1 What is a nonclassical logic?

A nonclassical logic is a logic that's not classical. Think: nonclassical music, or nonbread food.

1.1 What's a logic?

For my purposes, a logic (or consequence relation) is a *binary relation* between *premise-collection-thingies* and *conclusion-collection-thingies*. Much of the time, we can take a premise-collection-thingy to be a *set* of premises, and a conclusion-collection-thingy to be a *set* of conclusions; sometimes, we will take one or the other to be a *single* premise or conclusion. (Other options: multisets, lists, trees, richer structures.)

This is a distinctively liberal conception; it's really not much of a concrete conception at all. That's fine; this is intended to be the umbrella idea that lots of more particular cases will fall under. (Think, again: music, or food. But it still rules out, say, 3-place relations; hmm. . .)

1.1.1 Important case 1: RMT single-conclusion

One important kind of case: consequence as a relation $\vdash$ between *sets* of premises and *single* conclusions, subject to the following assumptions:

(R)eflexivity: $A \vdash A$

(M)onotonicity: If $\Gamma \vdash A$, then $\Gamma, \Gamma' \vdash A$

(T)ransitivity: If $\Gamma \vdash A$ for every $A \in \Gamma'$, and $\Gamma, \Gamma' \vdash B$, then $\Gamma \vdash B$

If $\vdash$ is compact, then (T) can be weakened to (T_{fin}): if $\Gamma \vdash A$ and $\Gamma, A \vdash B$, then $\Gamma \vdash B$. Both are strictly stronger than what we might usually call 'transitivity'.

1.1.2 Important case 2: RMT set-set

Another important kind of case: consequence as a relation $\vdash$ between *sets* of premises and *sets* of conclusions, subject to the following assumptions:

(R)eflexivity: $A \vdash A$

(M)onotonicity: If $\Gamma \vdash \Delta$, then $\Gamma, \Gamma' \vdash \Delta, \Delta'$

(T)ransitivity: If there is some set Σ such that $\Gamma, \Sigma_1 \vdash \Sigma_2, \Delta$ for every partition $\langle \Sigma_1, \Sigma_2 \rangle$ of Σ, then $\Gamma \vdash \Delta$

If $\vdash$ is compact, then (T) can be weakened to (T_{fin}): if $\Gamma \vdash A, \Delta$ and $\Gamma, A \vdash \Delta$, then $\Gamma \vdash \Delta$.

1.1.3 Relations between single-conclusion and set-set

Say that an RMT single-conclusion relation $\vdash_1$ and an RMT set-set relation $\vdash_2$ *agree* iff: $\Gamma \vdash_1 A$ iff $\Gamma \vdash_2 A$. Different set-set relations can agree with the same single-conclusion relation. (Indeed, *every* RMT single-conclusion relation agrees with more than one RMT set-set relation.)

The *weakest* RMT set-set relation that agrees with an RMT single-conclusion relation $\vdash_1$ is always $\{\langle \Gamma, \Delta \rangle : \exists \delta \in \Delta \text{ such that } \Gamma \vdash_1 \delta\}$. There may be no strongest. Every RMT set-set relation that agrees with an RMT single-conclusion relation $\vdash_1$ is contained in $\{\langle \Gamma, \Delta \rangle : \forall \Sigma, C \text{ if } \Sigma, \delta \vdash_1 C \text{ for every } \delta \in \Delta, \text{ then } \Sigma, \Gamma \vdash_1 C\}$. (Indeed, this is the union of all such relations.) But this may not *itself* be an RMT set-set relation: it always obeys (T_{fin}), but maybe not (T). If $\vdash_1$ is compact, then this relation obeys (T), and so is the strongest RMT set-set relation agreeing with $\vdash_1$—even in this case, it itself is *not* guaranteed to be compact.

1.2 What's a *classical* logic?

This is a harder question than it gets credit for. There are easy *paradigm cases*: classical propositional logic, or classical first-order logic (with or without equality), classical higher-order logics, etc. And there are easy *paradigm noncases*: intuitionist logics, or relevant logics, or paraconsistent logics, etc. Some distinctive nonclassicalities: in intuitionist logic, $A \vee \neg A$ is not a theorem, and $\neg(A \wedge B) \nvdash \neg A \vee \neg B$; in relevant logics, $A \rightarrow (B \rightarrow A)$ is not a theorem; in paraconsistent logics, $A \wedge \neg A \nvdash B$. But in between these paradigm cases and noncases there are a number of distinctions often overlooked.

The trouble comes around *metainferences*: closure principles that a consequence relation may or may not obey. (M) and (T) above are good examples. So is 'proof by cases', the principle that if $\Gamma, A \vdash C$ and $\Gamma, B \vdash C$, then $\Gamma, A \vee B \vdash C$. Proof by cases can *seem* to be required by classicality. But there are RMT single-conclusion relations that validate all classically-valid arguments and still are not closed under proof by cases. These are sometimes called 'weakly classical'. (The same is not true of RMT set-set relations, owing to $A \vee B \vdash A, B$, (M), and (T).) Outside the safety net provided by RMT, things are richer still, as we'll see when we look at monotonic logics.

For most of this week, then, 'classical logic' (and so 'nonclassical logic') will no more be precisely specified than 'logic' itself is. For particular purposes, it often helps to pin down a particular distinction, but overall, we're looking at something more akin to genres or cuisines here, with the associated fuzz around the edges.

1.3 Weakenings

Many nonclassical logics are (or can be seen as) *weakenings* of classical logic; they declare invalid some arguments that are classically valid. (Some exceptions: syllogistic ($\forall/\exists$), connexive logics ($\neg(A \to \neg A)$).)

Why weaken classical logic? Four main (related) motivations:

Correctness: Classical logic validates things that just ain't valid! Familiar complaints:

- explosion: $A \wedge \neg A \vdash B$
- excluded middle: $\vdash A \vee \neg A$
- almost everything to do with the conditional: eg $(A \wedge B) \to (C \vee D) \vdash (A \to C) \vee (B \to D)$ or $\neg(A \to B) \vdash A \wedge \neg B$

Strengthening: Weaken *here* in order to strengthen *there*. Develop theories that would classically collapse. Examples:

- Smooth infinitesimal analysis has axioms that are classically inconsistent, but intuitionistically consistent, and is pursued in intuitionist logic. (Infinitesimals are not not zero, but it must not follow that they *are* zero.)
- Naive set theory entails everything classically, via Russell or Burali-Forti. But it can be developed in paraconsistent settings.

Distinctions: Logics are often 'built in at the bottom' of other theories; say, of propositional attitudes, or of necessity. But there are many highly nontrivial classical validities—we don't want to wire *these* in too deeply, at the cost of missing the distinction between, say, believing A and believing $A \wedge B$ where B is some incredibly complicated classical tautology.

Explanatory oomph: Even when we know something holds, we still want to know *why* it holds. Did it depend on distribution of conjunction over disjunction, or did that not matter? A natural way to proceed: explore systems in which distribution fails.

2 Proofs and models

There are two main ways to present a consequence relation: via *proofs* and via *models*. Each comes in more varieties than you can shake a stick at. But the distinction is reasonably firm: *proofs* are things such that *validity* is demonstrated by getting one of them (and so invalidity amounts to a universal claim), while *models* are things such that *invalidity* is demonstrated by getting one of them (and so validity amounts to a universal claim).

Soundness and completeness relate a proof theory to a model theory. When they both hold between a proof theory and a model theory, the two notions of validity (and invalidity) are extensionally equivalent. Soundness is a *no-clash* condition: there aren't too many proofs/models; no proof shows something to be valid that a model shows to be invalid. Completeness is a, well, *completeness* condition: there are enough proofs/models; every argument either has a proof (and so is valid) or a countermodel (and so is invalid).

2.1 Proofs

To give a proof theory is to specify a bunch of objects called *proofs*, and to say which of them establish the validity of which arguments.

$$\text{Id: } \frac{}{A : A} \qquad \text{K: } \frac{\Gamma : \Delta}{\Gamma, \Gamma' : \Delta, \Delta'} \qquad \neg\text{L: } \frac{\Gamma : A, \Delta}{\Gamma, \neg A : \Delta} \qquad \neg\text{R: } \frac{\Gamma, A : \Delta}{\Gamma : \neg A, \Delta}$$

$$\wedge\text{L: } \frac{\Gamma, A_i : \Delta}{\Gamma, A_0 \wedge A_1 : \Delta} \qquad \wedge\text{R: } \frac{\Gamma : A, \Delta \qquad \Gamma : B, \Delta}{\Gamma : A \wedge B, \Delta} \qquad \vee\text{L: } \frac{\Gamma, A : \Delta \qquad \Gamma, B : \Delta}{\Gamma, A \vee B : \Delta} \qquad \vee\text{R: } \frac{\Gamma : A_i, \Delta}{\Gamma : A_0 \vee A_1, \Delta}$$

Figure 1: Classical logic in sequent form

2.1.1 Axiomatic

Axiomatic proof systems consist of axioms and rules; an axiomatic proof of a formula shows thatit follows from the axioms, given the rules. For example, here is an axiomatic presentation of the $\rightarrow$ fragment of the relevant logic R:

- $A \rightarrow A$
- $(A \rightarrow B) \rightarrow ((C \rightarrow A) \rightarrow (C \rightarrow B))$
- $(A \rightarrow (B \rightarrow C)) \rightarrow (B \rightarrow (A \rightarrow C))$
- $(A \rightarrow (A \rightarrow B)) \rightarrow (A \rightarrow B)$
- From A and $A \rightarrow B$, infer B

Here the axioms and rule are *schemas*: any formula of the right form counts as an axiom, and any inferential move of the right form is permitted by the rule. The $\rightarrow$ fragment of intuitionist logic can be reached by adding $A \rightarrow (B \rightarrow A)$, and the $\rightarrow$ fragment of classical logic can be reached from there by adding Peirce's Law $((A \rightarrow B) \rightarrow A) \rightarrow A$.

Two ways to get *argument validity* (rather than just formula validity) from axiomatic systems: 1) encode arguments as formulas in some way—so, say, $A, B, C \vdash D$ iff $(A \wedge B \wedge C) \rightarrow D$ has an axiomatic proof (or if $A \rightarrow (B \rightarrow (C \rightarrow D))$ does; or 2) consider expansions of the axiomatic system—so, say, $A, B, C \vdash D$ iff the system expanded with axioms A, B, C can prove D. Whether or not these ways are equivalent is obviously a system-by-system question. Moreover, whether an axiomatic system determines an RMT consequence relation can depend on how this choice is made. The second option will always result in an RMT single-conclusion relation, but the first need not.

2.1.2 Sequents

A sequent system is like an axiomatic system, but 1) its objects are *sequents* rather than individual formulas, and 2) it tends to have few axioms and many rules, while axiomatic systems tend to have many axioms and few rules.

A sequent represents an entire argument. If we're dealing with a consequence relation on *sets* of premises and conclusions, then our sequents will be of the form $\Gamma : \Delta$, where Γ and Δ are sets. If we have a set-formula consequence relation, sequents will be $\Gamma : C$, where Γ is a set and C a formula. Consequence relations on more complex relata will use sequents with more complex components.

Figure 1 contains a sequent calculus for classical logic, using set-set sequents. Here, the horizontal line tells you that if you have derived a sequent of the form given above the line, you may derive a sequent of the form given below the line. (Derivations themselves will be tree-shaped, as in the axiomatic case.)

Sequents are closely connected to the *trees* or *tableaux* that are perhaps still more popular. The main differences are: 1) sequent derivations are usually written with the root at the *bottom*, while trees are usualy written with the root at the *top*; 2) trees often have individual formulas at nodes, rather than entire sequents. This latter difference is the main one, and exactly how the difference is realized depends on the particular system in play. Note that, in the classical case, the $\neg$ rules allow for dispensing with the left/right distinction; anything on one side is equivalent to its negation on the other.

Any sequent calculus can determine consequence relations in multiple ways. The most common notion of validity for a sequent calculus, called *internal* validity, is as follows: an argument is valid according to a sequent calculus iff the sequent that represents it is derivable in the calculus. Another notion, *external* validity, is sometimes considered: an argument from premises A, B, C to a conclusion D is externally valid in a sequent calculus iff the sequent $: D$ can be derived in the calculus obtained by adding $: A$, $: B$, and $: C$ as axioms. Note that external validity is always an RMT single-conclusion relation; not so for internal validity.

2.2 Models

To give a model theory is to give a bunch of objects called *models*, and to say which of them establish the invalidity of which arguments (are *countermodels* or *counterexamples* to those arguments).

2.2.1 Matrices

One common approach is based on *matrices*. To give a matrix-based model theory, one specifies a particular set V of *values* and an $n$ary operation Op_C on V for every $n$ary connective C of the language in question. This is a matrix. A *model*, on this approach, is a function ν from the language to V such that, for every complex sentence $C(A_1, \ldots, A_n)$: $\nu(C(A_1, \ldots, A_n)) = Op_C(\nu(A_1), \ldots, \nu(A_n)))$. (That is, ν is a *homomorphism* from the language to the matrix.) There are a number of options for saying when a model is a countermodel to a particular argument. Let's consider arguments with *sets* of premises and conclusions.

On one option, there is a set $D \subseteq V$ of *designated values*; then a model ν is a counterexample to an argument $\Gamma : \Delta$ iff: $\nu(\gamma) \in D$ for every $\gamma \in \Gamma$ and $\nu(\delta) \notin \Delta$ for every $\delta \in \Delta$. This always produces an RMT relation.

On another option, there is a lattice order $\leq$ on V, which has *greatest lower bounds* and *least upper bounds* w/r/t the order; then a model ν is a counterexample to an argument $\Gamma : \Delta$ iff $\text{glb}\{\nu(\gamma)\}_{\gamma\in\Gamma} \not\leq \text{lub}\{\nu(\delta)\}_{\delta\in\Delta}$. Again, this always produces an RMT relation (now modulo concerns about infinitary cases).

On yet another option, there is a set $D_P \subseteq V$ of *premise-designated values* and a set $D_C \subseteq V$ of *conclusion-designated values*; then a model ν is a counterexample to an argument $\Gamma : \Delta$ iff: $\nu(\gamma) \in D_P$ for all $\gamma \in \Gamma$ and $\nu(\delta) \notin D_C$ for all $\delta \in \Delta$. This does *not* always produce an RMT relation, although it is always monotonic full stop, always reflexive if $D_P \subseteq D_C$, always transitive if $D_C \subseteq D_P$, and reduces to the simple designated-value approach if $D_P = D_C$.

There are other options available as well. For example, Weir's 'neoclassical' consequence (Tuesday) involves a more complicated condition.

2.2.2 Algebras

A more general approach is the *algebraic* one. On the algebraic approach, instead of specifying a particular set of values and operations, we instead specify a class of *algebras* of the same similarity type as the language in question. A model is now a homomorphism ν from the language to an algebra in the class.

Again, there are a number of ways to define consequence; most of these are natural generalizations of the matrix strategies surveyed above.

2.2.3 Frames

A *frame* gives us a way of bundling up a bunch of structures together. Usually, a frame-based model consists of a set W of matrix models (call them *points* in this context) built on the same matrix, together with some relation or relations between these models. Some of the vocabulary in the language can be interpreted or constrained point-by-point; its behaviour can be dtermined by the underlying matrix. The fun of frames comes with other vocabulary, which is interpreted or constrained in ways related to the relation or relations in question.

Counterexamples can again be handled in any of the ways they are for matrices. Sometimes frames come with a distinguished point, or a distinguished set of points; these may or may not play a role in what it takes to be a countermodel.

2.2.4 Abstract valuations

So long as we stay within the realm of RMT relations, there is a kind of universal model theory available via *abstract valuations*. A valuation is a function from the language to $\{1, 0\}$, and a valuation v is a counterexample to an argument $\Gamma : \Delta$ iff $v(\gamma) = 1$ for all $\gamma \in G$ and $v(\delta) = 0$ for all $\delta \in \Delta$. A consequence relation can then be given by giving a particular set of valuations; such relations are always RMT.

Note that valuations are allowed to entirely ignore the connectives of the language. This added flexibility ensures that *every* RMT consequence relation is sound and complete for some set of valuations. If the relation is set-set, the set of valuations is unique; if it is single-conclusion, there will be multiple such sets.

1 Vagueness

A simple binary choice between true and false does not seem to get the right kind of grip on how we use vague predicates. This is one motivation for developing *many-valued* logics.

1.1 Strong kleene models

Strong kleene models are matrix-based models; the matrix in question has three values, which I'll call 1, $\frac{1}{2}$, and 0. For connectives $\wedge$, $\vee$, and $\neg$, $Op_\wedge = \min$; $Op_\vee = \max$; and $Op_\neg(x) = 1 - x$. The idea here is that 1 and 0 behave just like classical truth and falsity, while $\frac{1}{2}$ sits in between, providing a value for borderline-case sentences to take.

There are a number of different consequence relations associated with these models, depending on what it takes to be a counterexample to an argument.

1.1.1 K3

A strong kleene model ν is a K3-counterexample to an argument $\Gamma : \Delta$ iff $\nu(\gamma) = 1$ for all $\gamma \in \Gamma$ and $\nu(\delta) \neq 1$ for all $\delta \in \Delta$. This is a designated-value approach, with 1 the only designated value. (The 'K' in 'K3' stands for 'Kleene'; Stephen Kleene developed this logic to help think about partial recursive functions. The '3' is for its three-valued matrix presentation.)

In K3, a set of sentences is satisfiable iff it is classically satisfiable (since the value 1—which is just like classical truth—is the only designated value). But argument validity differs. You can think of K3 as:

- The usual conjunction-disjunction fragment of classical logic, plus
- all de morgan and double-negation equivalences, plus
- explosion ($A, \neg A \vdash B$).

Crucially, K3 does *not* validate excluded middle ($\vdash A \vee \neg A$); that is, it is *paracomplete*. In fact, K3 has no *theorems* at all. The nonclassicalities of K3 can be seen as following from this. For example:

- $A \nvdash A \wedge B, A \wedge \neg B$
- $A \nvdash B, \neg(A \supset B)$
- K3 validity does not obey the sequent rule ¬R

If we take designation to be directly linked to assertibility or truth, then a K3-based theory of vagueness will predict that borderline sentences are not assertible, not true, etc. Moreover, it will predict that conjunctions, disjunctions, and negations of borderline sentences will share this status.

1.1.2 LP

A strong kleene model ν is an LP-counterexample to an argument $\Gamma : \Delta$ iff $\nu(\gamma) \neq 0$ for all $\gamma \in \Gamma$ and $\nu(\delta) = 0$ for all $\delta \in \Delta$. This is a designated-value approach, with 1 and $\frac{1}{2}$ as designated values. ('LP' stands for 'Logic of Paradox'; it was given this name by Graham Priest, and earlier developed by FG Asenjo.)

In LP, a set of sentences is a conseqeunce of no premises iff it is classically (since the value 0—which is just like classical falsity—is the only undesignated value). But argument validity, again, differs. You can think of LP as:

- The usual conjunction-disjunction fragment of classical logic, plus
- all de morgan and double-negation equivalences, plus
- excluded middle.

Crucially, LP does *not* validate explosion; that is, it is *paraconsistent*. In fact, *no* set of sentences is unsatisfiable in LP. The nonclassicalities of LP can be seen as following from this. For example:

- $A \vee B, A \vee \neg B \nvdash A$
- $A, A \supset B \nvdash B$.

- LP validity does not obey the sequent rule ¬L

K3 and LP are mirror images (duals) in the following sense: $\Gamma \vdash_{K3} \Delta$ iff $\neg\Delta \vdash_{LP} \neg\Gamma$, and $\Gamma \vdash_{LP} \Delta$ iff $\neg\Delta \vdash_{K3} \neg G$. So it's easy to transfer facts and intuitions back and forth between them.

If we take designation to be directly linked to assertibility or truth, then an LP-based theory of vagueness will predict that borderline sentences are assertible, true, etc. Moreover, it will predict that conjunctions, disjunctions, and negations of borderline sentences will share this status.

1.1.3 FDRM

A strong kleene model ν is an FDRM-counterexample to an argument $\Gamma : \Delta$ iff $\min\{\nu(\gamma)\}_{\gamma\in\Gamma} \not\leq \max\{\nu(\delta)\}_{\delta\in\Delta}$. ('FDRM' is so-called for being the first degree fragment of the logic RM, which itself is so-called for extending the relevant logic R with an axiom known as Mingle.)

Note that a strong kleene model is an FDRM-counterexample to an argument iff it is either a K3- or an LP-counterexample to that argument; as a result, an argument is FDRM-valid iff it is both K3- and LP-valid. You can think of FDRM as:

- The usual conjunction-disjunction fragment of classical logic, plus
- all de morgan and double-negation equivalences, plus
- $A \wedge \neg A \vdash B \vee \neg B$.

Without a notion of designation in play, it is not 100% clear that FDRM requires any particular predictions about the assertibility or truth of borderline sentences; certainly nothing is directly suggested by the matrix models.

1.1.4 ST

A strong kleene model ν is an ST-counterexample to an argument $\Gamma : \Delta$ iff: $\nu(\gamma) = 1$ for all $\gamma \in \Gamma$ and $\nu(\delta) = 0$ for all $\delta \in \Delta$. ('ST' comes from 'Strict' and 'Tolerant', the idea being that each model sets *two* consequence-relevant thresholds, a more strict one ($= 1$) and a more tolerant one (> 0).)

As things stand so far, ST determines exactly the usual set-set consequence relation of classical logic. But because of the two thresholds in the model, it is not guaranteed to remain transitive when extended with other vocabulary. (It is guaranteed to remain reflexive and monotonic.) This nontransitivity was inspired by the way soritical reasoning seems to work: every step is fine, but they cannot be all safely linked together.

1.2 Belnap-Dunn models

All the logics we've looked at so far contain FDRM. But there is a slightly weaker logic than FDRM that is better-known and in some ways more natural. This is the logic FDE, so-called because it is the first-degree fragment of the relevant logic E (and indeed of all the usual relevant logics).

You can think of FDE as:

- The usual conjunction-disjunction fragment of classical logic, plus
- all de morgan and double-negation equivalences.

This is very slightly weaker than FDRM. The usual model theory for FDE works with *Belnap-Dunn models*. These are matrix-based models with four values instead of three. One way to think of this: truth and falsity are independent of each other in FDE. This creates four possible combinations: true only, false only, both true and false, and neither true nor false.

For the connectives, say that in each model: a conjunction is true iff both conjuncts are, and false iff at least one conjunct is; a disjunction is true iff at least one disjunct is, and false iff both disjuncts are; and a negation is true iff its negatum is false, and false iff its negatum is true.

Finally, let a model be a counterexample to an argument iff all the argument's premises are true in the model and no conclusion is true in the model.

1.3 Łukasiewicz models

Just as the divide between the heaps and non-heaps is vague, so is the divide between the heaps and the borderline heaps. For this reason, you might think that three values aren't enough, and indeed that no finite number of values is enough. One nice place to explore here is *continuum-valued Łukasiewicz logic*.

In this logic, the matrix has as values all real numbers from 0 to 1 inclusive, usually though of as degrees of truth. Conjunction, disjunction, and negation are just as in strong kleene models: minimum, maximum, and $1-$. But there is a new connective to consider as well: a separate conditional $\rightarrow$, not reducible to $\neg$ and $\vee$. This works as follows: $A \rightarrow B$ falls as far short of complete truth as B's degree of truth falls short of A's. That is, for any model v: $v(A \rightarrow B) = min(1, 1 - (v(A) - v(B)))$.

Using this conditional, we can define a new conjunction $A \circ B$ as $\neg(A \rightarrow \neg B)$. This is like familiar conjunction in some ways: for example, $v(A \circ B) = v(B \circ A)$ in any model v, and $v(A \circ B) \leq v(A)$ in any model v. But it is also unlike familiar conjunction in some ways; for example, it does not always hold that $v(A \circ A) = v(A)$.

2 Truth-based paradoxes

Another motivation for exploring some of these logics comes from truth-based paradoxes, like 'This sentence is not true', or 'If this sentence is true, then Earth has nine moons'. We'd like to be able to have languages that contain their own truth predicates, and also be such that A and $T\langle A\rangle$ are completely equivalent, for any sentence A. (Here $T\langle A\rangle$ is the claim—within the object language itself—that A is true.)

But this leads us into trouble; if a sentence says of itself that it is not true, then it should be equivalent to its own negation. If our models are classical, this is not possible: 0 is never 1. But both strong kleene models and Belnap-Dunn models can work here: the value $\frac{1}{2}$ can be assigned to such sentences in the former, and there are two workable options in the latter (both true and false, and neither true nor false—note that here these are *metalanguage* truth and falsity).

2.1 The Kripke construction

The Kripke construction allows us to verify that this approach works. The idea is this: we start with a strong kleene model for the whole language *except* the predicate T. We then build up stage-by-stage to an interpretation for the full language *including* T as follows:

- At stage 0, every T sentence takes value $\frac{1}{2}$.
- At stage $n+1$, $T\langle A\rangle$ takes whatever value A took at stage n.
- At limit stages, $T\langle A\rangle$ takes value 1 if it took value 1 somewhere below, value 0 if it took value 0 somewhere below, and $\frac{1}{2}$ otherwise.

This way of proceeding ensures that once a sentence takes value 1 or 0, it stays there forever. (This depends on the details of the strong kleene matrix.) Because of this, we can eventually find a fixed point: a stage n that is just the same as stage $n+1$. This stage gives us our target model for the full language.

Just as in the case of vagueness, there are options for defining the countermodel relation, and so consequence itself, in this case. But these fixed points always assign the same value to A and $T\langle A\rangle$; so as long as which things are countermodels to which arguments depends only on the values taken by sentences, A and $T\langle A\rangle$ will be totally intersubstitutable for each other according to the eventual consequence relation, whether it is K3, LP, FDRM, or something else.

Moreover, the resulting consequence relation will always be a *conservative extension* of the original, T-free one. That is, an argument *not* containing T will come out as valid in the consequence relation based on fixed points iff the same argument is valid in the consequence relation based on strong kleene models in the same way. This is because any model for the T-free language can serve as the input for the Kripke construction.

1 Modal logics

1.1 Normal modal logics

Normal modal logics (this is a technical term) are built on top of classical logic. They involve two additional one-place connectives $\Box$ and $\Diamond$. These can be read as 'necessary' and 'possible', or as 'must' and 'can', or as 'should' and 'may', or in a number of other ways (although certain interpretations fit better than others).

What makes something a normal modal logic is that it obeys the following restrictions (in addition to full classical logic):

- $\Box A$ is equivalent to $\neg\Diamond\neg A$
- $\Box(A \to B) \to (\Box A \to \Box B)$ is a theorem, and
- Whenever A is a theorem, $\Box A$ is a theorem too.

The most usual model theory for normal modal logics is one based on *frames*.

1.1.1 Frames and models for NMLs

A *frame* is a set W, together with a binary relation R on the set. You can think of the members of W as possible worlds, with the fascinating and contentious metaphysics and epistemology that comes along, or in some less fraught way. But the idea of possible worlds gets at something: the members of W in some sense represent *alternatives*. R gives us the alternativeness relation: the worlds x such that wRx are those worlds that are in some sense *possible from* w's perspective, the worlds w sees as possible.

Given a frame $\langle W, R\rangle$, a *model* on that frame assigns classical values (1 or 0) to the sentences in our language—but there is a (possibly) different assignment at each world. As usual, atomic sentences obey no constraints; they can be assigned their values willy-nilly. Compounds formed from the usual classical connectives obey the usual classical constraints *world by world*: so a conjunction takes value 1 *at a world* iff both its conjuncts take value 1 *at that same world*, and so on.

The relation R is involved in giving constraints for $\Box$ and $\Diamond$. A sentence $\Box A$ takes value 1 at a world w iff A takes value 1 at *every* world x such that wRx. Dually, a sentence, $\Diamond A$ takes value 1 at a world w iff A takes value 1 at *some* world x such that wRx. (Because of the way 'every' and 'some' occur here, $\Box$ and $\Diamond$ end up behaving in some ways like simplified quantifiers.)

Finally, a countermodel to an argument $\Gamma : \Delta$ is a model M that contains a world w such that w assigns value 1 to everything in Γ and 0 to everything in Δ.

1.1.2 Variations

The smallest normal modal logic is called K. This is the logic determined by the full class of models given above, or by the axiomatic description in the bullet points above. K already exhibits rich and nontrivial behavior. For example, $\Box(A \wedge B)$ is equivalent to $\Box A \wedge \Box B$; and $\Box A \vee \Box B$ entails $\Box(A \vee B)$, but not vice versa.

There are a number of standard ways to extend K, either by adding axioms or by restricting the space of models under consideration. The next table gives some usual suspects, although there are plenty more:

Logic	Axiom	Restriction
D	$\Box A \to \Diamond A$	Every point sees some point.
T	$\Box A \to A$	R is reflexive.
S4	T + $\Box A \to \Box\Box A$	T + R is transitive.
S5	S4 + $\Diamond A \to \Box\Diamond A$	R is an equivalence relation.

1.1.3 Conditionals

The original motivation for developing modal logics was to avoid some of the problems with the classical $\to$ as a model of conditionality. The original development added a new two-place conditional connective $\Rightarrow$ rather than the one-place connectives that are more common today. We can think of $A \Rightarrow B$ as $\Box(A \to B)$; this is called a *strict conditional*.

This exhibits better behavior as a model of conditionality than $\to$. For example, we don't have $\neg(A \Rightarrow B)$ entailing $A \wedge \neg B$; it only entails $\Diamond(A \wedge \neg B)$. Similarly, we don't have $(A \wedge B) \Rightarrow C$ entailing $(A \Rightarrow C) \vee (B \Rightarrow C)$. $\Rightarrow$ also preserves some of the conditional-like behavior we do want. So long as we have the T axiom, modus ponens is still valid; if the argument from A to B is valid, then $A \Rightarrow B$ is a theorem, etc.

1.2 Kicking the tires

We can change things up, however. One way is to change the conditions that constrain $\Box$ and $\Diamond$. Indeed, this is required to capture some of the original modal systems. (Of CI Lewis's original S1–S5, only S4 and S5 are normal.) Another way is to swap the underlying classical logic out for something else entirely. And of course, these ideas can be combined.

1.2.1 Nonnormal modal logic

One way to change things up is to add 'nonnormal worlds'. So a frame will now consist of a set W of worlds, where each world is either normal or nonnormal (but not both), and a binary relation R on W. When we build models on these frames: at normal worlds, everything is as before; but at nonnormal worlds, *no* sentence of the form $\Box A$ is ever true, and *every* sentence of the form $\Diamond A$ is true. That is, at these worlds, anything is possible, and nothing is necessary. (They still obey classical logic for the classical connectives, however.)

There are two natural ways to say what a counterexample is: we can either take it to be a model that contains a *normal* world that satisfies the premises but none of the conclusions, or we can take it to be a model that contains *any* world that does so. This gives different consequence relations. (The former is the more usual approach.)

1.2.2 Neighborhood modal logic

Another way to vary things is to use, not a relation R on W, but instead a relation N (a *neighborhood* relation) between W and its powerset. That is, where R connected worlds to worlds, N connects worlds to *sets* of worlds. A neighborhood frame is a set W together with such an N. In building models on these frames, $\Box$ is constrained as follows: $\Box A$ is true at a point w iff the set $|A|$ of points in the model where A is true is such that $wN|A|$.

Every relational frame can be perfectly imitated by a neighborhood frame, but not vice versa. As a result, neighborhood frames allow for more freedom in interpretation. For example, $\Box(A \wedge B)$ and $\Box A \wedge \Box B$ are totally independent in general in these systems, and we can make either entail the other, to taste, by adjusting constraints on N. For example, if $\Box A$ is to mean 'A is highly probable', we should expect $\Box(A \wedge B)$ to entail $\Box A \wedge \Box B$, but not vice versa. This is not possible in relational frames, but is easy in neighborhood frames; we simply require that if wNX and $X \subseteq Y$, then wNY. (These are the 'monotonic' modal logics; the weakest is called M.)

1.2.3 Changing the underlying logic

Finally, we can include nonclassical worlds in our frames. This can be done either by swapping classical logic out entirely for something else (say, K3), or else we can repeat the kinds-of-worlds approach we looked at in the nonnormal case, dividing up worlds into chunks and imposing different requirements on the members of different chunks.

This can help with some of the remaining difficulties involving conditionals. For example, even with the strict conditional, we often have $(A \wedge \neg A) \Rightarrow B$, or $B \Rightarrow (A \vee \neg A)$. Changing the underlying logic of $\wedge, \vee$, and $\neg$ to allow for true contradictions or failures of excluded middle is a quick way to get rid of these.

Suppose we have an 'inner core' of worlds that obey classical logic, that worlds in general only obey FDE, and that a counterexample needs a *classical* world to do the counterexampling. Then, even with $\Box$ as in normal systems, we have full classical logic for the nonmodal language, but can avoid $\vdash (A \wedge \neg A) \Rightarrow B$ and the like. No classical world can fail to satisfy $(A \wedge \neg A) \rightarrow B$, but some may see FDE worlds that do.

1.2.4 Extending to first-order

For modal logics, even normal modal logics, the extension to the first-order case is not straightforward. Much of the fuss comes from the 'Barcan formula' $\forall x \Box A \rightarrow \Box \forall x A$. This seems to rule out possibilia: things that might have existed (and failed to satisfy A) but don't. Indeed, the most familiar models for logics in which the Barcan formula fails have a different domain of objects for each world; while the most familiar models for logics in which it holds have a single domain for the whole model.

Another source of fuss comes from *cross-world identity*. Some philosophers have thought that, owing to the nature of possible worlds, no individual *can* exist at more than one. Such philosophers understand 'Richard Nixon might not have been elected' not as claiming that there is a world containing *Nixon himself* in which he wasn't elected, but instead as claiming that there is a world containing *a counterpart of Nixon* in which that counterpart wasn't elected. Of course the details of this, both philosophical and technical, will depend on the nature of the counterpart relation. Can something have no counterpart at a world? More than one? Can something have a counterpart other than itself at the world where it does exist?

A third source of fuss is about equality and distinctness. Does $a = b$ entail $\Box(a = b)$? There is a quick argument that it does: if $a = b$, then a has every property b has. But surely $\Box(a = a)$; a has the property that a is necessarily equal to it. So b must have this property as well—thus, $\Box(a = b)$. This way of putting it seems to involve a modalized lambda-calculus; part of the question, then, is what principles such a calculus should obey.

1.2.5 The Routley star

We can use the resources modal models provide to think about familiar connectives as well as new ones. Suppose we have a set W of worlds with a unary operation $\star$ that takes one world to another; and suppose as well that $w^{\star\star} = w$, for every $w \in W$. Each world, then, comes paired with a 'star mate'. Let things like this be our frames.

To build a model on such a frame, allow atomic sentences to be arbitrary, and assign $\wedge$ and $\vee$ their usual world-by-world truth conditions, but handle negation differently. Say that $\neg A$ is true at a world w iff A is *not* true at $w^\star$. So instead of checking whether A is true at the world we care about, we check it at A's star mate.

If we define validity as truth-preservation at every point in every model in the usual way, this determines the logic FDE. Here, we do not consider it as four-valued, but instead as built from models on top of frames in the usual modal way.

We can recover K3 and LP this way too. Within a model, let $w \sqsubseteq w'$ iff every atom true at w is also true at w'. It turns out that when $w \sqsubseteq w'$, every *sentence* true at w is also true at w'. Now we can think about four kinds of worlds.

- A world w is *classical* iff $w \sqsubseteq w^\star$ and $w^\star \sqsubseteq w$,
- A world w is *complete* iff $w^\star \sqsubseteq w$, and
- A world w is *consistent* iff $w \sqsubseteq w^\star$.

If we require a counterexample to use a *complete* world, we get LP; if we require a *consistent* world, we get K3; and if we require a *classical* world, we get classical logic. This is a bit of a roundabout way of getting these logics, but it turns out to drastically simplify the frame-based models for full relevant logics, of which FDE is a part. (This is the so-called 'Australian plan' for relevant logics; the four-valued approach is part of the 'American plan'. There is also a 'Scottish plan', which is more proof-theoretic.)

2 Applications

2.1 Vagueness

Some approaches to vagueness are based on *supervaluations* or *subvaluations*. The idea is this: vagueness is construed as some kind of indeterminacy between precise languages. The vagueness of 'blue', on such an approach, is a matter of it not being settled *which* precise extension (from a certain range) 'blue' has. Modal logic is not far away.

We can build models in which the worlds represent precise ways our language could have worked. In one world, *this* is the set of blue things, while in another, *that* is. On a supervaluational approach, to be true is to be true in *every* precisification of the language; this is nicely handled by $\Box$. On a subvaluational approach, to be true is to be true in *some* precisification of the language; this is nicely handled by $\Diamond$.

If we understand validity as truth-preservation, then these theories suggest yet further consequence relations for us to consider. Say that a model is a *supervaluational* counterexample to $\Gamma : \Delta$ iff it contains a world at which $\Box A$ is true for every $A \in \Gamma$, but $\Box B$ is not true for any $B \in \Delta$; mutatis mutandis for *subvaluational* (with $\Diamond$). Depending on what our models are like, this yields a rich variety of consequence relations.

For example, suppose we are working with normal models for S5. Then we have a subvaluational counterexample to $A, B : A \wedge B$; we can easily get a world at which $\Diamond A$ and $\Diamond B$ are both true, but where $\Diamond(A \wedge B)$ fails. Similarly, we have a supervaluational counterexample to $A \vee B : A, B$. Interestingly, there are strong connections to classical logic in some restricted cases. Suppose we are only interested in arguments that don't themselves contain $\Box$ or $\Diamond$. Then single-conclusion supervaluational consequence matches single-conclusion classical consequence, even though the set-set formulations are different. Similarly, single-premise subvaluational consequence matches single-premise classical consequence, even though the set-set formulations are different.

$$\text{Id: } \frac{}{A : A} \qquad \bot\text{: } \frac{}{\bot : A} \qquad \text{K: } \frac{\Gamma : A}{\Gamma, \Gamma' : A}$$

$$\wedge\text{L: } \frac{\Gamma, A_i : B}{\Gamma, A_0 \wedge A_1 : B} \qquad \wedge\text{R: } \frac{\Gamma : A \qquad \Gamma : B}{\Gamma : A \wedge B} \qquad \vee\text{L: } \frac{\Gamma, A : C \qquad \Gamma, B : C}{\Gamma, A \vee B : C} \qquad \vee\text{R: } \frac{\Gamma : A_i}{\Gamma : A_0 \vee A_1}$$

$$\to\text{L: } \frac{\Gamma : A \qquad \Gamma', B : C}{\Gamma, \Gamma', A \to B : C} \qquad \to\text{R: } \frac{\Gamma, A : B}{\Gamma : A \to B}$$

Figure 1: Sequents for intuitionist logic

2.2 Intuitionist logic

Intuitionist logic is often motivated by thinking directly of proofs; the idea is to assert A only when a proof of A is in hand. Although it is often more convenient to work indirectly, we can think directly in terms of proofs:

- A proof of $A \wedge B$ is a pair $\langle \pi_0, \pi_1 \rangle$ of a proof π_0 of A and a proof π_1 of B.
- A proof of $A \vee B$ is a pair $\langle \pi, i \rangle$ such that $i \in \{0, 1\}$ and: if $i = 0$, then π is a proof of A; and if $i = 1$, then π is a proof of B.
- A proof of $A \to B$ is a method for transforming proofs of A into proofs of B.
- There are no proofs of $\bot$.

In this setting, $\neg A$ is typically defined as $A \to \bot$.

2.2.1 Sequents

Intuitionist logic is almost always considered in a single-conclusion formulation. With single conclusion sequents, we can present it as in Figure 2.2.1.

Intuitionist logic has the *disjunction property*: $\vdash A \vee B$ only if $\vdash A$ or $\vdash B$. As a result, excluded middle is not a theorem of intuitionist logic. However, it is not much like K3. For example, $\neg\neg A \not\vdash A$ in intuitionist logic; and $\neg(A \wedge B) \not\vdash \neg A \vee \neg B$. An easy way to remember how intuitionist negation works: it behaves like 'never'.

2.2.2 Models

In fact, this analogy with 'never' suggests a way to build models for intuitionist logic (so long as you're ok with branching futures). Let a frame be a set W together with a partial order $\leq$: a reflexive, antisymmetric, transitive relation on W. Now a model assigns value 1 or 0 to each sentence at each point w in a frame, subject to the following restrictions:

- If $w \leq w'$, then the atomic sentences that get 1 at w are a subset of those that get 1 at w',
- $A \wedge B$ gets 1 at w iff A and B both do,
- $A \vee B$ gets 1 at w iff either A or B does,
- $A \to B$ gets 1 at w iff: for every w' such that $w \leq w'$, if A gets 1 at w', then so does B, and
- $\bot$ never gets 1 anywhere.

In fact, the first condition (called *persistence*) can be shown to hold for all sentences. Now, let a model be a counterexample to $\Gamma : A$ iff it contains a point w such that B gets 1 at w for every $B \in \Gamma$, but A gets 0 at w. This determines intuitionist logic as well. These models are sometimes interpreted as representing the *development of knowledge*. But note that this must be *absolutely certain* knowledge, or else persistence makes no sense.

$$\text{Id: } \frac{}{A : A} \qquad \text{KL: } \frac{\Gamma : \Delta}{A, \Gamma : \Delta} \qquad \text{KR: } \frac{\Gamma : \Delta}{\Gamma : \Delta, A} \qquad \text{Cut: } \frac{\Gamma : \Delta, A \qquad A, \Gamma' : \Delta'}{\Gamma, \Gamma' : \Delta, \Delta'}$$

$$\text{WL: } \frac{A, A, \Gamma : \Delta}{A, \Gamma : \Delta} \qquad \text{WR: } \frac{\Gamma : \Delta, A, A}{\Gamma : \Delta, A} \qquad \text{CL: } \frac{\Gamma, A, B, \Gamma' : \Delta}{\Gamma, B, A, \Gamma' : \Delta} \qquad \text{CR: } \frac{\Gamma : \Delta, A, B, \Delta'}{\Gamma : \Delta, B, A, \Delta'}$$

$$\neg\text{L: } \frac{\Gamma : \Delta, A}{\neg A, \Gamma : \Delta} \qquad \neg\text{R: } \frac{A, \Gamma : \Delta}{\Gamma : \Delta, \neg A} \qquad \to\text{L: } \frac{\Gamma : \Delta, A \qquad B, \Gamma' : \Delta'}{A \to B, \Gamma, \Gamma' : \Delta, \Delta'} \qquad \to\text{R: } \frac{A, \Gamma : \Delta, B}{\Gamma : \Delta, A \to B}$$

$$\wedge\text{L: } \frac{A_i, \Gamma : \Delta}{A_0 \wedge A_1, \Gamma : \Delta} \qquad \wedge\text{R: } \frac{\Gamma : \Delta, A \qquad \Gamma : \Delta, B}{\Gamma : \Delta, A \wedge B} \qquad \vee\text{L: } \frac{A, \Gamma : \Delta \qquad B, \Gamma : \Delta}{A \vee B, \Gamma : \Delta} \qquad \vee\text{R: } \frac{\Gamma : \Delta, A_i}{\Gamma : \Delta, A_0 \vee A_1}$$

Figure 1: Gentzen's LK

1 Substructural logics

So far, almost everything we've seen has been an RMT relation, either single-conclusion or set-set. But the basic tools we have can adapt to other settings as well.

1.1 Gentzen's LK

Figure 1.1 contains (the propositional fragment of) Gentzen's original sequent calculus LK for classical logic (the names for the rules are not his). LK uses *lists* of premises and conclusions, rather than sets, as we've been doing. A list is unlike a set in two important ways: the *number of times* something occurs in a list matters, and the *order* in which things occur in a list matters. So while the set $\{A, A, B, A, B, B, A\}$ is *the same set* as the set $\{A, B\}$, the list A, A, B, A, B, B, A is a completely different list from the list A, B. The comma that appears in the rules is now *concatenation*, rather than union.

Note the assortment of rules above the separator: these are the *structural rules*; they do not pay attention to any particular vocabulary, but merely tell us how we can manipulate the structure of an argument. Given these structural rules, it would not affect derivability at all if Gentzen had simply used sets instead: the C rules ensure that order doesn't matter, and the W rules and their converses (which are special cases of the K rules) ensure that number of occurrences doesn't matter. Gentzen's own reasons for using lists rather than sets was to ensure that he had very tight control indeed over the structures of his derivations: his 'Hauptsatz' (main theorem) was that the rule of Cut is *eliminable* from LK: everything derivable with it is also derivable without it, and his proof depends on having such fine-grained control.

1.2 Ditching structural rules

So we can remove the structural rule of Cut from LK without affecting its internal consequence relation. But none of the other structural rules are like this; if we start removing them we arrive at new logics. These are *substructural logics*. For example, if you try to derive the classically-valid $A \wedge (B \vee C) : (A \wedge B) \vee (A \wedge C)$, you will find that you need all of K, C, and W to get the job done.

If we remove any of the rules K, C, or W from LK, we will no longer be able to derive distribution of $\wedge$ over $\vee$. In fact, end up at a weak system indeed, and one that is not much studied, as far as I know. (It's a relative of the 'classical associative Lambek calculus', but with an odd choice of vocabulary.)

1.2.1 MAAL

A much stronger substructural logic, and one that reveals some of the texture in play here, is multiplicative-additive affine logic, or MAAL. Sequents for MAAL can get away with using *multisets*: these are like lists in that number of occurrences matters, but like sets in that order does not matter. There is no need for C rules if we are using multisets; their effect is built in already. Affine logic also includes the K rules (this is what distinguishes it from

$$\otimes\text{L: } \frac{A, B, \Gamma : \Delta}{A \otimes B, \Gamma : \Delta} \qquad \otimes\text{R: } \frac{\Gamma : \Delta, A \qquad \Gamma' : \Delta', B}{\Gamma, \Gamma' : \Delta, \Delta', A \otimes B}$$

$$\oplus\text{L: } \frac{A, \Gamma : \Delta \qquad B, \Gamma' : \Delta'}{A \oplus B, \Gamma, \Gamma' : \Delta, \Delta'} \qquad \oplus\text{R: } \frac{\Gamma : \Delta, A, B}{\Gamma : \Delta, A \oplus B}$$

$$\mapsto\text{L: } \frac{\Gamma : \Delta, A \qquad B, \Gamma : \Delta}{A \mapsto B, \Gamma : \Delta} \qquad \mapsto\text{R: } \frac{\Gamma, A : \Delta}{\Gamma : \Delta, A \mapsto B} \qquad \mapsto\text{R: } \frac{\Gamma : \Delta, B}{\Gamma : \Delta, A \mapsto B}$$

Figure 2: Extra rules for MAAL

linear logic). But—very crucially—it does *not* include the W rules. Contraction is not guaranteed to preserve MAAL-validity.

MAAL is 'multiplicative-additive' because it includes *two* variants of each binary connective. We can take the operational rules from LK on board as they are (reinterpreting them to involve multisets rather than lists): $\wedge$ and $\vee$ are the *additive* conjunction and disjunction, and $\rightarrow$ is the *multiplicative* conditional. To these, we add the rules in Figure 2, which given the *multiplicative* conjunction and disjunction, and the *additive* conditional. (Girard's original linear logic notation differs importantly!)

In the presence of K and W, there is no difference between $\wedge$ and $\otimes$, or between $\vee$ and $\oplus$, or between $\rightarrow$ and $\mapsto$. But in the absence of either, important differences emerge. For example, suppose we want to ask about explosion in MAAL. $A \otimes \neg A \vdash$; and $A \wedge \neg A, A \wedge \neg A \vdash$; but $A \wedge \neg A \not\vdash$; this is a case where contraction fails. (Similarly if we ask about excluded middle.)

One usual motivation for thinking about affine logic is *resource sensitivity*. We want to see if the premises give us enough to get to our conclusions—where if we need to appeal to something twice, we'd better have it present twice. In this setting, $A \wedge B$ is something that can be used either as an A or as a B, as we choose, while $A \otimes B$ can deliver *both* an A and a B. To get both an A and a B from $A \wedge B$, we would have to use it twice. (Multiplicity on the right is a bit harder to interpret this way; there are also 'intuitionist' affine logics that only have single conclusions and so avoid the question.)

But MAAL has another nice twist to it: it can accommodate a naive truth predicate without blowing up. Every derivation of trouble from a liar or curry sentence requires *double use*: showing something once and appealing to it twice. By banning contraction, such troubles are blocked.

1.2.2 Relevant logic

Relevant logic began from the idea that the premises of a valid argument should have something to do with its conclusion. This is sometimes given as a *variable-sharing* criterion: if $\Gamma \vdash \Delta$, then Γ and Δ must have at least one propositional variable in common. Classical logic (and K3, LP, FDRM, intuitionist, Łukasiewicz, affine, etc) do not respect this condition. FDE, however, does; it is a relevant logic.

In fact, FDE is the conditional-free fragment of the relevant logic E (and of all the usual relevant logics). Relevant logics usually contain connectives $\wedge$, $\vee$, $\neg$, and a conditional $\rightarrow$, and are usually defined axiomatically. One way to get an idea of these systems is to read $\rightarrow$ as a kind of absolute sufficiency connective, expressing a very strong connection between its antecedent and its consequent. For example, here is an axiomatization of the relevant logic B (a particularly weak relevant logic):

- $A \rightarrow A$
- $A \rightarrow A \vee B$
- $B \rightarrow A \vee B$
- $A \wedge B \rightarrow A$
- $A \wedge B \rightarrow B$
- $(A \rightarrow B) \wedge (A \rightarrow C) \rightarrow (A \rightarrow B \wedge C)$
- $(A \rightarrow C) \wedge (B \rightarrow C) \rightarrow (A \vee B \rightarrow C)$

- $A \wedge (B \vee C) \to (A \wedge B) \vee (A \wedge C)$
- $A \to \neg\neg A$
- $\neg\neg A \to A$
- From A and $A \to B$, infer B
- From A and B, infer $A \wedge B$
- From $A \to B$ and $C \to D$, infer $(B \to C) \to (A \to D)$
- From $A \to B$, infer $\neg B \to \neg A$

Key principles that are *not* theorems of DW, but occur in stronger systems:

- contraction: $(A \to (A \to B)) \to (A \to B)$
- prefixing: $(A \to B) \to ((C \to A) \to (C \to B))$
- assertion, $A \to ((A \to B) \to B)$
- conjunctive syllogism, $(A \to B) \wedge (B \to C) \to (A \to C)$
- 'pseudo modus ponens', $(A \wedge (A \to B)) \to B$

There is also another connective sometimes added to relevant logics, called 'fusion', and added by the rule: From $(A \circ B) \to C$, infer $A \to (B \to C)$, and vice versa.

1.2.3 Drawing the threads together

In fact, there are close (but not perfect) connections between the Gentzen-inspired substructural logics and the relevant logics; it is the working out of these connections that provides much of the interest of substructural logics. Consider the following derivation:

$$\dfrac{\dfrac{\dfrac{A : A \qquad B : B}{A \to B, A : B}\ {\to}\text{L}}{A, A \to B : B}\ \text{CL}}{A : (A \to B) \to B}\ {\to}\text{R}$$

This derivation shows a connection between the structural rule CL and the principle $A \to ((A \to B) \to B)$, which B lacks. There is a sense, then, in which B does not allow for exchange. Other ways to see this: $(A \to (B \to C)) \to (B \to (A \to C))$ is not a B theorem; neither, if we add fusion, is $((A \circ B) \to C) \to ((B \circ A) \to C)$. (I've cheated a bit here, swapping out Gentzen's original $\to$R rule for one that interacts with order differently; this is the way that turns out to be more useful.)

However, even without any of the structural rules we've considered, we can still derive prefixing:

$$\dfrac{\dfrac{\dfrac{\dfrac{C : C \qquad A : A}{C \to A, C : A}\ {\to}\text{L} \qquad B : B}{A \to B, C \to A, C : B}\ {\to}\text{L}}{A \to B, C \to A : C \to B}\ {\to}\text{R}}{A \to B : (C \to A) \to (C \to B)}\ {\to}\text{R}$$

So if we want to get all the way down to B, we need to weaken yet again. Let a *structure* be inductively defined as follows: any formula is a structure, and if X and Y are both structures, then (X, Y) is a structure too. Our remaining rules, both structural and operational, need to be tweaked to accomodate this shift. But it turns out that once this is accomplished, it is *associativity* of the structural comma that gives rise to prefixing, just as it is exchange that gives rise to assertion. There is kind of a 'hidden step' of reassociating in the above derivation after the second $\to$L and before the $\to$Rs.

This is not quite enough to get us relevant logics directly; as we've seen, distribution of $\wedge$ over $\vee$ depends on structural rules, and many relevant logics have $\to$s that require very weak structural behavior indeed, while still requiring distribution. The usual solution to this puzzle involves having *two* structural connectives, one with very weak structural behavior, linked to $\to$ and $\circ$; and one obeying all the usual structural rules, linked to $\wedge$ and $\vee$.

1.3 Routley-Meyer models

One way to give a model theory for substructural logics is based on *Routley-Meyer models*. These are frame-based models like the ones we saw for modal logic, but with some modifications. The originals had some bells & whistles that we can get rid of for some purposes, so let's work with 'simplified' models.

A *simplified frame* is built on a set W with a distinguished subset N. (Think of worlds with distinguished 'normal' worlds, although this doesn't mean the same as in normal modal logic.) To these, we need to add a one-place operation $\star$ such that $w = w^{\star\star}$ for all w to handle negation (this is the Routley star from yesterday), and a three-place relation R to handle the conditional.

Every world assigns value 1 or value 0 to every sentence. Atomic sentences, as usual, are unconstrained. Conjunction and disjunction work world-by-world in a classical way. Negation works via the Routley star: $v(\neg A, w) = 1$ iff $v(A, w^\star) = 0$.

The conditional works differently at normal worlds and other worlds. For all normal worlds w, $v(A \to B, w) = 1$ iff *every* world x in the model such that $v(A, x) = 1$ is also such that $v(B, x) = 1$. At the other worlds, the relation R is involved: for these worlds w, $v(A \to B, w) = 1$ iff for all y, z such that $Rwyz$, if $v(A, y) = 1$ then $v(B, z) = 1$.

A countermodel to an argument is then a model containing a *normal* world at which the premises all get value 1 and the conclusions value 0. As relevant logics are usually presented axiomatically, we usually only care about the special case with one conclusion and no premises.

2 Paradox again

2.1 Truth again

Substructural logics turn out to be useful for addressing paradoxes of various sorts. Suppose that we add the following rules for truth to Gentzen's LK:

$$T\langle\rangle\text{L:}\quad \frac{A, \Gamma : \Delta}{T\langle A\rangle, \Gamma : \Delta} \qquad T\langle\rangle\text{R:}\quad \frac{\Gamma : \Delta, A}{\Gamma : \Delta, T\langle A\rangle}$$

Then we can get ourselves in trouble with a sentence λ that is $\neg T\langle\lambda\rangle$ as follows:

$$\dfrac{\dfrac{\dfrac{\dfrac{\lambda : \lambda}{T\langle\lambda\rangle : \lambda}\ T\langle\rangle\text{L}}{: \lambda, \neg T\langle\lambda\rangle}\ \neg\text{R}}{: \lambda}\ \text{WR} \qquad \dfrac{\dfrac{\dfrac{\lambda : \lambda}{\lambda : T\langle\lambda\rangle}\ T\langle\rangle\text{R}}{\neg T\langle\lambda\rangle, \lambda :}\ \neg\text{L}}{\lambda :}\ \text{WL}}{:}\ \text{Cut}$$

But, as it turns out, without contraction, no trouble will come from the truth rules. A naive truth predicate can be added to MAAL without blowing the whole thing up. (We don't yet know whether this is compatible with keeping to a standard model of arithmetic; in Łukasiewicz logic, which is closely related (but distinct!) it is not.)

2.2 Sets

Just as we might expect truth to behave in a transparent way (at least until we run into the paradoxes), so might we expect the same of sets. That is, it's natural to think that, for any open formula $A(x)$ and any thing a: the claim that a is a member of the set of things x such that $A(x)$ is equivalent to the claim $A(a)$. That is, you might expect $a \in \{x : A(x)\}$ to be equivalent to $A(a)$. This is sometimes called *naive comprehension*. Classically, this cannot work: consider for example the Russell set $\{x : x \notin x\}$. The claim that it is a member of itself is equivalent to the claim that it is not.

So far, this is not so different from the liar paradox. (Similarly, the set $\{x : x \in x \to \bot\}$ works just like the curry paradox.) But there is a twist: there is a second intuitive principle that a naive set theory ought to obey. This is the principle of *extensionality*: if sets x and y have the same members, they ought to be identical. This is something genuinely different from what happens in the truth case, and it can cause plenty of trouble on its own.

For example, affine logic has no problem accomodating naive comprehension, but it *cannot* also handle extensionality; everything becomes provable when extensionality is added. On the other hand, a number of relevant logics, including B, are known safe for the combination of naive comprehension and extensionality.

1 Monotonic logics

Today, we will return to a set-based conception of premises and conclusions; order and number of occurrences wll once again not matter. Our target now is *monotonic logic*.

A set-set *monotonic logic* is a consequence relation $\vdash$ such that whenever $\Gamma \vdash \Delta$, then $\Gamma, \Gamma' \vdash \Delta, \Delta'$. In other words, we retain *only* the 'M' from our 'RMT' formulation. In some sense, these are surely substructural logics: they allow for failures of the usual structural properties of reflexivity and transitivity. But they are not often studied, even in studies of substructural logics.

1.1 STT

1.1.1 STT on strong kleene models

Think in strong kleene models for a moment, and let an *ST-counterexample* to an argument $\Gamma : \Delta$ be a model that assigns 1 to everything in Γ and 0 to everything in Δ. It turns out that this determines the usual classical set-set consequence relation (which is, of course, fully RMT).

But we have seen that, via the Kripke construction, we can add a naive truth predicate to strong kleene models without getting ourselves in any trouble. Suppose we do this: our models are now just the fixed points for truth built on strong kleene models, and we use the ST understanding of counterexample. This determines the logic STT.

We have only *restricted* our space of models, so we must still validate every classically-valid argument. But we now have a truth predicate T in play such that $T\langle A\rangle$ is always fully equivalent to A, and we can allow for the formation of liars, curries, etc. Because the Kripke construction works from any starting point, we will not validate any new arguments that do not involve the truth predicate; we can thus be sure that the result hasn't exploded on us.

So it turns out to be possible after all to add a naive truth predicate to classical logic! Well, yes and no. While STT does validate every classically-valid argument, it is *not* RMT; in particular, it is not *transitive*. Where λ is a liar sentence, it must always take value $\frac{1}{2}$ in fixed points, since it must be equivalent to its own negation, and only the value $\frac{1}{2}$ allows for that. So in STT, $p \vdash \lambda$ and $\lambda \vdash q$; but still $p \not\vdash q$. We have a failure even of the simplest transitivity property.

1.1.2 STT via proofs

Another way to get to STT is to think proof-theoretically. For example, the sequent calculus for classical logic we considered on the first day includes no rule of *Cut*. Cut is the sequent rule with the tightest relation to transitivity. It comes in a number of forms; the simplest is probably this one:

$$\text{Cut:} \quad \frac{\Gamma : \Delta, A \qquad A, \Gamma : \Delta}{\Gamma : \Delta}$$

Cut is *admissible* in the sequent calculus we looked at (since it preserves classical validity); adding it would introduce new derivations, to be sure, but no new sequents would be derivable. The new derivations would lead only to destinations we already had other ways to. Since it would make no difference, it is often simpler not to add it, although from a lot of perspectives the choice is unimportant. But if we add the rules for truth we considered yesterday, something interesting happens.

$$T\langle\rangle\text{L:} \quad \frac{A, \Gamma : \Delta}{T\langle A\rangle, \Gamma : \Delta} \qquad T\langle\rangle\text{R:} \quad \frac{\Gamma : \Delta, A}{\Gamma : \Delta, T\langle A\rangle}$$

We still have all the derivations we have before; we have only *added* rules. So every classical validity remains derivable. But the new truth rules are enough to allow us to swap A for $T\langle A\rangle$ willy-nilly, as we please, even in the presence of paradoxical sentences. And we do not derive anything unwanted, like $p : q$ or its ilk.

But Cut is no longer admissible; if we were to impose cut over this extended system, anything would be derivable. (Similarly, if we add the truth rules to a classical sequent calculus that *does* include a rule of Cut, anything will turn out derivable.) In fact, this system is STT again.

1.1.3 Understanding nontransitive consequence

For lots of intended applications of consequence relations, STT seems to be a complete nonstarter. For example, if consequence is meant to tell us about truth-preservation or warrant-preservation or commitment-preservation, then we should probably expect it to be transitive, since *preservation* is.

One way to understand nontransitive relations is to see the two sides of an argument as telling us about *different* statuses. Suppose, for example, that an argument is valid iff it is out of bounds to assert all its premises while denying all its conclusions. We can explore different stories about what 'out of bounds' might amount to. If it turns out to be out of bounds to assert and deny the very same thing, that's reflexivity. If, having gone out of bounds, we can't come back in by making further assertions and denials, that's monotonicity. And if, so long as we're in bounds, there is always some stand on A we can take while remaining in bounds, that's Cut.

Paradoxical sentences provide natural objections to the first and third conditions: maybe they are the sort of thing it's ok to both assert and deny, or maybe they're the sort of thing it's not ok to take a stand on at all. STT is a picture you might arrive at by taking the latter choice.

2 Abstract valuations

2.1 RMT

Every RMT set-set relation has a two-valued model theory, although these are often not the most informative or interesting. The idea is this: let a *bivaluation* be a function from the language to $\{1, 0\}$. These are subject to *no* restrictions; every such function is a bivaluation. A bivaluation is a counterexample to an argument $\Gamma : \Delta$ iff it assigns 1 to everything in Γ and 0 to everything in Δ.

Now, given a set-set consequence relation $\vdash$, let $B(\vdash)$ be the set of bivaluations determined as follows: a bivaluation b is in $B(\vdash)$ iff it is not a counterexample to any argument that is $\vdash$-valid. Similarly, given a set V of bivaluations, let $C(V)$ be the consequence relation determined as follows: $\Gamma : \Delta$ is $C(V)$-valid iff no bivaluation in V is a counterexample to $\Gamma : \Delta$.

$C(V)$ is always RMT. But *every* RMT set-set relation can be gotten at this way. In fact, if $\vdash$ is an RMT set-set relation, then $C(B(\vdash)) = \vdash$. This is the 'abstract soundness and completeness theorem'. So every RMT set-set consequence relation has a two-valued model theory: the models are the bivaluations in $B(\vdash)$, and counterexamples are understood in the usual way.

2.2 M only

This will not work as is for all monotonic logics, since $C(V)$, as defined above, is always RMT. But there is a slight generalization that will work: all monotonic logics have a *four*-valued model theory. Let a *tetravaluation* be a function from the language to $\{\langle 1, 1\rangle, \langle 1, 0\rangle, \langle 0, 1\rangle, \langle 0, 0\rangle\}$, and let a tetravaluation be a counterexample to an argument $\Gamma : \Delta$ iff everything in Γ gets a value whose *first* coordinate is 1 and everything in Δ gets a value whose *second* coordinate is 0. Given this understanding of counterexample, bivaluations are special case where only $\langle 1, 1\rangle$ and $\langle 0, 0\rangle$ are used.

Now, given a set-set consequence relation $\vdash$, define $T(\vdash)$ as the set of tetravaluations that don't counterexample anything $\vdash$-valid; and given a set V of tetravaluations, let $C(V)$ be the consequence relation that declares an argument valid iff it has no counterexamples in V.

$C(V)$ is always monotonic. And indeed, every monotonic set-set relation can be gotten at in this way; if $\vdash$ is a monotonic set-set consequence relation, then $C(T(\vdash)) = \vdash$. Moreover, if a set V of tetravaluations completely avoids the value $\langle 1, 0\rangle$, then $C(V)$ is reflexive; and if is completely avoids $\langle 0, 1\rangle$, then $C(V)$ is transitive (in the strong way required for RMT).

General

[1] J. Michael Dunn and Gary M. Hardegree. *Algebraic Methods in Philosophical Logic*. Oxford University Press, Oxford, 2001.

[2] Susan Haack. *Philosophy of Logics*. Cambridge University Press, Cambridge, 1978.

[3] Lloyd Humberstone. *The Connectives*. MIT Press, Cambridge, Massachusetts, 2012.

[4] Sara Negri and Jan von Plato. *Structural Proof Theory*. Cambridge University Press, Cambridge, 2001.

[5] Graham Priest. *An Introduction to Non-Classical Logic: From If to Is*. Cambridge University Press, Cambridge, 2nd edition, 2008.

[6] Helena Rasiowa. *An Algebraic Approach to Non-Classical Logics*, volume 78 of *Studies in Logic and the Foundations of Mathematics*. North-Holland, Amsterdam, 1974.

[7] Yann-Joachim Ringard. Mustard watches: An integrated approach to time and food, 1990. Équipe de Logique Mathématique, Paris VII, Preprint.

[8] D. J. Shoesmith and T. J. Smiley. *Multiple-conclusion Logic*. Cambridge University Press, Cambridge, 1978.

Many-valued

[9] Jc Beall and Bas C. van Fraassen. *Possibilities and Paradox: An Introduction to Modal and Many-valued Logic*. Oxford University Press, Oxford, 2003.

[10] Petr Hájek, Jeff Paris, and John Shepherdson. The liar paradox and fuzzy logic. *Journal of Symbolic Logic*, 65(1):339–346, 2000.

[11] Saul Kripke. Outline of a theory of truth. *Journal of Philosophy*, 72(19):690–716, 1975.

[12] Hannes Leitgeb. Truth and the liar in De Morgan-valued models. *Notre Dame Journal of Formal Logic*, 40(4):496–514, 1999.

[13] Graham Priest. *In Contradiction*. Oxford University Press, Oxford, 2006.

[14] Nicholas J. J. Smith. *Vagueness and Degrees of Truth*. Oxford University Press, Oxford, 2008.

Modal

[15] Roberta Ballarin. Modern origins of modal logic. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Winter 2010 edition, 2010.

[16] Patrick Blackburn, Maarten de Rijke, and Yde Venema. *Modal Logic*. Cambridge University Press, Cambridge, 2001.

[17] Aldo Bressan. *A General Interpreted Modal Calculus*. Yale University Press, New Haven, 1972.

[18] Brian F. Chellas. *Modal Logic: an Introduction*. Cambridge University Press, Cambridge, 1980.

[19] Robert Goldblatt. *Quantifiers, Propositions, and Identity: Admissible Semantics for Quantified Modal and Substructural Logics*. Number 38 in Lecture Notes in Logic. Cambridge University Press, Cambridge, 2011.

[20] Helle Hvid Hansen. Monotonic modal logics. Master's thesis, University of Amsterdam, Amsterdam, 2003.

[21] G. E. Hughes and M. J. Cresswell. *A New Introduction to Modal Logic*. Routledge, New York, 1996.

[22] Graham Priest. *Towards Non-Being: The Logic and Metaphysics of Intentionality*. Oxford University Press, Oxford, 2005.

Substructural

[23] Katalin Bimbó and J. Michael Dunn. *Generalized Galois Logics: Relational Semantics of Nonclassical Logical Calculi*. Number 188 in CSLI Lecture Notes. CSLI Publications, 2008.

[24] Ross T. Brady. *Universal Logic*. CSLI Publications, Stanford, California, 2006.

[25] Ole Hjortland. *The Structure of Logical Consequence: Proof-theoretic Conceptions*. PhD thesis, University of St. Andrews, 2009.

[26] Francesco Paoli. *Substructural Logics: A Primer*. Kluwer Academic Publishing, Dordrecht, The Netherlands, 2002.

[27] Greg Restall. *An Introduction to Substructural Logics*. Routledge, London, 2000.

[28] Richard Routley, Robert K. Meyer, Val Plumwood, and Ross T. Brady. *Relevant Logics and their Rivals 1*. Ridgeview, Atascadero, California, 1982.

[29] Lionel Shapiro. Deflating logical consequence. *Philosophical Quarterly*, 61(243):320–342, 2011.

[30] Zach Weber. Transfinite numbers in paraconsistent set theory. *Review of Symbolic Logic*, 3(1):71–92, 2010.

[31] Elia Zardini. Truth without contra(di)ction. *Review of Symbolic Logic*, 4(4):498–535, 2011.

Monotonic

[32] Pablo Cobreros, Paul Égré, David Ripley, and Robert van Rooij. Tolerant, classical, strict. *Journal of Philosophical Logic*, 41(2):347–385, 2012.

[33] Lloyd Humberstone. Heterogeneous logic. *Erkenntnis*, 29:395–435, 1988.

[34] David Ripley. Paradoxes and failures of cut. *Australasian Journal of Philosophy*, 91(1):139–164, 2013.

[35] Peter Schroeder-Heister. The categorical and the hypothetical: A critique of some fundamental assumptions of standard semantics. *Synthese*, 187(3):925–942, 2012.

[36] Alan Weir. Naïve truth and sophisticated logic. In Jc Beall and Bradley Armour-Garb, editors, *Deflationism and Paradox*, pages 218–249. Oxford University Press, Oxford, 2005.

[37] Elia Zardini. A model of tolerance. *Studia Logica*, 90(3):337–368, 2008.