‘Transitivity’ of Consequence Relations

David Ripley

University of Connecticut, Storrs CT 06268, USA

Abstract. A binary relation R on a set S is *transitive* iff for all $a, b, c \in S$, if aRb and bRc, then aRc. This almost never applies to the relations logicians tend to think of as *consequence relations*; where such relations are relations on a set at all, they are rarely transitive. Yet it is common to hear consequence relations described as ‘transitive’, and to see rules imposed to ensure ‘transitivity’ of these relations. This paper attempts to clarify the situation.

1 Introduction

After briefly substantiating the claims in the abstract, this paper focuses on exploring a number of different properties of consequence relations that have traveled under the name ‘transitivity’, mapping the implications among them. From here forward, I will use ‘transitive’ and ‘transitivity’ very little, and only in their standard relation-theoretic sense. To reiterate: to be transitive, a relation R must be a binary relation on a set S, and it must be such that for any $a, b, c \in S$, if aRb and bRc, then aRc.

Many familiar consequence relations are not relations on a set at all, but instead relate *sets* of formulas (collections of premises) to *single* formulas (conclusions). That is, where $\mathcal{F}$ is the set of formulas under consideration, such a relation is a relation between $\wp(\mathcal{F})$ and $\mathcal{F}$. Following [6], I'll say these relations work in the ‘Set-Form framework’. Such a relation is not the right kind of thing to be transitive. Of course, these relations can, and frequently do, exhibit a number of properties more and less closely connected to transitivity. But I will not explore this here; I mention Set-Form relations to set them aside.

In what follows, I work entirely in the Set-Set framework. In this framework, consequence relations really are binary relations on a single set: the set $\wp(\mathcal{F})$. That is, they relate sets of formulas to sets of formulas. So they are at least the right *kind* of relation to be transitive.

Much research into Set-Set consequence relations (see eg [4, 13, 7, 11, 14, 6]) interprets the members of the set of conclusions as (in some sense) *different possibilities*. On this interpretation, arguments with *fewer* conclusions are stronger than those with more, since they *narrow down* more finely on a result. This is the interpretation I'll focus on in what follows.

These relations, too, are almost never transitive. Consider, for example, the Set-Set consequence relation $\vdash$ determined by classical logic, explored and defended in [7], among other places. This relation relates $\{A \vee B\}$ to $\{A, B\}$, and relates $\{A, B\}$ to $\{A \wedge B\}$, but does not relate $\{A \vee B\}$ to $\{A \wedge B\}$; it is thus

W. van der Hoek et al. (Eds.): LORI 2015, LNCS 9394, pp. 328–340, 2015.
DOI: 10.1007/978-3-662-48561-3_27

not transitive. The reason is nothing particularly to do with classical logic; it is instead to do with how *sets* of formulas are interpreted. As premises, they are meant *conjunctively*: as all available to be drawn on together in establishing conclusions. As conclusions, they are meant *disjunctively*: as jointly exhausting the space where the truth must lie, given the premises. This difference in interpretation prevents linking valid SET-SET arguments together in the simple way guaranteed by transitivity.

2 A Catalog of Linking Properties

In this section, I lay out the assumptions that will frame the paper, and then present a catalog of ten properties that a SET-SET consequence relation might exhibit, all of which, I think, are recognizable as related to what logicians often mean by 'transitivity'. These ten properties form the basis of the paper, which fully maps the implications among arbitrary conjunctions of these properties.

Some notational preliminaries: I use capital Roman letters for formulas, and capital Greek letters (that are not also capital Romans) for sets of formulas. $\mathcal{F}$ is the set of formulas in the language under consideration; each SET-SET consequence relation, then, is a binary relation on $\wp(\mathcal{F})$. (As above, I restrict attention entirely to SET-SET relations.) I abbreviate freely in usual sequent-calculus ways, so, for example, '$\Gamma, A, \Sigma \vdash$' abbreviates '$\Gamma \cup \{A\} \cup \Sigma \vdash \emptyset$'. When I talk of 'partitions' of a set, this should be understood to *include* partitions with an empty entry; for example, $\langle \emptyset, \Sigma \rangle$ *is* a partition of Σ, on this usage.

2.1 Assumptions

I assume in places that the language $\mathcal{F}$ contains infinitely many formulas; its cardinality does not otherwise matter. I make no assumptions about the nature or structure of formulas; $\mathcal{F}$ can be any infinite set.

Consequence relations are often defined as relations that are '*reflexive*, *monotonic*, and *transitive*'. The final condition, of course, is the subject of this paper, so I am certainly not assuming it. Nor will I assume reflexivity, although this turns out not to matter; all the results of the paper remain unchanged with such an assumption in place.[1]

[1] 'Reflexive' here is like 'transitive'; it does not have, in its usual application to SET-SET consequence relations, its usual relation-theoretic sense. In the usual sense, a relation R on a set S is *reflexive* iff for all $x \in S$, xRx. For consequence relations, this would require that for every set Γ of formulas, $\Gamma \vdash \Gamma$. As it happens, this is almost never the case; at the very least, the empty set does not entail itself in any familiar setting. There are two usual things one might mean by 'reflexivity' here: that $\Gamma \vdash \Gamma$ for all *singleton* Γ, or all *nonempty* Γ; these are the assumptions that would not change anything in what follows. To show this, I take care to make sure that all the examples I discuss are reflexive (in both of these senses), and that no proof of any claim depends on reflexivity (in any sense).

I will, however, assume throughout the paper that all consequence relations are *monotonic*: that whenever $\Gamma \vdash \Delta$, then $\Gamma, \Gamma' \vdash \Delta, \Delta'$.[2] This matters a great deal; the situation is very different if this assumption is not imposed, and many of the results to follow would not hold without it.

A consequence relation $\vdash$ is *compact* iff whenever $\Gamma \vdash \Delta$, then there are finite $\Gamma_{\text{fin}} \subseteq \Gamma$ and $\Delta_{\text{fin}} \subseteq \Delta$ such that $\Gamma_{\text{fin}} \vdash \Delta_{\text{fin}}$. In what follows, I will not require compactness in general, but I will keep track of compactness, and show what the effects of requiring compactness are.

2.2 The Catalog

Table 1 gives ten properties that a consequence relation $\vdash$ may or may not exhibit. Each of the properties is a *closure* property: they are all of the form 'if these things stand in the relation, then those things must also stand in the relation'. These should be understood as universally quantified; for example, $\vdash$ has the property KS iff whenever $\Gamma \vdash A$ and $A \vdash \Delta$, then $\Gamma \vdash \Delta$, for all choices of Γ, Δ, and A. The properties to be considered in this paper are the ten in Table 1, and arbitrary conjunctions of these.

Table 1. Linking properties

Name:	If	and	then
S	$C \vdash A$	$A \vdash D$	$C \vdash D$
KS	$\Gamma \vdash A$	$A \vdash \Delta$	$\Gamma \vdash \Delta$
/F	$\Gamma \vdash A$	$A, \Gamma \vdash \Delta$	$\Gamma \vdash \Delta$
F/	$\Gamma \vdash \Delta, A$	$A \vdash \Delta$	$\Gamma \vdash \Delta$
FG	$\Gamma \vdash \Delta, A$	$A, \Gamma \vdash \Delta$	$\Gamma \vdash \Delta$
/C	$\Gamma \vdash A$ for all $A \in \Sigma$	$\Sigma, \Gamma \vdash \Delta$	$\Gamma \vdash \Delta$
C/	$\Gamma \vdash \Delta, \Sigma$	$A \vdash \Delta$ for all $A \in \Sigma$	$\Gamma \vdash \Delta$
/C$^+$	$\Gamma \vdash \Delta, A$ for all $A \in \Sigma$	$\Sigma, \Gamma \vdash \Delta$	$\Gamma \vdash \Delta$
C$^+$/	$\Gamma \vdash \Delta, \Sigma$	$A, \Gamma \vdash \Delta$ for all $A \in \Sigma$	$\Gamma \vdash \Delta$
CG	$\Sigma^+, \Gamma \vdash \Delta, \Sigma^-$ for all partitions $\langle \Sigma^+, \Sigma^- \rangle$ of Σ		$\Gamma \vdash \Delta$

Each allows valid arguments to be *linked* in a specific way; in the antecedent of these properties, the formula A and/or the set Σ of formulas figures among the conclusions of the left conjunct and the premises of the right conjunct, but does not appear in the consequent at all. (CG is the only exception to this, as its antecedent does not have left and right conjuncts.) Two of these properties—S and KS—are special cases of transitivity. The others, however, are not.

The abbreviations for the properties are intended to be (at least somewhat) mnemonic without taking up too much space. The properties that have received the most attention are S for 'simple', FG for 'finite generalized', and CG for 'complete generalized'.[3] The remaining properties are lopsided; each focusses in on

[2] Unlike 'reflexive' and 'transitive', 'monotonic' here *does* have its usual relation-theoretic sense, w/r/t the order $\subseteq$ on sets of formulas.

[3] I take the terms 'simple' and 'generalized' from [16]. Weir's 'simple transitivity' is my S; his 'generalized transitivity' is my FG. (He does not consider CG.)

either the premise or conclusion side of the relation in question. The abbreviations for these properties include a '/'; where the property focusses on the premise side, a letter appears before '/', and where it focusses on the conclusion side, a letter appears after '/'. The 'F' and 'C' are for 'finite' and 'complete'.

Each property on the list has a *dual* also on the list. Properties P and P' are duals, in the sense relevant here, iff: for a consequence relation $\vdash$ to have P is for its converse $\dashv$ to have P'. The properties S, KS, FG, and CG are all self-dual. For the remaining properties, the names indicate duality; for example, /F and F/ are duals. Also, the assumptions in play about consequence relations (that they are monotonic SET-SET relations) are self-dual; a relation $\vdash$ meets them iff its converse $\dashv$ does. So too is compactness self-dual, in this sense. Noting these symmetries will allow for some of the following proofs to get away with only half the work they would otherwise take. For example, once we see that FG implies /F, we can immediately conclude that it implies F/ as well; and once we see that /C$^+$ does not imply C$^+$/, we can immediately conclude that C$^+$/ does not imply C$^+$/ either. I will use this style of reasoning frequently in what follows.

3 Previous Work

3.1 FG and Cut

In the present setting, FG is equivalent to the following property: if $\Gamma \vdash \Delta, A$ and $A, \Gamma' \vdash \Delta$, then $\Gamma, \Gamma' \vdash \Delta, \Delta'$. This property, in turn, is closely connected to [4]'s rule of *cut* in the sequent calculus LK. (Just like 'transitivity', 'cut' means many different things in different contexts. Most of them, however, are related to Gentzen's use of 'cut'.)

Cut looms large in many proof-theoretic investigations; FG, then, has real proof-theoretic import. But it also, at times, has philsophical import. For example, [7, 8] understand FG (as a condition on a particular consequence relation) as encoding the following constraint on certain conversational norms: if a certain combination of assertions and denials is within the norms, then for any formula A, either adding an assertion of A to that combination remains within the norms, or else adding a denial of A to that combination remains within the norms. [7, 8] endorse this constraint; [9, 10] dispute it.

3.2 CG and Bivaluations

One way to present a consequence relation on a language $\mathcal{F}$ is via *bivaluations*: binary partitions $\langle T, F\rangle$ of $\mathcal{F}$. By specifying a set $\mathfrak{M}$ of such partitions, one specifies a consequence relation $\vdash_{\mathfrak{M}}$ in the following way: $\Gamma \vdash_{\mathfrak{M}} \Delta$ iff there is no $\langle T, F\rangle \in \mathfrak{M}$ such that $\Gamma \subseteq T$ and $\Delta \subseteq F$. (Informally, you might think: the argument is valid iff there is no model on which all the premises are true and all the conclusions false.) This way of thinking is stressed in [13, 6], but even where it is not stressed it is often applicable. For example, any way of presenting a consequence relation using models with *designated values* in the usual way fits this mould directly: we can understand each model as partitioning the language into those formulas that receive a designated value and those that do not.

Any consequence relation arrived at in this way will have certain structural properties: it will be *reflexive* (in the senses of footnote 1), *monotonic*, and it will have the property CG. (For proof, see [13, p. 30].) As we will shortly see, CG in fact implies all the other properties in Table 1. This means that bivaluations will not prove useful in what follows; they obscure the relations between the linking properties under consideration, by forcing them all to hold.[4]

Many monotonic SET-SET consequence relations encountered in the wild can be presented in terms of bivaluations, and so exhibit CG and thus all the linking properties to be considered here. (Note, however, that [11, p. 83] complains that CG is overstrong, claiming that it requires "much more than the transitivity of consequence".) It is only in cases where CG fails that the distinctions explored here are revealed.

3.3 Quantum Logic

[3, p. 44] and [1] both consider forms of quantum logic, and attribute to it the conjunction of /F and F/, which I will call F/F. In quantum logic, distribution of conjunction over disjunction fails; as it happens, there are important connections between distribution and FG, which I do not have space to explore here (but see [6, p. 10], particularly Exercise 0.13.7(i)). In these authors' settings, quantum logic does not obey FG, which they take to be a default expression of transitivity; F/F is substituted to "reflect the transitivity of implication" [1, p. 247].

In both cases, the authors restrict their attention to compact relations, for which the conjunction of /C and C/, which I will call C/C, is equivalent to F/F.[5] (More on compactness presently.) Neither source discusses /F or F/ on their own.

3.4 Neo-Classical Logic

The 'neo-classical' logic explored in [15, 16], among other places, is another consequence relation that exhibits some of these properties but not others. As [16, p. 100] points out, this consequence relation obeys S. In fact it also obeys KS; as we will see, this is stronger. However, it does not exhibit any of the other properties in Table 1. Weir claims that S "should be incorporated in any genuine notion of logical consequence", but does not elaborate.

3.5 Cut$_3$

There is one other property not listed in Table 1 I'm aware of that has been considered a form of 'transitivity' for SET-SET consequence relations. This is

[4] Related techniques from [5], however, can avoid imposing CG.

[5] In fact, Dummett (but not Cutland & Gibbins) only considers finite sequents. Note as well that the discussion in [3] in support of F/F, if cogent, in fact supports the full strength of C/C, even for noncompact relations.

the property called 'Cut$_3$' in [13, p. 32]. A consequence relation $\vdash$ has Cut$_3$ iff whenever $\Gamma \vdash \Delta, A$ for all $A \in \Sigma_1$, and $B, \Gamma \vdash \Delta$ for all $B \in \Sigma_2$, and $\Sigma_1, \Gamma \vdash \Delta, \Sigma_2$, then $\Gamma \vdash \Delta$. But as Shoesmith and Smiley immediately show, Cut$_3$ is equivalent to the conjunction of $/\textsc{c}^+$ and $\textsc{c}^+/$; I will later call this conjunction $\textsc{c}^+/\textsc{c}^+$. (Their proof depends on monotonicity.)[6]

[12, p. 37], oddly, calls this property (there defined directly as the conjunction of $/\textsc{c}^+$ and $\textsc{c}^+/$) 'Cut', and takes it to be of some import. In particular, Segerberg points to FG, claims that it is not sufficient when infinite sets of premises and conclusions are considered, and then offers this property as the appropriate replacement. (He also points out that S, which he calls 'transitivity', is a 'very special case' of this property (p. 38).) I know of no other sources that have attended to this property.

4 Implications

There are ten properties listed in Table 1, and this paper will consider arbitrary conjunctions of these. Our exploration begins, then, with $2^{10} = 1024$ property-specifications to consider. Fortunately, there are many fewer distinct properties actually in play. In this section, I explore implications among these properties, and show that from our 1024, there are at most 21 distinct properties, and at most 7 if compactness is assumed. (I identify properties iff they imply each other.) In fact, these counts are exact, but the 'at least' part of the claim will not be proved until §5. First, I will lay out these implications in three categories: implications by special case, implications by monotonicity, and implications by semilattice properties. Then, I will consider the effects of compactness, and show additional implications among our properties that hold when compactness is assumed.

4.1 Three Kinds of Implications

Some implications from one property to another happen in the easiest possible way: when one property covers only certain special cases of another. These implications can be verified directly by inspection. In this way, five implications are secured: KS implies S; $/\textsc{c}$ implies $/\textsc{f}$; $\textsc{c}/$ implies $\textsc{f}/$; and each of $/\textsc{c}^+$ and $\textsc{c}^+/$ implies FG.

Other implications are not so direct; these require some appeal to monotonicity. The needed appeals to monotonicity, however, are quite formulaic: when one property's antecedent follows by monotonicity from another property's antecedent, then the first property implies the second. This gives eight more implications: each of $/\textsc{f}$ and $\textsc{f}/$ implies KS; FG implies both $/\textsc{f}$ and $\textsc{f}/$; $/\textsc{c}^+$ implies $/\textsc{c}$; $\textsc{c}^+/$ implies $\textsc{c}/$; and CG implies both $/\textsc{c}^+$ and $\textsc{c}^+/$.

[6] [13, p. 30ff.] considers FG, $/\textsc{c}^+$, $\textsc{c}^+/$, Cut$_3$, and CG; the implications and nonimplications among these properties shown there are among what is shown in the present paper.

Finally, implication among properties forms a semilattice with conjunction as the meet.[7] That is, implication is transitive, and the conjunction of two properties is their greatest lower bound w/r/t the implication order. Together with the implications recorded above, this secures a large range of additional implications among the properties under consideration. For example, since FG implies both /F and F/, it follows that it implies their conjunction. Since /C$^+$ implies /C, and /C implies /F, then /C$^+$ implies /F. And so on.

4.2 Twenty-One Properties

These implications narrow the space of properties under consideration to twenty-one: the ten properties that appear in Table 1, plus the eleven additional properties given in Table 2, generated from the original ten by conjunction.

Table 2. Additional linking properties formed by conjunction

Name:	Definition:	Name:	Definition:
F/F	F/ and /F.	C/C	C/ and /C.
F/C	F/ and /C.	C/F	C/ and /F.
/FG/C	/C and FG.	C/FG/	C/ and FG.
C/FG/C	C/, /C, and FG.	C$^+$/C$^+$	C$^+$/ and /C$^+$.
C/C$^+$	C/ and /C$^+$.	C$^+$/C	C$^+$/ and /C.
$\top$	The empty conjunction, exhibited by every consequence relation.		

Given the implications already recorded, each of the $2^{10} = 1024$ property-specifications we can generate from Table 1 by conjunction specifies one of these twenty-one properties. For example, for a consequence relation to exhibit the properties FG, F/, and /C is just for it to exhibit /FG/C, since FG already implies F/. Similarly, for a consequence relation to exhibit KS, /C$^+$, and F/ is just for it to exhibit /C$^+$, which implies the other two properties. And so on, for every combination.

4.3 Compactness

For compact relations, there are more implications to take account of among the properties in play; this section records these and takes account of their impact.

Proposition 1. *If $\vdash$ is compact and has* FG, *then it has* CG.

Proof. See [13, p. 37] for proof. (Their 'cut for formulae' is exactly FG, and their 'cut for sets' is exactly CG.)

Proposition 2. *If $\vdash$ is compact and has* /F, *then it has* /C.

[7] For semilattices (and lattices), see [2].

Proof. Suppose $\vdash$ is compact and has /F, that $\Gamma \vdash A$ for all $A \in \Sigma$, and that $\Sigma, \Gamma \vdash \Delta$. Since $\vdash$ is compact, this gives $\Sigma_{\text{fin}}, \Gamma_{\text{fin}} \vdash \Delta_{\text{fin}}$ for some finite $\Sigma_{\text{fin}} \subseteq \Sigma$, $\Gamma_{\text{fin}} \subseteq \Gamma$, and $\Delta_{\text{fin}} \subseteq \Delta$. By monotonicity, $\Sigma_{\text{fin}}, \Gamma \vdash \Delta$. Since $\Sigma_{\text{fin}} \subseteq \Sigma$, we have $\Gamma \vdash A$ for all $A \in \Sigma_{\text{fin}}$. Now, where n is the cardinality of Σ_{fin}, let $\Sigma_{\text{fin}} = \{\sigma_0, \ldots, \sigma_{n-1}\}$, and for $m \leq n$, let $\Sigma^m_{\text{fin}} = \{\sigma_m, \ldots, \sigma_{n-1}\}$. Thus, $\Sigma^0_{\text{fin}} = \Sigma_{\text{fin}}$, and $\Sigma^n_{\text{fin}} = \emptyset$.

I claim that for any i from 0 to n (inclusive), $\Sigma^i_{\text{fin}}, \Gamma \vdash \Delta$; when $i = n$, this is $\Gamma \vdash \Delta$, and the proposition follows. This can be shown by induction. The case where $i = 0$ is already shown. So suppose the claim is true for $i < n$; then $\Sigma_i, \Gamma \vdash \Delta$, which is to say $\sigma_i, \Sigma_{i+1}, \Gamma \vdash \Delta$. By assumption, $\Gamma \vdash \sigma_i$; monotonicity gives $\Sigma_{i+1}, \Gamma \vdash \sigma_i$. Now, applying /F, $\Sigma_{i+1}, \Gamma \vdash \Delta$.

Proposition 3. *If $\vdash$ is compact and has* F/, *then it has* C/.

Proof. From Proposition 2, by duality.

For compact relations, then, /F and /C are equivalent to each other, as are F/ and C/. This also means that F/F, F/C, C/F, and C/C are all equivalent to each other for such relations. In addition, Since FG implies every property under consideration for compact relations, all of FG, /C$^+$, C$^+$/, CG, /FG/C, C/FG/, C/FG/C, C$^+$/C, C/C$^+$, and C$^+$/C$^+$ are equivalent to each other for these relations. This leaves (at most) seven distinct properties: $\top$, S, KS, /F (= /C), F/ (= C/), F/F (= C/C), and FG (= CG).

The situation so far is recorded in Figure 1. In this figure, each arrow is an implication already recorded; the double-thickness arrows are implications that we have seen become equivalences in the presence of compactness. (For now, you can ignore the letters that label the arrows.) When compactness is assumed, only FG and the six other nodes implied by it remain distinct; each of the other fourteen nodes is connected to one of these seven by a path containing only double-thickness arrows.

5 Nonimplications

So far, only implications have been recorded. So while we know there are *at most* twenty-one distinct properties in play here, and at most seven if compactness is assumed, it's still possible, for all I've said so far, that there are fewer. In fact, there are not; the implications so far recorded exhaust the implications among these properties. This section shows that the remaining potential implications do not hold. In each case, I will show this by counterexample.

5.1 Presenting Consequence Relations

I will present consequence relations using a very simple kind of 'proof system'. I work with *sequents*; a sequent for a language $\mathcal{F}$ is a pair $\langle \Gamma, \Delta \rangle$ of subsets of $\mathcal{F}$;

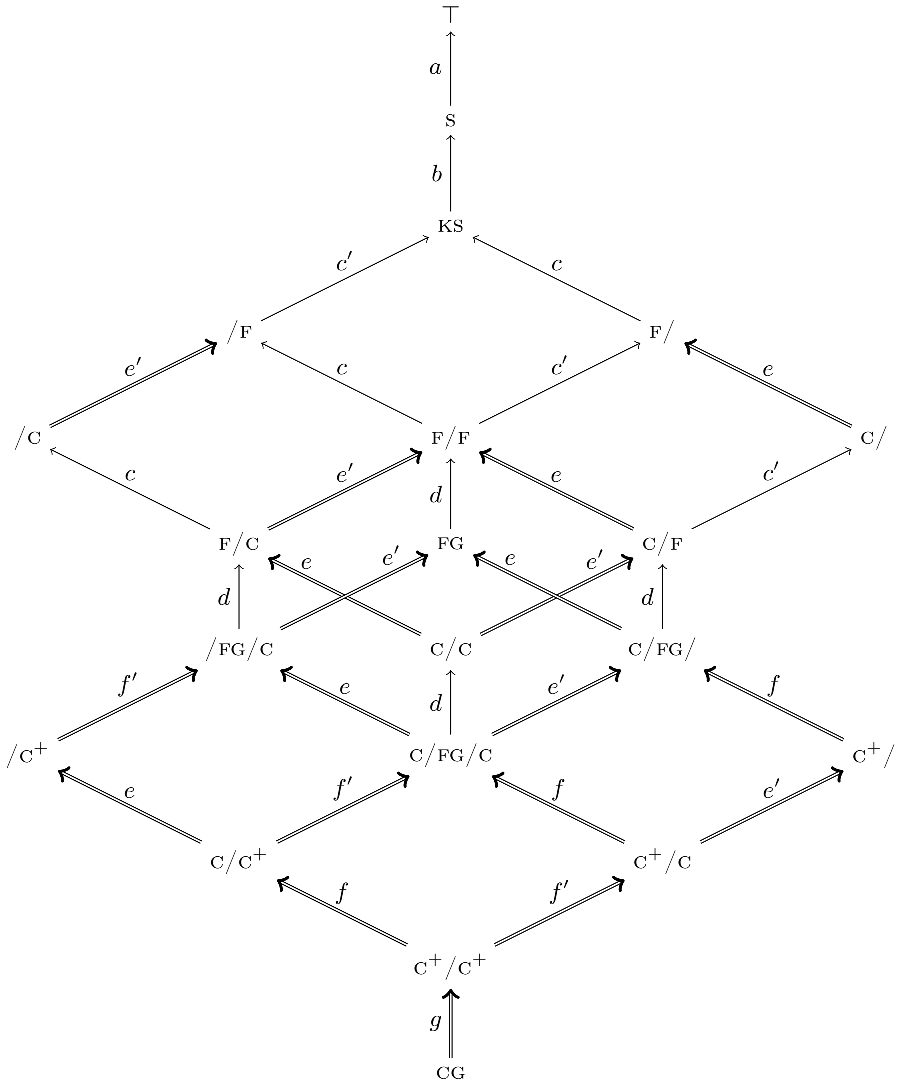

Fig. 1. Implications

I will write such a pair $[\Gamma \therefore \Delta]$. It is handy to consider the *subsequent* relation $\sqsubseteq$, defined: $[\Gamma' \therefore \Delta'] \sqsubseteq [\Gamma \therefore \Delta]$ iff $\Gamma' \subseteq \Gamma$ and $\Delta' \subseteq \Delta$.

A sequent-based proof system involves two components: some set of *initial sequents*, which are simply given as valid, and some *rules* that allow new validities to be generated from old. The proof systems I will draw on here are all quite simple. For each of them, I will specify a set $\mathfrak{P}$ of sequents; the initial sequents of the system are then all those sequents in $\mathfrak{P}$, together with all sequents of the form $[A \mathrel{.\,\cdot.} A]$, for any $A \in \mathcal{F}$. There is only a single rule in any of these systems: the rule of *infinitary weakening*, which allows us to derive $[\Gamma, \Gamma' \therefore \Delta, \Delta']$ from $[\Gamma \therefore \Delta]$, for any $\Gamma, \Gamma', \Delta, \Delta' \subseteq \mathcal{F}$.

So for any set $\mathfrak{P}$ of sequents, we have a consequence relation $\vdash_{\mathfrak{P}}$ determined as follows: $\Gamma \vdash_{\mathfrak{P}} \Delta$ iff either 1) $\Gamma \cap \Delta \neq \emptyset$, or 2) there is some $[\Gamma' \therefore \Delta'] \in \mathfrak{P}$ such that $[\Gamma' \therefore \Delta'] \sqsubseteq [\Gamma \therefore \Delta]$. A binary relation on $\wp(\mathcal{F})$ is $\vdash_{\mathfrak{P}}$ for some $\mathfrak{P}$ iff it is monotonic and reflexive (in the senses of footnote 1, which are equivalent given monotonicity), so this approach works at the right level of generality for present purposes. It also gives a tractable way to explore compactness: note that $\vdash_{\mathfrak{P}}$ is compact iff every infinite sequent in $\mathfrak{P}$ has a finite subsequent in $\mathfrak{P}$.

5.2 A Menagerie of Consequence Relations

Table 3 presents seven distinct consequence relations, *a*–*g*. For each, it notes two of the twenty-one properties: one that the consequence relation has and one that it lacks. These two properties are chosen so that the implications already recorded suffice to settle the situation as regards the remaining nineteen: each other property is either implied by the property the relation has, or else implies the property the relation lacks. (I find this easiest to see by referring to Figure 1.) Let B, C, D, E, F be five distinct formulas, and let $\Theta \subseteq \mathcal{F}$ be infinite. For each of these, Table 4 gives a counterexample to the property that it is listed in Table 3 as lacking; these are easy to check.[8]

Table 3. Seven consequence relations

Name:	$\mathfrak{P}$	Has:	Lacks:
a	$\{[B \therefore C], [C \therefore D]\}$	$\top$	S
b	$\{[\Gamma \therefore \Delta] : \max(\lvert\Gamma\rvert, \lvert\Delta\rvert) > 2 \text{ and } B \in \Gamma \cup \Delta\}$	S	KS
c	$\{[E \therefore B, C, D], [B \therefore C, D]\}$	/C	F/
d	$\{[C \therefore D, B], [B, C \therefore D]\}$	C/C	FG
e	$\{[\Gamma \therefore \Delta] : \Delta \text{ is infinite or } \Gamma \cap \Theta \neq \emptyset\}$	/C$^+$	C/
f	$\{[\Gamma \therefore \Delta] : \Delta \text{ is infinite or } \lvert\Gamma\rvert \geq 2\}$	C/C$^+$	C$^+$/
g	$\{[\Gamma \therefore \Delta] : \Gamma \cup \Delta \text{ is infinite}\}$	C$^+$/C$^+$	CG

For space reasons, I do not prove here that every relation in Table 3 has the property it is there claimed to have; none of the needed proofs is particularly devious. Here are two examples to give the flavour.

Proposition 4. *Relation d has* C/C.

Proof. Suppose it lacks /C; then there are Γ, Δ, Σ such that $\Gamma \nvdash_d \Delta$ while $\Gamma \vdash_d A$ for every $A \in \Sigma$ and $\Sigma, \Gamma \vdash_d \Delta$. If $\Sigma \subseteq \Gamma$, then $\Gamma \vdash_d \Delta$, contrary

[8] [13, p. 31] gives the relation here called g, for the same purpose: to show that C$^+$/C$^+$ and CG are distinct. See their Theorem 2.7 (p. 32).

Table 4. Counterexamples

Name:	Lacks:	Validates:	And:	But:
a	S	$B \vdash C$	$C \vdash D$	$B \nvdash D$
b	KS	$C, D \vdash B$	$B \vdash E, F$	$C, D \nvdash E, F$
c	F/	$E \vdash B, C, D$	$B \vdash C, D$	$E \nvdash C, D$
d	FG	$C \vdash D, B$	$B, C \vdash D$	$C \nvdash D$
e	C/	$\vdash \Theta$	$A \vdash$ for all $A \in \Theta$	$\nvdash$
f	C$^+$/	$B \vdash \mathcal{F} \setminus \{B\}$	$B, A \vdash$ for all $A \in \mathcal{F} \setminus \{B\}$	$B \nvdash$
g	CG	$\mathcal{F}^+ \vdash \mathcal{F}^-$ for every partition $\langle \mathcal{F}^+, \mathcal{F}^- \rangle$ of $\mathcal{F}$		$\nvdash$

to supposition. So there must be some $A \in \Sigma$ with $A \notin \Gamma$. Since $\Gamma \vdash_d A$ and $A \notin \Gamma$, it must be that $B, C \in \Gamma$ and $A = D$. Since $\Gamma \nvdash_d \Delta$ while $B, C \in \Gamma$, it must be that $D \notin \Delta$. Now, suppose $E \in \Sigma \cap \Delta$; since $\Gamma \vdash_d E$ and E is not D, we must have $E \in \Gamma$. Then $\Gamma \vdash_d \Delta$, contrary to supposition. So $\Sigma \cap \Delta$ is empty. But then $(\Sigma \cup \Gamma) \cap \Delta$ is empty, and since $D \notin \Delta$, it follows that $\Sigma, \Gamma \nvdash_d \Delta$. Contradiction.

For C/, the argument is dual, reversing the roles of C and D.

Proposition 5. *Relation f has* C/C$^+$.

Proof. First, that it has C/. If $A \vdash_f \Delta$ for each $A \in \Sigma$, then either Δ is infinite, in which case $\Gamma \vdash_f \Delta$ directly, or else $\Sigma \subseteq \Delta$; the only valid arguments with finitely many conclusions and a single premise are those where the premise is among the conclusions. But if $\Sigma \subseteq \Delta$, then if $\Gamma \vdash_f \Delta, \Sigma$, this is already $\Gamma \vdash_f \Delta$.

Second, that it has /C$^+$. Suppose $\Sigma, \Gamma \vdash_f \Delta$ and $\Gamma \vdash_f \Delta, A$ for each $A \in \Sigma$, to show $\Gamma \vdash_f \Delta$. If Δ is infinite, we're done; if $|\Gamma| \geq 2$ we're done; if $\Gamma \cap \Delta \neq \emptyset$ we're done. So suppose Δ is finite, $|\Gamma| < 2$, and $\Gamma \cap \Delta = \emptyset$. Since $\Sigma, \Gamma \vdash_f \Delta$, either $\Sigma \cap \Delta \neq \emptyset$ or $|\Sigma \cup \Gamma| \geq 2$. In the first case, take some $A \in \Sigma \cap \Delta$; since $\Gamma \vdash_f \Delta, A$ and $\Delta \cup \{A\} = \Delta$, we're done. In the second case, there must be some $A \in \Sigma$ but $A \notin \Gamma$; we then have $\Gamma \vdash_f \Delta, A$. But $|\Gamma| < 2$ and $\Gamma \cap (\Delta \cup \{A\}) = \emptyset$, so this is impossible.

5.3 No Further Implications

To see that there are no further implications, return to Figure 1, now attending to the letters that label the arrows. These letters correspond to consequence relations from §5.2; letters with $'$ pick out converse relations. The indicated consequence relation, in each case, is a counterexample to the claim that the implication in question is an equivalence. Moreover, where the implication is a single-line arrow—that is, where it is not already known to become an equivalence in the presence of compactness—the indicated counterexample is compact; this shows that no additional implications collapse to equivalences in the presence of compactness.[9]

[9] a–d are compact. For a, c, and d, $\mathfrak{P}$ contains only finite sequents. For b, $\mathfrak{P}$ contains a finite subsequent of each infinite sequent it contains. (e–g are not compact, and could not be, given the combinations of properties they exhibit.)

Because the properties under consideration are closed under conjunction, this suffices to rule out any additional implications. Any additional implication would bring with it an additional equivalence (if P implies Q, then P is equivalent to the conjunction of P and Q; this conjunction is already known to imply P); so the fact that there are no additional equivalences suffices to show that there are no additional implications.

As a result, the implications between these properties are now completely characterized. By taking arbitrary conjunctions of the ten properties in Table 1, there are exactly twenty-one distinct properties we can reach, twenty of which (all but $\top$) are linking properties, properties of the sort that can plausibly travel under the name 'transitivity'. For compact relations, these twenty-one collapse to seven, six of which (again, all but $\top$) are such linking properties.

6 Conclusion

'Transitivity', as applied to consequence relations, can conceal more than it reveals. When someone says a consequence relation is 'transitive', then, it is worth asking just what is meant. It is almost never the case that they mean that it is *transitive*, in the usual relation-theoretic sense. But then what can they mean?

This paper has explored some possible answers. It's a safe bet that nobody means $\top$ by 'transitivity', but the remaining twenty properties (in the general case) or six properties (in the presence of compactness) are all possible ways to fill in the idea. When *we* call consequence relations 'transitive', then, it behooves us to make clear exactly what we are saying; there is no single thing we must obviously mean.

Acknowledgements. For helpful discussion and comments, thanks to the Melbourne Logic Group (especially Rohan French and Lloyd Humberstone), audiences at the University of Groningen (GroLog) and the Australasian Association for Logic 2015 meeting, and three anonymous referees for LORI. This research was partially supported by the grant "Non-Transitive Logics", number FFI2013-46451-P, from the Ministerio de Economía y Competitividad, Government of Spain.

References

[1] Cutland, N.J., Gibbins, P.F.: A regular sequent calculus for quantum logic in which $\wedge$ and $\vee$ are dual. Logique et Analyse 25(99), 221–248 (1982)
[2] Davey, B.A., Priestley, H.A.: Introduction to Lattices and Order. Cambridge University Press, Cambridge (2002)
[3] Dummett, M.: The Logical Basis of Metaphysics. Duckworth, London (1991)
[4] Gentzen, G.: Investigations into logical deduction. In: Szabo, M.E. (ed.) The Collected Papers of Gerhard Gentzen, pp. 68–131. North-Holland Publishing Company, Amsterdam (1969)

[5] Humberstone, L.: Heterogeneous logic. Erkenntnis 29, 395–435 (1988)
[6] Humberstone, L.: The Connectives. MIT Press, Cambridge (2012)
[7] Restall, G.: Multiple conclusions. In: Hajek, P., Valdes-Villanueva, L., Westerståhl, D. (eds.) Logic, Methodology, and Philosophy of Science: Proceedings of the Twelfth International Congress, pp. 189–205. Kings' College Publications, London (2005)
[8] Restall, G.: Truth values and proof theory. Studia Logica 92(2), 241–264 (2009)
[9] Ripley, D.: Paradoxes and failures of cut. Australasian Journal of Philosophy 91(1), 139–164 (2013)
[10] Ripley, D.: Anything goes. Topoi 34(1), 25–36 (2015)
[11] Rumfitt, I.: Knowledge by deduction. Grazer Philosophische Studien 77(1), 61–84 (2008)
[12] Segerberg, K.: Classical Propositional Operators: An Exercise in the Foundations of Logic. Clarendon Press, Oxford (1982)
[13] Shoesmith, D.J., Smiley, T.J.: Multiple-conclusion Logic. Cambridge University Press, Cambridge (1978)
[14] Steinberger, F.: Why conclusions should remain single. Journal of Philosophical Logic 40(3), 333–355 (2011)
[15] Weir, A.: Naive truth and sophisticated logic. In: Beall, J.C., Armour-Garb, B. (eds.) Deflationism and Paradox, pp. 218–249. Oxford University Press, Oxford (2005)
[16] Weir, A.: A robust non-transitive logic. Topoi 34(1), 99–107 (2015)