(to appear in *Philosophical Issues* **16** (2006))

Indicative versus Subjunctive Conditionals, Congruential versus Non-Hyperintensional Contexts

Timothy Williamson

§0. A familiar if obscure idea: an indicative conditional presents its consequent as holding in the actual world on the supposition that its antecedent so holds, whereas a subjunctive conditional merely presents its consequent as holding in a world, typically counterfactual, in which its antecedent holds. Consider this pair:

(1) If you had starved as a child, you would have been as tall now as you actually are.

(2) If you starved as a child, you are as tall now as you actually are.

I do not know how tall you are; I doubt that you starved as a child. My empirical evidence tells strongly against (1), by indicating that childhood malnutrition tends to reduce adult height. In contrast, (2) is logically trivial. How things actually are includes your actual height, whether or not it also includes your childhood starvation; how things would have been if you had starved as a child does not include your actual height. A similar contrast arises between metaphysical and epistemic modalities. As a claim of metaphysical possibility, (3) is obviously true, while as a claim of epistemic possibility (4) is obviously false:

(3) You could have been less tall now than you actually are.

(4) You may be less tall now than you actually are.

Such differences are no mere curiosities of natural language. They exemplify a very general difference between two kinds of sentential context. This paper expounds the general difference, and uses it to derive constraints on the semantics of indicative conditionals.

§1. We first construct a taxonomy of sentential contexts, according to their fineness or coarseness of grain.

By a 'sentential context' is meant an ordered pair <**α**, **β**>, **α** being any formula and **β** any atomic formula, both in a given object-language L: if it helps, think of **β** as occurring in **α**, although that is not formally required. If **C** is a sentential context <**α**, **p**> and **β** a formula of L, **C(β)** is the result of uniformly substituting **β** for **p** throughout **α**. For instance, if **C** is <**~p ⊃ q**, **p**> then **C(p & r)** is **~(p & r) ⊃ q**. This is just a formal version of the informal notion of a sentential context.

L is assumed to contain atomic formulas **p**, **q**, … and the sentential operators ~ (negation), ⊃ (the material conditional) and □ (informally interpreted as an operator for metaphysical necessity). Other connectives are, as usual, introduced as metalinguistic abbreviations; in particular, ≡ is the material biconditional and ◊ is the dual of □. L may contain many other constructions too, such as quantifiers and propositional attitude operators, but they do not matter for present purposes.

We work with a fairly standard model-theoretic conception of logical truth. Each formula of L is evaluated as true or false (never both) in each model for L; although the main points to follow do not depend on that assumption, dropping it would considerably complicate the exposition. To determine the truth-values of formulas in a model M, each formula is first evaluated as true or false in M relative to various *points* in a set P(M) associated with M. Each point in P(M) assigns values to all semantic parameters of L. It assigns values to all the free variables of L; it determines values for the parameters that fix the circumstance with respect to which sentences of L are to be evaluated; perhaps it also determines values for parameters that fix the context in which sentences of L are considered as uttered, and so on. Most importantly for our purposes, each point x determines an item w(x) which we can informally think of as the world of evaluation. The model M is associated with a set of such 'worlds' W(M) = {w(x): x ∈ P(M)}. If x ∈ P(M) and u ∈ W(M), then x[u] is the point just like x except that it assigns u as the value of the world of evaluation parameter; as is orthodox, we assume that x[u] is always in P(M). 'M,x ⊨ **α**' means that **α** is true in M relative to x ∈ P(M).

The semantic clauses for ~, ⊃ and □ are these:

M,x ⊨ ~**α** iff not M,x ⊨ **α**.

$M,x \models \boldsymbol{\alpha} \supset \boldsymbol{\beta}$ iff not $M,x \models \boldsymbol{\alpha}$ or $M,x \models \boldsymbol{\beta}$.

$M,x \models \Box\boldsymbol{\alpha}$ iff for all $u \in W(M)$, $M,x[u] \models \boldsymbol{\alpha}$.

The clauses for ~ and $\supset$ are standard; thus all theorems of classical propositional logic are true in any model relative to any point, and modus ponens preserves truth in any given model relative to any given point. The clause for □ treats necessity as truth in all worlds of the model, holding fixed all other semantic parameters. It does *not* embody a restriction to those worlds to which the world component w(x) of point x bears a designated accessibility relation R. Consequently, all theorems of the propositional modal system S5 (including $\Box\mathbf{p} \supset \mathbf{p}$, $\Box\mathbf{p} \supset \Box\Box\mathbf{p}$ and $\Diamond\mathbf{p} \supset \Box\Diamond\mathbf{p}$) are true in any model relative to any point. S5 is a plausible but not uncontroversial logic for metaphysical necessity, and dropping the accessibility relation simplifies the semantics; a further rationale will emerge below.

To get from truth in a model relative to a point to truth in a model *simpliciter*, we assume that each model M distinguishes some point $a(M) \in P(M)$ as its *actual point*; correspondingly, w(a(M)) is the *actual world* of M. A sentence $\boldsymbol{\alpha}$ is *true in a model M* ($M \models \boldsymbol{\alpha}$) iff $\boldsymbol{\alpha}$ is true in M relative to the actual point of M ($M,a(M) \models \boldsymbol{\alpha}$); $\boldsymbol{\alpha}$ is *logically true* ($\models \boldsymbol{\alpha}$) iff $\boldsymbol{\alpha}$ is true in every model. More generally, $\boldsymbol{\alpha}$ is a logical consequence of a set of sentences $\boldsymbol{\Gamma}$ ($\boldsymbol{\Gamma} \models \boldsymbol{\alpha}$) iff $\boldsymbol{\alpha}$ is true in every model in which every member of $\boldsymbol{\Gamma}$ is true. We assume that the semantics of L provides enough models to falsify any non-theorem of S5; more precisely, if $\boldsymbol{\alpha}$ is constituted of atomic formulas, truth-functors and □ then $\models \boldsymbol{\alpha}$ only if $\boldsymbol{\alpha}$ is a theorem of S5.

We now have enough semantic apparatus to define some useful kinds of context. A sentential context **C** is *extensional* iff for all sentences **α** and **β** this holds:

EXTENSIONAL $\models (\boldsymbol{\alpha} \equiv \boldsymbol{\beta}) \supset (\mathbf{C}(\boldsymbol{\alpha}) \equiv \mathbf{C}(\boldsymbol{\beta}))$

C is *intensional* iff it is not extensional. For example, negation creates an extensional context: $<{\sim}\mathbf{p}, \mathbf{p}>$ is extensional because $\models (\boldsymbol{\alpha} \equiv \boldsymbol{\beta}) \supset ({\sim}\boldsymbol{\alpha} \equiv {\sim}\boldsymbol{\beta})$ for all **α** and **β**. More generally, we can show that for any sentence **α** constructed solely out of atomic sentences, ~ and ⊃ , the context $<\boldsymbol{\alpha}, \mathbf{p}>$ is extensional (by induction on the complexity of **α**). By contrast, a modal operator such as □ creates an intensional context: $<\Box\mathbf{p}, \mathbf{p}>$ is intensional because if $\models (\mathbf{p} \equiv \top) \supset (\Box\mathbf{p} \equiv \Box\top)$, where $\top$ is a truth-functional tautology, then, since $\models \Box\top$, $\models \mathbf{p} \supset \Box\mathbf{p}$, contrary to fact (in S5). As we should expect, the result **C1*C2** of composing a pair of extensional contexts **C1** and **C2** is itself an extensional context, where **C1*C2** is defined by the equation **C1*C2(α)** = **C1(C2(α))**. We cannot quite identify an extensional context of a sentence with a formula that expresses a truth-function of that sentence; for example, $<\mathbf{p} \,\&\, \mathbf{q}, \mathbf{p}>$ is extensional (since $\models (\boldsymbol{\alpha} \equiv \boldsymbol{\beta}) \supset ((\boldsymbol{\alpha} \,\&\, \mathbf{q}) \equiv (\boldsymbol{\beta} \,\&\, \mathbf{q}))$) even though **p & q** is no truth-function of **p**. Informally, for an extensional context **C**, at any point in any model the truth-value (extension) of **C(α)** is a function of the truth-value (extension) of **α**, although which function it is may depend on the model or the point (see further Humberstone 1986).

We can define a new notion by replacing material equivalence in EXTENSIONALITY with strict equivalence. A sentential context **C** is *non-hyperintensional* iff for all sentences **α** and **β** this holds:

NON-HYPERINTENSIONAL $\models \Box(\boldsymbol{\alpha} \equiv \boldsymbol{\beta}) \supset \Box(\mathbf{C}(\boldsymbol{\alpha}) \equiv \mathbf{C}(\boldsymbol{\beta}))$

C is *hyperintensional* iff it is not non-hyperintensional. For example, □ is non-hyperintensional because $\Box(\boldsymbol{\alpha} \equiv \boldsymbol{\beta}) \supset \Box(\Box\boldsymbol{\alpha} \equiv \Box\boldsymbol{\beta})$ is a theorem of S5 (indeed of S4). More generally, we can show that for any sentence **α** constructed solely out of atomic sentences, ~, ⊃ and □, the context <**α**, **p**> is non-hyperintensional (by induction on the complexity of **α**). By contrast, belief operators create hyperintensional contexts, given that one necessary truth may be believed while another is not. As expected, composing non-hyperintensional contexts gives a non-hyperintensional context. Informally, if the intension of a sentence **α** at a point x in a model M is the function that takes each world $u \in W(M)$ to the truth-value of **α** in M relative to x[u], a non-hyperintensional context **C** is one for which, relative to the actual point in any given model, the intension of **C(α)** is a function of the intension of **α**, although which function it is may depend on the model or the point. Still more informally, non-hyperintensional contexts concern only what could have been.[1]

Had we restricted the quantification over worlds in the semantic clause for □ by an arbitrary accessibility relation, we could not have shown that □ creates a non-hyperintensional context, for $\Box(\mathbf{p} \equiv \top) \supset \Box(\Box\mathbf{p} \equiv \Box\top)$ is equivalent in any normal modal logic to $\Box\mathbf{p} \supset \Box\Box\mathbf{p}$, the characteristic axiom of S4, which is invalid when the accessibility relation is non-transitive. The problem is that the strict equivalence of two sentences at a point in a model no longer requires their intensions to coincide at that point, and 'hyperintensionality' becomes correspondingly less interesting.[2] It is more perspicuous to define 'hyperintensional' in terms of a □ unrestricted by an accessibility relation, which is not to deny that semantic clauses with an accessibility relation suit other readings of □, such as easy possibility. Those (notably Salmon 1989) who deny that metaphysical necessity satisfies the S4 axiom should think of the unrestricted □ as stronger than metaphysical necessity, and therefore as at least as strong as the infinite conjunction of metaphysical necessity, the metaphysical necessity of metaphysical necessity, and so on.

We can define a third notion by replacing material equivalence with logical equivalence. A sentential context **C** is *congruential* iff for all sentences **α** and **β** this holds:

CONGRUENTIAL $\qquad$ If $\models \boldsymbol{\alpha} \equiv \boldsymbol{\beta}$ then $\models \mathbf{C}(\boldsymbol{\alpha}) \equiv \mathbf{C}(\boldsymbol{\beta})$

Since logical truth is closed under modus ponens for ⊃ , all extensional contexts are congruential. As expected, composing congruential contexts gives a congruential context. Belief operators create non-congruential contexts, given that one logical truth may be believed while another is not. Informally, for a congruential context **C**, the logical consequences of **C(α)** are a function of the logical consequences of **α**.

Assume that the rule of necessitation is valid: whenever $\models \boldsymbol{\alpha}$, $\models \Box\boldsymbol{\alpha}$. Then all extensional contexts are non-hyperintensional, and all non-hyperintensional contexts are congruential. For if **C** is extensional then $\models \Box((\boldsymbol{\alpha} \equiv \boldsymbol{\beta}) \supset (\mathbf{C}(\boldsymbol{\alpha}) \equiv \mathbf{C}(\boldsymbol{\beta})))$ by necessitation; we can distribute □ over ⊃ because we have the K schema

$\models \Box(\boldsymbol{\gamma} \supset \boldsymbol{\delta}) \supset (\Box\boldsymbol{\gamma} \supset \Box\boldsymbol{\delta})$, so **C** is non-hyperintensional. Similarly, if **C** is non-hyperintensional and $\models \boldsymbol{\alpha} \equiv \boldsymbol{\beta}$, then $\models \Box(\boldsymbol{\alpha} \equiv \boldsymbol{\beta})$ by necessitation, so $\models \Box(\mathbf{C}(\boldsymbol{\alpha}) \equiv \mathbf{C}(\boldsymbol{\beta}))$; we can eliminate the □ because we have the T schema $\models \Box\boldsymbol{\gamma} \supset \boldsymbol{\gamma}$, so **C** is congruential. Both inclusions are strict if L is rich enough. As already noted, □ creates a non-extensional non-hyperintensional context. To construct a hyperintensional congruential context, define a logical truth operator **O** such that $\mathbf{O}\boldsymbol{\alpha}$ is true everywhere if $\models \boldsymbol{\alpha}$, and $\mathbf{O}\boldsymbol{\alpha}$ is false everywhere otherwise. The context <**Op**, **p**> is congruential, for if $\models \boldsymbol{\alpha} \equiv \boldsymbol{\beta}$ then $\models \boldsymbol{\alpha}$ iff $\models \boldsymbol{\beta}$, so either $\models \mathbf{O}\boldsymbol{\alpha}$ and $\models \mathbf{O}\boldsymbol{\beta}$ or $\models \sim\mathbf{O}\boldsymbol{\alpha}$ and $\models \sim\mathbf{O}\boldsymbol{\beta}$; hence $\models \mathbf{O}\boldsymbol{\alpha} \equiv \mathbf{O}\boldsymbol{\beta}$. Yet <**Op**, **p**> is hyperintensional, for otherwise $\models \Box(\mathbf{p} \equiv \top) \supset \Box(\mathbf{Op} \equiv \mathbf{O}\top)$, which in S5 implies that $\models \Box\mathbf{p} \supset (\mathbf{Op} \equiv \mathbf{O}\top)$; since $\models \top$ but not $\models \mathbf{p}$ (**p** being just an atomic sentence), $\models \mathbf{O}\top$ and $\models \sim\mathbf{Op}$, so $\models \sim\Box\mathbf{p}$, contrary to fact.[3] Under the preceding assumptions, therefore, the three features are ordered in decreasing logical strength: extensionality; non-hyperintensionality; congruentiality.

§2. The picture becomes more complex once we consider languages in which necessitation does not preserve logical truth. To construct an example, we need only add to L an 'actually' operator **A** of the usual sort: it rigidly returns the evaluation to the actual world. Recall that for any point x in a model M, x[w(a(M))] is the point just like x except that it assigns w(a(M)), the actual world of the model, to the world of evaluation parameter. Thus the semantic clause for **A** is this:

$$M,x \models \mathbf{A}\boldsymbol{\alpha} \text{ iff } M,x[w(a(M))] \models \boldsymbol{\alpha}.$$

We use x[w(a(M))] rather than a(M), the actual point of the model itself, because 'actually' should not interfere with other aspects of the evaluation, such as the assignment of values to variables.

How does necessitation fail in this setting? Since truth in a model M is truth in M relative to its actual point, M $\models$ **Ap** iff M,a(M) $\models$ **Ap** iff M,a(M)[w(a(M))] $\models$ **p**; but a(M)[w(a(M))] is just a(M), so M $\models$ **Ap** iff M,a(M) $\models$ **p** iff M $\models$ **p**. Thus $\models$ **Ap** $\equiv$ **p**. If necessitation preserves logical truth, $\models$ □**(Ap** $\equiv$ **p)**, so $\models$ □**Ap** $\equiv$ □**p**. But $\models$ □**Ap** $\equiv$ **Ap**, for M,x $\models$ □**Ap** iff for all u $\in$ W(M), M,x[u] $\models$ **Ap** iff for all u $\in$ W(M), M,x[u][w(a(M))] $\models$ **p**; but x[u][w(a(M))] is just x[w(a(M))], so M,x $\models$ □**Ap** iff x[w(a(M))] $\models$ **p** iff M,x $\models$ **Ap**. Intuitively, **A** is a rigidifier, so prefixing □ to it is redundant. Consequently, if necessitation preserves logical truth then $\models$ **Ap** $\equiv$ □**p**, so $\models$ **p** $\equiv$ □**p**, contrary to fact.[4] Therefore, necessitation does not preserve logical truth in L. Informally: logical truth depends only on what happens in the actual world of each model, where **Ap** and **p** coincide in truth-value; they may differ in truth-value in a counterfactual world, which suffices to falsify □**(Ap** $\equiv$ **p)** in the actual world.

In the terminology of Crossley and Humberstone (1977), we defined logical truth as *real world validity*, truth in each model relative to its actual point, by contrast with *general validity*, truth relative to each point in each model.[5] Although necessitation does not preserve real world validity, it does preserve general validity.[6] In particular, **Ap** $\equiv$ **p** is not generally valid. For some purposes, general validity is a more useful notion than real world validity, as indicated below. For present purposes, however, real world validity constitutes a better standard of logical truth. Transitions between 'It's raining' and 'It's actually raining' are broadly logical in character. Moreover, real world validity is truth in all models, and so fits the model-theoretic account of logical truth, whereas general validity does not (Zalta 1988). We therefore continue to work with real world validity.

The argument above that necessitation does not preserve logical truth also shows that $\langle \Box\mathbf{p}, \mathbf{p}\rangle$ is a counterexample to the claim that every non-hyperintensional context is congruential. We have already noted that it is non-hyperintensional. It is non-congruential because $\models \mathbf{Ap} \equiv \mathbf{p}$ but not $\models \Box\mathbf{Ap} \equiv \Box\mathbf{p}$.

To construct a counterexample to the claim that no extensional context is hyperintensional, recall the logical truth operator **O** from the end of §1. Let **C** be the context $\langle(\mathbf{Aq} \equiv \mathbf{q}) \vee \mathbf{Op}, \mathbf{p}\rangle$. **C** is trivially extensional because $\models \mathbf{C}(\boldsymbol{\alpha})$ for all $\boldsymbol{\alpha}$. **C** is hyperintensional: otherwise $\models \Box(\mathbf{p} \equiv \top) \supset \Box(((\mathbf{Aq} \equiv \mathbf{q}) \vee \mathbf{Op}) \equiv ((\mathbf{Aq} \equiv \mathbf{q}) \vee \mathbf{O}\top))$; since $\models \Box{\sim}\mathbf{Op}$ and $\models \Box\mathbf{O}\top$, by normal modal logic $\models \Box\mathbf{p} \supset \Box(\mathbf{Aq} \equiv \mathbf{q})$, so $\models \Box\mathbf{p} \supset (\Box\mathbf{Aq} \supset \Box\mathbf{q})$; but $\models \mathbf{q} \supset \Box\mathbf{Aq}$, so $\models \Box\mathbf{p} \supset (\mathbf{q} \supset \Box\mathbf{q})$, contrary to fact.

In the presence of an 'actually' operator, the interrelations between the three kinds of context are therefore more complex. Of course, all extensional contexts are still congruential, but the logical features of extensionality and congruentiality are less tidily related to the modal feature of non-hyperintensionality.

Perhaps more surprisingly, the presence of an 'actually' operator forces new positive interrelations between the three features. For example, in the language of propositional modal logic without additional operators such as **A**, modal operators are congruential and non-hyperintensional without being extensional. Even when we have added **A**, we might be tempted to think of congruentiality and non-hyperintensionality as approximately equivalent categorisations, even though each applies to a few curiosities to which the other does not. On this picture, the two categories agree on excluding very fine-grained contexts and including moderately coarse-grained ones, while disagreeing on a marginal range of intermediate contexts. But that picture is wrong. On a more accurate view, the two categories exclude each other, but for an overlap on very coarse-grained contexts. For, if L contains **A**, *every congruential non-hyperintensional context is extensional*.

Here is the proof. Let **C** be a congruential non-hyperintensional context and **α** and **β** sentences of L. As usual:

(i) $\models \mathbf{A\alpha} \equiv \boldsymbol{\alpha}$

(ii) $\models \mathbf{A\beta} \equiv \boldsymbol{\beta}$

From (i) and (ii):

(iii) $\models (\boldsymbol{\alpha} \equiv \boldsymbol{\beta}) \supset (\mathbf{A\alpha} \equiv \mathbf{A\beta})$

The rigidifying effect of **A** makes □ redundant as applied to the consequent of (iii), as one can easily check by a semantic argument:

(iv) $\models (\mathbf{A\alpha} \equiv \mathbf{A\beta}) \supset \Box(\mathbf{A\alpha} \equiv \mathbf{A\beta})$

An instance of the non-hyperintensionality of **C** is:

(v) $\models \Box(\mathbf{A\alpha} \equiv \mathbf{A\beta}) \supset \Box(\mathbf{C(A\alpha)} \equiv \mathbf{C(A\beta)})$

The semantics for □ validates the T schema, and in particular:

(vi) $\models \Box(\mathbf{C(A\alpha)} \equiv \mathbf{C(A\beta)}) \supset (\mathbf{C(A\alpha)} \equiv \mathbf{C(A\beta)})$

Chaining (iii)-(vi) together we have:

(vii) $\models (\boldsymbol{\alpha} \equiv \boldsymbol{\beta}) \supset (\mathbf{C(A\alpha)} \equiv \mathbf{C(A\beta)})$

From (i) and (ii) respectively the congruentiality of **C** yields:

(viii) $\models \mathbf{C(A\alpha)} \equiv \mathbf{C(\alpha)}$

(ix) $\models \mathbf{C(A\beta) \equiv C(\beta)}$

Together, (vii)-(ix) yield:

(x) $\models \mathbf{(\alpha \equiv \beta) \supset (C(\alpha) \equiv C(\beta))}$

That is, **C** is extensional.

How should we interpret this result? *Very* roughly: extensional contexts are insensitive to differences between sentences with the same truth-value; non-hyperintensional contexts are insensitive to differences between sentences with the same intension; congruential contexts are insensitive to differences between sentences with the same real-world logical consequences. The proof works because any two sentences **α** and **β** with the same truth-value are linked by a chain in which successive members have either the same intension or the same real-world logical consequences: **α** has the same real-world logical consequences as **Aα**, which has the same intension as **Aβ**, which has the same real-world logical consequences as **β**. A context that is both congruential and non-hyperintensional must be insensitive to differences between the sentences at the opposite ends of such a chain, and therefore between any two sentences with the same truth-value; in other words, it is extensional. That explanation is rough, because 'insensitive' does not mean exactly the same in the three cases (hence the need for a rigorous proof); nevertheless, it gives the feel of what is going on.

§3. Some examples may help. Consider 'The probability that **p** is d' as a context of **p**. Standard axioms for probability guarantee that logical equivalents are equiprobable, but what is the appropriate notion of logical equivalence? A probabilistic context is non-extensional: sameness of truth-value does not imply sameness of probability. Therefore it is either hyperintensional or non-congruential. Which?

Suppose that the probability in question is objective chance. The resultant context is hardly hyperintensional; if two outcomes could have differed in their chances, one of them could have occurred without the other. Therefore, since talk of chances creates non-extensional contexts (otherwise everything has a chance of 0 or 1), it must create non-congruential ones, by the result above. For example, the real-world logical equivalence of (5) and (6) does not imply the real-world logical equivalence of (7) and (8):

(5) The next toss will be heads.

(6) Actually the next toss will be heads.

(7) The chance that the next toss will be heads is 0.5.

(8) The chance that actually the next toss will be heads is 0.5.

Whatever world u is actual, it is non-contingent whether (5) holds in u. Therefore the chance that (5) holds in u is either 0 or 1. Thus the chance that (5) holds in the actual world is not 0.5; (8) is false. Nevertheless, given indeterminism, (7) may very well be true. Unlike general logical equivalence, real-world logical equivalence does not imply equality of chance.

Suppose, by contrast, that the probability in question is (coherent) subjective credence or evidential probability. The resultant context is surely congruential. The rationale for assigning the same coherent subjective credence or evidential probability to ordinary logical equivalents extends straightforwardly to pairs such as (5) and (6); the coincidence in truth-value is easily accessible to reflection. Therefore, since talk of credences or evidential probabilities creates non-extensional contexts (otherwise everything has such a probability of 0 or 1), it must create hyperintensional ones by the result above. Suppose that the next toss will in fact be heads. Then (6) has the same intension as (9):

(9) The next toss will be heads or not.

Nevertheless, even for a perfectly coherent subject, (10) may be true while (11) is false:

(10) My credence that actually the next toss will be heads is 0.5.

(11) My credence that the next toss will be heads or not is 0.5.

For the credence for (6) should be the same as for (5), while the credence for (9) should be 1. Coincidence in intension is not always accessible to reflection. Real-world logical equivalence is a better candidate than metaphysically necessary equivalence for implying sameness of coherent credence or evidential probability.

Metaphysical and epistemic possibility pattern the same way as chance and credence respectively. Both create non-extensional contexts, for some falsehoods are both metaphysically and epistemically possible while others are neither.

Talk of metaphysical possibilities creates non-hyperintensional, non-congruential contexts. Non-hyperintensionality: If two sentences have the same intension, so have the results of prefixing them with 'It could have been that' (on the metaphysical reading).[7] For example, if you are in fact six feet tall, (12) has the same intension as (13), and (14) (= (3)) correspondingly the same intension as (15):

(12) You are less tall now than you actually are.

(13) You are less than six feet tall now.

(14) You could have been less tall now than you actually are

(15) You could have been less than six feet tall now.

Non-congruentiality: The real-world logical equivalence of (16) and (17) does not imply even the material equivalence of (18) and (19):

(16) The last toss was tails.

(17) The last toss was actually tails.

(18) It could have been that the last toss was tails.

(19) It could have been that the last toss was actually tails.

For (18) is true and (19) false if (16) is false. Similarly, despite the real-world logical equivalence of (12) and (20), (14) is true and (21) false:

(20) You are self-distinct.

(21) You could have been self-distinct.

By contrast, talk of epistemic possibilities creates a hyperintensional, congruential context. Hyperintensionality: If (16) is false, then (17) has the same intension as (20), but (22) is true (if nobody has checked) while (23) is false, so they have different intensions:

(22) It may be that the last toss was actually tails.

(23) It may be that you are self-distinct.

Similarly, despite the sameness in intension of (12) and (13), (24) can be true while (25) (= (4)) is false:

(24) It may be that you are less than six feet tall now.

(25) It may be that you are less tall now than you actually are.

Congruentiality: If **α** and **β** are logically equivalent, then **α** is consistent with what is known iff **β** is consistent with what is known. For instance, the real-world logical equivalence of (12) and (20) implies the real-world logical equivalence of (25) and (23).

Such examples may suggest that the distinction between non-hyperintentional non-congruential contexts and hyperintensional congruential contexts is fundamentally a distinction between the metaphysical and the epistemic. That thought might encourage attempts to explain the distinction within some sort of generalized two-dimensional semantics, perhaps in the manner of Chalmers (2004) or Jackson (2004). Such attempts would not be well motivated by the considerations of this paper, for at least two reasons.

First, the definitions in §1 and the argument of §2 nowhere used two-dimensional semantics. They make perfect sense even if each point in a model consists of a single world of evaluation (the 'metaphysical' dimension) and nothing else (in particular, no 'epistemic' dimension). Even if each point fixes other parameters too, such as times, they may all pertain to the circumstance of evaluation rather than the context of utterance in the sense of Kaplan (1989), and therefore still not pertain to the epistemic side of generalized two-dimensional semantics. Although the 'actually' operator was crucial, its semantic clause relied simply on the stipulation that each model has a distinguished actual point. The rationale for that stipulation was in turn simply to obtain a sensible notion of truth in a model. If semantic clauses for epistemic modal operators or the like involve a new 'epistemic' dimension of models, that reflects on those operators, not on the general notions of extensionality, non-hyperintensionality and congruentiality.[8]

Second, the category of hyperintensional congruential contexts is not limited to epistemic contexts. For instance, the logical truth operator **O** in §1 created a hyperintensional congruential context without being epistemic in any natural sense: logical truth was understood simply as truth in all models. Although 'All red things are coloured' may express a truth knowable *a priori*, the sentence is not logically true; $\forall x(Rx \supset Cx)$ is false in some models. Conversely, we have no general guarantee that every sentence true in all models expresses a truth knowable *a priori*. Of course, the completeness theorem for first-order logic tells us that a sentence of first-order logic is true in all models only if it is provable from standard axiomatizations and therefore expresses a truth knowable *a priori* (once interpreted).[9] But second-order logic is incomplete; we cannot assume that the second-order sentence that expresses a version of the Continuum Hypothesis is true in all models only if it is knowable *a priori*.

Do such examples undermine the idea that talk of epistemic probability or possibility creates congruential contexts? Let 'P' abbreviate an epistemically inaccessible logical truth while 'Q' abbreviates an epistemically accessible logical truth; thus 'P iff Q' also abbreviates an epistemically inaccessible logical truth. The epistemic probability that Q is 1; isn't the epistemic probability that P less than 1? It is not epistemically possible that not Q; isn't it epistemically possible that not P? The same queries might even be raised about a pair of epistemically accessible logical truths, one that will never in fact be accessed because it is too complex, the other simple and already accessed.

In the case of probability, its insensitivity to such differences is built into the technical notion, otherwise it would not satisfy the standard axioms of probability theory, on which most of its usefulness depends. The case of possibility is not very different. If a complex truth-functional tautology and a simple one can differ in epistemic possibility, the latter notion does not satisfy basic principles of weak normal propositional modal logic (specifically, the weakest normal system K), on which much of its usefulness depends. Of course, we might read the English 'It may be that not P' as expressing only that it is not known that P, rather than the stronger claim that it is consistent with the totality of what is known that P, for something tautologically equivalent to something known need not itself be known: but a mere denial of knowledge is not usefully regarded as a modality. Thus we legitimately impose closure under at least some forms of logical equivalence by brute force on both epistemic probability and epistemic possibility in order to achieve a perspicuous mathematical and logical structure. It may well be best to extend the imposition to closure under all forms of logical equivalence for the same reason. We can still register finer-grained epistemic differences elsewhere in our conceptual framework. In any case, the forms of logical equivalence of most interest for this paper, in particular that between 'P' and 'Actually P', are epistemically almost trivial.

The distinction between congruential and non-congruential contexts is more accurately conceived as purely logical (because it concerns closure under real-world logical equivalence) than as epistemic, and the distinction between non-hyperintensional and hyperintensional contexts as metaphysical (because it concerns closure under metaphysically necessary equivalence).[10] The term 'real-world' should not mislead one into regarding the former distinction as metaphysical too, for real-world validity was defined in terms of the actual *point* of a model; the actual *world* of the model may be only one component amongst many of the actual point, relevant only for the evaluation of sentences with operators that engage with metaphysical modality.

Epistemic probability and possibility operators have exhibited differences in purely logical behaviour from corresponding operators for metaphysical probability and possibility. The former are closed under real-world logical equivalence, the latter not (both kinds are closed under general logical equivalence). Less remarkably, the metaphysical operators are closed under (metaphysically) necessary equivalence, whereas the epistemic operators are not. Does the contrast apply to conditionals too?

§4. Let → be a conditional operator. We must consider two contexts, that of the antecedent, <**p** → **q**, **p**>, and that of the consequent, <**p** → **q**, **q**>. Say that → is *extensional in its antecedent* (respectively, *consequent*) iff <**p** → **q**, **p**> (respectively, <**p** → **q**, **q**>) is an extensional context; likewise for non-hyperintensionality and congruentiality. Usually, a conditional operator is extensional (non-hyperintensional, congruential) in its antecedent iff it is extensional (non-hyperintensional, congruential) in its consequent. The contrast between (1) and (2) (with 'actually' in the consequent) is similar to the following contrast (with 'actually' in the antecedent):

(26) If you had eaten much less as a child than you actually did, you would have been six feet tall now.

(27) If you ate much less as a child than you actually did, you are six feet tall now.

We might have strong empirical evidence against (26), concerning the effects of childhood malnutrition. Such evidence is irrelevant to (27), since its antecedent is logically false; perhaps (27) is vacuously true. An example with an epistemically possible antecedent is:

(28) If you had eaten much less as a child than I actually did, you would have been six feet tall now.

(29) If you ate much less as a child than I actually did, you are six feet tall now.

We might have the same sort of empirical evidence against (28) that we have against (26), while knowing that (29) is true because you have been measured as six feet tall now.

By the result of §2, any conditional that is both non-hyperintensional and congruential in its antecedent is extensional in its antecedent, and any conditional that is both non-hyperintensional and congruential in its consequent is extensional in its consequent. Indeed, given weak assumptions, we can establish something stronger: any conditional that is both non-hyperintensional and congruential in *either* its antecedent *or* its consequent is extensional in *both* its antecedent *and* its consequent. For we can show, on the assumptions, that a conditional is extensional in its antecedent iff it is extensional in its consequent. Moreover, on the same assumptions, we can show that a conditional is extensional in its antecedent or consequent iff it is logically equivalent to the material (truth-functional) conditional (see Appendix). Thus any conditional that is both non-hyperintensional and congruential in its antecedent or consequent is logically equivalent to $\supset$.

Of course, logical equivalence is weaker than synonymy. For example, if $\boldsymbol{\alpha \rightarrow \beta}$ is $\mathbf{A}(\boldsymbol{\alpha} \supset \boldsymbol{\beta})$, then $\models (\boldsymbol{\alpha} \rightarrow \boldsymbol{\beta}) \equiv (\boldsymbol{\alpha} \supset \boldsymbol{\beta})$ (this $\rightarrow$ is both non-hyperintensional and congruential). Nevertheless, $\rightarrow$ and $\supset$ differ in meaning because they embed differently. In particular, $\models (\mathbf{p} \rightarrow \mathbf{q}) \supset \Box(\mathbf{p} \rightarrow \mathbf{q})$ because $\models \mathbf{A}\boldsymbol{\gamma} \supset \Box\mathbf{A}\boldsymbol{\gamma}$ more generally, but of course not $\models (\mathbf{p} \supset \mathbf{q}) \supset \Box(\mathbf{p} \supset \mathbf{q})$. However, the most basic disputes about conditionals concern the truth-values of unembedded occurrences. To say that that $\boldsymbol{\alpha \rightarrow \beta}$ is always logically equivalent to $\boldsymbol{\alpha} \supset \boldsymbol{\beta}$ is to answer the most urgent questions about $\rightarrow$.

The strict conditional ($\Box(\boldsymbol{\alpha} \supset \boldsymbol{\beta})$) is not logically equivalent to the material conditional. It is non-hyperintensional in both its antecedent and its consequent because $\Box$ creates a non-hyperintensional context (in S4 and S5). Therefore the strict conditional is non-congruential in both its antecedent and its consequent. Specifically, although $\models \mathbf{Ap} \equiv \mathbf{p}$, neither $\models \Box(\mathbf{Ap} \supset \mathbf{p}) \equiv \Box(\mathbf{p} \supset \mathbf{p})$ nor $\models \Box(\mathbf{p} \supset \mathbf{Ap}) \equiv \Box(\mathbf{p} \supset \mathbf{p})$, for otherwise, since $\models \Box(\mathbf{p} \supset \mathbf{p})$, either $\models \Box(\mathbf{Ap} \supset \mathbf{p})$ or $\models \Box(\mathbf{p} \supset \mathbf{Ap})$; models in which $\mathbf{p}$ is actually contingently true refute the first alternative and models in which $\mathbf{p}$ is actually contingently false refute the second.

An example of the reverse case is the binary logical consequence operator defined by $\mathbf{O}(\boldsymbol{\alpha} \supset \boldsymbol{\beta})$, where $\mathbf{O}$ is the logical truth operator from the end of §1. It is not logically equivalent to the material conditional. It is congruential in both its antecedent and its consequent because **O** creates a congruential context. Therefore the logical consequence operator is hyperintensional in both its antecedent and its consequent. Specifically, if $\models \Box(\mathbf{p} \equiv \mathbf{q}) \supset \Box(\mathbf{O}(\mathbf{p} \supset \mathbf{p}) \equiv \mathbf{O}(\mathbf{q} \supset \mathbf{p}))$ then $\models \Box(\mathbf{p} \equiv \mathbf{q}) \supset (\mathbf{O}(\mathbf{p} \supset \mathbf{p}) \equiv \mathbf{O}(\mathbf{q} \supset \mathbf{p}))$ (by the T axiom); but $\models \mathbf{O}(\mathbf{p} \supset \mathbf{p})$ and not $\models \sim\mathbf{O}(\mathbf{q} \supset \mathbf{p})$, so $\models \sim\Box(\mathbf{p} \equiv \mathbf{q})$, contrary to fact. Hyperintensionality in the consequent is similar.

In these respects, the subjunctive conditional patterns like the strict conditional. Its non-extensionality is clear, but let us check it anyway. Suppose that John is in China. All three of these conditionals have false antecedents and false consequents:

(30) If John were in London, he would be in England.

(31) If John were in Paris, he would be in England.

(32) If John were in London, he would be in France.

The subjunctive conditional is non-extensional in its antecedent because (30) is true and (31) false, even though (30) and (31) have the same consequent and antecedents with the same truth-value. It is non-extensional in its consequent because (30) is true and (32) false, even though (30) and (32) have the same antecedents and consequents with the same truth-value. Thus the subjunctive conditional is either hyperintensional or non-congruential in its antecedent and in its consequent. Intuitively, it is non-hyperintensional in both.[11] For what possibilities a subjunctive conditional obtains in seems to supervene on what possibilities its antecedent and consequent obtain in; thus replacing the antecedent or consequent by another that obtains in exactly the same possibilities should make no difference to what possibilities the conditional itself obtains in.[12] For example, any possible worlds semantics of the usual Stalnaker-Lewis type for subjunctive conditionals excludes hyperintensionality, provided that □ ranges over all the worlds in a model. Whether the consequent is true in the most similar worlds to the actual one in which the antecedent is fixed by what worlds the antecedent and consequent are true in.

An intriguing objection is that hyperintensionality infects the subjunctive conditional because logically equivalent antecedents sometimes trigger evaluation by different similarity relations. For example, when we counterfactually suppose that Bizet had been the same nationality as Verdi, we tend to keep Verdi's nationality fixed and imagine both of them as Italian, whereas when we counterfactually suppose that Verdi had been the same nationality as Bizet, we tend to keep Bizet's nationality fixed and imagine both of them as French. Even though co-nationality is symmetric, we might assent to (33) and dissent from (34):

(33) If Bizet had been the same nationality as Verdi, Bizet would have been Italian.

(34) If Verdi had been the same nationality as Bizet, Bizet would have been Italian.

However, one can explain such examples without invoking hyperintensionality. We tend to read (33) as (35) and (34) as (36):

(35) If Bizet had been the same nationality as Verdi actually was, Bizet would have

been Italian.

(36) If Verdi had been the same nationality as Bizet actually was, Bizet would have been Italian.

Since the antecedents of (35) and (36) are not strictly equivalent, a difference in truth-value between (35) and (36) does not imply hyperintensionality. This explanation of the difference between (33) and (34) is confirmed by the observation that the difference disappears once it is explicit that what counts is the nationality of the second-mentioned composer in the counterfactual circumstances, not the actual ones:

(37) If Bizet had been the same nationality as Verdi had been, Bizet would have been Italian.

(38) If Verdi had been the same nationality as Bizet had been, Bizet would have been Italian.

Thus the objection fails to identify hyperintensionality in the subjunctive conditional.

Given that the subjunctive conditional is non-hyperintensional in both antecedent and consequent, it is also non-congruential in both. Examples are easy to find. If the subjunctive conditional were congruential in its antecedent, we could drop ‘actually’ and the associated rigidification from the antecedent of (39) (= (26)):

(39) If you had eaten much less as a child than you actually did, you would have been six feet tall now.

But that would transform the contingently false (39) into a vacuously true conditional with an impossible antecedent on the unrigidified reading. Similarly, if the subjunctive conditional were congruential in its consequent, we could drop 'actually' and the associated rigidification from the consequent of (40) (= (1))

(40) If you had starved as a child, you would have been as tall now as you actually are.

But that would transform the contingently false (40) into a necessarily true conditional with a trivial consequent.

We can also argue more formally, making two weak assumptions about the subjunctive conditional, symbolized as $\Box\!\!\rightarrow$. First, it is reflexive:

REFLEXIVITY $\models \boldsymbol{\alpha \Box\!\!\rightarrow \alpha}$

If something were so, it would indeed be so. Second, the subjunctive conditional transmits metaphysical possibility:

POSSIBILITY $\models \boldsymbol{(\alpha \Box\!\!\rightarrow \beta) \supset (\Diamond\alpha \supset \Diamond\beta)}$

If something could be so, whatever would be so if it were so could also be so.[13] Suppose that **p** is contingently false (**~p &** ◊**p**). By the semantics of **A**, **Ap** is non-contingently false (~◊**Ap**). Thus, by POSSIBILITY, **p** □→ **Ap** is false. By REFLEXIVITY, **Ap** □→ **Ap** is true, so □→ is non-congruential in its antecedent, and **p** □→ **p** is true, so □→ is non-congruential in its consequent.

Such interactions between the subjunctive conditional and 'actually' should not surprise us once we consider the functions that subjunctive conditionals perform in our lives. For example, they help us learn from our mistakes. The subjunctive (41) gives one reason to eat less in future; the vacuous indicative (42) does not:

(41) If I had eaten much less than I actually did, I would not have a stomach ache now.

(42) If I ate much less than I actually did, I do not have a stomach ache now.

Subjunctive conditionals also help us in abductive reasoning. Given (43), we may provisionally judge that John has the disease; by contrast, the indicative (44) is an uninformative triviality:

(43) If John had the disease, he would appear just as he actually does.

(44) If John has the disease, he appears just as he actually does.

Such uses of subjunctive conditionals involve comparisons within the scope of the conditional between how things would be and how they actually are. For those comparisons, the rigidifying effect of ‘actually’ is very convenient. We might try to avoid it by quantification, as in (45), but why deprive ourselves of the option of referring back to actuality whenever we feel like it from within the scope of a counterfactual supposition that we are in the process of exploring?

(45) For some q, is the quantity that I ate and if I had eaten much less than q, I would not have a stomach ache now.

We need **p □→ Ap** to fail when **p** is contingently false for such comparisons of the counterfactual with the actual to make non-trivial sense. Equally, we need REFLEXIVITY to hold for making a counterfactual supposition to have its point. Thus the non-congruentiality of the subjunctive conditional in both its antecedent and its consequent is no accidental by-product of the interaction between two otherwise unrelated operators. Rather, it goes to the heart of thought about how things could have been otherwise.

If the contrast between indicative and subjunctive conditionals follows those in §3, the indicative conditional should be non-extensional and congruential, and therefore hyperintensional, in both antecedent and consequent. It certainly seems to be congruential.[14] If ‘P’ and ‘P*’ are (real-world) logically equivalent, surely ‘If P, Q’ is logically equivalent to ‘If P*, Q’ and ‘If Q, P’ is logically equivalent to ‘If Q, P*’. For example, (46) is equivalent to (47) and to (48):

(46) If it rained, the match was cancelled.

(47) If it actually rained, the match was cancelled.

(48) If it rained, the match was actually cancelled.

Such closure under real-world equivalence explains the triviality of (2) and the vacuity of (27). We therefore assume that the indicative conditional is congruential in both antecedent and consequent.[15]

Given that the indicative conditional is non-extensional in both antecedent and consequent, we can conclude that it is hyperintensional in both antecedent and consequent. But is the indicative conditional really non-extensional?

A crude argument for non-extensionality mechanically follows the argument for the non-extensionality of the subjunctive conditional, based on (30)-(32). The indicative analogues of (30)-(32) are:

(49) If John is in London, he is in England.

(50) If John is in Paris, he is in England.

(51) If John is in London, he is in France.

The argument insists that (49) is true but (50) and (51) false, wherever John in fact is. Suppose that John is in China, although we do not know that. Thus (49)-(51) all have false antecedents and false consequents. Then, according to the argument, the indicative conditional is non-extensional in its antecedent because (49) is true and (50) false, even though (49) and (50) have the same consequent and antecedents with the same truth-value. It is non-extensional in its consequent because (49) is true and (50) false, even though (49) and (50) have the same antecedents and consequents with the same truth-value.

However, the arguments for non-extensionality are not perfectly parallel for the two conditionals. For the subjunctive conditional, the argument for non-extensionality works even if we are certain in evaluating (30)-(32) that John is in China; that does not prevent us from making sensible judgements as to what would be the case were he elsewhere. By contrast, for the indicative conditional, if we are certain in trying to evaluate (49)-(51) that John is in China, it is hard to know what to make of them. What is clear is that when we have no idea of John's whereabouts, it is fine to assert (49) but far from fine to assert (50) or (51). But that does not exclude the hypothesis that (49) differs from (50) and (51) in epistemic status rather than truth-value. If indicative conditionals have truth-values that are independent of their epistemic status, it remains unclear what is wrong with arguing for (50) and (51) as follows. Since John is in China, he is not in Paris, so either he is not in Paris or he is in England, so if he is in Paris he is in England. Moreover, he is not in London, so either he is not in London or he is in France, so if he is in London he is in France.

On a rival hypothesis, the truth-values of indicative conditionals depend on their epistemic status.[16] For instance, (54) is false while (53) is true, even though (52) is true (where definite descriptions take narrow scope throughout):

(52) The masked man is my father.

(53) If the masked man is not my father, the masked man is not my father.

(54) If the masked man is not my father, the masked man is not the masked man.

Consequently, the argument from the identity (52) and (53) (or from (53) alone) to (54) is invalid, because it fails to preserve truth. But the argument from (52) and (53) to (54) looks valid, unlike the arguments from (52) and (55) to (56) and from (52) and (57) to (58):

(55) The evidential probability that the masked man is not my father is non-zero.

(56) The evidential probability that the masked man is not the masked man is non-zero.

(57) It is epistemically possible that the masked man is not my father.

(58) It is epistemically possible that the masked man is not the masked man.

Of course, once (52) is known, the differences in truth-value may disappear between (55) and (56) and between (57) and (58). But the validity of an argument is a matter of preserving truth from premises to conclusion, not of preserving known truth. It is far more plausible that (52) and (55) can be true while (56) is false, or that (52) and (57) can be true while (58) is false, than that (52) and (53) can be true while (54) is false. That the arguments from (52) and (55) to (56) and from (54) and (57) to (58) preserve certainty does not trick us into falsely thinking that they preserve truth. Thus one cannot adequately explain the appearance of validity in the argument from (52) and (53) to (54) merely by claiming that, because it preserves certainty, we are tricked into falsely thinking that it preserves truth. Surely indicative 'if' does not create an opaque context.

Once the validity of the argument from (52) and (53) to (54) is conceded, the hyperintensionality of the indicative conditional is hard to maintain. The argument from (59) and (60) (or from (59) alone) to (61) seems valid:

(59) John is six feet tall.

(60) If John is six feet tall, John is six feet tall.

(61) If John is his actual height, John is six feet tall.

It preserves truth from premises to conclusion, even though we do not know whether premise (59) is true. Our ignorance is no ground to exclude the truth of (61). Although we may answer ‘No, he may be only six feet tall’ to someone equally ignorant who asserts (61), we might thereby merely be denying his warrant to make the assertion rather than denying the truth of its content. After all, (62) is quite plausible even when we are in no position to assert any verifying instance:

(62) For some height *h*, if John is his actual height, John is *h* tall.

Thus the assignment of epistemic truth-conditions to indicative ‘if’ looks ill-grounded.

More generally, the forms of inference from (63) and (64) to (65) and from (63) and (66) to (67) seem valid (where ‘iff’ is read as the material biconditional):

(63) Necessarily, P iff P*.

(64) If P, Q.

(65) If P*, Q.

(66) If Q, P.

(67) If Q, P*.

But then indicative ‘if’ is non-hyperintensional. Since it is also congruential, it is extensional. Therefore, by the argument of the Appendix, it is logically equivalent to the truth-functional conditional.

Throughout, we have worked within a framework of truth-conditional semantics. That may seem unfair to those, such as Edgington (1986), who deny that indicative conditionals have truth-conditions. The critical definitions and arguments of this paper were stated in terms of logical truth ($\models$). However, theorists who deny that indicative conditionals have truth-conditions propose non-truth-conditional notions of validity for them. The preceding definitions and arguments may be reinterpretable in terms of such a reading of ‘$\models$’. Of course, the application of the definitions of extensionality, congruentiality and non-hyperintensionality to indicative ‘if’ does require the intelligibility of embedding it in the scope of the material biconditional and other operators, but to deny that intelligibility might be a further cost of the position. If the argument for the extensionality of indicative ‘if’ is accepted, understood in terms of EXTENSIONALITY even on a non-truth-conditional reading of ‘$\models$’, it seems strained to continue to insist that indicative ‘if’ is not truth-functional.

Caution is in order, for we cannot hope to discuss here all the extant challenges to a truth-functional reading of indicative ‘if’. We leave those who reject the truth-functional reading with the challenge of making the required hyperintensionality of ‘if’ plausible.

Probably *no* semantic theory of the indicative conditional can smoothly explain all the data. Speakers often react to sample sentences under the influence of their inchoate theorizing about the indicative conditional, partly because its embedded occurrences tend to produce unnatural sentences. We may need to recall a moral from the epistemic theory of vagueness: understanding an expression as a native speaker does not always put one in a position to recognize a correct perspicuous non-homophonic description of its reference as correct. Our use of the indicative conditional may determine an assignment of reference to it in a way that is not fully cognitively accessible to us: some of our dispositions in using it may turn out to involve errors that can be recognized as such from the third-person perspective of the theorist but not from the first-person perspective of the speaker. That is especially likely if the practice embodies internal tensions.

§5. One might have expected a category of *mildly non-extensional* contexts that would be congruential but neither extensional nor hyperintensional. As we have seen, in sufficiently rich languages that category is empty. This constraint on the possibilities for sentential operators may help us classify them more accurately.[17]

APPENDIX

Assume these two principles, for all formulas **α** and **β** of L:

(I) $\models (\boldsymbol{\alpha} \ \& \ \boldsymbol{\beta}) \rightarrow \boldsymbol{\beta}$

(II) $\models \boldsymbol{\alpha} \rightarrow (\boldsymbol{\alpha} \vee \boldsymbol{\beta})$

(III) $\models (\boldsymbol{\alpha} \rightarrow \bot) \supset {\sim}\boldsymbol{\alpha}$

(IV) $\models (\top \rightarrow \boldsymbol{\beta}) \supset \boldsymbol{\beta}$

(I) and (II) are special cases of a more general principle: a conditional whose consequent follows by truth-functional logic from its antecedent is true. Similarly, (III) and (IV) are tantamount to special cases of modus ponens for →, since modus ponens requires a true conditional with a true antecedent to have a true consequent. But even someone who rejected one or both of those more general principles might accept all of (I)-(IV).

Suppose that → is extensional in its antecedent. To show that → is extensional in its consequent, it suffices to show that **α** → **β** is logically equivalent to the material conditional **α** ⊃ **β** (which is of course extensional in its consequent). We argue thus:

(i) $\models (\boldsymbol{\alpha} \supset \boldsymbol{\beta}) \supset (\boldsymbol{\alpha} \equiv (\boldsymbol{\alpha} \ \& \ \boldsymbol{\beta}))$ truth-functional logic

(ii) $\models (\boldsymbol{\alpha} \equiv (\boldsymbol{\alpha} \ \& \ \boldsymbol{\beta})) \supset ((\boldsymbol{\alpha} \rightarrow \boldsymbol{\beta}) \equiv ((\boldsymbol{\alpha} \ \& \ \boldsymbol{\beta}) \rightarrow \boldsymbol{\beta})$ extensionality in antecedent

(iii) $\models (\boldsymbol{\alpha} \equiv (\boldsymbol{\alpha} \ \& \ \boldsymbol{\beta})) \supset (\boldsymbol{\alpha} \rightarrow \boldsymbol{\beta})$ (ii), (I)

(iv)	$\models (\alpha \supset \beta) \supset (\alpha \rightarrow \beta)$	(i), (iii)
(v)	$\models \sim(\alpha \supset \beta) \supset (\alpha \equiv \top)$	truth-functional logic
(vi)	$\models (\alpha \equiv \top) \supset ((\alpha \rightarrow \beta) \equiv (\top \rightarrow \beta))$	extensionality in antecedent
(vii)	$\models (\alpha \equiv \top) \supset ((\alpha \rightarrow \beta) \supset \beta)$	(vi), (IV)
(viii)	$\models \sim(\alpha \supset \beta) \supset ((\alpha \rightarrow \beta) \supset \beta)$	(v), (vii)
(ix)	$\models \sim(\alpha \supset \beta) \supset \sim\beta$	truth-functional logic
(x)	$\models \sim(\alpha \supset \beta) \supset \sim(\alpha \rightarrow \beta)$	(viii), (ix)
(xi)	$\models (\alpha \rightarrow \beta) \equiv (\alpha \supset \beta)$	(iv), (x)

Conversely, suppose that that → is extensional in its consequent. To show that → is extensional in its antecedent, it suffices to show that $\alpha \rightarrow \beta$ is logically equivalent to the material conditional $\alpha \supset \beta$ (which is extensional in its antecedent). We argue thus:

(i)	$\models (\alpha \supset \beta) \supset (\beta \equiv (\alpha \vee \beta))$	truth-functional logic
(ii)	$\models (\beta \equiv (\alpha \vee \beta)) \supset ((\alpha \rightarrow \beta) \equiv (\alpha \rightarrow (\alpha \vee \beta))$	extensionality in consequent
(iii)	$\models (\beta \equiv (\alpha \vee \beta)) \supset (\alpha \rightarrow \beta)$	(ii), (II)
(iv)	$\models (\alpha \supset \beta) \supset (\alpha \rightarrow \beta)$	(i), (iii)
(v)	$\models \sim(\alpha \supset \beta) \supset (\beta \equiv \bot)$	truth-functional logic
(vi)	$\models (\beta \equiv \bot) \supset ((\alpha \rightarrow \beta) \equiv (\alpha \rightarrow \bot))$	extensionality in consequent
(vii)	$\models (\beta \equiv \bot) \supset ((\alpha \rightarrow \beta) \supset \sim\alpha)$	(vi), (III)
(viii)	$\models \sim(\alpha \supset \beta) \supset ((\alpha \rightarrow \beta) \supset \sim\alpha)$	(v), (vii)
(ix)	$\models \sim(\alpha \supset \beta) \supset \alpha$	truth-functional logic
(x)	$\models \sim(\alpha \supset \beta) \supset \sim(\alpha \rightarrow \beta)$	(viii), (ix)
(xi)	$\models (\alpha \rightarrow \beta) \equiv (\alpha \supset \beta)$	(iv), (x)

Thus → is extensional in its antecedent iff it is extensional in its consequent, and it is extensional in either position only if it is logically equivalent to the material conditional. Conversely, it is trivial that if → is logically equivalent to the material conditional, then it is extensional in both its antecedent and its consequent.

Notes

1 The term ‘hyperintensional’ is sometimes used with the sense given to ‘congruential’ below, but the present definition is more closely related to the usual understanding of ‘intension’.

2 Dropping the initial $\Box$ in the consequent of NON-HYPERINTENSIONAL evades the immediate problem, for $\Box(\boldsymbol{\alpha} \equiv \boldsymbol{\beta}) \supset (\Box\boldsymbol{\alpha} \equiv \Box\boldsymbol{\beta})$ is a theorem of all normal modal systems. However, that makes worse problems in other directions, for $\Box(\mathbf{p} \equiv \top) \supset (\Box\Box\mathbf{p} \equiv \Box\Box\top)$ fails when accessibility is non-transitive, so $\langle\Box\Box\mathbf{p}, \mathbf{p}\rangle$ fails the new condition. Thus the composition of contexts that satisfy the new condition need not satisfy the new condition.

3 **O** feels somewhat metalinguistic, since its logic is not closed under uniform substitution. §3 provides more natural examples of hyperintensional congruential contexts.

4 More generally, $\models \Box\ldots\Box\mathbf{A}\boldsymbol{\alpha} \equiv \mathbf{A}\boldsymbol{\alpha}$ for any number of iterations of $\Box$, so we need not worry about objections to S4 here.

5 Humberstone 2004 traces the distinction back to Meredith and Prior 1965. Incidentally, the generality validity of $\boldsymbol{\alpha}$ need not be recoverable as the real-world validity of $\Box\boldsymbol{\alpha}$, for although the general validity of $\boldsymbol{\alpha}$ implies the general validity of $\Box\boldsymbol{\alpha}$ and

therefore its real-world validity, the latter means only that for every model M and every $u \in W(M)$, $M,a(M)[u] \models$ **α**, which does not imply the general validity of **α** because there may be points $x \in P(M)$ not of the form a(M)[u]. The latter coincides with x in its world component when u is w(x), but even a(M)[w(x)] may differ from x in some other component. Operators generalizing over all those other dimensions would need to be combined with □ to capture general validity as the real world validity of the result of prefixing operators to the relevant formula. Similarly, the general validity of **Aα** implies the real world validity of **α** but not vice versa.

6 The rule of necessitation is standard in axiomatizations of S5. Since it does not preserve real-world validity, one cannot argue directly by induction on the length of proofs that all theorems of S5 are real-world valid. Rather, one argues by induction that all theorems of S5 are generally valid, and then deduces that they are all real-world valid.

7 Many of the English sentences quoted in the text have several readings. It should be clear which ones are relevant.

8 One might try to quantify over the 'epistemic' dimension of our models with an operator ■ such that ■**α** is true at a point in a model M iff **α** is true (at the actual point) in every model that differs from M at most in which world is actual; ■ is equivalent to the 'Fixedly actually' operator that Davies and Humberstone (1980) propose as corresponding to the notion of 'deep necessity' (Evans 1979; □ corresponds to 'superficial necessity'). Call the result of replacing □ by ■ throughout the definition of

‘non-hyperintensionality’ ‘deep non-hyperintensionality’. Since, for ■, all theorems of S5 are real-world valid and necessitation preserves real-world validity, irrespective of what special operators the language contains, reasoning like that towards the end of §1 shows that extensionality, deep non-hyperintensionality and congruentiality are so ordered in decreasing logical strength. For present purposes, however, the contrast between ‘superficial’ non-hyperintensionality and congruentiality is exactly the point of interest.

9 The situation for first-order logic is subtler than the main text suggests, although not in ways that undermine the present argument; see Rayo and Williamson 2003. Strictly speaking, the talk of ‘models’ throughout this paper should be modified as explained there.

10 The latter distinction also has a logical element, represented by ‘ $\models$’ in NON-HYPERINTENSIONAL.

11 Non-hyperintensionality in the consequent of subjunctive implication is a special case of the principle that if **α** subjunctively implies **β** and **β** strictly implies **γ** then **α** subjunctively implies **γ**, which is valid on standard possible worlds semantics for subjunctive conditionals. By contrast, there is no such plausible generalization of non-hyperintensionality in the antecedent of subjunctive implication: if **α** strictly implies **β** and **β** subjunctively implies **γ**, still **α** may not subjunctively imply **γ**. Perhaps, if it had rained, it would not have rained hard; necessarily, if it rained hard then it rained; it does not follow that if it had rained hard, it would not have rained hard.

12 The supervenience principle must be rejected by those who allow some but not all subjunctive conditionals with impossible antecedents to be false (or at least contingent). They may include those who reject the S4 schema for possibility, for if X is contingently impossible while Y is necessarily impossible, so that X and Y are contingently necessarily equivalent, it is tempting to evaluate 'If X had happened, X would have happened' as true and 'If X had happened, Y would have happened' as false. As already noted, it is best for present purposes to work with a notion of possibility that satisfies S4.

13 POSSIBILITY (combined with REFLEXIVITY) is discussed and applied in Williamson 2005.

14 Objection: 'If the intuitionists are right, that problem is not not decidable' does not imply 'If the intuitionists are right, that problem is decidable', because according to the intuitionists 'That problem is not not decidable' does not imply 'That problem is decidable'. Reply: Allowing such examples leaves almost nothing of the logic of indicative conditionals. We are not obliged to build a non-standard probability theory on which the probability of 'That problem is not not decidable' conditional on 'The intuitionists are right' is higher than the probability of 'That problem is decidable' conditional on 'The intuitionists are right'. Similarly, we are not obliged to build a non-standard logic of indicative conditionals on which 'If the intuitionists are right, that problem is not not decidable' does not imply 'If the intuitionists are right, that problem is decidable'.

15 Congruentiality in the consequent of indicative implication is a special case of the plausible principle that if **α** indicatively implies **β** and **β** logically implies **γ** then **α** indicatively implies **γ**. The corresponding generalization of congruentiality in the antecedent of indicative implication is less clear: if **α** logically implies **β** and **β** indicatively implies **γ**, does **α** indicatively imply **γ**? That it will rain hard logically implies that it will rain; if it rains, it will not rain hard (we may assume); does it follow that if it rains hard, it will not rain hard? On the truth-functional reading, the conclusion *does* follow and is vacuously true.

16 Weatherson 2001 gives a theory of indicative conditionals as epistemic and subjunctive conditionals as 'metaphysical', in a two-dimensional semantic setting. He emphasizes the difference between the two types of conditional in their interaction with 'actually'. Jackson (1981: 129; 1987: 74-75) first called attention to the difference as an argument against attempts by Stalnaker 1976 and Davis 1979 to assimilate indicative 'if' semantically to subjunctive 'if'.

17 Thanks to audiences in Oxford, including a meeting of the Luxemburg group, and especially to Dorothy Edgington and John Hawthorne, for discussion of earlier versions of this material, and to Lloyd Humberstone and Brian Weatherson for written comments.

References

Chalmers, D. 2004. ‘Epistemic two-dimensional semantics’. *Philosophical Studies* 118: 153-226.

Crossley, J., and Humberstone, L. 1977. ‘The logic of “actually”’. *Reports on Mathematical Logic* 8: 11-29.

Davies, M., and Humberstone, L. 1980. ‘Two notions of necessity’. *Philosophical Studies* 38: 1-30.

Davis, W. 1979. ‘Indicative and subjunctive conditionals’. *Philosophical Review* 87: 544-564.

Edgington, D. 1986. ‘Do conditionals have truth-conditions?’, *Crítica* 18: 3-30. Reprinted in Jackson 1991.

Evans, G. 1979. ‘Reference and contingency’. *Monist* 62: 161-189.

Humberstone, L. 1986. ‘Extensionality in sentence position’. *Journal of Philosophical Logic* 15: 27-54.

Humberstone, L. 2004. ‘Two-dimensional adventures’. *Philosophical Studies* 118: 17 -65.

Jackson, F. 1981. ‘Conditionals and possibilia’. *Proceedings of the Aristotelian Society* 81: 125-137.

Jackson, F. 1987. *Conditionals*. Oxford: Blackwell.

Jackson, F. (ed.) 1991. *Conditionals*. Oxford: Oxford University Press.

Jackson, F. 2004. ‘Why we need A-intensions’. *Philosophical Studies* 118: 257-277.

Kaplan, D. 1989. ‘Demonstratives’, in J. Almog, J. Perry and H. Wettstein (eds.), *Themes* *from Kaplan*. Oxford: Oxford University Press.

Meredith, C., and Prior, A. 1965. ‘Modal logic with functorial variables and a contingent constant’. *Notre Dame Journal of Formal Logic* 6: 99-109.

Rayo, A., and Williamson, T. 2003. ‘A completeness theorem for unrestricted first-order languages’, in J.C. Beall (ed.), *Liars and Heaps*. Oxford: Clarendon Press.

Salmon, N. 1989. ‘The logic of what might have been’. *Philosophical Review* 98: 3-34.

Stalnaker, R. 1976 ‘Indicative conditionals’, in A. Kasher (ed.), *Language in Focus*. Dordrecht: Reidel.

Weatherson, B. 2001. ‘Indicative and subjunctive conditionals’. *Philosophical Quarterly* 51: 200-216.

Williamson, T. 2005. ‘Armchair philosophy, metaphysical modality and counterfactual thinking’. *Proceedings of the Aristotelian Society* 105: 1-23.

Zalta, E. 1988. ‘Logical and analytic truths that are not necessary’. *Journal of Philosophy* 85: 57-74.