

On the Logical Content of Knowledge Bases

Alexander V. Gheorghiu^{1,2}[0000–0002–7144–6910] and
Tao Gu^{1,2}[0000–0001–5749–0758]

¹ University of Southampton, Southampton SO17 1BJ, UK

² University College London, London WC1E 6BT, UK
{a.v.gheorghiu, t.gu}@soton.ac.uk

Abstract. Standard epistemic logics introduce a modal operator \mathbf{K} to represent knowledge, but in doing so presuppose the logical apparatus they aim to explain. By contrast, this paper explores how logic may be derived from the structure of knowledge itself. We begin from a pre-logical notion of a knowledge base understood as a network of inferential connections between atomic propositions. Logical constants are then defined in terms of what is supported by such a base: intrinsically, by relations that hold within it, and extrinsically, by the behaviour of those relations under extension. This yields a general semantic framework in which familiar systems — classical, intuitionistic, and various intermediate logics — arise naturally from different assumptions about the form of knowledge. This offers a reversal of the traditional explanatory order: rather than treating logic as a precondition for the articulation of knowledge, it shows how logical structure can emerge from epistemic organisation.

Keywords: epistemic logic · meaning and knowledge · inferentialism · proof-theoretic semantics

1 Introduction

The standard treatment of knowledge in logic is modal: one introduces an operator \mathbf{K} such that $\mathbf{K}\phi$ expresses that ϕ is known. Systems axiomatizing \mathbf{K} are known as *epistemic logics* [7]. This framework has proved fruitful, enabling the analysis of common knowledge and multi-agent reasoning, and supplying formal methods for AI, game theory, economics, distributed computing, and natural-language semantics. Philosophical work in this tradition typically concerns the status of principles such as self-reflection ($\mathbf{K}\phi \rightarrow \mathbf{K}\mathbf{K}\phi$), rather than the *epistemic content* of formulae themselves. In this sense, such systems are best viewed as accounts of *knowability* rather than of knowledge proper.

By contrast, we argue that knowledge should be taken as conceptually prior to logic. By this we mean that we might begin with a concept of knowledge and determine a logic from it. This is the process typically used in machine intelligence. For example, it is the process used in the traditional approach to AI known as *logic programming* (LP) — see Kowalski [9, 10]. This pragmatic approach to knowledge and reasoning motivates an alternative understanding of ‘epistemic’ logic that offer a deeper theoretical understanding of how logic

can emerge from knowledge. The logics we derive according to this plan are not epistemic logics in the traditional sense, but rather a logics whose semantics is given in terms of ‘knowledge’ (rather than, say, possible worlds).

We begin by introducing a primitive notion of a knowledge base \mathcal{B} . This is a structure in which items of information are directly connected. We then define logical signs depending on how one interprets its *logical content*. To illustrate, consider conjunction: What should it mean to say that ‘ ϕ and ψ holds for \mathcal{B} ?’ Intuitively, it requires that both ϕ and ψ are independently known by \mathcal{B} . This case is straightforward, perhaps even trivial.

More interesting is the case of the conditional: what should it mean to assert that ‘if ϕ , then ψ ’? Intuitively, this says that ψ is supported by \mathcal{B} whenever ϕ is. Taken at face value this account collapses into vacuity: the conditional would be satisfied whenever ϕ is not supported by \mathcal{B} , thereby undermining the hypothetical force of the statement. To properly represent the hypothetical-force of the if-then statement we require something more. We consider extensions \mathcal{C} of \mathcal{B} such that ψ is supported in them whenever ϕ is.

To make this analysis precise, in this paper we proceed relative to a simple account of knowledge as an *atomic system*. That is, a collection of atomic propositions structured by some inferential connections. The intuition is developed in several places; see, for example, Piecha and Schroeder-Heister [21, 14]. To know ‘All humans are mortal’ is to have in mind the inferential connection,

$$\frac{\text{‘}x \text{ is a human’}}{\text{‘}x \text{ is mortal’}}$$

as x ranges over the names of all individuals. (Returning to LP, in PROLOG the above example is captured by a program containing

$$M(x) :- H(x)$$

where M is the predicate ‘is mortal’, H is the predicate ‘is human’.) We structure knowledge as this set of inferential connections because some knowledge may be implicit and accessed through internal reasoning. Suppose that an agent Aristotle who knows both the above rule and that ‘Socrates is a human’ (i.e., $H(s)$ where s is ‘Socrates’). These commitments mean that they also know that ‘Socrates is mortal’.

One might object that these inferential connections already presuppose some form of *logic* since they are deductive in character. Importantly though no logical signs are in play here: none have yet been defined. The horizontal bar represents deduction at a meta-theoretic level. Our aim is to *define* logical operators — implication, conjunction, disjunction, and negation — by appeal to the *logical content* encoded in these knowledge bases. This is the central work of the paper.

Given a knowledge base \mathcal{B} and an atomic proposition P , we write $\vdash_{\mathcal{B}} P$ when the commitments contained in \mathcal{B} suffice to derive P — the details are given in Section 4. An example of an atomic system following the discussion of Aristotle above is as follows:

$$\frac{H(x)}{M(x)} \quad \overline{H(s)}$$

as x ranges over the names of individuals. These are natural deduction rules in the sense of Gentzen [24], but are read *per se* so composed without substitution.

The logical signs are defined by a *support* judgment (\Vdash). In the simplest case, suppose an atomic proposition P follows directly from the commitments of the knowledge base, then P is supported. Thus,

$$\Vdash_{\mathcal{B}} P \quad \text{iff} \quad \vdash_{\mathcal{B}} P$$

The challenge is to extend this to characterize $\Vdash_{\mathcal{B}} \phi$ where ϕ is a logically-complex formula involving, say, implications, disjunctions, conjunctions, or negations.

In this formal setup, we may revisit the analysis of implication above to give a formal definition. Recall that intuitively, ‘if ϕ , then ψ ’ is supported by a knowledge base \mathcal{B} just in case extension $\mathcal{C} \supseteq \mathcal{B}$ that supports ϕ also supports ψ . What are these extensions? There are two possibilities. An initial account uses a translation ($\llbracket - \rrbracket$) from the logical syntax into inferential form of the knowledge base. Using this, we may represent if-then as the logical sign \rightarrow defined as follows:

$$\Vdash_{\mathcal{B}} \phi \rightarrow \psi \quad \text{iff} \quad \Vdash_{\mathcal{B}, \llbracket \phi \rrbracket} \psi$$

The problem with this approach is that such a translation ($\llbracket \phi \rrbracket$) might only exist when ϕ already has some particular structure (e.g., it is in some clausal form). This restriction makes this clause, perhaps, limited in expressiveness. We may call this the *intrinsic* implication because it speaks of the knowledge itself; for example, the definition of $\phi \rightarrow \psi$ considers the situation in which ϕ is itself parts of the knowledge base.

A more cautious alternative is to define a contrasting extrinsic implication $\dashv\vdash$. Intuitively, $\phi \dashv\vdash \psi$ is supported by \mathcal{B} whenever an arbitrary extension $\mathcal{C} \supseteq \mathcal{B}$ supports ϕ it also supports ψ :

$$\Vdash_{\mathcal{B}} \phi \dashv\vdash \psi \quad \text{iff} \quad \text{for any } \mathcal{C} \supseteq \mathcal{B}, \text{ if } \Vdash_{\mathcal{C}} \phi, \text{ then } \Vdash_{\mathcal{C}} \psi$$

Observe that this speaks about the behaviour of the knowledge. That is, $\phi \dashv\vdash \psi$ declares the inferential commitment that in any hypothetical scenario where the knowledge is extended in such a way to support ϕ one is also committed to ψ . Of course, one may expect some parity between intrinsic \rightarrow and extrinsic $\dashv\vdash$ — for example, one may expect the latter to be a conservative extension of the former.

In Section 3 we systematically provide intrinsic and extrinsic versions of all the standard connectives — conjunction, disjunction, and negation. We then analyse the resulting semantics in Section 4. Perhaps surprisingly, in one setting we recover intuitionistic logic, in another classical logic, and in others intermediate systems that lie between them. This then means that we may read these systems as logics of knowledge; thus, we may say that they are *epistemic* logics in the sense that they express the logical content of knowledge bases.

This strategy described here of constructing a logic out of justificatory structure (i.e., knowledge represented as a network of inferential links) closely follows

Dummett’s [2] programme. Indeed it places us within the tradition of *proof-theoretic semantics*. This is a long-standing project in modern logic; we defer to Schroeder-Heister [20] for a concise summary of its history and development. The central idea is this: the meaning of a logical sign is not fixed by pointing to some external model or mathematical structure; rather, it is fixed by the inferential role that the sign plays in reasoning.

2 Knowledge Bases as Atomic Systems

We begin with the notion of an *atomic system*. An atomic system records an agent’s commitments concerning the inferential relations between atomic sentences. These are pre-logical in the sense that they contain no logical operators. Formally, atomic systems are sets of atomic rules which are the inferential relations and are represented as natural deduction rules in the sense of Gentzen [24] and generalized by Schroeder-Heister [19].

To illustrate how this works, we consider an example. An agent may or may not know whether or not A(lice) is a parent of B(ob). However, because they understand the relationship between parenthood and relative age they can work that if ‘A is a parent of B’, it must be that ‘A is older than B’. They can do this because whatever statement P follows from if A is B’s father or B’s mother also follows A being B’s parent. This is represented as an atomic rule as follows:

$$\frac{\text{A is a parent of B} \quad \frac{[\text{A is B's father}] \quad P}{P} \quad \frac{[\text{A is B's mother}] \quad P}{P}}{P}$$

The $[-]$ marks discharge as usual. Since the agent also knows

$$\frac{\text{A is B's father}}{\text{A is older than B}} \quad \text{and} \quad \frac{\text{A is B's mother}}{\text{A is older than B}}$$

they can make the requisite deduction.

Working in a predicate language raises a choice. One may include variables in atomic systems or restrict to closed atoms. Gheorghiu [4] has shown that both formalisms have the same expressive power. To keep things simple we will *not* include variables.

We assume that we have denumerably many atomic propositions \mathbb{A} .

Definition 1 (Atomic Rule). *An n th-level atomic rule is defined as follows, where $P_1, \dots, P_k, C \in \mathbb{A}$ and $\Sigma_1, \dots, \Sigma_k$ are (possibly empty) sets of n th-level atomic rules:*

- A zeroth-level atomic rule is a rule

$$\frac{}{C}$$

- A first-level atomic rule

$$\frac{P_1 \quad \dots \quad P_k}{C}$$

- An $(n + 2)$ th-level atomic rule is a rule

$$\frac{[\Sigma_1] \quad \dots \quad [\Sigma_k]}{P_1 \quad \dots \quad P_k} C$$

where Σ_i are finite sets of n th-level rules.

To be clear: premisses may be empty, so that an m th-level atomic rule also counts as an n th-level atomic rule for any $n > m$.

Definition 2 (Atomic System). An atomic system is a set of atomic rules.

An atomic system is a set of inferential commitments, but we must be precise about what this amounts to. Heuristically, commitment to the premisses entails commitment to the conclusion; this is the reading of inference adopted by Gentzen [24] — see also Piecha and Schroeder-Heister [21, 14]. We formalise this idea by defining a derivability judgement, which specifies when a conclusion is supported by an atomic system together with a set of assumptions. In the general case we proceed as follows:

Definition 3 (Derivation in an Atomic System). Let \mathcal{A} be an atomic system. The set of \mathcal{A} -derivations is defined inductively as follows:

- **BASE CASE.** If \mathcal{A} contains a zeroth-level rule concluding C , then the natural deduction argument consisting only of the node C is an \mathcal{A} -derivation. Likewise, if \mathcal{A} contains a first-level rule

$$\frac{P_1 \quad \dots \quad P_k}{C}$$

and there is an \mathcal{A} -derivation \mathcal{D}_i with hypotheses Γ_i and conclusion P_i . Then the natural deduction argument with root C and immediate sub-trees $\mathcal{D}_1, \dots, \mathcal{D}_k$ is an \mathcal{A} -argument from $\Gamma_1 \cup \dots \cup \Gamma_k$ to C .

- **INDUCTION STEP.** Suppose \mathcal{A} contains an $(n + 2)$ th-level rule

$$\frac{[\Sigma_1] \quad \dots \quad [\Sigma_k]}{P_1 \quad \dots \quad P_k} C$$

and suppose that, for each $1 \leq i \leq k$, there is an \mathcal{A} -derivation \mathcal{D}_i with hypotheses $\Gamma_i \cup \Sigma_i$ and conclusion P_i . Then the natural deduction argument with root C and immediate sub-trees $\mathcal{D}_1, \dots, \mathcal{D}_k$ is an \mathcal{A} -argument from $\Gamma_1 \cup \dots \cup \Gamma_k$ to C .

An atom C is derivable from Γ in \mathcal{A} — denoted $\Gamma \vdash_{\mathcal{A}} C$ — iff there is a \mathcal{A} -derivation from Γ to C .

For those familiar with natural deduction after Gentzen [24], second-level rules, which include zeroth- and first-level rules, will be familiar and unsurprising. Meanwhile, third-level rules and beyond are much less common and may benefit from presentation by an example.

Example 1 (Promissory Estoppel). We have the follow propositions:

- a — ‘promisor made a promise’
- b — ‘promisee reasonably relied on the promise’
- c — ‘injustice can be avoided only by enforcement’
- d — ‘promise is legally binding’.

The doctrine of *promissory estoppel* is the statement:

“If a promise was made (a), and if the fact that promises of this type induce reasonable reliance (i.e., that making such a promise would lead to reasonable reliance, $[a] \triangleright b$) is itself what makes enforcement necessary to avoid injustice (c), then the promise is legally binding (d).”

This is represented by the following third-level rule:

$$\frac{\frac{a}{\frac{b}{c}}}{d}$$

A promise becomes enforceable not simply because:

- A promise was made (a), or
- Reliance occurred (b), or
- Injustice would result without enforcement (c)

Rather, the promise is binding because the *reliance-inducing character* of the promise

$$\frac{a}{b}$$

is precisely what makes non-enforcement unjust (c).

Not every form of atomic system needs to be admitted, corresponding to different conceptions of knowledge. To capture this, we introduce the notion of a *basis*, which specifies the class of atomic systems under consideration. Given a fixed basis \mathfrak{B} , its elements are called *bases* \mathcal{B} . Once a basis is fixed, we always work relative to its bases. Accordingly we introduce the notion of *base-extension*, a restricted version of set-theoretic extension that respects the chosen basis.

Definition 4 (Basis). A basis \mathfrak{B} is a set of atomic systems.

Definition 5 (Base-extension). Given a basis \mathfrak{B} , base-extension is the smallest relation satisfying

$$\mathcal{Y} \supseteq_{\mathfrak{B}} \mathcal{X} \quad \text{iff} \quad \mathcal{X}, \mathcal{Y} \in \mathfrak{B} \text{ and } \mathcal{Y} \supseteq \mathcal{X}.$$

Our semantics for the logical constants will be parameterised by the notion of basis. In particular, we shall consider the following cases, where \mathcal{A} ranges over atomic systems:

- $\mathfrak{B}_1 := \{\mathcal{A} \mid \mathcal{A} \text{ is zeroth- or first-level}\}$,
- $\mathfrak{B}_2 := \{\mathcal{A} \mid \mathcal{A} \text{ is zeroth-, first-, or second-level}\}$
- \mathfrak{B}_ω comprises all atomic systems.

In fact, we will restrict ourselves to the first two only as \mathfrak{B}_2 and \mathfrak{B}_ω are equivalent in the sense any atomic system in \mathfrak{B}_ω has an equivalent atomic system in \mathfrak{B}_2 in terms of derivability. More formally, Stafford et al. [23] has shown that if one allows an indefinite supply of fresh atoms, then for each $\mathcal{B} \in \mathfrak{B}_\omega$ one can construct a $\mathcal{B}' \in \mathfrak{B}_2$ such that $\vdash_{\mathcal{B}} P$ iff $\vdash_{\mathcal{B}'} P$ for arbitrary P in the original language (i.e., excluding the fresh atoms).

Example 2 (Example 1 cont'd). Consider the rule for promissory estoppel. It can be represented by the following second-level rules:

$$\frac{[a]}{b} \quad \frac{a \quad x}{b} \quad \text{and} \quad \frac{[x]}{c} \quad \frac{c}{d}$$

where x is a fresh atom alien to the domain of discourse. This device of introducing x has been used by Mints [13] within computational logic — see also Gheorghiu [3]

We shall henceforth take atomic systems as our model of knowledge bases. A base records the atomic commitments of an agent together with the inferences they acknowledge between them. In this way it represents the informational state of the agent prior to the introduction of logical constants. The subsequent task is to investigate what logical content such a base supports.

3 Logical Content

We have now given a mathematical account of knowledge bases as atomic systems. What remains is to draw out their logical significance. Our aim is to define the logical constants in terms of knowledge itself — not by appealing to an external semantics of models or valuations, but by seeing how logic arises from the structure and behaviour of knowledge bases. In Section 1, we saw that there are two distinct attitudes one might take toward what the logical signs express.

One may regard them as describing *intrinsic* relations within the body of knowledge: how the contents of a knowledge base are internally ordered — this is the view captured by the connective \rightarrow . This concerns the structure of the current state of knowledge and so the logical signs explain the relations that hold within that state. It thus mirrors the model-theoretic conception based on truth-*in-a-model* after Tarski [25, 26].

Alternatively, one may regard them as expressing *extrinsic* relations: how a body of knowledge behaves when it is extended, revised, or placed in interaction with new information — this corresponds to the connective \multimap . By ‘behaves’ we mean that these connectives are declarations about one’s commitments as one’s

knowledge changes. Taking knowledge to be monotonic (cf. Pollock [15]), the dynamics we have in mind here are *extensions* of the current knowledge base.

Let $\Gamma \Vdash_{\mathcal{B}} \phi$ denote that ϕ is supported when the knowledge \mathcal{B} is extended with inferential commitments that licence one to support the formulae in Γ . Following the meaning of the extrinsic implication $\dashv\vdash$, we may define it as follows:

$$\Gamma \Vdash_{\mathcal{B}} \phi \quad \text{iff} \quad \text{for any } \mathcal{C} \supseteq \mathcal{B}, \text{ if } \Vdash_{\mathcal{C}} \psi \text{ for all } \psi \in \Gamma, \text{ then } \Vdash_{\mathcal{C}} \phi$$

This definition forms the basis on which we may properly define the extrinsic connectives.

The *extrinsic* connectives articulate ones *inferential commitments* under the assumption that their constituent formulas hold. They are distinguished by the way those constituents are structured. The ‘structures’ we have in mind may be organized by the corresponding intrinsic connectives precisely because they are declarations about the state of knowledge. Heuristically, the extrinsic connective E matches the intrinsic connective I in the sense that $\phi E \psi$ holds exactly when, across all possible extensions of the current knowledge base, any atomic proposition P that would be supported given $\phi I \psi$ is also supported *in its own right*. This is expressed formally as follows:

$$\Vdash_{\mathcal{B}} \phi E \psi \quad \text{iff} \quad \text{for all } \mathcal{C} \supseteq \mathcal{B} \text{ and all atoms } P, \text{ if } \phi I \psi \Vdash_{\mathcal{C}} P, \text{ then } \Vdash_{\mathcal{C}} P$$

The extrinsic connectives are therefore characterised not by the *content* of the knowledge base but by the *inferential role* they impose upon it.

Conjunction. A conjunction ‘ ϕ and ψ ’ represents that ϕ and ψ are both known (relative to a knowledge base \mathcal{B}). In Section 1, we gave the intrinsic reading of this as the familiar form in which ϕ and ψ are both independently supported by the knowledge, expressed as follows:

$$\Vdash_{\mathcal{B}} \phi \wedge \psi \quad \text{iff} \quad \Vdash_{\mathcal{B}} \phi \text{ and } \Vdash_{\mathcal{B}} \psi$$

This leaves open the question of what would be the so-called extrinsic conjunction (\otimes). Following the explanation above, $\phi \otimes \psi$ is to be defined by the inferential behaviour of $\phi \wedge \psi$. This may be expressed as follows:

$$\Vdash_{\mathcal{B}} \phi \otimes \psi \quad \text{iff} \quad \text{for any } \mathcal{C} \supseteq \mathcal{B} \text{ and atom } P, \text{ if } \phi \wedge \psi \Vdash_{\mathcal{C}} P \text{ then } \Vdash_{\mathcal{C}} P$$

Unpacking \wedge , this can be expressed as follows:

$$\Vdash_{\mathcal{B}} \phi \otimes \psi \quad \text{iff} \quad \text{for any } \mathcal{C} \supseteq \mathcal{B} \text{ and atom } P, \text{ if } \phi, \psi \Vdash_{\mathcal{C}} P \text{ then } \Vdash_{\mathcal{C}} P$$

This formulation captures the idea that \oplus represents a form of conjunction where we acknowledge that the conjunctions hold but without committing to them. In Section 4, we observe that under certain settings these conjunctions collapse into each other.

Disjunction. Intuitively, a disjunction ‘ ϕ or ψ ’ represents a collection of alternatives. The intrinsic attitude says that one of the disjuncts must hold:

$$\Vdash_{\mathcal{B}} \phi \vee \psi \quad \text{iff} \quad \Vdash_{\mathcal{B}} \phi \text{ or } \Vdash_{\mathcal{B}} \psi$$

Following the approach above, the extrinsic disjunction (\oplus) is then defined as follows:

$$\Vdash_{\mathcal{B}} \phi \oplus \psi \quad \text{iff} \quad \text{for any } \mathcal{C} \supseteq \mathcal{B} \text{ and atom } P, \text{ if } \phi \vee \psi \Vdash_{\mathcal{C}} P \text{ then } \Vdash_{\mathcal{C}} P$$

Unpacking \vee , this can be expressed as follows:

$$\Vdash_{\mathcal{B}} \phi \oplus \psi \quad \text{iff} \quad \text{for any } \mathcal{C} \supseteq \mathcal{B} \text{ and atom } P, \text{ if } \phi \Vdash_{\mathcal{C}} P \text{ and } \psi \Vdash_{\mathcal{C}} P, \text{ then } \Vdash_{\mathcal{C}} P.$$

This formulation captures the idea that \oplus represents a form of disjunction where we acknowledge that one of a set of alternatives holds but without committing to any specific one.

There are many natural examples of cases modelled by such extrinsic disjunctions. Consider, for instance, the following scenario:

There is a lottery with 1000 tickets. Let the statement that the i th ticket wins be T_i . Since there are exactly 1000 tickets, we can assert the disjunction

$$T_1 \text{ or } T_2 \text{ or } \dots \text{ or } T_{1000}.$$

The ‘or’ here cannot be the intrinsic disjunction (\vee) as we, of course, do not know which ticket will win. Rather it is the extrinsic one; that is, our knowledge is indeed $T_1 \oplus \dots \oplus T_{1000}$.

What the ‘or’ here means is best understood in terms of behaviour. Suppose that the participants have agreed that, whichever ticket wins, the prize money will be used to buy supplies for an office party. Let P be the statement that the office party will be supplied. Then whoever ticket wins (i.e., whichever T_i holds) it follows that P holds — that is, $T_i \Vdash_{\mathcal{B}} P$ for any i . The definition of \oplus is exactly such that from this together with the knowledge that at least one ticket wins (i.e., $T_1 \oplus \dots \oplus T_{1000}$) we know that P : the office party will be supplied.

Negation. We now turn to negation. To begin, we introduce both an intrinsic and an extrinsic unit for absurdity. Following the pattern established above, the intrinsic one is given by

$$\Vdash_{\mathcal{B}} \perp \quad \text{never holds}$$

and the extrinsic one by

$$\Vdash_{\mathcal{B}} 0 \quad \text{iff} \quad \text{for all } P, \Vdash_{\mathcal{B}} P.$$

These serve as the respective units for negation: a statement is negated by asserting that its assumption leads to absurdity.

Since there are two implications and two units, we obtain four possible negations:

$$\phi \rightarrow \perp, \quad \phi \rightarrow 0, \quad \phi \multimap \perp, \quad \phi \multimap 0.$$

The first of these is vacuous as it never obtains, but all three of the remainders are meaningful.

The second (i.e., $\phi \rightarrow 0$) expresses that when ϕ is added to the knowledge base \mathcal{B} , any proposition can be proved. It thus represents an intrinsic notion of negation in the proof-theoretic sense — that is, as given by *ex falso quodlibet*.

The third (i.e., $\phi \multimap \perp$) expresses that ϕ is not supported by the current knowledge base. It therefore validates the disjunctive syllogism with respect to the intrinsic disjunction:

$$\text{if } \Vdash_{\mathcal{B}} \phi \vee \psi \text{ and } \Vdash_{\mathcal{B}} \phi \multimap \perp, \text{ then } \Vdash_{\mathcal{B}} \psi.$$

Finally, the fourth (i.e., $\phi \multimap 0$) gives an *extrinsic* form of disjunctive syllogism. Given $\Vdash_{\mathcal{B}} \phi \oplus \psi$ and $\Vdash_{\mathcal{B}} \phi \multimap 0$, we obtain the following conditional:

$$\text{for any } \mathcal{C} \supseteq \mathcal{B} \text{ and any atom } P, \text{ if } \psi \Vdash_{\mathcal{C}} P, \text{ then } \Vdash_{\mathcal{C}} P$$

This expresses the extrinsic reading of ϕ in precisely the sense used to define the extrinsic connectives. It does not ask whether ϕ itself is supported, but rather concerns its inferential behaviour: the status of some arbitrary datum (an atom P) in a hypothetical situation in which knowing ϕ would lead us to support that datum.

4 Base-extension Semantics

Assume we have a function $(-)$ from some class of formulae to bases. The logical signs discussed in Section 3 are collected in Figure 1, where Δ is non-empty and for (\rightarrow) we have the caveat that ϕ belongs to the domain of $(-)$.

What would be a suitable $(-)$? Given the form of knowledge base, we can choose it as the inverse of the intuitive mapping from bases to sets of formulae. Let $\llbracket - \rrbracket$ be an injection defined as follows:

$$\left\{ \begin{array}{l} \llbracket \overline{C} \rrbracket := C \\ \llbracket \frac{P_1 \dots P_n}{C} \rrbracket := (P_1 \wedge \dots \wedge P_n) \rightarrow C \\ \llbracket \frac{[\Sigma_1] \dots [\Sigma_n]}{C} \rrbracket := ((\llbracket \Sigma_1 \rrbracket \rightarrow P_1) \wedge \dots \wedge (\llbracket \Sigma_n \rrbracket \rightarrow P_n)) \rightarrow C \end{array} \right.$$

The choice of \wedge and \rightarrow over \otimes and \multimap reflects their intrinsic nature as the formula is put, literally, inside the knowledge base. We set $(-)$ as the left-inverse of this injection.

We may now ask the following: under this choice of knowledge base as atomic systems, $(-)$ as the left inverse of $\llbracket - \rrbracket$, and these clauses, does this semantics correspond to any known logic? There are several possible answers depending on further details:

$\Vdash_{\mathcal{B}} P$	iff $\vdash_{\mathcal{B}} P$	(At)
$\Vdash_{\mathcal{B}} \phi \rightarrow \psi$	iff $\Vdash_{\mathcal{B}, \{\phi\}} \psi$	(\rightarrow)
$\Vdash_{\mathcal{B}} \phi \multimap \psi$	iff $\phi \Vdash_{\mathcal{B}} \psi$	(\multimap)
$\Vdash_{\mathcal{B}} \phi \wedge \psi$	iff $\Vdash_{\mathcal{B}} \phi$ and $\Vdash_{\mathcal{B}} \psi$	(\wedge)
$\Vdash_{\mathcal{B}} \phi \otimes \psi$	iff for any $\mathcal{C} \supseteq \mathcal{B}$ and atom P , if $\phi, \psi \Vdash_{\mathcal{C}} P$, then $\Vdash_{\mathcal{C}} P$	(\otimes)
$\Vdash_{\mathcal{B}} \perp$	never	(\perp)
$\Vdash_{\mathcal{B}} 0$	iff for any atom $P, \Vdash_{\mathcal{B}} P$	(0)
$\Vdash_{\mathcal{B}} \phi \vee \psi$	iff $\Vdash_{\mathcal{B}} \phi$ or $\Vdash_{\mathcal{B}} \psi$	(\vee)
$\Vdash_{\mathcal{B}} \phi \oplus \psi$	iff for any $\mathcal{C} \supseteq \mathcal{B}$ and atom P , if $\phi \Vdash_{\mathcal{C}} P$ and $\psi \Vdash_{\mathcal{C}} P$, then $\Vdash_{\mathcal{C}} P$	(\oplus)
$\Delta \Vdash_{\mathcal{B}} \phi$	iff $\forall \mathcal{C} \supseteq \mathcal{B}$, if $\Vdash_{\mathcal{C}} \psi$ for any $\psi \in \Delta$, then $\Vdash_{\mathcal{C}} \phi$	(Inf)

Fig. 1. Base-extension Semantics

The intrinsic fragment of Figure 1 corresponds to logic programming in the style of Miller [11, 12] and Harland [6] based on the Harrop formulae. This is a programming language paradigm based on predicate logic: certain clausal formulae (in this case, hereditary Harrop formulae) are read as instructions and computation proceeds via proof-search — see Kowalski [8, 9]. In this case, the bases \mathcal{B} are *programs* and the judgment $\Vdash_{\mathcal{B}} \phi$ says that ϕ may be computed from \mathcal{B} . This makes sense on the knowledge reading as such programs are sometimes regarded as knowledge bases — see, for example, Baral and Gelfond [1].

The extrinsic fragment is, perhaps, more surprising. Depending on the basis, one recovers two different but much celebrated logics. Using \mathfrak{B}_1 one recovers classical logic but using \mathfrak{B}_2 one recovers intuitionistic logic — see Sandqvist [17, 18] and Gheorghiu [4]. In this setting, the extrinsic conjunction collapses into the intrinsic one:

$$\Vdash_{\mathcal{B}} \phi \otimes \psi \quad \text{iff} \quad \Vdash_{\mathcal{B}} \phi \text{ and } \Vdash_{\mathcal{B}} \psi$$

— see Gheorghiu et al. [5]. That the change of basis while fixing the semantic clauses should have this effect remains mysterious.

We may also consider certain hybrid systems that comprise both intrinsic and extrinsic connectives. Stafford et al. [22, 23] have investigated the system comprised of the connectives \multimap , \wedge , \vee , and \perp . If one restricts to \mathfrak{B}_1 , they define Stafford’s Logic; and, if one takes \mathfrak{B}_2 , they determine general inquisitive logic (see Punčochář [16]) — in short, intuitionistic logic together with the generalized Kreisel-Putnam axiom:

$$(\chi \multimap (\phi \vee \psi)) \multimap ((\chi \multimap \phi) \vee (\chi \multimap \psi))$$

as χ ranges over \vee -free formulae.

5 Conclusion

We have presented a framework for analysing the logical content of knowledge bases conceived as atomic systems. Within this setting, the logical constants are defined by the inferential structure of knowledge itself: intrinsically, through relations that hold within a base, and extrinsically, through the behaviour of those relations under extension. The resulting base-extension semantics yields a range of well-known logics — classical, intuitionistic, and intermediate systems — depending on the basis adopted. Under this purview, these logics may be understood as epistemic: each expresses a distinct conception of how knowledge is structured and how it persists under revision. Thus, rather than taking logic as prior to knowledge, we may view the traditional logical systems as arising from it as formal articulations of its inferential and dynamic organisation.

This analysis has begun with a deliberately simple conception of a knowledge base as an atomic system. In reality, the structure of knowledge — whether in human cognition or in artificial agents — is likely to be far more intricate, or at least could well be conceived as such. Yet even within this minimal setting, we have seen that a seemingly small variation — the presence or absence of hypotheses in the atomic rules, the move from \mathfrak{B}_2 to \mathfrak{B}_1 — already gives rise to distinct logics, namely the intuitionistic and the classical. It is therefore natural to ask what further logics might emerge when the intrinsic and extrinsic connectives are interpreted over richer, more sophisticated and representative models of knowledge.

Acknowledgments. The authors would like to thank David J. Pym for many thoughtful discussions about proof-theoretic semantics and computational logic that influenced this work, and to Michael Rawson for his input in the formulation of the ideas presented herein.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Baral, C., Gelfond, M.: Logic Programming and Knowledge Representation. The Journal of Logic Programming **19-20**, 73–148 (1994), special Issue: Ten Years of Logic Programming
2. Dummett, M.: The Logical Basis of Metaphysics. Harvard University Press (1991)
3. Gheorghiu, A.V.: On an inferential semantics for intuitionistic sentential logic. Journal of Logic and Computation **35**(4), exaf029 (June 2025). <https://doi.org/10.1093/logcom/exaf029>, <https://doi.org/10.1093/logcom/exaf029>
4. Gheorghiu, A.V.: Proof-theoretic Semantics for First-order Logic. Logic Journal of the IGPL **33**(5) (2025)
5. Gheorghiu, A.V., Gu, T., Pym, D.J.: Proof-theoretic Semantics for Intuitionistic Multiplicative Linear Logic. Studia Logica (2024). <https://doi.org/10.1007/s11225-024-10158-6>

6. Harland, J.: On Hereditary Harrop Formulae as a Basis for Logic Programming. Ph.D. thesis, University of Edinburgh (1991)
7. Hintikka, K.J.J.: Knowledge and Belief: An Introduction to the Logic of the Two Notions. Cornell University Press (1962)
8. Kowalski, R.: Predicate Logic as a Programming Language. In: IFIP Congress. pp. 569–574 (1974)
9. Kowalski, R.: Logic for Problem Solving. North-Holland Publishing Co. (1986)
10. Kowalski, R.: Logic Without Model Theory. In: What Is a Logical System?, pp. 35–71. Oxford University Press (1994)
11. Miller, D.: A Logical Analysis of Modules in Logic Programming. The Journal of Logic Programming **6**(1-2), 79–108 (1989)
12. Miller, D., Nadathur, G., Pfenning, F., Scedrov, A.: Uniform Proofs as a Foundation for Logic Programming. Annals of Pure and Applied Logic **51**, 125–157 (1991)
13. Mints, G.: Gentzen-type systems and resolution rules Part I: Propositional logic. In: Martin-Löf, P., Mints, G. (eds.) International Conference on Computer Logic — COLOG, pp. 198–231. Springer (1988). https://doi.org/10.1007/3-540-52335-9_55
14. Piecha, T., Schroeder-Heister, P.: The Definitional View of Atomic Systems in Proof-theoretic Semantics. In: The Logica Yearbook 2016, pp. 185–200. College Publications London (2017)
15. Pollock, J.L., Cruz, J.: Contemporary Theories of Knowledge, vol. 35. Bloomsbury Publishing (1999)
16. Punčochář, V.: A generalization of inquisitive semantics. Journal of Philosophical Logic **45**(4), 399–428 (2016). <https://doi.org/10.1007/s10992-015-9379-1>
17. Sandqvist, T.: Classical Logic without Bivalence. Analysis **69**(2), 211–218 (2009)
18. Sandqvist, T.: Base-extension Semantics for Intuitionistic Sentential Logic. Logic Journal of the IGPL **23**(5), 719–731 (2015). <https://doi.org/10.1093/jigpal/jzv021>
19. Schroeder-Heister, P.: A natural extension of natural deduction. The Journal of Symbolic Logic **49**(4), 1284–1300 (1984)
20. Schroeder-Heister, P.: Proof-theoretic Semantics. In: Zalta, E.N., Nodelman, U. (eds.) The Stanford Encyclopedia of Philosophy. Metaphysics Research Lab, Stanford University, Summer 2024 edn. (2024)
21. Schroeder-Heister, P., Piecha, T.: Atomic Systems in Proof-theoretic Semantics: Two Approaches. In: Ángel Nepomuceno Fernández, Martins, O.P., Redmond, J. (eds.) Epistemology, Knowledge and the Impact of Interaction. Springer Verlag (2016)
22. Stafford, W.: Proof-theoretic Semantics and Inquisitive Logic. Journal of Philosophical Logic (2021)
23. Stafford, W., Piecha, T., Schroeder-Heister, P.: Logics of proof-theoretic semantics. Topoi (2025), *In press*
24. Szabo, M. (ed.): The Collected Papers of Gerhard Gentzen. Studies in Logic and Foundations of Mathematics, College Publications (1969)
25. Tarski, A.: O pojęciu wynikania logicznego. Przegląd Filozoficzny **39** (1936)
26. Tarski, A.: On the Concept of Following Logically. History and Philosophy of Logic **23**(3), 155–196 (2002). <https://doi.org/10.1080/0144534021000036683>