SJÄLVSTÄNDIGA ARBETEN I MATEMATIK

MATEMATISKA INSTITUTIONEN, STOCKHOLMS UNIVERSITET

IpC^2 as a Foundation of Mathematics

av

Mattias Wikström

2016 - No 8

MATEMATISKA INSTITUTIONEN, STOCKHOLMS UNIVERSITET, 106 91 STOCKHOLM

IpC^2 as a Foundation of Mathematics

Mattias Wikström

Självständigt arbete i matematik 15 högskolepoäng, grundnivå

Handledare: Peter LeFanu Lumsdaine

2016

Abstract

This paper discusses quantified intuitionistic propositional logic (IpC^2) and suggests that it may be able to serve as a simple and yet powerful foundation of mathematics. The logic is understood topologically, as a theory for reasoning about parts of objects, and it is shown how it has the expressive power for saying how the parts of an object with finitely many parts are structured. It is shown how a conventional first-order theory (whose logic may be classical logic, intuitionistic logic, or minimal logic) for reasoning about parthood can be translated into IpC^2. The paper also shows how IpC^2 allows us to define a description operator, further highlighting the power of IpC^2, and it is shown how the operator in question is related to well-known definitions of conjunctions, disjunctions, and the existential quantifier out of implication and the universal quantifier. The paper suggests three ways in which IpC^2 may be extended with existence axioms, a topic that matters for any foundation of mathematics. The existence axioms in question turn out to be related to three different fragments of IpC^2 which are also discussed in the paper, fragments where quantifiers are restricted from above and/or below.

Contents

1 Introduction

Second-order intuitionistic propositional logic IpC^2 extends ordinary intuitionistic propositional logic by allowing quantification over propositions. Rather remarkably, as was noted in Prawitz (1965)[1], no primitives other than universal quantifiers and implication are needed in this system: Conjunction, disjunction, and existential quantifiers can all be defined out of universal quantifiers and implication. This does perhaps give a hint that IpC^2 is very powerful in spite of its simplicity.

When combined with the Curry-Howard isomorphism (see (Howard, 1980) and (Girard, 1989, p. 93)), IpC^2 gives rise to what is known as 'System F,' which was independently introduced by Girard (1972) and Reynolds (1974). In accordance with the Curry-Howard isomorphism, System F treats any proposition P as a type (in the type-theoretic sense). Specifically, the proposition P is treated as the type of proofs of P. A proof of $P \to Q$ is a function that maps any proof of P into a proof of Q while a proof of $\forall X.P_X$ is a function which maps any type X into a proof of P_X. And since the functions involved here are the kind of functions that you can execute on a computer, System F may also be regarded as a typed programming language.

But IpC^2 remains an expressive system even when we do not distinguish between different proofs of the same proposition but merely ask what is provable and what is not. Instead of encoding things (truth values, numbers, lists, and so on) as proofs we may encode them as propositions. What is the expressive power of IpC^2 when used in this way? Löb (1976) observes that '[s]ome syntactically simple fragments of intuitionistic logic possess considerable expressive power compared with their classical counterparts' and then goes on to show how classical first-order logic with identity may be embedded in IpC^2. Gabbay (1981) describes IpC^2 as 'a system which is essentially as strong as HPC [the Heyting Predicate Calculus]' (p. 4) and demonstrates the power of IpC^2 by showing how the (undecidable) classical theory of a reflexive and symmetric relation may be faithfully interpreted in IpC^2. The reader is referred to (Sørensen and Urzyczyn, 2010) for further information on the translations that have been set up.

But if IpC^2 has passed the strength test then it would seem that it has earned the right to further tests. Is IpC^2 good or bad as a foundation of mathematics? Can it be used in practice in place of more well-known theories (such as first-order predicate logic) for the formalization of mathematical theories? If so, what are the pros and cons of using it in this way? This is what the present paper aims to explore.

So how can one formalize things in IpC^2? To begin with, while the objects of IpC^2 are

[1] See also (Scott, 1979, p. 692) and (Ruitenburg, 1991, p. 18).

customarily called "propositions," we will use them for any objects whatever. For example, a set or a natural number could end up being a "proposition." Along with this reinterpretation of propositions we will also use a reinterpretation of $a \to b$: We will often think of this as saying that b is *part of* a (or that a *contains* b; we will use $b \leftarrow a$ as a synonym for $a \to b$). While this reading of $\to$ may seem novel, it is actually well-known that intuitionistic logic admits a topological interpretation.

Although we will consider topological models in section 10, for the most part we will not be formal about the connection with topology. We will also be informal about the connection with mereology[2], the branch of philosophy that deals with wholes and parts, but these connections will nevertheless play a role in motivating what we are doing as well as in motivating the terminology we are using.

Reading $a \to b$ as saying that b is a part of a has the consequence that $\top$ is a part of everything. The reader may find this odd: Is it not $\bot$ that is conventionally part of everything? Well, note that if we accept $a \wedge b$ then a and b are part of what we accept[3]. More generally, when we accept a proposition then all its consequences are part of what we accept. It thus looks as if the logical consequences of a proposition – including $\top$ – are part of the proposition.

The reader may object at this point that $\bot$ plays an important role in intuitionistic logic: It is what allows us to negate propositions. Without negations our ability to form predicates seems severely limited. But it is possible to take a propositional constant $\mathcal{R}$ and agree that by convention $\mathcal{R}$ is an absurd proposition so that $x \to \mathcal{R}$ expresses the fact that x is not true. This is precisely what is done in minimal logic[4], and it will play an important role in the present paper.

The reader may still feel, though, that the convention that $\bot$ is part of everything is a natural one that has its uses. Indeed, even with the present approach, as we ask the question 'Which parts x of object o satisfy the predicate $\phi(x)$?' we find that $\phi(x) \leftrightarrow_{def.} \top$ corresponds to 'everything' rather than 'nothing.' Even as we have accepted that $a \to b$ expresses the fact that b is part of a, the fact that $a \to b$ is equivalent to $\forall x.(b \to x) \to (a \to x)$ means that it can sometimes seem as if what stands to the left of $\to$ is the part rather than the whole.[5]

[2]See (Varzi, 2016).

[3]Note also that if a proof of $a \wedge b$ is an ordered pair whose first component is a proof of a and whose second component is a proof of b (in accordance with the Brouwer–Heyting–Kolmogorov interpretation of intutionistic logic) then proofs of a and b are part of what we have when we have a proof of $a \wedge b$.

[4]See (Johansson, 1937).

[5]And since $(a \vee b) \to x$ is equivalent to $(a \to x) \wedge (b \to x)$, it can sometimes seem as if it is $\wedge$ rather than $\vee$ that expresses 'the common part of a and b.' Similarly, we have that $(a \wedge b) \to x$ is implied by (but not in general equivalent to) $(a \to x) \vee (b \to x)$. We can also substitute a particular constant, say $\mathcal{R}$, for x and observe the same phenomenon: $(a \vee b) \to \mathcal{R}$ is equivalent to $(a \to \mathcal{R}) \wedge (b \to \mathcal{R})$. Above it was observed

Consider now the following problem: If $a \to b$ says that b is part of a, how do we say that b is *not* part of a? The 'obvious' solution is to use $(a \to b) \to \bot$, but it follows from this that $b \leftrightarrow \bot$, and that is hardly what we want. Instead, in order to say that b is not part of a we will use the formula $(a \to b) \to \mathcal{R}$, where $\mathcal{R}$ is as described above.

As it turns out, there are some limitations on how we can use $\mathcal{R}$ to express the fact that one object is not part of another. In particular, we cannot use $\mathcal{R}$ to reason about the parts of $\mathcal{R}$. However, if for some reason we want to reason about the parts of $\mathcal{R}$ then we can do so by adding another object $\mathcal{R}'$ which we use just like $\mathcal{R}$ but which unlike $\mathcal{R}$ allows us to reason about the parts of $\mathcal{R}$. The bottom line is that in practice we do have a way to say that one object is not part of another.

Suppose now that we want to say what the parts of some particular object a are. It seems we can do so by combining three things: Formulas of the form 'x is part of y', formulas of the form 'x is not part of y,' and universal quantification. As a simple example, we can say that:

1) for all x, x is part of a if and only if $x = a$, $x = \top$, $x = p_1$, or $x = p_2$, and

2) a is not part of p_1, a is not part of p_2, p_1 is not part of $\top$, and p_2 is not part of $\top$.

But how can we be sure that our formulas really say what we think they say? What we will do is to define a semantics for IpC^2 in terms of Heyting algebras which possess infima of a certain specified type (see section 7), and we will then be able to prove precise theorems on what our formulas say for such Heyting algebras (see section 12).

There is a surprise of sorts, though, in that the parts of an object always form a Heyting algebra. Although some philosophers have made the even stronger assumption that the parts of any object form a Boolean algebra,[6] we want to avoid making some controversial philosophical assumption about parthood.

However, there is a connection between Heyting algebras and partially ordered sets (posets) which makes any finite Heyting algebra correspond to a finite poset and vice versa: The upward-closed subsets of a poset form a Heyting algebra and the meet irreducible elements of a Heyting algebra form a poset, and in the finite case these two operations are each other's inverses. By only looking at meet irreducible elements we can thus get a 'poset

that if we accept $a \wedge b$ then a and b are part of what we accept, but with *rejections* it is the other way around: If we reject $a \vee b$ (perhaps by asserting $(a \vee b) \to \mathcal{R}$) then a and b are part of what we reject. Setting $r(x) =_{def.} x \to \mathcal{R}$ we find that $x_1 \to x_2$ ('x_1 contains x_2') implies $r(x_1) \leftarrow r(x_2)$ ('$r(x_1)$ is part of $r(x_2)$') and that r has a 'reversing' effect on the arrow. To sum this up: If an acceptance of x_2 is contained in an acceptance of x_1 then a rejection of x_1 is contained in a rejection of x_2.

[6] See (Varzi, 2016).

view' of objects, and suddenly it seems that the parts of an object can form any poset, at least in the finite case. This is discussed in section 13.

At this point, reasoning in IpC^2 may seem rather strange and unusual: We negate formulas using $\mathcal{R}$ and we look at meet irreducible elements when we want a 'poset view' on things. We will, however, hide away these technical details by setting up a translation from a certain first-order language into IpC^2. The language will have a single relation symbol that expresses parthood, and it will have one sort that gives us a 'poset view' of things and another sort that gives us a 'Heyting algebra view' of things. See section 13 for the details.

At this point the question arises of how we can understand mathematical structures such as sets, vector spaces, topological spaces, Hilbert spaces, and so on in terms of our basic theory for reasoning about objects and their parts. Can we do things in such a way that the substructures of a mathematical structure become literal parts of it? Can we do things in such a way that the elements of a set become parts of the set? The present paper will stop short of actually addressing these questions. It is hoped, though, that it makes these questions seem worth exploring.

Another question that arises is what the 'universe' we are quantifying over in IpC^2 should look like? We clearly do not want for $\forall x.x = \top$ to hold (which would mean that exactly one object exists, trivializing the whole theory), but what should we assume instead concerning what we are quantifying over?

Section 14 of this paper will propose three alternative axiom schemata, each of which seems to ensure that lots and lots of objects exist. Informally, one may think of the schemata as vindicating the idea that whenever an axiom system is free from contradictions then objects of the kind described by the axioms actually exist. In fact, the axiom schemata take this idea a step further by using conservativity rather than freedom of contradictions as the criterion of existence. See section 14 for the details.

2 Axioms for IpC^2

Some minor variations exist in the way that IpC^2 is formalized, and I will use a formalization where we distinguish between 'propositional constants' (or 'atomic propositions') and 'propositional variables.' The distinction is meant to be analogous to that between constants and variables in first-order predicate logic, something that seems quite appropriate when our goal is to use IpC^2 where first-order predicate logic has traditionally been used. Just as in first-order logic, variables may either occur free or be bound by quantifiers. Constants, by contrast, cannot be bound by quantifiers, but a practical application may instead add assumptions/axioms that limit what constants can stand for. This is exactly analogous to the way that a first-order theory may contain not only constants but also axioms that involve those constants. In fact, just as one speaks of theories in first-order logic, so we will speak of theories in IpC^2.

The syntax for the formation of formulas in our system is given by the following BNF grammar[7]:

variable := *(any of the letters u, v, w, x, y, z, possibly decorated with an index or with primes)*
constant := *(any other roman letter, possibly decorated with an index or with primes)*
formula := *variable* | *constant* | (∀ *variable*.(*formula*)) | (*formula* → *formula*)

Moreover:

- We will follow the usual conventions for the omission of parentheses. In particular, instead of (∀ *variable*.(*formula*)) we may write (∀ *variable.formula*) or ∀ *variable.formula*. Implication ('→') associates to the right, and $a \to b \to c$ is thus another way of writing $a \to (b \to c)$.
- Instead of $\forall x_1.\forall x_2...., \forall x_n.\phi_{x_1,x_2,..,x_n}$ we may write $\forall x_1, x_2, ..., x_n.\phi_{x_1,x_2,..,x_n}$.
- Uppercase and lowercase letters are both acceptable as constants (and a is not the same constant as A).
- Greek letters will be used as metavariables to stand for formulas. For example, in $\phi \to \psi$ we may set $\phi = a \to b$ and $\psi = c$ to obtain $(a \to b) \to c$.
- In practice we use variable and constant symbols in a *schematic* way. For example, we may say that for any variable x we can bind x in a formula through the universal quantifier, and we may speak in a general way about a theory T with constants c_1, ..., c_n. Moreover, when using symbols in this schematic way (as we tend to do all the time in practice) we will not always follow the above conventions on which letters stand for constants and which letters stand for variables. Instead, as we introduce a new symbol we make it clear what it stands for.
- Free and bound variables are defined in the usual way. $\phi[\psi/x]$ denotes the result of substituting ψ for x in ϕ. Similarly, we define $\phi[\psi_1/x_1, ..., \psi_n/x_n]$ to be the result of simultaneously substituting ψ_1 for x_1, ..., ψ_n for x_n in ϕ. Two formulas are said to be α-equivalent if they arise from each other by the renaming of bound variables.
- We will use the following definitions (see section 3 for an explanation of how the first three definitions actually work):

[7]'BNF' stands for 'Backus Normal Form' or 'Backus Naur Form' and is a widely used notation for the presentation of the syntax of programming languages as well as other formal languages.

$\phi \wedge \psi$ is defined as $\forall x.(\phi \to \psi \to x) \to x$ and referred to as 'ϕ and ψ' or 'the conjunction of ϕ and ψ,' $\phi \vee \psi$ is defined as $\forall x.(\phi \to x) \to (\psi \to x) \to x$ and referred to as 'ϕ or ψ' or 'the disjunction of ϕ and ψ,'

$\exists x.\phi$ is defined as $\forall y.(\forall x.(\phi \to y)) \to y$ and expresses 'existential quantification,'[8]

$\bot$ is defined as $\forall y.y$ and referred to as 'bottom,'

$\top$ is defined as $\bot \to \bot$ and referred to as 'top,'

$\neg\phi$ is defined as $\phi \to \bot$ and referred to as 'the negation of ϕ,'

$\phi \leftarrow \psi$ is defined as $\psi \to \phi$,

$\phi \leftrightarrow \psi$ is defined as $(\phi \to \psi) \wedge (\psi \to \phi)$ and read 'ϕ is equivalent to ψ,' and

$\phi = \psi$ is defined as $\phi \leftrightarrow \psi$.

- As suggested by the last of the above definitions, we regard equivalent propositions as 'the same'/'identical'/'equal.' We have no reason to distinguish between them, and regarding them as the same makes a difference when it comes to counting how many objects we have.

Our rules of inference will be:

I1. From ϕ and $\phi \to \psi$ we may infer ψ (*modus ponens*).

I2. If x and y are not free in ϕ then $\phi \to \forall x.\psi$ may be inferred from $\phi \to \psi[y/x]$.

We will use the following axioms (note that A3.-A4. are actually schemas with infinitely many axioms as instances):

A1. $\forall x, y.x \to y \to x$

A2. $\forall x, y, z.(x \to y \to z) \to (x \to y) \to x \to z$

A3. $(\forall x.\phi) \to \phi[y/x]$

A4. $\exists x.x \leftrightarrow \phi$ (x not free in ϕ)

That A1. and A2. suffice to axiomatize intuitionistic implication is well-known. The rule I2. and the axiom A3. tell us how we may introduce and eliminate universal quantifiers. Finally, A4. is a comprehension schema which helps determine what we are quantifying over. It plays an important role in making the system what it is. For example, without it we cannot take for granted that $\exists x.x \leftrightarrow \phi_1 \wedge \phi_2$ and $\exists x.x \leftrightarrow \phi_1 \vee \phi_2$ hold.

[8] When applying this definition, let y be a variable that is not free in ϕ.

(Gabbay, 1981, p. 159) considers an additional axiom which with the present notation may be written: $(\forall x.\phi \vee \psi) \to (\phi \vee \forall x.\psi)$ (x not free in ϕ). As the expression 'IpC^2' is understood here (and presumably elsewhere as well), this axiom is *not* included. As natural as it may seem, one can find many topological models of IpC^2 where it fails to hold.[9]

By a *theory* T in IpC^2 we mean a set Sig_T of constants along with a set of formulas Ax_T that contain no constants apart from those in Sig_T. Sig_T may be referred to as the *signature of* T.

When two theories T_1 and T_2 are such that all the constants in Sig_{T_1} are included in Sig_{T_2} and such that all the axioms in Ax_{T_1} are included in Ax_{T_2}, then we say that T_2 is an extension of T_1.

A formula ϕ is considered to be deducible/provable in a theory T if it can be deduced from the axioms of IpC^2 along with the axioms of T.

Note that we are not requiring theories to be deductively closed.

Theorem 1. (Deduction theorem) If a theory T with finitely many constants $c_1, ..., c_m$ and finitely many axioms $\chi_1, ..., \chi_{n-1}$ proves χ_n and if (for each i) ${\chi_i}^*$ is a constant-free formula such that ${\chi_i}^*[c_1/x_1, ..., c_m/x_m]$ is χ_i then IpC^2 proves $\forall x_1, ..., x_m.({\chi_1}^* \to ...{\chi_{n-1}}^* \to {\chi_n}^*)$.

Proof. Omitted. □

3 A description operator; infima

When reading the previous section, the reader may have been puzzled by the definitions of $\wedge$, $\vee$, and $\exists$ in terms of $\to$ and $\forall$. How do these definitions actually work? The present section will explain this and more.

An important feature of IpC^2 is that it allows us to define the following operator:
$\inf x \,\text{s.t.}\, \phi \leftrightarrow_{def.} \forall x.\phi \to x$ ('*s.t.*' is read as 'subject to' or 'such that'). As the notation may suggest, this operator has the property that $\inf x \,\text{s.t.}\, \phi$ is the infimum of all objects x such that ϕ holds. To be precise:

[9]Topological models are discussed in section 10. Using the interpretation of intuitionistic logic in terms of open sets in a topological space, we can get a counterexample by using the real line with the standard topology as our topological space, by letting ϕ stand for $\{t \in \mathbb{R} : t \neq 0\}$ and by letting ψ be such that it takes on exactly the values $\{t \in \mathbb{R} : |t| < r\}$, where r can be any positive real number, as x ranges over all open sets.

Theorem 2. 1) For any formula ω such that $\phi[\omega/x]$ holds, $\inf x \,\text{s.t.}\, \phi$ implies ω, and 2) if μ is any such formula (for any formula ω such that $\phi[\omega/x]$ holds, μ implies ω) then μ implies $\inf x \,\text{s.t.}\, \phi$.

Proof. To see that 1) is true, note that $\forall x.\phi \to x$ (which is what $\inf x \,\text{s.t.}\, \phi$ stands for) implies $\phi[\omega/x] \to \omega$, which implies ω under the assumption that $\phi[\omega/x]$ holds. So $\inf x \,\text{s.t.}\, \phi$ implies ω under the assumptions of 1).

To see that 2) is true, assume that for any formula ω such that $\phi[\omega/x]$ holds, μ implies ω. We want to show that μ implies $\inf x \,\text{s.t.}\, \phi$, so assume that μ holds. Because of our second assumption, our first assumption can be simplified to: For any object ω such that $\phi[\omega/x]$ holds, ω holds. But this is just another way of saying $\forall x.\phi \to x$, which is also what $\inf x \,\text{s.t.}\, \phi$ stands for. Hence μ implies $\inf x \,\text{s.t.}\, \phi$. □

It is now easy to see how the definitions of $\wedge$, $\vee$, and $\exists$ work. The definitions we used before are equivalent to:

$x_1 \wedge x_2 \leftrightarrow_{def.} \inf y \,\text{s.t.}\, (x_1 \to (x_2 \to y))$,
$x_1 \vee x_2 \leftrightarrow_{def.} \inf y \,\text{s.t.}\, ((x_1 \to y) \wedge (x_2 \to y))$, and
$\exists x.\phi \leftrightarrow_{def.} \inf y \,\text{s.t.}\, \forall x.(\phi \to y)$.

Theorem 3. 1) $\alpha \wedge \beta$ implies α as well as β,

2) $\alpha \to (\beta \to (\alpha \wedge \beta))$,

3) $\alpha \to (\alpha \vee \beta)$ and $\beta \to (\alpha \vee \beta)$,

4) $[(\alpha \to \omega) \wedge (\beta \to \omega)] \to [(\alpha \vee \beta) \to \omega]$,

5) $\psi[\chi/x] \to \exists x.\psi$, and

6) if $\chi \to \psi$ holds and if x does not occur free in ψ then $(\exists x.\chi) \to \psi$ holds.

Proof. 1) To prove that $\alpha \wedge \beta$ ($= \inf x \,\text{s.t.}\, (\alpha \to (\beta \to x))$) implies α, apply part 1) of theorem 2 with $(\alpha \to (\beta \to x))$ in place of ϕ and with α in place of ω. This requires us to show that $\alpha \to (\beta \to \alpha)$ holds, but this is trivial.

That $\alpha \wedge \beta$ implies β can be shown analogously. This time, we are required to show that $\alpha \to (\beta \to \beta)$ holds, which is trivial.

2) Assume that α holds. Then β is a formula such that if $\alpha \to (\beta \to \omega)$ holds then β implies ω (regardless of what ω is). By part 2) of theorem 2 we therefore get $\beta \to \inf x \,\text{s.t.}\, (a \to (\beta \to x))$. Without the assumption α, we end up with $\alpha \to \beta \to \inf x \,\text{s.t.}\, (\alpha \to (\beta \to x))$.

3) Apply part 2) of theorem 2 with $((\alpha \to x) \wedge (\beta \to x))$ substituted for ϕ and with μ set to α. This requires us to show that if $(\alpha \to \omega) \wedge (\beta \to \omega)$ holds then $a \to \omega$ holds, which is trivial. We end up with the conclusion that $\alpha \to \inf x \, \text{s.t.} \, ((\alpha \to x) \wedge (\beta \to x))$.

The proof of $\beta \to (\alpha \vee \beta)$ is similar. It requires us to show that if $(\alpha \to \omega) \wedge (\beta \to \omega)$ holds then $\beta \to \omega$ holds, which is trivial.

4) This is what part 1) of theorem 2 says when that theorem is applied with ϕ set to $(\alpha \to x) \wedge (\beta \to x)$.

5) Apply part 2) of theorem 2 with ϕ set to $\forall z.(\psi \to x)$ and μ set to $\psi[\chi/x]$. This requires us to show that for any formula ω such that $\forall z.(\psi \to \omega)$ holds, $\psi[\chi/x]$ implies ω. However, this is a simple matter of substituting χ for z in $\forall z.(\psi \to \omega)$.

6) Assume $\chi \to \psi$. By setting ϕ equal to $\forall z.(\chi \to x)$ and ω equal to ψ in part 1) of theorem 2, we get that $\inf x \, \text{s.t.} \, \forall z.(\chi \to x) \to \psi$. This requires us to prove that $\forall z.(\chi \to \psi)$ holds, but this is a simple consequence of the assumption $\chi \to \psi$. □

Another application of the operator $\inf x \, \text{s.t.} \, \phi$ is that we are able to *define* a description operator $\imath$. The defining characteristic of such an operator is that if there is exactly one object a such that $\phi[a/x]$ holds then $\imath x.\phi$ denotes that object. Here is one way that we may define such an operator in IpC^2:

$\imath x.\phi \leftrightarrow_{def.} \inf x \, \text{s.t.} \, \phi$

As long as there is a unique x satisfying ϕ, this definition succeeds in picking that object out, and this is all that is required of a description operator. But even when $\exists! x.\phi$ is not provable we still have $(\exists! x.\phi) \to \phi[(\imath x.\phi)/x]$ (just assume $\exists! x.\phi$ and note that $\phi[(\imath x.\phi)/x]$ becomes provable). Here, $\exists! x.\phi$ is defined as usual:

$\exists! x.\phi \leftrightarrow_{def.} (\exists x.\phi) \wedge \forall y, z.(\phi[y/x] \leftrightarrow \phi[z/x]) \to y = z$ (let y nor z be two variables that do not occur in ϕ).[10]

There is also a generalized version of the operator $\inf x \, \text{s.t.} \, \phi$ that is worth mentioning[11]:

$\inf \psi[x] \, \text{s.t.} \, \phi \leftrightarrow_{def.} \forall x.\phi \to \psi$.

The notation has been chosen to suggest that $\inf \psi[x] \, \text{s.t.} \, \phi$ can be seen as giving us the solution to an optimization problem, and the following theorem confirms this:

[10] It is worth noticing that $\inf x \, \text{s.t.} \, \phi$ implies $\exists! x.\phi$: If we assume $\inf x \, \text{s.t.} \, \phi$ then we have $\inf x \, \text{s.t.} \, \phi = \top$, which means that ϕ only holds for $\top$, from which $\exists! x.\phi$ follows. Note also that setting $\imath x.\phi \leftrightarrow_{def.} (\exists! x.\phi \to \inf x \, \text{s.t.} \, \phi)$ would work as an alternative (non-equivalent) definition of $\imath x.\phi$: We would still get $(\exists! x.\phi) \to \phi[(\imath x.\phi)/x]$.

[11] Russell referred to expressions of the form $\forall x.\phi \to \psi$ as 'formal implications' and Church gave them the notation $\phi \supset_x \psi$ (Church, 1956, p. 44). We could write this as $\phi \to_x \psi$, but we will instead use $\phi \to_x \psi$ to mean $x \wedge (\phi \to \psi)$. See section 6.

Theorem 4. 1) For any formula ω such that $\phi[\omega/x]$ holds, $\inf \psi[x] \,\text{s.t.}\, \phi$ implies $\psi[\omega/x]$, and 2) if μ is any such formula (for any object ω such that $\phi[\omega/x]$ holds, μ implies $\psi[\omega/x]$) then m implies $\inf \psi[x] \,\text{s.t.}\, \phi$.

Proof. The proof is simply a generalized version of the proof of theorem 2. $\forall x.\phi \to \psi$ (=$\inf \psi[x] \,\text{s.t.}\, \phi$) implies $\phi[\omega/x] \to \psi[\omega/x]$, which implies $\psi[\omega/x]$ under the assumption that $\phi[\omega/x]$ holds. So $\inf \psi[x] \,\text{s.t.}\, \phi$ implies $\psi[\omega/x]$ under the assumptions of 1).

To prove 2), assume that for any formula ω such that $\phi[\omega/x]$ holds, μ implies $\psi[\omega/x]$. We want to show that μ implies $\inf \psi[x] \,\text{s.t.}\, \phi$, so assume that μ holds. Because of our second assumption, our first assumption can be simplified to: For any object ω such that $\phi[\omega/x]$ holds, $\psi[\omega/x]$ holds. But this is just another way of saying $\forall x.\phi \to \psi$, which is also what $\inf \psi[x] \,\text{s.t.}\, \phi$ stands for. Hence μ implies $\inf \psi[x] \,\text{s.t.}\, \phi$.

□

Note that $\inf \psi[x] \,\text{s.t.}\, \top$ is equivalent to $\forall x.\psi$ and that $\inf \psi[x] \,\text{s.t.}\, \phi$ is equivalent to $\phi \to \psi$ in the case when x does not occur in ϕ or ψ. The two primitives of IpC^2 may thus both be defined in terms of $\inf \psi[x] \,\text{s.t.}\, \phi$ (although the definition of $\to$ is somewhat impractical in that it requires us to find a fresh variable each time we want to express $\to$).

By the following theorem, it is also possible to define the universal quantifier out of the infimum operator and the existential quantifier[12]:

Theorem 5. $\forall x.\phi \leftrightarrow \inf y \,\text{s.t.}\, \exists x.y = \phi$

[12] Could we set up an alternative axiomatization of IpC^2 where we use the infimum operator and the existential quantifier as primitives and where we do not include the universal quantifier as a primitive? There is a problem in that we used universal quantifiers to define conjunctions which were in turn used to define '='. However, if the primitives were '$\to$,' '=,' existential quantification, and the infimum operator, then we could of course use the definition $\forall x.\phi \leftrightarrow_{def.} \inf y \,\text{s.t.}\, \exists x.y = \phi$ to define universal quantifiers (the definition is to be understood as being valid only when y does not occur free in ϕ). We are clearly able to translate formulas expressed using each set of primitives into formulas expressed using the other set of primitives, but could we find an explicit set of axioms which employs the alternative set of primitives? Well, we can certainly translate our usual axioms and inference rules (see section 2) so that they come to employ the alternative set of primitives. For example, axiom A3. becomes (note that we need to add an extra clause to the axiom about z not being free in ϕ):
A3#. $((\inf z \,\text{s.t.}\, \exists x.z = \phi) \to \phi[y/x]$ (z not free in ϕ).
However, this will not quite suffice since we we can only prove formulas where the infimum operator and the existential quantifier occur together (for example, we would not be able to prove $\exists x.x$). Things will work, though, if we turn our old definitions of '=,' the existential quantifier, and the infimum operator into axioms. For example, we would add the following axioms to ensure that $\exists x.\phi$ is interderivable with $\forall y.(\forall x.\phi \to y) \to y$:
A∃1. $\exists x.\phi \to [\forall y.(\forall x.\phi \to y) \to y]$, and
A∃2. $[\forall y.(\forall x.\phi \to y) \to y] \to \exists x.\phi$.

Proof. Assume that $\forall x.\phi$ is provable. Then ϕ is also provable as is $\phi = \top$. Hence $\exists x.y = \phi$ can be rewritten as $\exists x.y = \top$ or just $y = \top$. $\inf y \,\text{s.t.}\, \exists x.y = \phi$ therefore becomes $\inf y \,\text{s.t.}\, y = \top$, which is obviously $\top$. To prove the other direction, assume instead that $\inf y \,\text{s.t.}\, \exists x.y = \phi$ holds. This is by definition the same thing as $\forall y.(\exists x.y = \phi) \to y$. By the law that $(\exists x.\chi) \to \psi$ is equivalent to $\forall x.\chi \to \psi$ (which is related to the rule that $(a \vee b) \to c$ is equivalent to $(a \to c) \wedge (b \to c)$), this may be rewritten as $\forall y.\forall x.(y = \phi) \to y$. By swapping the quantifiers we get $\forall x.\forall y.(y = \phi) \to y$, which is evidently equivalent to $\forall x.\phi$. □

Let us finally note that it is possible to define a *supremum* operator that is 'dual' to the infimum operator (let y be a variable that does not occur free in ϕ or ψ):

$\sup \psi[x] \,\text{s.t.}\, \phi \leftrightarrow_{def.} \exists y.\phi[y/x] \wedge \psi[y/x]$, and
$\sup x \,\text{s.t.}\, \phi \leftrightarrow_{def.} \sup x[x] \,\text{s.t.}\, \phi$.

We then have:

Theorem 6. 1) For any formula ω such that $\phi[\omega/x]$ holds, $\psi[\omega/x]$ implies $\sup \psi[x] \,\text{s.t.}\, \phi$, and 2) if μ is any such formula (for any object ω such that $\phi[\omega/x]$ holds, $\psi[\omega/x]$ implies μ) then $\sup \psi[x] \,\text{s.t.}\, \phi$ implies μ.

Proof. To prove 1), let ω be a formula such that $\phi[\omega/x]$. Assume $\psi[\omega/x]$. From this we can clearly infer $\exists y.\phi[y/x] \wedge \psi[y/x]$, which is what $\sup \psi[x] \,\text{s.t.}\, \phi$ says. Hence $\psi[\omega/x]$ implies $\sup \psi[x] \,\text{s.t.}\, \phi$.

To prove 2), let μ be such that for any formula ω such that $\phi[\omega/x]$ holds, $\psi[\omega/x]$ implies μ. Substitute y for ω to get $\forall y.\phi[y/x] \to (\psi[y/x] \to \mu)$. Now assume $\sup \psi[x] \,\text{s.t.}\, \phi$. That is, we are assuming $\exists y.\phi[y/x] \wedge \psi[y/x]$, which we may also write (using the definition of the existential quantifier) as $\forall z.[\forall y.\phi[y/x] \to (\psi[y/x] \to z)] \to z$. Substitute μ for z to get $[\forall y.\phi[y/x] \to (\psi[y/x] \to \mu)] \to \mu$. By Modus Ponens we get μ. Hence we have that $\sup \psi[x] \,\text{s.t.}\, \phi$ implies μ. □

4 Theories that only employ restricted quantification; the logic $\mathbf{IpC^2}\!\restriction$

As was mentioned in the introduction, we will often think of implication as expressing parthood. When doing so, we think of $a \to b$ (which we may equally well write $b \leftarrow a$) as saying that b is part of a. This means that $\top$ is part of everything ($\forall x.\top \leftarrow x$) while $\bot$ contains everything as a part ($\forall x.\bot \to x$). It also means that we may want to define versions of the quantifiers which only quantify over the parts of a certain object. Let us thus make the following definitions:

$\forall x \leftarrow \chi.\phi$ is defined as $\forall x.(x \leftarrow \chi) \rightarrow \phi$ (where χ is an expression where x does not occur)

$\exists x \leftarrow \chi.\phi$ is defined as $\exists x.(x \leftarrow \chi) \wedge \phi$ (where χ is an expression where x does not occur)

Note that restricted quantification becomes equivalent to unrestricted quantification in the special case where $\chi = \bot$; by using $\bot$ as our domain of quantification we are quantifying over everything.

Let us call a formula *restricted* if it is built up from $\rightarrow$, $\wedge$, $\vee$, $\top$, restricted quantifiers ($\forall x \leftarrow \chi.\phi$ and $\exists x \leftarrow \chi.\phi$), constants, and variables, or if is equivalent to such a formula (note the exclusion of $\bot$ from the list). Let us also call a theory T in IpC^2 *restricted* if its axioms are all restricted (or equivalent to restricted formulas). Finally, let us say that a theory T' which extends a theory T is a *conservative extension of T with respect to restricted formulas* if a restricted formula of T is provable in T' if and only if it is provable in T.

When dealing with restricted formulas, the following theorem can be useful:

Theorem 7. If a formula ϕ is such that for some constant c every constant c_i in ϕ satisfies $c_i \leftarrow c$ then $\phi \leftarrow c$ holds.

Proof. Assume c. Under this assumption we can prove $c = \top$ as well as $c_i = \top$ for each of the constants. Now use structural induction on the formula ϕ to see that $\phi = \top$ has to be provable. Since $\phi = \top$ is provable from the assumption c, it follows that $c \rightarrow \phi$ holds. □

Note that the inclusion of $\vee$, $\top$, and $\exists x \leftarrow \chi.\phi$ in the definition of restricted formulas is actually unnecessary since we have $a \vee b \leftrightarrow [\forall x \leftarrow (a \wedge b).(a \rightarrow x) \rightarrow (b \rightarrow x) \rightarrow x]$, $\top \leftrightarrow (\phi \rightarrow \phi)$ (where ϕ is an arbitrary formula), and $\exists x \leftarrow \chi.\phi \leftrightarrow [\forall x \leftarrow \chi.(\phi \rightarrow x) \rightarrow x]$. We could also dispense with $\wedge$ at the cost of rewriting our theories (just add an extra constant u along with the axiom $u \rightarrow c_i$ for each existing constant c_i and then use $\forall x \leftarrow u.a \rightarrow b \rightarrow x$ in place of $a \wedge b$).

We may go even further and define a logical system $\mathrm{IpC}^2\!\restriction$ (read this as 'IpC^2 with restricted quantification' or 'IpC^2 restricted') which is like IpC^2 except that no unrestricted quantification is ever allowed.[13] The expression $\forall x \leftarrow \chi.\phi$ is thus no longer to be seen

[13] It could be argued that $\mathrm{IpC}^2\!\restriction$ ought to be described as a variant of 'minimal logic' rather than 'intuitionistic logic' since we do not have an object $\bot$ such that $\bot \rightarrow \phi$ holds regardless of what ϕ happens to be. But the difference between minimal logic and intuitionistic logic is slight in any case, and the notation $\mathrm{IpC}^2\!\restriction$ emphasizes the fact that the system is closely related to IpC^2.

as an abbreviation of $\forall x.(x \leftarrow \chi) \to \phi$, but restricted quantification has instead replaced unrestricted quantification as a basic primitive.

In order to avoid unnecessary problems I will define $\text{IpC}^2\!\restriction$ in such a way that it includes $\wedge$ as a primitive. We saw above how it could be dispensed with in restricted theories, but the rearranging of theories that is required in order to make this work is not very convenient. As axioms for $\wedge$ we may use:

$(\phi \wedge \psi) \to \phi$,
$(\phi \wedge \psi) \to \psi$, and
$\phi \to \psi \to (\phi \wedge \psi)$.

I will also add $\top$ as a primitive (although one may prefer to define it as $c \to c$ for an arbitrary constant c) along with the following axiom:

$\phi \to \top$.

$\vee$ and $\exists$ may be defined as suggested above:

$a \vee b \leftrightarrow_{def.} [\forall x \leftarrow (a \wedge b).(a \to x) \to (b \to x) \to x]$
$\exists x \leftarrow \chi.\phi \leftrightarrow_{def.} [\forall x \leftarrow \chi.(\phi \to x) \to x]$

The axioms A1. and A2. that we used for IpC^2 work equally well with $\text{IpC}^2\!\restriction$. Axioms A3.-A4. and the inference rule I2. need some changes, though (as before, $a \leftarrow b$ is to be understood as another way of writing $b \to a$):

I2'. $\phi \to [(y \leftarrow \chi) \to \psi[y/x]]$ implies $\phi \to \forall x \leftarrow \chi.\psi$ (x and y not free in ϕ)

A3'. $(\forall x \leftarrow \chi.\phi) \to (y \leftarrow \chi) \to \phi[y/x]$

A4'. $\exists x \leftarrow \chi.x \leftrightarrow [(y \leftarrow \chi) \to \phi]$ (x not free in ϕ)

Let us also make sure we have restricted versions of the definitions of section 3:

$\inf \psi[x \leftarrow \chi]\,\text{s.t.}\,\phi \leftrightarrow_{def.} \forall x \leftarrow \chi.\phi \to \psi$
$\imath x \leftarrow \chi.\phi \leftrightarrow_{def.} \inf x[x \leftarrow \chi]\,\text{s.t.}\,\phi$

Note that a theory in $\text{IpC}^2\!\restriction$ is a restricted theory in IpC^2 and that instead of saying 'conservative extension of T with respect to restricted formulas' we can simply say 'conservative extension of T' as long as the logic is $\text{IpC}^2\!\restriction$.

5 An example of a theory that describes the parts of an object

Let us now look at an example of how a restricted theory may be used to describe the parts of an object. We will consider a theory S with four constants a, b, c, and d along with the following axioms (when reading these axioms, keep in mind that we are treating equivalent propositions as identical; see Section 2):

$b \leftarrow a$ ('b is part of a')

$c \leftarrow b$ ('c is part of b')

$d \leftarrow b$ ('d is part of b')

$\forall x \leftarrow a.(x \leftrightarrow a) \vee (x \leftarrow b)$ ('any part of a is either a itself or part of b')

$\forall x \leftarrow b.(x \leftrightarrow b) \vee (x \leftarrow c) \vee (x \leftarrow d)$ ('any part of b is either b itself, part of c, or part of d')

$\forall x \leftarrow c.(x \leftrightarrow c) \vee (x \leftrightarrow \top)$ ('any part of c is either c itself or $\top$')

$\forall x \leftarrow d.(x \leftrightarrow d) \vee (x \leftrightarrow \top)$ ('any part of d is either d itself or $\top$')

In spite of being a restricted theory, theory S seems able to tell us a lot about the parts of its objects. It cannot say anything about what exists outside those constants, though (it would no longer be a restricted theory if it did).

Note that S is consistent with a, b, c, and d all being $\top$. In fact, theorem 7 shows that any restricted theory is consistent with all constants being $\top$.

Some other possibilities are: 1) $a = \bot, b = c = d = \top$, 2) $a = b = \bot, c = d = \top$, 3) $a = b = c = \bot, d = \top$, 4) $a = b = d = \bot, c = \top$, and 5) $a = b = c = d = \bot$. We can get more possibilities by not limiting ourselves to $\top$ and $\bot$, but to properly reason about such 'possibilities' we had better give a formal definition of a 'model' of IpC^2. We will look at this in section 7.

Models will also help us address doubts as to whether the axioms of S really succeed in saying what they have here been taken to say. For example, how do we know that it is correct to think of $\forall x \leftarrow a.(x \leftrightarrow a) \vee (x \leftarrow b)$ as saying 'any part of a is either a itself or part of b'? Section 12 tries to address such doubts through some precise theorems.

6 A 'dual' way of restricting formulas of IpC^2; the logic $\mathrm{IpC}^2{}_{\equiv}$

In section 4 we ended up with a 'restricted' version of IpC^2 in which $\bot$ was forbidden from appearing in any formula. We will now look at a 'dual' restriction where it is instead '$\top$'

that is forbidden from appearing in any formula.

Since both $\exists x.x$ and $\phi \to \phi$ are equivalent to $\top$, we must evidently do more than simply forbidding '$\top$' from appearing in any formula. As for quantifiers, while section 4 considered quantifiers that were restricted 'from above,' we will now instead consider quantifiers restricted 'from below':

$\forall \rho \leftarrow x.\phi \leftrightarrow_{def.} \forall x.(\rho \leftarrow x) \to \phi$, and
$\exists \rho \leftarrow x.\phi \leftrightarrow_{def.} \exists x.(\rho \leftarrow x) \wedge \phi$.

As for '$\to$,' its role will be taken by a connective that I will denote '$\to_z$,' which can be defined as follows in ordinary IpC^2:

$x \to_z y \leftrightarrow_{def.} [x \to y]_z$, where the operation '$[-]_z$' is simply:
$[x]_z =_{def.} z \wedge x$.

I will also make the following, closely related definition:

$x =_z y \leftrightarrow_{def.} [x = y]_z$
Note that $x =_z y$ is equivalent to $(x \to_z y) \wedge (y \to_z x)$.

Of course, $x \to y$ is equivalent to $x \to_\top y$. Thus, in the presence of '$\top$' and '$\wedge$' we have that the connective '$\bullet \to_- \star$' is interdefinable with the connective '$\bullet \to \star$.'

We may now consider the fragment of IpC^2 consisting of formulas which are either built up using the following constructs or provably equivalent to such formulas[14]:

- $\to_z$,
- $\forall \rho \leftarrow x.\phi$,
- $\exists \rho \leftarrow x.\phi$,
- $\wedge$,
- $\vee$, and
- $\bot$.

This gives us a sublanguage of our original language, and the sublanguage inherits the entailment relation of the original language. I will denote the resulting logic '$\text{IpC}^2{}_{\equiv}$' and I will write entailments in the usual way using a turnstile symbol ('$\vdash$').

[14] By structural induction on formulas one can easily prove that this subset of IpC^2 is closed under substitution: If ϕ is in this fragment, so is $\phi[\psi/x]$.

Note that with $\mathrm{IpC}^2\equiv$ we can never have $\vdash \psi$ for any formula ψ in $\mathrm{IpC}^2\equiv$ (as that would make ϕ a synonym for $\top$). Instead, what $\mathrm{IpC}^2\equiv$ gives us are always entailments, expressions of the form '$\phi \vdash \chi$.'

In order to make $\mathrm{IpC}^2\equiv$ look more like a typical logic, we can define '$\Gamma \vdash_\psi \chi$' to mean the same thing as '$\Gamma, \psi \vdash \chi$' (where $\Gamma = \gamma_1, ..., \gamma_n$ is a sequence of 0 or more formulas).

We then get the following rule (the verification is trivial):
$\Gamma, \phi \vdash_\psi \chi$ is equivalent to $\Gamma \vdash_\psi \phi \to_\psi \chi$.
Note that this allows us to rewrite '$\gamma_1, ..., \gamma_n \vdash_\psi$' as '$\vdash_\psi \gamma_1 \to ... \to \gamma_n \to \phi$.'
Other rules involving '$\vdash_\psi$' are (the verifications are again trivial):

(i) $\phi_1, \phi_2 \vdash_\psi \phi_1 \wedge \phi_2$,

(ii) $\phi \to_\psi \chi_1, \phi \to_\psi \chi_2 \vdash_\psi \phi \to_\psi (\chi_1 \wedge \chi_2)$,

(iii) $\phi_1 \vdash_\psi \phi_1 \vee \phi_2$,

(iv) $(\phi_1 \vee \phi_2), \phi_1 \to_\psi \chi, \phi_2 \to_\psi \chi \vdash_\psi \chi$, and

(v) $\bot \vdash_\psi \phi$.

One additional rule that may be of some interest is:
$\vdash_\psi \chi$ and $\Gamma \vdash_\chi \phi$ we may infer $\Gamma \vdash_\psi \phi$.

I will not try to give an explicit axiomatization of $\mathrm{IpC}^2\equiv$, but it looks to me as if the above rules could serve as a starting point for such an axiomatization.

Let us next consider the problem of defining 'theories' in $\mathrm{IpC}^2\equiv$. I will define a theory T in $\mathrm{IpC}^2\equiv$ to be a theory in IpC^2 where all axioms have (or can be rewritten in) the form $\phi \to \chi$, where the formulas ϕ and χ are formulas in $\mathrm{IpC}^2\equiv$.[15] I will refer to an IpC^2 formula of this form as an '$\mathrm{IpC}^2\equiv$ implication.' Since $\mathrm{IpC}^2\equiv$ implications are formulas in IpC^2, we can ask which ones follow from which (in the logic IpC^2), and we may ask, in particular, which $\mathrm{IpC}^2\equiv$ implications follow from the theory T.

In a particular application of $\mathrm{IpC}^2\equiv$, one could decide to think of some particular formula ψ as 'true' even though the logic will not allow us to add ψ as an axiom of any theory. One would then conclude from $\vdash_\psi \chi$ that χ is true, and one would conclude from $\phi \vdash_\psi \chi$ that if ϕ is true, so is χ.

[15] We can think of $\phi \to \chi$ as expressing what we expressed as $\vdash_\phi \chi$ above. The advantage of using $\to$ rather than the notation $\Gamma \vdash_\psi \chi$ is that we are able to stay within IpC^2. While it seems impossible to use formulas of $\mathrm{IpC}^2\equiv$ as axioms, at least we do not have to use anything other than formulas of IpC^2.

In parthood terms, we may decide (in a particular application of IpC$^2{\equiv}$) to think of some particular formula ψ as 'part of everything,' even though IpC$^2{\equiv}$ will not allow us to add ψ as an axiom. We would then conclude from $\phi \vdash_\psi \chi$ that χ is part of ϕ.

To make this more formal, we need only remember that IpC$^2{\equiv}$ can be seen as a fragment of IpC^2 and that the latter allows us to say that ψ is 'true'/'part of everything.' In effect, we can make sense of what we are doing in IpC$^2{\equiv}$ by regarding our formulas as formulas of IpC^2 and by treating ψ as an axiom.

But note that as far as IpC$^2{\equiv}$ is concerned there is nothing special about the formula ψ. Instead, we are always free to take any formula and regard it as what is 'true'/'part of everything.' One is reminded of relativity theory and the principle that any observer who is falling freely can be regarded as unmoving.

Let us finally note that we may look at the set of formulas of IpC^2 that are included in *both* IpC$^2{\upharpoonright}$ and IpC$^2{\equiv}$. We may denote this logic 'IpC$^2{\upharpoonright}{\equiv}$.'

7 Different Types of Semantics of IpC^2

In order to better understand theories of IpC^2, we may use models of one sort or another. Since IpC^2 is an extension of quantifier-free intuitionistic propositional logic, any semantics for IpC^2 must also be a semantics for intuitionistic propositional logic. On top of that, it also needs to have a semantics for the universal quantifier.

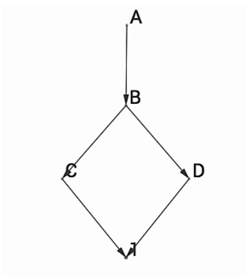

Figure 1: A Heyting algebra model of the theory of section 5.

Let us look at the basic idea before we introduce models in full generality. In the case of the theory S of section 5 we could model $\top$, a, b, c, and d as elements 1, A, B, C, and D in a partially ordered set with exactly these elements. We stipulate that the order $\geq$ is to be the relation $\{(1,1),(1,C),(C,C),\ (1,D),(D,D),(1,B),\ (C,B),(D,B),(B,B),$

$(1, A), (C, A), (D, A), (B, A), (A, A)\}$ (this is illustrated in the figure above). We formally define a function $[\![-]\!]$ from the constants of S to their interpretations: $[\![\top]\!] = 1$, $[\![a]\!] = A$, $[\![b]\!] = B$, $[\![c]\!] = C$, and $[\![d]\!] = D$. We can then extend $[\![-]\!]$ to arbitrary expressions of S through rules such as $[\![x \to y]\!] = [\![x]\!] \Rightarrow [\![y]\!]$ and $[\![x \vee y]\!] = [\![x]\!] \vee [\![y]\!]$ (the operations $\Rightarrow$ and $\vee$ (which make sense in posets with the right properties) will be properly introduced below). Finally, we can verify that ϕ is provable in S if and only if $[\![\phi]\!] = 1$ holds (for example, $[\![a \to a]\!] = ([\![a]\!] \Rightarrow [\![a]\!]) = (A \Rightarrow A) = 1$ while $[\![\top \to a]\!] = ([\![\top]\!] \Rightarrow [\![a]\!]) = (1 \Rightarrow A) = A$).

Instead of modeling the objects of IpC^2 as elements of a poset, we may also try other objects. Since we are regarding IpC^2 as a theory for reasoning about parts, and since one often uses topology to reason about the parts of objects, why not try to model the objects of IpC^2 as closed sets in a topological space (using open sets is more common but does not make $a \to b$ say that b is part of a; see the introduction)? This turns out to work, and the result is that we get *topological models* of IpC^2. These will be discussed in section 10.

A third type of semantics, frame semantics, will not be considered in this paper. One way to think of it is as a 'dual' of algebraic semantics. See (Blackburn, de Rijke, and Venema, 2001, section 5.4) for more on this topic.

Each of the three types of semantics that have been mentioned here may be seen as (important) specializations of semantics for modal logics. This is explained by the existence of a faithful translation of intuitionistic logic into S4 modal logic[16] given by the following rules (we write ϕ^{tr} for the translation of ϕ):

- The translation commutes with finite conjunctions and disjunctions (and with universal and existential quantifiers when these are present). (In detail: $(\phi \wedge \psi)^{tr} = \phi^{tr} \wedge \psi^{tr}$, $(\phi \vee \psi)^{tr} = \phi^{tr} \vee \psi^{tr}$, $(\top)^{tr} = \top$, and $(\bot)^{tr} = \bot$.)
- Intuitionistic implications translate into strict implications: $(\phi \to \psi)^{tr} = \Box(\phi^{tr} \to \psi^{tr})$. As derived rules we get $(\phi \leftrightarrow \psi)^{tr} = \Box(\phi^{tr} \leftrightarrow \psi^{tr})$, $(\top \to \phi)^{tr} = \Box(\phi^{tr})$, and $(\phi \to \bot)^{tr} = \Box(\phi^{tr} \to \bot)$.
- A propositional constant c gets translated into $\Box c$ (that is, $c^{tr} = \Box c$).

[16] The axioms of S4 are $\Box(A \to B) \to (\Box A \to \Box B)$, $\Box A \to A$, and $\Box A \to \Box\Box A$, and the translation of intuitionistic logic into S4 modal logic is known as the Gödel–McKinsey–Tarski translation. It was first given in (Gödel, 1933) and its faithfulness was proven in (McKinsey and Tarski, 1948). The logic can actually be taken to be S4.Grz (S4 modal logic extended with Grzegorczyk's axiom: $\Box((\Box(p \to \Box p) \to p) \to p) \to \Box p)$ and this observation leads to a connection with provability logic. See (Esakia, 2004) for more on this.

8 Algebraic Semantics: Basic Definition

Let us now look more closely at the problem of giving an algebraic semantics for IpC^2. Since IpC^2 extends ordinary intuitionistic propositional logic, we may begin by looking at the problem of how to model ordinary intuitionistic logic in a poset. Somehow the poset needs to have counterparts of $\top$, $\bot$, $\vee$, $\wedge$, and $\rightarrow$. A poset with counterparts for the first four of these is known as a *lattice*. Formally, a lattice is a poset with finite meets and finite joins. 0 is the smallest element while 1 is the largest element, and for two elements x and y we have that $x \wedge y$ is maximal among elements less than x and y while $x \vee y$ is minimal among elements greater than x and y.

To model $\rightarrow$ we need to require the lattice to be a *Heyting algebra*.[17] A Heyting algebra is a lattice which has a *relative pseudo-complement* (also known as an *exponential*) $x \Rightarrow y$ for any two elements x and y. The defining property of $x \Rightarrow y$ is that it is maximal among elements z such that $(x \wedge z) \leq y$. It follows from this that $x \Rightarrow y = 1$ if and only if $x \leq y$ (we can thus eliminate $\leq$ in favor of $\Rightarrow$ and equalities of the form $\phi = 1$).

When interpreting intuitionistic propositional logic in a Heyting algebra, one defines (note that the symbols '$\wedge$' and '$\vee$' are used here for connectives in intuitionistic logic as well as for lattice operations):

- $[\![\top]\!] = 1$
- $[\![\bot]\!] = 0$
- $[\![\phi \wedge \psi]\!] = [\![\phi]\!] \wedge [\![\psi]\!]$
- $[\![\phi \vee \psi]\!] = [\![\phi]\!] \vee [\![\psi]\!]$
- $[\![\phi \rightarrow \psi]\!] = [\![\phi]\!] \Rightarrow [\![\psi]\!]$

It is customary to use $x \Leftrightarrow y$ to denote the same element as $(x \Rightarrow y) \wedge (y \Rightarrow x)$, and we then naturally get:

- $[\![\phi \leftrightarrow \psi]\!] = [\![\phi]\!] \Leftrightarrow [\![\psi]\!]$.

A formula is considered to hold in a Heyting algebra model if and only if its interpretation is 1. For example, $\top \rightarrow \top$ holds in any Heyting algebra since $[\![\top \rightarrow \top]\!] = [\![\top]\!] \Rightarrow [\![\top]\!] =$

[17] At least, this seems the most straightforward way of doing things. (Pavičić and Megill, 1999) showed that classical logic can actually be modelled in lattices that need not be Heyting algebras or even distributive lattices ($a \wedge (b \vee c) = (a \wedge b) \vee (a \wedge c)$ need not hold).

$1 \Rightarrow 1 = 1$. As a derived rule we get that $\phi \to \psi$ holds in a Heyting algebra if and only if $[\![\phi]\!] \le [\![\psi]\!]$. (So while one could say that the semantics makes '$\to$' correspond to '$\Rightarrow$,' one could also say with some justification that it makes '$\to$' correspond to '$\le$.')

In the case of IpC^2 we do not need to explicitly state the interpretations of $\top$, $\bot$, $\wedge$, and $\vee$ since these can be defined out of $\to$ and $\forall$.[18] However, we still need a poset equipped with the operation $\Rightarrow$ (since this operation is what we use for the interpretation of $\to$), and we may as well assume from the start that the poset is actually a Heyting algebra.

In order to be able to handle $\forall$ we will assign interpretations not only to closed formulas of IpC^2, but also to formulas with free variables. A formula with free variables will be interpreted as a function whose arguments correspond to the free variables. If an expression ϕ contains n free variables x_1, ..., x_n then we write $[\![x_1, ..., x_n \rhd \phi]\!]$ for its interpretation.[19] This will be a function f of n arguments such that $f([\![\chi_1]\!], ..., [\![\chi_n]\!]) = [\![\phi[\chi_1/x_1, ..., \chi_n/x_n]]\!]$. We identify a function of 0 arguments with its value.

We will use $\bar{x}$ as a shorthand notation for a sequence such as $x_1, ..., x_k$. Thus, $x_1, ..., x_k \rhd \phi$ will typically be written $\bar{x} \rhd \phi$ and $f(a_1, ..., a_m)$ will typically be written $f(\bar{a})$.

I will often write functions using lambda notation. Thus, $\lambda x.\phi$ denotes a function whose value for the argument x is ϕ (in another commonly used notation: $x \mapsto \phi$) and $\lambda x_1, ..., x_n.\chi$ (or $\lambda \bar{x}.\chi$) denotes a function of n arguments whose value for the arguments $x_1, ..., x_n$ is χ. For example, $\lambda x.x$ is the identity function and $\lambda x.\top$ is the constant function of one argument whose value is always $\top$. Since all the functions we deal with will take their arguments in the Heyting algebra H, we will not write this out explicitly (for example, we write $\lambda x.x$ rather than $\lambda x : H.x$).

We extend Heyting algebra operations such as $\wedge$, $\vee$, $\Rightarrow$, 1, and 0 to functions in a pointwise manner. For example, we have $(\lambda \bar{x}.\, \phi) \wedge (\lambda \bar{y}.\, \chi) = \lambda \bar{x}.\, \phi \wedge \chi$ and $1 = \lambda \bar{x}.\, 1$. Moreover, we understand $\lambda \bar{x}.\, \phi \le \lambda \bar{y}.\, \chi$ to mean that we have $(\lambda \bar{x}.\, \phi)(\bar{e}) \le (\lambda \bar{y}.\, \chi)(\bar{e})$ for any choice of elements $\bar{e} = e_1, ..., e_n$.

To model $[\![\bar{x} \rhd \forall x_{n+1}.\phi]\!]$ we take the infimum of $[\![\bar{x}, x_{n+1} \rhd \phi]\!](\bar{a}, a_{n+1})$ as a_{n+1} ranges over the Heyting algebra. To be precise, $[\![\bar{x} \rhd \forall x_{n+1}.\phi]\!] = \lambda \bar{e}.\, \inf\{[\![\bar{x}, x_{n+1} \rhd \phi]\!](\bar{e}, e_{n+1}) : e_{n+1} \in H\}$ (the letter 'e' will frequently be used for elements of the Heyting algebra H).

But why model $\forall$ as an infimum? Well, we found in section 3 that $\forall x.\psi$ is equivalent to

[18] We should verify, though, that the interpretations of $\top$, $\bot$, $\wedge$, and $\vee$ are as expected. The verifications for $\top$ and $\bot$ will be trivial and the verification for $\wedge$ and $\vee$ will be done in lemma 9 and theorem 11, respectively, below.

[19] Strictly speaking, an expression with one or more free variables is assigned no interpretation at all. Interpretations are instead assigned to constructions of the form $x_1, ..., x_n \rhd \phi$, where the order of the variables $x_1, ..., x_n$ matters and where it is not required that the variables $x_1, ..., x_n$ actually occur in ϕ.

$\inf \psi[x]$ s.t. $\top$ and this would seem to suggest rather strongly that $\forall$ can be understood as an infimum. The formal justification lies in the fact that we get a semantics that is sound and complete (see below), and it may be noted that it is actually possible to set up a Heyting algebra semantics for IpC^2 in which $\forall$ is *not* modeled through infima (such a semantics is given in (Pitts, 1992)).

Note that $[\![\forall x.\phi]\!] = 1$ if and only if $[\![x \rhd \phi]\!]$ is the constant function whose value is always 1. In the very special case where the Heyting algebra consists of exactly two elements 0 and 1, this property is sufficient to fix the behavior of $\forall$.

Note also that the functions $[\![\bar{x} \rhd \phi]\!]$ that arise in this semantics include functions in addition to those we would get with unquantified intuitionistic propositional logic (functions that can be formed out of $\to$, $\wedge$, $\vee$, $\top$, $\bot$, constants, and variables). This reflects the fact that IpC^2 is far more expressive than unquantified intuitionistic propositional logic.[20]

Let us now state all rules from scratch:
By a *structure* S for a theory T in IpC^2 we mean a tuple $(H, [\![-]\!])$ consisting of a Heyting algebra H and a function $[\![-]\!]$ that maps formulas of T to their interpretations. More exactly, $[\![-]\!]$ is defined for constructions of the form $x_1, ..., x_n \rhd \phi$ (typically written $\bar{x} \rhd \phi$), where x_1, ..., x_n include all the free variables of ϕ. When ϕ lacks free variables we also use $[\![\phi]\!]$ as a synonym for $[\![\rhd \phi]\!]$. The interpretation of $\bar{x} \rhd \phi$ is to be a function from H^n to H. As was mentioned above, we identify a function of 0 arguments with its value.

$[\![-]\!]$ may be any function that satisfies the following rules:

- $[\![\bar{x} \rhd \phi]\!]([\![\psi_1]\!], ..., [\![\psi_n]\!]) = [\![\rhd \phi[\psi_1/x_1, ..., \psi_n/x_n]]\!]$,
- $[\![\bar{x} \rhd \phi \to \psi]\!] = \lambda \bar{e}. [\![\bar{x} \rhd \phi]\!](\bar{e}) \Rightarrow [\![\bar{x} \rhd \psi]\!](\bar{e})$, and
- $[\![\bar{x} \rhd \forall x_{n+1}\phi]\!] = \lambda \bar{e}. \inf\{[\![\bar{x}, x_{n+1} \rhd \phi]\!](\bar{e}, e_{n+1}) : e_{n+1} \in H\}$.

The structure S is called a *model* of T if $[\![\phi]\!] = 1$ whenever ϕ is provable in T. When $[\![\phi]\!] = 1$ we say that ϕ *holds in* S or that it is *true in* S. More generally, when ϕ contains free variables $x_1, ..., x_n$ we say that it holds/is true in S if $[\![x_1, ..., x_n \rhd \phi]\!]$ is a constant function that is always 1.

[20] The expressive power of IpC^2 was discussed in the introduction of this paper.

9 Algebraic Semantics: Soundness, Completeness, and Homomorphisms

Theorem 8. (Soundness for the Heyting algebra semantics) If a structure S validates all the axioms of a theory T, then it is a model of T.

Proof. (Depends on lemma 10.)

What we need to show is that S validates axioms A1.-A4. and that it is closed under the rules of inference I1.-I2.

For our inference rule I1. (modus ponens) we find that if $[\![\bar{z} \triangleright \phi]\!] = 1$ and $[\![\bar{z} \triangleright \phi \rightarrow \chi]\!] = [\![\bar{z} \triangleright \phi]\!] \Rightarrow [\![\bar{z} \triangleright \chi]\!] = 1$ then $[\![\bar{z} \triangleright \chi]\!] = 1$.

For axiom A1. $(\forall x, y.x \rightarrow y \rightarrow x)$ we get $[\![\bar{z} \triangleright \forall x, y.x \rightarrow y \rightarrow x]\!](\bar{e})$
$= \inf\{[\![\bar{z}, x \triangleright \forall y.x \rightarrow y \rightarrow x]\!](\bar{e}, a) : a \in H\}$
$= \inf\{\inf\{[\![\bar{z}, x, y \triangleright x \rightarrow y \rightarrow x]\!](\bar{e}, a, b) : b \in H\} : a \in H\}$
$= \inf\{\inf\{[\![\bar{z}, x, y \triangleright x]\!](\bar{e}, a, b) \Rightarrow [\![\bar{z}, x, y \triangleright y]\!](\bar{e}, a, b) \Rightarrow [\![\bar{z}, x, y \triangleright x]\!](\bar{e}, a, b) : b \in H\} : a \in H\}$
$= \inf\{\inf\{[\![\bar{z}, x, y \triangleright \top]\!](\bar{e}, a, b) : b \in H\} : a \in H\}$
$= 1.$

The verification for axiom A2. is similar to the verification for axiom A1.

For axiom A3. $((\forall x.\phi) \rightarrow \phi[y/x])$ we get $[\![\bar{z}, y \triangleright (\forall x.\phi) \rightarrow \phi[y/x]]\!](\bar{c}, b)$
$= [\![\bar{z}, y \triangleright (\forall x.\phi)]\!](\bar{c}, b) \Rightarrow [\![\bar{z}, y \triangleright \phi[y/x]]\!](\bar{c}, b)$
$= \inf\{[\![\bar{z}, y, x \triangleright \phi]\!](\bar{c}, b, a) : a \in H\} \Rightarrow [\![\bar{z}, y, x \triangleright \phi]\!](\bar{c}, b, b)$
$= 1.$

For axiom A4. $(\exists x.\, x \leftrightarrow \phi$ (x not free in ϕ)) we get $[\![\bar{z} \triangleright \exists x.\, x \leftrightarrow \phi]\!](\bar{e})$
$= [\![\bar{z} \triangleright \forall y.\, [\forall x.\, (x \leftrightarrow \phi) \rightarrow y] \rightarrow y]\!](\bar{e})$
$= \inf\{\inf\{([\![\bar{z}, x \triangleright x \leftrightarrow \phi]\!](\bar{e}, a) \Rightarrow [\![y \triangleright y]\!](b) : a \in H\} \Rightarrow [\![y \triangleright y]\!](b) : b \in H\}$
$= \inf\{\inf\{([\![\bar{z}, x \triangleright x \leftrightarrow \phi]\!](\bar{e}, a) \Rightarrow b : a \in H\} \Rightarrow b : b \in H\}$
$= \inf\{\inf\{[a \Leftrightarrow [\![\bar{x} \triangleright \phi]\!](\bar{e})] \Rightarrow b : a \in H\} \Rightarrow b : b \in H\}.$
Here the last step uses lemma 10. At this point, we find that the inner infimum simplifies to b (since the expression $[a \Leftrightarrow [\![\bar{x} \triangleright \phi]\!](\bar{e})] \Rightarrow b$ attains the value b when $a = [\![\bar{x} \triangleright \phi]\!](\bar{e})$), and since it cannot exceed b), and the whole expression therefore simplifies to $\inf\{b \Rightarrow b : b \in H\}$, which is obviously 1.

Let us finally turn to the verification of our second rule of inference:
I2. If x and y are not free in ϕ then $\phi \rightarrow \forall x.\psi$ may be inferred from $\phi \rightarrow \psi[y/x]$.

What we need to show is that if $[\![\top]\!] = [\![\bar{z} \rhd \phi]\!](\bar{e}) \Rightarrow [\![\bar{z}, y \rhd \psi[y/x]]\!](\bar{e}, b)$ holds for any choice of $\bar{e}, b$ then $[\![\top]\!] = [\![\bar{z} \rhd \phi]\!](\bar{e}) \Rightarrow [\![\bar{z} \rhd \forall x.\psi]\!](\bar{e})$ holds (for any choice of $\bar{e}$). That is, we need to show that if $[\![\bar{z} \rhd \phi]\!](\bar{e}) \leq [\![\bar{z}, y \rhd \psi[y/x]]\!](\bar{e}, b)$ holds for any choice of $\bar{e}, b$ then $[\![\bar{z} \rhd \phi]\!](\bar{e}) \leq [\![\bar{z} \rhd \forall x.\psi]\!](\bar{e})$ holds. However, this is a direct consequence of the fact that we are modeling the universal quantifier as an infimum. □

Lemma 9. $[\![\bar{x} \rhd \phi \wedge \chi]\!] = [\![\bar{x} \rhd \phi]\!] \wedge [\![\bar{x} \rhd \chi]\!]$.

Proof. $\phi \wedge \chi$ is defined to be $\forall y.[\phi \to \chi \to y] \to y$ (where y is not free in ϕ or χ). We therefore have $[\![\bar{x} \rhd \phi \wedge \chi]\!](\bar{e})$
$= [\![\bar{x} \rhd \forall y.[\phi \to \chi \to y] \to y]\!](\bar{e})$
$= \inf\{[\![\bar{x}, y \rhd [\phi \to \chi \to y] \to y]\!](\bar{e}, e') : e' \in H\}$
$= \inf\{[[\![\bar{x} \rhd \phi]\!](\bar{e}) \Rightarrow [\![\bar{x} \rhd \chi]\!](\bar{e}) \Rightarrow [\![y \rhd y]\!](e')] \Rightarrow [\![y \rhd y]\!](e') : e' \in H\}$
$= \inf\{[[\![\bar{x} \rhd \phi]\!](\bar{e}) \Rightarrow [\![\bar{x} \rhd \chi]\!](\bar{e}) \Rightarrow e'] \Rightarrow e' : e' \in H\}$.

At this point, the desired result follows once it is noticed that $\inf\{[e_1 \Rightarrow e_2 \Rightarrow e_3] \Rightarrow e_3 : e_3 \in H\} = e_1 \wedge e_2$ holds in any Heyting algebra (the proof of this fact about Heyting algebras is left out). □

Lemma 10. $[\![\bar{x} \rhd \phi \leftrightarrow \chi]\!] = [\![\bar{x} \rhd \phi]\!] \Leftrightarrow [\![\bar{x} \rhd \chi]\!]$.

Proof. By definition, $\phi \leftrightarrow \chi$ is the same thing as $(\phi \to \chi) \wedge (\chi \to \phi)$. We therefore have $[\![\bar{x} \rhd \phi \leftrightarrow \chi]\!] = [\![\bar{x} \rhd (\phi \to \chi) \wedge (\chi \to \phi)]\!] = [\![\bar{x} \rhd (\phi \to \chi)]\!] \wedge [\![\bar{x} \rhd (\chi \to \phi)]\!]$, where the last identity comes from lemma 9. This is clearly equivalent to $([\![\bar{x} \rhd \phi]\!] \Rightarrow [\![\bar{x} \rhd \chi]\!]) \wedge [\![\bar{x} \rhd \chi]\!] \Rightarrow [\![\bar{x} \rhd \phi]\!])$, which is by definition the same thing as $[\![\bar{x} \rhd \phi]\!] \Leftrightarrow [\![\bar{x} \rhd \chi]\!]$. □

The preceding lemmas show that $\wedge$ and $\leftrightarrow$ get interpreted as expected. The following theorem confirms that $\vee$ and $\exists$ get interpreted in the expected way. This time the proof will use theorem 3 in combination with the soundness theorem we have just proved (although a more direct verification might work equally well):

Theorem 11. (1) $[\![\bar{x} \rhd \phi \vee \chi]\!] = [\![\bar{x} \rhd \phi]\!] \vee [\![\bar{x} \rhd \chi]\!]$

(2) $[\![\bar{x} \rhd \exists y \phi]\!] = \lambda \bar{e}.\ \sup\{[\![\bar{x}, y \rhd \phi]\!](\bar{e}, e') : e' \in H\}$

Proof. We have from theorem 3 that $\phi \to (\phi \vee \chi)$ holds. Thus, we have $[\![\bar{x} \rhd \top]\!] = [\![\bar{x} \rhd \phi \to (\phi \vee \chi)]\!]$, which gives us $[\![\bar{x} \rhd \phi]\!] \leq [\![\bar{x} \rhd \phi \vee \chi]\!]$. We similarly get $[\![\bar{x} \rhd \chi]\!] \leq [\![\bar{x} \rhd \phi \vee \chi]\!]$.

Theorem 3 also gives us $[(\alpha \to \omega) \wedge (\beta \to \omega)] \to [(\alpha \vee \beta) \to \omega]$, from which we can get $[[\![\bar{x} \rhd \alpha]\!](\bar{e}) \Rightarrow [\![y \rhd y]\!](e')] \wedge [[\![\bar{x} \rhd \beta]\!](\bar{e}) \Rightarrow [\![y \rhd y]\!](e')] \leq [\![\bar{x} \rhd (\alpha \vee \beta)]\!](\bar{e}) \Rightarrow [\![y \rhd y]\!](e')$. Or more simply: $[[\![\bar{x} \rhd \alpha]\!](\bar{e}) \Rightarrow e'] \wedge [[\![\bar{x} \rhd \beta]\!](\bar{e}) \Rightarrow e'] \leq [\![\bar{x} \rhd (\alpha \vee \beta)]\!](\bar{e}) \Rightarrow e'$.

By combining the above observations, we see that $[\![\bar{x} \rhd \phi \vee \chi]\!](\bar{e})$ has to be $[\![\bar{x} \rhd \phi]\!](\bar{e}) \vee [\![\bar{x} \rhd \chi]\!](\bar{e})$, from which we get $[\![\bar{x} \rhd \phi \vee \chi]\!] = [\![\bar{x} \rhd \phi]\!] \vee [\![\bar{x} \rhd \chi]\!]$.

To prove 2), we first use the part of theorem 3 which says that $\psi[\chi/x] \to \exists x.\psi$, which gives us $[\![\top]\!] = [\![\bar{x}, y \rhd \psi[y/x]]\!](\bar{e}, e') \Rightarrow [\![\bar{x} \rhd \exists x.\psi]\!](\bar{e})$ if we set $\chi = y$. We may also write this as $\lambda\bar{e}.\, [\![\bar{x}, x \rhd \psi]\!](\bar{e}, e') \leq [\![\bar{x} \rhd \exists x.\psi]\!]$.

Theorem 3 also tells us that if $\chi \to \psi$ holds and if x does not occur free in ψ then $(\exists x.\chi) \to \psi$ holds. That is, if $1 = [\![\bar{x}, x \rhd \chi]\!](\bar{e}, e') \Rightarrow [\![\bar{x} \rhd \psi]\!](\bar{e})$ holds for any choice of $\bar{e}, e'$ then $1 = [\![\bar{x} \rhd (\exists x.\chi)]\!] \Rightarrow [\![\bar{x} \rhd \psi]\!]$. Equivalently, if $\lambda\bar{e}.\, [\![\bar{x}, x \rhd \chi]\!](\bar{e}, e') \leq [\![\bar{x} \rhd \psi]\!]$ holds for any choice of e' then $[\![\bar{x} \rhd (\exists x.\chi)]\!] \leq [\![\bar{x} \rhd \psi]\!]$.

By combining the above observations, we see that $[\![\bar{x} \rhd \exists x.\psi]\!](\bar{e}) = \sup [\![\bar{x}, x \rhd \psi]\!](\bar{e}, e') : e' \in H$. □

Do we have a similar theorem for the interpretation of $\inf \psi[x]\,\text{s.t.}\,\chi$? We do have the following:

Theorem 12. $[\![\bar{x} \rhd \inf \psi[x]\,\text{s.t.}\,\chi]\!](\bar{a}) \leq \inf\{[\![\bar{x}, x \rhd \psi]\!](\bar{e}, e) : e \in H \;\&\; [\![\bar{x}, x \rhd \chi]\!](\bar{e}, e) = [\![\top]\!]\}$

Proof. To claim that $o \leq \inf S$ is equivalent to claiming that for each element x of S we have $o \leq x$. Thus, what needs to be shown is that for any element e of H such that $[\![\bar{x}, x \rhd \chi]\!](\bar{e}, e) = [\![\top]\!]$ we have $[\![\bar{x} \rhd \inf \psi[x]\,\text{s.t.}\,\chi]\!](\bar{e}) \leq [\![\bar{x}, x \rhd \psi]\!](\bar{e}, e)$.

We have that $[\![\bar{x} \rhd \inf \psi[x]\,\text{s.t.}\,\chi]\!](\bar{e}) = [\![\bar{x} \rhd \forall x.\, \chi \to \psi]\!](\bar{e}) = \inf\{[\![\bar{x}, x \rhd \chi]\!](\bar{e}, e') \Rightarrow [\![x \rhd \psi]\!](\bar{e}, e') : e' \in H\} \leq [\![\bar{x}, x \rhd \chi]\!](\bar{e}, e) \Rightarrow [\![\bar{x}, x \rhd \psi]\!](\bar{e}, e)$, which reduces to $[\![\bar{x}, x \rhd \psi]\!](\bar{e}, e)$ whenever $[\![\bar{x}, x \rhd \chi]\!](\bar{e}, e) = [\![\top]\!]$. □

That the above inequality cannot be strengthened to an equality can be shown by a simple counterexample. Let χ be a constant c such that $[\![c]\!]$ differs from $[\![\top]\!]$ as well as $[\![\bot]\!]$ (pick any model with more than two elements) and let ψ be any expression that lacks free variables other than x and which is such that $\forall x.\psi$ is not implied by c (for example, we can set $\phi = x$ so that $\forall x.\psi$ becomes $\bot$). By definition, $\inf \psi[x]\,\text{s.t.}\,\chi$ will then be $\forall x.c \to \psi$, which is equivalent to $c \to \forall x.\psi$. At the same time, $\inf\{[\![x \rhd \psi]\!](e) : e \in H \;\&\; [\![x \rhd \chi]\!](e) = [\![\top]\!]\}$ will be the infimum of the empty set and hence equal to $[\![\top]\!]$. As a result, we could only get an equality if $[\![c \to \forall x.\psi]\!]$ were $[\![\top]\!]$, but this will not be the case since we assumed that $\forall x.\psi$ was not implied by c.

Theorem 13. (Completeness for the Heyting algebra semantics) Every theory T has a model. Moreover, the model can be taken to have the property that the only formulas that hold in the model are such as are probable in T.

Proof. We construct a term model for T. Consider formula ϕ to be 'equivalent' to formula ψ when T proves $\phi = \psi$.

For a formula ϕ (we allow free variables to occur in ϕ) we define $\langle\phi\rangle$ to be its equivalence class. Let H be the set of all such equivalence classes.

We turn H into a Heyting algebra by setting $\langle\phi\rangle \leq \langle\psi\rangle$ if and only if $\phi \to \psi$ is provable in T. It is a routine task to verify that this makes H a Heyting algebra with $1 = \langle\top\rangle$, $\langle\phi\rangle \wedge \langle\psi\rangle = \langle\phi \wedge \psi\rangle$, $0 = \langle\bot\rangle$, $\langle\phi\rangle \vee \langle\psi\rangle = \langle\phi \vee \psi\rangle$, $\langle\phi\rangle \Rightarrow \langle\psi\rangle = \langle\phi \to \psi\rangle$.

Define $[\![x_1, ..., x_n \rhd \phi]\!]$ to be the function from H^n to H which maps $\langle\psi_1\rangle, ..., \langle\psi_n\rangle$ to $\langle\phi[\psi_1/x_1, ..., \psi_n/x_n]\rangle$. This is well-defined since $\phi = \phi'$ and $\psi_1 = \psi_1'$, ..., $\psi_n = \psi_n'$ entail $\phi[\psi_1/x_1, ..., \psi_n/x_n] = \phi'[\psi_1'/x_1, ..., \psi_n'/x_n]$. Note that we get $[\![\phi]\!] = \langle\phi\rangle$ in the special case where ϕ lacks free variables.

We need to verify that the above rules make universal quantifiers correspond to infima. Since T proves $\forall y.\phi \to \phi[\chi/y]$ (which is axiom A4), we have that $\langle\forall y.\phi\rangle \leq \langle\phi[\chi/y]\rangle$. The element $\langle\forall y.\phi\rangle$ is therefore a lower bound for elements of the form $\langle\phi[\chi/y]\rangle$.

To show that it is maximal among such lower bounds, we consider an arbitrary lower bound $\langle\psi\rangle$ for the elements $\langle\phi[\chi/y]\rangle$. $\langle\psi\rangle$ will then, in particular, be a lower bound for elements of the form $\langle\phi[z/y]\rangle$, where z is a variable that does not occur free in ψ or ϕ. From $\langle\psi\rangle \leq \langle\phi[z/y]\rangle$ we get that $\psi \to \phi[z/y]$ is provable in T. We will therefore also be able to prove $\psi \to \forall y.\phi$. This gives us $\langle\psi\rangle \leq \langle\forall y.\phi\rangle$, from which we see that the element $\langle\forall y.\phi\rangle$ is maximal among lower bounds for elements of the form $\langle\phi[\chi/y]\rangle$.

As long as ϵ_1, ..., ϵ_n do not contain y as a free variable, we have that $[\![x_1, ..., x_n \rhd \forall y.\phi]\!](\langle\epsilon_1\rangle, ..., \langle\epsilon_n\rangle) = \langle(\forall y.\phi)[\epsilon_1/x_1, ..., \epsilon_n/x_n]\rangle = \langle\forall y.\phi[\epsilon_1/x_1, ..., \epsilon_n/x_n]\rangle$. By what we have just shown, this is the greatest lower bound (the infimum) for elements of the form $\langle\phi[\epsilon_1/x_1, ..., \epsilon_n/x_n][\chi/y]\rangle$. But this element may also be written $[\![x_1, ..., x_n, y \rhd \phi]\!](\langle\epsilon_1\rangle, ..., \langle\epsilon_n\rangle, \langle\chi\rangle)$. We thus have that $[\![x_1, ..., x_n \rhd \forall y.\phi]\!](\langle\epsilon_1\rangle, ..., \langle\epsilon_n\rangle) = \inf\{[\![x_1, ..., x_n, y \rhd \phi]\!](\langle\epsilon_1\rangle, ..., \langle\epsilon_n\rangle, \langle\chi\rangle) : \langle\chi\rangle \in H\}$.

Finally, since $[\![\phi]\!] = 1$ is equivalent to $[\![\phi]\!] = [\![\top]\!]$ we see that the only formulas that hold in the model are ones for which $\phi = \top$ is provable. And since $\phi = \top$ is equivalent to ϕ we get the result that a formula cannot hold in the model unless it is provable. □

The model constructed for a theory T in the proof of theorem 13 will be called its *initial* model of T and denoted $\mathcal{M}(T)$.

Conversely, any Heyting algebra H gives rise to a theory $\mathcal{T}(H)$ which has as its constants the elements of H and which includes ϕ as an axiom just in case $[\![\phi]\!] = [\![1]\!]$ when each constant is interpreted by itself.

Since theories in the logics $\mathrm{IpC}^2\!\restriction$ (see section 4), $\mathrm{IpC}^2\!\equiv$, and $\mathrm{IpC}^2\!\restriction\!\equiv$ (see section 6) are also theories in IpC^2, in defining models for theories in IpC^2 we have also defined models for theories in these fragments of IpC^2.

By contrast, when we turn a model into a theory we need to make sure we end up with a theory in the right fragment of IpC^2. Since I have not given any rigorous definition of theories in $\mathrm{IpC}^2\!\equiv$ or $\mathrm{IpC}^2\!\restriction\!\equiv$, I will ignore these logics here. However, for $\mathrm{IpC}^2\!\restriction$ I will use the notation $\mathcal{T}_{\restriction}(H)$ to denote the theory which has the same constants as $\mathcal{T}(H)$ but which only includes axioms that belong to $\mathrm{IpC}^2\!\restriction$. As a result, while $\mathcal{T}_{\restriction}(H)$ will include a constant C_0 corresponding to the bottom element of the Heyting algebra H, it will not allow us to prove $C_0 = \bot$.

It remains to say something about how different models may be related to each other. A mapping $h : H \to H'$ between Heyting algebras will of course relate the theory $\mathcal{T}(H)$ to the theory $\mathcal{T}(H')$ in some way, and we will be interested in the case where the theories get related in a particularly 'nice' way.

To begin with, we may require that whenever a formula holds in $\mathcal{T}(H)$ then the corresponding formula must hold in $\mathcal{T}(H')$. In this case I will refer to h as an 'IpC^2 homomorphism' (or a 'homomorphism of IpC^2 structures'). We can analogously define '$\mathrm{IpC}^2\!\restriction$ homomorphism,' '$\mathrm{IpC}^2\!\equiv$ homomorphism,' and '$\mathrm{IpC}^2\!\restriction\!\equiv$ homomorphism' by requiring formulas in the respective fragments of IpC^2 to be preserved.

Note that an IpC^2 homomorphism is required to preserve $\bot$ while an $\mathrm{IpC}^2\!\restriction$ homomorphism is not.

Two particular types of homomorphisms seem worth singling out: Those that are injective and those that are surjective. If the mapping $h : H \to H'$ is injective we may write it as $h : H \rightarrowtail H'$ and if it is surjective we may write it as $h : H \twoheadrightarrow H'$.

A homomorphism $h : S \to S'$ for which we can find a homomorphism $k : S' \to S$ such that h and k are inverses of each other will be referred to an an *isomorphism*. We write $h : S \xrightarrow{\sim} S'$ in this case.

By a *substructure* of a structure S, I will mean an equivalence class of injective homomorphisms $h : S' \rightarrowtail S$, where $k : S'' \rightarrowtail S$ is equivalent to h if there is an isomorphism $i : S' \xrightarrow{\sim} S''$ such that $h = ki$.

Dually, a *quotient structure* of a structure S will mean an equivalence class of surjective homomorphisms $h : S \twoheadrightarrow S'$, where $k : S \twoheadrightarrow S''$ is equivalent to h if there is an isomorphism $i : S' \xrightarrow{\sim} S''$ such that $k = ih$.

Note that when we extend a theory T into a theory T' by adding constants without adding any axioms then we get a corresponding *injection* $h : \mathcal{M}(T) \rightarrowtail \mathcal{M}(T')$. (The equivalence class of h will be a subobject of $\mathcal{M}(T')$.)

Note also that if we extend a theory T into a theory T^* by adding axioms without adding any constants then we get a corresponding *surjection* $k : \mathcal{M}(T) \twoheadrightarrow \mathcal{M}(T^*)$. (The equivalence class of k will be a quotient object of $\mathcal{M}(T')$.)

10 Topological Semantics

Let us now look at topological models. From a topological space S we get a model as follows:

- The elements of the Heyting algebra are the open sets of S.
- We set $e_1 \wedge e_2 =_{def.} e_1 \cap e_2$, $e_1 \vee e_2 =_{def.} e_1 \cup e_2$, $0 =_{def.} \{\}$, $1 =_{def.} S$.
- $e_1 \Rightarrow e_2 =_{def.} \Box(\complement e_1 \cup e_2)$, where $\Box(-) = int(-)$ is the *interior* operation on S.

With this interpretation it becomes very natural to say that $e_1 \vee e_2$ contains e_1 and e_2 as parts. How can this be reconciled with the understanding that what $a \to b$ tells us is that b is part of a? Well, there is a dual translation that makes the elements of the Heyting algebra correspond to *closed* sets and which does make $a \to b$ say that b is part of a:

- The elements of the Heyting algebra are the closed sets of S.
- We set $e_1 \wedge e_2 =_{def.} e_1 \cup e_2$, $e_1 \vee e_2 =_{def.} e_1 \cap e_2$, $0 =_{def.} S$, $1 =_{def.} \{\}$.
- $e_1 \Rightarrow e_2 =_{def.} = \Box(\complement e_1 \cap e_2)$, where $\Box(-) = cl(-)$ is the *closure* operation on S.

As the reader may have guessed, the operation $\Box$ that we are using here can be seen as corresponding to the necessity operator ('$\Box$') of S4 modal logic. We saw above how intuitionistic logic can be translated into S4 modal logic, and this may be followed by a topological interpretation of modal logic where the necessity operator comes to correspond to either the interior operation (if open sets are used) or the closure operation (if closed sets are used). Modal logic is thus very relevant to the topological interpretation of intuitionistic logic.

Note that while it is true that the open sets of a topological space S stand in a one-to-one correspondence with the closed sets of S, if we embed S in a larger topological space S' then the correspondence between open and closed sets changes: In S the closed set $\{\}$ corresponds to the open set S whereas in S' we instead find that the closed set $\{\}$ corresponds to S' (and we will not have $S = S'$ unless the embedding is trivial). This has the consequence that when we use the interpretation in terms of open sets then the meaning of $\top$ changes if we decide to use S' rather than S as our topological space. With the interpretation in terms of closed sets, it is instead the meaning of $\bot$ that is sensitive to whether the space is S or S'. This would seem to suggest that $\mathrm{IpC}^2\!\upharpoonright$ (see section 4) gives us reason for preferring the interpretation in terms of closed sets while $\mathrm{IpC}^2\!\equiv$ (see section 6) gives us reason for preferring the interpretation in terms of open sets.

It does not seem to be known whether the topological semantics described here is complete.

11 How to say that one object is *not* part of another

The theory S presented in section 5 leaves some questions open. For example, it is neither provable nor disprovable in S that $a \leftarrow b$. Suppose now that we want an extension S' of S where $a \leftarrow b$ is denied. How can we achieve this?

At first the answer seems very simple: S' needs to include the axiom $\neg(b \to a)$ (which we may also write as $(b \to a) \to \bot$ or $(b \to a) \leftrightarrow \bot$), and that is it. But $(b \to a) \to \bot$ is equivalent to $((b \to \bot) \to \bot) \wedge (a \to \bot)$, which is hardly what we wanted to say:

Theorem 14. $(b \to a) \to \bot$ is equivalent to $((b \to \bot) \to \bot) \wedge (a = \bot)$.

Proof.
Set $c = \bot$ in the following theorem. □

Theorem 15. If $c \to a$ holds then $(b \to a) \to c$ is equivalent to $((b \to c) \to c) \wedge (a = c)$.

Proof. For the left-to-right direction, combine parts (i) and (iii) of the following lemma. The right-to-left direction is trivial. □

Lemma 16. (i) $(b \to a) \to c$ implies $a \to c$,

(ii) if $a' \to a$ holds then $(b \to a) \to c$ implies $(b \to a') \to c$, and

(iii) if $c \to a$ holds then $(b \to a) \to c$ implies $(b \to c) \to c$.

Proof.

(i) Use the fact that $a \to (b \to a)$ along with the transitivity of implication.

(ii) Assume $a' \to a$, $(b \to a) \to c$, and $b \to a'$. Combine the first and the last of these implications to get $b \to a$ and then use the second implication.

(iii) This may be seen either as a consequence of (i) (the combination of $a \to c$ and $c \to a$ gives us $a = c$) or as a consequence of (ii) (set $c = a'$).

□

And just as $\neg(b \to a)$ has unwanted consequences, the same can easily be the case with other formulas that seem unproblematic at first. For example, if we try to deny $b \to a$ through $(b \to a) \to b$ or $(b \to a) \to (a \wedge b)$ (which are equivalent), we end up asserting $a \to b$ (cf. part (i) of the above lemma).

So what is the solution? One thing we can do is to define theories in such a way that a theory may not only *accept* certain formulas as correct but may also *reject* certain formulas as incorrect. Gabbay (1981) defines theories of IpC^2 in precisely this way: A theory consists of two sets of formulas, those that are accepted by the theory and those that are rejected (see (Gabbay, 1981, p. 160)). A theory counts as inconsistent if some formula is accepted and rejected at the same time, and models are defined in such a way that rejected formulas cannot hold in them. We can then ensure that $b \to a$ does not hold by simply rejecting this formula.

But the rejection here is not expressed inside IpC^2 itself, and we could theoretically end up with a very long list of formulas that we reject. What can we do if our goal is to do as much as possible within IpC^2 itself? Well, note that if $\phi \to \psi$ holds and we want for neither ϕ nor ψ to hold, then it is enough to reject ψ. By exploiting this fact we can make sure that no more than one formula needs to be rejected outside of IpC^2 (perhaps using the method of (Gabbay, 1981)).

So if we have expressively rejected ψ then all we need to do in order to get an effective rejection of ϕ is to assert $\phi \to \psi$.

We should be wary, though, that in asserting $\phi \to \psi$ we are really asserting that ψ is part of ϕ. Our sole intention may be to reject ψ, but we are doing more than that. For example, if ψ is $\bot$ then $\phi \to \psi$ is another way of saying that ϕ is $\bot$.

This is not always a problem, but it is something we really want to be aware of when our goal is to reject $b \to a$. In asserting $(b \to a) \to \psi$ we are indirectly asserting $a \to \psi$ (by part (i) of the above lemma). We have already discussed the extreme case where ψ is $\bot$: We wanted to deny $b \to a$ (that b contains a) but ended up asserting $a \to \bot$ (that everything is part of a).

However, $a \to \psi$ may be more acceptable if ψ is an object that 'contains very little' (as opposed to the object $\bot$ that contains everything). We cannot set $\psi = \top$ (in which case $a \to \psi$ would be absolutely harmless) since we want to reject ψ and cannot reject $\top$, but we can set ψ equal to an object that is 'close' to $\top$, an object that contains little.

What we will do in practice is to set $\psi = \mathcal{R}$, where $\mathcal{R}$ is a constant that we introduce in order to have a way to reject formulas. For example, to reject $b \to a$ we assert $(b \to a) \to \mathcal{R}$.[21]

Readers familiar with minimal logic will notice that $\mathcal{R}$ is used much like the absurd proposition in minimal logic.[22] To emphasize the similarities between $\phi \to \bot$ and $\phi \to \mathcal{R}$, I will introduce the following notation:

$$\neg_{\mathcal{R}}\phi \leftrightarrow_{def.} \phi \to \mathcal{R}.$$

As already pointed out, in asserting $(b \to a) \to \mathcal{R}$ we are asserting that a contains $\mathcal{R}$. In fact, as soon as we assert $(\top \to a) \to \mathcal{R}$ (which is a way of rejecting $a = \top$) we are asserting that a contains $\mathcal{R}$. The result is that $\mathcal{R}$ comes to be part of every object x for which we use $\mathcal{R}$ to reject $x = \top$. This gives $\mathcal{R}$ the character of a very 'tiny' object, an object much like $\top$.

At the same time, when we use $\phi \to \mathcal{R}$ to reject ϕ, then we are using $\mathcal{R}$ in the same way that $\bot$ is traditionally used, as an absurd proposition, something we know will never be provable. Moreover, it becomes natural to think of $\phi \vee \mathcal{R}$ as 'saying the same thing as' ϕ[23],

[21] Note that although our motivating question has been how to say that one object is not part of another, we can equally well think of our problem as one of saying that one *proposition* fails to *entail* another. To say that b entails a we may assert $b \to a$, but to deny that b entails a we would normally assert $\nvDash (b \to a)$ rather than $(b \to a) \to \bot$ (if we assert the latter then we are indirectly asserting $a = \bot$). But the '$\nvDash$' notation does not belong to the same language as '$b \to a$,' and so we may prefer to add a constant $\mathcal{R}$ to the language and use $(b \to a) \to \mathcal{R}$ to reject $b \to a$.

[22] See (Johansson, 1937).

[23] If we understand $\phi \vee \mathcal{R}$ as meaning 'either ϕ or $\mathcal{R}$' and if we have rejected $\mathcal{R}$ then it seems that ϕ has to hold whenever $\phi \vee \mathcal{R}$ holds. This informal idea can be substantiated by looking at models where $\phi_1 \vee \phi_2$ cannot hold unless either ϕ_1 or ϕ_2 holds. That we can get a complete semantics for IpC^2 using only models with this property was shown in (Gabbay, 1981).

and since $\mathcal{R}$ entails $\phi \vee \mathcal{R}$ we see that $\mathcal{R}$ gets the character of a proposition that entails every proposition: Regardless of what ϕ is, $\mathcal{R}$ entails a proposition $\phi \vee \mathcal{R}$ which we can think of as 'saying the same thing as' ϕ.

To sum up, $\mathcal{R}$ tends to be contained in the objects whose parts we are reasoning about (such as the object a considered above), but it simultaneously contains propositions that we use to reason about those objects. From one point of view $\mathcal{R}$ looks like $\top$; from another point of view it looks like $\bot$.

In fact, when reasoning with $\mathcal{R}$ we may often want to compare objects using one of the following equivalence relations, the first of which makes $\mathcal{R}$ equivalent to $\top$ and the second of which makes it equivalent to $\bot$:
$\chi_1 \sim_{\mathcal{R}} \chi_2 \leftrightarrow_{def.} \chi_1 \wedge \mathcal{R} = \chi_2 \wedge \mathcal{R}$, and
$\chi_1 \equiv_{\mathcal{R}} \chi_2 \leftrightarrow_{def.} \chi_1 \vee \mathcal{R} = \chi_2 \vee \mathcal{R}$.

Note that $\chi_1 \sim_{\mathcal{R}} \chi_2$ may be rewritten as $\mathcal{R} \to (\chi_1 = \chi_2)$. To determine whether $\chi_1 \sim_{\mathcal{R}} \chi_2$ holds is thus a matter of determining whether $\chi_1 = \chi_2$ holds under the assumption that $\mathcal{R} = \top$. In practice we will not be able to prove $\mathcal{R} = \top$ (after all, we are rejecting $\mathcal{R}$), but when counting how many parts an object has we will normally want to compare objects using the relation $\sim_{\mathcal{R}}$ so that $\mathcal{R}$ and $\top$ (and any other parts of $\mathcal{R}$) count as one and the same thing.

This makes sense if we think of $\mathcal{R}$ as an 'artificial object' that we have added in order to be able to reason effectively about which object is part of which. If $\mathcal{R}$ is like a camera then we do not want for any traces of $\mathcal{R}$ to remain in the final pictures.

If, for some reason, we do want to reason about the parts of $\mathcal{R}$, then we can do so by adding a second constant $\mathcal{R}'$ that we use just like $\mathcal{R}$ (we may then assert $(\top \to \mathcal{R}) \to \mathcal{R}'$). To reason about the parts of $\mathcal{R}'$ we would use a third constant $\mathcal{R}''$. In intuitive terms, we are using one 'camera' to photograph another.[24]

Let us now consider the question of how we can be sure that nothing bad will result if we start rejecting things using $\mathcal{R}$. The following theorem seems to be what we need:

Theorem 17.

[24] It is also instructive to compare what we are doing here to cases where one is using one formal language L' to reason about another formal language L. We can suppose that L proves the law of excluded middle – any proposition P is equivalent to either *truth* or *falsity* – but that L' nevertheless allows us to get a more 'fine-grained' view on the sentences of L according to which some sentences of L have a semantic value that differs from 1 (*truth*) as well as 0 (*falsity*). (We could make the semantic value 1 match with 'provable,' make the semantic value 0 match with 'disprovable,' and make other semantic values match with 'true in some models and false in others.') This is to be compared to the way that the equivalence relation $\sim_{\mathcal{R}'}$ may be more 'fine-grained' than the equivalence relation $\sim_{\mathcal{R}}$.

Given a theory T which does not contain the constant $\mathcal{R}$, a theory T' which extends T with the constant $\mathcal{R}$ and with axioms of the form $\phi_i \to \mathcal{R}$ is a conservative extension of T. Moreover, in any model of T' where $[\![\mathcal{R}]\!]$ is distinct from $[\![\top]\!]$, $[\![\phi_i]\!]$ is distinct from $[\![\top]\!]$ for each i.

Proof.
That T' is a conservative extension follows since it is possible to create a further extension T'' which adds $\top = \mathcal{R}$ as an axiom (but does not otherwise add anything). It is easy to see that the axioms of T' all follow from $\top = \mathcal{R}$ and so T'' is essentially just T along with the constant $\mathcal{R}$ and the axiom $\top = \mathcal{R}$. That is evidently a conservative extension of T, and so T' must also be a conservative extension of T. This proves the first part of the theorem. To prove the second part, suppose that for some i we have $[\![\phi_i]\!] = [\![\top]\!]$ in our model. Because of the axiom $\phi_i \to \mathcal{R}$ we must then also have $[\![\mathcal{R}]\!] = [\![\top]\!]$, contrary to the assumption that $[\![\mathcal{R}]\!] \neq [\![\top]\!]$. □

Further reassurance that things work as we intend them to is given by the following theorem:

Theorem 18. In any structure where $[\![\mathcal{R}]\!] \neq [\![\top]\!]$ we have

1) $[\![x \to y]\!] = [\![\top]\!]$ if and only if $[\![x]\!] \leq [\![y]\!]$,
2) $[\![(x \to y) \to \mathcal{R}]\!] = [\![\top]\!]$ implies $[\![x]\!] \nleq [\![y]\!]$, and
3) if a) $[\![\mathcal{R}]\!] \leq v$ only holds for $v = [\![\top]\!]$ and $v = [\![\mathcal{R}]\!]$ and b) the structure satisfies the property that when $v \vee w = [\![\top]\!]$ holds then $v = [\![\top]\!]$ or $w = [\![\top]\!]$, then $[\![(x \to y) \to \mathcal{R}]\!] = [\![\top]\!]$ is equivalent to $[\![x]\!] \nleq [\![y]\!]$.

Proof.
1) follows directly from the definition: $[\![x \to y]\!] = [\![x]\!] \Rightarrow [\![y]\!]$. This will be $[\![\top]\!]$ $(= 1)$ just in case $[\![x]\!] \leq [\![y]\!]$.

For the proof of 2), note that $[\![(x \to y) \to \mathcal{R}]\!] = [\![\top]\!]$ is equivalent to $[\![x \to y]\!] \leq [\![\mathcal{R}]\!]$. But (since $[\![\mathcal{R}]\!]$ was assumed to be distinct from $[\![\top]\!]$) this is incompatible with $[\![x \to y]\!] = [\![\top]\!]$. Combined with 1), this establishes 2).

The left-to-right direction of 3) follows from 2), so assume $[\![x]\!] \nleq [\![y]\!]$. We want to show $[\![(x \to y) \to \mathcal{R}]\!] = [\![\top]\!]$.

Note first that $(x \to y) \to \mathcal{R}$ may be written as $[(x \to y) \vee \mathcal{R}] \to \mathcal{R}$. It will therefore suffice to show that $[\![[(x \to y) \vee \mathcal{R}] \to \mathcal{R}]\!] = [\![\top]\!]$.

Since $[\![\mathcal{R}]\!] \leq [\![(x \to y) \vee \mathcal{R}]\!]$, assumption a) means that we have one of $[\![(x \to y) \vee \mathcal{R}]\!] = [\![\top]\!]$ and $[\![(x \to y) \vee \mathcal{R}]\!] = [\![\mathcal{R}]\!]$, and since $[\![\mathcal{R}]\!] \neq [\![\top]\!]$ we cannot have both at the same time.

We want to exclude the case $[\![(x \to y) \vee \mathcal{R}]\!] = [\![\top]\!]$. This can be written as $[\![(x \to y)]\!] \vee [\![\mathcal{R}]\!] = [\![\top]\!]$, and assumption b) then tells us that we either have $[\![(x \to y)]\!] = [\![\top]\!]$ or $[\![\mathcal{R}]\!] = [\![\top]\!]$. But the former is excluded by the assumption $[\![x]\!] \nleq [\![y]\!]$ and the latter is excluded by the assumption $[\![\mathcal{R}]\!] \neq [\![\top]\!]$.

We must therefore have that $[\![(x \to y) \vee \mathcal{R}]\!] = [\![\mathcal{R}]\!]$, from which $[\![(x \to y) \to \mathcal{R}]\!] = [\![\top]\!]$ follows. □

One thing we may not have expected is that if we assert $a \to \mathcal{R}$ as well as $b \to \mathcal{R}$ then $(a \vee b) \to \mathcal{R}$ follows. In a way this is as it should be: Someone who rejects a and b seems to be implicitly rejecting $a \vee b$. But if all we want to do as we 'reject' a and b is to make sure neither is provable, then we may still want to embrace $a \vee b$.

It seems possible to get what we want, though, by considering a formula ϕ to be 'embraced' whenever $\mathcal{E} \to \phi$ is provable (here $\mathcal{E}$ is a new constant that we are adding to our theories much like $\mathcal{R}$). To be more precise, for a formula ϕ involving neither $\mathcal{E}$ nor $\mathcal{R}$, we consider ϕ to be embraced whenever $\mathcal{E} \to \phi$ is provable. We can then use $(\mathcal{E} \to a) \to \mathcal{R}$ and $(\mathcal{E} \to b) \to \mathcal{R}$ to make it clear we are embracing neither a nor b, and at the same time we can use $\mathcal{E} \to (a \vee b)$ to embrace $a \vee b$.[25]

Let me finally give two examples of how one may use $\mathcal{R}$ to reason about parts. Here are two ways of saying that x is an 'atom'[26]:

$IsAtom_1(x) \leftrightarrow_{def.} [\forall y \leftarrow x.([y]_\mathcal{R} = [x]_\mathcal{R}) \vee ([y]_\mathcal{R} = [\top]_\mathcal{R})] \wedge \neg_\mathcal{R}([x]_\mathcal{R} = [\top]_\mathcal{R})$.
$IsAtom_2(x) \leftrightarrow_{def.} (x \to \mathcal{R}) \wedge [\forall \mathcal{R} \leftarrow y \leftarrow x.(y = x) \vee (y = \mathcal{R})] \wedge \neg_\mathcal{R}(x = \mathcal{R})$, and

With the first definition, x can be any object whatever, but what matters is what $[x]_\mathcal{R}$ is. With the second definition, x is required to be such that $x \to \mathcal{R}$ holds (which we may also express through $x = [x]_\mathcal{R}$). We expect to have $IsAtom_1(x) \leftrightarrow IsAtom_2([x]_\mathcal{R})$ for any x.

Note that there is a sense in which both definitions fail to ensure that x is an atom: They are consistent with x having lots and lots of parts as long as all those parts are parts of $\mathcal{R}$. We can, however, make parts of $\mathcal{R}$ 'count for nothing' by using $\sim_\mathcal{R}$ rather than $=$ as our standard for comparing objects (see above), and when we view things this way then the definitions seem to do precisely what we want them to do.

[25] How do we know the axioms that we are adding involving $\mathcal{E}$ are not causing any problems? They all follow from $\mathcal{E} \to \bot$, and if we extend a theory T with the constant $\mathcal{E}$ and the axiom $\mathcal{E} \to \bot$ then the resulting extension of T will evidently be a conservative one (since $\mathcal{E}$ merely functions as a synonym for $\bot$).

[26] The operation '$[-]_-$' was defined in section 6.

12 Lattice-describing formulas

Section 5 used formulas in IpC^2 that seemed to describe the parts of a certain object, but it was not clear at that point that the formulas actually said what they seemed to be saying. We will now prove theorems that give us some reassurance that our formulas really say what they seem to say.

The next two theorems deal with the way that H is related to the models of $\mathcal{T}(H)$ (see the definition in section 9).

Theorem 19. For a Heyting algebra H which is a structure for IpC^2 and any Heyting algebra model H' of the theory $Th(H)$, the mapping $x \mapsto [\![C(x)]\!]$ is a Heyting algebra homomorphism.

Proof. It is easy to check that this mapping indeed preserves $\leq$, $\Rightarrow$, $\vee$, and $\wedge$. For example, from $e_1 \leq e_2$ in H we get that $Th(H)$ proves $C(e_1) \to C(e_2)$, from which we get that $[\![C(e_1)]\!] \leq [\![C(e_2)]\!]$. □

For a theory T, let $Strict(T)$ be the theory which extends T by adding the constant $\mathcal{R}$ (if T already contains this constant, we first rename the old one) and by adding the axiom $\neg_{\mathcal{R}}(a = b)$ for any two distinct constants a and b of T.

Theorem 20. For a Heyting algebra H which is a structure for IpC^2 and any Heyting algebra model H' of the theory $Strict(Th(H))$ where $[\![\mathcal{R}]\!] \neq [\![\top]\!]$, the mapping $x \mapsto [\![C(x)]\!]$ is an injective Heyting algebra homomorphism.

Proof. That the mapping $x \mapsto [\![C(x)]\!]$ is a Heyting algebra homomorphism is the content of theorem 19. To prove injectivity, suppose e_1 and e_2 are two different elements of H. We need to show that $[\![C(e_1)]\!] \neq [\![C(e_2)]\!]$. From $e_1 \neq e_2$ we get that $C(e_1)$ and $C(e_2)$ are two different constants. By the definition of $Strict(-)$ we therefore have $\neg_{\mathcal{R}}(C(e_1) = C(e_2))$ (which is an abbreviation of $(C(e_1) = C(e_2)) \to \mathcal{R})$ as an axiom of $Strict(Th(H))$. We will therefore have $[\![C(e_1) = C(e_2)]\!] \leq [\![\mathcal{R}]\!]$ which means we cannot have $[\![C(e_1)]\!] = [\![C(e_2)]\!]$ unless $[\![\mathcal{R}]\!] = [\![\top]\!]$. But according to the assumptions we have $[\![\mathcal{R}]\!] \neq [\![\top]\!]$ and so we must have $[\![C(e_1)]\!] \neq [\![C(e_2)]\!]$. □

An important special case of the above theorems is the one where the Heyting algebra H is finite. In this case we can also prove the following:

Theorem 21. For a Heyting algebra H with finitely many elements $e_1 = 1, e_2, ..., e_n = 0$ all the axioms of $Th(H)$ follow from those of the form $C(e_i) \to C(e_j) = C(e_k)$, those of the form $C(e_i) \wedge C(e_j) = C(e_k)$, and the axiom $\forall x \leftarrow C(e_n).x = C(e_1) \vee x = C(e_2) \vee ... \vee x = C(e_n)$.

Proof. Let us write $Th'(H)$ for the subtheory of $Th(H)$ that only includes axioms of the form $C(e_i) \to C(e_j) = C(e_k)$, those of the form $C(e_i) \wedge C(e_j) = C(e_k)$, and the axiom $\forall x \leftarrow C(e_n).x = C(e_1) \vee x = C(e_2) \vee ... \vee x = C(e_n)$.

Let χ be an arbitrary axiom of $Th(H)$. We will show that $Th'(H)$ proves $\chi = C(e_i)$ for some element e_i of H. The theorem then follows trivially (of course, $C(e_i)$ is provable in $Th'(H)$ if and only if it is provable in $Th(H)$).

Assuming only the axioms of $Th'(H)$, we can transform χ into one of the constants $C(e_1), ..., C(e_n)$ in a step-by-step manner as follows:

1. a formula of the form $\forall x \leftarrow C(e_i).\phi$ is (because of the axiom $C(e_n) \to C(e_i)$) provably equivalent to $\forall x \leftarrow C(e_n).(x \leftarrow e_i) \to \phi$ and we can therefore rewrite χ so that all quantification is over $C(e_n)$.
2. $\forall x \leftarrow C(e_n).\phi$ is equivalent to $\phi([C(e_1)/x]) \wedge \phi([C(e_2)/x]) \wedge ... \wedge \phi([C(e_n)/x])$. By applying this equivalence we can therefore rewrite any universal quantifiers of χ in terms of $\wedge$ (since we have defined $\wedge$ in terms of universal quantifiers, it is not quite right to say that we are eliminating universal quantifiers).
3. By the definition of $\vee$ in terms of universal quantifiers we can eliminate any occurrence of $\vee$, and we can thus rewrite χ in terms of $\wedge$, $\to$, and the constants $C(e_1), C(e_2), ..., C(e_n)$.
4. By using the axioms of the form $C(e_i) \to C(e_j) = C(e_k)$ and those of the form $C(e_i) \wedge C(e_j) = C(e_k)$ we can simplify χ to one of the constants $C(e_1), C(e_2), ..., C(e_n)$.

□

By taking the conjunction of the finite set of axioms that we get from the previous theorem and subsequently replacing the constants $C(e_1), C(e_2), ..., C(e_n)$ with variables $x_1, ..., x_n$ we obtain a formula which we may denote $LikeHeytingAlgebra_H(x_1, ..., x_n)$ which tells us that the objects $x_1, ..., x_n$ are related as the elements of the finite Heyting algebra H.

13 A first-order theory for reasoning about parthood

As we reason about the parts of objects, we are likely to find ourselves informally distinguishing between 'objects' and 'propositions about objects.' For example, if we use $a \to b$ to state that object b is a part of object a, then $a \to b$ seems to play the role of a 'proposition' about the 'objects' a and b. We will now take this idea seriously by setting up a translation from a theory in ordinary predicate logic (many-sorted first-order predicate logic with a description operator, to be exact) into IpC^2. We will set up one version that uses minimal logic (mostly for demonstrative purposes), a second version that uses intuitionistic logic, and a third version that uses classical logic. The theory will have a predicate for saying that one object is part of another, and we will interpret this through the connective $\rightarrow$ of IpC^2. The theory will actually be a thin layer on top of IpC^2. Its role is not only to give IpC^2 a 'user-friendly' surface but (more importantly) to simplify the interpretation of traditional theories in IpC^2.

In addition to getting the customary division between 'objects' and 'propositions about objects' we will also be able to get two 'views' on each object. This will help us deal with an issue we have avoided so far: The parts of an object in IpC^2 form a Heyting algebra, but we may well find ourselves wanting to describe an object whose parts do not seem to form a Heyting algebra. As it turns out, there is a one-to-one relationship between finite Heyting algebras and finite posets: *Any finite poset occurs (up to isomorphism) as the meet irreducible elements of a unique (up to isomophism) finite Heyting algebra* (this is proved in theorem 22 below). It is therefore not at all clear that it is a limitation of IpC^2 that parts always form Heyting algebras: By taking a different 'view' of things we can instead see those Heyting algebras as posets with few constraints (in the finite case, there will be no constraints at all on what the posets may look like). Formally, we will achieve this by having a separate sort for 'irreducibles' (to be contrasted with the sort for 'objects').

An element x in a poset P is said to be meet irreducible if $x = y \wedge z$ implies either $x = y$ or $x = z$.[27]

If we want a restricted formula in IpC^2 which says this, the closest we can get would seem to be:

$$Irreducible(x) \leftrightarrow_{def.} \forall y \leftarrow x, z \leftarrow x.(x = y \wedge z) \rightarrow (x = y \vee x = z).$$

In a saturated Heyting algebra model[28] $[\![Irreducible(e)]\!] = [\![\top]\!]$ implies that $[\![e]\!]$ is meet irreducible.

Theorem 22. Any finite Heyting algebra H is isomorphic to the lattice of upper sets of the partial order formed by the meet irreducible elements of H.

Proof.
For an element x of H, let $U(x) = \{y \in H : x \leq y\}$ and let $I(x) = \{y \in U(x) : y$ is meet

[27] In a Heyting algebra this is equivalent to requiring the element e to be meet prime: If $y \wedge z \leq x$ then either $y \leq x$ or $z \leq x$. Dually, one may also talk about elements that are join irreducible or join prime.

[28] A model is said to be saturated if $[\![\phi \vee \psi]\!] = 1$ implies that either $[\![\phi]\!] = 1$ or $[\![\psi]\!] = 1$ and if $[\![\exists x.\phi]\!] = 1$ implies that there is an element e in the model such that $[\![x \rhd \phi]\!](e) = 1$.

irreducible}. It then follows from the lemma below that $x = i_1 \wedge ... \wedge i_n$, where $i_1, ..., i_n$ are the elements of $I(x)$. We may also express this in the form $x = \wedge(I(x))$, where $\wedge(F)$ denotes the meet of all the elements of a finite set F.

Suppose next that S is an upper set of the partial order formed by the meet irreducible elements of H (that is, the elements of S are meet irreducible elements of H and from $x \in S$, $x \leq y$ we get $y \in S$). It is then trivial that $I(\wedge(S)) = S$.

Taking the above observations together, we see that the operations I and $\wedge$ are mutually inverse. Since they are also obviously order-preserving, we get an isomorphism of Heyting algebras. □

Lemma 23. An element e in a finite Heyting algebra H can be written as $e = i_1 \wedge ... \wedge i_n$, where $i_1, ..., i_n$ are the meet irreducible elements in H which are greater than or equal to e ($e \leq i_1$, ..., $e \leq i_n$).

Proof.
This can be shown by induction on the number of elements x that satisfy $e < x$.

Assume the theorem has been proved for all elements e' that satisfy $e < e'$. We want to show that it holds for e. If e is meet irreducible then the theorem clearly holds for e, so assume e can be written as $j \wedge k$, where $e < j$ and $e < k$. By the inductive assumption we can write $j = j_1 \wedge ... \wedge j_m$ and that $k = k_1 \wedge ... \wedge k_n$, where $j_1 ... j_m$ and $k_1 ... k_n$ are meet irreducible. we then have $i = j_1 \wedge ... \wedge j_m \wedge k_1 \wedge ... \wedge k_n$. The theorem thus holds for e.

Since the theorem holds when no element x satisfies $e < x$ (that is, when $e = 1$), the theorem holds for all elements e of H. □

To see why it is false that an arbitrary Heyting algebra is determined by the poset of its meet irreducible elements, it is enough to consider a free Heyting algebra with a single generator a. Such a Heyting algebra lacks meet irreducible elements. For example, a is the meet of $a \vee \neg a$ and $\neg\neg a$, and $a \vee \neg a$ is the meet of $(a \vee \neg a) \vee \neg(a \vee \neg a)$ and $\neg\neg(a \vee \neg a)$. In terms of parthood, a has infinitely many parts, all of which can be written as meets of two of their proper parts.

So let us now set up a two-sorted theory in predicate logic where we have two sorts: *Irreducible* and *Object*. We will use order-sorted logic so that we can make the sort *Irreducible* a subsort of the sort *Object* (in symbols we may write this as *Irreducible* $\subseteq$ *Object*).

The one relation we really need is the parthood relation, which we will write P. This relation is defined between arbitrary objects (including irreducibles), something that is customarily written $P : Object \times Object$ (the notation is merely meant to be suggestive; we are not actually forming 'product types' through an operator $\times$). We also have the equality relation $= : Object \times Object$. To say that a variable or a constant is of a particular sort we use the colon notation as in $x{:}Irreducible$ and $x{:}Object$. We also use this notation in combination with the quantifiers (for example, it will be the case that $\exists x{:}Irreducible. \forall y{:}Object.yPx$).

The expression $\imath x : Object. \phi$ will stand for the unique object x such that ϕ holds and it may stand for any object whatever if no such object exists. We also use $\imath x : Irreducible. \phi$ to stand for the unique *Irreducible* such that ϕ holds, but the result need not be an *Irreducible* in the case when the requirement of unique existence fails.

Let us now turn to the problem of interpreting this language inside IpC^2. We will begin by setting up the version of the system that uses minimal logic[29]. The connectives $\wedge$, $\rightarrow$, and $\vee$ are then all translated in the most straight-forward way possible: $(\phi \wedge \psi)^{tr} = (\phi)^{tr} \wedge (\psi)^{tr}$, $(\phi \rightarrow \psi)^{tr} = (\phi)^{tr} \rightarrow (\psi)^{tr}$, and $(\phi \vee \psi)^{tr} = (\phi)^{tr} \vee (\psi)^{tr}$.

For the interpretation of negation we use a constant $\mathcal{R}$ (as in section 11): $(\neg\phi)^{tr} = (\phi)^{tr} \rightarrow \mathcal{R}$. The symbol $\mathcal{R}$ is to be understood in a schematic way: The actual constant used will vary from one application to another. In one application we may have $\mathcal{R} = \mathcal{R}_1$, in another $\mathcal{R} = \mathcal{R}_2$, and so on.

We will interpret objects in such a way that objects are parts of a certain domain D (we treat D as a constant alongside $\mathcal{R}$) and so that they contain $\mathcal{R}$ as a part. Having a domain D is in accordance with the standard practice in predicate logic and requiring $\mathcal{R} \leftarrow x$ to hold for anything that interprets an *Object* makes $\sim_{\mathcal{R}}$ (defined in section 11) coincide with $=$ for *Objects* and additionally provides some separation between objects and propositions[30]. In accordance with this, we make the following definition:

$$Object(x) \leftrightarrow_{def.} (\mathcal{R} \leftarrow x) \wedge (x \leftarrow D).$$

When $x{:}Object$ and $x{:}Irreducible$ are used apart from the quantifiers and descriptions, they function just like predicates and must therefore be handled as such. We set $(x{:}Object)^{tr} = Object((x)^{tr})$ and $(x{:}Irreducible)^{tr} = Object((x)^{tr}) \wedge Irreducible((x)^{tr})$. With these requirements on objects, we can interpret P using $\leftarrow$ and $=$ using $=$. That is: $(xPy)^{tr} = (x)^{tr} \leftarrow (y)^{tr}$ and $(x = y)^{tr} = [(x)^{tr} = (y)^{tr}]$.

[29] By 'minimal logic' is meant the logic described in (Johansson, 1937). As in intuitionistic logic we have $(\neg a) = (a \rightarrow \bot)$, but $\bot$ is no longer a neutral element with respect to $\vee$ (that is, we do not have $(\bot \vee \phi) = \phi$) and may instead be introduced as a propositional constant like any other. Since an object is a neutral element with respect to $\vee$ if and only if it implies everything, we may equivalently say that we do not have the rule $\bot \rightarrow \phi$. In place of the rule $(\bot \vee \phi) \rightarrow \phi$ we can use the weaker $(\bot \vee \phi) \rightarrow \neg\neg\phi$ and we may (depending on the circumstances) also be able to reason at the metalevel that if $\bot \vee \phi$ if provable then ϕ must be provable. In place of the rule $\bot \rightarrow \phi$ we may use the weaker $\bot \rightarrow \neg\phi$.

[30] To be exact, this ensures $\top$ is not an *Object*, and except in the version of the system that uses minimal logic it ensures that no proposition is an object.

We set $(\forall x{:}\mathit{Object}.\,\phi)^{tr} = \forall x \leftarrow D.\, x{:}\mathit{Object} \rightarrow (\phi)^{tr}$. For the existential quantifier we similarly get (recall that we are as yet only considering the minimal logic version) $(\exists x{:}\mathit{Object}.\,\phi)^{tr} = \exists x \leftarrow D.\, x{:}\mathit{Object} \wedge (\phi)^{tr}$. For the expression $(\imath x : \mathit{Object})\phi$ we use $\imath x.\, x{:}\mathit{Object} \wedge \phi$.

To handle quantification and description in connection with the subsort *Irreducible* we apply an extra translation prior to all other translations: $\forall x{:}\mathit{Irreducible}.\,\phi$ translates into $\forall x{:}\mathit{Object}.x{:}\mathit{Irreducible} \rightarrow \phi$, $\exists x{:}\mathit{Irreducible}.\,\phi$ translates into $\exists x{:}\mathit{Object}.x{:}\mathit{Irreducible} \wedge \phi$, and $\imath x{:}Irreducible.\,\phi$ translates into $\imath x{:}Object.\, x{:}\mathit{Irreducible} \wedge \phi$.

For the classical case, we apply the double-negation translation before applying the translation described above.[31] This translates $\phi \vee \psi$ into the classically equivalent $\neg(\neg\phi \wedge \neg\psi)$, which will then translate into $\neg_{\mathcal{R}}(\neg_{\mathcal{R}}\phi \wedge \neg_{\mathcal{R}}\psi)$. For $\exists x : \mathit{Object}.\phi$ we likewise end up with $\neg_{\mathcal{R}}\neg_{\mathcal{R}}[\exists x \leftarrow D.(\mathcal{R} \leftarrow x) \wedge (\phi)^{tr}]$, and double negations will appear in front of all predicates as well as the claim $x : \mathit{Irreducible}$ (the translation mentioned in the previous paragraph is to be applied before the double-negation translation).

The intuitionistic case is similarly handled by a translation. Unlike in the classical case, disjunctions and existential quantifiers are left as they are, but we still need to apply a translation to predicates and to the claim $x : \mathit{Irreducible}$. This time we do not replace ϕ by $\neg\neg\phi$, but instead we replace ϕ with $\bot \vee \phi$. For example, $x : \mathit{Irreducible}$ becomes $\bot \vee (x : \mathit{Irreducible})$.

Note that $(xPy)^{tr}$ and $(x = y)^{tr}$ will differ depending on what logic we are using. One might have naively thought that the formula '$x = y$' had the same meaning in any logic, namely that x is identical to y, but neither with classical nor intuitionistic logic will the truth of '$x = y$' quite succeed in ensuring that x really is identical to y in any model. In the classical case the problem is that IpC^2 will not allow us to infer $x = y$ from $[(x = y) \rightarrow \mathcal{R}] \rightarrow \mathcal{R}$; in the intuitionistic case the problem is that we cannot infer $x = y$ from $\mathcal{R} \vee (x = y)$.

It looks to me as if all that can be said in the case of classical logic is that if we succeed in proving '$x = y$' then x and y will be indistinguishable from the viewpoint of the logic that we are using. In the intuitionistic case we can do better by only looking at models which satisfy the following requirements:

1) $[\![\mathcal{R}]\!] \neq [\![\top]\!]$ and
2) if $[\![\phi_1 \vee \phi_2]\!] = [\![\top]\!]$ then either $[\![\phi_1]\!] = [\![\top]\!]$ or $[\![\phi_2]\!] = [\![\top]\!]$.

With these assumptions we can infer $[\![x = y]\!] = [\![\top]\!]$ from $[\![\mathcal{R} \vee (x = y)]\!] = [\![\top]\!]$, which is

[31] Alternatively, we can assume that $\mathcal{R}$ satisfies $\forall x \leftarrow \mathcal{R}.x \vee (x \rightarrow \mathcal{R})$ and then use the same translation as for intuitionistic logic.

precisely what is needed.[32]

14 Existence Schemata

The axioms of IpC^2 give us few guarantees on what exists. As far as they are concerned, $\forall x.[x = \top \lor x = \bot]$ or $\forall x.x = \top$ might hold.

Things look a bit different, though, when we turn to the 'restricted' fragments of IpC^2 that have been defined elsewhere in this paper (the fragments $\mathrm{IpC}^2\restriction$, $\mathrm{IpC}^2{\equiv}$, and $\mathrm{IpC}^2\restriction{\equiv}$; see sections 4 and 6). After all, neither of the aforementioned formulas ($\forall x.[x = \top \lor x = \bot]$ and $\forall x.x = \top$) can be expressed in any of those fragments.

The difference becomes apparent when we consider a theory T with infinitely many constants and no axioms. With IpC^2, adding the formula $\forall x.x = \top$ as an axiom to T will ensure that all the constants of T are equal, but with the fragments $\mathrm{IpC}^2\restriction$, $\mathrm{IpC}^2{\equiv}$, and $\mathrm{IpC}^2\restriction{\equiv}$ there can be no finite set of axioms which ensures that all the constants of T are equal.

In intuitive terms, it is as if IpC^2 is a more 'dangerous' environment than $\mathrm{IpC}^2\restriction$, $\mathrm{IpC}^2{\equiv}$, and $\mathrm{IpC}^2\restriction{\equiv}$: The truth of a single formula such as $\forall x.[x = \top \lor x = \bot]$ or $\forall x.x = \top$ can cause the whole universe to collapse. We may, however, look for a way to ensure that no such collapse can occur in practice.

In fact, if we add $\forall x.x = \top$ as an axiom then we get what is standardly considered an 'inconsistent' theory, and we may ignore such theories for many purposes. A theory would not normally be considered inconsistent, though, just because it proves $\forall x.[x = \top \lor x = \bot]$, and so it will not be enough for us to restrict our attention to (what we ordinarily count as) consistent theories.

The remedy? We will look at three different 'existence schemata,' all meant to ensure that we are quantifying over lots and lots of objects and not just the two objects suggested by $\forall x.[x = \top \lor x = \bot]$. Each existence schema will be an axiom schema with infinitely many axioms expressed in IpC^2 extended with one extra predicate/connective (more on this below). Moreover, each schema will be closely related to a specific fragment of IpC^2: We will get an existence schema for each of the three fragments $\mathrm{IpC}^2\restriction$, $\mathrm{IpC}^2{\equiv}$, and $\mathrm{IpC}^2\restriction{\equiv}$.

But if we extend IpC^2 with an extra predicate/connective, are we not moving into a stronger logic? Technically yes, but section 15 will discuss how the extra predicates/connectives can

[32] That we can get a complete semantics of IpC^2 using only models that satisfy 2) is proved in (Gabbay, 1981).

be defined in IpC^2 itself (all we need are ordinary constants).

In order to make the presentation of the existence schemata look neater, I will state them using quantifiers restricted to objects satisfying a particular predicate:
$\forall_P x.\phi \leftrightarrow_{def.} \forall x.P(x) \rightarrow \phi$, and
$\exists_P x.\phi \leftrightarrow_{def.} \exists x.P(x) \wedge \phi$.

Let me now give the existence schema corresponding to $\text{IpC}^2\!\restriction$. Each axiom of the schema will take the following form:
$\forall_{Small} x_1, ..., x_m.\phi \rightarrow \exists_{Small} y_1, ..., y_n.\chi$.

Here the formula $\forall_{Small} x_1, ..., x_m.\phi \rightarrow \exists_{Small} y_1, ..., y_n.\chi$ must include no constants[33] or free variables.

Small is the extra connective/predicate that we are adding to IpC^2 (see above), and it is required to satisfy the following axioms[34]:
$\forall x, y.[Small(x) \wedge (x \rightarrow y)] \rightarrow Small(y)$,
$\forall x, y.[Small(x) \wedge Small(y)] \rightarrow Small(x \wedge y)$, and
$\forall x, y.[Small(x) \wedge Small(y)] \rightarrow Small(x \vee y)$.

To see whether the formula $\forall_{Small} x_1, ..., x_m.\phi \rightarrow \exists_{Small} y_1, ..., y_n.\chi$ is actually an axiom, we form two theories T and T'. The theory T is to have m constants $c_1, ..., c_m$ while the theory T' is to extend the theory T with n constants $d_1, ..., d_n$. The sole axiom of T is to be $\phi[c_1/x_1, ..., c_m/x_m]$ and theory T' is to additionally include the axiom $\chi[c_1/x_1, ..., c_m/x_m, d_1/y_1, ..., d_n/y_n]$. At this point we ask ourselves whether the theories T and T' are theories in the logic $\text{IpC}^2\!\restriction$ and whether T' is a *conservative* extension of T in $\text{IpC}^2\!\restriction$. When the answer to both questions is yes, $\forall_{Small} x_1, ..., x_m.\phi \rightarrow \exists_{Small} y_1, ..., y_n.\chi$ is to be included as an axiom of the existence schema.[35]

[33] An exception to this rule must be made if constants are used to define the predicate *Small*; see section 15.

[34] In section 15 both axioms will be turned into provable theorems. The reason for the redundancy among the axioms will become clear below when we consider the other two existence schemata.

[35] Note also that the existence schema here described involves an infinite set S of axioms which is not a recursively enumerable set. To see that S is not recursively enumerable it is enough to consider the set S' of axioms of the form $\forall_{Small} x_1.[\forall x_2 \leftarrow x_1.x_2 = x_1 \vee x_2 = \top] \rightarrow \exists_{Small} y.\psi$. It easy to see that if S were recursively enumerable then S' would also be recursively enumerable, so it suffices to show that S' is not recursively enumerable. But the axiom $\forall_{Small} x_1.[\forall x_2 \leftarrow x_1.x_2 = x_1 \vee x_2 = \top] \rightarrow \exists_{Small} y.\psi$ is in S' if and only if it is impossible to derive $x_1 = \top$ from ψ.

From (Löb, 1976) we know that classical predicate logic can be faithfully embedded in IpC^2, and the question of which axioms are in S' can therefore be turned into the question of which theories in classical predicate logic are consistent and which ones are not (although the formula ψ is limited to the fragment $\text{IpC}^2\!\restriction$ and cannot refer to $\bot$, this is no real problem here since we can use the variable y in place of $\bot$ as 'the object that contains everything'). But it is well-known that the set of consistent theories in classical predicate logic is not recursively enumerable (if it were, that would give us a computer program that solves

This completes the specification of the existence schema corresponding to $\text{IpC}^2\restriction$, but what does it amount to? We may begin by considering the case where $m = 0$. The theory T will then lack constants, and ϕ will have to be $\top$ (after all, it can involve no constants or variables and it has to be a formula in $\text{IpC}^2\restriction$). An axiom with $m = 0$ can thus be expressed as follows:

$\exists_{Small} y_1, ..., y_n . \chi$.

Here χ must involve no constants and no variables other than $y_1, ..., y_n$, but apart from this it can be any formula that is expressible in $\text{IpC}^2\restriction$ (since T is the empty theory, we need not worry that T' may fail to be a conservative extension of T). For example, it might be $\top$, in which case we are simply told (in the axiom $\exists_{Small} y_1, ..., y_n . \chi$) of the existence of n (small) objects $y_1, ..., y_n$, or it might be $(y_1 \to y_1) \wedge (y_1 \to y_2) \wedge ... \wedge (y_1 \to y_n)$, in which we are told of the existence of n (small) objects $y_1, ..., y_n$ such that one of the objects (namely y_1) contains the others.

Let us now turn to the case where m is not 0 and where ϕ says something non-trivial about $x_1, ..., x_m$. χ must then be such that T' becomes a conservative extension of T in the logic $\text{IpC}^2\restriction$. For example, if $m = 2$ and $\phi = \top$ then χ must not say that $x_1 = x_2$ as T' would then come to say that $c_1 = c_2$, violating conservativity. Nor can it say, for example, that $x_1 \to x_2$ or $x_1 \vee x_2$ holds; this would again lead to a violation of conservativity.

Note, though, that the conservativity requirement applies to the logic $\text{IpC}^2\restriction$. It is *not* required for T' to be a conservative extension of T in the logic IpC^2. This seems to allow the existence schema to tell us something non-trivial about the parts of $\bot$, an object which we are unable to refer to in the logic $\text{IpC}^2\restriction$. For example, from the axiom $\forall_{Small} x_1, x_2 . \exists_{Small} y_1 . (y_1 \to x_1) \wedge (y_1 \to x_2)$ we learn that for any two small parts x_1 and x_2 of $\bot$ we can find a small part y_1 of $\bot$ which contains x_1 as well as x_2.

So what distinguishes a 'small' object from an arbitrary object? The existence schema is meant to be such that $\bot$ cannot be proved to be 'small.' In set theory it is common to distinguish between small and large sets[36] and we are here making a somewhat analogous distinction in IpC^2. With the existence schema any small object seems to be such that one can find a larger object that is still small, but with $\bot$ we can prove that any object which includes it as a part must be equal to it.

To ensure that the objects guaranteed to exist through the existence schema are actually different (rather than one and the same object), we may use the methods of section 11. This may at first seem to require us to relax the requirement that ϕ and χ contain no constants – we want to be able to refer to the constant $\mathcal{R}$ – but we can easily work around

the halting problem, but it is well-known that no such computer program exists).

[36]This distinction is important in category theory, where a category is called 'small' if it has a small set of objects as well as a small set of morphisms. See (Mac Lane, 1998).

this restriction by using a variable (say x_1) in place of $\mathcal{R}$ in the formulas ϕ and χ. When time comes to apply the axiom involving ϕ and χ, we substitute $\mathcal{R}$ for this variable, and everything will be as if the axiom had been able to refer directly to $\mathcal{R}$.[37]

So far we have only considered the existence schema that is based on the logic $\mathrm{IpC}^2\!\restriction$, but the remaining existence schemata are to a large extent analogous.

With the existence schema corresponding to $\mathrm{IpC}^2\!\equiv$, an axiom will take the following form:
$\forall_{Large} x_1, ..., x_m.\phi \to \exists_{Large} y_1, ..., y_n.\chi$.

As before, this formula must include no constants[38] or free variables.

This time, the primitive that we are adding to IpC^2 is called *Large*, and it is required to satisfy the following axioms:
$\forall x, y.[Large(x) \wedge (x \leftarrow y)] \to Large(y)$,
$\forall x, y.[Large(x) \wedge Large(y)] \to Large(x \wedge y)$, and
$\forall x, y.[Large(x) \wedge Large(y)] \to Large(x \vee y)$.

To check whether $\forall_{Large} x_1, ..., x_m.\phi \to \exists_{Large} y_1, ..., y_n.\chi$ is indeed an axiom, we again consider a theory T with constants $c_1, ..., c_m$ and a theory T' which adds the constants $d_1, ..., d_n$. T is to include $\phi[c_1/x_1, ..., c_m/x_m]$ as its sole axiom and T' is to additionally include the axiom $\chi[c_1/x_1, ..., c_m/x_m, d_1/y_1, ..., d_n/y_n]$. What has changed is that T and T' now need to be theories in the logic $\mathrm{IpC}^2\!\equiv$ (so they are allowed to mention $\bot$ but must not mention $\top$; see section 6 for the exact definition), and that T' is to be a conservative extension of T in $\mathrm{IpC}^2\!\equiv$ (rather than $\mathrm{IpC}^2\!\restriction$).

For the existence schema corresponding to $\mathrm{IpC}^2\!\restriction\!\equiv$, an axiom will take the following form:
$\forall_{MiddleSized} x_1, ..., x_m.\phi \to \exists_{MiddleSized} y_1, ..., y_n.\chi$.

Once again, there are to be no constants[39] or free variables.

For the predicate/connective *MiddleSized* we have the following axioms:
$\forall x, y, z.[MiddleSized(x) \wedge MiddleSized(y) \wedge (z \to x) \wedge (z \leftarrow y)] \to MiddleSized(z)]$,
$\forall x, y.[MiddleSized(x) \wedge MiddleSized(y)] \to MiddleSized(x \wedge y)$, and
$\forall x, y.[MiddleSized(x) \wedge MiddleSized(y)] \to MiddleSized(x \vee y)$.

We form theories T and T' as before, the only difference being that they are now to be theories in the logic $\mathrm{IpC}^2\!\restriction\!\equiv$ (where neither $\top$ nor $\bot$ can be referred to) and that T' must

[37] Note that this method requires us to assume that $\mathcal{R}$ is a small object.

[38] Again, an exception to this rule must be made if constants are used to define the predicate *Large*; see section 15.

[39] Once again, an exception needs to be made when constants are used in the definition of *MiddleSized* itself.

be a conservative extension of T in this logic.

This time there will be no axiom of the form $\exists_{MiddleSized} y_1, ..., y_n.\chi$ (no axiom that unconditionally asserts the existence of an object), since there is nothing ϕ could be in that case (without constants, variables, $\top$ and $\bot$, there is nothing left for ϕ to be).

An important question that has not been considered in this section is whether the axiom schemata here described are non-trivial. Could it not be that all small/large/middle-sized objects are provably equal?

This paper will leave the non-triviality of the existence schemata as an unproved conjecture.

The question of which existence schema to prefer will be discussed in the concluding section of this paper.

15 Tagged objects

The axiom schemata presented in the previous section all made use of a predicate/connective which we did nothing to define (the predicates in question were called *Small*, *Large*, and *MiddleSized*, respectively). Instead, we thought of the predicate as an undefined primitive that we added to the primitives of IpC^2.

We will now look at how predicates satisfying the required axioms can be defined in IpC^2 itself. In fact, we will only look in detail at the predicate *Small*, but definitions for the other two predicates will be given, and the proofs needed for these predicates will be sketched.

So what exactly is a small object? Well, one thing we do not want is for $\bot$ to come out small, and we also do not want for a small object to be the same thing as a part of some particular object o. After all, the existence schema of the previous section seemed to tell us that for any small object one could find a larger object which was itself small.

Our definition of $Small(x)$ will be as follows:
$Small(x) \leftrightarrow_{def.} o_1 \vee (o_2 \rightarrow x)$.

For the predicates *Large* and *MiddleSized* we may instead use the following defintions[40]:
$Large(x) \leftrightarrow_{def.} o_1 \vee (x \rightarrow o_3)$.

[40] Note that the definitions are not meant to be used together. In that case one would want to rename *Small* as *NotLarge* and *Large* as *NotSmall*. It seems odd to say that an object is middle-sized if it is large and small at the same time, but reasonable to say that an object is middle-sized if it satisfies *NotLarge* as well as *NotSmall*.

$$MiddleSized(x) \leftrightarrow_{def.} o_1 \vee [(o_2 \to x) \wedge (x \to o_3)].$$

The definition of $Small(x)$ depends on two constants o_1 and o_2, which need to be in theories that employ the predicate $Small(x)$.

Note that if we were to add the axiom $o_1 = \top \wedge o_2 = \bot$ then $Small(x)$ would hold for any object x. Thus, if we start from a theory T without the constants o_1 and o_2 and then extend it into a theory T' which is like T except that it adds the constants o_1 and o_2 and axioms of the form $Small(\phi)$, then we have that the theory T' is a conservative extension of T.

Two basic properties of *Small* are as follows:

Theorem 24. $[Small(x) \wedge (x \to y)] \to Small(y)$

Proof.
Assume $Small(x)$ and $x \to y$. By definition, $Small(x)$ means $o_1 \vee (o_2 \to x)$. It easy to see that $Small(y)$ follows from o_1 as well as well as $o_2 \to x$, and so it has to hold. □

Theorem 25. $[Small(x) \wedge Small(y)] \leftrightarrow Small(x \wedge y)$

Proof. The right-to-left direction is a consequence of theorem 24.

In order to prove the left-to-right direction, assume $Small(x)$ and $Small(y)$. That is, we have $o_1 \vee (o_2 \to x)$ and $o_1 \vee (o_2 \to y)$. Through reasoning by cases we obtain $o_1 \vee [(o_2 \to x) \wedge (o_2 \to y)]$. But $(o_2 \to x) \wedge (o_2 \to y)$ is equivalent to $o_2 \to (x \wedge y)$. Thus, we get $o_1 \vee [o_2 \to (x \wedge y)]$, which is precisely what $Small(x \wedge y)$ says. □

To get a rough idea of how *Small* works we may note that $o_1 = \bot$ gives us $Small(x) \leftrightarrow (o_2 \to x)$ while $o_1 = \top$ instead gives us $Small(x) \leftrightarrow \top$ (which we can also think of as $\bot \to x$). From this we may expect that before we have added one of $o_1 = \bot$ and $o_1 = \top$, the objects that are provably small can include objects that are not part of o_2 while failing to include any object whatever. Theorem 30 below confirms this.

Theorem 26. $Small(o_2)$

Proof. Trivial. □

Lemma 27. $Small(\inf x \,\text{s.t.}\, Small(x))$ is equivalent to $o_1 \vee \forall x.[Small(x) \leftrightarrow (o_2 \to x)]$.

Proof. By definition, $\mathit{Small}(\inf x \text{ s.t. } \mathit{Small}(x))$ is equivalent to $o_1 \vee (o_2 \rightarrow \inf x \text{ s.t. } \mathit{Small}(x))$. It therefore suffices to prove that $o_2 \rightarrow \inf x \text{ s.t. } \mathit{Small}(x)$ is equivalent to $\forall x.[\mathit{Small}(x) \leftrightarrow (o_2 \rightarrow x)]$.

Assume $o_2 \rightarrow \inf x \text{ s.t. } \mathit{Small}(x)$. Then theorem 2 gives us $\mathit{Small}(x) \rightarrow (o_2 \rightarrow x)$. But since we have $\mathit{Small}(o_2)$ (by theorem 26), we get $\forall x.[\mathit{Small}(x) \leftrightarrow (o_2 \rightarrow x)]$.

Assume instead that $\forall x.[\mathit{Small}(x) \leftrightarrow (o_2 \rightarrow x)]$ holds. If, additionally, we assume o_2 then we clearly have $\mathit{Small}(x) \leftrightarrow x$ from which we get that $\inf x \text{ s.t. } \mathit{Small}(x)$ is equivalent to $\inf x \text{ s.t. } x$, which is simply $\top$. Thus, $o_2 \rightarrow \inf x \text{ s.t. } \mathit{Small}(x)$. □

Lemma 28. $\forall x.[\mathit{Small}(x) \leftrightarrow (o_2 \rightarrow x)]$ is equivalent to $\neg(o_1 \wedge o_2)$.

Proof. It is easily seen that $\forall x.[\mathit{Small}(x) \leftrightarrow (o_2 \rightarrow x)]$ is equivalent to $\forall x.[\mathit{Small}(x) \rightarrow (o_2 \rightarrow x)]$ (use theorem 26 and theorem 24).

Using the definition of *Small* we get that $\forall x.[\mathit{Small}(x) \rightarrow (o_2 \rightarrow x)]$ is equivalent to $\forall x.[(o_1 \vee (o_2 \rightarrow x)) \rightarrow (o_2 \rightarrow x)]$. This can clearly be rewritten as $\forall x.[o_1 \rightarrow (o_2 \rightarrow x)] \wedge [(o_2 \rightarrow x) \rightarrow (o_2 \rightarrow x)]$, which simplifies to $\forall x.[(o_1 \wedge o_2) \rightarrow x]$. This in turn is easily seen to be equivalent to $(o_1 \wedge o_2) \rightarrow \forall x.x$, which we may also write as $\neg(o_1 \wedge o_2)$. □

Theorem 29. $\mathit{Small}(\inf x \text{ s.t. } \mathit{Small}(x))$ is equivalent to $o_1 \vee \neg(o_1 \wedge o_2)$.

Proof. Combine lemma 27 with lemma 28. □

Theorem 30. If a theory T of $\mathrm{IpC}^2\restriction$ that lacks the constants o_1 and o_2 is extended into a theory T' that adds the constants o_1 and o_2 and axioms of the form $\mathit{Small}(\phi)$, then the object $(\inf x \text{ s.t. } \mathit{Small}(x))$ is not provably small in T'.

Proof. Suppose on the contrary that the formula $\mathit{Small}(\inf x \text{ s.t. } \mathit{Small}(x))$ could be proved in T'.

The proof can only use a finite subset of the axioms of T'. Thus, only finitely many of the axioms of the form $o_1 \vee (o_2 \rightarrow e)$ could be involved in the proof. We may write these axioms as $\mathit{Small}(e_1)$, ..., $\mathit{Small}(e_n)$.

But these axioms are obviously all consequences of $o_2 \rightarrow (e_1 \wedge ... \wedge e_n)$ (first derive $o_2 \rightarrow e_k$ and then $\mathit{Small}(e_k)$), and so $\mathit{Small}(\inf x \text{ s.t. } \mathit{Small}(x))$ would have to be provable in the theory T^* which combines the axiom $o_2 \rightarrow (e_1 \wedge ... \wedge e_n)$ with the axioms of T.

Now note that if $\mathit{Small}(\inf x \text{ s.t. } \mathit{Small}(x))$ were provable in T^* then it would also be provable in the theory T^{**} which adds the axiom $\top \rightarrow o_2$. But theorem 29 tells us that

$Small(\inf x \text{ s.t. } Small(x))$ is equivalent to $o_1 \vee \neg(o_1 \wedge o_2)$, which is provably equivalent to $o_1 \vee \neg o_1$ in T^{**}.

Thus, if $Small(\inf x \text{ s.t. } Small(x))$ were provable in T^* then $o_1 \vee \neg o_1$ would be provable in T^{**}. But T^{**} is a theory in the logic $\text{IpC}^2\restriction$ in which no axiom mentions o_1, and so there is obviously no way that $o_1 \vee \neg o_1$ could be provable in it (use induction on the length of the proof). Contradiction. □

Let us finally consider how things differ when we use one of the predicates *Large* or *MiddleSized* instead of the predicate *Small*. For *Large* we have:

Theorem 31. 1) $[Large(x) \wedge (y \to x)] \to Large(y)$,

2) $[Large(x) \wedge Large(y)] \leftrightarrow [Large(x \vee y)]$,

3) $Large(o_3)$,

4) $\forall x.[Large(x) \leftrightarrow (x \to o_3)]$ is equivalent to $o_1 \to o_3$,

5) $Large(\sup x \text{ s.t. } Large(x))$ is equivalent to $o_1 \vee \forall x.[Large(x) \leftrightarrow (x \to o_3)]$,

6) $Large(\sup x \text{ s.t. } Large(x))$ is equivalent to $o_1 \vee (o_1 \to o_3)$,

7) If a theory T of $\text{IpC}^2{\equiv}$ that lacks the constants o_1 and o_3 is extended into a theory T' that adds the constants o_1 and o_3 and axioms of the form $Large(\phi)$, then the object $(\sup x \text{ s.t. } Large(x))$ is not provably large in T'.

Proof. 1) Analogous to the proof of theorem 24.

2) Analogous to the proof of theorem 25.

3) Trivial.

4) Like the proof of lemma 28. However, while $\forall x.[o_1 \to (o_2 \to x)]$ simplifies to $\neg(o_1 \wedge o_2)$, $\forall x.[o_1 \to (x \to o_3)]$ instead simplifies to $o_1 \to o_3$.

5) Use theorem 6 along with 3).

6) Combine 4) with 5).

7) We can use the proof of theorem 30 with a few changes:

- Note that only finitely many axioms of the form $Large(\phi)$ could be involved in a proof of $Large(\sup x \text{ s.t. } Large(x))$ and that $Large(e_1)$, ..., $Large(e_n)$ are all consequences of $(e_1 \vee ... \vee e_n) \to o_3$,

- let T^* extend T with the axiom $(e_1 \vee ... \vee e_n) \rightarrow o_3$,
- let T^{**} be the theory that extends T^* with the axiom $o_3 \rightarrow \bot$,
- note that $Large(\sup x \,\text{s.t.}\, Large(x))$ simplifies to $o_1 \vee \neg o_1$ with the theory T^{**}, and
- note that if $Large(\sup x \,\text{s.t.}\, Large(x))$ were provable in T^* then it would be provable in T^{**}, which is absurd since T^{**} is a theory in $\text{IpC}^2{\equiv}$ in which no axiom mentions o_1.

□

Since $MiddleSized(x)$ is just the conjunction of $Small(x)$ and $Large(x)$, we can apply all of the preceding theorems to middle-sized objects. In particular, by combining theorem 24 with part 1) of theorem 31 we get:
$\forall x, y, z.[MiddleSized(x) \wedge MiddleSized(y) \wedge (z \rightarrow x) \wedge (z \leftarrow y)] \rightarrow MiddleSized(z)]$.

Similarly, we need only combine our results for *Small* and *Large* to get:
$\forall x, y.[MiddleSized(x) \wedge MiddleSized(y)] \rightarrow MiddleSized(x \wedge y)$, and
$\forall x, y.[MiddleSized(x) \wedge MiddleSized(y)] \rightarrow MiddleSized(x \vee y)$.

16 Conclusions

This paper has discussed IpC^2 from a number of viewpoints:

1. Syntax (section 2) and semantics (sections 7-10),
2. the infimum/description operator (defined in section 3 and used in sections 13 and 15),
3. fragments of IpC^2 where $\bot$ and/or $\top$ cannot be referred to (introduced in sections 4 and 6, further discussed in section 9, and of central importance in section 14),
4. the question of what exists (sections 14 and 15), and
5. how it can be used to reason about the parts of objects (sections 5, 11, 12, and 13).

However, the main goal of the paper has been to explore IpC^2 as a potential foundation for mathematics. Let us therefore now try to assess the merits and demerits of using IpC^2 as a foundation of mathematics.

With any foundation of mathematics we can ask what basic concepts it starts from. As axiomatized in this paper, IpC^2 is based on two primitives: $\rightarrow$ and $\forall$. Throughout the paper we have regarded $\rightarrow$ as expressing parthood, and we have encountered no problems at all when doing so. That we have sometimes read $\rightarrow$ as 'implies' does in no way contradict this, but merely means that proposition a implies proposition b if and only if b is part of a, an idea that seems to work excellently.

It thus looks as if IpC^2 is a theory which has parthood as one of its primitives. The reader who finds this hard to believe should note that it is well-known that intuitionistic logic can be given a topological interpretation (see section 10) and that it is natural to see topology as a subject that helps us reason about parthood.

The parthood relation would seem to be a very natural relation to put at the bottom of ones theorizing about mathematics. It may be thought to underlie not only the idea of logical consequence, but also the idea of a collection (it seems a natural thought that the elements and subsets of a set are parts of the set). It may be noted in this connection that Bertrand Russell tried to give the parthood relation a fundamental role in the foundations of mathematics before he learned how to do without it.[41]

The parthood relation by itself is not enough, though, if we want to say what the parts of an object are. For this we also need the other primitive of IpC^2, the universal quantifier ($\forall$). We have seen in this paper how these two primitives are enough to say what exactly the parts of an object are (see sections 12 and 13, for example).

Perhaps it would be more natural to replace the single primitive $\rightarrow$ with the pair of primitives $=$ and $\wedge$ (are these not simpler?).[42] But $\rightarrow$ has the nice feature that it helps us express at the object level what $\vdash$ expresses at the metalevel, and it is also nice in that it can be seen as expressing 'the parthood relation.'

The reader may prefer, though, to replace $\forall$ and $\rightarrow$ with the construction $\inf \psi[x]\,\text{s.t.}\,\phi$ that was defined in section 3. The construction $\inf \psi[x]\,\text{s.t.}\,\phi$ seems fairly simple and natural (to some extent, if only that, it would seem to be grasped by anyone who knows how to solve optimization problems) and regarding it as basic seems in line with the idea that

[41] Russell's 1899-1900 draft of *The Principles of Mathematics* consists of seven parts, the second of which is entitled 'Whole and Part' (see (Russell, 1993)). That Russell considered 'the theory of whole and part' to be fundamental is shown by the following passage: 'The theory of elementary arithmetical addition and ratio [...] is not prior to the elementary theory of whole and part, but coordinate with it.' (*ibid.*, p. 35). In a revealing note from October 1900 Russell writes that 'Peano's distinction of $\supset$ and ϵ shows whole and part to be different from implication. The former is primarily implication, the latter gives the relation of simple part to whole.' (*ibid.*, p. 10). The word 'primarily' here is perhaps explained by the fact that Russell recognized not only 'material implications' but also 'formal implications'; formal implications were encountered in section 3 of the present paper (see footnote 11).

[42] In the setting of higher order logic, a proposal of this kind is considered in (Scott, 1979, pp. 692-693).

adjoints/Galois connections play a basic role in logic and mathematics.[43]

In any case, out of our basic primitives we are able to define $\wedge$, $\vee$, $\exists$, and $=$ (which we might also write $\leftrightarrow$) as well as a description operator (see section 3).[44] IpC^2 is also economical in its use of a single type where other systems have one sort for propositions and another sort for 'ordinary objects.' A consequence of this is that there is no difference between functions and relations in IpC^2 (except, perhaps, a change in point of view), and so we are again able to keep down the number of basic concepts.

So much for the question of basic concepts. Another important question to ask about a foundation of mathematics is whether it is based on classical logic, intuitionistic logic, or some other logic. We saw in section 13 that as we reason about parts of objects in IpC^2 we can get either classical logic or intuitionistic logic depending on how we do things. Of course, 'IpC^2' is an abbreviation for 'second-order *intuitionistic* propositional logic,' but this paper has stayed away from the assumption that all objects quantified over in IpC^2 are to be thought of as 'propositions' as well as the (closely related) assumption that $\bot$ is to be understood as an absurd proposition (instead, $\bot$ has been understood as an object that 'contains everything'). Sections 11 and 13 used a perfectly ordinary constant $\mathcal{R}$ (rather than $\bot$) where an 'absurd proposition' was needed, and depending on how we choose $\mathcal{R}$ we may or may not get the law $\forall x \leftarrow \mathcal{R}. (x = \top) \vee (x = \mathcal{R})$ (which we can think of as a version of 'the law of excluded middle' (which holds in classical logic but which need not hold in intuitionistic logic)).

So as understood here, IpC^2 is agnostic on whether we should perform our reasoning using classical logic or intuitionistic logic. How about the fact that IpC^2 is a 'second-order' system? Do we really want to use second-order logic in our foundation of mathematics? I believe the naming of IpC^2 as 'second-order' is misleading since we are not able to quantify over functions or relations. Section 13 showed how we can approach IpC^2 using the syntax of ordinary first-order logic, and it seems arguable that IpC^2 is more 'first-order' than 'second-order' in nature.

This leads us to a different question: Is it not an important limitation of IpC^2 that we are unable to quantify over functions and relations? Perhaps, but it may be possible to find reasonable replacements in practice. For example, an idea that seems worth exploring is what can be done in IpC^2 using objects of the form $\inf x \,\text{s.t.}\, \mathit{Small}(x) \wedge \phi(x)$ (in words: 'the smallest object that contains all small objects satisfying property ϕ'; see section 15 for the definition of *Small*), where ϕ can be any property. Unlike the property ϕ, this object can be quantified over in IpC^2, and it could sometimes be a useable replacement for the

[43] Cf. (Lawvere, 1969).

[44] One thing that will not work is to try to use $\exists$ instead of $\forall$ as a basic primitive of IpC^2. (Zdanowski, 2009) contains a proof that the system one gets with $\exists$ in place of $\forall$ is strictly weaker than IpC^2.

property ϕ.

Let us now turn from the logic we use to reason about objects to those objects themselves. Can we represent objects encountered in mathematics in a natural way using IpC^2? As mentioned in the introduction, this paper will stop short of addressing this question. It may be hoped, though, that the 'topological' nature of IpC^2 (see section 10) in combination with the fact that topology has lots of applications throughout mathematics can help make IpC^2 a very natural system for the formalization of mathematics.

One aspect of IpC^2 that may be a strength and a weakness at the same time is that all objects have their parts arranged as the elements of a Heyting algebra. This does not always seem to match actual mathematical objects very well. For example, the subspaces of a vector space do not form a Heyting algebra but a non-distributive lattice.

However, one can find ways to work around the restriction to Heyting algebras (see section 13), and the fact that $a \Rightarrow b$ is defined regardless of what a and b are makes Heyting algebras much easier to work with than non-distributive lattices.

By itself, IpC^2 may be considered too weak to serve as a foundation of mathematics. Far from guaranteeing the existence of a 'mathematical universe,' IpC^2 is consistent with only two objects existing (just add $\forall x.(x = \top) \vee (x = \bot)$ as an axiom). However, section 14 showed how one may extend IpC^2 with existence schemata that seem to give us universes of the kind one may want from a foundation of mathematics.

I think any reader will agree with me after a careful read of section 14 that each existence schema arises in a very natural way from a fragment of IpC^2 (fragments where $\top$ and/or $\bot$ cannot be referred to).[45] There would seem to be nothing at all 'arbitrary' about them.

But how can we make sense of the fact that we are getting three existence schemata rather than one? The difference between the schemata can be seen as lying in the question of whether objects should be infinitely extensible, infinitely divisible, or both. With the existence schema corresponding to the logic $\mathrm{IpC}^2\restriction$, objects are infinitely extensible, but it is possible for an object to be an indivisible atom. By contrast, with the other two schemata there is no such thing as an atom. I personally believe the most useful and basic of the existence schemata to be the one where objects are always extensible as well as divisible (the one corresponding to the logic $\mathrm{IpC}^2\restriction\equiv$).

[45] We can perhaps informally think of the existence schema corresponding to $\mathrm{IpC}^2\restriction$ as being what we get if we start with $\mathrm{IpC}^2\restriction$ and then add $\bot$ in a 'non-destructive' way. Similarly, we could think of the existence schema corresponding to $\mathrm{IpC}^2\equiv$ as being what we get when we start with $\mathrm{IpC}^2\equiv$ and then add $\top$ in a non-destructive way, and we could think of the existence schema corresponding to $\mathrm{IpC}^2\restriction\equiv$ as being what we get when we start with $\mathrm{IpC}^2\restriction\equiv$ and add $\top$ as well as $\bot$ in a non-destructive way.

Note that it remains to be proved that the existence schemata are actually non-trivial.

An area where IpC^2 may score poorly is in 'user-friendliness.' However, as illustrated in section 13, one need not reason in IpC^2 directly but can use another language on top of it. In this way, one can hopefully get the best of two worlds.

IpC^2 may also score poorly when it comes to 'constructivity.' The axiom schemata discussed in section 14 are not recursively enumerable (see footnote 35), and (at least as I have done things in this paper) we do not have the existence property: The fact that we can prove $\exists x.\phi$ does not mean that we are able to refer to a specific object χ such that $\phi[\chi/x]$ holds.

17 References

Blackburn, P, de Rijke, M., and Venema, Y. (2001), *Modal Logic, vol. 53 of Cambridge Tracts in Theoretical Computer Science.* Cambridge University Press.

Church, A. (1956). *Introduction to Mathematical Logic.* Princeton University Press, New Jersey.

Esakia, L. (2004). Intuitionistic logic and modality via topology. *Annals of Pure and Applied Logic*, **127**, 155–170.

Gabbay, D. M. (1981). Semantical investigations in Heyting's intuitionistic logic. *Synthese Library*, **148**. Reidel, Dordrecht.

Gaifman, H. (1964). Infinite Boolean polynomials. I. *Fund. Math.*, **54**, 230-250.

Girard, J.-Y. (1972). *Interprétation fonctionnelle et élimination des coupures de l'arithmétique d'ordre supérieur* (Ph.D. thesis) (in French). Université Paris 7.

Girard, J.-Y., Lafont, Y., and Taylor, P. (1989). *Proofs and Types.* Cambridge University Press.

Gödel, K. (1933). Eine Interpretation des intuitionischen Aussagenkalküls. *Ergebnisse Math. Colloq.*, **4**, 39–40.

Hindley, J. and Seldon, J. (eds) (1980). *To H.B. Curry: Essays on Combinatorial Logic, Lambda Calculus and Formalism.* Academic Press, London.

Howard, W. (1980). *The formulae-as-types notion of construction*, in J. Hindley and J. Seldin (eds), (Hindley and Seldon, 1980), 479–491.

Johansson, I. (1937). Der Minimalkalkül, ein reduzierter intuitionistischer Formalismus. *Compositio mathematica 4*, 119-136.

Johnstone, P. T. (2002). *Sketches of an Elephant: A Topos Theory Compendium.* Oxford University Press.

Lawvere, F. W. (1969). Adjointness in Foundations. *Dialectica*, **23 (3-4)**, 281-296.

Löb, H. M. (1976). Embedding first order predicate logic in fragments of intuitionistic logic. *Journal of Symbolic Logic*, **41**, 705–18.

Mac Lane, S. (1998). *Categories for the Working Mathematician.* Graduate Texts in Mathematics 5 (second ed.). Springer.

McKinsey, J., Tarski, A. (1948). Some theorems about the sentential calculi of Lewis and Heyting. *Journal of Symbolic Logic*, **13**, 1–15.

Pavičić, M., Megill, N. D. (1999). Non-Orthomodular Models for Both Quantum Logic and Standard Classical Logic: Repercussions for Quantum Computers, *Helv. Phys. Acta*, **72**, 189-210.

Pitts, A. M. (1992). On an Interpretation of Second Order Quantification in First Order Intuitionistic Propositional Logic. *Journal of Symbolic Logic.* **57(1)**, 33-52.

Prawitz, D. (1965). *Natural deduction.* Almqvist & Wiksell, Stockholm.

Reynolds, J. (1974). Towards a Theory of Type Structure. *Programming Symposium, Proceedings Colloque sur la Programmation*, 408-423.

Ruitenburg, W. (1991). The unintended interpretations of intuitionistic logic. *Perspectives on the History of Mathematical Logic.* Birkhäuser, Boston. 134-160.

Russell, B. (1993). *The Collected Papers of Bertrand Russell, Volume 3: Toward the "Principles of Mathematics," 1900–02*, ed. G. H. Moore. Routledge, London and New York.

Scott, D. S. (1979). Identity and existence in intuitionistic logic, in M. P. Fourman, C. J. Mulvey, D. S. Scott (eds.), Applications of Sheaves, *Lecture Notes in Mathematics* **753**, Springer, 660–696.

Scott, D. S. (2008). The Algebraic Interpretation of Quantifiers: Intuitionistic and Classical. *Andrzej Mostowski and Foundational Studies*, ed. Ehrenfeucht, A., Mostowski, A., Marek, V. W., and Srebrny, M., 289-312.

Solovay, R. (1966). New proof of a theorem of Gaifman and Hales. *Bull. Amer. Math. Soc.*, **72(2)**, 282-284.

Sørensen, M. H. and Urzyczyn, P. (2010). A Syntactic Embedding of Predicate Logic into Second-Order Propositional Logic. *Notre Dame Journal of Formal Logic*, **51(4)**, 457-473.

Varzi, A. (2016). Mereology, in *The Stanford Encyclopedia of Philosophy (Spring 2016 Edition)*, ed. E. N. Zalta, URL = http://plato.stanford.edu/archives/spring2016/entries/mereology/.

Zdanowski, K. (2009). On second order intuitionistic propositional logic without a universal quantifier. *The Journal of Symbolic Logic*, **74**, 157–67.