4

Discursive Role Semantics

4.1 Introduction

In the previous two chapters, I argued against wordly semantics in its two most prominent forms—what I called "extra-worldly semantics" and "intra-worldly" semantics. I claimed that worldly semantic theories of both sorts are not able to explain our knowledge of meaning because the worldly knowledge to which they appeal can only be understood as depending on our knowledge of meaning. I now turn to the positve task of explicating the sort of semantic theory that can do what worldly semantics cannot: explain our knowledge of meaning and, along with it, our knowledge of the "worldly" entities, such as possible worlds and properties, to which worldly semantic theories centrally appeal. This "worldly" knowledge, on the account I develop, is conceived of as nothing other than semantic knowledge, transposed in the worldly mode. The task of this chapter it to lay out the altnerate, non-worldly semantic theory—based on the semantic theory proposed by Sellars and developed by Robert Brandom (1994, 2000)—that sets the ground for this account of "worldly" knowledge put forward in the following chapter. On this semantic theory, which I call "discursive role semantics," the meaning of an expression is understood directly in terms of its role in discourse, rather than this role being understood in terms of the sentence's having the worldly meaning that it does.

4.2 A Different Kind of Semantic Theory

Discursive role semantics is an alternative to truth-conditional semantics. As such, perhaps the best way of introducing it is to introduce it as a member of a wider class of alternatives to truth-conditional semantic theories that have gained some traction in the past few decades: *dynamic* semantic theories. A dynamic semantic theory is one in which, rather thinking of the meaning of a sentence in terms of the conditions under which it is true, we think of the meaning of a sentence in terms of its potential, when employed in a given context, to change (or "update") that context. In a slogan, the meaning of a sentence is its context change potential. The meaning of a subsentential expression is the contribution that it makes to the context change potential of sentences in which it can occur.

On a standard dynamic theory, we think of the contexts that get updated when sentences are employed as information states.[1] The basic idea is that discourse participants are in certain information states at a given point in discourse, and the use of a particular sentence by some discourse participant will change the informational states of all the participants in that context who accept that sentence. If the sentence is informative, the participants will posses information that they had previously not possessed. In a simple sort of update semantics, we might model the information common to all parties as the set of worlds that are epistemically possible given their information (the set of worlds that, so far as these participants know, could be actual). We can then think about an update, potentially effected upon the assertive utterance of some sentence φ in some context σ, as a mapping from the set of worlds that are taken to be epistemically possible by the participants of σ before the utterance of φ to the set of worlds that are taken to be epistemically possible after. Assuming that everyone in the context is in the same information state at a given point in discourse (an assumption we may eventually want to drop), we have a semantic theory in which the value assigned to φ is a function $[\varphi]$ that maps each discursive context σ in which φ can be employed to the context $\sigma[\varphi]$

[1] See Veltman (1996), Groenendijk, Stokhof and Veltman (1996), Gillies (2004), and Willer (2013).

that would result upon its being employed in that context.

A dynamic theory of this sort is a possible worlds semantics, but one in which the values of sentences are not sets of worlds, but, rather, functions from sets of worlds to sets of worlds. For an atomic sentence p, updating σ with $[p]$ results in a context $\sigma[p]$ that contains only the worlds in σ in which p is true. That is:

1. $\sigma[p] = \{w \in \sigma : p \text{ is true in } w\}$

The update effected by the assertoic use of a sentence of the form $(\neg\varphi)$ in context σ has the opposite effect, subtracting the $\sigma[\varphi]$, the context that would result from the assertoric use of φ in σ, from σ. That is,

2. $\sigma[\neg\varphi] = \sigma - \sigma[\varphi]$

Conjunction is treated as sequential update. So, the update effected by the assertoric use of a sentence of the form $(\varphi \wedge \psi)$ in context σ is one in which σ is first updated with $[\varphi]$ and then updated with $[\psi]$. That is,

3. $\sigma[\varphi \wedge \psi] = (\sigma[\varphi])[\psi]$

Disjunction can be defined in terms of negation and conjunction, exploiting the idea that we can think of $(\varphi \vee \psi)$ as $\neg(\neg\varphi \wedge \neg\psi)$. So, though it's a bit unwieldy, we have:

4. $\sigma[\varphi \vee \psi] = \sigma - ((\sigma - \sigma[\varphi]) - (\sigma - \sigma[\varphi])[\psi])$.[2]

Recursively defining updates of logically complex sentences in this way enables us to specify the update function that is the semantic value of any logically complex sentence.

[2]The derrivation is as follows:

$$\begin{aligned}
\sigma[\varphi \vee \psi] &= \sigma[\neg(\neg\varphi \wedge \neg\psi)] \\
&= \sigma - \sigma[(\neg\varphi) \wedge (\neg\psi)] \\
&= \sigma - ((\sigma[\neg\varphi])[\neg\psi]) \\
&= \sigma - ((\sigma - \sigma[\varphi])[\neg\psi]) \\
&= \sigma - ((\sigma - \sigma[\varphi]) - (\sigma - \sigma[\varphi])[\psi])
\end{aligned}$$

More complex dyanmic frameworks, for instance, billateral frameworks that assign both positive and negative updates such as that proposed by Willer (2021) are able to provide more elegant characterizations of disjunction.

A standard "informational" dynamic theory of this sort takes it for granted that atomic sentences of the language encode pieces of information about how the world is. The assertoric utterance of an atomic sentence rules out a particular set of worlds in a given context in virtue of the fact that such a sentence encodes such a piece of information, one that can be modeled as a set of possible worlds—the worlds that are consistent with this information about how the world is. Standard informational dynamic semantics is thus, despite its differences from a "static" possible worlds semantics, still a worldly semantics. Speakers' knowledge of worldly states of affairs and their relations is taken as basic with respect to their knowledge of the relations that obtain between sentences of their language. Consider just the notion of incompatibility between sentences. On a standard informational dynamic semantics, two sentences φ and ψ are incompatible just in case updating any context σ with $[\varphi]$ and then $[\psi]$ results in the absurd context, consisting in the null set of words. That is, φ and ψ are incompatible just in case, for all contexts σ, $(\sigma[\varphi])[\psi] = \varnothing$. So, for instance, "*a* is white" is incompatible with "*a* is black" just in case if you update any context with "*a* is white" and then "*a* is black," you end up with the absurd context, consisting in the null set of worlds. This will be the case just in case there is no world in which both "*a* is white" is true and "*a* is black" is true. So, if we want a theory of this sort to account for speakers' knowledge of the semantic relation between these two sentences, we must, in thinking of the incompatibility between these sentences in these terms, take it that speakers antecedently have knowledge of the fact that the set of worlds in which *a* is white and the set of worlds in which *a* is black are disjoint. In Chapter Two, I argued that we cannot appeal to knowledge of this sort in accounting for speakers' knowledge of meaning. So, endorsing a dynamic semantics of this informational variety does not evade the arguments against worldly semantics put forward in the previous two chapters.

The version of dynamic semantics I'll propose here, which does not appeal to a notion of information at all and so is not subject to the criticisms of worldly semantics put forward in the previous chapters, will be a formalization of the semantic theory proposed by Brandom (1994).[3] Brandom's basic idea is to start with a notion of discourse, understood along game-playing model and then give an account of the propositional content of a sentence that can be used in that discourse has by thinking of the use of that sentence as a certain type of "discursive move," the significance of which can be understood entirely in terms of the change in score that making of such a move would bring about, as this change in score is assessed from the players of the game. Characterizing this semantic theory as a dynamic theory, it is one in which, rather than thinking of contexts as sets of possible worlds, we think of what a context is in terms of the "score" that characterizes a particular stage in discourse, and we think of the meaning of a sentence in terms of its potential to change that score.[4] The resulting framework is what Bernhard Nickel (2013) calls a "normative" rather than "informational" dynamic semantics, where the contexts that get updated are understood not in terms of the informational content they contain (modeled by a set of worlds), but in terms of the normative statuses that have been assigned to the discursive participants.

4.3 The Game-Playing Model of Discursive Practice

Following Brandom (1994, 2000) the basic idea of discursive role semantics is that we can model discourse on what he calls "the game of giving and asking for reasons." Uttering a sentence is conceived of as making a move in the game. Like any game, the game of giving and asking for reasons has rules. The basic rule in the game is that you can only make a move if you're *entitled* to make it. An entitlement is a sort of move-making license, something that's acknowledged by the players of the game as making a move available for one to make. There are a few ways in which one can

[3]Perhaps more accurately, it's formalization of the *normative pragmatic* theory, put forth by Brandom in the first half of *Making It Explicit*, that is supposed to ground the *inferentialist semantic* theory, put forth in the second half of the book. Essentially, I'm avoiding this two-stage order of explanation and doing semantics directly in terms of the pragmatics. This is how Nickel (2013), who I'm drawing from

[4]Lewis (1978) proposed thinking of various aspects of meaning in terms of this sort of scorekeeping, but this sort of scorekeeping was taken to be a supplement to possible worlds semantics, not a replacement of it. Before Lewis, however, Wilfrid Sellars proposed to think of linguistic meaning entirely in terms of this sort of scorekeeping model, and it is this Sellarsian idea that gets taken up in Brandom's *Making It Explicit*.

acquire entitlement to a move. One way is to be attributed entitlement by another player who takes you to have made a move as the exercise of a reliable observational capacity, or instance, the capacity to see that something is the case. Such a capacity is entitlement-conferring in the sense that, insofar as one is taken to exercise it, one will be taken to be entitled to the claim one comes to endorse upon that exercise. Another way to come to be entitled to a move is by *inheriting* this entitlement from some other player who has licitly made the move. One of the main functions of actually making a move (as opposed to merely being entitled to make it) is that, in making a move to which you are entitled, you pass the entitlement that you have to make that move on to others, who are then able to make that move themselves.

What makes the game a game is that players can *challenge* each other's moves, calling into question the entitlement they have to a move that they've made. To respond to a challenge, you must demonstrate your entitlement to the move that was called into question. One way of responding to a challenge to some move of yours is to demonstrate how the making of that move was made available by way of your exercise of an entitlement-conferring capacity. If you're able to do such a thing, you've done what you need to do in order to secure your entitlement to that move, successfully defending your move against that challenge. In a case in which you've made your move on the basis of another player's making it, you can respond to a challenge by deferring back to this other player. It is then this player who must respond to the challenge, once again, either by demonstrating how the move that they made was made available to them by the exercise of an entitlement-conferring capacity, or by deferring to the authority of another player. Somewhere along the line, some player must have exercised an entitlement-conferring capacity, or else no one in that chain of deference is entitled to the move. The authority that you have in making a move, entitling other players to make it, corresponds to the responsibility that you are able to bear in responding to challenges to your making of that move. This is why, if you continually fail to be able to respond to challenges to moves that you've made, failing to live up to the responsibility that you've undertaken in making the moves that you have, other players will stop taking your to have any authority. Eventually, your moves will no longer be taken to have the significance that moves generally do, functioning to entitle others to make them.

Given this general structure of the game of giving and asking for reasons, we can think of what it is that you do in making some move as undertaking a particular sort of *commitment*—a commitment to demonstrate your entitlement to that move if challenged, either by showing how the making of the move was the product of the execution of an entitlement-granting procedure, or by deferring to another player from whom you've inherited the entitlement to it.[5] To undertake a commitment of this sort, in making some move, is to take on the responsibility that underwrites the authority that one has, in making a move, to entitle others to make that move. Since one has this authority only if one takes on this responsibility, undertaking a commitment in making a move is necessary in order for move-making to serves its basic function—entitling others to make the move that was made.

The key idea that enables us to construct a semantic theory on the basis of this conception of discursive practice is that when one undertakes a commitment to some move, say p, one will generally not commit oneself to that move and only that move. Rather, commitment to that move will bring with it commitments to certain other moves. These other commitments that one takes on in committing oneself to p are the *committive consequences* of p. If a move q is a committive consequence of a move p, then a player who commits herself to p is not only committed to p, but also committed to q. So, this player is not only responsible for defending p against potential challenges, but also responsible for defending q against potential challenges. A second, directly related relation is that of *permissive consequence*. Roughly, if q is a permissive consequence of

[5]John MacFarlane (2010, 91) claims that this way of construing what it is to be committed to a move is circular. The worry is that, if we think of what it is to be committed to a move p in terms of a commitment to demonstrate entitlement to that commitment, then the account of commitment to p must appeal to the notion of being committed to p. However, given the way I'm construing it here, there is no circularity involved. McFarlane's charge of circularity hinges on the claim that what one must be committed to demo. What one is committed to doing, on this construal, is demonstrating one's entitlement to *make the move*, not demonstrating one's entitlement to *be committed to the move*. When one makes a move one does, of course, undertake a commitment to it, but undertaking that commitment just is committing oneself to demonstrating ones entitlement to make that move.

p, a player who is committed and entitled to p is (prima facie) entitled to q and so can appeal to her entitlement to p in response to a challenge to q. Generally, if q is a committive consequence of p, q will also be a permissive consequence of p, though the converse might not hold.[6] So, though, in making some move, one commits oneself to more than just that one move, if one is entitled to the move one makes, one will also be able to appeal to this entitlement in response to a challenge to one of these other moves. Finally, commitment to some moves will preclude entitlement to others. These are what I'll call the *preclusive consequences* of p.[7] If q is a preclusive consequence of p, then, if some player is committed to p, they are precluded from being scored as entitled to q, insofar as they remain committed to p.

I've specified three consequence relations. We now have the raw materials to specify, in broad outline, how the semantic theory constructed on the basis of this conception of discursive practice will work. To model the players' attribution of normative statuses to one another, we'll say that each player has a "scorecard" wherein she keeps track of all the normative statuses that she's attributed to all of the other players of the game. We'll call a player's attitudes of taking some moves to be committive, permissive, and preclusive consequences of others her "scorekeeping principles." When we model the meaning of some sentence, we'll do it from the perspective of some player who has a certain set of scorekeeping principles. We define a subset of the total set of sets of normative assignments, which are the scorecards that each player might have, given their scorekeeping principles. We can then define the semantic value of a sentence φ, from the perspective of some scorekeeper m, as a function that takes any player n and any scorecard σ that m might have, and returns another scorecad, $\sigma[\checkmark_n\langle\varphi\rangle]$, which is the result of m's updating σ with n's making the move φ. These are the candidates for models of the meanings of sentences that I think Brandom's semantic

[6]For instance, consider how certain inductive lines of reasoning might entitle one to some move on the basis of some other move, but not commit one to it. For instance, observing a red sky at night might, on inductive grounds, entitle one to the claim that the next day's whether will be fair, but one is presumably not *committed* to this claim upon being committed to the claim that the sky is red.

[7]Brandom calls these moves the moves that are "incompatible" with p. I avoid this terminology here because to employ it is to assume that the relation of preclusive consequence is symmetric, which, if it is true, is something that should be shown, not assumed from the outset.

theory gives us, and what we have, defining semantic values in this way, is a kind of dynamic semantics. We take the semantic value of a sentence φ to be its context change potential. However, unlike standard informational dynamic theories in which contexts are understood as sets of epistemically possible worlds, contexts are be understood in terms of sets of normative assignments that conform to a given speakers' scorekeeping principles—scorecards that speaker might have.

Before getting into the details of the proposal, we are now in a position to see how framework that Nickel proposes to systematize Brandomian semantics is inadequate. Nickel tries to model Brandomian contexts as sets of sentences; a given context, he says, can be understood as the set of sentences to which everyone in that context is committed (2013, 340). While this does lead to a simple semantic framework, it is not at all adequate to enable us to model a Brandomian context. Let me point out just two key problems with it. The first problem with this way of modeling a Brandomian context is that it assumes that there is some single set of claims to which everyone in the discursive context is committed. It is crucial, however, that it need not be the case that everyone in the context is committed to the same things. Indeed, on the Brandomian picture, communication requires that particular participants, who are uniquely entitled to claims, are able to uniquely commit themselves to claims, bearing the responsibility for those commitments such that players are able to pass the buck back to them. Only by seeing how commitments and entitlements are such as to vary from player to player are we able to see how the game can function as a model for communication at all. The second problem with this way of thinking about a Brandomian context is that only one normative status is considered, and, as I am sketching the framework, we cannot do without at least a few. This will be clear when it comes to giving a semantics for logical operators. For instance, negation will be defined in terms of preclusive consequence, where this notion requires the interplay between distinct normative statuses, the idea that *commitment* to some claim can preclude *entitlement* to another. So, though Nickel has the right idea for how to think about Brandomian semantics, the actual framework he proposes is utterly inadequate to put the Brandom's basic resources to work. With that noted, let us now flesh out the formal details of a more adequate formal semantic framework.

4.4 The Basic Framework

To start, it will be helpful to introduce a set of special symbols to express the normative statuses that one might bear towards a move:

1. "$\checkmark$" expresses the status of *having made* a move. The formula $\checkmark_n\langle\varphi\rangle$ says that player n has made move φ. When this formula shows up in some player m's scorecard, this means that m scores n as having actually made the move φ.

2. "$\oplus$" expresses the status of *being committed* to a move. The formula $\oplus_n\langle\varphi\rangle$ says that n is committed to φ. When this formula shows up in some player m's scorecard, this means that m scores things in such a way that n is *obligated* to make φ if they're called upon to do so in the context of an appropriate challenge.

3. "$\bigcirc$" expresses the status of *being entitled* to a move. The formula $\Diamond_n\langle\varphi\rangle$ says that n is entitled to φ. When this formula shows up in some player m's scorecard, this means that m scores things in such a way that n is *permitted* to make φ insofar they recognize that it's a move that they're in a position to make.

4. "$\ominus$" expresses the status of *being precluded from being entitled* to a move. The formula $\ominus_n\langle\varphi\rangle$ says that n is precluded from being entitled to φ. When this formula shows up in some player m's scorecard, this means that m scores things in such a way that n is *precluded* from making φ, given the other moves that they've made.

Given a language $\mathcal{L}$, the set of moves appealed to in providing the discursive role semantics for $\mathcal{L}$ will simply be the set of sentences of $\mathcal{L}$, and the "players" appealed to in providing the discursive role semantics will be the speakers of $\mathcal{L}$. Thu, following Kukla, Lance, and Retall (2009) we can define the following:

Field of Play: A *field of play* is a triple consisting of

1. A non-empty set of players (PLAYER)
2. A non-empty set of moves (MOVE)
3. A non-empty set of normative statuses (STATUS)

For our toy language, PLAYER is the set $\{A, B, C\}$, our three discursive participants, our three players of the game, and MOVE is the set of all the claims that can be made by any of our three players, either by employing one of the 27 atomic sentences of our language or employing one infinite of the logically complex sentences, and STATUS is the set of statuses just defined: $\{\checkmark, \oplus, \bigcirc, \ominus\}$. The basic way in which these three elements come together is in the form of a *normative assigment*, defined as follows:

Normative Assignments: A *normative assignment* is any formula consisting of the specification of an $\boldsymbol{s} \in$ STATUS, a $\varphi \in$ MOVE, and an $n \in$ PLAYER that is written as $\boldsymbol{s}_n\langle\varphi\rangle$.

Now, to speak about normative positions that one might occupy, such as "being committed to p," in abstraction from anyone's actually occupying that position, we'll introduce what we'll call a "player place-holder," for which we'll use the greek letter α. With this, we can define two more things:

Normative Positions: A *normative position* is any formula consisting in an $\boldsymbol{s} \in$ STATUS, a $\varphi \in$ MOVE, and a player place-holder, which is written as $\boldsymbol{s}_\alpha\langle\varphi\rangle$.

Scorekeeping Principles: A *scorekeeping principle* is a sequence of the form $\Gamma \vdash A$, where Γ is a (possibly null) sequence of normative positions, and A is a single normative position.

It needs to be emphasized that the use of the player place-holder α in the specification of scorekeeping principles is not to be understood in terms of universal quantification over the elements of PLAYER on the part of the speaker who has that scorekeeping principle. It will turn out to be the case that, for any scorekeeping principle that a speaker has, expressible with the use of this place-holder, there will correspond to a practice describable with the use of universal quantification. However, it important to be clear that one's *having* a scorekeeping principle, keeping score *in accordance* with it, is distinct from one's *making* a corresponding quantificational claim, *explicitly acknowledging* a commitment to the practice of keeping score, the specification of which would require universal quantification. One will only be able to do such a thing insofar as one's language contains quantificational vocabulary, and the semantic theory works by enabling us to comprehend such vocabulary as functioning to make *explicit* what is must already be *implicit* in the practice of keeping score. So, it is worth emphasizing that the "α" in a scorekeeping principle is to be understood, in the first instance, as the generic "one" not the universal "everyone."

We can now define two fundamental sorts of consequence relations—committive and preclusive—as two different sorts of scorekeeping principles. Where Γ is any sequence of normative positions and φ is any element of MOVE, a principle of *committive consequence* is a scorekeeping principle of the form $\Gamma \vdash \oplus_\alpha\langle\varphi\rangle$, and a principle of *preclusive consequence* is a scorekeeping principle of the form $\Gamma \vdash \ominus_\alpha\langle\psi\rangle$.[8] We will work on the simplifying assumption, which is fine for our toy language but will need to be reconsidered for a genuine natural language, that it is at least possible for there to be a language in which there are only principles of committive and preclusive consequence, and entitlement just comes along for the ride, being attributed whenever one commits oneself to something to which one is not precluded from being entitled. We'll say that a *material scorekeeping principle* is a scorekeeping principle containing only atomic sentences in the move spot of the normative positions it contains. Material scorekeeping principles are what determine the semantic significances of the atomic sentences. Sample material scorekeeping principles from our toy language include the following:

$\oplus_\alpha\langle a \text{ is white}\rangle \vdash \ominus_\alpha\langle a \text{ is black}\rangle$
$\oplus_\alpha\langle a \text{ is gray}\rangle, \oplus_\alpha\langle b \text{ is white}\rangle \vdash \oplus_\alpha\langle a \text{ is darker than } b\rangle$
$\oplus_\alpha\langle a \text{ is darker than } b\rangle, \oplus_\alpha\langle b \text{ is darker than } c\rangle \vdash \oplus_\alpha\langle a \text{ is darker than } c\rangle$

Clearly, if we tried to enumerate all of the scorekeeping principles for our toy language in this way, it'd be quite a long list! The number of scorekeeping principles we'll have

[8]Note that this a generalization of Brandom's own definitions as we allow both commitments and preclusions on the left of the turnstile.

will be reduced once we articulate the theory at a subsentential level. Once we do that, our principles will be general, not just with respect to the player expressions they contain, but with respect to the singular terms occurring in the specifications of the moves, and with that sort of generality we will then be able to easily specify all the principles we need for this simple toy language. However, we'll stay at the level of sentences for the moment to get an initial grip on how the semantic theory is meant to work.

Scorecards get updated through the application of scorekeeping principles like these ones. Consider just the first: $\oplus_\alpha\langle a \text{ is white}\rangle \vdash \ominus_\alpha\langle a \text{ is black}\rangle$. The turnstile here can be informally understood as saying that if some player is scored as occupying the positions on the left, then they should be scored as occupying the position on the right. So, applying this scorekeeping principle to a scorecard σ amounts to adding to σ, $\ominus_n\langle a \text{ is black}\rangle$ for any player n such that $\oplus_n\langle a \text{ is white}\rangle$ is in σ. To officially state this idea, where A is some normative position of the form $\boldsymbol{s}_\alpha\langle\varphi\rangle$, let us use the notation A_n to denote the result of substituting the player place-holder α with a player n. Likewise, for a set of positions Γ, let Γ_n denote the result of substituting the player place-holder in each position in Γ with n. We can then define the application of principles as follows:

> **Application of Principles:** The result of applying a set of scorekeeping principles π to a scorecard σ, which we denote $\pi(\sigma)$, is the smallest superset of σ such that for every principle of the form $\Gamma \vdash A \in \pi$ and every player n, if $\Gamma_n \in \pi(\sigma)$, then $A_n \in \pi(\sigma)$

This definition of application of scorekeeping principles ensures that the operation of applying a set of scorekeeping principles to a scorecard is a closure operation. That is, for any scorecards σ and τ, the following facts hold:

> **Extensivity:** $\sigma \subseteq \pi(\sigma)$
> **Monotonicity:** If $\sigma \subseteq \tau$, then $\pi(\sigma) \subseteq \pi(\tau)$
> **Idempotency:** $\pi(\pi(\sigma)) = \pi(\sigma)$

Thus, a set of scorekeeping principles can be understood much like a classical consequence relation, under which a set of sentences, or in this case, normative assignments, can be closed. We'll say that a scorecard σ *conforms* to a set of scorekeeping principles π just in case $\pi(\sigma) = \sigma$.

We can now define two things, relative to one another: a set of scorecards that each player m might have, and the effect of any player n's making some move φ, relative to any scorecard that m might have:

> **Scorecards Players Might Have:** Let m be any player with a set of scorekeeping principles π. The set of scorecards Σ_m that m might have can be recursively defined as follows:
>
> 1. $\varnothing \in \Sigma_m$
> 2. For any $\sigma \in \Sigma_m$, any $n \in$ PLAYER, and any $\varphi \in$ MOVE, $\sigma[\checkmark_n\langle\varphi\rangle] \in \Sigma_m$
>
> **Updates:** Let n be any other player. The result of updating σ with $\checkmark_n\langle\varphi\rangle$, which we write as "$\sigma[\checkmark_n\langle\varphi\rangle]$," is defined as the final step in the following three step process:
>
> 1. $\sigma[\checkmark_n\langle\varphi\rangle]_1 = \sigma \cup \{\checkmark_n\langle\varphi\rangle, \oplus_n\langle\varphi\rangle\}$
> 2. $\sigma[\checkmark_n\langle\varphi\rangle]_2 = \pi(\sigma[\checkmark_n\langle\varphi\rangle]_1)$
> 3. $\sigma[\checkmark_n\langle\varphi\rangle]$ is the smallest superset of $\sigma[\checkmark_n\langle\varphi\rangle]_2$ such that, for any $\psi \in$ MOVE, if $\oplus_n\langle\psi\rangle \in \sigma[\checkmark_n\langle\varphi\rangle]_2$, and neither $\ominus_n\langle\psi\rangle \in \sigma[\checkmark_n\langle\varphi\rangle]_2$ nor $\ominus_m\langle\psi\rangle \in \sigma[\checkmark_n\langle\varphi\rangle]_2$, then $\bigcirc_n\langle\psi\rangle \in \sigma[\checkmark_n\langle\varphi\rangle]$

So, supposing we are m, we assume that one way that we might score the game is to have it such that no one has played any moves at all, and so no one is committed, entitled, or precluded from being entitled to anything. When some player n makes some move φ, we add n's having made φ and being committed to φ to our scorecard. We then apply our scorekeeping principles to that scorecard, assigning to n any positions that follow from our scorekeeping principles. Finally, we attribute entitlement to any move ψ to which we now score as n as committed, unless we take n to be precluded from being entitled to ψ or we take ourselves to be precluded from being entitled to ψ. This last step amounts to Brandom's (1994, 176-178) principle of "default entitlement," according to which when one makes a claim one is generally taken to be entitled to it by default, unless there's some specific reason to challenge it, such as incompatibility with the claimant's commitments or our own, and that's how entitlement figures into the system here. So, the way we are doing things here, scorekeeping principles fundamentally involve the attributions of commitments and preclusions of entitlements, and entitlement just comes along for the ride by default wherever it can.

Defining updates and scorecards players might have in this way lets us define the semantic value of a sentence φ, relative to a player m, as a function that takes any scorecard m might have and any other player n and returns the scorecard that is the result of m's updating their scorecard with n's making move the φ:

Semantic Values:

$$[\![\varphi]\!]^m = f : (\Sigma_m \times \text{PLAYER}) \rightarrow \Sigma_m$$
$$f(\sigma, n) = \sigma[\checkmark_n\langle\varphi\rangle]$$

The task of defining semantic values for the total set of sentences of the language amounts to the task of recursively specifying rules for keeping score, such that, given a base set of scorekeeping principles, which determine the semantic significance of the simple expressions of the language, one can specify a set of scorekeeping principles sufficient for determining the semantic significance of any of the complex expressions belonging to the language. Let us first consider how, given a set of scorekeeping principles that relate positions involving atomic sentences, we can determine the set of scorekeeping principles that relate positions involving any of the logically complex sentences.

4.5 Introducing Logical Operators

If we're giving a discursive role semantics for some language, we'll start by specifying our set of scorekeeping principles concerning only the atomic sentences of that language. This enables us to specify the update function that is the semantic value of each atomic sentences. In order to extend the semantics to logically complex sentences, we need a way of extending our set of scorekeeping principles so that we can specify the update function for logically complex sentences. Nickel, who thinks of Brandomian contexts in terms of sets of sentences that everyone is committed to, thinks that specifying such updates will be quite difficult. He writes,

> Conjunction is easy: a speaker who asserts a conjunction ($p \wedge q$) and thus commits herself to it just commits herself to each of the conjuncts p and q. Negation is trickier: committing oneself to ($\neg p$) is not the same as not committing oneself to p—the latter, but not the former, is compatible with agnosticism about p, (345).

Nickel is right that, if the only status we have is commitment, defining negation in normative terms is tricky, indeed, probably impossible.[9] But if we have multiple normative statuses—particularly, the statuses of commitment and preclusion of entitlement—it is relatively straightforward.

Brandom's way of thinking about negation is to start with the notion of incompatibility, understood in terms of the relation of preclusive consequence—where Γ is incompatible with φ just in case commitment to all the claims in Γ precludes entitlement to φ—and to then think the negation of a claim as the minimal incompatible, the claim that is entailed by every claim that is incompatible with that claim. So, for instance, "It's not the case that *a* is darker than *b*" is entailed by {"*a* is lighter than *b*" }, {"*a* is the same shade as *b*" }, {"*a* is white *b*," "*b* is gray" } and every other set of claims such that commitment to all the members of that set precludes entitlement to "*a* is darker than *b*," and it is in virtue of this fact that it can be identified as the negation of the claim "*a* is darker than *b*."[10] While still thinking of negation in basically these terms, I want to propose a somewhat different way of introducing negation and other classical con-

[9] This is something that Nickel himself doesn't seem to realize. When he tries to consider Brandom's (2008) incompatibility semantics for logical operators, he isn't even able to define the notion of incompatibility on which Brandom's semantics is based. Nickel tells us, purporting to speak for Brandom, "two sentences are incompatible just in case commitment to one precludes commitment to the other," (345). This is crucially not Brandom's definition of incompatibility. For Brandom, two sentences are incompatibile just in case commitment to one precludes *entitlement* to the other. One surely *can* be *committed* to two incompatible sentences. What one *can't* be is *committed and entitled* to both sentences, since commitment to one precludes entitlement to the other. Taking there to be two normative statuses that interact in this way is one of the fundamental technical innovations of Brandom's semantic theory that distinguishes it over predecessor theories of a similar theoretical orientation, most notably Dummett's assertability semantics, and this interplay between the statues of commitment and entitlement is absolutely essential to Brandom's definition of incompatibility and, accordingly, negation.

[10] Here, I am quoting sentences and, in doing so, I mean to be mentioning not the sentence itself but the move that is made in assertorically uttering that sentence.

nectives into this dynamic framework, since doing so will illustrate how a certain class of existing proof-theoretic semantic proposals for the logical connectives, developed in independence of Brandom, can be intuitively interpreted in this framework. Here's the thought: rather than introducing negation by way of a notion of incompatibility defined in terms of preclusive consequence, as all of Brandom's proposals have us do, we can introduce it directly in terms of the statuses of commitment and preclusion of entitlement. Doing this enables us to provide an interpretation of *bilateral* logic, of the sort proposed by Smiley (1996) and Rumfit (2000), directly in the context of this framework that enables us to provide a semantics for the logical connectives.

The basic formal idea of Smiley/Rumfit bilateral logic is that we associate each sentence of the language with a sign, positive or negative. Smiley and Rumfitt think of these two signs in terms of two opposite acts, acceptance or affirmation, on the one hand and rejection or denial, on the other. The basic interpretive idea, as it's spelled out by Rumfitt, is to think of a unsigned formula φ as the raising of the question "Is it the case that φ?" In response to the question, one can say "Yes," affirming φ or "No," denying φ. A signed formula, on Rumfit's interpretation of his logic, expresses one of these two types of acts. We'll provide a different interpretation of these signs, thinking of the "two ways" of bilateral logic in terms two opposing normative statuses that one might have with respect to some move. Where φ is a move, there are two opposing non-neutral ways in which one might stand, normatively, with respect to it: one might be *committed* to it, or one might be *precluded from being entitled* to it. It is not hard to see why there is this correspondence between these normative statuses and the two signs of Rumfitt's logic. If one is committed to a move, then, when prompted to affirm or deny the move, one is committed to affirming it. If one is precluded from being entitled to a move, then, when prompted to affirm or deny the move, one is committed to denying it

Now, there are different bilateral systems for the logical connectives that will fulfill our purpose, and any off-the-shelf bilateral system for classical logic, such as the natural deduction system proposed by Smiley (1996) or Rumfit (2000), will do.[11] However, since what we need is a way to *expand* a set of scorekeeping principles relating atomic sentences to a set of scorekeeping principles relating logically complex sentences, our purposes are really better fulfilled by a sequent calculus, along the lines of Gentzen's (1935/1969) classical sequent calculus LK, where we only have introduction rules.[12] The idea of a sequent calculus of this sort being *bilateral*, in the Smiley/Rumfit sense, is as far as I'm aware, a new one.[13] Bilateral logic has standardly been proposed in the form of natural deduction systems, which have both introduction and elimination rules (Smiley 1996, Rumfit 2000, Francez 2015). Rather than having introduction and elimination rules, Gentzen's LK is a *multiple conclusion* sequent calculus which has rules for introducing connectives on the *left* and the *right* side of the turnstile. What's significant about going bilateral here, is that, for any multiple conclusoin sequent calculus such as Gentzen's LK, we can specify a bilateral sequent system that has only *single conclusion* sequents. In such a system, rather than having left and right rules, we have positive and negative rules. This not only provides a more intuitive calculus, from the perspective of a traditional understanding of consequence, avoiding, for instance, Rumfit's (2008) criticisms of multiple conclusion sequent calculi, but it is also technically crucial here, since we want to understand sequents in terms of their role in *updating* scorecards. Only a single conclusion sequent of the form $\Gamma \vdash A$ positively provides a scorekeeper with an instruction of *what to do* when they score someone as occupying all of the positions in Γ: score them as occupying the position A. A mutliple conclusion sequent of the form $\Gamma \vdash \Delta$, where the conclusions are interpreted

[11] Rumfit's system (2000, 800-802), is arguably more natural Smiley's, but contains twice as many rules. It's worth pointing out in connection that one benefit of defining semantic values in the dynamic way that we have, rather than as they are defined in proof-theoretic semantics (Francez 2015, Stovall 2021), is that we won't have differences in meaning depending on which of two systems that both determine the same consequence relation we pick. On our approach, it is the consequence relation determined by the logical rules, which determines updates, that matters in defining semantic values, rather than the logical rules themselves.

[12] For a discussion of how Gentzen-style systems can be put to this use, see Brandom (2018), and for some examples, see Hlobil (2017) and Kaplan (2018). The specific version of this approach presented here was developed in collaboration with the ROLE (Research on Logical Expressivism) group, led by Brandom and Hlobil, as I discuss in the next chapter.

[13] Gentzen's multiple conclusion sequent calculus have notably been *interpreted* bilaterally (Restall 2006, Ripley 2015), but

disjunctively, provides no such instruction.[14] If we want intuitive rules that define the classical connectives, and function to expand scorekeeping principles from atomic to logically complex sentences, a bilateral sequent calculus is just what we need.

To introduce such a system, note first that our definition of the application of scorekeeping principles as a closure operation imposes on sets of scorekeeping principles validates the standard structural rules of Gentzen's sequent calculus:[15]

$$\frac{\Gamma, A, B, \Delta \vdash C}{\Gamma, B, A, \Delta \vdash C}\ \text{Exchange} \qquad \frac{\Gamma, A, A \vdash B}{\Gamma, A \vdash B}\ \text{Contraction} \qquad \frac{\Gamma, A \vdash B}{\Gamma, A, A \vdash B}\ \text{Expansion}$$

$$\frac{}{\Gamma, A \vdash A}\ \text{Containment} \qquad \frac{\Gamma \vdash A}{\Gamma, B \vdash A}\ \text{Monotonicity} \qquad \frac{\Gamma \vdash A \quad \Gamma, A \vdash B}{\Gamma \vdash B}\ \text{Cumulative Transitivity}$$

Now, if we're serious about doing natural language semantics in this framework, we're going to have to modify our definitions so as to allow for exceptions to Monotonicity and Cumulative Transitivity, as there are semantic implications and incompatibilities in natural language that violate these rules. For a substructural development of this framework, see the Appendix. For our purposes here, however, it does no harm to keep things simple by setting things up in a way that validates the standard structural rules. For operational rules in this bilateral setting, rather than having rules for getting logically complex formulas on the right and left of the turnstile, we only right rules, but which come in a positive and negative form. The following calculus is a translation of Ketonen's (1944) formulation of Genzen's LK:[16]

[14]Such a sequent might be understood as saying that, if one is committed to *everything* in Γ, then one's committed to *something* in Δ, in the precise sense that one is precluded from being entitled to *denying* everything in Δ. A scorekeeper who has such a "principle," and who scores someone as committed to everything in Γ, thus knows what they *can't* score such a player as precluded from being entitled to, but they don't know what they *should* score them.

[15]Note that Cumulative Transitivity, in the presence of Monotonicity, yields what may be the more familiar

[16]The Ketonen system of which this is a translation has the same rules as the system Negri and von Plato (2008) call "G3cp," but with the standard negation rules of LK. The proof of the equivalence of Ketonen's system and the bilateral calculus presented here is provided in the Appendix. Note also that, while our definition of the application of scorekeeping principles imposes MO and CT, the sequent calculus presented here, which is equivalent to G3cp with LK negation, needs neither MO nor CT to function (See Kaplan 2018, Hlobil 2018).

$$\frac{\Gamma \vdash \ominus_\alpha\langle\varphi\rangle}{\Gamma \vdash \oplus_\alpha\langle\neg\varphi\rangle}\ \oplus_\neg \qquad \frac{\Gamma \vdash \oplus_\alpha\langle\varphi\rangle}{\Gamma \vdash \ominus_\alpha\langle\neg\varphi\rangle}\ \ominus_\neg$$

$$\frac{\Gamma \vdash \oplus_\alpha\langle\varphi\rangle \quad \Gamma \vdash \oplus_\alpha\langle\psi\rangle}{\Gamma \vdash \oplus_\alpha\langle\varphi \wedge \psi\rangle}\ \oplus_\wedge \qquad \frac{\Gamma, \oplus_\alpha\langle\varphi\rangle \vdash \ominus_\alpha\langle\psi\rangle}{\Gamma \vdash \ominus_\alpha\langle\varphi \wedge \psi\rangle}\ \ominus_\wedge$$

$$\frac{\Gamma, \ominus_\alpha\langle\varphi\rangle \vdash \oplus_\alpha\langle\psi\rangle}{\Gamma \vdash \oplus_\alpha\langle\varphi \vee \psi\rangle}\ \oplus_\vee \qquad \frac{\Gamma \vdash \ominus_\alpha\langle\varphi\rangle \quad \Gamma \vdash \ominus_\alpha\langle\psi\rangle}{\Gamma \vdash \ominus_\alpha\langle\varphi \vee \psi\rangle}\ \ominus_\vee$$

For these rules to yield a complete system of propositional logic, we must impose an additional bilateral structural rule, dubbed *Reversal* by Smiley. Where A and B are normative positions and starring a normative position yeilds the opposite signed normative position (such that, if A is of the form $\oplus_\alpha\langle\varphi\rangle$, then A^* is $\ominus_\alpha\langle\varphi\rangle$ if vice versa), the rule can be stated as follows:

$$\frac{\Gamma, A \vdash B}{\Gamma, B^* \vdash A^*}\ \text{Reversal}$$

Reversal is a generalized contraposition principle. A key consequence having it as a structural rule is that it amounts to treating incompatibility as symmetric. Thus, if commitment to "*a* is black" precludes entitlement to "*a* is white," we'll also have that commitment to "*a* is white" precludes entitlement to "*a* is black."[17]

Let us go through these rules to get a sense of this normative interpretation of the classical connectives. First, consider the negation rules, which are just the introduction rules of Rumfitt's (2000) natural deduction system, translated into this context. The committive negation rule says that if occupying a set of normative positions Γ *precludes one from being entitled* to some sentence φ, then Γ *commits one* to its negation, $\neg\varphi$. As promised, this captures Brandom's (1994, 2008) definition of the negation of a sentence φ as its "minimal incompatible," the sentence implied by every set of sentences incompatible with φ. Alternately, the preclusive negation rule says that if Γ commits one to φ, Γ precludes one from being entitled to $\neg\varphi$. Consider now the committive

[17]The symmetry of incompatibility is presupposed by Brandom's (1994) definition of incompatibility. It is also explicitly assumed in incompatibility-based semantics for non-classical logics proposed by Restall (1999) and Berto (2015), though this assumption has been questioned by De and Omori (2018). I've argued elsewhere (Simonelli M.S.) that we need not take this fact as simply primitive; we can actually give a pragmatic argument for why incompatibility *must* be symmetric.

conjunction and preclusive disjunction rules. If a set of normative positions Γ commits one to φ, and Γ also commits one to ψ, then Γ commits one to $\varphi \wedge \psi$. Dually, if a set of normative positions Γ precludes one from being entitled to ψ, and Γ also precludes one from being entitled to ψ, then Γ precludes one from being entitled to $\varphi \vee \psi$. Once again, these are just the introduction rules of Rumfitt's natural deduction system. The preclusive conjunction and committive disjunction are the new rules to this calculus, and they are not only essential to its technical workings, but also conceptually significant, in offering a new way of thinking about conjunction and disjunction. The preclusive conjunction rule says that if, relative to a set of normative positions Γ, φ and ψ are *incompatible*, or, as an Aristotelian logician would put it, *contraries* in the sense that commitment to one precludes entitlement to other, then Γ precludes one from being entitled to $\varphi \wedge \psi$. Dually, the committive disjucntion says that if, relative to Γ, φ and ψ are such *subcontraries* in the sense that being precluded from being entitled to one commits one to the other, then Γ commits one to $\varphi \vee \psi$.

This set of operational rules, with only the structural rules of Containment and Exchange, give us classical logic. More importantly for our purposes, with these rules, a speaker's basic set of scorekeeping principles, involving only logically simple moves, can be expanded to include principles of committive and preclusive consequence with respect to logically complex claims. The basic idea underlying this way of introducing logical vocabulary is that to grasp this bit of vocabulary—to understand conjunction, disjunction, and negation—is to grasp how making a move in which it is used situates a player in the game. Accordingly, we can model the meaning of this bit of vocabulary by way of a set of rules which enable a player to expand their scorekeeping principles such that we can specify the update that takes place when a move in which this vocabulary is used is made. To see how these rules work, let's consider an example. We should want our rules for logical operators to combine with our basic scorekeeping principles in such a way that, since commitment to "*a* is black" precludes entitlement to "*a* is white," and commitment to "*a* is gray" precludes entitlement to "*a* is white," we'll have that commitment to "*a* is black or *a* is gray" commits one to "*a* is not white." We get this as follows:

$$\dfrac{\dfrac{\dfrac{\oplus\langle b\rangle \vdash \ominus_\alpha\langle w\rangle}{\oplus_\alpha\langle w\rangle \vdash \ominus_\alpha\langle b\rangle}\text{RV} \quad \dfrac{\oplus_\alpha\langle g\rangle \vdash \ominus_\alpha\langle w\rangle}{\oplus_\alpha\langle w\rangle \vdash \ominus_\alpha\langle g\rangle}\text{RV}}{\dfrac{\oplus_\alpha\langle w\rangle \vdash \ominus_\alpha\langle b \vee g\rangle}{\oplus_\alpha\langle b \vee g\rangle \vdash \ominus_\alpha\langle w\rangle}\text{RV}}\ominus_\vee}{\oplus_\alpha\langle b \vee g\rangle \vdash \oplus_\alpha\langle \neg w\rangle}\oplus_\neg$$

In this way, scorekeeping principles relating logically complex sentences can be generated by rules that expand a set of scorekeeping principles relating atomic sentences. In this way, we can make sense of one's ability to grasp the updates imposed by a potentially infinite number of complex sentences on the basis of a finite amount of knowledge—knowledge of the scorekeeping principles relating atomic sentences and knowledge of the rules for generating scorekeeping principles relating logically complex sentences.

Before getting to the issue of subsentential structure, it is worth pausing for a moment to take note of the fact that this sort of semantic theory provides an alternative account of the fact that speakers can understand the meanings of a potentially infinite number of sentences. Rather than accounting for this fact by thinking of the meanings of these sentences as composed out of meanings of the parts, we account for this fact by thinking of the rules for determining the semantic significance of a sentence as recursively iterable, such that rules for keeping score on the utterances of expressions of arbitrary complexity can be determined by the rules for keeping score on simple expressions. The recursive determination of meanings is all that's necessary to account for the fact that speakers can understand a potentially infinite number of sentences. We need not think of meanings themselves as compositional, in the sense of being composed out of the meanings of their parts.

4.6 Predicative Structure

In our very simple toy language, there is a basic syntactic distinction between two types of subsentential expressions: singular terms and predicates. This is, of course, is corresponds more closely to the syntactic structure of a simple formal langauge, like first-order logic, rather than a natural language like English, but the basic strategy of accounting for subsentential structure here can be extended to other syntactic categories. Let us start with singular terms.

The semantic significance of singular terms can be understood in terms of the way in which co-referential terms can be substituted in for one another. If a and b are taken to be co-referential, then, for any good inference in which a figures in a certain spot, either in the premises or the conclusions, there will be a corresponding good inference in which b figures in that spot. That is, we'll have the following rules:[18]

$$\frac{\Gamma, {}^{\oplus}/_{\ominus_\alpha}\langle\Phi(a)\rangle \vdash A}{\Gamma, {}^{\oplus}/_{\ominus_\alpha}\langle\Phi(b)\rangle \vdash A} \qquad \frac{\Gamma, {}^{\oplus}/_{\ominus_\alpha}\langle\Phi(b)\rangle \vdash A}{\Gamma, {}^{\oplus}/_{\ominus_\alpha}\langle\Phi(a)\rangle \vdash A}$$

Of course, this yields a very simple account of the meaning of proper names, a Millian one that does not take into account anything like Fregean sense. That is, indeed, something that this framework can accommodate due to its multi-perspectival nature. Following the account proposed in Chapter Eight of *Making It Explicit*, . However, rather than substantially complicating the formal framework, since our main concern here is with the meanings of predicates, this simple account will do for our purposes.

Thinking about rules for substitution of this sort enables us to think abstractly about the roles of predicates in abstraction from any particular singular terms to which those predicates are attached. So, for instance, in making the move that one makes in uttering an "a is black," one precludes oneself from being entitled to make the move that one makes in uttering an "a is white," and if one also commits oneself to "b is gray," then one commits oneself "a is darker than b," and so on. To arrive at the roles of predicates, we isolate the element of those sentential roles that stays constant as we substitute different singular terms into the sentences in accord with a rule of the above form. So, we notice that, if we take another singular term, say "c," and substitute it for the utterance of "a" in any of these utterances, the normative relations between the moves made by the utterances are preserved. Thus, we can characterize the sentences "a is black" and

[18]Because we have Reversal, these rules also give us the rules where the relevant formulas occur in the conclusion.

"*c* is black," as both sentences of the form "*x* is black," and we can say, for instance, commitment to a sentence of the form "*x* is black," precludes one from being entitled to an sentence of the form "*x* is white," and, if one is additionally committed to a sentence of the form "*y* is".[19] Talk of scorekeeping principles involving these "sentence forms" or, as Brandom puts it, *sentence frames* is intelligible through considering how the predicative aspect of sentential scorekeeping principles remains stable as different singular terms are substituted for one another into those principles. It's important to emphasize, in this regard,

So, finally, we can think of atomic scorekeeping principles as derived from scorekeeping principles relating predicate frames and rules for substituting singular terms into those frames. For instance, we can think of scorekeeping principles on sentences like $\oplus_\alpha\langle a \text{ is black}\rangle, \oplus_\alpha\langle b \text{ is gray}\rangle \vdash \oplus_\alpha\langle a \text{ is darker than } b\rangle$ as resulting from the saturation with a singular term of a scorekeeping principle on predicate frames like $\oplus_\alpha\langle x \text{ is black}\rangle, \oplus_\alpha\langle y \text{ is gray}\rangle \vdash \oplus_\alpha\langle x \text{ is darker than } y\rangle$. To spell this out, let us suppose we can use a set of variables $x_1, x_2 \ldots x_n$, and we can think of the any variable x_i as replaceable with a singular term by way of the following rule, where $\Phi_1, \Phi_2 \ldots \Phi_n$ and Ψ are any predicative contexts (which may or may not contain variables) and τ is any singular term belonging to the language:

$$\frac{\Gamma, {}^{\oplus}\!/_{\ominus_\alpha}\langle\Phi_1(x_i)\rangle \ldots {}^{\oplus}\!/_{\ominus_\alpha}\langle\Phi_n(x_i)\rangle \vdash {}^{\oplus}\!/_{\ominus_\alpha}\langle\Psi(x_i)\rangle}{\Gamma, {}^{\oplus}\!/_{\ominus_\alpha}\langle\Phi_1(\tau)\rangle \ldots {}^{\oplus}\!/_{\ominus_\alpha}\langle\Phi_n(\tau)\rangle \vdash {}^{\oplus}\!/_{\ominus_\alpha}\langle\Psi(\tau)\rangle}$$

Thus, we have, for instance, the following double application of this rule:

$$\cfrac{\cfrac{\oplus_\alpha\langle x_1 \text{ is black}\rangle, \oplus_\alpha\langle x_2 \text{ is gray}\rangle \vdash \oplus_\alpha\langle x_1 \text{ is darker than } x_2\rangle}{\oplus_\alpha\langle a \text{ is black}\rangle, \oplus_\alpha\langle x_2 \text{ is gray}\rangle \vdash \oplus_\alpha\langle a \text{ is darker than } x_2\rangle}}{\oplus_\alpha\langle a \text{ is black}\rangle, \oplus_\alpha\langle b \text{ is gray}\rangle \vdash \oplus_\alpha\langle a \text{ is darker than } b\rangle}$$

In this way, we can think of the semantic significance of predicates in terms of the rules governing sentence frames that can be saturated with any singular terms, and this explains how, for instance, when a novel singular term, for instance "*d*," is used,

[19]Talk of commitment to sentences, I should be clear, is just shorthand for commitment to the moves one makes in uttering those sentences.

a speaker will grasp the significance of saying "*d* is black," even though this is not a sentence they would have previously considered.

4.7 Providing the Full Lexical Semantics

With this formal machinery on the table, we can finally provide the full lexical semantics for our very simple toy language. To make this a bit easier on ourselves, let us first add one more bilateral structural rule, which I'll call *Bilateral Reductio* (BR).[20] Once again, where A and B are any normative positions, and starring a normative position yields the oppositely signed position, the rule can be put as follows:

$$\frac{\Gamma, A \vdash B \quad \Gamma, A \vdash B^*}{\Gamma \vdash A^*}\ \text{BR}$$

The idea is that if being committed to φ would leave one a situation in which one is both committed and precluded from being entitled to some sentence ψ, then one is precluded from being entitled to φ. Likewise, if being precluded from being entitled to φ would leave one in such a situation, then one is committed to φ. Given BR, one can treat the Reversal rule specified above as a derived structural rule, derived as follows:

$$\frac{\dfrac{\dfrac{\Gamma, A \vdash B}{\Gamma, A, B^* \vdash B}\ \text{MO}}{\Gamma, B^*, A \vdash B}\ \text{EX} \qquad \Gamma, B^*, A \vdash B^*\ \text{CO}}{\Gamma, B^* \vdash A^*}\ \text{BR}$$

It will also be helpful to have the structural rule of "Simple Transitivity" at hand:

$$\frac{\Gamma \vdash A \quad A \vdash B}{\Gamma \vdash B}\ \text{ST}$$

[20] Rumfit (2000) calls it *Smileian Reductio* (855) in reference to Smiley (1996) who first proposes the rule. It is perhaps worth noting that, given the translation procedure for going from this bilateral sequent calculus to a multiple conclusion sequent calculus, that this principle corresponds directly to Gentzen's *Cut* rule:

$$\frac{\Gamma, \varphi \vdash \Delta \quad \Gamma \vdash \varphi, \Delta}{\Gamma \vdash \Delta}\ \text{Cut}$$

which can be derived from our Cumulative Transitivity rule along with Monotonicity simply follows:

$$\dfrac{\Gamma \vdash A \quad \dfrac{A \vdash B}{\Gamma, A \vdash B}\,\text{MO}}{\Gamma \vdash B}\,\text{CT}$$

There will be reason to not treat things in this way if we decide to go substructural, since we will want to reject the structural rule of Monotonicity. For our purposes here, however, things are simplified by treating things in this way.[21]

Now we can articulate a "kernel" from which we can derive the full lexical semantics for our toy language, articulating sixteen scorekeeping principles that a speaker of this language has from which all others can be derived:

1. $\oplus_\alpha\langle x$ is darker than $y\rangle \vdash \oplus_\alpha\langle y$ is lighter than $x\rangle$
2. $\oplus_\alpha\langle x$ is lighter than $y\rangle \vdash \oplus_\alpha\langle y$ is darker than $x\rangle$
3. $\oplus_\alpha\langle x$ is darker than $y\rangle, \oplus_\alpha\langle y$ is darker than $z\rangle \vdash \oplus_\alpha\langle x$ is darker than $z\rangle$
4. $\vdash \ominus_\alpha\langle x$ is darker than $x\rangle$
5. $\oplus_\alpha\langle x$ is the same shade as $y\rangle, \oplus_\alpha\langle y$ is the same shade as $z\rangle \vdash \oplus_\alpha\langle x$ is the same shade as $z\rangle$
6. $\oplus_\alpha\langle x$ is the same shade as $y\rangle \vdash \oplus\langle y$ is the same shade as $x\rangle$
7. $\vdash \oplus_\alpha\langle x$ is the same shade as $x\rangle$
8. $\oplus_\alpha\langle x$ is the same shade as $y\rangle \vdash \ominus\langle x$ is darker than $y\rangle$
9. $\oplus_\alpha\langle x$ is gray$\rangle, \oplus_\alpha\langle y$ is white$\rangle \vdash \oplus_\alpha\langle x$ is darker than $y\rangle$
10. $\oplus_\alpha\langle x$ is black$\rangle, \oplus_\alpha\langle y$ is gray$\rangle \vdash \oplus_\alpha\langle x$ is darker than $y\rangle$
11. $\oplus_\alpha\langle x$ is black$\rangle, \oplus_\alpha\langle y$ is white$\rangle \vdash \oplus_\alpha\langle x$ is darker than $y\rangle$
12. $\oplus_\alpha\langle x$ is white$\rangle, \oplus_\alpha\langle y$ is white$\rangle \vdash \oplus_\alpha\langle x$ is the same shade as $y\rangle$
13. $\oplus_\alpha\langle x$ is gray$\rangle, \oplus_\alpha\langle y$ is gray$\rangle \vdash \oplus_\alpha\langle x$ is the same sahde as $y\rangle$
14. $\oplus_\alpha\langle x$ is black$\rangle, \oplus_\alpha\langle y$ is black$\rangle \vdash \oplus_\alpha\langle x$ is the same shade as $y\rangle$
15. $\oplus_\alpha\langle x$ is black$\rangle \vdash \ominus_\alpha\langle y$ is darker than $x\rangle$

[21] Even if we do go substructrual, however, we may nevertheless *locally* treat things thins way, since, for this particular bit of vocabulary, the various structural rules *do* apply. See Hlobil (2017) for a discussion of this notion of structural rules holding locally.

16. $\oplus_\alpha\langle x \text{ is white}\rangle \vdash \ominus_\alpha\langle y \text{ is lighter than } x\rangle$

For instance, with this set of scorekeeping principles, we can derive the principle

$$\oplus_\alpha\langle x \text{ is black}\rangle \vdash \ominus_\alpha\langle x \text{ is white}\rangle$$

from (11) and (4) as follows:

$$\dfrac{\oplus_\alpha\langle Bx\rangle,\oplus_\alpha\langle Wx\rangle \vdash \oplus_\alpha\langle Dxx\rangle \qquad \dfrac{\vdash \ominus_\alpha\langle Dxx\rangle}{\oplus_\alpha\langle Bx\rangle,\oplus_\alpha\langle Wx\rangle \vdash \ominus_\alpha\langle Dxx\rangle}\,\text{MO}}{\oplus_\alpha\langle Bx\rangle \vdash \ominus_\alpha\langle Wx\rangle}\,\text{BR}$$

and we can then derive the principle

$$\oplus_\alpha\langle a \text{ is black}\rangle \vdash \ominus_\alpha\langle a \text{ is white}\rangle$$

as a particular instance of this general one, the instance in which the singular term "a" has been substituted into the open spot marked by "x." Similarly, we can derive:

$$\oplus_\alpha\langle x \text{ is darker than } y\rangle \vdash \ominus_\alpha\langle x \text{ is lighter than } y\rangle$$

from (2), (3), and (4) as follows:

$$\dfrac{\dfrac{\oplus\langle Dxy\rangle,\oplus\langle Dyx\rangle \vdash \oplus\langle Dxx\rangle \qquad \dfrac{\vdash \ominus\langle Dxx\rangle}{\oplus\langle Dxy\rangle,\oplus\langle Dyx\rangle \vdash \ominus\langle Dxx\rangle}\,\text{MO}}{\oplus\langle Dxy\rangle \vdash \ominus\langle Dyx\rangle}\,\text{BR} \qquad \dfrac{\oplus\langle Lxy\rangle \vdash \oplus\langle Dyx\rangle}{\ominus\langle Dyx\rangle \vdash \ominus\langle Lxy\rangle}\,\text{RV}}{\oplus\langle Dxy\rangle \vdash \ominus\langle Lxy\rangle}\,\text{ST}$$

And, though it's a bit tedious, we can derrive

$$\oplus_\alpha\langle x \text{ is lighter than } y\rangle,\oplus_\alpha\langle y \text{ is lighter than } z\rangle \vdash \oplus_\alpha\langle x \text{ is lighter than } z\rangle$$

from (1), (2), and (3) as follows:

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\oplus\langle Dzy\rangle,\oplus\langle Dyx\rangle \vdash \oplus\langle Dzx\rangle \qquad \oplus\langle Dzx\rangle \vdash \oplus\langle Lxz\rangle}{\oplus\langle Dzy\rangle,\oplus\langle Dyx\rangle \vdash \oplus\langle Lxz\rangle}\,\text{ST}}{\oplus\langle Dzy\rangle,\ominus\langle Lxz\rangle \vdash \ominus\langle Dyx\rangle}\,\text{RV} \qquad \dfrac{\oplus\langle Lxy\rangle \vdash \oplus\langle Dyx\rangle}{\ominus\langle Dyx\rangle \vdash \ominus\langle Lxy\rangle}\,\text{RV}}{\oplus\langle Dzy\rangle,\ominus\langle Lxz\rangle \vdash \ominus\langle Lxy\rangle}\,\text{ST}}{\oplus\langle Lxy\rangle,\ominus\langle Lxz\rangle \vdash \ominus\langle Dzy\rangle}\,\text{RV} \qquad \dfrac{\oplus\langle Lyz\rangle \vdash \oplus\langle Dzy\rangle}{\ominus\langle Dzy\rangle \vdash \ominus\langle Lyz\rangle}\,\text{RV}}{\oplus\langle Lxy\rangle,\ominus\langle Lxz\rangle \vdash \ominus\langle Lyz\rangle}\,\text{ST}}{\oplus\langle Lxy\rangle,\oplus\langle Lyz\rangle \vdash \oplus\langle Lzx\rangle}\,\text{RV}$$

There is likely a simpler axiomatization of our toy language. Perhaps some of the scorekeeping principles included in this kernal can be derrived from others in it, or perhaps there is a smaller set of material scorekeeping principles from which all of those included here can be derrived. The point here is not to provide the simplest kernel, but simply to show that there is some not only finite but relatively manageable set of material scorekeeping principles from which the total set of material scorekeeping principles for this toy language can be derived.

Now, it is not clear whether we can actually provide anything like a full lexical semantics for natural language as we did with respect to our toy language. Nevertheless, several semantics, perhaps most notably Barbara Partee (2005), have suggested that the project of lexical semantics can be undertaken by doing something of this sort, laying down "meaning postulates." Such postulates, often formulated in first-order logic, can be interpreted straightforwardly in discursive role semantics as expressing basic scorekeeping principles. For instance, consider the postulate:

$$\forall x : \mathbf{black}(x) \rightarrow \neg\mathbf{white}(x)$$

On a model-theoretic way of thinking about meaning postulates, we might think of this postulate as saying, informally, that everything in the domain of discourse is such that if the predicate "black" is correctly applied to it, it is not the case that the predicate "white" is correctly applied to it. In model-theoretic semantics, we might take these postulates, so interpreted, to constrain the models that are considered for the purpose of semantic theorizing. In the context of discursive role semantics, however, we can interpret it as expressing the following scorekeeping principle:

$$\oplus_\alpha\langle x \text{ is black}\rangle \vdash \ominus_\alpha\langle x \text{ is white}\rangle$$

The idea of such a quantificational formula expressing such a material scorekeeping principle will be made precise in the next chapter. The point for now is just that a lexical semantics consisting in a set of meaning postulates for the atomic sentences of a natural language can be interpreted in this way, or, better, could be directly done in this framework in terms of the explicit laying down of material scorekeeping principles. Discursive role semantics thus promises to provide a unified framework for both the lexical semantics of atomic sentences, understood in terms of a set of basic scorekeeping principles (which enable us to specify updates for the atomic sentences), and the proof-theoretic semantics for non-atomic sentences, understood in terms of rules for deriving scorekeeping principles (which enable us to specify updates for the non-atomic sentences).

4.8 Conclusion

I have presented a semantic theory that, in principle, presupposes no worldly knowledge. I am actually prepared to make this claim unrestrictedly, but, for the purposes of the present project, the relevant sort of worldly knowledge that this semantic theory does not presuppose is knowledge of such things as properties and relations, and modal relations among them, or such things as possible worlds and set-theoretic relations among them. Rather than the worldly contents expressed by predicates or sentences determining the relations of entailment and incompatibility that these predicates or sentences stand to one another, the relations of entailment and incompatibility between sentences are understood directly in pragmatic terms, in terms of commitment to some sentence committing one to others or precluding one from being entitled to others. This opens up the door for thinking about the "worldly" knowledge appealed to in the semantic theories previously considered—knowledge of modal relations between properties or set-theoretic relations between sets of worlds—as nothing other than knowledge of the norms governing the use of various expressions, but *reified*, transposed into a "worldly" mode. Spelling out such a conception of the relevant sort of worldly knowledge is the task to which I now turn.