984) notes:

‘Here it is a matter of taste whether the human prover wishes to see this performance done by the machine, in all its frequently repulsive detail, or wishes only to see the highlights, or is merely content to let the machine announce the result (a theorem!).’

Various kinds of composition can be given operational semantics. For example, we can extend the language of actions with symbols $;$, $+$, and $(-)^*$:

$$a:: = \rho \in A \mid a ; a \mid a + a \mid a^*$$

By adding appropriate rules to \mathcal{O} , we define $;$ as ‘sequential composition of reduction operators’, $+$ as ‘non-deterministic choice of reduction operators’, and $*$ as ‘arbitrary number of applications’ by the following rules:

$$\frac{S \xrightarrow{a} S' \quad S' \xrightarrow{b} S''}{S \xrightarrow{a;b} S''} \uparrow; \quad \frac{S \xrightarrow{a} S'}{S \xrightarrow{a+b} S'} \uparrow_+^1 \quad \frac{S \xrightarrow{b} S'}{S \xrightarrow{a+b} S'} \uparrow_+^2$$

$$\frac{S \xrightarrow{a} S' \quad S' \xrightarrow{a^*} S''}{S \xrightarrow{a^*} S''} \uparrow_*^1 \quad \frac{}{S \xrightarrow{a^*} S} \uparrow_*^2$$

This language is a relatively refined way of expressing sequences of reductions in terms of process algebra—see, for example, Milner (1983, 1989) and Bergstra and Klop (1984).

We have thus provided an operational semantics for reduction, but what determines its correctness with respect to a formal system? We can apply the soundness and completeness criteria from Sect. 3, suitably translated.

(Soundness) RED is *sound* for \mathcal{L} if, for any goal $G \in \text{GOALS}$, if $G \longrightarrow \square$, then G is valid in \mathcal{L} .

(Completeness) RED is *complete* for \mathcal{L} if, for any goal $G \in \text{GOALS}$, if G is valid in \mathcal{L} , then $G \longrightarrow \square$.

This operational semantics satisfactorily explains what reduction operators are. However, understanding what makes them *valid* is crucial. We can assert that reduction operators are valid when there exists a corresponding admissible rule for the logic in which we wish to reason, following the approach outlined in Sect. 3. In the next section, we refine this validity condition by relating it to the *theory of tactical proof*, as introduced by Milner (1984)

4.2 The Theory of Tactical Proof

A foundational framework that supports the mechanization of reductive logic is the theory of *tactical proof*. It systematically underpins the proof-assistants mentioned in Sect. 1. The centrality of tactical proof to computer assisted formal reasoning cannot be understated—see, for example, Gordon (2015). Its efficacy is that it allows users to articulate insight into reasoning methods and delegate routine but possibly error-prone work to machines. The theory clarifies how concepts such as ‘goal’, ‘strategy’, ‘achievement’, ‘failure’, etc., interrelate during reductive reasoning. In this section, we observe that tactical proof can be viewed as a semantics of reduction operators.

To simplify the operational semantics, we define a set of primary entities called goals (GOALS) and reduction operators given by mappings:

$$\rho : \text{GOALS} \rightarrow \text{LIST}(\text{GOALS})$$

Milner (1984) extends this set-up by introducing another set of primary entities called *events* (EVENTS) and an *achievement* relation α , which specifies how events witness goals:

$$\alpha \subseteq \text{GOALS} \times \text{EVENTS}$$

This allows us to express that certain events satisfy certain goals.

Example 14 Consider G from Example 1, asserting that Alice and Bob meet before noon on Saturday. Let E be the event of ‘Alice and Bob meeting under the clock at Waterloo station at 11:53 on Saturday.’ Here, E achieves G as it satisfies the designated description of the goal.

We now ask: *What makes a reduction operator invalid?* Intuitively, events that satisfy subgoals should establish events that satisfy the original goal. As noted by Milner (1984), in Example 1, if ‘noon’ were replaced by ‘evening’ in both G_1 and G_2 , then ρ should be deemed invalid. Therefore, a reduction operator is *valid* when any events achieving subgoals can be justified to suffice for an event achieving the goal. Milner (1984) refers to this justification as a *procedure*:

$$\pi : \text{LIST}(\text{EVENTS}) \rightarrow \text{EVENTS}$$

Note that, like reduction operators, procedures are partial functions.

Example 15 To complement G_1 and G_2 from Example 1, asserting that Alice and Bob arrive at Waterloo before noon on Saturday, consider the events:

- E_1 : Alice arrives at Waterloo Station at 11:57 on Saturday

- E_2 : Bob arrives at Waterloo Station at 11:53 on Saturday

Here, E_1 achieves G_1 and E_2 achieves G_2 . Moreover, the event E in Example 14 arises from E_1 and E_2 via the *wait* procedure, which requires Alice and Bob to wait for each other.

Milner (1984) defines a *tactic* as a reduction operator that also provides the procedure justifying the subgoals. This formalizes the validity condition described above. Let Π denote the set of procedures, a subset of all functions $\pi : \text{LIST}(\text{EVENTS}) \rightarrow \text{EVENTS}$.

Definition 16 (Tactic) A *tactic* is a partial mapping $\tau : \text{GOALS} \rightarrow \text{LIST}(\text{GOALS}) \times \Pi$.

Definition 17 (Validity for Tactics) A tactic τ is *valid* if, for any $G \in \text{GOALS}$, whenever $\tau : G \mapsto \langle [G_1, \dots, G_n], \pi \rangle$ and $E_1, \dots, E_n \in \text{EVENTS}$ are such that $G_i \alpha E_i$ for $i = 1 \dots n$, then $G \alpha \pi(E_1, \dots, E_n)$.

A tactic is a reduction operator that upon application to a goal produces a list of subgoals as well as a witness that the reduction is valid (i.e., the procedure). This provides a robust account of reductive reasoning widely applied in mechanized reasoning (see Gordon 2015). However, we require only that procedures *exist* to account of the validity of the reduction itself. Thus:

Definition 18 (Validity for Reduction Operators) A reduction operator ρ is *valid* if, for any $G \in \text{GOALS}$, whenever $\rho : G \mapsto [G_1, \dots, G_n]$ and $E_1, \dots, E_n \in \text{EVENTS}$ are such that $G_i \alpha E_i$ for $i = 1 \dots n$, then there exists $\pi \in \Pi$ such that $G \alpha \pi(E_1, \dots, E_n)$.

This formulation of validity provides a satisfactory account of reduction and reduction operators as it separates the action of ‘reduction’ from the validity of such an action.

Example 19 (Milner’s Proof-search Machine, Milner 1984) Let \vdash be a *Tarskian* (Tarski 1936) consequence relations—in particular, it is monotonic. The set-up is as follows:

- the set of *goals* is defined to be the set of sequents $\Gamma \triangleright \varphi$, in which Γ is a list of formulas and φ is a formula
- the set of *events* is defined to be the set of consequences—that is, sequents $\Delta \triangleright \psi$ such that $\Delta \vdash \psi$
- the *achievement* relation α is defined as follows:

$$(\Gamma \triangleright \varphi) \alpha (\Delta \triangleright \psi) \text{ if and only if } \varphi = \psi \text{ and } \Delta \sqsubseteq \Gamma,$$

where $\Delta \sqsubseteq \Gamma$ means that every element in Δ is also in Γ .

In this set-up, a tactic or reduction operator is *valid* if it has a procedure corresponding to an admissible rule for the logic.

Observe that the semantics here is one based on ‘proof’ rather than ‘truth’. That is, according the notion of validity thus provided, what makes a reduction (or sequence of reductions) valid is, ultimately, that it represents a valid forwards inference. This is not a trivial observation as it stands in contrast to the dominant notion of semantics within logic; namely, *model-theoretic semantics* (M-tS).

In M-tS, validity is set up using an abstract notion of *model* \mathfrak{M} , satisfying a range of conditions, together with a satisfaction relation \models between models and formulas. A formula φ is *true* in a model \mathfrak{M} if the model satisfies the formula; that is, $\mathfrak{M} \models \varphi$. A formula is *φ valid* iff it is true in all models. This is the paradigm in which one has, for example, the possible-worlds semantics for IL by Beth (1955) and Kripke (1965).

Following van Atten (2022), the paradigm of meaning based on ‘proof’ is known as *proof-theoretic semantics* (P-tS) (Schroeder-Heister 2018; Wansing 2000; Francez 2015). Importantly, here ‘proof’ does not mean proof in a fixed formal system, but rather constructions based on some *a priori* notion of inference. In the semantics for reductive logic above, this notion of inference is the notion of *procedure*. While P-tS has historically received little attention, most of which restricted to IL (cf. Schroeder-Heister 2018, 2006), the P-tS of various logical systems has recently seen rapid development; see, for example, Makinson (2014), Nascimento et al. (Nascimento et al. 2023; Stafford and Nascimento 2023; Nascimento 2023), Piecha et al. (Piecha et al. 2015; Piecha 2016; Piecha and Schroeder-Heister 2019), Pym et al. (Gheorghiu et al. 2023b; Gu et al. 2023; Gheorghiu et al. 2024; Eckhardt and Pym 2024; Pym et al. 2024), Sandqvist (Sandqvist 2005, 2009, 2015), Schroeder-Heister (Schroeder-Heister 2006, 2008, 1984), and Stafford (2021).

In Gheorghiu and Pym (2023a), P-tS is seen to give a semantics for tactical proof. In this paper, tactical proof arises as a special instance of a P-tS of reductive logic. The two accounts are mutually coherent. Accordingly, we observe that P-tS as a paradigm of meaning in logic is both historically and practically important for the understanding of logic as a mathematics of reasoning, though rarely expressed as such.

5 Control Régimes

While in the space of proofs, a rule takes a list of premisses and derives a single conclusion, navigating the space of reductions involves many choices that impact the validity of the resulting construction. We refer to all the strategic aspects of managing search in reductive logic as ‘control’.

Example 20 Consider the following introduction rule for a multiplicative conjunction $*$ as found in substructural logics such as *linear logic* (Girard 1987) and the *logic of bunched*

implications (O’Hearn and Pym 1999):

$$\frac{\Gamma_1 \triangleright \varphi_1 \quad \Gamma_2 \triangleright \varphi_2}{\Gamma_1, \Gamma_2 \triangleright \varphi_1 * \varphi_2}$$

It may be read as a reduction operator as follows:

$$\frac{\Gamma_1 \triangleright \varphi_1 \quad \Gamma_2 \triangleright \varphi_2}{\Gamma \triangleright \varphi_1 * \varphi_2} \quad (\Gamma = \Gamma_1, \Gamma_2) \quad \Uparrow$$

Here, it is necessary to divide Γ into two parts, Γ_1 and Γ_2 . The reduction operator doesn’t specify which splitting to use; thus, in applying it, one must *choose* a particular one.

The control problems may actually be summarized by just two kinds of choices:

- C1** The choices, which are mutually dependent, of which reduction operator to apply and to which premiss
- C2** The choice to *backtrack* - that is, to return to a previous point in the search where a choice was made and make a different choice

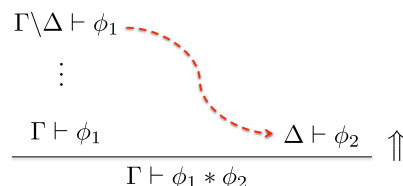
A *control régime* is a specification governing the use of the reduction operators; it is the policy to which we refer when considering these choices.

The application of a control régime delivers a *proof-search* algorithm. Instantiating Kowalski’s (Kowalski 1979) celebrated maxim, ‘*Algorithm = Logic + Control*’, in the context of reductive logic, Pym and Ritter (2004) gave the following slogan:

Proof-search = Reductive Logic + Control

It is control that determines the efficacy of proof-search procedures: some procedures will be complete, some not; some will affect the shape of proofs being found, and some will affect the complexity of the procedure. The more control structure provided, the more work that is delegated: mechanical problem-solving begets algorithmic theorem-proving techniques, which beget a programming language paradigm known as logic programming (LP). A well-known example of this is provided by implementations of the logic programming language Prolog, in which the typical strategy is ‘leftmost literal first’ in SLD resolution—see, for example, Kowalski and Kuehner (1971) and Plaisted and Zhu (1997). So central are these problems to the use of reductive logic that control is surely a first-class citizen.

How can we represent and reason about control régimes? We have seen that it is possible to extend the operational semantics for reduction with a process calculus, allowing for a flexible way to express control flow during reduction. This setup is more expressive than what can be achieved through the composition of reduction operators (or tactics):



We leave these questions for future work. While important to reductive logic and its applications, they are beyond the semantic foundations of reduction operators, demanding a broader perspective of reductive reasoning.

While deductive logic provides a formal foundation for defining proofs through ‘forwards’ inference, the dual consideration of how one constructs proofs using ‘backwards’ inference is essential for understanding the use of logic when reasoning about a putative conclusion. This is reductive logic. It has profound implications across various fields, includ-

Second, a given reduction operator is valid if there corresponds to it a pre-approved *process* such that whenever the premisses are achieved, the conclusion is also achieved. The notion of achievement here is flexible, enabling many diverse situations to be encapsulated. A process corresponds to an accepted forwards inference. This version of validity is essentially Milner’s theory of tactical proof (Milner 1984). Thus, we regard tactical proof as ‘proof-theoretic semantics’ (Schroeder-Heister 2018) for reductive logic, and this explains the effectiveness of the framework for proof-assistants.

As we have remarked in Sect. 1, this paper should be seen as a companion to Gheorghiu and Pym (2023a). Whereas in Gheorghiu and Pym (2023a) proof-theoretic semantics is used to provide a semantics for tactical proof, this paper provides a general foundation for reductive reasoning in terms of proof-theoretic semantics, establishes direct connections with foundational abstract work in the semantics of computation, and obtains tactical proof as a significant instance.

The scope of this work is the semantic foundations of reduction operators. It has included the semantic foundations of the things constructed by reduction operators. Future work includes giving the semantic foundations for the method by which those constructions (i.e., reductions) are built using reduction operators. This involves, among other things, expressing control regimes, determining the strategy used when exploring the space of reductions, and expressing how one may leverage information about explored portions of the space of reductions to inform future choices. Bundy (1998) has introduced the notion of a *proof plan* as a meta-theoretic tool to represent and reason about the pattern of proof-search and the structure of proofs. Future work also includes developing a semantic account of proof plans as an extension of this paper.

Acknowledgements We are grateful to Simon Docherty, Tao Gu, and Eike Ritter for discussions of various aspects of this work.

Funding This work has been partially supported by the UK EPSRC grants EP/S013008/1 and EP/R006865/1 and by the EU MOSAIC MCSA-RISE project.

Data availability There is not data associated with this article.

Declarations

Conflict of interest The authors have no conflict of interest to declare that are relevant to the content of this article.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

van Atten M (2022) The Development of Intuitionistic Logic. The Stanford Encyclopedia of Philosophy

- Barendregt H (1991) Introduction to Generalized Type Systems. Journal of Functional Programming 1(2):125–154
- Barendregt H, Dekkers W, Statman R (2013) Lambda Calculus with Types. Perspectives in Logic, Cambridge University Press
- Barendregt HP (1993) Lambda calculi with types. Oxford University Press Inc, USA, pp 117–309
- Barwise J, Moss L (1996) Vicious Circles: On the Mathematics of Non-Wellfounded Phenomena. Cambridge University Press
- Beaney M, Raysmith T (2024) Analysis. In: Zalta EN, Nodelman U (eds) The Stanford Encyclopedia of Philosophy, Fall, 2024th edn. Stanford University, Metaphysics Research Lab
- Bergstra JA, Klop JW (1984) The algebra of recursively defined processes and the algebra of regular processes. International Colloquium on Automata, Languages, and Programming pp 82–94
- Bertot Y, Castéran P (2004) Interactive Theorem Proving and Program Development. Springer. <https://doi.org/10.1007/978-3-662-07964-5>
- Beth EW (1955) Semantic Construction of Intuitionistic Logic. Indagationes Mathematicae 17(4):327–338. [https://doi.org/10.1016/S1385-7258\(55\)50014-6](https://doi.org/10.1016/S1385-7258(55)50014-6)
- Bloom B, Istrail S, Meyer AR (1995) Bisimulation can't be traced. Journal of the ACM 42(1):232–268
- Bonchi F, Zanasi F (2015) Bialgebraic Semantics for Logic Programming. Logical Methods in Computer Science 11(1). [https://doi.org/10.2168/LMCS-11\(1:14\)2015](https://doi.org/10.2168/LMCS-11(1:14)2015)
- Brandom R (2000) Articulating Reasons: An Introduction to Inferentialism. Harvard University Press
- Brouwer LEJ (1913) Intuitionism and Formalism. Bulletin of the American Mathematical Society 20(2):81–96
- Bundy A (1985) The Computer Modelling of Mathematical Reasoning. Academic Press Professional, Inc
- Bundy A (1998) Proof planning. University of Edinburgh, Department of Artificial Intelligence, Tech. rep
- Cartmell J (1978) Generalised Algebraic Theories and Contextual Categories. PhD thesis, Oxford University
- Cheng E (2022) The Joy of Abstraction. An exploration of math, category theory, and life. Cambridge University Press
- Clark KL (1977) Negation as Failure. Logic and Data Bases pp 293–322
- Coq Team (2024) The Coq Proof Assistant. <https://coq.inria.fr/>. Accessed July 2024
- van Dalen D (2012) Logic and Structure. Universitext, Springer
- Eckhardt T, Pym DJ (2024) Proof-theoretic Semantics for Modal Logics. Logic Journal of the IGPL In Press - [arXiv:2401.13597](https://arxiv.org/abs/2401.13597)
- Farooque M, Graham-Lengrand S, Mahboubi A (2013) A Bisimulation between DPLL(T) and a Proof-search Strategy for the Focused Sequent Calculus, Association for Computing Machinery, p 3-14. <https://doi.org/10.1145/2503887.2503892>
- Fitch FB (1952) Symbolic Logic: An Introduction. Ronald Press Co
- Francez N (2015) Proof-theoretic Semantics. College Publications
- Frank Pfenning, Carsten Schürmann, Brigitte Pientka, Roberta Vago, and Kevin Watkins (2024) The Twelf Project. http://twelf.org/wiki/Main_Page. Accessed July 2024
- Friedman H (1975) Some systems of second order arithmetic and their use. In: Proceedings of the International Congress of Mathematicians, vol 1. Canadian Mathematical Congress, pp 235–242
- Gheorghiu A, Pym DJ (2023a) Proof-theoretic Semantics and Tactical Proof. [arXiv:2301.02302](https://arxiv.org/abs/2301.02302) Accessed July 2024
- Gheorghiu AV, Pym DJ (2023) Defining Logical Systems via Algebraic Constraints on Proofs. Journal of Logic and Computation. <https://doi.org/10.1093/logcom/exad065>
- Gheorghiu AV, Pym DJ (2023) Semantical Analysis of the Logic of Bunched Implications. Studia Logica 111(4):525–571
- Gheorghiu AV, Docherty S, Pym DJ (2023a) Reductive Logic, Proof-Search, and Coalgebra: A Perspective from Resource Semantics. Samson Abramsky on Logic and Structure in Computer Science and Beyond pp 833–875

- Gheorghiu AV, Gu T, Pym DJ (2023b) Proof-theoretic Semantics for Intuitionistic Multiplicative Linear Logic. In: Automated Reasoning with Analytic Tableaux and Related Methods — TABLEAUX. Springer, pp 367–385
- Gheorghiu AV, Gu T, Pym DJ (2024) A Note on the Practice of Logical Inferentialism. In: 2nd Logic and Philosophy conference, in Press — [arXiv:2403.10546](https://arxiv.org/abs/2403.10546)
- Gillies D (1996) Artificial Intelligence and Scientific Method. Oxford University Press
- Girard JY (1987) Linear Logic. Theoretical Computer Science 50(1):1–101
- Gordon MJ, Melham TF (1993) Introduction to HOL. Cambridge University Press, A theorem proving environment for higher order logic
- Gordon MJ, Wadsworth CP, Milner R (1979) Edinburgh LCF. Springer, A mechanised logic of computation. <https://doi.org/10.1007/3-540-09724-4>
- Gordon MJC (2015) Tactics for mechanized reasoning: a commentary on milner (1984) ‘the use of machines to assist in rigorous proof’. Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences 373(2039)
- Gu T, Gheorghiu AV, Pym DJ (2023) Proof-theoretic Semantics for the Logic of Bunched Implications. [arXiv:2311.16719](https://arxiv.org/abs/2311.16719) Accessed February 2024
- Gupta G, Bansal A, Min R, et al. (2007) Coinductive Logic Programming and its Applications. In: International Conference Logic Programming — ICLP. Springer, pp 27–44
- Harland J, Pym DJ (1997) Resource-distribution via Boolean Constraints. In: Automated Deduction — CADE. Springer, pp 222–236
- Harland J, Pym DJ (2003) Resource-distribution via Boolean Constraints. ACM Transactions on Computational Logic 4(1):56–90
- Heyting A (1989) Intuitionism: An Introduction. Cambridge University Press
- Hodas JS, Miller D (1994) Logic Programming in a Fragment of Intuitionistic Linear Logic. Information and Computation 110(2):327–365
- Hofmann M (1997) Syntax and Semantics of Dependent Types. Semantics and Logics of Computation p 79–130. <https://doi.org/10.1017/CBO9780511526619.004>
- Howard WA (1980) The Formulae-as-Types Notion of Construction. To H B Curry: Essays on Combinatory Logic, Lambda Calculus and Formalism 44:479–490
- Jacobs B (1991) Categorical Type Theory. PhD thesis, The University of Nijmegen
- Jacobs B (2017) Introduction to Coalgebra, vol 59. Cambridge University Press
- Kleene SC (1961) Mathematical Logic. Wiley and Sons
- Knaster B, Tarski A (1928) Un Théorème sur les Fonctions d’Ensembles. Annales de la Societe Polonaise de Mathematique 6:133–134
- Kolmogorov A (1932) Zur Deutung der Intuitionistischen Logik. Mathematische Zeitschrift 35
- Komendantskaya E, Power J (2011) Coalgebraic Semantics for Derivations in Logic Programming. International Conference on Algebra and Coalgebra in Computer Science pp 268–282
- Komendantskaya E, McCusker G, Power J (2011) Coalgebraic Semantics for Parallel Derivation Strategies in Logic Programming. International Conference on Algebraic Methodology and Software Technology - AMAST 13:111–127
- Komendantskaya E, Power J, Schmidt M (2016) Coalgebraic Logic Programming: from Semantics to Implementation. Journal of Logic and Computation 26(2):745–783
- Kowalski R (1979) Algorithm = Logic + Control. Communications of the ACM 22(7):424–436
- Kowalski R (1986) Logic for Problem-Solving. North-Holland Publishing Co
- Kowalski R, Kuehner D (1971) Linear resolution with selection function. Artificial Intelligence 2(3):227–260. [https://doi.org/10.1016/0004-3702\(71\)90012-9](https://doi.org/10.1016/0004-3702(71)90012-9)
- Kripke SA (1965) Semantical Analysis of Intuitionistic Logic I, vol 40, Elsevier, pp 92–130
- Lambek J (1980) From λ -calculus to Cartesian Closed Categories. Essays on Combinatory Logic, Lambda Calculus and Formalism, To H B Curry, pp 375–402
- Lemmon EJ (1978) Beginning Logic. Hackett Publishing
- Makinson D (2014) On an Inferential Semantics for Classical Logic. Logic Journal of IGPL 22(1):147–154
- Martin-Löf P (1975) An Intuitionistic Theory of Types: Predicative Part. Studies in Logic and the Foundations of Mathematics 80:73–118
- Miller D (1989) A Logical Analysis of Modules in Logic Programming. The Journal of Logic Programming 6(1–2):79–108
- Miller D, Nadathur G, Pfenning F et al (1991) Uniform Proofs as a Foundation for Logic Programming. Annals of Pure and Applied logic 51(1–2):125–157
- Milner R (1983) Calculi for Synchrony and Asynchrony. Theoretical Computer Science 25(3):267–310
- Milner R (1984) The Use of Machines to Assist in Rigorous Proof. Philosophical Transactions of the Royal Society of London Series A, Mathematical and Physical Sciences 312(1522):411–422
- Milner R (1989) Communication and Concurrency. Prentice Hall
- Nascimento V (2023) Foundational studies in proof-theoretic semantics. PhD thesis, Universidade do Estado do Rio de Janeiro
- Nascimento V, Pereira LC, Pimentel E (2023) An Ecumenical view of Proof-theoretic Semantics. [arXiv:2306.03656](https://arxiv.org/abs/2306.03656) Accessed July 2024
- Negri S (2005) Proof Analysis in Modal Logic. Journal of Philosophical Logic 34(5):507–544
- nLab (2021) free functor. <https://ncatlab.org/nlab/show/free+functor>. Accessed July 2024
- O’Hearn PW, Pym DJ (1999) The Logic of Bunched Implications. Bulletin of Symbolic Logic 5(2):215–244
- Paulson LC (1994) Isabelle: A Generic Theorem Prover. Springer. <https://doi.org/10.1007/BFb0030541>
- Pavlović D (1990) Predicates and Fibrations. PhD thesis, University of Utrecht
- Pfenning F, Schürmann C (1999) System description: Twelf — a meta-logical framework for deductive systems. In: Automated Deduction — CADE, pp 202–206
- Piecha T (2016) Completeness in Proof-theoretic Semantics. Advances in Proof-theoretic Semantics pp 231–251
- Piecha T, Schroeder-Heister P (2019) Incompleteness of Intuitionistic Propositional Logic with Respect to Proof-theoretic Semantics. Studia Logica 107(1):233–246
- Piecha T, de Campos Sanz W, Schroeder-Heister P (2015) Failure of Completeness in Proof-theoretic Semantics. Journal of Philosophical Logic 44(3):321–335
- Plaisted DA, Zhu Y (1997) The Efficiency of Theorem Proving Strategies. Springer
- Plotkin GD (1981) A Structural approach to Operational Semantics. Aarhus University, Tech. rep
- Portoraro F (2024) Automated Reasoning. In: Zalta EN, Nodelman U (eds) The Stanford Encyclopedia of Philosophy, Spring, 2024th edn. Stanford University, Metaphysics Research Lab
- Pym DJ, Ritter E (2004) Reductive Logic and Proof-search: Proof Theory, Semantics, and Control, Oxford Logic Guides, vol 45. Oxford University Press
- Pym DJ, Wallen LA (1993) Logic programming via proof-valued computations. In: Proceedings of the UK Conference on Logic Programming — ALPUK. Springer, pp 253–262

- Pym DJ, Ritter E, Robinson E (2024) Categorical proof-theoretic semantics. *Studia Logica* pp 1–38. <https://doi.org/10.1007/s11225-024-10101-9>
- Reitzig R (2012) Discussion on Computer Science Stack Exchange. <https://cs.stackexchange.com/questions/525/what-is-coinduction>. Accessed June 2024
- Rutten J (2000) Universal Coalgebra: a theory of systems. *Theoretical Computer Science* 249(1):3–80. [https://doi.org/10.1016/S0304-3975\(00\)00056-6](https://doi.org/10.1016/S0304-3975(00)00056-6)
- Sandqvist T (2005) An inferentialist Interpretation of Classical Logic. PhD thesis, Uppsala University
- Sandqvist T (2009) Classical logic without Bivalence. *Analysis* 69(2):211–218
- Sandqvist T (2015) Base-extension Semantics for Intuitionistic Sentential Logic. *Logic Journal of the IGPL* 23(5):719–731
- MacLane Saunders (1971) *Categories for the Working Mathematician*. Springer
- Schroeder-Heister P (1984) A natural extension of natural deduction. *The Journal of Symbolic Logic* 49(4):1284–1300
- Schroeder-Heister P (2006) Validity Concepts in Proof-theoretic Semantics. *Synthese* 148(3):525–571
- Schroeder-Heister P (2008) Proof-Theoretic versus Model-Theoretic Consequence. *The Logica Yearbook 2007*
- Schroeder-Heister P (2018) Proof-theoretic Semantics. *The Stanford Encyclopedia of Philosophy*
- Seely RA (1983) Hyperdoctrines, Natural Deduction and the Beck condition. *Mathematical Logic Quarterly* 29(10):505–542
- Simon L, Bansal A, Mallya A, et al. (2007) Co-Logic Programming: extending logic programming with coinduction. In: *International Colloquium on Automata, Languages and Programming — ICALP*. Springer, pp 472–483
- Smullyan RM (1968) *Analytic Tableaux*, Springer, pp 15–30. https://doi.org/10.1007/978-3-642-86718-7_2
- Stafford W (2021) Proof-theoretic Semantics and Inquisitive Logic. *Journal of Philosophical Logic*
- Stafford W, Nascimento V (2023) Following all the Rules: Intuitionistic Completeness for Generalized Proof-theoretic Validity. *Analysis*. <https://doi.org/10.1093/analys/anac100>
- Streicher T (1988) Correctness and Completeness of a Categorical Semantics of the Calculus of Constructions. PhD thesis, University of Passau
- Szabo ME (ed) (1969) *The Collected Papers of Gerhard Gentzen*. North-Holland Publishing Company
- Tarski A (1936) On the Concept of Logical Consequence. *Logic, semantics, metamathematics* 52:409–420
- Tarski A (1955) A Lattice-theoretical Fixpoint Theorem and Its Applications. *Pacific Journal of Mathematics* 5(2):285–309
- Troelstra AS, Schwichtenberg H (2000) *Basic Proof Theory*, Cambridge Tracts in Theoretical Computer Science, vol 43. Cambridge University Press
- Turi D, Plotkin GD (1997) Towards a Mathematical Operational Semantics. In: *Logic in Computer Science — LICS*. IEEE, pp 280–291, <https://doi.org/10.1109/LICS.1997.614955>
- Wallen LA (1987) Automated proof search in non-classical logics : efficient matrix proof methods for modal and intuitionistic logics. PhD thesis, University of Edinburgh
- Wansing H (2000) The Idea of a Proof-theoretic Semantics and the Meaning of the Logical Operations. *Studia Logica* 64:3–20
- Wikipedia Proof Assistant. https://en.wikipedia.org/wiki/Proof_assistant
- Worrell J (2005) On the Final Sequence of a Finitary Set Functor. *Theoretical Computer Science* 338(1):184–199. <https://doi.org/10.1016/j.tcs.2004.12.009>