The Consequentiality of Explicitation

Ryan Simonelli

June 16, 2025

0 Introduction

Saying things—making assertions—often commits one to making other assertions. For instance, asserting "The ball is crimson" commits one to asserting "The ball is red." Of course, being committed to an assertion does not mean that one must actually make it unprompted (if it did, we'd be stuck saying things non-stop all day). But it does mean that one must be *prepared* to make it (and defend it) if appropriately prompted to do so. This notion of committive consequence is a core semantic notion appealed to in the two leading alternatives to truth-conditional semantics: dynamic semantics and inferentialist semantics. In both traditions, however, structural constraints are generally placed on the basic framework that preclude the following possibility: by *making commitments explicit*, actually *making* assertions to which one was previously only *committed*, one can come to take on *new commitments*. Drawing a phrase from Brandom [3], in both traditions, explicitation is generally taken to be inconsequential.

In this paper, I'll argue from three seemingly quite different cases, that explicitation is, at least sometimes, consequential. The first case has to do with defeasible material inferences of the sort of concern to Sellars [25] and Brandom [1] [2] [3], the second has to do with Frege's [9] puzzle about co-reference and belief, and the third has to do with McGee's [18] puzzle about modus ponens. After motivating the consequentiality of explicitation with these three cases, I'll put forward a formal framework that combines elements of both the dynamic tradition and the inferentialist tradition that formally explicates the consequentiality of explicitation that obtains in these cases. Beyond enabling us to make formal sense of this motivating data, this dynamic inferentialist framework yields a new conception of the basic notion of consequence in natural language, differing both from truth-conditional theories and standard non-truth-conditional alternatives.

1 The Inconsequentiality of Explicitation in Two Traditions

In the current semantic landscape, there are many formal frameworks that aim, for various reasons, to understand meaning in non-truth-conditional terms. Two leading alternatives are *dynamic* semantic approaches and *inferentialist* semantic approaches. There are fundamental differences between both traditions. Most notably, dynamic semantics is standardly carried out in a *model-theoretic* framework, whereas inferentialist semantics is standardly carried out in a *proof-theoretic* framework. However, both traditions can be understood as broadly *use-theoretic* traditions, thinking of the meaning of a sentence in terms of what its use functions to do or what the rules governing its use are, rather than what it stands for or represents. Accordingly, in such use-theoretic semantic frameworks, one notion in terms of which it is natural to articulate semantic properties is the normative notion of *commitment*. If a speaker is committed to some sentence A, then it is appropriate for other speakers to hold them to asserting A if prompted. A crucial semantic relation between sentences, then, is that of *committive consequence*: asserting some sentences can commit speakers to other sentences. My concern here will be with a basic structural feature of committive consequence that is assumed to hold in both traditions. In both traditions, actually *making* assertions to which one was antecedently only *committed*—making these commitments *explicit*—cannot result in one's taking on *new commitments*. Following Brandom [3] I'll refer to this structural feature as the "inconsequentiality of explicitation."

I'll start with dynamic semantics. The basic idea of dynamic semantics, as put by Groenendijk, Stokhof, and Veltman [32], is that "the meaning of a sentence is the change an utterance of it brings about." A bit more officially, a dynamic meaning is a *context change potential*: a function from each context in which it might be uttered to the context that would result upon its utterance. In general, the semantic value of a sentence A is a function that maps each context or state σ in which A might be uttered to the state $\sigma[A]$, which would result upon the utterance (and acceptance) of A in that state. In standard informational versions of dynamic semantics a context or state σ is taken to be a set of possible worlds—those that are possible, given the information possessed by the agent(s) in that context.[1] In the basic paradigm cases, the utterance of a sentence functions to eliminate worlds from an information

[1] Or, if one wants a bit more structure in one's contexts, sets of sets of possible worlds. Such differences are not important for our purposes.

state. However, this informational specification of the change brought about by the utterance of a sentence is not necessary to get the basic idea of dynamic semantics and the basic structure of the core dynamic framework into view. At a very high level of abstraction, we can note that standard dynamic frameworks are such that explicitation is inconsequential.

The standard way of defining commitment in the context of a dynamic framework is to take it that "the participants of a context are committed to something just in case uttering it leaves their state unchanged," [11, 303]. Officially, we define commitment relative to a state and a sentence as follows:

> **Commitment:** An agent in state σ is committed to A, $\sigma \Vdash A$, just in case $\sigma[A] = \sigma$.

It is easy to easy to see that, if an agent in state σ is committed to some sentence, then updating σ with that sentence will not result in any new commitments. Officially, defining commitment in this way has the following consequence:[2]

> **Inconsequentiality of Explicitation:** For any state σ, if $\sigma \Vdash A$ and $\sigma[A] \Vdash B$, then $\sigma \Vdash B$.

The inconsequentiality of explicitation is closely linked to a structure feature of *entailment*, defined as committive consequence. Sentences A_1, $A_2 \ldots A_n$ entail some sentence B just in case updating a state with all of the premise sentences, in that order, yields a state that is committed to B. Following Veltman [34] we might define two such notions of entailment: a *strict* notion, that involves universal quantification over *all* states, and a *weak* notion, which is committive consequence relative to the least opinionated *minimal* state, μ. Thus:

> **Strict Committive Entailment:** A_1, $A_2 \ldots A_n$ *strictly* entails B just in case, for any state σ, $\sigma[A_1][A_2] \ldots [A_n] \Vdash B$.
>
> **Weak Committive Entailment:** A_1, $A_2 \ldots A_n$ *weakly* entails B just in case, for the minimal state μ, $\mu[A_1][A_2] \ldots [A_n] \Vdash B$.

It is easy to see that, for both notions of entailment, the following fact holds:[3]

[2] The proof is trivial. Suppose $\sigma \Vdash A$ and $\sigma[A] \Vdash B$. Since $\sigma \Vdash A$, by Commitment $\sigma[A] = \sigma$, and so, since $\sigma[A] \Vdash B$, it follows by substitution that $\sigma \Vdash B$. So, if a state σ commits one to a sentence A and updating σ with A commits one to a sentence B, then σ already commits one to B.

[3] Consider an arbitrary state σ and suppose that $A \vDash B$ and $A, B \vDash C$. Since $A \vDash B$, $\sigma[A] \Vdash B$ and so $\sigma[A] = \sigma[A][B]$, and, since $A, B \vDash C$, $\sigma[A][B] \Vdash C$. So, $\sigma[A][B][C] = \sigma[A][B]$, so, by substitution, $\sigma[A][C] = \sigma[A]$, and so $\sigma[A] \Vdash C$. Since σ was arbitrary, $A \vDash C$. $\square$

$$\frac{A \vDash B \quad A, B \vDash C}{A \vDash C}$$

This feature of entailment is, in some sense, just a different statement of the inconsequentiality of explicitation. If updating with A commits one to B, and updating with A along with B commits one to C, then updating with A *already* commits one to C. Though there are many developments and refinements of dynamic frameworks, this basic structural feature of committive consequence remains in place across all standard variants of dynamic semantics. The fundamental notion of commitment defined in standard dynamic frameworks is such that making one's commitments explicit is always inconsequential.

Let me now turn to the inferentialist tradition. Inferentialism aims to understand the meaning of a linguistic expression in terms of the inferential rules governing its use. Originally developed by Sellars [25] [26] [27] and Brandom [1] [2] (with influence from Wittgenstein [38]), in the past few decades there has been substantial efforts to turn inferentialism into a formal semantic theory, capable of standing as a genuine rival to traditional model-theoretic semantics. The basic formal framework in which inferentialist accounts of meaning have been developed is *proof-theoretic semantics*, an idea owed primarily to Gentzen [10] and notably developed by Dummett and Prawitz [20]. The core idea of proof-theoretic semantics is that proof rules can themselves be understood as semantic clauses. Consider the standard rules for conjunction:

$$\frac{A \quad B}{A \wedge B}\,{\scriptstyle \wedge_I} \qquad\qquad \frac{A \wedge B}{A}\,{\scriptstyle \wedge_{E_1}} \qquad\qquad \frac{A \wedge B}{B}\,{\scriptstyle \wedge_{E_2}}$$

Here, the horizontal line can be understood as expressing a primitive relation of committive consequence. Thus, the introduction rule can be understood as saying when someone asserts A and asserts B, then they are committed to $A \wedge B$, and the elimination rules can be understood as saying that if one asserts $A \wedge B$, then they are committed to A, and they are also committed to B. In general, expressing such a rule horizontally in sequent notation, a sequent of the form $X \vdash A$ (where X is a set of sentences), can be understood as saying that asserting everything in X commits one to A. Here, X might be understood as the set of *explicit* commitments, whereas A is an *implicit* commitment. However, in many developments of inferentialism, this distinction between explicit and implicit commitments is lost.

As before, the question of whether explicitation is consequential essentially amounts to whether the following structural rule, known as *Cumulative Transitivity*, holds of the consequence relation:

$$\frac{X \vdash A \quad X, A \vdash B}{X \vdash B}$$

Once again, this rule says that if asserting everything in X commits one to A, and asserting everything in X along with asserting A commits one to B, then asserting everything in X already commits one to B. Thus, if this rule holds, explicitating the commitment to a sentence A, implicit in the set of assertions X, will always be inconsequential. Now, in a proof-theoretic context, whether or not Cumulative Transitivity holds of the proof-theoretically-defined consequence relation will depend on what sort of proof system ones opts for in spelling out one's semantic theory. Inferentialist theories of meaning, following Gentzen [10] and Prawitz [20], are typically formulated in *natural deduction* proof systems, and the consequence relations imposed by such systems are Transitive by default, since proofs can be strung together as follows:

$$\begin{array}{cc} & X \\ & \vdots \\ X & A \\ \vdots & \\ B & \end{array}$$

Beyond just the set-up of the proof-system, the operational rules for the connectives also typically impose Transitivity. Consider, for instance, the standard rules for the material conditional, written now is sequent notation:

$$\frac{X, A \vdash B}{X \vdash A \supset B}\supset_I \qquad\qquad \frac{X \vdash A \supset B \quad X \vdash A}{X \vdash B}\supset_E$$

Given these rules, Cumulative Transitivity can be derived as follows:

$$\frac{\dfrac{X, A \vdash B}{X \vdash A \supset B}\supset_I \quad X \vdash A}{X \vdash B}\supset_E$$

Now, there are ways of restricting the structure of proofs so as to try to rule out proofs of this sort as legitimate (e.g. [33]), however, insofar as the rules themselves are taken to be valid, such restrictions are hard to motivate in an inferentialist context. As far as I'm aware, all proof-theoretic developments of inferentialism in natural deduction systems impose Cumulative Transitivity, and thus preclude the possibility of the consequentiality of explicitation.

Recently, some inferentialists have proposed to proof-theoretic developments of inferentialism in the context of a sequent calculus. In such a system, rather than providing introduction and elimination rules codifying the role of a logically complex sentence as a conclusion and a premise, one provides right rules and left rules. For instance, standard rules for conjunction in the context of a sequent calculus are the following:

$$\frac{X \vdash A \quad X \vdash B}{X \vdash A \wedge B} \qquad\qquad \frac{X, A, B \vdash C}{X, A \wedge B \vdash C}$$

Unlike natural deduction, these sequent rules do not enable one to derive Transitivity, and it is a consequence of Gentzen's Cut Elimination theorem that the sequent calculus is complete without a Transitivity principle. This formal fact has led some inferentialists to question Transitivity. Perhaps most notably, in response to paradoxes such as the liar and Sorites, Ripley [22] draws on Restall's [21] bilateral interpretation of the sequent calculus to call Transitivity into doubt.[4] On Restall's interpretation, a sequent of the form $X \vdash Y$ is read as saying that asserting everything in X and denying everything in Y is incoherent or "out of bounds." With this interpretation in mind, consider again the principle of Cumulative Transitivity for unsigned sentences:

$$\frac{X \vdash A \quad X, A \vdash B}{X \vdash B}$$

To reject this principle, on Restall's bilateral interpretation, is to maintain that there might be some set of sentences X and sentences A and B such that (1) asserting everything in X and denying A is out of bounds, (2) asserting everything in X along with asserting A and denying B is out of bounds, but (3) asserting everything in X and denying B is not out of bounds. To see that there might be such a case on a certain plausible approach to paradoxes, suppose A is a sentence such as the liar that's always out of bounds to assert or deny, or a sentence involving a vague predicate which should neither be asserted nor denied in a borderline case. Then,

[4] See also [6], [5]

regardless of what X and B are, we will have the two premise sequents, but the conclusion sequent, which says that asserting everything in X along with denying B is incoherent, shouldn't follow.

While this sort of approach may be a plausible way of dealing with the paradoxes, such cases, while formally modeled as transitivity failures in the context of a sequent calculus, are not conceived of as genuine failures of the transitivity of committive consequence by their proponents. Indeed, following Restall, Ripley [23] explicitly eschews a notion of committive consequence in spelling out an inferentialist approach to meaning, focusing instead solely on the notion of the coherence or incoherence of sets of assertions and denials.[5] According the "transitivity" failures that arise with paradoxical sentences, on this approach, are not really failures of transitivity of *consequence* at all. Rather, as Restall spells it out, these are failures of the principle that Restall dubs "Extensibility," which claims that, for every coherent set of assertions and denials Γ and every sentence A, Γ can be coherently extended to include either the assertion of A or the denial of A. In slightly different terminology, these failures might be understood as failures of a principle we might call *bilateral excluded middle* (c.f. [7], [28]), maintaining that there are some sentences in some contexts that are neither coherent to assert nor coherent to deny.[6] By contrast, the cases that follow, which are independent of paradoxes such as the heap and the liar, are presented as genuine failures of the transitivity of committive consequence.

2 Three Cases of the Consequentiality of Explicitation

Before I go on to introduce the cases of the consequentiality of explicitation on which I'll focus for the remainder of the paper, let me introduce one modification in approach, following a recent trend. In recent years, both in the dynamic tradition and the inferentialist tradition, *bilateral* approaches to meaning, according to which assertion and denial are treated as equally basic, have become prominent.[7] I will follow suit here, adopting a bilateral approach both in presenting the cases and in developing the framework that will enable us to make sense of them. Thus, rather than just speaking of commitments to *sentences*, I will speak of commitments to

[5] See [31] for a criticism of this feature of Restall and Ripley's approach.

[6] See [28] for a detailed development of this line of thought.

[7] In the dynamic tradition, see [35] [37], and, in the inferentialist tradition, see [24], [**?**] [8], [**?**], [14], [13]

the opposite *speech acts* of *asserting* or *denying* sentences. I'll refer to such speech acts as "discursive moves," or just "moves" for short. Unlike Restall's version of bilateralism mentioned above, according to which sequents are *interpreted* in bilateral terms, I will follow Rumfitt [24] in making the bilateralism explicit in the notation itself, positively or negatively signing sentences to express their assertion or denial. Thus, $+\langle A\rangle$ expresses the assertion of A, whereas $-\langle A\rangle$ expresses the denial of A. Continuing to use capital letters such as A and B for sentences, I will use Greek letters such as φ and ψ for moves which may be assertions or denials, and, continuing to X and Y for sets of sentences, I will use capital Greek letters such as Γ and Δ for sets of moves, which, once again, may be assertions and denials.

The cases of the consequentiality of explicitation I will present in this section will all take the following form:

1. Given that one has made some set moves Γ, one is committed to φ.
2. If one actually makes the move φ, explicitly acknowledging what one was previously only implicitly committed to, one is committed to some further move, ψ.
3. However, originally, having just made the moves in Γ, one is not committed to ψ.

So, reading $\Gamma \vdash \varphi$ as saying that making all the moves in Γ commits one to φ, these are all cases in which in the following structural rule, known as *Cumulative Transitivity* (now construed as a structural rule pertaining to bilateral consequence), fails:

$$\frac{\Gamma \vdash \varphi \quad \Gamma, \varphi \vdash \psi}{\Gamma \vdash \psi}$$

The failure of this principle, in the cases that follow, will have to do with the essentially *dynamic* character of commitment attribution and acknowledgment. That is, one can *attribute* a set of commitments to some speaker, and if that speaker explicitly *acknowledges* those commitments, this can lead to one's attribution of further commitments that one had not previously attributed.

Now, for each case, there are possible alternative ways of accommodating it. I am not going to go try the rule out these possible alternatives for each case. Rather, by giving these three different cases and proposing a simple unified explanation of what is going on in each of them, I hope to show that there is indeed a unified phenomenon here: these are all cases in which explicitation is consequential.

2.1 The Maddy Case

Let me start with cases having to do with *material* inferences. Such inferential relations are of crucial importance in the inferentialist tradition. Consider what Brandom [1] says about material inferential relations of committive consequence:

> Inheritance of commitment (being committed to one claim as a consequence of commitment to another) is what will be called a committive, or commitment-preserving, inferential relation. Deductive, logically good inferences exploit relations of this genus. But so do materially good inferences, such as inferences of the form: A is to the West of B, so B is to the East of A; This monochromatic patch is green, so it is not red; Thunder now, so lightning earlier. Anyone committed to the premises of such inferences is committed thereby to the conclusions.

Brandom here is following the lead of Sellars, who argues that "Material rules are as essential to meaning (and hence to language and thought) as formal rules," ([25]). Sellars, however, also draws attention to the fact that it's a feature of material rules of inference that they are generally *defeasible*. Though the particular examples that Brandom gives here are examples of strict conceptual or nomological necessities, the vast majority of material inferences are not like this. Consider, for instance, the following inferences: It's a beer, so it's alcoholic; It's a rose, so it's red; It's a fish, so it doesn't walk on land; It's bird, so it flies. Each of these inferences are materially good, but the addition of certain premises can *defeat* their material goodness: the beer might be an O'doul's, the rose might be a white rose, the fish might be a mudskipper, and the bird might be a penguin. Though each of these inferences is, by default, a good one, if any of these circumstances obtain, the inference is no longer good.

Transposing talk from inference to the attribution of commitments, the basic idea of a defeasible relation of committive consequence is that, *in general*, upon making a move φ, one is scored as committed to ψ. Yet the attribution of this commitment may be *defeated* upon one's making some other moves. For instance, to use an example from Simonelli [29], in general, asserting "Maddy's drinking a beer" commits one to "Maddy's drinking an alcoholic beverage." However, if one asserts "Maddy's drinking a beer" along with asserting "Maddy's drinking an O'Doul's," one does commit oneself to "Maddy's drinking an alcoholic beverage." Now, as just articulated, this case of a defeasible relation of committive consequence involves a failure of the structural rule of *Monotonicity:*

$$\frac{\Gamma \vdash \varphi}{\Gamma, \psi \vdash \varphi}$$

However, as Simonelli [29] points out, such failures of Monotonicity generally also give rise to failures of Cumulative Transitivity. Consider the following failure of Cumulative Transitivity corresponding to the failure of Monotonicity just cited:

The Maddy Triad:

1. Asserting "Maddy's drinking a beer" commits one to asserting "Maddy's drinking an alcoholic beverage."
2. Asserting "Maddy's drinking a beer" along with "Maddy's drinking an alcoholic beverage," commits one to denying "Maddy's drinking an O'Doul's."
3. However, asserting "Maddy's drinking a beer" by itself does not commit one to denying "Maddy's drinking an O'Doul's."

Insofar as we countenance defeasible relations of committive consequence at all, it seems that we should accept (1) as a genuine relation of committive consequence. Moreover, (2) is undeniable. It seems, however, that we shouldn't say that asserting "Maddy's drinking a beer" commits one to denying "Maddy's drinking an O'Doul's." Just because one has asserted that someone's drinking a beer, one is under no compulsion to defend the claim that they're not drinking an O'Doul's; after all, they might be a teetotaler.

Rather than rejecting one of the claims in the triad, the suggestion here is simply that we accept this as a case in which explicitation is consequential. That is, when someone says "Maddy's drinking a a beer," I will take them to be (defeasibly) committed to assert "Maddy's drinking an alcoholic beverage." If, however, they explicitly *acknowledge* that commitment, actually going on to *assert* "Maddy's drinking an alcoholic beverage," I'm then going to take them to be committed to denying "Maddy's drinking an O'Doul's," a commitment I had not previously attributed to them. In this way, by explicitly acknowledging commitments to which one was previously only implicitly attributed, one comes to take on new implicit commitments. Having stated the suggestion for making sense of this case, let me turn to the second case.

2.2 The Superman Case

In the history of the development of semantic frameworks, since Frege himself, it has been taken to be a basic criterion of adequacy for a semantic framework that it be able to handle the sort of cases involving referential opacity known as "Frege puzzles," where two different terms, with differing conceptual significances, refer to the same individual. To take a classic contemporary case, consider the names "Superman" and "Clarke Kent." At least to someone like Lois Lane, these two names have different conceptual significances, yet there is one individual to which they both refer. There are a number of philosophical issues that such cases give rise to. Here, I want to show how these cases can be understood as involving another instance of the consequentiality of explicitation.

Suppose Lois Lane asserts "Superman flies" and then, shortly latter, denies "Clarke Kent flies." Given the bilateralist account of denial and negation, Lois's denial of "Clarke Kent flies" has the same significance as the assertion of "Clarke Kent doesn't fly," and, since we know that Superman and Clarke Kent are the very same person, we take it that, in committing herself to "Clarke Kent doesn't fly," Lois has committed herself to "Superman doesn't fly." Of course, Lois Lane herself doesn't know that Superman is Clarke Kent, and so she wouldn't actually *say* "Superman doesn't fly." However, *we* know this, and it's on the basis of *our* knowledge of entailment relations, not on the basis of *her* knowledge, that we attribute commitments to her. So, we take her to be committed to "Superman flies" (given that she's asserted this very sentence) and "Superman doesn't fly" (given that she's asserted "Clarke Kent doesn't fly") and, so, insofar as conjunction introduction is commitment preserving, we take her to be committed to the contradictory claim "Superman flies and doesn't fly."

Now, one crucial desideratum for many dynamic and proof-theoretic semantic frameworks is that these systems are able to preserve the strength of classical logic, at least for modal and conditional-free fragments of a language. In a bilateral context, this means maintaining the following:[8]

> **Classicality:** If $A_1, A_2 \ldots A_n$ classically entail B, then $+\langle A_1\rangle, +\langle A_2\rangle \ldots + \langle A_n\rangle \vdash +\langle B\rangle$ (where neither $A_1, A_2 \ldots A_n$ nor B contain any modal or conditional operators).

[8] See [17] for a discussion of failures of classicalility in standard dynamic theories of epistemic modality.

Given that, for any sentences A and B, $A \wedge \neg A$ classically entails B, maintaining Classicality requires taking it that asserting "Superman flies and doesn't fly" commits one to asserting "The moon is made of cheese." Of course, one might object to Classicality, opting instead for a paraconsistent logic in which the principle of explosion fails to hold. However, I'll work on the assumption that Classicality, in the sense just defined, is a non-negotiable requirement.

And yet, it seems that Lois's performing the speech acts that she does, asserting "Superman flies" and denying "Clark Kent flies," should not commit her to asserting everything. For instance, this set of speech acts should not commit her to asserting "The moon is made of cheese." Thus we have the following case of an apparent failure of Cumulative Transitivity:

The Superman Case:

1. Asserting "Superman flies" along with denying "Clark Kent flies" commits one to asserting "Superman flies and doesn't fly."
2. Asserting "Superman flies," denying "Clarke Kent flies," and asserting "Superman flies and doesn't fly" commits one to asserting "The moon is made of cheese."
3. However, asserting "Superman flies" and denying "Clark Kent flies" doesn't, by itself, commit one to asserting "The moon is made of cheese."

The acceptance of (1) is the consequence of basic logical principles concerning negation and conjunction along with a principle regarding the substitutability of co-referential names in the context of commitment attribution. As we've said, (2) follows from the classicality. Given that we are treating classicality as a non-negotiable here, this leaves us with the option rejecting (3). The problem with rejecting (3), however, is that it seems crucial to our account of discourse, that, when someone has committed themself to a contradiction in virtue of making two incompatible moves, they have committed themself to a *specific* contradiction. Bringing out the specific contradiction to which they've committed themself in the course of discourse is a crucial dialectical move, requiring them to revise the retract one of the moves that are jointly incompatible. If, upon making two incompatible claims, one comes to be committed to *everything*, we are precluded from being able to make sense of this possibility.

I suggest, then, that we understand the Superman Case as another case in which explicitation is consequential. Upon making the moves that she does, Lois commits herself to the contradiction "Superman flies and doesn't fly." Now, if she actually goes on to *assert* this contradiction, then we might plausibly throw up our hands and say "I don't know how to score her now," an attitude reflected in our formal system by attributing to her a commitment to everything. Her merely being *committed* to this contradiction, however, does not commit her to everything. On the other hand, if This strikes me as the only plausible way to deal with this case while retaining Classicality, and it requires rejecting Cumulative Transitivity.

2.3 The Reagan Case

The final example is a variation on McGee's proposed counterexample to modus ponens.

> Opinion polls taken just before the 1980 election show the Republican Ronald Reagan decisively ahead of the Democrat Jimmy Carter, with the other Republican in the race, John Anderson, a distant third.

Though McGee uses this example to propose a counterexample to modus ponens, we can just as well see it as giving us a counterexample to Cumulative Transitivity:

> **The Reagan Case:**
>
> 1. Asserting all of the polling data commits one to asserting "A Republican will win the election."
> 2. Asserting all of the polling data along with asserting "A Republican will win the election" commits one to asserting "If it's not Reagan who wins, it will be Anderson."
> 3. However, asserting all of the polling data doesn't, by itself, commit one to asserting "If it's not Reagan who wins, it will be Anderson."

The proposal here, once again, is simply that one accepts all of these claims and maintains that this is a case of the consequentiality of explicitation. That is, asserting the polling data *commits* one to the claim "A Republican will win the election." However, actually *asserting* "A Repuplican will win the election," making this commitment *explicit*, one commits oneself to the conditional "If it's not Reagan who wins, it will be Anderson," something to which one was not antecedently committed to in just asserting the polling data.

Now, it's worth noting that, as I've presented it here, this example in fact doesn't hinge on features of natural language indicative conditionals. One might treat the conditional as a material conditional and inter-substitutable with "Either Reagan will win or Anderson will win," and the case still seems to hold up. However, in what follows, I will work on the assumption that natural language indicatives are distinct from material conditionals

3 A Scorekeeping Framework (and Case 1)

I've now given three intuitive cases in which it seems plausible to say that explicitation is consequential. Having motivated this idea, let me now turn to the task of articulating a formal framework that enables us to accommodate it. This framework will involve a combination of ideas from the dynamic framework and the inferentialist framework. In dynamic fashion, the meaning of a sentence will be understood as its *context change potential*: a function mapping any context in which it might be uttered to the context that would result upon its being uttered. However, rather than taking a context to be a set of possible worlds, as standard in informational dynamic frameworks, a context will be a characterization of the "social deontic score," the set of moves a player has made and those to which they are committed. Brandom [1] in fact suggests just such a framework, though never formally spells it out:[9]

> The significance of an assertion of p can be thought of as a mapping that associates with one social deontic score—characterizing the stage before that speech act is performed, according to some scorekeeper—the set of scores for the conversational stage that results from the assertion, according to the same scorekeeper (1994, 190).

In this way, the framework is dynamic. In line with the proof-theoretic approach to inferentialism, however, the proof-theory (a particular sequent system) will govern the way in which contexts get updated. I have suggested that we can interpret a sequent of the form $\Gamma \vdash \varphi$ as saying that whenever one has *made* the moves in Γ, one is *committed* to φ. This distinction in status between the moves on the left-side of the turnstile and the moves on the right-side of the turnstile enables us to make sense of idea that, by actually *making* a move to which one was previously only *committed*, one comes to take on *new implicit commitments*. Thus far, I have left this

[9] See [19] for an earlier attempt to formalize Brandom's theory in a dynamic framework.

thought informal. Our scorekeeping framework will enable us to make this intuitive interpretation formally precise, enabling us to concretely model the way in which commitments are attributed to a speaker upon their making of various discursive moves. In the context of such a framework, a sequent of the form $\Gamma \vdash \varphi$ can be understood as a "scorekeeping principle," through which a scorekeeper determines the update in discursive score that is to take place upon some speaker's assertion or denial of some sentence. Let me now lay out the details.

For our purposes, we can take the "scorekeeping game" to involve just two participants, a designated scorekeeper, with a certain set of scorekeeping principles, and a designated move-maker, who makes the moves. We now define:

> **Scorecards:** A *scorecard* σ is a pair of sets of signed formulas of the form $\langle \sigma_m, \sigma_c \rangle$, where σ_m is the set of moves that a player has *made* and σ_c is the set of moves to which a player is *committed*.
>
> **Scorekeeping Principles:** A *scorekeeping principle* is a bilateral sequent of the form $\Gamma \vdash \varphi$, where Γ is a set of signed formulas and φ is a single signed formula.

The basic idea of commitment attribution is that we have a set of scorekeeping principles, and we can apply those scorekeeping principles to a set of moves that have been made in order to determine the committive consequences of those moves. Thus, where M is a set of moves, we can define the result of applying a set of scorekeeping principles π to M, arriving at the set of commitments we attribute to someone who has made those moves, as follows:

> **Application of Scorekeeping Principles (first pass):** The result of applying a set of scorekeeping principles π to a set of moves M, which we denote π(M) is the smallest superset of M such that, for every scorekeeping principle of the form $\Gamma \vdash \varphi \in \pi$ if $\Gamma \subseteq$ M then $\varphi \in \pi$(M).

Note this definition is neutral as to whether, when scorekeeping principles are applied, consequences are extracted from the set of moves one has *made* or the set of moves to which one is *committed*. However, interpreting $\Gamma \vdash \varphi$ as saying that one who has *made* the moves in Γ is *committed* to A, the set of moves to which we want to apply our scorekeeping principles, in updating our scorecard, is the set of moves *made*. Thus, we arrive at the following definition of updates:

Updates: The result of updating a scorecard σ with a move φ, $\sigma[\varphi]$, is $\langle \sigma_m \cup \{\varphi\}, \pi(\sigma_m \cup \{\varphi\}) \rangle$

So, when we update our scorecard upon someone's making the move φ, the set of moves *made*, in the updated scorecard, is simply be the set of moves made in the original scorecard plus the additional move φ, and the set of *commitments* is the result of applying the scorekeeping principles to this updated set of moves made.

Note that the result of updating some scorecard with the making of some move is always another scorecard. This lets us define, for a scorekeeper with a certain set of scorekeeping principles, the set of scorecards that are *possible*, relative to that player:

Possible scorecards: For a scorekeeper with any set of scorekeeping principles, π, the set of possible scorecards, relative to that scorekeeper, is recursively defined as follows:

1. $\langle \varnothing, \pi(\varnothing) \rangle \in \Sigma$
2. For any $\sigma \in \Sigma$, any move φ and $\sigma[\varphi] \in \Sigma$

And thus, the *discursive significance* of a move φ, relative to a scorekeeper with a certain set of scorekeeping principles, can be defined as a function that maps any possible scorecard they might have to the one that would result upon the move-maker's making the move φ. That is:

Discursive Significance: The *discursive significance* of a move φ, is a function $f : \Sigma \to \Sigma$ such that, for all $\sigma \in \Sigma$, $f(\sigma) = \sigma[\varphi]$

The discursive significance of a *sentence* can then be understood just as the pair consisting in the discursive significance of the assertion of that sentence along with the discursive significance of its denial.

This framework is clearly in the spirit of familiar dynamic frameworks briefly discussed in Section 1. It doesn't, however, impose the structure regarding Cumulative Transitivity that such frameworks impose. Before showing that officially, one modfication is in order. Though this framework does not validate Cumulative Transitivity, as things stand, it does not quite suffice to capture the structure of scorekeeping practices when it comes to *defeasible* attributions of commitments, as it validates Monotonicity. Consider, for instance, someone who asserts "Maddy's drinking a beer" and "Maddy's drinking an O'Doul's." Given that we have $+\langle \textbf{beer} \rangle \vdash +\langle \textbf{alcoholic} \rangle$ in our set of scorekeeping principles and $\sigma_m = \{+\langle \textbf{beer} \rangle, +\langle \textbf{O'Doul's} \rangle\}$, the condition for attributing commitment to $+\langle\textbf{alcoholic}\rangle$ is triggered by Application of Principles. However, we shouldn't attribute this commitment. So, we should modify our framework such that the attribution of a commitment to $+\langle\textbf{alcoholic}\rangle$, which is usually triggered upon a scorecard's containing $+\langle\textbf{beer}\rangle$, can be *defeated* by additional presence of $+\langle\textbf{O'Doul's}\rangle$. In general, if we have some scorekeeping principle of the form $\Gamma \vdash \varphi$, and a player has made the moves in Γ, the attribution of commitment to φ of to this player can be *defeated* if that player has also made the moves Δ, and we do not have a principle of the form $\Delta, \Gamma \vdash \varphi$. Moreover, such defeat can itself be defeated. Officially, we can revise Application of Scorekeeping Principles as follows:

> **Application of Scorekeeping Principles (revised):** The result of applying a set of scorekeeping principles π to a set of moves M, which we denote $\pi(\mathrm{M})$ is the smallest superset of M such that, for every scorekeeping principle of the form $\Gamma \vdash A \in \pi$ if $\Gamma \subseteq \mathrm{M}$ then either $A \in \pi(\mathrm{M})$ or the following conjunctive condition, the *defeat condition* holds (where Δ and Θ are non-empty):
>
> 1. There is some set of moves $\Delta \subseteq \mathrm{M}$ such that $\Delta, \Gamma \vdash \varphi \notin \pi$
> 2. There is no set of moves $\Theta \subseteq \mathrm{M}$ such that $\Theta, \Delta, \Gamma \vdash \varphi \in \pi$

Writing $\Gamma : \varphi$ to express an argument with premises Γ and conclusion φ, this definition let us define two notions of validity, relative to a set of scorekeeping principles π:

> **Strict Validity:** $\Gamma : \varphi$ is *strictly* valid, $\Gamma \vDash_s \varphi$, just in case, for *every* scorecard σ, $\varphi \in \sigma[\Gamma]_c$.

> **General Validity:** $\Gamma : \varphi$ is *generally* valid, $\Gamma \vDash_g \varphi$, just in case, for the *minimal* scorecard $\mu := \langle \varnothing, \pi(\varnothing)\rangle$, $\varphi \in \mu[\Gamma]_c$.

So $\Gamma : \varphi$ is *strictly* valid just in case someone who is scored as having made the moves in Γ, no matter which other moves they are scored as having made, is scored as committed to φ, whereas $\Gamma : \varphi$ is *generally* valid just in case someone who is scored as having made the moves in Γ, and is not previously scored as having made any moves, is scored as committed to φ. Thus, $+\langle\textbf{red}\rangle \vdash +\langle\textbf{colored}\rangle$ is strictly valid, since, no matter what other moves one has made, one who asserts "*a* is red" is committed to asserting "*a* is colored," whereas $+\langle\textbf{bird}\rangle \vdash +\langle\textbf{flies}\rangle$ is only generally valid, since someone who asserts "Bella's a bird," without having antecedently been scored as having made any moves, is committed to asserting "Bella flies," but someone who asserts "Bella's a bird" and "Bella's a penguin" is not.

We can now state a few basic facts about this framework regarding its ability to accommodate substructural relations of committive consequence. Let us say that a rule of the form:

$$\frac{\Gamma_1 : \varphi_1, \Gamma_2 : \varphi_2 \ldots \Gamma_n : \varphi_n}{\Delta : \psi}$$

holds with respect to strict/general validity just in case there is no set of scorekeeping principles π such that all of the premises are strictly/generally valid and the conclusion is not strictly/generally valid. Two claims are of note:

> **Proposition 1:** Monotonicity holds with respect to strict validity, but not general validity.
>
> *Proof:* For strict validity, suppose for contradiction $\Gamma \vDash_s \varphi$ but not $\Gamma, \psi \vDash_s \varphi$. So, for every set of moves M, $\varphi \in \pi(\mathrm{M} \cup \Gamma)$, and there is some set of moves M such that $\varphi \notin \pi(\mathrm{M}' \cup \Gamma \cup \{\psi\})$. But, of course, $\mathrm{M}' \cup \{\psi\}$ is some set of moves, and so $\varphi \in \pi(\mathrm{M}' \cup \Gamma \cup \{\psi\})$. Contradiction, so if $\Gamma \vDash_s \varphi$, then $\Gamma, \psi \vDash_s \varphi$. For general validity, just suppose $\Gamma \vdash \varphi \in \pi$ but $\Gamma, \psi \vdash \varphi \notin \pi$. Then $\varphi \in \pi(\Gamma)$, and so $\Gamma \vDash_g \varphi$, but $\varphi \notin \pi(\Gamma \cup \{\psi\})$, and so $\Gamma, \psi \nvDash_g \varphi$. □
>
> **Proposition 2:** Cumulative Transitivity holds with respect to neither general nor strict validity.
>
> *Proof:* For general validity, just suppose $\Gamma \vdash \varphi \in \pi$ and $\Gamma, \varphi \vdash \psi \in \pi$, but $\Gamma \vdash \psi \notin \pi$. Then $\varphi \in \pi(\Gamma)$, and so $\Gamma \vDash_g \varphi$ and $\psi \in \pi(\Gamma \cup \{\varphi\})$ so $\Gamma, \varphi \vDash_g \psi$, but $\psi \notin \pi(\Gamma)$ so $\Gamma \nvDash_g \psi$. For strict validity, construct the same example, supposing that, for all Δ, $\Delta, \Gamma \vdash \varphi \in \pi$ and $\Delta, \Gamma, \varphi \vdash \psi \in \pi$, but $\Gamma \vdash \psi \notin \pi$. □

The fact that Cumulative Transitivity holds for neither general nor strict validities is important, since, although the failure of Cumulative Transitivity in Case 1 involves the consideration of defeasible inferential relations, Case 2, which we'll turn to officially in the following section, involves only *strict* validities.

We can now precisely define the consequentiality of explicitation. Informally, the thought was that explicitation is consequential just in case making moves to which one is already scored committed can have an effect on the commitments one is attributed. Officially, now, this thought can be stated as follows:

The Consequentiality of Explicitation: Explicitation is *consequential* just in case, for some scorecard $\sigma \in \Sigma$ and move φ, $\varphi \in \sigma_c$ and $\sigma[\varphi]_c \neq \sigma_c$.

That is, explicitation is consequential, relative to some scorekeeper, if there's some possible scorecard σ such that φ is among the commitments of that scorecard, and updating σ with the move φ—making this commitment explicit—yields a scorecard containing a different set of commitments than σ. We can now explicitly note that this could be the case either if *new commitments are added* or if *existing commitments be subtracted*. These two possibilities correspond, respectively, to violations of the following two structural rules:[10]

$$\frac{\Gamma \vdash \varphi \quad \Gamma, \varphi \vdash \psi}{\Gamma \vdash \psi}\ {\scriptstyle \text{Cumulative Transitivity}} \qquad\qquad \frac{\Gamma \vdash \varphi \quad \Gamma \vdash \psi}{\Gamma, \varphi \vdash \psi}\ {\scriptstyle \text{Cautious Monotonicity}}$$

Though I've only given examples of the first sort of failure (and I'm not aware of any similarly intuitive examples in which Cautious Monotonicity fails), it is easy to see that this framework permits both possibilities for the consequentiality of explicitation.

For the first possibility, let us consider how this framework handles our first example of the consequentiality of explicitation. Let us suppose we have the following set of scorekeeping principles π:

$$\begin{array}{cc} +\langle \textbf{beer} \rangle \vdash +\langle \textbf{alcoholic} \rangle & +\langle \textbf{alcoholic} \rangle \vdash -\langle \textbf{O'Doul's} \rangle \\ +\langle \textbf{O'Doul's} \rangle \vdash +\langle \textbf{beer} \rangle & +\langle \textbf{beer} \rangle, +\langle \textbf{alcoholic} \rangle \vdash -\langle \textbf{O'Doul's} \rangle \\ +\langle \textbf{O'Doul's} \rangle \vdash -\langle \textbf{alcoholic} \rangle & +\langle \textbf{beer} \rangle, +\langle \textbf{O'Doul's} \rangle \vdash -\langle \textbf{alcoholic} \rangle \end{array}$$

Now, suppose we start with the empty scorecard $\langle \varnothing, \varnothing \rangle$ and then update it with $+\langle \textbf{beer} \rangle$. Our updated scorecard will be $\langle \varnothing \cup \{+\langle \textbf{beer} \rangle\}, \pi(\varnothing \cup \{+\langle \textbf{beer} \rangle\}) \rangle$, and so it will be $\langle \{+\langle \textbf{beer} \rangle\}, \{+\langle \textbf{beer} \rangle, +\langle \textbf{alcoholic} \rangle\} \rangle$. That is, the set of moves we score the speaker as having *made* is that of asserting **beer** and the set of moves to which we score the speaker as being *committed* is that of asserting **beer** and asserting **alcoholic**. If, then, the speaker goes on to *make* the move of asserting **alcoholic**, our updated scorecard will be $\langle \{+\langle \textbf{beer} \rangle\} \cup \{+\langle \textbf{alcoholic} \rangle\}, \pi(\{+\langle \textbf{beer} \rangle\} \cup \{+\langle \textbf{alcoholic} \rangle\}) \rangle$, and so it will be $\langle \{+\langle \textbf{beer} \rangle, +\langle \textbf{alcoholic} \rangle\}, \{+\langle \textbf{beer} \rangle, +\langle \textbf{alcoholic} \rangle, -\textbf{O'Doul's}\} \rangle$. In this way, the formal framework make precises the idea that by making *explicit* commitments that were previously only *implicit* one can come to take on *new implicit commitments*.

[10]See [3] for an informal statement of this fact.

It should be clear that this example is possibly only insofar as the set of scorekeeping principles is *not* closed under Cumulative Transitivity.

Now, as I mentioned above, I am not aware of any similarly intuitive examples in which Cautious Monotonicity fails.[11] Nevertheless, it is still worth noting explicitly how this formal framework permits the possibility of the other way in which explicitation may be consequential. Consider a set of scorekeeping principles π such that $\varphi \vdash \psi \in \pi$ and $\varphi \vdash \chi \in \pi$, but $\varphi, \psi \vdash \chi \notin \pi$. Suppose we start with the empty scorecard $\langle \emptyset, \emptyset \rangle$, and then update it with φ. Adding this to the moves made and applying our scorekeeping principles, we get $\langle \{\varphi\}, \{\varphi, \psi, \chi\} \rangle$. Now suppose we update this scorecard with ψ, which is contained among the commitments of this scorecard. Since $\varphi, \psi \vdash \chi \notin \pi$, the exception condition is triggered, and so the resulting scorecard is $\langle \{\varphi, \psi\}, \{\varphi, \psi\} \rangle$. Thus, by making the commitment to ψ explicit, actually making the move ψ, one *loses* the implicit commitment to χ. I will leave it as an open question whether there are any natural language examples of this sort.

4 A Bilateral Sequent Calculus (and Case 2)

This scorekeeping framework provides a determinate role for system of bilateral logic in the context of a broader inferentialist theory of discursive significance. The function of a logical system is to take a set of *base* scorekeeping principles, specifying relations of committive consequence for assertions and denials of *atomic* sentences, and extend this set of scorekeeping principles to a set specifying relations of committive consequence for assertions and denials of *logically complex* sentences. In the context of this scorekeeping framework, the base consequence relation will suffice to determines the discursive significance of all atomic sentences whereas the extended consequence relation will suffice to determine the discursive significance of the infinite number of logically complex sentences the language contains. Thus, for instance, given that we have $+\langle \textbf{beer} \rangle, +\langle \textbf{O'Doul's} \rangle \vdash -\langle \textbf{alcoholic} \rangle$ in the base set of scorekeeping principles, we should have $+\langle \textbf{beer} \wedge \textbf{O'Doul's} \rangle \vdash +\langle \neg\textbf{alcoholic} \rangle$, which, in the context of this scorekeeping framework, will make it such that, if someone asserts "Maddy's drinking a beer and she's drinking an O'Doul's" they are thereby scored as committed to asserting "Maddy's not drinking an alcoholic beverage."

[11]Hlobil [somewhere?] proposes an example of a failure of Cautious Monotonicity, however, the example relies on potentially contentious disjunction principles.

The desideratum for this system is that it be *classical*, in the sense defined above, while also allowing failures of both Montonicity or Cumulative Transitivity. Standard versions of bilateral logic do not meet this desideratum. As we've already seen, natural deduction systems typically impose such structural rules and so will not be suitable for our purposes. Beyond this, however, standard versions of bilateral logic impose certain *coordination principles*, distinctively bilateral structural rules which "coordinate" the opposite signs expressing assertion and denial, and these also impose structure on consequence that is problematic in the current setting. In Rumfitt's [24] original formulation of bilateral logic, widely followed by bilateral logicians such as Incurvati and Schloder [14], we have two basic coordination principles. Where φ and ψ are signed formulas, and starring a signed formula yields the oppositely signed formula, these two basic principles can be put as follows:

$$\frac{\Gamma \vdash \varphi \quad \Delta \vdash \varphi^*}{\Gamma, \Delta \vdash \bot}\ \text{Incoherence} \qquad\qquad \frac{\Gamma, \varphi \vdash \bot}{\Gamma \vdash \varphi^*}\ \text{Reductio}$$

Incoherence says that if making all the moves in Γ commits one to making some move φ (asserting or denying some sentence) and making all the moves in Δ commits one to the opposite move φ^* (denying or asserting that sentence), then making all the moves in Γ along with all the moves in Δ is incoherent. Reductio says that if Γ along with making some φ is incoherent, then Γ commits one to making the opposite move φ^*. While these coordination principles might seem plausible, given Reflexivity (which simply maintains that making some move commits one to that very move), they enable one to derive Cumulative Transitivity as follows:

$$\dfrac{\Gamma \vdash \varphi \qquad \dfrac{\dfrac{\Gamma, \varphi \vdash \psi \qquad \overline{\psi^* \vdash \psi^*}\ \text{Reflex.}}{\Gamma, \varphi, \psi^* \vdash \bot}\ \text{Inc.}}{\Gamma, \psi^* \vdash \varphi^*}\ \text{Reduc.}}{\dfrac{\Gamma, \psi^* \vdash \bot}{\Gamma \vdash \psi}\ \text{Reduc.}}\ \text{Inc.}$$

Note, I here use the *mixed* context version of the Incoherence rule, as, though they are not explicit about it, this is the rule bilateralists such as Incurvati and Schloder appeal to in develop an inferentialist theory of meaning in a bilateralist setting.[12] One might wonder about whether this issue could be avoided by adopting the *shared* context version of the Incoherence rule:

[12] See

$$\frac{\Gamma \vdash \varphi \quad \Gamma \vdash \varphi^*}{\Gamma \vdash \bot} \text{ Incoherence (shared context)}$$

However, if consider our second case, it is easy to see that there is a direct problem with this rule as well, at least in combination with Reductio. Letting $\mathbf{c}m$ be "The moon is made of cheese," $\mathbf{f}s$ be "Superman flies," and $\mathbf{f}c$ be "Clarke Kent flies," we can see that even the shared context version of the incoherence rule, in conjunction with the Reductio enables us to derive the problematic conclusion of the Superman Triad:

$$\dfrac{\dfrac{+\langle \mathbf{f}s\rangle, -\langle \mathbf{f}c\rangle, -\langle \mathbf{c}m\rangle \vdash +\langle \mathbf{f}s\rangle \quad +\langle \mathbf{f}s\rangle, -\langle \mathbf{f}c\rangle, -\langle \mathbf{c}m\rangle \vdash -\langle \mathbf{f}s\rangle}{+\langle \mathbf{f}s\rangle, -\langle \mathbf{f}c\rangle, -\langle \mathbf{c}m\rangle \vdash}}{+\langle \mathbf{f}s\rangle, -\langle \mathbf{f}c\rangle \vdash +\langle \mathbf{c}m\rangle} \text{ Reduc.}$$

Thus, given that asserting "Superman flies," denying "Clarke Kent flies," and denying "The moon is made of cheese" is incoherent (this set of moves commits one to asserting "Superman flies" and to denying "Superman flies"), Reductio lets us conclude that asserting "Superman flies" and denying "Clarke Kent flies" commits one to asserting "The moon is made of cheese." But, when we consider the case of Lois, that's just the conclusion that we want to resist.

With all these considerations in view, it seems clear that, if we want a bilateral logic to that enables us to accommodate these cases of the consequentiality of explicitation, we can retain neither the operational rules of standard natural deduction systems nor can we retain the coordination principles of standard bilateral systems. One might wonder whether a classical bilateral system that meets these desiderata is even possible. Simonelli [30] shows that, if we switch from a bilateral *natural deduction* system, which contains (positive and negative) introduction and elimination rules to a bilateral *sequent calculus*, which contains (positive and negative) left and right rules, it is. In particular, the following sequent calculus, defended at length by Simonelli, does the job perfectly:[13]

$$\frac{}{\Gamma, \varphi \vdash \varphi} \text{ Containment} \qquad\qquad \frac{}{\Gamma, \varphi, \varphi^* \vdash \psi} \text{ B. Explo.}$$

where Γ contains only atomics and φ and ψ are atomic

[13] See also [12], [15], [16], and [13] for related unilateral sequent calculi and developments of this approach for accommodating defeasible inferences in the context of a sequent calculus.

$$\frac{\Gamma \vdash -\langle A \rangle}{\Gamma \vdash +\langle \neg A \rangle} +_{\neg_R} \qquad \frac{\Gamma \vdash +\langle A \rangle}{\Gamma \vdash -\langle \neg A \rangle} -_{\neg_R} \qquad \frac{\Gamma, -\langle A \rangle \vdash \varphi}{\Gamma, +\langle \neg A \rangle \vdash \varphi} +_{\neg_L} \qquad \frac{\Gamma, +\langle A \rangle \vdash \varphi}{\Gamma, -\langle \neg A \rangle \vdash \varphi} -_{\neg_L}$$

$$\frac{\Gamma \vdash +\langle A \rangle \quad \Gamma \vdash +\langle B \rangle}{\Gamma \vdash +\langle A \wedge B \rangle} +_{\wedge_R} \qquad \frac{\Gamma, +\langle A \rangle \vdash -\langle B \rangle}{\Gamma \vdash -\langle A \wedge B \rangle} -_{\wedge_{R_1}} \qquad \frac{\Gamma, +\langle B \rangle \vdash -\langle A \rangle}{\Gamma \vdash -\langle A \wedge B \rangle} -_{\wedge_{R_2}}$$

$$\frac{\Gamma, +\langle A \rangle, +\langle B \rangle \vdash \varphi}{\Gamma, +\langle A \wedge B \rangle \vdash \varphi} +_{\wedge_L} \qquad \frac{\Gamma, -\langle A \rangle \vdash \varphi \quad \Gamma, -\langle B \rangle \vdash \varphi}{\Gamma, -\langle A \wedge B \rangle \vdash \varphi} -_{\wedge_L}$$

Note that rather than containing the standard axiom of Reflexivity, $\varphi \vdash \varphi$, along with the structural rule of Monotonicity, this sequent calculus features the axiom of *contexted* Reflexivity, or *Containment*, $\Gamma, \varphi \vdash \varphi$, which says that, given any set of assertions and denials Γ, making some move commits one to making that very move. It is easy to see that this is strictly valid. Additionally, this sequent calculus contains an axiom of *bilateral explosion*, which says that asserting and denying the very same sentence, commits one to everything. While this principle is potentially contentious, it is a requirement for classicality.[14] Now, it is important that these axiom schemas are restricted to *atomic* signed formulas, since the basic application of these sequent calculus is to take a set of atomic scorekeeping principles encoding material relations of committive consequence, along with these logical axioms, and extend this set of scorekeeping principles relating assertions and denials of atomic to a set of scorekeeping relating assertions and denials of logically complex sentences.

Generalizing the notion of classicality stated above to a bilateral context, let us say that a bilateral argument of the form $\Gamma : \varphi$ is *classically valid* just in case there is no classical valuation v such that all positively signed sentences in Γ are true, all of the negatively signed sentences in Γ are false, and A is false if φ is of the form $+\langle A \rangle$ or A is true if φ is of the form $-\langle A \rangle$. Letting all sets of scorekeeping principles be closed under the rules of this sequent calculus, we can now state the following:

> **Proposition 3:** Any classical validity is strictly valid.
>
> *Proof:* As shown in [31], this sequent calculus is complete with respect to classical validity: for any classically valid argument $\Gamma : \varphi$, $\Gamma \vdash \varphi$ has a proof from instances of Containing. So, if $\Gamma : \varphi$ is classically valid, then any set of scorekeeping principles π contains $\Gamma \vdash \varphi$, and, moreover, contains $\Delta, \Gamma \vdash \varphi$ for any Δ since, if $\Gamma : \varphi$ is classically valid, then $\Delta, \Gamma : \varphi$ is classically valid. Accordingly, if $\Gamma : \varphi$ is classically valid, then for every set of moves M, $\varphi \in \pi(\mathrm{M} \cup \Gamma)$, and so $\Gamma \vDash_s \varphi$. □

[14] In fact, if one keeps this sequent calculus and drops this principle, one gets a bilateral system for the paraconsistent logic LP.

We thus obtain a classical system without structural rules of Monotonicity, Cumulative Transitivity, or any bilateral coordination principles.

Having laid out a logic system capable of accommodating failures of Cumulative Transitivity, let us return to our second example of the consequentiality of explicitation, and show how we can officially capture the intuitions spelled out above in this framework. The fact that "Superman" and "Clarke Kent" are inferentially equivalent, from our perspective, can be captured with the following rules, under which our set of scorekeeping principles are closed:

$$\frac{\Gamma \vdash \varphi(s)}{\Gamma \vdash \varphi(c)}\ {\scriptstyle s=c} \qquad \frac{\Gamma \vdash \varphi(c)}{\Gamma \vdash \varphi(s)}\ {\scriptstyle s=c} \qquad \frac{\Gamma, \varphi(s) \vdash \psi}{\Gamma, \varphi(c) \vdash \psi}\ {\scriptstyle s=c} \qquad \frac{\Gamma, \varphi(c) \vdash \psi}{\Gamma, \varphi(s) \vdash \psi}\ {\scriptstyle s=c}$$

where no substitutions can be made in the scope of intensional operators.

Given our definition of discursive significance, the result is that, from our perspective, given *our* scorekeeping principles, "Superman flies" and "Clarke Kent flies" have the very same discursive significance—one takes on just the same commitments in asserting or denying either sentence. However, such principles will *not* be operative in the scorekeeping principles of Lois Lane. Accordingly, we can say that, from *Lois's* perspective, given *her* scorekeeping principles, "Superman flies" and "Clarke Kent flies" have distinct discursive significances—for instance, in uttering the latter but not the former one takes on a commitment to "The Man of Steel flies." A full response to Frege's puzzle would require a full account of the discursive significance of proper names and belief reports, and such an account is beyond the scope of the present paper. However, what we have on the table at the moment is enough to officially show how this framework handles our second case of the consequentiality of explicitation.

Let us first show officially that asserting "Superman flies" and denying "Clarke Kent flies" commits one to asserting "Superman flies and doesn't fly":

$$\dfrac{\dfrac{}{+\langle \mathbf{f}s\rangle, -\langle \mathbf{f}c\rangle \vdash +\langle \mathbf{f}s\rangle}\ \text{Co.} \qquad \dfrac{\dfrac{\dfrac{}{+\langle \mathbf{f}s\rangle, -\langle \mathbf{f}c\rangle \vdash -\langle \mathbf{f}(c)\rangle}\ \text{Co.}}{+\langle \mathbf{f}s\rangle, -\langle \mathbf{f}c\rangle \vdash +\langle \neg\mathbf{f}c\rangle}\ +_{\neg_R}}{+\langle \mathbf{f}s\rangle, -\langle \mathbf{f}c\rangle \vdash +\langle \neg\mathbf{f}s\rangle}\ {\scriptstyle s=c}}{+\langle \mathbf{f}s\rangle, -\langle \mathbf{f}c\rangle \vdash +\langle \mathbf{f}s \wedge \neg\mathbf{f}s\rangle}\ +_{\wedge_R}$$

Let us now show that asserting "Superman flies," denying "Clarke Kent flies," and asserting "Superman flies and doesn't fly" commits one to asserting "The moon is made of cheese," which we already know that the system proves on account of its classicality:[15]

$$\dfrac{\dfrac{\dfrac{}{+\langle \mathbf{f}\rangle, -\langle \mathbf{f}c\rangle, +\langle \mathbf{f}s\rangle, -\langle \mathbf{f}s\rangle \vdash +\langle \mathbf{c}m\rangle}\,\text{B. Explo.}}{+\langle \mathbf{f}\rangle, -\langle \mathbf{f}c\rangle, +\langle \mathbf{f}s\rangle, +\langle \neg\mathbf{f}s\rangle \vdash +\langle \mathbf{c}m\rangle}\,+_{\neg_L}}{+\langle \mathbf{f}s\rangle, -\langle \mathbf{f}c\rangle, +\langle \mathbf{f}s \wedge \neg\mathbf{f}s\rangle \vdash +\langle \mathbf{c}m\rangle}\,+_{\wedge_L}$$

And yet, though we have the scorekeeping principles $+\langle \mathbf{f}s\rangle, -\langle \mathbf{f}c\rangle \vdash +\langle \mathbf{f}s \wedge \neg\mathbf{f}s\rangle$ and $+\langle \mathbf{f}s \wedge \neg\mathbf{f}s\rangle \vdash +\langle \mathbf{c}m\rangle$, our scorekeeping principles are not closed under Cumulative Transitivity, and so we *don't* get $+\langle \mathbf{f}\rangle, -\langle \mathbf{f}c\rangle \vdash +\langle \mathbf{c}m\rangle$, that asserting "Superman flies" and denying "Clarke Kent flies" commits one to asserting "The moon is made of cheese."

Applying our scorekeeping framework to make sense of this case, when Lois asserts "Superman flies" and denies "Clarke Kent flies," we take her to be committed to asserting the contradiction, "Superman flies and doesn't fly." There is a crucial distinction, however, between her being *committed* to asserting "Superman flies and doesn't fly," given the moves she's made, and her actually *making* this assertion. If one is committed to an assertion, then, if this commitment is appropriately brought out in the course of discourse, one is obliged to make (and defend) this assertion *or* retract some of the moves one has made so that one no longer is so committed. In the case of a commitment to the assertion of a *contradictory* assertion, it is clearly the second option that one must go in for, rectifying the incoherence in one's moves by retracting some of them. For instance, in the case of Lois, upon finding out that Superman and Clarke Kent are the same person and so she has committed herself to asserting contradictory sentence "Superman flies and doesn't flies," the obvious thing to do is to retract her denial of "Clarke Kent flies." The intuitive thought underlying the formalism here is that if instead she just goes ahead explicitly asserts the contradiction "Superman flies and doesn't fly," then it's reasonable for us to throw up our hands and say we don't know how to score her, and in *this* case (and *only* this case) does it make sense to simply score her as committed to everything.

[15]Note that the starting sequent here says that asserting "Superman flies" and denying "Superman flies" commits one to asserting "The moon is made of cheese." Crucially, here, it is the *same sentence* both asserted and denied, and this is distinct from asserting "Superman flies" and denying the inferentially equivalent sentence "Clarke Kent flies."

5 An Epistemic Conditional (and Case 3)

I have an introduced a logical system that enables us to model the discursive significance of sentences containing conjunctions and negations, such as those featured in our second example. Our third case, however, features conditional sentences. So it is time to introduce conditionals into our framework. Let us note first that it is straightforward to introduce a material conditional $\supset$ (such that $A \supset B$ is equivalent to $\neg(A \wedge \neg B)$) with rules of the same form as the conjunction operator:

$$\frac{\Gamma, +\langle A\rangle \vdash +\langle B\rangle}{\Gamma \vdash +\langle A \supset B\rangle}\ +_{\supset_{R_1}} \qquad \frac{\Gamma, -\langle B\rangle \vdash -\langle A\rangle}{\Gamma \vdash +\langle A \supset B\rangle}\ +_{\supset_{R_2}} \qquad \frac{\Gamma, -\langle A\rangle \vdash \varphi \quad \Gamma, +\langle B\rangle \vdash \varphi}{\Gamma, +\langle A \supset B\rangle \vdash \varphi}\ +_{\supset_L}$$

$$\frac{\Gamma \vdash +\langle A\rangle \quad \Gamma \vdash -\langle B\rangle}{\Gamma \vdash -\langle A \supset B\rangle}\ -_{\supset_R} \qquad \frac{\Gamma, +\langle A\rangle, -\langle B\rangle \vdash \varphi}{\Gamma, -\langle A \supset B\rangle \vdash \varphi}\ -_{\supset_L}$$

While a material conditional surely can be introduced into the formal system in this way, and indeed the third case of the consequentiality of explicitation can be straightforwardly dealt with by introducing a such a conditional, there are some reasons to want to introduce a conditional of a different sort. Note, for instance, given that we take it that asserting "Superman flies" commits one to asserting "Clarke Kent flies," these rules would mean that we consider everyone to be committed to asserting "If Superman flies, then Clarke Kent flies." While it is surely possible to get into a mindset where this seems right, we may plausibly want to attribute conditional commitments to individuals in accordance with *their* scorekeeping principles, rather than *our* scorekeeping principles. That is, we may want to introduce a conditional operator that does not have the result that Lois is committed to to "If Superman flies, then Clarke Kent flies," since she doesn't have the scorekeeping principle of taking someone who asserts "Superman flies" to be committed to asserting "Clarke Kent flies."

Now, just as I did not put forward a full theory of proper names to sketch the solution to Frege's puzzle above, I will not put forward a full theory of conditionals here. There are a number of conditionals that can be defined in this framework, and I will just put forward a simple proposal here to illustrate one possibility that enables us to accommodate our final case of the consequentiality of explicitation having to do with indicative conditionals. The proposal is to simply identify the conditions under which one is committed to asserting or denying a conditional as follows:[16]

[16]Compare the dynamic proposal of [36, 242-247] and the inferentialist proposal of [14], both of

Epistemic Conditional: A speaker s with a set of scorekeeping principles π and self-kept scorecard σ is committed to

- $+\langle A \to B\rangle$ just in case $+\langle B\rangle \in \pi(\sigma_m \cup \{+\langle A\rangle\})$
- $-\langle A \to B\rangle$ just in case $-\langle B\rangle \in \pi(\sigma_m \cup \{+\langle A\rangle\})$

Thus, one is committed to asserting a conditional just in case they take that, were they to assert the antecedent, they'd be committed to asserting the consequent, and one is committed to denying a conditional just in case they take it that, were they to assert the antecedent, they would be committed to denying the consequent. On this way of defining them, conditional commitments are *epistemic*, reflecting *one's own* commitments and scorekeeping principles. Thus, though *we* are committed to the conditional, "If Clarke Kent is wearing a purple shirt, then Superman is wearing a purple shirt," given *our* scorekeeping principles, *Lois* is not committed to this conditional, given *her* scorekeeping principles. On the other hand, she is of course, committed to the conditional, "If Clarke is wearing a purple shirt, he's wearing a colored shirt."

To see how this conditional enables us to accommodate our final case of the consequentiality of explicitation, let Γ be the set of assertions sufficient to characterize the polling data, and let us hypothetically consider our own conditional commitments. Since the polling data shows Reagan to be decisively ahead, we have $\Gamma \vdash +\langle\textbf{republican}\rangle$ in our set of scorekeeping principles π, and so given our acceptance of the polling data, we're committed to asserting "A republican will win the election." Now, given that Anderson is the only republican running other than Reagan, we have $\Gamma, +\langle\textbf{republican}\rangle, -\langle\textbf{Reagan}\rangle \vdash +\langle\textbf{Anderson}\rangle$, and thus, given our negation rules, $\Gamma, +\langle\textbf{Republican}\rangle, +\langle\neg\textbf{Reagan}\rangle \vdash +\textbf{Anderson}$. So, $+\langle\textbf{Anderson}\rangle \in \pi(\Gamma, +\textbf{republican}, +\neg\textbf{Reagan}\rangle$. Accordingly, if we *were* to explicitly assert "A republican will win the election," we would thereby be committed to asserting "If it's not Reagan who wins, it will be Anderson." However, simply asserting the polling data, we are not committed to asserting "If it's not Reagan who wins, it will be Anderson." In this way, the example is an instance of the consequentiality of explicitation.

Now, McGee presents his examples as failures of modus ponens, and it's worth

which exploit bilateralism in a similar way. There is obviously room for disputes and refinements (for instance, adding to the clause a presupposition that there be no φ such that $\varphi \in \pi(\sigma_m \cup \{+\langle A\rangle\})$ and $\varphi^* \in \pi(\sigma_m \cup \{+\langle A\rangle\})$). I present this clause just as a simple way to spell out the basic approach.

saying a word on how this framework both does not validate modus ponens, depending on how one formulates the rule. In the absence of Transitivity, we can distinction between to different kinds of modus ponens principles, known respectively, as *internal* modus ponens, and *external* modus ponens. On this interpretation here, internal modus ponens can be understood as saying that, no matter what other moves one has made, if one *asserts* A and also *asserts* $A \to B$, then one is committed to asserting B. Given that asserting $A \to B$ expresses that, were one to assert A, they'd be committed to asserting B, clearly internal modus ponens holds. External modus ponens, on the other hand, says that, if given the moves one has made, one is *committed* to asserting A, and one is also *committed* to asserting $A \to B$, then one is committed to asserting B. Now, in McGee's original example, accepting the polling data, one is committed to asserting "A republican will win the election," and one is also committed to asserting, "If a republican will win the election, then, if it's not Reagan, it will be Anderson," but one is not committed to asserting "If it's not Regan, it will be Anderson." On this framework, we can accept this example as a counter-example to *external* modus ponens, while maintaining the validity of *internal* modus ponens.

To formally implement these informal claims into the scorekeeping framework that we've been introduced, we need to connect this conditional that we've defined (through which one expresses *one's own* scorekeeping commitments) to the notion of validity that we've defined (which pertains to our attribution of commitments *to others*). To do this, we must expand our scorekeeping framework such that, in addition to the framework involving our scorecard and scorekeeping principles, it must also involve a model of our representation of the possible scorecards the other player might have, keeping track of their own commitments, and the possible sets of scorekeeping principles they might have through which they take themselves to undertake commitments. So, lets the define the following:

> **Scorekeeping possibility:** A *scorekeeping possibility* for the other player, θ, be a pair $\langle \sigma^*, \pi^* \rangle$ consisting in a (representation of a) self-kept scorecard σ^* and a set of scorekeeping principles π^* (where $\sigma^* := \langle {\sigma^*}_m, \pi^*({\sigma^*}_m) \rangle$). Θ is the set of live scorekeeping possibilities.

We will assume that each set of scorekeeping principles in each scorekeeping possibility is closed under our logical rules. We will also assume that if we score the player as having made a move, they score themself as having made the move as well (and so being committed to its consequences, given their scorekeeping principles):

Self-Scoring 1: If $\varphi \in \sigma_m$, then for all $\langle \sigma^*, \pi^* \rangle$ in Θ: $\varphi \in {\sigma^*}_m$.

Additionally, we will assume that if we take the player to score themself as being committed to a move, we score them to be committed to that move as well:

Self-Scoring 2: If $\varphi \in {\sigma^*}_c$ for all $\langle \sigma^*, \pi^* \rangle$ in Θ, then $\varphi \in \sigma_c$.

Now, an assertion of a conditional filters the set of live scorekeeping possibilities as follows:

Conditional Updates: $\Theta[+\langle A \rightarrow B \rangle] = \{\langle \sigma^*, \pi^* \rangle \in \Theta \mid +\langle B \rangle \in \pi^*({\sigma^*}_m \cup \{+\langle A \rangle\})$

We can now state and prove the following

Proposition 4: Internal modus ponens is strictly valid: $+\langle A \rightarrow B \rangle, +\langle A \rangle \vDash_s +\langle B \rangle$

Proof: Take a player with an arbitrary set of moves made and commitments and suppose they assert $+\langle A \rightarrow B \rangle$. By the conditional update rule, we filter our set of scorekeeping possibilities Θ such that, for all $\langle \sigma^*, \pi^* \rangle \in \Theta : +\langle B \rangle \in \pi^*({\sigma^*}_m \cup \{+\langle A \rangle\}$. Now, suppose they assert $+\langle A \rangle$. Given, Self-Scoring 1, for all $\langle \sigma^*, \pi^* \rangle$ in Θ: $+\langle A \rangle \in {\sigma^*}_m$, and, so given the filter on Θ imposed by the update with $+\langle A \rightarrow B \rangle$, for all $\langle \sigma^*, \pi^* \rangle$ in Θ: $+\langle B \rangle \in {\sigma^*}_c$. Given Self-Scoring 2, $+\langle B \rangle \in \sigma_c$. Since this player had arbitrary set of moves made, $+\langle A \rightarrow B \rangle, +\langle A \rangle \vDash_s +\langle B \rangle$. □

We can also state and prove the following:

Proposition 5: External modus ponens:

$$\frac{\Gamma : +\langle A \rangle \quad \Gamma + \langle A \rightarrow B \rangle}{\Gamma : +\langle B \rangle}$$

holds for neither general nor strict validities.

Proof: Follows straightforwardly from the fact that Cumulative Transitivity holds for neither general nor strict validities (Proposition 2). Suppose we and the player on which we're keeping score have the same set of scorekeeping principles, and suppose these include $\Gamma \vdash +\langle A \rangle$ (and $\Delta, \Gamma \vdash +\langle A \rangle$ for all Δ), and $\Gamma, +\langle A \rangle \vdash +\langle B \rangle$ (and $\Delta, \Gamma, +\langle A \rangle \vdash +\langle B \rangle$ for all Δ), but not $\Gamma \vdash +\langle B \rangle$. It follows that $\Gamma \vDash_s +\langle A \rangle$ and $\Gamma \vDash_s +\langle A \rightarrow B \rangle$, but not $\Gamma \vDash_g +\langle B \rangle$. □

The fact that external modus ponens and Cumulative Transitivity are so tightly linked in this way should not come as a suprise.[17] McGee's original example, a counterexample to external modus ponens is the failure of Cumulative Transitivity conditionalized, and, in fact, in fact, for any of our examples of failures of Cumulative Transitivity can be conditionalized to be transformed into failures of external modus ponens. Thus, we may transform Case 1 and Case 2 as follows:

The Maddy Triad, conditionalized:

1. Asserting "Maddy's drinking a beer" commits one to asserting "Maddy's drinking an alcoholic beverage."
2. Asserting "Maddy's drinking a beer" commits one to asserting "If Maddy's drinking an alcoholic beverage, then she's not drinking an O'Doul's."
3. However, asserting "Maddy's drinking a beer" by itself does not commit one to "Maddy's not drinking an O'Doul's."

The Superman Triad, conditionalized:

1. Asserting "Superman flies" along with denying "Clark Kent flies" commits one to asserting "Superman flies and doesn't fly."
2. Asserting "Superman flies" along with denying "Clarke Kent flies," commits one to asserting "If Superman flies and doesn't fly, then the moon is made of cheese."
3. However, asserting "Superman flies" along with denying "Clark Kent flies" doesn't, by itself, commit one to asserting "The moon is made of cheese."

Insofar as the assertion of a conditional expresses commitment to the consequent, conditional upon assertion of the antecedent, we should expect just such a correspondence between failures of Cumulative Transitivity and External Modus Ponens. We can accept these failures while still maintaining that modus ponens, understood as the principle that asserting a conditional and asserting its antecedent commits one to its consequent, is strictly valid.

[17] Though it is not surprising once pointed out, I am indebted to Ulf Hlobil and Matthew Fox Burley [4] for originally pointing it out to me.

6 Conclusion

The guiding thought of this paper has been that *making one's commitments explicit* can, in some cases, result in one's taking on *new implicit commitments*. Surveying both dynamic-semantic and inferentialist traditions, we saw that their core frameworks typically build in structural principles—above all, Cumulative Transitivity—that preclude this possibility. I have presented three different cases to illustrate that this possibility is sometimes actual; explicitation is sometimes consequential. To formally explicate this idea of the consequentiality of explicitation and accommodate this motivating data, I have put forward a novel dynamic inferentialist framework in which contexts are modeled not as sets of worlds but as scorecards tracking both moves actually made and commitments thereby incurred. Beyond accommodating the data, this framework yields a new fundamentally dynamic conception of the nature of consequence in natural language, challenging the very idea of *closing* a set of claims under their consequences, since adding the consequences of one's assertions to those assertions can result in *new* consequences.

References

[1] Robert Brandom. *Making It Explicit: Reasoning, Representing, and Discursive Commitment*. Harvard University Press, Cambridge, Mass., 1994.

[2] Robert Brandom. *Articulating Reasons: An Introduction to Inferentialism*. Harvard University Press, Cambridge, 2000.

[3] Robert Brandom. From logical expressivism to expressivist logic: Sketch of a program and some implementations1. *Philosophical Issues*, 28(1):70–88, 2018.

[4] Mathew Fox Burley. Indicative conditionals and non-transitivity. Master's thesis, Concordia University, 2021.

[5] Sam Carter. Keeping vague score. *Journal of Philosophy*, forthcoming.

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

[7] Pedro del Valle-Inclan. Harmony and normalisation in bilateral logic. *Bulletin of the Section of Logic*, 52(3):377–409, 2023.

[8] Nissim Francez. Bilateralism in proof-theoretic semantics. *Journal of Philosophical Logic*, (2-3):1–21, 2013.

[9] Gottlob Frege. Sense and reference. *Philosophical Review*, 57(3):209–230, 1948.

[10] Gerhard Gentzen. Untersuchungen Über das logische schließen. i. *Mathematische Zeitschrift*, 35:176–210, 1935.

[11] Simon Goldstein. A theory of conditional assertion. *Journal of Philosophy*, 116(6):293–318, 2019.

[12] Ulf Hlobil. A nonmonotonic sequent calculus for inferentialist expressivists. In Pavel Arazim and Michal Danč'ak, editors, *The Logica Yearbook 2015*, pages 87–105. College Publications, 2016.

[13] Ulf Hlobil and Robert Brandom. *Reasons for Logic, Logic for Reasons: Pragmatics, Semantics, and Conceptual Roles*. Routledge, New York, Forthcoming.

[14] Luca Incurvati and Julian J. Schlöder. *Reasoning with Attitude*. Oxford University Press USA, New York, 2023.

[15] Daniel Kaplan. A multi-succident sequent calculus for logical expressivists. In P. Arazim and T. Lákiva, editors, *The Logica Yearbook 2017*, pages 139–154. College Publications, 2018.

[16] Daniel Kaplan. *Substructural Content*. PhD thesis, University of Pittsburgh, 2021.

[17] Matthew Mandelkern. Dynamic non-classicality. *Australasian Journal of Philosophy*, 98(2):382–392, 2020.

[18] Vann McGee. A counterexample to modus ponens. *Journal of Philosophy*, 82(9):462–471, 1985.

[19] Bernhard Nickel. Dynamics, brandom-style. *Philosophical Studies*, 162(2):333–354, 2013.

[20] Dag Prawitz. *Natural Deduction: A Proof-Theoretical Study*. Dover Publications, Mineola, N.Y., 1965.

[21] Greg Restall. Multiple conclusions. In Petr Hájek, Luis Valdés-Villanueva, and Dag Westerstahl, editors, *Logic, Methodology and Philosophy of Science*. College Publications, 2005.

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

[23] David Ripley. On the 'transitivity' of consequence relations. *Journal of Logic and Computation*, 28(2):433–450, 12 2017.

[24] Ian Rumfitt. "yes" and "no". *Mind*, 109:781–823, 2000.

[25] Wilfrid Sellars. Inference and meaning. *Mind*, 62(247):313–338, 1953.

[26] Wilfrid Sellars. Some reflections on language games. *Philosophy of Science*, 21(3):204–228, 1954. Reprinted in Sellars, 1963, *Science, Perception, and Reality*, 321–358.

[27] Wilfrid Sellars. Meaning as functional classification. *Synthese*, 27(3):417–437, 1974.

[28] Ryan Simonelli. "Yes," "no," neither, and both. manuscript, available at ryansimonelli.com/research.

[29] Ryan Simonelli. Considering the exceptions: On the failure of cumulative transitivity for indicative conditionals. *Synthese*, 200(5):1–21, 2022.

[30] Ryan Simonelli. Bringing bilateralisms together. *Australasian Journal of Philosophy*, forthcoming.

[31] Ryan Simonelli. Generalized bilateral harmony. In *The 2023 Logica Yearbook*. College Publications, Forthcoming.

[32] Martin Stokhof, Jeroen Groenendijk, and Frank Veltman. Coreference and modality. In Shalom Lappin, editor, *The handbook of contemporary semantic theory*, pages 179–216. Blackwell Reference, 1996.

[33] Neil Tennant. Natural deduction and sequent calculus for intuitionistic relevant logic. *Journal of Symbolic Logic*, 52(3):665–680, 1987.

[34] Frank Veltman. Defaults in update semantics. *Journal of Philosophical Logic*, 25(3):221–261, 1996.

[35] Malte Willer. Two puzzles about ability can. *Linguistics and Philosophy*, 44(3):551–586, 2021.

[36] Malte Willer. Negating conditionals. In *Oxford Studies in Philosophy of Language Volume 2*. Oxford University Press, 02 2022.

[37] Malte Willer. Credences for strict conditionals. *Philosophy and Phenomenological Research*, 110(1):23–50, 2025.

[38] Ludwig Wittgenstein. *Philosophical Investigations*. Wiley-Blackwell, 4th edition, 2009. Originally published in 1953.