On sets of OT rankings

Igor Yanovich
MIT
yanovich@mit.edu

October 5, 2011

Contents

1 Introduction

Starting with its very formulation in [Prince and Smolensky, 1993], Optimality Theory (OT) has been a remarkably precise and formally well-defined theory, easily allowing abstract formal reasoning — in fact, many of the results of [Prince and Smolensky, 1993] are already given in such abstract form. Those initial advances were followed by a subsequent body of formal work exemplified, for instance, by [Tesar, 1995], [Tesar and Smolensky, 1996], [Moreton, 1996], [Samek-Lodovici and Prince, 1999], to name just a few relatively early references. However, one of the formal aspects of OT remained understudied, even though it would have become handy in a number of OT applications, most prominently in grammar learning: the structure of sets of all rankings compatible with a given tableau.

Let me give an example first. For the tableau in 1,[1] the set of all rankings compatible with it is given in 2. While it is possible to compute the set in 2 by hand, even for the small tableau in 1 that is a daunting task. Moreover, the resulting set is unwieldy: just looking at 2, it is very hard to figure out what is going on, and which pairwise rankings between constraints are crucially required by the data.

(1)

$C1$	$C2$	$C3$	$C4$	$C5$
W	W	e	e	L
W	e	W	e	L
e	e	W	W	L

[1] All tableaux in the paper are given in the comparative format introduced by [Prince, 2000]. The ordering of the constraints in the tableau does *not* represent the ranking.

$$
\left\{
\begin{array}{ll}
C1 \gg C3 \gg C2 \gg C4 \gg C5 & C3 \gg C1 \gg C2 \gg C4 \gg C5 \\
C1 \gg C2 \gg C3 \gg C4 \gg C5 & C3 \gg C2 \gg C1 \gg C4 \gg C5 \\
C1 \gg C2 \gg C4 \gg C3 \gg C5 & C3 \gg C2 \gg C4 \gg C1 \gg C5 \\
C2 \gg C1 \gg C4 \gg C3 \gg C5 & C2 \gg C3 \gg C4 \gg C1 \gg C5 \\
C2 \gg C4 \gg C1 \gg C3 \gg C5 & C2 \gg C4 \gg C3 \gg C1 \gg C5 \\
C2 \gg C1 \gg C3 \gg C4 \gg C5 & C2 \gg C3 \gg C1 \gg C4 \gg C5 \\
C1 \gg C3 \gg C4 \gg C2 \gg C5 & C3 \gg C1 \gg C4 \gg C2 \gg C5 \\
C1 \gg C4 \gg C3 \gg C2 \gg C5 & C3 \gg C4 \gg C1 \gg C2 \gg C5 \\
C1 \gg C4 \gg C2 \gg C3 \gg C5 & C3 \gg C4 \gg C2 \gg C1 \gg C5 \\
C4 \gg C1 \gg C2 \gg C3 \gg C5 & C4 \gg C3 \gg C2 \gg C1 \gg C5 \\
C4 \gg C2 \gg C1 \gg C3 \gg C5 & C4 \gg C2 \gg C3 \gg C1 \gg C5 \\
C4 \gg C1 \gg C3 \gg C2 \gg C5 & C4 \gg C3 \gg C1 \gg C2 \gg C5 \\
C1 \gg C3 \gg C2 \gg C5 \gg C4 & C3 \gg C1 \gg C2 \gg C5 \gg C4 \\
C1 \gg C2 \gg C3 \gg C5 \gg C4 & C3 \gg C2 \gg C1 \gg C5 \gg C4 \\
C2 \gg C1 \gg C3 \gg C5 \gg C4 & C2 \gg C3 \gg C1 \gg C5 \gg C4 \\
C1 \gg C3 \gg C4 \gg C5 \gg C2 & C3 \gg C1 \gg C4 \gg C5 \gg C2 \\
C1 \gg C4 \gg C3 \gg C5 \gg C2 & C3 \gg C4 \gg C1 \gg C5 \gg C2 \\
C4 \gg C1 \gg C3 \gg C5 \gg C2 & C4 \gg C3 \gg C1 \gg C5 \gg C2 \\
C1 \gg C4 \gg C2 \gg C5 \gg C3 & C4 \gg C1 \gg C2 \gg C5 \gg C3 \\
C1 \gg C2 \gg C4 \gg C5 \gg C3 & C4 \gg C2 \gg C1 \gg C5 \gg C3 \\
C2 \gg C1 \gg C4 \gg C5 \gg C3 & C2 \gg C4 \gg C1 \gg C5 \gg C3 \\
C2 \gg C3 \gg C4 \gg C5 \gg C1 & C3 \gg C2 \gg C4 \gg C5 \gg C1 \\
C2 \gg C4 \gg C3 \gg C5 \gg C1 & C3 \gg C4 \gg C2 \gg C5 \gg C1 \\
C4 \gg C2 \gg C3 \gg C5 \gg C1 & C4 \gg C3 \gg C2 \gg C5 \gg C1 \\
C1 \gg C3 \gg C5 \gg C2 \gg C4 & C3 \gg C1 \gg C5 \gg C2 \gg C4 \\
C1 \gg C3 \gg C5 \gg C4 \gg C2 & C3 \gg C1 \gg C5 \gg C4 \gg C2 \\
C1 \gg C4 \gg C5 \gg C2 \gg C3 & C4 \gg C1 \gg C5 \gg C2 \gg C3 \\
C1 \gg C4 \gg C5 \gg C3 \gg C2 & C4 \gg C1 \gg C5 \gg C3 \gg C2 \\
C2 \gg C3 \gg C5 \gg C1 \gg C4 & C3 \gg C2 \gg C5 \gg C1 \gg C4 \\
C2 \gg C3 \gg C5 \gg C4 \gg C1 & C3 \gg C2 \gg C5 \gg C4 \gg C1
\end{array}
\right\} \tag{2}
$$

Moreover, suppose we have computed such a set. Can we tell if we lose any information — that is, can we compute the original tableau (modulo OT-equivalence) from the set we have? The answer is yes, but without the apparatus we will develop in this paper, it is hard to see how one could obtain that answer — and even less how one could actually compute back the tableau from the set 2.

Our example shows why sets of all rankings compatible with a tableau were not studied: it is very hard to study directly objects like the set in 2. This paper overcomes the problem by way of finding simpler objects which can represent big sets such as 2. Those objects are *partial rankings* and *sets of partial rankings*. We define OT compatibility for them, and show that they can be safely used as substitutes when reasoning about sets of total rankings. For instance, instead of 2, we can use the following set of partial rankings:

$$
\left\{
\begin{array}{l}
(C1 \gg C5) \wedge (C3 \gg C5) \\
(C1 \gg C5) \wedge (C4 \gg C5) \\
(C2 \gg C5) \wedge (C3 \gg C5)
\end{array}
\right\} \tag{3}
$$

3 contains exactly the same information as 2, and we can compute one from the other (in fact, that is exactly how I built the set in 2: I did not have to check all the 5! = 120 total orderings of the 5-constraint **Con** for compatibility with 1.)

The goal of the paper is thus to develop the methods for working with simpler partial-ranking-based objects, and with their help to analyze the space of sets of total rankings. By the end of the paper, we will be able to build a set of all rankings compatible with a tableau, and then turn it back into the original tableau, going back and forth between the two without loss of information; to tell if a given set of rankings is such that it contains all and only rankings compatible with a given tableau, and thus constitutes a proper OT grammar hypothesis (we will show that not all sets of rankings have that property); and much more.

The plan of the paper is as follows. Section 2 discusses some of the previous work on which we build in this paper, and the issue of overcommitment in OT learning which was unavoidable for the learning algorithms proposed, and which the current paper overcomes.

After that, we give in Section 3 without proof the results concerning OT-equivalence classes of tableaux: before we start looking at sets of rankings, it is useful to get all the instruments needed in order not to pay any attention to differences between OT-equivalent tableaux. The full theory of tableau equivalence classes, which generalizes the earlier work by generalizes the work of [Prince, 2002] and [Brasoveanu and Prince, 2005], a.o., is given in Appendix A: though the main corollaries of that theory are crucial for the main paper, the actual mechanics of proofs is not, and can be safely skipped. The fundamental new result of the Appendix which will be particularly useful is the existence and uniqueness of a normal form tableau for an equivalence class of OT tableaux.[2] Moreover, it is shown that any (finite) tableau can be (effectively) transformed into an arbitrary equivalent (finite) tableau by a sequence of applications of five elementary operations and their inverses. Thus any equivalence-preserving transformation of tableaux can always be translated into a sequence of such elementary operations. That allows us to use the normal form as a representative for the whole equivalence class, and exploit its properties in the proofs in the main part of the paper.

With the theory of tableau equivalence classes in place, we can start our investigation of sets of rankings. But as we already noted, attacking the properties of sets of rankings head on is not easy. This is why we develop the general theory of sets of rankings in two subsequent steps. We start with the theory of partial rankings in Section 4. First in Section 4.1 we define OT-compatibility for partial rankings, and develop a logical framework

[2]The normal form we use is similar to Brasoveanu and Prince's basis for a tableau — an equivalent tableau which represents the information in a more convenient way. The novelty of our approach is mainly in its greater generality. For instance, [Brasoveanu and Prince, 2011] claim that they prove in the manuscript [Prince and Brasoveanu, 2010] the fact that for a *single* given tableau, their basis is unique. In the Appendix, we prove that for a *whole equivalence class* of tableaux, the normal form (a close relative of Brasoveanu and Prince's basis) is unique. This allows us to use the normal form as the "name" for that class.

for working with them, viewing rankings as formulas. That allows us to easily analyze the relation of entailment, and show partial rankings' equipotence to certain sets of total rankings in Section 4.2. That, in turn, helps us to show in Section 4.3 that for each partial ranking, there exists a unique normal form tableau, called **representative tableau**, which contains exactly the same information about domination relations between constraints. In Section 4.4, we show that the representative tableau for a partial ranking allows us to easily find the set of all tableaux with which the ranking is compatible. That result will prove to be extremely useful for our analysis of equivalence classes of sets of rankings further on. Finally, in Section 4.5, we define the operations of ranking-intersection $\cap_r$ and ranking-union $\cup_r$ on partial rankings, which, again, will later help us to define more complex operations on sets of rankings.

In Section 5, we at last turn to sets of rankings. We introduce the notion in Section 5.1. No two individual partial rankings are equivalent, but that is not so for sets of rankings. In fact, equivalence classes of sets of rankings turn out to be quite large, and we provide some useful notions for working with such equivalence classes in the same Section 5.1. Then in Section 5.2 we discover one-one correspondence between arbitrary tableaux with special sets of rankings, called **proper sets**. The total refinements of a proper set of rankings are all and only total rankings compatible with its corresponding tableau. Thus the one-one correspondence essentially provides the answer for the general form of the OT Ranking problem, and allows one to build an OT learner which does not overcommit. In Section 5.3 we show that each proper set can be used as the "name" for its equivalence class, and develop the theory of those classes, in particular showing how to compute for any set of rankings its equivalent proper set, and how to determine if two sets of rankings are equivalent. Finally, in Section 5.4 we discuss the algebraic structure of the domain of proper sets of rankings, within the larger domain of all sets of rankings.

Section 6 concludes the paper, discussing future directions for research.

Presenting the material, I tried to simultaneously have in mind two categories of readers: those who are eager to grasp the useful practical results, but are not particularly interested in formal details, and those who find the methods not less valuable than the results themselves. The material which is aimed mostly at the second kind of reader is given in smaller font size. Hopefully, skipping it should not prevent a reader of the first kind from enjoying the fruits of the investigation. Trying to find the right balance between density and verbosity, I have tried to be more explicit in simple proofs closer to the beginning of the paper, and then gradually increase the level of density. Each of the two sections forming the core of the paper, Sections 4 and 5, ends with a short review of its main results.

2 Harvesting compatible rankings, and overcommitment

It is trivial that for any non-contradictory OT tableau there exists a number of OT rankings compatible with it. Each such ranking is viewed by the classical OT as a possible grammar. Given that in the usual analytical situation of analyzing a language, only data are directly accessible, the task of finding a ranking — an OT grammar — which is consistent with the data becomes imperative, and was addressed from the early days on. In particular, a number of different algorithms were proposed for finding some total ranking consistent with a given tableau, most well-known of which are perhaps the Recursive Constraint Demotion (RCD) of [Tesar and Smolensky, 1996] and the (Minimal) Gradual Learning Algorithm (GLA) of [Boersma, 1997]. But there were no algorithms proposed for finding the set of all rankings compatible with the given set of data — the task which can be called the general form of the OT Ranking problem, with the narrower version of the problem being finding at least one compatible ranking.

It is not that the solution for the general form of the Ranking problem was simply not needed. Consider the following tableau:

(4)

	$C1$	$C2$	$C3$
a ~ b	W	W	L

The RCD algorithm, if run on this tableau, produces the stratified hierarchy $\{C1, C2\} \gg C3$. Whichever total ranking this stratified hierarchy is refined to, that ranking will have to be compatible with the following tableau as well:

(5)

	$C1$	$C2$	$C3$
a ~ c	W	e	L
a ~ d	e	W	L

Unfortunately, this constitutes overcommitment: a ranking like $C2 \gg C3 \gg C1$ is compatible with 4, but not with 5. The RCD algorithm thus unnecessarily restricts our search for the right grammar compatible with the data, and misses the difference between non-equivalent tableaux 4 and 5.

In fact, no stratified hierarchy in general can distinguish between the tableaux in 4 and 5. So RCD is not at fault: any algorithm which outputs a single stratified hierarchy is bound to overcommit. Only a set of rankings can represent faithfully the information content of 4.

This problem was discovered early on, and a practical solution emerged. [Tesar, 1997] already recognizes the need to store (what we would have now called) a comparative tableau with the data as the representation of the current knowledge about the underlying grammar, as opposed to storing a ranking. A specific ranking compatible with the stored tableau can then be generated from scratch by running RCD on the tableau at each step of learning, after new data are added. Conceptually, that might seem a strange choice: on the one hand, we say that the grammar is a ranking, but on the other, we cannot actually use anything like a ranking as our grammar hypothesis. From the practical point of view, however, there was no better alternative at hand, because any single stratified hierarchy could commit to more domination relations than the data actually required.

The development of the theory of entailment and equivalence for OT rows and tableaux (see an early attempt to address the problem in [Hayes, 1997], and [Prince, 2006]'s extensive commentary on, and improvement of that early result) has allowed one to manipulate tableaux effectively and thus to get more mileage out of the idea that tableaux should be stored instead of rankings. In particular, the development of the ERC theory of [Prince, 2002] was a radical leap forward, and created the groundwork for the breakthrough of [Brasoveanu and Prince, 2005], who discovered the Fusional Reduction (FRed) algorithm transforming an arbitrary comparative tableau into a particularly convenient, compact equivalent form (a *basis*, in their terms.) Thus any collection of comparative rows could be distilled into a convenient representation easy to work with.

Having such technical inventory in hand, switching from the view "an OT grammar is a total ranking", whose deficiencies from the practical perspective of grammar learning were evident, to the novel view "an OT grammar is a comparative tableau, and rankings compatible with it can be computed as tools, not as a goal in itself" was only logical. The latter position, while not mainstream at the moment, is advocated by Prince, [Prince, 2010]. More broadly, using a tableau essentially as the grammar hypothesis, without necessarily committing to the position that the tableau, or ERC set, *is* the grammar, is a frequently made choice within the OT learning literature.

These historical developments underscore the fact that while it is in principle clear that for a tableau there is a set of all OT rankings compatible with it, working directly with such sets — for instance, using them as grammar hypotheses in learning, — has never been seriously attempted,[3] though certain properties of such sets have been hunted for. As an example of such a property, it is very natural in the OT setting to ask which pairwise domination relations between constraints are necessary given the data. For a tableau M, $C1$ necessarily dominates $C2$ if and only if any ranking compatible with M says that $C1 \gg C2$. As another example, [Riggle, 2008] uses the cardinality of the set of rankings compatible with a tableau to determine which way his learning algorithm should overcommit.

There are at least two good reasons for why it has not been tried to work with sets of all rankings compatible with the data directly. One reason, already discussed above

[3] That is not to say that every proposed OT learner was designed to output a single total ranking or a stratified hierarchy instead of a complete hypothesis given the data. There have been learners proposed which output something else, e.g., the probabilistic learner of [Jarosz, 2006] which outputs a probability distribution over all rankings.

[Brasoveanu and Prince, 2005] briefly discuss how their FRed algorithm can be used to harvest the set of all rankings compatible with the input tableau as a side result, but do not further study the properties of such sets — for instance, they do not distinguish the sets of rankings which can be the output of such algorithm, and those which cannot.

with the help of the example in 2, is that such sets are very complex objects, and unless their representations can be simplified, it is hard for humans to think in their terms. The other reason is that it is very hard to approach the task of finding the set of all rankings compatible with the data without being in full control of what the data are: namely, without being able to identify which datasets are OT-equivalent and which are not.

The current paper overcomes both problems, and develops the theory of sets of all rankings compatible with a given tableau. We have already hinted in Section 1 how we will overcome the first problem. The solution of the second problem was enabled by the pioneering work of Prince and Brasoveanu, on which we heavily build when we develop a full theory of equivalence classes of OT tableaux in Appendix A. The next section gives without proof the main results of that theory which we will heavily use in the subsequent sections.

The theoretical machinery developed in this paper makes it a trivial task to build for a given set of data the set of all OT rankings compatible with them. Therefore we can replace the cherry-picking algorithms outputting only a single ranking or stratified hierarchy, such as RCD or GLA, with algorithms which do not overcommit. Using the methods to be developed below, we can go back and forth between tableaux (modulo OT-equivalence) and the sets of all rankings compatible with them, essentially treating the two kinds of objects as two sides of the same coin. Building a hypothesis from the data is no longer harvesting, it is simply translation from one form into the other.

Of course, we still may want to design learning algorithms which overcommit in certain ways — e.g., place markedness constraints on top of faithfulness ones rather than vice versa, etc. But now the choice will never be only between learning with "wrong bias" and "right bias", as it was the case when overcommitment was inevitable for all known algorithms. Since it is possible to never overcommit, while studying bias, we will also have to justify with evidence the choice between learning guided by some bias and baseline, conservative learning without any overcommitment.

3 Equivalence classes of OT tableaux: a quick tour

We will often use the notions of a "possible row" and a "possible tableau". A possible row for the purposes of this paper is simply any tuple of W-s, L-s and e-s of the appropriate length, and similarly for tableaux. Of course, many, if not most, actual choices of Con and Gen would make some of our "possible rows" impossible: for instance, if whenever some $C1$ is violated, then $C2$ must be violated as well, and $C2$ can only contain a single violation, then we can never have a row with an L in $C1$, but a W in $C2$. We thus work with the space of all *logically possible*, all *conceivable* tableaux rather than with all tableaux possible given a fixed choice of Con and Gen. Our analysis creates a baseline with which the way a particular Con and Gen restrict the range of possible tableaux can be compared.

We will frequently use "inline notation" for rows, where (W, L, L, L) is a 4-constraint comparative row. For a row r, we denote by $L(r)$ the set of constraints where r has an L; $W(r)$ is the set of constraints with a W in them. Thus we can define a new row r_1 as follows: $W(r_1) := \{C1, C3\}$, $L(r_1) := \{C4\}$, for a constraint set $\{C1, C2, C3, C4\}$, defines the row $r_1 = (W, e, W, L)$.

Since we will be studying primarily rankings and sets of rankings in this paper, we will not be interested in the distinctions between tableaux which are OT-equivalent — that is, such tableaux with which exactly the same rankings are compatible. Rather than dealing with individual tableaux, we will be mostly dealing with **equivalence classes of tableaux**. However, in order to do that, we need to be able to handle such classes. As a bare minimum, we need to be able to check if two tableaux are equivalent or not. Of course, this can be checked by brute force: we can check whether two tableaux disagree on any particular ranking by trying all of them. But a less brute force way of checking, and, moreover, one not relying on checking compatibility with individual rankings. Furthermore, working with equivalence classes becomes much easier if each class has a name, or a representative.

The fact that some tableaux are equivalent is, of course, not new, and has generated a lot of work on equivalency-preserving transformations, see, e.g., [Hayes, 1997] and [Brasoveanu and Prince, 2005]. But the focus of such work has so far been on simplifying tableaux emerging from analytical practice, with the goal of making the information in the tableau more accessible to a human analyst (and Brasoveanu and Prince solve the problem, reaching what arguably is the best possible level of simplification). The transformations which do not necessarily simplify a tableau were not systematically studied. As for the general question of which arbitrary tableaux are equivalent and which are not, little has been done except for the basic semantic characterization saying that equivalent tableaux are compatible with the same set of rankings.

Appendix A fills the existing gap, and satisfies our needs in the current paper. First, we build an inventory of five operations which provably preserve OT equivalency and their inverses. Most of the results concerning those transformations as such are not new, and can be either taken from, or straightforwardly generalized from, the results already reported in the literature, most notably, in [Prince, 2002] and [Prince, 2006]. For some of the familiar results we provide alternative, semantic proofs which highlight the other side of the phenomena involved, and complement the original syntactic proofs from the literature, but in terms of the results themselves, there is little new.

What is new is how we put those operations to use: we prove that they are sufficient for transforming an arbitrary tableau into an arbitrary equivalent tableau.[4] The road to this result is through normalization. We define the following class of OT tableaux:

(6) **Normal form for OT tableaux**:

[4] Cf. that with the algorithms of [Brasoveanu and Prince, 2005] which transform an arbitrary tableau into an equivalent tableau in a certain convenient form.

1. The only contradictory tableau in the normal form is the tableau with a single L in the first constraint.
2. Each row has at most a single L.
3. There are no rows which can be inference-eliminated.
4. In multiple-W rows, there are no false W-s.
5. The rows are alphabetically ordered in the following manner:
   - The closer the W-constraints to the first constraint, the higher the row: $[\exists C_i : (C_i \in W(r)) \wedge (C_i \notin W(q)) \wedge \nexists C_j : (C_j \prec C_i) \wedge (C_j \in W(q)) \wedge (C_j \notin W(r))] \rightarrow r \prec q$
   - Among the rows with the same set of W-s (each such group is contiguous by the previous condition), ordering is by the position of the L: $W(r) = W(q) \rightarrow [[\exists C_i \in L(r) : \exists C_j \in L(q) : \wedge C_i \prec C_j] \rightarrow r \prec q]$

Here is an example of a normal form tableau:

(7)

$C1$	$C2$	$C3$	$C4$	$C5$
W	e	L	e	e
W	e	e	L	e
W	e	e	e	L
e	W	e	L	e

But just calling something a normal form does not make it one. The following two crucial results indeed establish the fact that the class defined in 6 has normal form properties:

(8) **Normal form existence theorem**.
An arbitrary (finite) tableau T can be transformed into an equivalent normal form tableau by a (finite) sequence of equivalence-preserving transformations including row swaps, row splittings, inference eliminations, false W eliminations, and contradictory jumps and backward contradictory jumps.

(9) **Normal form uniqueness theorem**.
In each equivalence class of OT tableaux, there is at most one normal form tableau.

Thus there is exactly one normal form tableau per equivalence class, so it can serve as a true representative. Moreover, any tableau may be converted into its normal form by a sequence of applications of our five transformations.

When we discuss those transformations, we introduce them in pairs, so any transformation has its inverse that undoes its result. This allows us to easily get, using 8, the result that any two equivalent tableaux can be converted one into another by a sequence of application of our five operations. The recipe is simple: we convert both tableaux into the normal form, which has to be the same since they are equivalent; then we invert one of the sequences that we used, and append it to the other. The resulting sequence converts one of the original tableaux into the other.

Thus our toolkit of five elementary transformations is *functionally complete*: any equivalence-preserving operation of a given tableau may be represented as a certain sequence of applications of those five operations. Among other things, this means that having a single tableau[5] from an equivalence class, we can easily enumerate the members of this class, since sequences of applications of the five transformations can be enumerated. Several other useful consequences of the normal form results are suggested in A.3.

Returning to our purposes in the rest of the paper, the result of Appendix A which will be used over and over is that a normal form tableau may be used as the representative of its equivalence class. Also, we will often sloppily talk of tableaux when we actually mean their equivalence classes. With the machinery introduced in Appendix A, this is harmless terminological sloppiness.

4 Partial rankings

Our ultimate goal is to study the OT behavior of sets of rankings, but as the first step we define and study individual partial rankings, which turn out to be equivalent to sets of total rankings of a certain kind. We will view rankings as formulas in a logical language **OTR**. Truth for these formulas coincides with the usual notion of OT compatibility, naturally and conservatively extended from total rankings to all rankings.

A partial ranking is true in a tableau iff all its total ranking refinements are true in it. Thus partial rankings may be viewed as abbreviations for sets of their total refinements, or, more accurately, the partial ranking and its set of total refinements turn out to be two different ways of representing the very same OT object.

Studying entailment between partial rankings, we find out that when Con and Gen are such that every logically possible comparative row can actually arise in a tableau, a ranking ϕ entails a ranking ψ only if ψ is ϕ's refinement. For any ranking ϕ, there exists a special representative tableau M_ϕ such that ϕ and its refinements are true in it, but no other ranking is. ϕ and M_ϕ contain exactly the same information (modulo OT equivalence), and can be effectively transformed one into another (if finite). Thus if M_ϕ is the set of data a learner observed, partial ranking ϕ is the faithful grammar hypothesis which rules out all the rankings not compatible with the data, but does not rule out more than is required by them. The set of total refinements of such ϕ includes all and only total rankings compatible with M_ϕ.

Representative tableaux for partial rankings form the domain $\mathcal{M}^-$, a subdomain of the whole domain of (equivalence classes of) OT tableaux, and there is a simple criterion for membership in $\mathcal{M}^-$: if the normal form representative of an equivalence class of tableaux only contains rows with at most one W and one L, then that equivalence class corresponds

[5] As long as that tableau is finite.

to a partial ranking. We call such normal form tableaux non-disjunctive. The domain of partial rankings Φ and the domain of non-disjunctive (equivalence classes of) tableaux $\mathcal{M}^-$ are in one-one correspondence with each other. On the other hand, all sets of data with genuinely disjunctive normal forms are such that a single partial ranking can never be a faithful grammar hypothesis for them. We will discuss such disjunctive tableaux and faithful hypotheses for them later, in Section 5.

Preparing the groundwork for the analysis in Section 5, we discuss two other sets of issues. First, we build a correspondence between non-disjunctive representative tableaux in $\mathcal{M}^-$ and sets of tableaux, so that the tableau M_ϕ corresponds to the set σ_ϕ of all tableaux where ϕ is true. This correspondence can be used right away to define which sets of data are compatible with a current grammar hypothesis in the form of a partial ranking, but the real reason we build that correspondence is that it can be easily extended to disjunctive representative tableaux as well, and that will become very useful in Section 5.3.

Second, we provide a simple analysis of the domain of partial rankings Φ as an algebraic structure, and define two important operations of ranking-intersection $\cap_r$ and ranking-union $\cup_r$. Again, $\cap_r$ and $\cup_r$ are interesting in their own right, but the real reason we introduce them is that they, especially $\cup_r$, will become extremely useful in Section 5.

Most of the results in the current section are quite straightforward, but we discuss most of them nevertheless, stating only a few of simplest results without proof. This allows us to demonstrate the techniques we will use when we deal with significantly more complex objects, sets of rankings, in Section 5. In general, this section serves two goals: first, it lays the foundation, providing us with simple results for partial rankings which will be used as building blocks in our analysis of sets of rankings; second, the form of the results of this section and the methods of analysis are often parallel to what we will develop in Section 5. After those easier results are derived, it will be easier to approach the results harder to understand concerning sets of partial rankings.

4.1 Logic for partial rankings

From the logical perspective on the matter, we can view rankings as formulas in a certain language, and OT tableaux as models in which those formulas may be true or false. (Hence throughout the paper we will use for tableaux names such as M, N, ..., which conventionally refer to models.) The notion of a ranking's truth at a tableau is simply the usual OT notion of the ranking's compatibility with the tableau (a ranking is compatible with a tableau iff it selects the designated winner candidate.) By the end of this subsection, we will have developed a formal logical analysis of OT rankings which will captures that intuition.[6]

[6]Strictly speaking, developing a logical analysis is not a necessity. Some formalization is needed, of course, but one could have done without the logical notions. One and the same mathematical object — e.g., a ranking in OT — may be viewed from different perspectives. But since the language of rankings may be analyzed as a logical language, it *is* one. The reason I chose to use a logical presentation is to highlight the place of the logic of OT rankings among other logical languages, and to draw on familiar logical notions.

But before we define formally the language of OT rankings, we discuss total and partial rankings as such.

Just as [Prince, 2002]'s investigation of the space of sets of ERCs (E(lementary) R(anking) C(ondition)s, essentially comparative rows) starts with entailment between ERCs, so will our discussion of OT (partial) rankings start with entailment between rankings. Informally, entailment is a relation which holds between some entities A and B (e.g., A and B may be rankings) when in every "situation" where A is "present" (true), B has to be "present" (true) as well. The "situation", in our case, is a tableau, and a ranking or a set of rankings A is "present" in a tableau M if A is compatible with M.

Given the logical interpretation where rankings are formulas, and tableaux are models, our informal notion of entailment is just the usual logical entailment: one ranking ϕ entails another ranking ψ iff in all models (that is, in all possible tableaux) in which ϕ is true ψ is true as well. Or, using the traditional notion of OT compatibility, ϕ entails ψ iff ψ is compatible with all tableaux ϕ is compatible with. (In what follows, we freely switch between the truth talk and the compatibility talk, for the two are equivalent.) If ϕ entails ψ, we write $\phi \vDash \psi$.

If $\phi \vDash \psi$, then ψ, the entailed ranking, is compatible with every possible tableau that ϕ is compatible with, and then maybe with some more. So the rankings which entail very many rankings are compatible with very few tableaux. In fact, a ranking which is not compatible with *any* tableau vacuously entails every other ranking.

In the classical OT [Prince and Smolensky, 1993], the grammar is a *total* ordering of a fixed set of constraints Con. Unfortunately, the set of total rankings is not very interesting to study — it has too little structure. In general, total OT rankings do not asymmetrically entail each other, unless the constraints in the specific Con we chose are formulated so that one constraint can only be violated if some other constraint is (or is not), but not vice versa. And as we have said earlier, the study of extra structure imposed by a specific choice of Con falls outside the scope of this paper.

Let us demonstrate why total rankings cannot asymmetrically entail one another. Take some arbitrary total rankings ϕ and ψ which disagree on how to rank constraints $C1$ and $C2$: ϕ says $C1 \gg C2$, and ψ says $C2 \gg C1$. Other constraints are ranked arbitrarily.

Then there may be some tableaux which are compatible with both ϕ and ψ, and some tableaux compatible with neither (neither ϕ nor ψ selects the winner right.) But since there is this disagreement between the two rankings, in the tableaux which crucially require the ranking $C1 \gg C2$ only one of ϕ, ψ will be true, and similarly for tableaux crucially requiring that $C2 \gg C1$. For instance, a comparative row r with a W in the cell corresponding to $C1$, an L in the $C2$ cell, and e-s in all other cells, crucially requires any ranking it is compatible with to say that $C1 \gg C2$. r is thus compatible with ϕ, but not with ψ, and similarly we can build a row q which is compatible with ψ, but not ϕ.

So neither of ϕ and ψ entails the other. As our choice of ϕ and ψ was arbitrary, this means any two distinct total rankings will not entail one another. All that we needed to show that was the possibility to build the rows r and q. Obviously not all constraint sets are such that for any two constraints, the row with a W in the first and an L in the second constraint may arise, but in the space of all logically possible tableaux, it is always possible to build such rows.

The domain of *partial* rankings is more interesting than the domain of total rankings: there may be more relations between its inhabitants. For instance, unlike total rankings, partial rankings can asymmetrically entail other rankings. A formal definition of compatibility for partial rankings we will use is given in 10:

(10) A ranking ϕ is **OT-compatible with** a row r iff for every L in r in cell Ci, there is a constraint Cj dominating Ci in ϕ such that Cj has a W in r.

10 is a straightforward, conservative generalization of the notion of compatibility of the classical OT. The only new thing which arises then we move to partial rankings is that while for a total ranking to be true in some row r there must be some *single* W covering all the L-s, for a partial ranking it is not guaranteed that there will be a single W covering all the L-s, even though every L still has to be covered by some W.

One kind of a partial ranking familiar from the literature is a *stratified hierarchy* (as opposed to a non-stratified hierarchy which is a total ranking), [Tesar and Smolensky, 1996]. In a stratified hierarchy, each constraint may be un-ranked only with respect to the constraints in the same stratum with it, and each constraint in a stratum has to be ranked exactly the same way with respect to all the constraints outside of the stratum. So a stratified hierarchy may be thought of as a total order on a set of subsets of CON given by some partition.

It is easy to construct a partial ranking which is not a stratified hierarchy: e.g., a ranking which ranks $C1$ over $C2$, but does not rank $C3$ with respect to either is not a stratified hierarchy (if we were to treat $C3$ as a part of some stratum A, than the stratum A would not be ranked with respect to the strata containing $C1$ and $C2$.)

Tesar and Smolensky show that if a stratified hierarchy ϕ is compatible with an arbitrary tableau, then all the total rankings which are its *refinements* — rankings which agree with ϕ on all the pairwise, atomic rankings which are in ϕ, and resolve one way or the other all pairwise rankings which were underspecified in ϕ — have to be also true at the same tableau. In logical terms, a stratified hierarchy entails all its total refinements. The proof generalizes easily to the whole class of partial rankings.

Obviously a partial ranking which is not a total ranking entails its refinements asymmetrically: take some refinement ϕ_1 of a partial ranking ϕ, and some pair of constraints $C1$ and $C2$ which were not ranked with respect to each other in ϕ. Without loss of generality, suppose ϕ_1 says that $C1 \gg C2$. Build a row r which has a W in $C1$, an L in $C2$, and e-s in all other cells. ϕ_1 is compatible with r: ϕ_1 specifically says that $C1$ dominates $C2$, so the L in $C2$ is dominated by the W in $C1$ in r. But ϕ is not compatible with the row: there is an L in $C2$, and there are no W-s to cover it other than the one in $C1$, but ϕ does not include $C1 \gg C2$ by assumption. Thus $\phi \vDash \phi_1$, but $\phi_1 \nvDash \phi$. So unlike in the domain of total rankings, in the domain of partial rankings there *exists* non-trivial structure imposed by asymmetric entailment relations.

But what ontological status do we assign to partial rankings? Should we treat them as legitimate OT grammars? In principle, we do not have to: later on we will show that truth for partial rankings is parasitic on truth for sets of their total refinements. So partial rankings are just abbreviations of such sets. The real question then is whether we are willing to say that OT grammars may be sets of total rankings rather than single rankings. The question should be settled empirically, but it is clear that at least as faithful *hypotheses* about the grammar in learning, sets of total rankings are indispensable, so we should learn how to work with sets of rankings irrespectively of whether we believe sets of rankings may be legitimate grammars.[7]

Note that our treatment of partial rankings is different from both treatments suggested in the literature, that of violation cancellation, and that of optionality.

According to our definition, if $C1$ and $C2$ are not ranked with respect to each other, it simply means that a W in either cannot cover an L in the other.

On the mark cancellation view, if $C1$ and $C2$ are not ranked, they are in a crucial tie, and basically behave as a single super-constraint. This view only makes sense for stratified hierarchies, but not for partial rankings in general. For consider a partial ranking like this:

(11) $C1 \gg C2 \gg C3$
$C4 \gg C5$

$C4$ in this ranking is not ranked with respect to $C1$, $C2$ and $C3$. If we are to interpret mutual non-ranking as imposing a crucial tie, then we have to say for the ranking in 11 there are crucial ties in all three pairs $C1$-$C4$, $C2$-$C4$, and $C3$-$C4$. Presuming that crucial ties are transitive, we derive a contradiction with the fact that $C1 \gg C2 \gg C3$ according to the ranking.

The other proposal is that of [Anttila, 1997], [Anttila and Cho, 1998], a.o., who treat unranked pairs of constraints as creating optionality. Under this view, a grammar with an unranked pair of constraints ranks them one way half the time, and the other way another half.[8] As the result, different output forms may be generated. This view does not derive a contradiction for a ranking like 11. But its treatment of non-ranking is different from ours.

[7] One argument against using partial rankings as grammars is given by [Tesar and Smolensky, 1996, pp. 28-29], and is based on the claim that unless the data are generated by a total ranking, the learning algorithm, meaning their RCD, would not converge. This argument does not apply. As is clear from their example of an endless loop, Tesar and Smolensky presuppose the crucial tie interpretation of mutual unranking between constraints. Our interpretation of mutual unranking is different. RCD run on data generated by a partial ranking in our system does converge. The tableau produced by a partial ranking has all L-s covered by some W or other, and moreover, the W constraints on the very top of domination chains should have no L-s, for otherwise it would be impossible to cover those L-s. Thus RCD will always be able to create a new stratum up to the point where all L-s are covered. Of course, the output of RCD may be quite different from the generating grammar, but that is so even when the data are generated by a total ranking. We will see how to learn faithfully, without overcommitment from the complete set of data describing the language in 57.

[8] When there are more than two unranked constraints, the arithmetics gets more complicated. Each total refinement of the underspecified grammar creates a different tableau for the same input. The frequency of a given output for that input is the ratio of the number of tableaux with that output to the overall number of tableaux generated.

So if unranked $C1$ and $C2$ have a W and an L which is not covered by any other W, on the crucial tie view (if the number of violations of $C1$ and $C2$ was the same), the winner is decided by other constraints ranked lower; on Anttila's view, the designated winner of the row will win some of the time; and on our view, it will always lose.

Our treatment is conservative in the sense that like in the classical OT with only total rankings, an uncovered L always leads to failure. The crucial tie and the optionality treatments of mutual unranking say that under certain circumstances, uncovered L-s are fine.

We are now ready to define the language of OT rankings **OTR**. We analyze rankings as formulas which are true or false at rows of comparative tableaux. A row of a comparative tableau (or equivalently, Prince's ERC) is a point in a model for us; a tableau is a model. We will use both OT and logical terminology interchangeably.

The truth should be set up as follows. A ranking is true at a row whenever it is OT-compatible (see 10) with it, or, to use another term, *explains* it, putting some W, a winner-preferring constraint, on top of every L, loser-preferring constraint. A ranking is true in a tableau when it is compatible with every individual row in the tableau; that is, when there are no unexplained L-s in the whole tableau.

The basic building blocks for our formulas are *atomic rankings* which relate just two constraints. Truth conditions for whole rankings are defined as a function of the values for atomic rankings mentioned in the whole ranking at a point — just as truth conditions for propositional logic formulas depend only on the valuation for propositional variables used.

An atomic ranking is not an object of OT proper, and should be distinguished from a whole partial ranking only ranking two constraints. If we know just that $C1 \gg C2$, and are not sure about how all other constraints are ranked in some ranking ϕ (where being not sure is not the same thing as being sure the other constraints are crucially not ordered), OT cannot predict much. Suppose, for instance, that we have a row like this:

(12)

$C1$	$C2$	$C3$
W	L	L

Definitely $C1 \gg C2$ explains the L in $C2$, but it does not help with the L in $C3$, and unless we know whether there is another atomic ranking in ϕ which takes care of that L, we cannot tell if ϕ is compatible with our row. On the other hand, knowing that $C1 \gg C2$ tells us something about the L in $C2$: we know we can no longer worry about it, as it is covered by a W in $C1$. So on the global level, we cannot say much, but on the local level, we can say something, as long as it only concerns $C1$ and $C2$.

In our logic, the global is defined in terms of the local. On the local level, we have atomic rankings which depend only on the exact content of the two constraints they relate to each other, and truth conditions for atomic rankings which tell us whether we should worry about the content of those two constraints in a given row. On the global level we use atomic rankings as atomic building blocks for whole rankings, or formulas. We will define formulas in such a way that for every constraint in a fixed Con, they tell us whether it is dominated at all, and if yes, then by what. So a well-formed formula will have to say precisely which relation holds between any two constraints: whether they are ranked one way or another, or not ranked. It is not allowed for a formula to be "not sure" about the relation of any two constraints. Truth conditions for full formulas will depend on, and only on, the truth of atomic rankings which are mentioned in them.

Here is our definition of the syntax of the language of OT rankings $\mathbf{OTR_{Con}}$, in a signature determined by the choice of the set of constants naming constraints **Con**:

(13) a. **Constraints (terms of first level)**: an at most denumerable fixed set **Con** of constraint symbols $C1, C2, ...$ and the special symbol $\varnothing$.

b. **Atomic rankings (terms of second level)**: all expressions of the form $Ci \gg Cj$, where Ci and Cj are in **Con**.

c. **Rankings (formulas)**:

1. $\Lambda = \bigwedge(Ci \gg \varnothing)^9$, where Ci ranges over all constraints in $\mathbf{Con} \setminus \{\varnothing\}$.
   Λ (or, more precisely, $\Lambda_{\mathbf{Con}}$) is the minimal possible formula in $OTR_{\mathbf{Con}}$: an empty ranking.
2. If ϕ is a formula, and $Cj \gg Ci$ does not appear in ϕ (a condition preventing contradictory rankings), then $Tr(\phi \wedge (Ci \gg Cj))$ is a formula, where $Tr(\phi)$ is the set of the atomic rankings transitively closing the ordering ϕ.[10]

According to this definition, if **Con** is $\{C1, C2, C3\}$, then 14 is a formula of $\mathbf{OTR_{Con}}$.

(14) $(C1 \gg C3) \wedge (C1 \gg \varnothing) \wedge (C2 \gg \varnothing) \wedge (C3 \gg \varnothing)$

Strictly speaking, formulas as objects of our logic are not rankings as such. But there is a natural correspondence between formulas and rankings: a formula has an atomic ranking $Ci \gg Cj$, if $Cj \neq \varnothing$, just in case its corresponding ranking says that Ci dominates Cj. Obviously, there will be many formulas corresponding to the same ranking (for instance, any formulas differing only in the order of occurrence of atomic rankings are mapped to the same ranking), but we are not interested in the syntactic differences between equivalent formulas; we will thus sloppily call our formulas partial rankings.

The only constructor we use to build formulas is $\wedge$, and we will set up truth for rankings so that it is sensitive only to the presence or absence of a particular atomic ranking, never to the presence of multiple instances of a ranking and to the linear positions of atomic rankings in the formula. Thus an atomic ranking occurring twice is just as good as the same atomic ranking occurring once, and all permutations of atomic rankings in a formula generate an equipotent formula. So formulas essentially correspond to *sets* of atomic rankings, and we can talk about atomic rankings being *in* the ranking when they are mentioned in it, which we will heavily exploit in our semi-formal notation.

[9]For a set of atomic rankings and formula Γ, $\bigwedge \Gamma$ means the big conjunction of all elements of Γ. For instance, if $\Gamma = \{\phi, \psi, \xi\}$, then $\bigwedge \Gamma = \phi \wedge \psi \wedge \xi$.

[10]For example, $Tr((C1 \gg C2) \wedge (C2 \gg C3)) = (C1 \gg C2) \wedge (C2 \gg C3) \wedge (C1 \gg C3)$

The definition of truth for atomic rankings is based on the following intuition: an atomic ranking is true at a row iff the two constraints related by it do not present a problem in the OT sense; they either have no offending L-s (are not loser-preferring), or the L that is assessed by the dominated constraint is covered by the W of the dominator constraint. If you do not find the definition natural, you are free to treat it as a technical instrument not supposed to be intuitively appealing: what matters is that this definition will derive the right results for whole rankings once we get to them. The definition is given in Table 1.

Table 1: Truth for atomic rankings

Ci	Cj	$Ci \gg Cj$
W	W	true
W	e	true
W	**L**	true
e	W	true
e	e	true
e	**L**	**false**
L	W	**false**
L	e	**false**
L	**L**	**false**

Truth for formulas in our logic, however, will be more complicated. The first reason for that has nothing to do with truth as such — it is just that even minimal formulas of our logic are quite long. A single atomic ranking like $C1 \gg C2$ is not a proper formula — unlike, say, in propositional logic, where any propositional variable A is a formula, albeit a small one.

A single atomic ranking in our logic is not enough to build a formula because we want the truth of our formulas to depend on the content of *all* cells in a row rather than just on just a *subset* of them: recall the role $C3$ plays in 12 above. So before we actually define truth, we first unwind the definition of formulas we have in 13c, explaining one by one its parts.

First, we need to have a handle on standalone constraints in every whole ranking — constraints which do not participate in any atomic rankings between the actual constraints from Con described in the formula. If such a standalone constraint has a W or e in its cell, it is fine; but if it has an L, it means that the ranking should not be true at the row: there will be an undominated loser-favoring constraint. We define a technical empty "bottom constraint" $\varnothing$. We can think of it as an extra virtual all-e column added to every tableau. We stipulate that any constraint in the constraint set dominates $\varnothing$ in any ranking (the first clause of 13c ensures that for the minimal ranking, and the second clause requires that the minimal ranking be a part of any non-minimal ranking.) Technically, we say that each formula contains an atomic ranking $Ci \gg \varnothing$ for any Ci in the constraint set. We also stipulate that $\varnothing$ has an e in every row. Then if there is an undominated constraint Cj which has an L, the atomic ranking $Cj \gg \varnothing$ will be false, according to the truth definition in Table 1. If Cj has an e or a W, $Cj \gg \varnothing$ will be true.

(15) The cell corresponding to the constraint named $\varnothing$ always has an e. Each constraint set **Con** includes $\varnothing$.

Second, since we are only interested in transitively closed orderings of the constraint set, we stipulate that we only deal with such. In the second clause of 13c, we require that whenever we add a new atomic ranking to a formula, we also transitively close the resulting ordering. We can formally define $Tr(\phi)$, the transitive closure of ϕ, as follows:

(16) $Tr(\phi)$ is the smallest set of atomic rankings such that 1) $\phi \subset Tr(\phi)$, and 2) for any pair of $Ci \gg Cj$, $Cj \gg Ck$ in $Tr(\phi)$, $Tr(\phi)$ also contains $Ci \gg Ck$.

The second clause of 13c ensures that each formula is transitively closed.[11]

Note that we can distinguish **meaningful atomic rankings** — those which cannot be restored if we omit them — and deducible atomic rankings in a formula ϕ. All atomic rankings with $\varnothing$ are included in any formula automatically, so they are deducible. An atomic ranking which is entailed by transitivity by two others also can be recovered if deleted. In what follows, we will heavily exploit that fact, omitting all deducible rankings when we write down formulas, to save space; it should be remembered, though, that our notation is just a shortcut for the real form of our formulas.

Finally, we have a restrictive condition in our definition for well-formed formulas: there is an explicit ban on formulas containing two contradictory atomic rankings. There can be no $Ci \gg Cj$ in ϕ if ϕ has $Cj \gg Ci$, as the second clause of 13c guarantees.

The second clause of 13c ensures that the added atomic ranking $(Ci \gg Cj)$ itself is not contradicted by $Cj \gg Ci$ already present in ϕ, but it does not require directly that atomic rankings added to $\phi \wedge (Ci \gg Cj)$ by transitive closure do not contradict what was in ϕ. In fact, it is not needed, as the precondition excludes the possibility of such contradictions implicitly.

The minimal ranking Λ is clearly non-contradictory. Assume that ϕ is a well-formed formula that is transitively closed and non-contradictory. Suppose a certain atomic ranking in $Tr(\phi \wedge (Ci \gg Cj))$ contradicts some atomic ranking already in ϕ. Since ϕ was transitively closed already, $Tr(\phi \wedge (Ci \gg Cj)) \setminus \phi$ can only have rankings of the form $Ci \gg Ck$, for each Ck such that $Cj \gg Ck$ was in ϕ, and $Cl \gg Cj$, for each Cl s.t. $Cl \gg Ci$ was in ϕ.

Suppose for some Ck, $(Ck \gg Ci)$ were in ϕ, contradicting the new atomic ranking $Ci \gg Ck$ from the transitive closure. But $Ci \gg Ck$ can only be in the closure if $Cj \gg Ck$ was in ϕ. Then from

[11] It is possible to conceive of an OT-like theory where transitivity of ranking would not be respected. In such a system we can have $Ci \gg Cj$ and $Cj \gg Ck$ in our ranking without having $Ci \gg Ck$, and an L in Ck will require a W in Cj, but a W in a higher Ci will not be able to explain the L in Ck. So far there seemed to be no need to stipulate such behavior in phonology, to my knowledge.

the fact $Cj \gg Ck$ and $Ck \gg Ci$ were in ϕ, $Cj \gg Ci$ should have also been in ϕ, as it is transitively closed by assumption; and we excluded the possibility that $(Cj \gg Ci) \in \phi$ by the precondition.

Similarly for Cl, if a contradictory $Cj \gg Cl$ were in ϕ, given that $Cl \gg Ci$ was in ϕ, we would have that $Cj \gg Ci$ were already in ϕ, which is impossible.

By induction, all well-formed formulas are non-contradictory.

The second reason truth for OT ranking formulas is more complex than, say, the truth for the usual propositional logic, is that we have to express a global notion of having each L covered with a W as a condition on the local truth of atomic rankings, which do not even directly refer to W-s, L-s and e-s. What is surprising is not that the definition is complex, but that it is possible at all to define truth for rankings even after collapsing all the information there was in the row — with 9 possible variants for each pair of constraints — to just two values **true** and **false** for atomic rankings.

We first give the definition in the form of an algorithm for determining if a formula ϕ of our language of partial rankings is true, and after that we give an equivalent definition in the usual static form. The algorithm is given to provide a better idea of what the static definition actually says.[12]

(17) **Truth for formulas of $\mathbf{OTR_{Con}}$, in the algorithmic form**
The algorithm takes a formula ϕ and a row r in the same signature **Con** as its input.

1. Take the set UD of all undominated constraints in ϕ, where Ck is undominated in ϕ iff there is no atomic ranking $Cl \gg Ck$ in ϕ.
   (When checking if a constraint is undominated, we should not forget to look at the rankings obtained by transitivity; the $\varnothing$ constraint is always dominated, because of the first clause of the definition for formulas.)
2. For each $Ci \in UD$, check the truth of all the atomic rankings $Ci \gg Cj$ in ϕ in row r. If all of them are false for some Ci, the formula is false. Otherwise proceed to the next step.
3. Forget about the atomic rankings with undominated constraints which are true, and turn to the ones which are false — they require attention.
4. For each false ranking $Ci \gg Cj$ which requires attention (that is, $Ci \in UD$, and $Ci \gg Cj$ is false in r), try to find a *covering ranking* $(Ck \gg Cj) \in \phi$ which is true in r. If there is one, forget the false ranking which required attention — we took care of it; if not successful, the whole formula is false.
5. If there are no false atomic rankings from the previous step left uncovered, the formula is true.

[12]The algorithm is not designed to be efficient. Its purpose is not to compute truth fast, but simply to define truth in a relatively reader-friendly manner. A more efficient equivalent algorithm is not hard to design, if need be.

And here is the static condition, which, as is easy to check, holds precisely when the algorithm in 17 says ϕ is true at r:

(18) **Truth for formulas of $\mathbf{OTR_{Con}}$**

For a formula ϕ and a row r, ϕ is true at r iff

$$\forall U \in \mathbf{Con} : (\neg\exists Ck \in \mathbf{Con} : (Ck \gg U) \in \phi) \rightarrow$$
$$(\exists Cl \in \mathbf{Con} : (U \gg Cl) \in \phi \ \wedge \ r(U \gg Cl) = 1) \ \wedge$$
$$(\forall Ci \in \mathbf{Con} : (U \gg Ci) \in \phi) \rightarrow$$
$$(\exists Cj \in \mathbf{Con} : (Cj \gg Ci) \in \phi \wedge r(Cj \gg Ci) = 1),$$

where $r(Ci \gg Cj) = 1$ iff the atomic ranking $Ci \gg Cj$ is true in row r.

In words, for all constraints U, if there is no Ck dominating U in ϕ (that is, for all U which are undominated in ϕ), there is always some Cl dominated by U in ϕ such that $U \gg Cl$ is true in r, and furthermore, for all Ci dominated by U, there is some Cj which dominates Ci in ϕ, and $Cj \gg Ci$ is true in r. (This Cj may be U itself, or some other constraint.)

It immediately follows from our definition that the truth of a formula is a function of the truth of the atomic rankings featured in it.

(19) A ranking ϕ is compatible with a row r according to 10 just in case the algorithm in 18 outputs **true** for the pair of ϕ and r.

Proof of 19. We fix an arbitrary ranking ϕ and an arbitrary row r, and show that ϕ is true in r iff each L in r is covered by a W.

Suppose there is a constraint which has an L in r. There are two cases: either it is undominated, or it is dominated in ϕ.

If there is an undominated constraint U with an L, the ranking ϕ is not compatible with r. In that case any atomic ranking of the form $U \gg X$ will be false in r[13], and thus the subformula $(\exists Cl \in \mathbf{Con} : (U \gg Cl) \in \phi \ \wedge \ r(U \gg Cl) = 1)$ of the definition is false, which makes the whole definition false. So 10 and 18 agree in this case.

Now suppose there is an L in a dominated constraint Ci. By 10, for ϕ to be true in r, there must be a Cj s.t. $(Cj \gg Ci) \in \phi$ and Cj has a W in r. Since Ci is dominated by hypothesis, there is some undominated U such that $U \gg Ci$ is in ϕ. But then the second conjunct of the consequent guarantees that there is a Cj dominating Ci such that $Cj \gg Ci$ is true in r. Since Ci has an L, $Cj \gg Ci$ can only be true if Cj has a W. Thus the L in Ci has to be covered by some W if 18 is true.

By induction on constraints with L-s, we get that when ϕ is declared true by 18 exactly when all L-s are covered with some W or other. This establishes the equivalence of 10 and 18, so our formal truth definition correctly captures the notion of compatibility. ⊣

[13]Note that in order for all atomic rankings with U to be false, $U \gg \varnothing$ has to be false — which can only be when U has an L. If U has an e, but all of its dominated "real" constraints have L-s, and thus all "real" atomic rankings with U will be false, the ranking $Ci \gg \varnothing$ will still be true.

The special case of a row with only e-s in it deserves a few words. Two candidates used to generate such a row are equally harmonic, they do equally on every constraint. [Tesar and Smolensky, 1996, p. 15, fn. 8] treat both of them as equally harmonic, and declare they should be in free alternation. Algorithms such as RCD view an all-e row as vacuously compatible with any ranking. This behavior agrees with the interpretation of free alternation: it means that any ranking at all does not prefer one candidate to the other, and hence both are equally fine, so they are expected to be in free alternation in case either of them is more optimal than any other candidate. Our 10 and 18 make the same prediction for such all-e rows as Tesar and Smolensky.

4.2 Entailment between partial rankings

As soon as we have defined the notion of truth, we can formally speak of entailment between our formulas. A ranking ϕ entails a ranking ψ iff at every possible comparative row/tableau where ϕ is true, ψ is true as well. If ϕ entails ψ, we write $\phi \vDash \psi$.

(20) For rankings ϕ, ψ, $\phi \vDash \psi$ iff ψ is true in (compatible with) every tableau ϕ is true in (compatible with).

If we have an arbitrary ϕ and some ϕ' which is a refinement of ϕ — that is, which has all atomic rankings that ϕ has, and then some more — then ϕ necessarily entails ϕ'.

(21) ϕ' is a **refinement** of ϕ iff $\forall Ci, Cj : (Ci \gg Cj) \in \phi \rightarrow (Ci \gg Cj) \in \phi'$.
If ϕ' is a refinement of ϕ, and $\phi \neq \phi'$, then ϕ' is a **proper refinement** of ϕ.

(22) If ϕ' is a refinement of ϕ, $\phi \vDash \phi'$.

Proof of 22. If $\phi' = \phi$, then trivially $\phi \vDash \phi'$. So consider the case when $\phi' \neq \phi$. By 21, there are some atomic rankings $Ci \gg Cj$ which are in ϕ', but not in ϕ, and otherwise ϕ' and ϕ agree. We show by induction that $\phi \vDash \phi'$ in this case.

Let ϕ be true at some fixed row r, and add an arbitrary atomic ranking $Ci \gg Cj$ to ϕ obtaining $Tr(\phi \wedge (Ci \gg Cj))$.

If both Ci and Cj are dominated in ϕ, the new atomic ranking cannot, by 18, make the resulting formula false, as that addition does not change the set of undominated constraints of ϕ. So $Tr(\phi \wedge (Ci \gg Cj))$ is true as well. (Had ϕ been *false* at r, adding $Ci \gg Cj$ might have changed the truth of $\phi \wedge (Ci \gg Cj)$: the newly added $Ci \gg Cj$ might have helped to cover an L in Cj or some of the constraints dominated by Cj.)

If we add $Ci \gg Cj$ where Cj is undominated in ϕ, we have made the set UD smaller, and have added by transitivity a new atomic ranking $Ck \gg Cj$ for each undominated Ck such that $Ck \gg Ci$ is in ϕ. Making the set UD smaller cannot make the formula false by 18 false, but newly added atomic rankings in principle could. However, as Cj is undominated in ϕ, and ϕ is true, Cj cannot have an L: otherwise ϕ would be false. Any $Ck \in UD$ also cannot have an L, for the same reason. But then all new rankings $Ck \gg Cj$ are true, so $Tr(\phi \wedge (Ci \gg Cj))$ is true.

If we add $Ci \gg Cj$ where Ci is undominated in ϕ, then from the assumption, Ci cannot contain an L: otherwise ϕ would be false. If Ci has a W, we are immediately fine. If it has an e, the new atomic ranking may be false in the case Cj has an L.

Suppose towards a contradiction that Cj has an L. If the offending Cj was undominated in ϕ, ϕ would have been false, contrary to our assumption. So a Cj with an L could not have been undominated in the original formula. On the other hand, if Cj was dominated in ϕ, there was at least one atomic ranking $Ck \gg Cj$ with Ck undominated. Now if $Ck \gg Cj$ is true, we are fine: it covers the false new ranking $Ci \gg Cj$. If $Ck \gg Cj$ is false, it required attention, and there must have been some true $Cl \gg Cj$ covering $Ck \gg Cj$. The same $Cl \gg Cj$ will cover the false $Ci \gg Cj$ in $Tr(\phi \wedge (Ci \gg Cj))$ as well. In any case, the new false atomic ranking receives cover, and $Tr(\phi \wedge (Ci \gg Cj))$ has to be true.

By simple induction on the difference between ϕ and ϕ', 22 follows. $\dashv$

Can a refinement ϕ' entail its (grand*)parent ϕ? In general, no, unless Con or Gen have special structure. Take some $Ci \gg Cj$ which is in ϕ', but not in ϕ. Build a counterexample row: put a W into Ci, an L into Cj, and e-s into all other cells. Now, ϕ' is true at the resulting row, but ϕ is not.

So if our Con and Gen do not exclude the possibility of building such a counterexample row, then if ϕ' is a proper refinement, $\phi' \nvDash \phi$.

So for any ϕ and its refinement ϕ', we have that $\phi \vDash \phi'$. What about ϕ and some ψ which is not ϕ's refinement? It turns out neither can asymmetrically entail the other, unless a particular choice of Con and Gen forces that. The argument we made earlier in Section 4.1 about the impossibility of asymmetric entailment among total rankings is essentially a specific case of the proof for 23.

(23) If ϕ and ψ are not refinements of each other, then $\phi \nvDash \psi$ and $\psi \nvDash \phi$.

Proof of 23. As ϕ and ψ are not refinements of each other, there is some $Ci \gg Cj$ in ϕ, but not in ψ. We can build a counterexample row $r_{Ci,Cj}$ for $Ci \gg Cj$ just as we did to show that $\phi' \nvDash \phi$: we put a W into Ci, an L into Cj, and e-s into all other cells. ϕ is true at $r_{Ci,Cj}$, but ψ is not, so $\phi \nvDash \psi$. Similarly for $\psi \nvDash \phi$. $\dashv$.

Again, the argument rests on the assumption we adopt throughout this paper that the specific Con and Gen we use do not preclude the possibility to build the counterexample rows.

Both relations of being a refinement and of being asymmetrically entailed create an ordering between rankings. 22 and 23 tell us that the two orderings coincide (again, unless the specific Con and Gen we use change the picture).

The more underspecified a ranking is, the more rankings are its refinements and are entailed by it. If $\phi \vDash \psi$, that essentially means that ϕ is true at a lesser or equal number of conceivable rows than ψ. In particular, properly partial rankings are true in a lesser number of possible rows than total rankings. This is so because the more atomic rankings we have, the more we have chances to cover the present L-s in a row. E.g., if we have an L in a constraint Ci, and there is no other constraint dominating Ci by our ranking ϕ, there is simply nothing ϕ can do about that L. If there is one constraint dominating Ci in ϕ, we have a chance that it has a W in the row we are considering, though it might not. If there are two such constraints, we have twice as many chances to cover the problematic Ci, and so forth.

Another way to put it is to say that the more underspecified ranking is also the more informative one: it is true in a lesser number of situations; hence if we learn it is true, we get more knowledge about the situation we are in (that is, about how our tableau looks like) than we would if we learned that a less underspecified ranking were true.

Finally, the following important fact holds, in two formulations:

(24) a. If every proper refinement of ϕ is true in a tableau M, then ϕ is also true in M.
b. If every proper refinement adding only one meaningful atomic ranking to ϕ is true in M, then ϕ is true in M.

It is easier to see why the second claim is true. If the refinements of ϕ by both $Ci \gg Cj$ and $Cj \gg Ci$ are true in tableau M, that means pairwise ranking of Ci and Cj is not crucial for accounting for M as long as all other atomic rankings from ϕ are present. Therefore if for each pair of Ci and Cj unordered in ϕ both ways of ordering them makes the result true in tableau M, then the atomic rankings which are already in ϕ are enough to make ϕ true in M. The first claim easily follows from the second by induction.

This is an important fact: 24 means that the truth of a partial ranking may be viewed as *dependent*, or *parasitic*, on the truth of its proper refinements. If all of them are true in some tableau, then the original partial ranking has to be true there. If at least one proper refinement is false, than the ancestor partial ranking is false, by 22.

The following is an important special case of 24 allowing us to view partial rankings as sets of their total refinements:

(25) A partial ranking ϕ is true in tableau M iff every total refinement of ϕ is true in M.

This means that a partial ranking and the set of its total refinements can be used interchangeably. They are equivalent objects, so we can replace everything we say about partial rankings can also be said about corresponding sets of total rankings. It is just that partial rankings are much intuitively easier than sets of their total refinements.

Note that while every partial ranking corresponds to the unique set of its total refinements, not every set of total rankings corresponds to a partial ranking: e.g., the set consisting of $C1 \gg C2 \gg C3$ and $C3 \gg C2 \gg C1$ does not correspond to one.

4.3 Correspondence between partial rankings and tableaux

In this section, we will show that for every partial ranking, there exists a special tableau, from a certain class of tableaux we will call $\mathcal{M}^{-}$, which contains exactly the same amount of information. Thus there is duality between partial rankings and such tableaux: they are essentially two different ways to represent the same object. In itself, this result is of limited usefulness, because the duality is restricted to a very special class of tableaux. But it will serve as a model for a similar result in Section 5, which puts an arbitrary tableau into a correspondence with the set of all rankings true in it.

We call a ranking ϕ *maximal* with respect to a tableau M, or M-*maximal*, iff ϕ is true in M, but none of ϕ's ancestors is.[14] (The mnemonic for the term is "a ranking which is *maximally underspecified*, but still true in a tableau M".)

Sometimes, there will be only one maximal ranking for a given tableau, but not always:

(26)

$C1$	$C2$	$C3$
W	L	e

$C1 \gg C2$ is maximal

(27)

$C1$	$C2$	$C3$
W	L	W

$C1 \gg C2$ is maximal, $C3 \gg C2$ is maximal

From the second tableau above it is clear that in the general case, one tableau may have more than one maximal ranking. But for every partial ranking ϕ, there exists a special tableau where ϕ is the only maximal ranking. In particular, for the ranking $C1 \gg C2$, 26 is such a tableau.

The significance of that tableau is that it is enough to recover the ranking which is maximal in it. So given a ranking, we can build such a tableau, and given that tableau, we can build the ranking again. Thus the two objects contain exactly the same information. We will call such a special tableau containing exactly the same amount of information as a set of rankings the **representative tableau** of that set. We will often write M_ϕ for the representative tableau of a ranking ϕ, or for the equivalence class defined by it. L

The construction procedure of the representative tableau for an arbitrary ranking ϕ is simple (and we have in fact already used its main element when constructing counterexample rows earlier.) For each meaningful atomic ranking $Ci \gg Cj$ in ϕ (that is, not entailed by transitivity, and not of the $Ci \gg \varnothing$ form), build a row $r_{Ci,Cj}$ such that its Ci cell has a W, its Cj cell has an L, and all other cells have e-s. $r_{Ci,Cj}$ is not compatible with any ranking which does not include the atomic ranking $Ci \gg Cj$: the L in the Cj cell must be dominated, and there is only one W to do that in $r_{Ci,Cj}$, the one in Ci, so no other atomic ranking but $Ci \gg Cj$ can make $r_{Ci,Cj}$ happy.

Combine all such rows in a single tableau. Now if we subtract any meaningful atomic ranking from ϕ, the result will not be compatible with the constructed tableau because the row corresponding to the subtracted ranking will become false. Of course, it is harmless to add more atomic rankings, constructing a refinement of ϕ, but we cannot take things away without making the resulting ranking false.

[14] We omit an explicit reference to the tableau with respect to which a ranking is maximal when it is clear from the context.

(28) **Representative tableau M_ϕ for a partial ranking ϕ:**
$$M_\phi := \{r \mid (|W(r)| = |L(r)| = 1) \wedge [\exists (Ci \gg Cj) \in \phi : Ci \in W(r) \wedge Cj \in L(r)]\}$$

The inverse of the procedure is easy to compute, too: given a tableau with only $r_{Ci,Cj}$ rows, we can recover for each such row the corresponding atomic ranking $Ci \gg Cj$. Combining all such atomic rankings, we will get the original ϕ.

There is thus a computable correspondence between partial rankings, and certain tableaux. But which tableaux are those? For instance, if we try to apply the procedure above to the following tableau, clearly we will not be able to build a ranking, because of a contradiction:

(29)

$C1$	$C2$	$C3$
W	L	e
e	W	L
L	e	W

It is easy to see that the range of the function from partial rankings to tableaux that we defined contains only normal form tableaux (see the definition in 6). Indeed, since partial rankings cannot contain contradictory atomic rankings, the resulting tableaux are not contradictory; since we only used meaningful atomic rankings, no rows in the resulting tableau are superfluous; all rows have just one L; finally, there can be no false W-s, for each row only has one W.[15]

Recall that each normal form tableau defines its OT-equivalence class: there is no other normal form tableau equivalent to it (see 139). Since there is only one normal form tableau per equivalence class, our function from partial rankings into equivalence classes is an injection: each partial ranking is mapped to a different equivalence class. Moreover, if we take those equivalence classes whose normal forms only have one-W-one-L rows, the function is also a surjection: every such non-contradictory tableau is the result of applying the procedure to some partial ranking.

Thus our function establish a one-one correspondence between partial rankings and those equivalence classes of tableaux whose normal forms only have one-W-one-L rows.

What does this mean? It is easier to grasp the significance of the correspondence using the learning perspective. Suppose we have a set of data which is equivalent to a normal form tableau M with only one-W-one-L rows. Then the correspondence immediately provides to us a grammar hypothesis which exactly matches the data: the partial ranking ϕ for which M is representative. Indeed, as ϕ is the only maximal ranking for M, ϕ's total refinements are the only total rankings true in M. Thus choosing ϕ as a hypothesis does not rule out any rankings compatible with the data, and does not erroneously rule in any which are not.

[15]Strictly speaking, the order of rows may be different than the order required by the normal form. In the main text, we completely abstract away from the issues of the order of rows, to spare the reader trivial steps in reasoning.

Of course, in a practical setting the problem of finding out whether the data are equivalent to such a special normal form is to be solved before we can actually capitalize on the correspondence. This can be done using tableau transformations based on the methods of [Brasoveanu and Prince, 2005] or Appendix A (though we will see in Section 5 that it is also possible to reason about faithful hypotheses using only rankings, not transforming tableaux.)

Nevertheless the correspondence provides a criterion of when exactly the faithful grammar hypothesis corresponding to a set of data cannot be a single partial ranking, and has to be a set of such rankings: it is precisely when the normal form of the set of data only have one-W-one-L rows. In a sense, it allows us to distinguish data sets which are simply silent on the mutual ranking of certain constraints, and those which imply true disjunctivity. It is not surprising that the criterion has to do with whether there are rows with more than one W: it has long been known that disjunctivity arises from such rows as (W, W, L). But not every tableau with multiple-W rows implies true disjunctivity, thus our finding is a genuine step forward. In what follows, we will call the special normal form tableaux with only one-W-one-L rows *non-disjunctive tableaux*, and their equivalence classes, non-disjunctive classes.

We thus take stock of what we have done in this subsection:

(30) **Non-disjunctive tableau domain $\mathcal{M}^-$**:
In the domain of all OT-equivalence classes of tableaux $\mathcal{M}$, there is a subdomain $\mathcal{M}^-$ of equivalence classes whose normal form tableaux have only rows with exactly one W and one L.

(31) **Correspondence between partial rankings and non-disjunctive tableaux**
There is a correspondence between the domain of all partial rankings Φ and the domain of non-disjunctive (equivalence classes of) tableaux $\mathcal{M}^-$ such that for ϕ and M_ϕ in correspondence, ϕ's total refinements are the only total rankings compatible with M_ϕ. The correspondence is computable both ways (assuming finite tableaux).

Thus the domains Φ and $\mathcal{M}^-$ are dual: ϕ and M_ϕ in correspondence are essentially the same object, only written differently: either as a partial ranking or set of total rankings, or as a tableau equivalent to a non-disjunctive normal form tableau.

The fact that the correspondence is computable, in finite settings, in principle allows us to use whichever of the two objects is more convenient for the present purposes even in practical tasks.

4.4 Duality between partial rankings and sets of tableaux

In the previous section, we have established that a non-disjunctive tableau M can be put in correspondence with the set of all total rankings true in it, which can also be viewed as the only partial ranking maximal in M. In the special case when the data are non-disjunctive (that is, have a non-disjunctive normal form), this allows us to compute the faithful grammar hypothesis which contains exactly the same information as the data did.

But given an adopted grammar hypothesis, we can also ask whether a given set of data is explained by it or not. For instance, the learner might need to check whether its analysis of a new datum is compatible with what it has learned up to this point (and if not, then, say, reanalyze the datum or modify the hypothesis). So it could be useful to find out which sets of data make the current hypothesis true, and which are not. However, if you are not convinced this is more than just an intellectual exercise, in fact there is a hidden agenda in this section: the result we are about to derive will turn out to be easily generalizable into the realm of genuinely disjunctive tableaux as well, and there it will become extremely handy, as we will see in Section 5.3. So now we are not only answering a question concerning partial rankings and tableaux they are true in, but also prepare a nice tool which we will heavily exploit further down the road.

We can trivially define for each partial ranking ϕ the set σ_ϕ of all (equivalence classes of) tableaux that ϕ is compatible with. (Recall that in terms of sets of total rankings, a partial ranking is true in a tableau whenever all total rankings from the corresponding ranking set are true in it.) We will now study what those σ_ϕ sets actually are.

First of all, for any distinct rankings ϕ and ψ, σ_ϕ is not equal to σ_ψ. Suppose ψ is not a refinement of ϕ. Then ψ is not compatible with M_ϕ, the representative tableau of ϕ, so $\sigma_\phi \neq \sigma_\psi$. If, on the other hand, ψ is a proper refinement of ϕ, then $\sigma_\phi \subseteq \sigma_\psi$ by definition of refinement, and we show that the inclusion is proper. We select some $(Ci \gg Cj) \in \psi$ which is not in ϕ, and build the row $r_{Ci,Cj}$. That row is compatible with ψ, but not with ϕ, and the one-row tableau consisting just of $r_{Ci,Cj}$ is in σ_ψ, but not in σ_ϕ. So for ψ a proper refinement of ϕ, $\sigma_\phi \subset \sigma_\psi$. Thus if $\phi \neq \psi$, we have $\sigma_\phi \neq \sigma_\psi$, and the function mapping a ranking to the set of all tableaux it is compatible with is an injection.

Obviously, not all sets of (equivalence classes of) tableaux are σ_ϕ for some ϕ. But we can characterize such σ_ϕ sets easily using the correspondence between partial rankings and their representative tableaux we established in the previous section. Let ϕ be an arbitrary ranking, M_ϕ its representative tableau. We will now characterize the set σ_ϕ of tableaux in which ϕ is true in terms of M_ϕ.

Take an arbitrary tableau N with only one W and one L in each row which has a row r which is not entailed by M_ϕ. By construction of M_ϕ, ϕ can only account for rows in M_ϕ or entailed by it (note that this is entailment between tableaux, not between rankings), so ϕ is not compatible with r, and thus with tableau N as a whole. Without loss of generality, suppose N is in the normal form. N then defines an equivalence class of tableaux, by 139. Then ϕ is not compatible with any equivalence class of tableaux whose normal form N has a row not entailed by M_ϕ.

On the other hand, any tableau N consisting only of rows from M_ϕ and rows entailed by it is bound to make ϕ true. We call M_ϕ^{En} (with En for "(tableau) closed under row (En)tailment") the tableau consisting of all rows entailed by M_ϕ. Then we can define σ_ϕ

in terms of M_ϕ through M_ϕ^{En}:

(32) The entailment closure tableau M_ϕ^{En} is such a tableau that $r \in M_\phi^{En}$ iff M_ϕ entails r.

(33) $\sigma_\phi := \{N \mid N^{No} \subseteq M_\phi^{En}\}$, where N^{No} is the normal form of N.

As an illustration, for the tableau M_ϕ in 34, which is the representative tableau of the ranking $\phi := (C1 \gg C2) \wedge (C2 \gg C3)$, the corresponding M_ϕ^{En} is given in 35.

(34) $M_\phi =$

$C1$	$C2$	$C3$
W	L	e
e	W	L

(35) $M_\phi^{En} =$

$C1$	$C2$	$C3$
W	L	e
e	W	L
W	e	L
W	L	L
W	L	W
W	W	L

Any subset of M_ϕ^{En} is a tableau in σ_ϕ, so even if we restrict ourselves to those which are in normal form, writing σ_ϕ down will take a lot of space.

A procedure for determining whether an arbitrary set of data N makes a hypothesis ϕ true may be as follows: first we can compute its normal form N^{No}, and then check whether it is a subset of M_ϕ^{En}.[16] M_ϕ^{En} can be called **the generating tableau** of the set of tableaux σ_ϕ. If computing the normal form is effective, and tableaux are finite, this is an effective procedure.

σ_ϕ sets contain an infinite number of individual tableaux, but only a small number of equivalence classes of tableaux. Therefore it is better to think of σ_ϕ sets as of sets of normal form tableaux: for a finite Con, M_ϕ and M_ϕ^{En} have to be finite, and thus the number of normal form tableaux in σ_ϕ is finite, too. This allows us to have an effective procedure computing the inverse function, which computes M_ϕ given σ_ϕ: if we have a representation of σ_ϕ in the form of all normal form tableaux in it, we can, for instance, combine all such tableaux into a single one, and then convert the resulting big tableau into a normal form. It will produce exactly M_ϕ.[17]

[16] Note that we cannot just check whether N^{No} is a subset of M_ϕ: for instance, if M_ϕ consists of two rows (W, L, e) and (e, W, L), and N^{No} consists of a single row (W, e, L) entailed by transitivity by M_ϕ, even though M_ϕ entails N^{No}, it is not the case that $N^{No} \subseteq M_\phi$.

[17] Simply combining all the rows of a set of tableaux may result in a contradictory tableau. But all tableaux in σ_ϕ are entailed by M_ϕ. As such, they must be consistent. Therefore no contradiction may arise.

While M_ϕ is a non-disjunctive tableau, the normal form members of σ_ϕ can be disjunctive as well. E.g., for the non-disjunctive tableau in 34, a genuinely disjunctive tableau with a single row (W, W, L) is within the corresponding σ set.

In terms of sets of all total rankings, any tableau N in σ_ϕ which is not equivalent to M_ϕ is compatible with all total rankings M_ϕ is compatible with, and then with some more. If N is non-disjunctive, then it must correspond to a partial ranking for which ϕ is a refinement. If N is disjunctive, then it corresponds to a set of rankings for which the set of total refinements of ϕ is a subset. Naturally, any datum which can be explained by the current hypothesis has to correspond itself to a not less wider hypothesis.

We can easily define the exact domain Σ^- of sets σ_ϕ of tableaux for which there exists a partial ranking ϕ such that all tableaux in σ_ϕ are compatible with it. For now, it does not serve an immediate practical purpose to do that, but in Section 5.3, we will see the importance of (a slightly generalized version of) this characterization.

(36) Set of tableaux σ is in Σ^- iff there is a tableau M^{En} such that:

- M^{En} belongs to a non-disjunctive equivalence class (that is, its normal form is in $\mathcal{M}^-$), and
- for every row r entailed by M^{En}, $r \in M^{En}$,[18] and
- for every normal form tableau $N \in \sigma$, every row $q \in N$ is also in M^{En}, and
- for every normal form tableau $P \subseteq M^{En}$, $P \in \sigma$, and
- σ is closed under OT-equivalency-preserving operations.

Let's work through the clauses of the right side of the equivalency statement. The first clause ensures that M^{En} is generated by some non-disjunctive normal form tableau M_ϕ. The second clause says that M^{En} is closed under row entailment. The third clause says that all rows we can find in the normal form part of σ are included in M^{En}. The fourth clause establishes the other direction: every normal form subset of M^{En} is required to be in σ. Finally, the last clause guarantees that not only normal form tableaux get into σ, but their whole equivalence classes, and at the same time that there are no "stray" non-normal form tableaux that crawled into σ without their normal form being there.

If we simply remove the first clause of 36, we would get a definition of the domain Σ of all sets σ of tableaux such that the tableaux in σ all make true a certain proper OT grammar hypothesis, as we will see in Section 5.

We can sum up our findings about the dualities between three domains Φ, $\mathcal{M}^-$ and Σ^-.

(37) The following three domains are in duality relations with each other, so that each ranking correspondent $\phi \in \Phi$ is maximal for the tableau correspondent $M_\phi \in \mathcal{M}^-$, and is compatible with all tableaux from the set of tableaux correspondent $\sigma_\phi \in \Sigma^-$, and with no other tableaux:

[18] Note that all tableaux in the equivalence class entail the same set of rows.

1. the domain Φ of all partial rankings;
2. the domain $\mathcal{M}^-$ of equivalence classes of tableaux with normal forms having only rows with one W and one L;
3. the domain Σ^- of sets of tableaux characterized in 36.

As we call the element of $\mathcal{M}^-$ in correspondence with a ranking ϕ its **(partial) representative tableau** (or (partial) representative equivalence class of tableaux), we will call the element of Σ^- in correspondence with ϕ its **(partial) representative set of tableaux**. The qualification "partial" does not mean that tableaux and sets are not full representatives; instead it is meant to show that their correspondent is a partial ranking, as opposed to a set of partial rankings (in Section 5 we will see how sets of rankings fit into the general picture.) We omit the word *partial* whenever the context makes it clear whether the correspondent is a ranking or a set of rankings.

Duality between the domains means that the corresponding objects from different domains contain exactly the same information, and that we can work with whichever of the three is most convenient, and then easily transfer the results into two other domains.

The duality results for partial rankings will serve as a model for similar duality results for sets of partial rankings we discuss in Section 5. The latter ones will be harder to get, which is why we have started from fairly straightforward results for partial rankings, in order to make it easier to approach the dualities of real interest for us, in Section 5.

In addition to sheer greater complexity of sets of partial rankings compared to individual partial rankings, there is another important difference between the two domains: as we have seen, every partial ranking has a representative tableau, and thus every partial ranking may serve as a valid OT grammar hypothesis. But for sets of rankings, it is not so: there are sets which do not correspond to any tableau at all. So within the domain of all sets of rankings, only a particular subdomain has direct significance for OT, and it will be useful to study the place of this subdomain within the greater domain. We will do that in Section 5, but before we finally turn to sets of rankings, we will take a closer look at the whole domain of single partial rankings and its internal structure: just as with duality, it will prepare us for a more complex analysis of the structure of the more complex domain of sets of rankings later on; moreover, the operations on individual partial rankings we will define in the next section will serve us well once we get to work with sets of partial rankings.

4.5 The domain of partial rankings Φ as an algebraic structure

We have already characterized the relations partial rankings bear to equivalence classes of tableaux and sets of tableaux in which they are true. But partial rankings (and their dual objects, of course) also have relations between themselves. One of those we have looked at already in some detail: entailment. In this section, we will look more closely into the internal structure of the domain Φ of partial rankings created by different relations that hold between them. First, that will allow us to understand better what sort of creatures partial rankings are. Second, when we take on sets of partial rankings in Section 5, the knowledge we will acquire by the end of this section will become handy.

Here is how the relation of entailment between rankings organizes their domain, for a constraint set with just 3 constraints $C1$, $C2$, $C3$ (recall that Λ is the minimal ranking which only has $Ci \gg \varnothing$ atomic rankings):

(38)

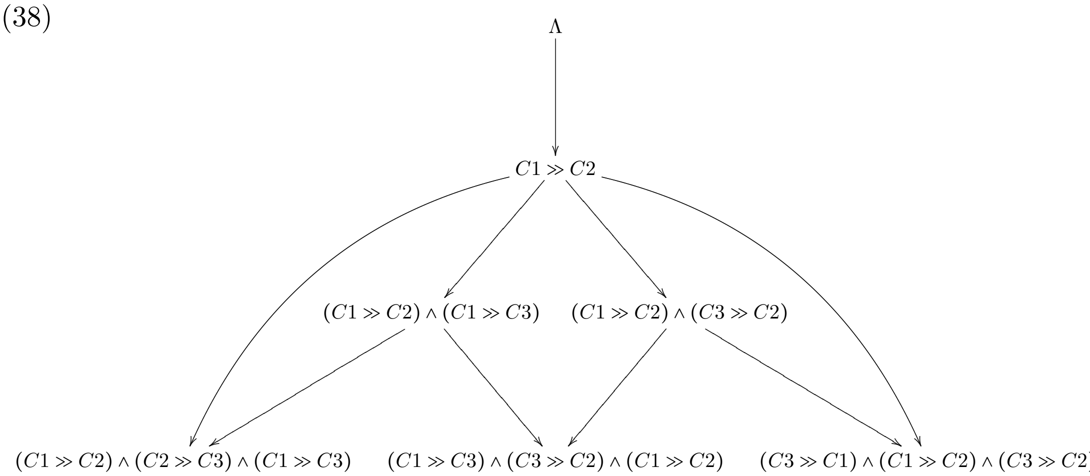

This picture shows only the small part of the structure, namely the part which is dominated by the partial ranking $(C1 \gg C2)$. Every other ranking in the picture is its refinement. Those refinements belong to one of the two levels: $(C1 \gg C2) \wedge (C1 \gg C3)$ and $(C1 \gg C2) \wedge (C3 \gg C2)$ are still underspecified, they do not decide, respectively, the ranking between $C2$ and $C3$, and between $C1$ and $C3$. The three rankings at the bottom of the picture are total rankings: they are fully specified in our small Con. If there is a direct arrow from ϕ to ψ in the picture, it means that we can get ψ from ϕ by adding one atomic ranking, and transitively closing the result. So each arrow connects two rankings which are at minimal distance from each other in the structure. The reason why a partial ranking can have immediate descendants at different levels is transitive closure. For instance, when we add to $C1 \gg C2$ an atomic ranking $C2 \gg C3$, we effectively add also the third atomic ranking $C1 \gg C3$ as well, by transitivity. Had partial rankings not been necessarily transitively closed, the structure would have looked differently. Note that there is no direct arrow from $(C1 \gg C2)$ to $(C1 \gg C3) \wedge (C3 \gg C2) \wedge (C1 \gg C2)$: there is no single atomic ranking that we can add to the former to immediately get the latter.

The entailment relation obeys reflexivity (every ranking entails itself), antisymmetry (if ranking ϕ entails ψ and they are not equal, then ψ does not entail ϕ) and transitivity (if ϕ entails ψ, and ψ entails ξ, then ϕ also entails ξ), and thus imposes a partial order on the rankings. The structure $\langle \Phi, \Lambda, \vDash \rangle$ (the set Φ of partial rankings, with a special element Λ, plus a relation of entailment between rankings) is a poset (partially ordered set) with a bound on one side.We can thus say that more underspecified rankings are greater than their refinements, treating $\vDash$ as $\geq$. The minimal ranking Λ is the maximal element, or 1.

We can call the refinements of a ranking its *daughters* or *descendants*, bringing in the tree terminology, as our poset structure can be viewed as a tree with the empty ranking as the root. Then we have a natural name for the inverse of the refinement relation: a ranking a is a *parent* for a ranking b iff b is a refinement of a.

Recall that an M-maximal ranking is such a ranking that it is true in M, but none of its immediate predecessors is. All descendants of an M-maximal ranking in the structure $\langle \Phi, \Lambda, \vDash \rangle$ have to be true in M, so the maximal ranking defines a “triangle” of rankings compatible with M, and is the greatest element in this “triangle”.

If $\phi \vDash \psi$, then all atomic rankings of ϕ are also in ψ. But it is not the only relation of interest possible between rankings. Take some arbitrary ϕ and ψ. First, they may have a common set of atomic rankings — a (ranking-)intersection. Second, when all the atomic rankings from both are combined together, they define a bigger ranking which contains both ϕ and ψ — a (ranking-)union.

Why should we be interested in those? There are two reasons. First and foremost, we will heavily use the operations on individual rankings we will now define in our analysis of sets of OT rankings. In that sense, we now simply build the foundation which will allow us to do much less trivial work further down the road. And the second reason is less grounded in the concerns relevant for the current paper, but instead can be appreciated right away, rather than after we see how we put those operations to use in Section 5: the proposed operations have nice potential practical applications which we illustrate using hypothetical scenarios involving grammar learners.

Imagine we have a learner comparing two partial ranking hypotheses ϕ and ψ about what the grammar is, knowing that each hypothesis may be wrong about any atomic ranking it contains. Then the intersection of ϕ and ψ contains the atomic rankings present in both hypotheses. Other things being equal, in our hypothetical scenario those atomic rankings are less likely to be wrong than the atomic rankings present in only one of ϕ and ψ.

For union, imagine a learner which acquired two partial ranking grammar hypotheses ϕ and ψ from different sources (e.g., from two sets of data processed in parallel threads of computation), and now wants to merge them into another partial ranking which is able to account for all the data ϕ or ψ can account for in the most conservative manner possible. Such a learner will have to find the minimal partial ranking which is entailed by both ϕ and ψ — their union.

Thus it makes sense to define the natural operations of ranking-intersection $\cap_r$ and ranking-union $\cup_r$.

(39) Let ϕ and ψ be well-formed formulas of **OTR**.

$\phi \cap_r \psi$ is the smallest well-formed formula ξ s.t.
$\forall Ci, Cj : [(Ci \gg Cj) \in \phi \wedge (Ci \gg Cj) \in \psi] \rightarrow (Ci \gg Cj) \in \xi.$

The ranking-intersection $\cap_r$ behaves simply as set intersection for sets of atomic rankings contained in the argument partial rankings. If there are contradictory atomic rankings in ϕ and ψ under intersection, they are simply excluded from $\phi \cap_r \psi$ because the intersection only includes common atomic rankings of the two. Atomic rankings obtained by transitivity are also fine: if there were two atomic rankings forcing a third by transitivity, in order to get into $\phi \cap \psi$ they had to be present in both ϕ and ψ, so then if there were a contradiction, it would have been to be one in individual ϕ and ψ already. Thus the result of $\cap_r$ is always defined.

But ranking-union is *not* a simple set union on the sets of atomic rankings corresponding to ϕ and ψ under union. Well-formed partial rankings must not contain contradictory atomic rankings, and must be transitively closed. But if we simply combine all the atomic rankings in ϕ and ψ, we may, first, create a contradictory ranking, in case ϕ and ψ contradicted each other, and second, the resulting set of atomic rankings may be not transitively closed. So our definition of ranking-union $\cup_r$, first, leaves the operation undefined on contradictory arguments, and second, ensures the result is transitively closed:

(40) Let ϕ and ψ be well-formed formulas of **OTR**.

$\phi \cup_r \psi$ is defined iff $\neg\exists Ci, Cj : (Ci \gg Cj) \in \phi \wedge (Cj \gg Ci) \in \psi$.

If $\phi \cup_r \psi$ is defined, then it is the smallest well-formed formula ξ s.t.
$\forall Ci, Cj : [(Ci \gg Cj) \in \phi \vee (Ci \gg Cj) \in \psi] \rightarrow (Ci \gg Cj) \in \xi$.

Both union and intersection are obviously associative, commutative, distributive, and obey absorption, provided that all unions are defined.

We can now view the domain of partial rankings as the structure $\langle \Phi, \Lambda, \vDash, \cap_r, \cup_r \rangle$.

Since our ranking-union $\cup_r$ is not set union, the structure $\langle \Phi, \Lambda, \vDash, \cap_r, \cup_r \rangle$ is not an algebraic substructure of the power set algebra generated by the set of atomic rankings with set-theoretic operations on it. In the powerset algebra, elements are really bags with atomic rankings in them, and since the requirements to be non-contradictory and to be transitively closed are not imposed, there are more elements ("pseudo-rankings") in it than in Φ. Because of the conditions of consistency and transitive closure, OT carves a non-trivial part of the powerset algebra, creating the domain Φ.

The intersection for partial rankings is just a restriction of the powerset algebra intersection to Φ, and it makes $\langle \Phi, \Lambda, \vDash, \cap_r, \cup_r \rangle$ a semilattice. But it is not a full lattice, because $\cup_r$ is not always defined. (The minimal element of a full lattice would have been the maximally contradictory ranking which contains all atomic rankings whatsoever.)

It might be interesting to characterize this structure $\langle \Phi, \Lambda, \vDash, \cap_r, \cup_r \rangle$ independently, not using the specific OT notions, as if it is just an abstract algebraic structure about whose members we do not know a thing. This would provide an independent characterization of what kind of structure OT actually creates in the single-ranking realm. I leave that for some other occasion.

$\cap_r$ and $\cup_r$ are straightforwardly connected to entailment:

(41)

$$\begin{aligned} \phi \cap_r \psi &\models \phi \\ \phi &\models \phi \cup_r \psi \end{aligned}$$

Thus $\phi \cap_r \psi$ is an ancestor for both ϕ and ψ, and $\phi \cup_r \psi$ (if it exists) is a descendant for both of them. Moreover, the partial ranking $\phi \cap_r \psi$ is the most specified common ancestor, and $\phi \cup_r \psi$ is the least specified common descendant of ϕ and ψ.

Entailment among partial rankings is definable in terms of $\cap_r$: $\phi \models \psi$ iff $\phi \cap_r \psi = \phi$.

We finish this section with a remark on how the relation of entailment $\models$ is interpreted in the domains of representative tableaux $\mathcal{M}^-$ and representative sets of tableaux Σ^-.[19]

For Σ^-, the notion of entailment imported from the domain of partial rankings Φ says that $\sigma_\phi \models \sigma_\psi$ when $\sigma_\phi \subseteq \sigma_\psi$: if $\phi \models \psi$, then ψ is compatible with all the tableaux ϕ is compatible with, and then possibly with some more, which exactly describes the situation when $\sigma_\phi \subseteq \sigma_\psi$.

For the domain of representative tableaux $\mathcal{M}^-$, entailment between equivalence classes M_ϕ and M_ψ is most easily characterized using the representatives of those classes which are entailment closures of their normal forms M_ϕ^{En} and M_ψ^{En} (see 32 above for the definition). $\phi \models \psi$ exactly when $M_\phi^{En} \subseteq M_\psi^{En}$. This is, of course, not surprising, given the relation between σ_ϕ and M_ϕ^{En}, 33.

4.6 Summary of Section 4

- Partial rankings over a specified Con may be viewed as a formulas in the logical language $\mathbf{OTR_{Con}}$.
- We can give a 2-valued semantics for $\mathbf{OTR_{Con}}$ which renders a formula ϕ representing a partial ranking true in a tableau M exactly when the partial ranking is compatible with M.
- If there are no restrictions on the space of logically possible rows imposed by Con and Gen, a ranking ϕ entails another ranking ψ exactly when ψ is a refinement of ϕ.
- A partial ranking is compatible with a tableau exactly when all its total refinements are compatible with it. Thus partial rankings may be viewed as notational shortcuts for sets of total rankings of a certain form. More precisely still, partial rankings, their sets of total refinements, and representations in between these two extremes are essentially different ways of representing the same OT object. In most cases, we use the partial ranking representative of this object, simply because it is most compact, and thus easy to work with.

[19] We do not discuss the interpretations of $\cap_r$ and $\cup_r$ in those domains. Those operations are very easy to understand in the domain of partial rankings, but their interpretations in the dual domains are quite complex, and do not necessarily allow one to better grasp the notion.

- There is a one-to-one correspondence between partial rankings and tableaux which only have rows with at most one W and one L. (More accurately, equivalence classes of tableaux whose normal form only contains rows with at most one W and one L.) Partial rankings and their representative tableaux contain exactly the same amount of information, and thus are interchangeable.

- There is also a one-to-one correspondence between partial rankings and sets of tableaux compatible with them. The set σ_ϕ of tableaux compatible with a partial ranking ϕ can be characterized in terms of the entailment closure of the representative tableau of ϕ, written M_ϕ^{En}. The normal form of any tableau in σ_ϕ is a subset of M_ϕ^{En}, all normal form subsets of M_ϕ^{En} are in σ_ϕ, and σ_ϕ is closed under transformations preserving OT-equivalency.

- There are natural operations of ranking-intersection $\cap_r$ and ranking-union $\cup_r$ of partial rankings. $\phi \cap_r \psi$ is the most specified common ancestor (=common entailer) of ϕ and ψ, or, in other words, the collection of all shared atomic rankings in ϕ and ψ. $\phi \cup_r \psi$ is the least specified common descendant (=common entailee) of ϕ and ψ, or the least collection of atomic rankings which contains all atomic rankings from ϕ and ψ and is transitively closed. $\cap_r$ is always defined, while $\cup_r$ is only defined for non-contradictory pairs of rankings. We can view the domain of partial rankings Φ not just as a set, but as a richer algebraic structure $\langle \Phi, \Lambda, \vDash, \cap_r, \cup_r \rangle$.

5 Sets of partial rankings

In Section 4, we introduced and studied partial rankings. Those are important because they are, in a certain sense explained above, equivalent to particular sets of total rankings. Namely, individual partial rankings are equivalents of sets which contain all and only total rankings true in a non-disjunctive tableau.

But some tableaux are disjunctive, and we have no means for working with sets of all rankings true in those. This section develops the necessary apparatus. We extend the notion of truth/OT-compatibility to sets of (partial) rankings. Obviously there are very many sets of rankings, but the notion of truth collapses them into equivalence classes relevant for OT. We find representatives for those equivalence classes, and study their structure. It turns out that these representatives exactly correspond to the sets of total rankings true in some tableau, the same way individual partial rankings corresponded to sets of total rankings true in some non-disjunctive tableau. We call sets of rankings corresponding to a tableau in this way proper sets.

We develop the methods which allow us to check if a given set of rankings is proper. We solve both the general form of the OT Ranking problem, and its inverse. Namely, given an arbitrary tableau, we can compute the faithful grammar hypothesis given that tableau as the data: the set of rankings which records all the rankings compatible with the data;

and conversely, given a proper set of rankings, we can compute the maximal set of data it is compatible with, and determine if a new set of data is consistent with such a grammar hypothesis.

The structure of this section is as follows. Section 5.1 sets the stage by defining the crucial basic notions. Section 5.2 solves the general form of the OT Ranking problem and provides the characterization of proper sets of rankings (in other words, of proper OT grammar hypotheses.) Section 5.3 develops an analysis of the relations between proper and all other sets of rankings, which shows that each proper set of rankings is the minimal proper representative S^{Min} of its equivalence class, and provides a method for finding an equivalent proper set from an arbitrary set of rankings. Finally, Section 5.4 discusses the algebraic structure of the domain of proper sets of rankings, which is shown to be generated by $\Cup$-closure of a small number of atoms, very simple sets of rankings.

5.1 OT-compatibility, entailment and equivalence for sets of partial rankings

Some tableaux, however, have more than just one maximal partial ranking — in fact, any tableau with a normal form containing multiple-W rows does. In particular, the following simple tableau has two maximal rankings, $C1 \gg C3$ and $C2 \gg C3$:

(42)

$C1$	$C2$	$C3$
W	W	L

Of course, it is trivially possible to construct the set of all total rankings true in a disjunctive tableau. Such a set would constitute a faithful grammar hypothesis in the same sense in which a partial ranking is one for a non-disjunctive tableau: such a set records the same information about constraint rankings as the tableau. But handling such big sets is hard. To make working with them easy, we will use the strategy developed for non-disjunctive tableaux. As a single partial ranking is a compact representative for sets of total rankings corresponding to non-disjunctive tableaux, a set of partial rankings will be a compact representative for sets of total rankings compatible with disjunctive tableaux.

A single partial ranking ϕ is true in a tableau M whenever every ranking from the set of all its total refinements M. So when ϕ is M-maximal, it is the faithful grammar hypothesis given M. If we consider the set of all total rankings compatible with a disjunctive tableau N, it cannot be equivalent to a single partial rankings. But we can extend the definition of OT-compatibility/truth for sets of rankings as well. After all, our earlier definition of compatibility for single partial rankings in 10 is already such a definition for certain sets, if we take into account the fact in 25. So the plan is this: we first define the notion of OT-compatibility for sets of (partial or total) rankings, which automatically creates equivalence classes of sets of rankings; then we find convenient representatives for such equivalence classes: compact and easy to work with.

We will write down sets of (total or partial) rankings in the usual set notation: $\{\phi, \psi\}$, where ϕ, ψ are rankings, and use uppercase letters such as S, T, ... as variables over such sets.

A set S of rankings is true in a row r (=compatible with r) iff all of the rankings in the set are true there. Similarly, S is true in a tableau M iff all rankings in S are true in M.

(43) A set of rankings $S = \{\phi_1, ..., \phi_n\}$ is true in a row r iff for all $\phi_i \in S$, ϕ_i is true in r.

(44) $S = \{\phi_1, ..., \phi_n\}$ is true in a tableau M iff for all $\phi_i \in S$, ϕ_i is true in M.

Entailment and equivalence for sets are standard:

(45) For sets of rankings S, T, $S \vDash T$ iff T is true in every tableau S is true in.

(46) For sets of rankings S, T, S and T are equivalent iff $S \vDash T$ and $T \vDash S$.

There are two distinguished sets of rankings: $\varnothing$ and $\{\Lambda\}$. $\varnothing$ is the empty set of rankings, and it is not compatible with any row whatsoever, including even L-less rows: there is simply no ranking in $\varnothing$ to satisfy the definition of truth above. $\{\Lambda\}$, on the other hand, does contain one ranking, though it is the minimal one. Thus $\{\Lambda\}$ is true in L-less rows. We call all other sets of rankings but those two *non-trivial*.

For singleton sets of rankings, the notions of truth and entailment coincide with the corresponding notions for single partial rankings. Thus we identify singleton sets of rankings with individual rankings (e.g., we do not distinguish between ϕ and $\{\phi\}$).

To figure out the general structure of equivalence classes of rankings, we will have to do a lot of work, but one thing we can note immediately: by the definition of truth for sets of rankings, if $\phi \vDash \psi$, then the set $\{..., \phi, \psi, ...\}$ is always equivalent to $\{..., \phi, ...\}$. In other words, it is always safe to omit the rankings which are refinements of other rankings in the set: it never changes OT-compatibility. It will be convenient for us to always bring sets of rankings we are working with to the standard form which does not include refinements. In what follows, we will always assume we work with sets which are in such a standard form.

The requirement for all rankings in the set to be true in M is very strong, and thus creates a very coarse-grained notion of equivalency. In particular, many different sets of total rankings turn out to be equivalent to each other. Consider the following example:

(47) $S := \{C1 \gg C2 \gg C3,\ C3 \gg C2 \gg C1\}$
$T := S \cup \{C1 \gg C3 \gg C2\}$

Suppose towards a contradiction there is a row r where S is true, but T is not. This means that in r, the two rankings from S are true, but the ranking $C1 \gg C3 \gg C2$ from T is not. This can only be if $C1 \gg C3 \gg C2$ lacks some atomic ranking crucial in accounting for r. What could this atomic ranking be? It cannot be $C3 \gg C1$ or $C2 \gg C1$, for the first ranking in S lacks them as well. It also cannot be $C2 \gg C3$, for the second ranking in

S lacks it. But all other 3 atomic rankings possible for a set of 3 constraints are already in $C1 \gg C3 \gg C2$. So there can be no such r in which S is true, but T is not. In fact, if we add to T yet another total ranking $C3 \gg C1 \gg C2$, by the same reasoning it will not change the OT-compatibility. But that four-ranking set will be the largest set of total rankings in its equivalence class. In fact, in each equivalence class there always exists a unique largest set of total rankings:

(48) **Largest total representative lemma**
For a non-trivial equivalence class of sets of rankings $\mathcal{C}$, there always exists a unique **largest total representative** S^{Tot} such that 1) each $\phi \in S^{Tot}$ is total, and 2) for every T in equivalence class $\mathcal{C}$ consisting only of total rankings, $T \subseteq S^{Tot}$.

Proof of 48. Existence of at least one member of the class which only contains total rankings follows from 25. Suppose there is a non-trivial class $\mathcal{C}$ of sets of rankings, pick some member U of it which contains a non-total ranking. By 25, that partial ranking is true iff all of its total refinements are true, so if we replace it with the set of such refinements, the resulting U' will be in the same equivalence class. By induction on the non-total rankings in U, we can build a member of $\mathcal{C}$ containing only total rankings.

For uniqueness, suppose some S and T are two largest distinct members of class $\mathcal{C}$. As they are distinct, their union $S \cup T$ has larger cardinality. As every ranking in $S \cup T$ was either in S or in T, $S \cup T$ is in $\mathcal{C}$ as well. So S and T were not largest sets of total rankings in $\mathcal{C}$, contrary to assumption. In general, S^{Tot} is the union of all sets of total rankings in $\mathcal{C}$. ⊣

Largest total ranking sets play an important role: if an equivalence class of sets of rankings is a faithful grammar hypothesis for some tableau (by the end of this section we will learn that in fact any class is), then its S^{Tot} representative contains all total rankings compatible with the data, providing the answer to the general form of the Ranking problem of OT.

On the other hand, sets of total rankings within a class that are *not* largest in it are deficient: there definitely cannot be a tableau M such that they contain all rankings compatible with M. This is an important result: not all sets of total rankings are born equal. Only some of them have direct OT significance.

Our notion of truth collapses such deficient sets into the same equivalence class with the corresponding S^{Tot}. But it is also possible to study the structure of the domain of sets of total rankings without that, using the notion of "weak compatibility" for sets of rankings which only requires one of the members of the set to be true in a tableau in order for the set to be declared true there. It is easy to see that under this notion, the sets S and U from 47 are not equivalent, and in fact no two distinct sets of total rankings are. An earlier version of this paper used this route, and it is quite possible to achieve all the results that way.

On the weak compatibility approach, the hard part is the unintuitive notion of truth. On the current approach, it is the coarse-grained equivalence classes of sets of rankings which are hard to understand and work with. But the main steps are the same anyway, whichever of the two routes we take. For instance, what looks like finding the right equivalence classes on the weak compatibility approach is finding the non-deficient representatives within an equivalence class on the current approach.

S^{Tot} representatives are largest in their equivalence class: they only contain total, that is, maximally specified rankings, and moreover they contain a maximal number of them. But of course for practical purposes it is convenient to have smallest, rather than largest, representatives for equivalence classes of sets of rankings. We will now define those.

(49) For an equivalence class of sets of rankings $\mathcal{C}$, for any $S \in \mathcal{C}$ and any $\phi \in S$, all total refinements of ϕ are in the largest total representative S^{Tot} of $\mathcal{C}$.

Proof of 49. Suppose there are S, ϕ for which 49 does not hold. Fix some total ranking ψ which is a refinement of ϕ and is not in S^{Tot}. By definition, for any tableau where S is true, ϕ and therefore by 25 ψ have to be true. But then $S^{Tot} \cup \{\psi\}$ is in the same equivalence class $\mathcal{C}$, and S^{Tot} was not the largest representative, contrary to assumption. ⊣

Thus any set of rankings in an equivalence class only contains ancestors of the total rankings in the S^{Tot} representative. This suggests where to look for a natural smallest representative of a class: it should include rankings as underspecified as possible, and cover all total rankings in S^{Tot}. We define the following notion:

(50) The **minimal proper representative** S^{Min} for an equivalence class of sets of rankings $\mathcal{C}$ is defined as follows:

Take the largest total representative S^{Tot} of $\mathcal{C}$. Add to S^{Tot} every ranking for which all total refinements are in S^{Tot}. By 25, the resulting set will also be in $\mathcal{C}$: adding such rankings preserves OT-compatibility. Then subtract from that set all rankings with ancestors in the set. Again, this preserves OT-compatibility. Call the resulting set S^{Min}.

S^{Min} is clearly minimal in the sense that it only contains maximally underspecified, smallest partial rankings. But is it truly minimal, that is, can there be some set of rankings T in the same equivalence class $\mathcal{C}$ which is a proper subset of S^{Min}? The answer to this question is no, there can be such T-s, though we will only be able to show this in Section 5.3.2, after more work on the structure of equivalence classes of sets of rankings is done. Still, S^{Min} will turn out to be a useful representative as all its subsets also in $\mathcal{C}$ have smaller sets of total refinements than S^{Tot}. Thus we call S^{Min} the minimal *proper* representative.[20]

Our inability to answer right away the question asked in the previous paragraph underscores an important point: even though we defined equivalence classes of sets of rankings,

[20]Later on, we will call certain sets of rankings *proper* sets. This usage of the same word is not accidental: it will turn out by the end of our analysis of sets of rankings that minimal proper representatives are indeed proper sets in that other sense.

at the moment we do not have the means to handle them. For instance, given an arbitrary set of rankings, we cannot tell if it is the S^{Min} set (or S^{Tot} set) for any equivalence class. Similarly, given an arbitrary set of rankings, we do not know how to build from it a representative of the equivalence class it is in. But the gaps cannot be filled without connecting sets of rankings and tableaux.

Once we take a look at sets of rankings from the individual tableau perspective, even more questions will immediately arise, and those will guide our investigation. Earlier we have defined the notion of an M-maximal partial ranking for a tableau M: the ranking which is true in M, and no ancestor of which is. Now we define another notion, that of M-maximal set of rankings:

(51) A set of rankings S is maximal in a tableau M (is M-maximal) iff all rankings in S are M-maximal, all M-maximal rankings are in S, and no rankings which are not M-maximal are in S.

As each ranking is either maximal in a tableau M or not, for each tableau there is a unique M-maximal set of rankings. M-maximal sets are important because they record the full set of rankings compatible with M.

(52) Ranking ϕ is true in tableau M iff ϕ is a refinement of some ψ in the M-maximal set of rankings.

Indeed, suppose it were not so. Then either the falsifying ranking ϕ is M-maximal itself, or not. If it is, it is supposed to be in the M-maximal set. If it is not, then some of ϕ's ancestors is M-maximal, and then that ancestor is in the M-maximal set of rankings.

Thus the set of total refinements of the M-maximal set is the set of all total rankings true in M. The M-maximal set is the faithful grammar hypothesis given M.

Let's review the two sides of the picture which we have to connect. On the one hand, the notion of truth organizes sets of rankings into equivalence classes, and there are S^{Tot} and S^{Min} representatives for those classes. On the other hand, for each tableau M, there is the M-maximal set of rankings, and the set of its total refinements. Are M-maximal sets of rankings S^{Min} representatives for some equivalence classes of sets of rankings? If yes, then are all S^{Min} representatives the maximal sets for some tableau? It will turn out that the answer to both questions is yes, but we will have to do a lot of work before we can show that.

First, in Section 5.2 we will study the sets which are maximal for some tableau, and will discover a way to tell whether an arbitrary set belongs to that group. We will also define an effective correspondence between tableaux and the sets which are maximal for them (thus essentially solving the general form of the Ranking problem of OT). After that, in Section 5.3 we will study the connection between those sets maximal for some tableau and minimal proper representatives of equivalence classes of sets of OT rankings. We will discover that in fact those two groups of sets of rankings are the same group, and thus minimal proper representatives are all and only sets maximal for some tableau, and each set maximal for a tableau defines its equivalence class of sets of rankings.

5.2 Correspondence between tableaux and sets of rankings

5.2.1 From tableaux to sets of rankings maximal for them

We will now define a procedure for building from an arbitrary tableau M the M-maximal set of rankings, the faithful grammar hypothesis for that tableau. We will prove that the resulting set of rankings is indeed maximal for the tableau using heavy decomposition. We will first define the operation of **pairwise ranking-union** $\Cup$ on sets of rankings and prove that if we apply $\Cup$ to two sets which are M_1- and M_2-maximal, the resulting set of rankings is the maximal for the tableau combining the two, $M_1 \cup M_2$. Then we will find the maximal sets for individual rows, and using $\Cup$, will become able to build maximal sets for arbitrary tableaux.

The operation of pairwise ranking-union $\Cup$ on sets of rankings is defined as follows:

(53) For sets of rankings S, T, $S \Cup T := \{\xi \mid \exists \phi_i \in S : \exists \psi_j \in T : \xi = \phi_i \cup_r \psi_j\}$

Thus $S \Cup T$ contains every ranking resulting from taking the ranking-union of a member of S and a member of T. Recall that for partial rankings ϕ and ψ, $\phi \cup_r \psi$ is the smallest ranking which is true in every tableau where either ϕ or ψ is true. Generalization of $\Cup$ to an arbitrary number of arguments is straightforward. We give an example of how $\Cup$ works:

(54)
$$\{(C1 \gg C2 \gg C3),\ (C4 \gg C5)\} \ \Cup\ \{(C1 \gg C5 \gg C4),\ (C3 \gg C4)\} =$$
$$= \{(C1 \gg C2 \gg C3) \wedge (C1 \gg C5 \gg C4),\ \varnothing,$$
$$(C1 \gg C2 \gg C3 \gg C4),\ (C3 \gg C4 \gg C5)\}$$

Note that as $(C4 \gg C5)$ and $(C1 \gg C5 \gg C4)$ are contradictory, their $\cup_r$ is not defined (we mark that by using $\varnothing$ where their ranking-union must have been.) In principle, it could have been that one of the other three partial rankings were a refinement of some other ranking in the set, but as it happens, none of the three rankings in the set is entailed by another.

$\Cup$ is commutative and associative, since it piggy-backs on $\cup_r$ which is both. The maximum number of rankings in $S \Cup T$ is $|S|$ times $|T|$, but in many cases, the actual number will be smaller: for some pairs of ϕ_i and ψ_j, the ranking-union is not defined. Note also that it is often the case that some $\phi_i \cup_r \psi_j$ is a refinement of some other ranking in $S \Cup T$, which means it can be omitted without offending OT-compatibility.

(55) For any tableaux M, M_1, M_2 s.t. $M = M_1 \cup M_2$, let S_1 be such that the total refinements of S_1 are all and only total rankings true in M_1, and similarly for S_2 and M_2.[21]

Then $S_1 \uplus S_2$ is true in M, and moreover, its total refinements are all and only total rankings true in M.

Proof of 55. For any i, j, $\phi_i \uplus \psi_j$, if defined, is compatible with both M_1 and M_2, and therefore with M as a whole. Thus all rankings in $S_1 \uplus S_2$ are true in M, and thus the set as a whole is.

Furthermore, suppose towards a contradiction that some total ranking ξ is true in M, but is not a refinement for any ranking in $S_1 \uplus S_2$. As ξ is true in M, it is also true in its parts M_1 and M_2. By assumption, all total rankings true in either are refinements of some ϕ from S_1 and some ψ from S_2. But then such a total ranking is also a refinement of $\phi \cup_r \psi$ which is in $S_1 \uplus S_2$ if defined. It cannot be that the ranking-union is defined for any pair of rankings from S_1 and S_2 which are ancestors to ξ: that would mean that ξ itself contains contradictory atomic rankings. Thus there can be no such ξ. ⊣

The proof trivially generalizes to an arbitrary division of M into parts.

Note that if one of M_1, M_2 is contradictory, no ranking is true in it, and the corresponding set of rankings is $\varnothing$. A pairwise ranking-union of any set with $\varnothing$ produces $\varnothing$, in symbols $\varnothing \uplus S = \varnothing$.

On the other hand, if one of M_1, M_2 does not have a single L, its corresponding set is $\{\Lambda\}$, and $\Lambda \uplus S = S$ for any S.

Now that we have $\uplus$ and the important lemma about it in 55, if we provide the way to build sets of rankings maximal for minimal tableaux — for individual rows — we will have the tools for building maximal sets for an arbitrary tableau.

(56) For an arbitrary row r, define set of rankings S_r as the biggest set of rankings (without duplicate rankings) where for any $\phi \in S_r$ two conditions hold:

(1) $\forall Cj \in L(r) : \exists Ci_{\phi,j} \in W(r) : (Ci_{\phi,j} \gg Cj) \in \phi \;\wedge\; \forall Ci : ((Ci \gg Cj) \in \phi) \rightarrow i = i_{\phi,j}$

(2) $\forall Ck \notin L(r) : \neg \exists Cm : (Cm \gg Ck) \in \phi$

Claim: S_r is r-maximal.

The first condition makes every L to be covered by some W, and prohibits covering it with any other constraint. The second condition ensures non-L constraints are not covered.

Proof of 56.

The proof is essentially a spelling out of what the two conditions on ϕ-s in S_r say.

Take an arbitrary row r. If there are no W-s in it, no ϕ satisfies the conditions, thus $S_r = \varnothing$. In a W-less row, no ranking is true, so indeed $\varnothing$ is the set of rankings maximal for it.

[21] In particular, S_1 can be the M_1-maximal set, and similarly for S_2 and M_2.

Suppose there is at least one W in r. Then it is easy to find a ranking ϕ satisfying the two conditions: fix some $Ck \in W(r)$, and build ϕ as $\bigwedge\{Ck \gg Cl \mid Cl \in L(r)\}$. Clearly such a ϕ cannot be contradictory, and thus S_r will always be non-trivial if there is at least one W in r.

We first show that every ranking ϕ satisfying the two conditions is r-maximal, and secondly, that every r-maximal ranking satisfies the two conditions and thus is in S_r.

The first condition guarantees that any $\phi \in S_r$ has for each L an atomic ranking covering it with a W, so ϕ must be true at r. Suppose some ϕ satisfying the two conditions is not r-maximal. Then there should be an atomic ranking which we can subtract from ϕ without making it false at r. By the second condition, only constraints from $L(r)$ can be dominated in ϕ. Take an arbitrary $Cj \in L(r)$. By the first condition, there is only one atomic ranking where Cj is dominated. If we subtract that ranking, there will be no way to account for the L in Cj. As Cj was arbitrary, no atomic ranking can be subtracted from ϕ without making it false at r, so ϕ has to be r-maximal.

Now suppose there exists some $\psi \notin S_r$ which is r-maximal. We show that then it satisfies both conditions, and thus must have been in S_r. Clearly an r-maximal ψ has to satisfy the first part of the first condition: $\forall Cj \in L(r) : \exists Ci_{\psi,j} \in W(r) : (Ci_{\psi,j} \gg Cj) \in \psi$, for otherwise ψ would be false in r. Suppose ψ does not satisfy the second part of the first condition. Then $\exists Cj \in L(r) : \exists Ck, Cl \in W(r) : k \neq l \wedge (Ck \gg Cj) \in \psi \wedge (Cl \gg Cj) \in \psi$. But then we can subtract either of $(Ck \gg Cj)$ and $(Cl \gg Cj)$ without making ψ false at r, so ψ is not maximal. Thus ψ has to validate the first condition. Suppose ψ does not validate the second condition. Then, similarly, we find $(Cm \gg Ck) \in \psi$ where $Ck \notin L(r)$, and note that subtracting it will not affect ψ's truth at r, so again ψ is not maximal, contrary to the assumption. Thus every r-maximal ranking satisfies the two conditions, and has to be in S_r. $\dashv$

Using 55 and 56, we can easily define a function from tableaux to their maximal sets.

(57) For an arbitrary tableau M, to compute the M-maximal set of rankings, do the following:

1. For each row $r_i \in M$, compute the r_i-maximal set of rankings S_i, using 56.
2. Compute $S'_M \coloneqq \biguplus_{i:r_i \in M} S_i$.
3. If there are distinct $\phi, \psi \in S'_M$ s.t. $\phi \vDash \psi$, delete ψ from the set. After all such refinements are deleted, call the resulting set S_M.

S_M is the M-maximal set of rankings, by 55.

As the M-maximal set records in a condensed way all the rankings compatible with M, it essentially provides the solution to the general form of the OT Ranking problem. If viewed in isolation from other work which we have done so far, this solution could have be viewed as a crowning achievement for the methods developed, a practical justification for all the fuss about partial rankings and sets of those. In contrast to the only previous full solution to the Ranking problem of [Brasoveanu and Prince, 2005], who compute the relevant set of rankings as a side-product of a syntactic manipulation of the tableau by the Fusional Reduction algorithm, 57 gives us a precise semantic characterization of what the solution set is and how it is related to the tableau; moreover, the apparatus just developed allows us to understand the connection between maximal sets for the parts of the tableaux and the whole tableau, something which the syntactic computation of [Brasoveanu and Prince, 2005] did not do.

However, after the work we have done so far, it is evident that the solution as such turns out to be not that impressive. Yes, 57 gives us a function from tableaux to sets of all rankings true in them. But many important questions naturally arise, given the framework we developed so far. For instance, if we are given an arbitrary set of rankings, we cannot even tell if it is maximal for some tableau or not. This is the question we take on next.

5.2.2 Proper sets of rankings

It is useful to divide sets of rankings into those which are maximal for some tableaux, and those which cannot be maximal for any particular tableau at all:

(58) A set of rankings S is **proper** iff S is maximal for some tableau M.

The task of finding out which sets are proper can be approached by a sort of reverse engineering: if we study how maximal sets look like, we will become able to tell whether a given set is maximal or not for any tableau.

Without loss of generality, we restrict ourselves to normal form tableaux. We can divide an arbitrary tableau M into two parts: the *non-disjunctive part* with single-W rows, and the *disjunctive part* with multiple-W rows.

From the non-disjunctive part, we will build the *core ranking* of the corresponding set of rankings: a partial ranking which contains all and only unconditionally necessary atomic rankings. Consider a non-disjunctive row r with a W in Ci and an L in Cj. Unless a ranking contains the atomic ranking $Ci \gg Cj$, it cannot be true in any tableau containing r. The core ranking will contain all such atomic rankings.

(59) For a set of rankings S, the **core ranking** S^{Co} is defined as $\bigwedge\{(Ci \gg Cj) \mid \forall \phi_k \in S : (Ci \gg Cj) \in \phi_k\}$.

We can easily show the following fact, which, in particular, allows us to factor out the core ranking out of any set of rankings. This will simplify the set, and leave us with a set of rankings $\{\phi_1, ..., \phi_n\}$ which, if the set was proper, should be maximal for the disjunctive part of the tableau.

(60) **Core addition theorem**
For any proper set S and a ranking ϕ, $\{\phi\} \Cup S$ is a proper set.

Proof of 60.

$\{\phi\}$ has a corresponding tableau non-disjunctive tableau M_ϕ, and by assumption S also has a corresponding tableau for which it is maximal, M_S. By 55, $\{\phi\} \Cup S$ is maximal for $M_\phi \cup M_S$. ⊣

The disjunctive part is more complex. For each row r in it, the L in r must be covered with a W, but there are several possible ways to do that, as there are multiple W-s in the row. We call different atomic rankings putting one of r's W-s on top of r's L *alternative solutions* for the *problem* that L constitutes. If for each row in the disjunctive part, we pick one of the alternative solutions and combine them into a ranking, that ranking will be true in the disjunctive part. Moreover, if we exhaust all combinatorial combinations of the alternative solutions for different rows, there will be no other possible ways to build a ranking true in the tableau: if a ranking does not contain one of the alternative solutions for some row, it cannot be compatible with the row.

So any set maximal for a tableau must be decomposable into a core ranking and a full set of combinatorial combinations of the alternative solutions. When we see an arbitrary set of rankings S, it is always possible to extract from it the core ranking: for any set, there is always a ranking containing all and only atomic rankings present in any $\phi \in S$. So in order to find out whether S is proper, we need to check, after extracting the core ranking, whether the rest of S is the full set of combinatory combinations of alternative solutions for some tableau. Thus the focus of this section is to find out how such full sets look like.

Clearly not any set of rankings is maximal for a tableau with only multiple-W rows. For instance, the set $\{C1 \gg C2,\ C3 \gg C4\}$ cannot be maximal for any non-trivial tableau whatsoever (or even true in any non-trivial tableau.) But the general task of finding out which sets are is not a trivial matter. We start taking it on by decomposing the problem.

An arbitrary normal form tableau with only disjunctive rows can be divided into (possibly empty) parts indexed by the members of **Con**. For each Ci, there will be a part O_{Ci} (the mnemonic is "O(nly) Ci (part)") containing all and only rows with an L in Ci. (Since we assumed the tableau is in the normal form, there are only single-L rows in it.)

Consider some O_{Ci}. It has a number of rows $r_1, \ldots\ r_n$ s.t $L(r_1) = \ldots = L(r_n) = \{Ci\}$, and each row has two or more W-s. Each O_{Ci} trivially has its corresponding maximal set of rankings which we will call S_{Ci}, or the **set of Ci-alternatives**. We will call each of S_{Ci}'s member rankings a **Ci-alternative**: it is a minimal recipe for covering the L in each row of O_{Ci} with a W without redundant coverings.

The following are an example of a possible O_{C5} tableau, and the set of rankings S_{C5} maximal in that O_{C5}:

(61) O_{C5} :

$C1$	$C2$	$C3$	$C4$	$C5$
W	W	e	e	L
W	e	W	e	L
e	e	W	W	L

The corresponding S_{C5} maximal for O_{C5} in 61 above:

(62)
$$S_{C5}: \left\{\begin{array}{l}(C1 \gg C5) \wedge (C3 \gg C5)\\ (C1 \gg C5) \wedge (C4 \gg C5)\\ (C2 \gg C5) \wedge (C3 \gg C5)\end{array}\right\}$$

Suppose we have an arbitrary set of rankings S such that every member ranking in it only has atomic rankings with a single constraint dominated, and furthermore, does not contain refinements or duplicates. Some of those, like the one in 62 above, are maximal for some O_{Ci} tableau, but can we tell if a randomly picked set is? Yes, we can. Perhaps surprisingly, it turns out that any set satisfying those simple conditions is maximal for some O_{Ci}, so distinct sets of that class are in one-one correspondence with O_{Ci} tableaux. We thus define the class of such sets of rankings in 63, and then prove that each of them is maximal for some O_{Ci}, and vice versa, for each O_{Ci}, there is its correspondent S_{Ci} set maximal for it.

(63) A set of rankings S is **set of Ci-alternatives** iff:[22]

(1) $$\exists Ci \forall \phi \in S : [\exists (Cx \gg Cy) \in \phi] \rightarrow [Cy = Ci]$$

(2) $$\forall \phi \in S : \neg \exists \psi \in S : (\phi \vDash \psi) \wedge (\phi \neq \psi)$$

The first condition only allows the rankings in a Ci-alternative set to have a single constraint as dominated, namely Ci. The second conditions simply rules out the presence of non-essential partial rankings within the set.

(64) For S_{Ci} a Ci-alternative set, there exists $O_{Ci} = \{r_1, ..., r_n\}$ s.t. $S_{r_1} \uplus ... \uplus S_{r_n} = S_{Ci}$. (That is, S_{Ci} is O_{Ci}-maximal.) Moreover, (for finite **Con**) there is an effective function computing O_{Ci} from an arbitrary S_{Ci}.

Proof of 64. We prove 64 for the finite case by actually providing the effective function which computes O_{Ci} from an arbitrary S_{Ci}. The same construction and proof work for the infinite case, but of course the function is not effective then.

First, we make sure that every combination of constraints dominating Ci in different ϕ_k in S_{Ci} has some useful job in the O_{Ci} we are constructing. For that, we form all combinatorially possible sets of constraints T such that for each ϕ_k, there is one dominator constraint from ϕ_k in T. Then we build for each such set T a row r with $W(r) = T$. If we combine those rows together, each ϕ_k will have to be maximal in it: for each atomic ranking within ϕ_k, there will be at least one row where that ranking is the only one in ϕ_k which can account for the row.

It is instrumental to give an example. Let **Con** $:= \{C1, ..., C6\}$. Consider the set of three $C6$-alternatives $S := \{(C1 \gg C6) \wedge (C2 \gg C6),\ (C3 \gg C6) \wedge (C4 \gg C6),\ (C1 \gg C6) \wedge (C4 \gg C6) \wedge (C5 \gg C6)\}$. Below we build the full combinatorial set of T-s for this S. (Recall that each T-set contains one dominator constraint from each of the 3 rankings in S; as some constraints are dominating in more than one ranking, some T sets will be equal, but we spell out all $2 \times 2 \times 3 = 12$ combinations.)

[22] Note that by this definition a singleton set can also be declared a set of Ci-alternatives, even though we are discussing only multiple-alternative sets at the moment. There is no harm in this generality, for a set with only one alternative can be viewed as a special simple case as well. In fact, later in Section 5.4 we will specifically use the general notion of S_{Ci} sets, both non-singleton and singleton.

$$T_1 = \{C1, C3, C1\} = \{C1, C3\}$$
$$T_2 = \{C1, C3, C4\}$$
$$T_3 = \{C1, C3, C5\}$$
$$T_4 = \{C1, C4, C1\} = \{C1, C4\}$$
$$T_5 = \{C1, C4, C4\} = \{C1, C4\}$$
$$T_6 = \{C1, C4, C5\}$$
$$T_7 = \{C2, C3, C1\} = \{C1, C2, C3\}$$
$$T_8 = \{C2, C3, C4\}$$
$$T_9 = \{C2, C3, C5\}$$
$$T_{10} = \{C2, C4, C1\} = \{C1, C2, C4\}$$
$$T_{11} = \{C2, C4, C4\} = \{C2, C4\}$$
$$T_{12} = \{C2, C4, C5\}$$

For each T_l, we build a row r s.t. $L(r) := Ci$, $W(r) := T_l$. Obviously we can omit the rows build from those T_i-s which are supersets of some T_j, as well as full duplicates — they will be superfluous when we combine the built rows into a single tableau. For instance, the row for T_1 will entail the rows for T_2, T_3, T_7; the rows for T_4 and T_5 are equal, so one of them is not needed, and entail the rows for T_6 and T_{10}. T_{11} entails T_8, T_{10} (also entailed by T_4/T_5), and T_{12}. The row for T_9 is independent and does not entail any other row we are supposed to build from T-s. This leaves us with T_1, T_4, T_9 and T_{11}. So here is what those four sets are, and the tableau built from them:

(65) Non-superfluous T_l sets:
$T_1 = \{C1, C3\}$ (T_2, T_3 and T_7 are supersets of this one)
$T_4 = \{C1, C4\}$ (T_5 is a duplicate, T_6 and T_{10} are supersets)
$T_9 = \{C2, C3, C5\}$ (no duplicates or supersets)
$T_{11} = \{C2, C4\}$ (T_{12} is a superset)

(66)

$C1$	$C2$	$C3$	$C4$	$C5$	$C6$
W	e	W	e	e	L
W	e	e	W	e	L
e	W	e	W	e	L
e	W	W	e	W	L

It is easy to check that each ranking in the original set $S := \{(C1 \gg C6) \wedge (C2 \gg C6), (C3 \gg C6) \wedge (C4 \gg C6), (C1 \gg C6) \wedge (C4 \gg C6) \wedge (C5 \gg C6)\}$ is maximal in the tableau in 66. Why did it happen? We give the argument for the general case now.

Take an arbitrary T_l set. For each $\phi_k \in S_{Ci}$, there is some $Cj \in T_l$ s.t. $Cj \gg Ci \in \phi_k$, by construction of T_l. So if we build a row with $W(r) := T_l$, $L(r) := Ci$, every $\phi_k \in S_{Ci}$ will be true in this row.

When we combine all rows built from T_l sets into a tableau O, clearly all rankings in S_{Ci} will be true in O. What we need to show is that they not only are true, but also are maximal there, and that all O-maximal rankings are in S_{Ci}. If that is so, then there is no ranking true in O which is not a refinement of one of the rankings in S_{Ci}.

Suppose towards a contradiction that there exists a ranking ψ true in O such that ψ is not a refinement of any $\phi_k \in S_{Ci}$. Then for each $\phi_k \in S_{Ci}$ there is at least one atomic ranking $Cj \gg Ci$ which is in ϕ_k, but not in ψ. We fix one such $Cj \gg Ci$ for each ϕ_k, and gather their dominator constraints into a set. This set is in fact one of the T_l sets we built when constructing O. So there must be a row r_l in O built from this set of constaints. But because of how we picked Cj constraints, they do not dominate Ci in ψ, so ψ cannot account for that row r_l, and hence for the tableau as a whole, contrary to assumption. (Note that if r_l we happened to select is entailed by another row in O, it does not make ψ's life better: the entailer row will have a subset of r_l's W-s, and ψ cannot use any of those W-s as it does not feature the relevant Cj constraints as dominators.) Thus by construction of O, all ψ-s compatible with O have to be entailed by one of $\phi_k \in S_{Ci}$.

From this fact, all ancestors of any $\phi_k \in S_{Ci}$ have to be false in O, which makes each ϕ_k O-maximal. Similarly, any other ranking not in S_{Ci} cannot be O-maximal because if it is not a refinement of some ϕ, then it is false at O. Thus S_{Ci} is indeed maximal for the constructed O. ⊣

The proof of 64 is quite simple, but it does not make the fact proven less astonishing. As an exercise helping to fully appreciate the significance of 64, the reader is advised to pick a couple of arbitrary sets of the S_{Ci} form, and first try to come up with tableaux for which those sets would be maximal by hand, and then apply the function in 64 if naïve attempts to find such a tableau fail.

(67) *Corollary to 64.*
There is a one-one correspondence between normal form tableaux O_{Ci} with L-s in the same constraint in all rows and sets of rankings which are sets of Ci-alternatives (as defined in 64), such that the corresponding S_{Ci} is O_{Ci}-maximal.

Proof of 67. 64 provides us with a function from sets S_{Ci} to the described domain of tableaux with L-s in a single constraint Ci, such that the value of the function is the tableau for which S_{Ci} is maximal. We need to show that the function is both into and onto that domain.

First, we prove injectivity. Suppose two sets S^1_{Ci} and S^2_{Ci} are mapped to the same tableau O_{Ci}. By 64, each ranking ψ true in O_{Ci} is either in S^1_{Ci} or a refinement of a ranking in it; and same for S^2_{Ci}. It follows that $S^1_{Ci} = S^2_{Ci}$.

Surjectivity is not much harder to prove. Take an arbitrary O_{Ci} and compute the pairwise ranking-union of the maximal sets for its rows $S_{O_{Ci}}$, then trim the refinements. By 55, the resulting set S is O_{Ci}-maximal. As any row in O_{Ci} only has an L in Ci, rankings in S only contain atomic rankings of the form $X \gg Ci$, and as a maximal set for some tableau, it does not contain superfluous refinements. Thus S is a proper set of Ci-alternatives, and is taken to an O tableau for which it is maximal by the function in 67. Since S is maximal for both O_{Ci} and O, $O_{Ci} = O$. As O_{Ci} was arbitrary, the range of the function in 67 is the whole domain of O_{Ci} tableaux. ⊣

Thus any set with the same constraint dominated in all of the rankings — any set of Ci-alternatives for some Ci 63 — is maximal for some tableau. By 55 and 67, this means that every pairwise ranking-union of a number of such S_{Ci} sets is the maximal set of rankings for some tableau, and conversely, any disjunctive tableau can be divided into O_{Ci} parts, from which we can compute the S_{Ci} sets of rankings, and take their $\Cup$ to get the set of rankings maximal for the whole disjunctive tableau. So we are one step closer to understanding which form the set of rankings maximal for an arbitrary (genuinely) disjunctive tableau have: they always are of the form $S_{Ci} \Cup S_{Cj} \Cup ... \Cup S_{Cn}$.

But this is not yet the full answer, for how can we tell if an arbitrary set of rankings is the result of applying $\Cup$ to a number of S_{Ci} sets? This is the question we answer next. But before we give the general solution, we consider, in the reverse engineering fashion, a specific example first, in order to build the intuition behind the answer.

Consider a set $S := \{(C1 \gg C3 \gg C2),\ (C4 \gg C2 \wedge C4 \gg C3)\}$. Is it proper or not? If it is, then it must be a pairwise ranking-union for some S_{Ci} sets.

Let's try to recover the specific S_{Ci} sets from S. We will assume it is possible, and try to figure out what those sets were. (In fact, the set is not proper, so we will in a sense fail, but instead we will find the next best thing: how to *extend* S to a proper set.)

First, there should be two S_{Ci} sets, for there are two constraints which appear as dominated in S: $C2$ and $C3$.

Consider S_{C2}. It must have within it atomic rankings $C3 \gg C2$ and $C4 \gg C2$, which must be in different members of S_{C2}: if they were a part of the same ranking in S_{C2}, they could not have been split in S. Furthermore, S_{C2} might also feature $C1 \gg C2$ in the same ranking with $C3 \gg C2$, but it might also not: we do not know if that atomic ranking is present in the first ranking of S because it follows by transitivity from $C1 \gg C3$ and $C3 \gg C2$, or if it were present in S_{C2}. Thus we have two possible candidates for S_{C2}, and their corresponding O_{C2} tableaux:

(68)

$C1$	$C2$	$C3$	$C4$
e	L	W	W

$\Leftrightarrow\ \{(C3 \gg C2),\ (C4 \gg C2)\} = S_{C2}^{Small}$

(69)

$C1$	$C2$	$C3$	$C4$
e	L	W	W
W	L	e	W

$\Leftrightarrow\ \{((C3 \gg C2) \wedge (C1 \gg C2)),\ (C4 \gg C2)\} = S_{C2}^{Large}$

By similar reasoning, S_{C3} has to be $\{(C1 \gg C3),\ (C4 \gg C3)\}$, with the following corresponding O_{C3} (this time, there is no uncertainty, as there are no atomic rankings with $C3$ dominated following by transitivity in S):

(70)

$C1$	$C2$	$C3$	$C4$
W	e	L	W

$\Leftrightarrow\ \{(C1 \gg C3),\ (C4 \gg C3)\}$

Any smaller S_{C2} and S_{C3} sets would not be able to produce our two rankings in S: we only included the bare minimum. But if we actually take their $\Cup$, the resulting set is greater than our initial S. First let's compute step by step $S_{C2}^{Small} \Cup S_{C3}$:

(71) $S_{C2}^{Small} \Cup S_{C3} =$
$= \{(C3 \gg C2),\ (C4 \gg C2)\} \ \Cup \ \{(C1 \gg C3),\ (C4 \gg C3)\} =$
$= \{(C3 \gg C2) \cup_r (C1 \gg C3),\ \ (C3 \gg C2) \cup_r (C4 \gg C3),$
$(C4 \gg C2) \cup_r (C1 \gg C3),\ \ (C4 \gg C2) \cup_r (C4 \gg C3)\} =$
$= \{C1 \gg C3 \gg C2,\ \ C4 \gg C3 \gg C2,\ \ (C4 \gg C2) \wedge (C1 \gg C3),$
$(C4 \gg C2) \wedge (C4 \gg C3)\}$

Recall that if a set of rankings contains a ranking ϕ and its refinement ψ, then ψ clearly does not do any work: if the set is true, that means ϕ is true, but then the refinement ψ automatically has to be true. In 71 above, the ranking $C4 \gg C3 \gg C2$ is a refinement of $(C4 \gg C2) \wedge (C4 \gg C3)$, so it can be deleted:

(72) $S_{C2}^{Small} \Cup S_{C3} =$
$= \{C1 \gg C3 \gg C2,\ \ (C4 \gg C2) \wedge (C1 \gg C3),\ \ (C4 \gg C2) \wedge (C4 \gg C3)\}$

If we compute $S_{C2}^{Large} \Cup S_{C3}$, we in fact get the same resulting set: $((C3 \gg C2) \wedge (C1 \gg C2)) \cup_r (C1 \gg C3)$ is the same ranking $C1 \gg C3 \gg C2$; and though $((C3 \gg C2) \wedge (C1 \gg C2)) \cup_r (C4 \gg C3)$ produces the ranking $(C4 \gg C3 \gg C2) \wedge (C1 \gg C2)$, which is different from the one we got in $S_{C2}^{Small} \Cup S_{C3}$, that ranking is still a refinement of $(C4 \gg C2) \wedge (C4 \gg C3)$ and thus is superfluous anyway, and can be deleted, resulting in the same 72.

Now, there is a tableau for which that set in 72 is maximal: we can obtain it if we combine the O_{C2} and O_{C3} tableaux from above.[23]

(73)

$C1$	$C2$	$C3$	$C4$
e	L	W	W
W	e	L	W

$\Leftrightarrow$

[23]The tableau in 73 looks like the union of 68 and 70. But we had another possibility for O_{C2}, the tableau in 69. What happened to it? In fact, if we take the union of 69 and 70, we will immediately see that the second row of 69 is superfluous in the result:

(1)

$C1$	$C2$	$C3$	$C4$
e	L	W	W
W	L	e	W
W	e	L	W

If we account for the first row with $C3 \gg C2$, and the third row with $C1 \gg C3$, we get by transitivity that $C1 \gg C2$, which takes care of the second row. If we use $C3 \gg C2$ for the first row and $C4 \gg C3$ for the third, again by transitivity we get $C4 \gg C2$, which accounts for the second row. Finally, if we use $C4 \gg C2$ for the first row, it immediately takes care of the second row. Thus the second row is entailed by the first and third rows taken together.

We will prove in 75 the general result that it never matters whether we include the atomic rankings followed by transitivity in the rankings of the original set when we recover S_{Ci}: their inclusion or omission do not change the end result.

$$\Leftrightarrow \quad S^{P+} = \{(C1 \gg C3 \gg C2),\ (C4 \gg C2 \wedge C1 \gg C3),\ (C4 \gg C2 \wedge C4 \gg C3)\}$$

Obviously $S \subset S^{P+}$. In fact, S^{P+} is the minimal possible extension of S to a proper set. S itself is not proper: whichever tableau may be accounted for by the rankings $(C1 \gg C3 \gg C2)$ and $(C4 \gg C2 \wedge C4 \gg C3)$ will be bound to be accounted for by $(C4 \gg C2 \wedge C1 \gg C3)$ as well.

When we tried to recover the S_{Ci} sets which could have generated S, we failed, but we only learned that when we actually gathered minimal such sets. At the same time, the result of our actions was valuable in its own right: it was a set of rankings including the original set, but unlike the original S, set S^{P+} was proper. On the other hand, if we could recover such sets S_{Ci} which produce the original S when we take their pairwise ranking-union, that would have meant that the original set was proper.

This is the strategy we will use for our attack on the problem of determining which sets are proper: the decision procedure for a set S will be to build the minimal proper extension of S, and then check if that extension is S itself or a bigger set. We have just done precisely that for a single example above, and now we proceed to define the general procedure. First, we define the notion of minimal proper extension itself:

(74) Take an arbitrary set $S = S^{Co} \Cup \{\phi_1, ..., \phi_n\}$ where $\forall j : S^{Co} \cap_r \phi_j = \Lambda$.[24]

Let $L(S)$ be the set of all constraints dominated in some ϕ_j. For each $Ci \in L(S)$, we define $S_{Ci} := \{\phi_1^{Ci}, ..., \phi_n^{Ci}\}$ where ϕ_j^{Ci} is a restriction of the corresponding ϕ_j including all and only atomic rankings with Ci as the dominated constraint. Formally, $\phi_j^{Ci} := \bigwedge_{(Ck \gg Cl) \in \phi_j} (Ck \gg Cl)$.[25]

The **minimal proper extension** S^{P+} of set S is as follows:

$$S^{P+} := \{S^{Co}\} \Cup (\Cup_{Ci \in \mathbf{Con}} \{S_{Ci}\})$$

We discuss the definition clause by clause. We start with an arbitrary set S, and disassemble it into the core ranking S^{Co} and the residue set of rankings $\{\phi_1, ..., \phi_n\}$. The question then is whether that residue set can be rewritten as a pairwise ranking-union of a number of S_{Ci} sets.

[24]Strictly speaking, it is not necessary to factor the core ranking out of S. An alternative equivalent definition would define the minimal proper extension of a proper set $\{\phi_1, ..., \phi_n\}$ as simply $\Cup_{Ci \in \mathbf{Con}} \{S_{Ci}\}$, which we will use that in section 5.4.

But the factoring out of the core ranking helps to stress the different roles which the non-disjunctive and the genuinely disjunctive parts of a tableau play: only for the latter we need the complex procedure of looking for the minimal proper extension.

[25]Note that ϕ_j^{Ci} may contain as meaningful atomic rankings which could have been derived by transitivity in ϕ_j, and thus might have been omitted from its written form.

For each constraint Ci which is dominated in some ϕ_i, we simply extract from each ϕ its part where that Ci is dominated, and combine them in a set: $S_{Ci} \coloneqq \{\phi_1^{Ci}, ..., \phi_n^{Ci}\}$.[26]

When we build the S_{Ci} sets, we essentially disassemble the original rankings into sets of alternative recipes for each dominated constraint. Thus we end up with the full set of solutions for "problems" (L-s) that the original set of ranking could have hope to handle. Finally, when we take $\Cup_{Ci \in \mathbf{Con}}\{S_{Ci}\}$, reassemble those recipes together, trying out all possible combinations — for all Ci, Cj, we try combining each ϕ_i^{Ci} with each ϕ_j^{Cj}. The resulting set will definitely include the rankings ϕ_i which were present in the original S, but it also may include more rankings beside that, recording the combination which were not there. In this sense the result S^{P+} is literally an extension of S: it fills up the gaps, the ways to combine different recipes which were not in S — using the parts that *were* given in S. The idea is that the newly included combinations are necessary other alternative ways to take care of every row the original rankings are true at.

In our example above, the new ranking $(C4 \gg C2) \wedge (C1 \gg C3)$ extending S to S^{P+} in 73 contains combines a solution for $C2$ coming from the second ranking of S with the solution for $C3$ coming from the first ranking of S. As we will prove, such a recombination leads to the emergence of rankings which are necessarily true in any tableau where the original rankings are true.

Thanks to our results in 55 and 64, we can immediately note that being a pairwise ranking-union of S_{Ci} sets (and a single core ranking), any minimal proper extension S^{P+} is indeed a proper itself: all S_{Ci} sets are themselves proper, corresponding to O_{Ci} tableaux, and applying $\Cup$ to proper sets produces a proper set.

OK, if a set is equal to its own minimal proper extension, it is clear that it is proper. But why is an arbitrary set *required* to be its own minimal proper extension to be proper? We will prove this formally shortly, but first we describe the general idea. Let's take some set S and assume, without any reason, that it is proper. How can somebody then convince us we are wrong? If S is proper, then it has a corresponding tableau. This tableau then can be divided into the core part and O_{Ci} parts. And for the O_{Ci} tableaux, there are S_{Ci} sets such that $\Cup_{Ci \in L(T)} S_{Ci} = \{\phi_1, ..., \phi_n\}$. So by existence of the corresponding tableau, S will be its own proper extension, granted that the S_{Ci} sets we recover from it are indeed the same as the S_{Ci} sets corresponding to the O_{Ci} parts.

It is not immediately obvious from the procedure in 74 that it should be so: by that definition, we recover S_{Ci} sets by gathering all the atomic rankings from different ϕ_i, including those that follow by transitivity from other atomic rankings. When we discussed the example above, we considered two variants for S_{C2}: one which included such atomic rankings and one which did not. In that case, the final result was the same. But in

[26]Some ϕ_i^{Ci} may well be equal to each other, in which case we automatically omit them, as we deal with sets, not multisets. But if some ϕ_i^{Ci} entails some ϕ_j^{Ci}, we do *not* delete the refinement from the set: even though such a deletion would have been preserving OT-compatibility, with regard to the gathered S_{Ci} sets we are interested not in their OT-compatibility, but in their ability to generate back the original $\{\phi_1, ..., \phi_n\}$ set.

the general case, how can we be sure we recover the right set of atomic rankings? What if sometimes it becomes crucial whether we include in it some atomic rankings that are transitively entailed? The following lemma shows that remarkably, we have nothing to worry about, there is no real choice here, as the result is always the same.

(75) S_{Ci} **recovery lemma**

Take an arbitrary set of rankings $\{\phi_1, ..., \phi_n\} = \Cup_{Ci \in \mathbf{Con}}\{S_{Ci}\}$.
Define for each ϕ_j, Ci the ranking ϕ_j^{Ci} which contains all independent atomic rankings of the form $Ck \gg Ci$ in ϕ_j, and in addition to them any number of $Ck \gg Ci$ entailed by transitivity in the same ϕ_j.

Claim: Regardless of the choice of each ϕ_j^{Ci}, for sets $T_{Ci} := \{\phi_1^{Ci}, ..., \phi_n^{Ci}\}$ it holds that $\Cup_{Ci \in \mathbf{Con}}\{S_{Ci}\} = \Cup_{Ci \in \mathbf{Con}}\{T_{Ci}\}$.

Proof of 75.

We prove the lemma in two steps: first, we show that taking $\Cup$ of the smallest sets, which do not contain any entailed atomic rankings, produces the original set. Then we show that adding to any alternative solution atomic rankings following by transitivity does not do any harm.

It is clear that the pairwise ranking-union of the extracted T_{Ci} sets, as long as we harvest all the non-entailed atomic rankings, will include all the rankings in $\{\phi_1, ..., \phi_n\}$. Indeed, any ranking ϕ_j was divided into ϕ_j^{Ci} parts, whatever the method. If we recombine those exact parts, we get the same ϕ_j, and the result does not depend on whether we included atomic rankings following by transitivity or not.

On the other hand, the S_{Ci} sets in $\Cup_{Ci \in \mathbf{Con}}\{S_{Ci}\}$ cannot be any smaller than the harvested sets which do not contain any entailed atomic rankings. Hence when T_{Ci} does not require any atomic rankings required by transitivity in the ϕ_j rankings, $\Cup_{Ci \in \mathbf{Con}}\{S_{Ci}\} = \Cup_{Ci \in \mathbf{Con}}\{T_{Ci}\}$.

Now suppose that the addition of some atomic ranking $Ck \gg Cl$ to some Cl-alternative ϕ_j^{Cl} in T_{Cl} makes the set $\Cup_{Ci \in \mathbf{Con}}\{T_{Ci}\}$ to include more rankings than we had in $\Cup_{Ci \in \mathbf{Con}}\{S_{Ci}\}$. Fix such a new ranking ψ. Compare it with ϕ_j from which the atomic ranking $Ck \gg Cl$. We know by induction hypothesis that $\phi_j = \cup_r \phi_j^{Ci}$ where ϕ_j^{Ci} only contain non-entailed atomic rankings. But as $Ck \gg Cl$ was entailed in ϕ_j, ψ must be equal to ϕ_j, contrary to assumption.[27] By induction on entailed atomic rankings, we derive the claim.

Thus it never matters how many atomic rankings followed by transitivity we include when we construct S_{Ci} sets according to a plan like the on in 74. In practical applications, it is probably easier to only harvest the independent atomic rankings; in our definition, we opt for the simplicity and define ϕ_j^{Ci} as including all atomic rankings with Ci dominated in the original ϕ_j. But there is no substantial difference.

Now that we have discussed how we build minimal proper extensions, we can return to the question of how one can convince us that a set S, which we think is proper, actually is not. *If*, as we pretended, S is proper, by S_{Ci} recovery lemma 75, we should have recovered

[27] In most cases, an identical ϕ_j^{Ci} belongs to more than one ranking in the original set. It is easy to see that the same argument applies to all of those.

exactly those S_{Ci} sets which will produce the distributive residue of S. Thus if, on the contrary, after building S^{P+} we see that $S \neq S^{P+}$, that simply means S was not proper in the first place.

(76) **Proper set theorem**
Set of rankings S is proper iff $S = S^{P+}$.

Proof of 76.
($\Rightarrow$) Suppose S is proper. There must be a tableau M_S witnessing that. M_S can be divided into the non-disjunctive part $M_{S^{Co}}$, the representative tableau for the core ranking S^{Co} of S, and the disjunctive part $M_S \setminus M_{S^{Co}}$. The disjunctive part may be further subdivided into O_{Ci} tableaux. By 67, there are S_{Ci} sets maximal for those O_{Ci} tableaux. By 55, $S = \{S^{Co}\} \Cup (\Cup_{Ci \in \mathbf{Con}} \{S_{Ci}\})$.

To prove that such S is equal to S^{P+}, we only need to show that the procedure defined in 74 extracts exactly the right S_{Ci} sets out of S. This is established by the S_{Ci} recovery lemma 75.

($\Leftarrow$) In the other direction, suppose $S = S^{P+}$. Given that $S^{P+} = \{S^{Co}\} \Cup (\Cup_{Ci \in \mathbf{Con}} \{S_{Ci}\})$, we can build little tableaux $M_{S^{Co}}$ and a number of O_{Ci} for which the sets whose pairwise ranking-union we took are maximal. By 55, the set of rankings $S^{P+} = S$ is maximal for the union of those tableaux. $\dashv$

We now have a procedure (in the finite case, an effective procedure) which decides whether an arbitrary set S is proper or not: it suffices to build the minimal proper extension, and then check whether it is equal to the original set or not.

The goal of this section is achieved, for now we can tell if a set is maximal for any tableau or not. We can then define a function from maximal sets into their corresponding tableaux, and this, in turn, will allow us to start working on the relationships between proper sets and S^{Min} sets.

5.2.3 From proper sets of rankings to their corresponding tableaux

Above, we have provided a procedure for computing from an arbitrary tableau M the set of rankings S_M maximal for it. But without being able to distinguish sets of rankings which are proper (that is, the range of the function from tableaux to corresponding sets of rankings), we could not define well the inverse function. Now, after we have established 76, we can:

(77) For an arbitrary proper set of rankings S, to compute its representative tableau M_S for which S is maximal, do the following:

1. Factor out the core ranking S^{Co} 59 out of S.
2. Recover S_{Ci} sets from the residue using the procedure in 74.
3. Build the corresponding tableaux $M_{S^{Co}}$ for the core ranking S^{Co} (as in 28), and O_{Ci} for the S_{Ci} sets (by 64).
4. Let $M_S := M_{S_{Co}} \cup (\bigcup_{Ci \in \mathbf{Con}} O_{Ci})$

M_S is the tableau for which S is maximal, by 55.

As 57 and 77 are inverses, and moreover both are effective for finite constraint sets (assuming without loss of generality that tableaux are in normal form), we have a one-one correspondence between arbitrary tableaux and proper sets of rankings, and moreover, the correspondence is effective.

We can now turn to the question of which position proper sets of rankings occupy within the equivalence classes of rankings. After some work, we will be able to show that proper sets of rankings are exactly minimal proper representatives S^{Min}.

5.3 Proper sets of rankings within their equivalence classes

On the one hand, we have learned how to distinguish proper sets of rankings — those maximal for some tableau — from other sets, and how to put them in effective correspondence with the tableaux they are maximal for. On the other hand, we know that many different sets of rankings collapse into a single equivalence class, and we identified minimal proper representatives of equivalence classes S^{Min}. But what is the connection between those two perspectives on the domain of sets of OT rankings? In particular, which position proper sets occupy within their equivalence class? We will see by the end of this section that in fact each proper set is the minimal proper representative S^{Min} for its equivalence class, and that each S^{Min} representative is a proper set.

The plan is as follows. First we extend the correspondence between $\mathcal{M}^-$ and Σ^- we established in Section 4.4 to the domain of all tableaux $\mathcal{M}$. This will allow us to immediately note that the σ set of tableaux given by the correspondence for some tableau M is precisely the sets of tableaux in which the M-maximal set of rankings is true. In other words, that σ set characterizes the equivalence class $\mathcal{C}$ of sets of rankings to which the M-maximal set of rankings belongs. Next, we show that for that equivalence class $\mathcal{C}$, the M-maximal set of rankings is indeed equal to the minimal proper representative S^{Min}. And finally, we will show that there are no other equivalence classes of sets of rankings than the ones characterized by σ sets, which means that every S^{Min} representative is maximal for some tableau, and thus that each equivalence class of sets of rankings precisely corresponds to an equivalence class of sets of tableaux.

5.3.1 Duality between sets of tableaux and representative tableaux

In Section 4.4, we defined a correspondence between representative non-disjunctive tableaux M_ϕ and sets σ_ϕ of all tableaux compatible with the corresponding ranking ϕ. We extend the correspondence between tableaux and sets of tableaux to include disjunctive tableaux as well.

The entailment closure of a tableau defined in 78 is essentially the same notion that we used in 32, and the difference between 33 and 79 is just that now we define a correspondence between arbitrary normal form tableaux and σ sets, not just between non-disjunctive tableaux representing a single partial ranking. For a tableau M, we write the corresponding set of tableaux as σ^M.

(78) M^{En} is such a tableau that $r \in M^{En}$ iff M entails r.

(79) For a tableau M, $\sigma^M := \{N \mid N^{No} \subseteq M^{En}\}$, where N^{No} is the normal form of N.

We can also easily define the domain of sets of tableaux in the range of our correspondence by relaxing the requirement for the representative tableau to be non-disjunctive. Later, we will relate that new domain Σ to the domain of equivalence classes of sets of rankings. (In fact, we will find out that Σ contains exactly sets of tableaux which characterize equivalence classes of sets of rankings.)

(80) Set of tableaux σ is in Σ iff there is a tableau M^{En} such that:

- for every row r entailed by M^{En}, $r \in M^{En}$, and
- for every normal form tableau $N \in \sigma$, every row $q \in N$ is also in M^{En}, and
- for every normal form tableau $P \subseteq M^{En}$, $P \in \sigma$, and
- σ is closed under OT-equivalency-preserving operations.

What is the relation between tableaux M and their corresponding sets of tableaux σ^M?

(81) For a tableau M and its corresponding set of tableaux σ^M, a ranking is true in M iff it is true in all $N \in \sigma^M$.

Proof of 81. Suppose some ϕ is true in M, but is false in some tableau $N \in \sigma^M$. But the normal form of N is a subset of M^{En}, and thus if ϕ is false in N, it is also false in M^{En}, and thus in M, contrary to assumption.

M itself is in σ^M, so if a ranking is false in M, it cannot be true in all tableaux in σ^M. $\dashv$

There is another important fact about σ^M: there is no tableau not in σ^M where all the rankings true in M were true. To prove that fact, we will need to employ our findings about the correspondence between tableaux and their maximal sets of rankings from Section 5.2.

(82) **σ set completeness theorem**
For a tableau M, there is no tableau where all rankings true in M are true, and which is not in σ^M.

Proof of 82. Suppose towards a contradiction there is a tableau N such that all rankings true in M are also true in N, and yet N is not in σ^M. Then there must be a row $r \in N$ which is not in M^{En}, and the M-maximal set of rankings S_M must be true in that r as it is true in N as a whole by assumption.

Consider that fixed r. Without loss of generality, we can assume r has a single L. Any ranking in S_M must put one of the constraints in $W(r)$ on top of the constraint Ci in $L(r)$.

If there is a single W in r, then the corresponding atomic ranking putting that W on top of the L must be in the core ranking of S_M. But then M contains precisely the same row r, contrary to assumption.

If there are multiple W-s in r, any ranking in S_M must put one of those W-s on top of that L in Ci. If it is the same W in any ranking in S_M, that means r is entailed by a row in the representative tableau of the core ranking of S_M. If it is different W-s, then the S_{Ci} set recovered from S_M minus its core ranking must be such that each alternative in it contains an atomic ranking putting one of the W-s in $W(r)$ on top of Ci. If so, then the corresponding O_{Ci} contains a row entailing r.

Thus there can be no row not in M^{En} where all the rankings true in M were true. Therefore any N where all rankings true in M are true is bound to be in σ^M. ⊣

The theorem 82 is a very important result: it follows from it that we can use σ sets to characterize certain equivalence classes of sets of rankings, as we do in the next section.

5.3.2 Equivalence classes of sets of rankings characterized by σ sets

For any proper set of rankings S and its corresponding tableau M_S, the set of tableaux σ^{M_S} by 82 contains all and only tableaux M_S is true in. We can then use the set of tableaux σ^{M_S} to define the equivalence class of S in the following way: a set of rankings T is equivalent with S iff σ^{M_S} is the set of all tableaux T is compatible with.

Now we can finally build a link between proper sets and minimal proper representatives S^{Min} we defined in 50:

(83) A proper set S is the minimal proper representative S^{Min} of its equivalence class $\mathcal{C}$.

Proof of 83. For an arbitrary proper set S, there exists a set of tableaux σ^{M_S} which contains all and only tableaux S is true in. Any tableau in the equivalence class $\mathcal{C}$ of S is true in exactly the same set of tableaux.

We compare the set S' of all total refinements of S and the S^{Tot} representative of $\mathcal{C}$ (see 48). It cannot be that $(S' \setminus S^{Tot}) \neq \varnothing$, for then S^{Tot} were not the largest set in $\mathcal{C}$ containing only total rankings. It cannot also be that $(S^{Tot} \setminus S') \neq \varnothing$: if that were the case, there would be a ranking $\phi \in S^{Tot}$ true in all tableaux in σ^{M_S} yet not in S'. In particular, such ϕ would have to be true in M_S. But then it would be a total refinement of one of the rankings in S, and thus included in S', contrary to assumption. Thus $S' = S^{Tot}$.

Thus S is the set with its set of total refinements being S^{Tot}. We compare S to S^{Min} (see 50). S^{Min} contains all the rankings for which all total refinements are in S^{Tot}. Those rankings are precisely the rankings maximal for the tableau M_S, by definition of tableau-maximal rankings. Furthermore, S^{Min} does not contain any other rankings than M_S-maximal ones. Hence $S = S^{Min}$. ⊣

Thus all proper sets are the minimal proper representatives of their equivalence classes. But are there any S^{Min} representatives which are *not* proper sets? 83 leaves open the possibility that there are more equivalence classes of sets of rankings than there are proper sets.

But in fact it can be shown that each S^{Min} representative is a proper set. We show that using σ sets: we prove that there can be no equivalence class of sets of rankings which is true in a set of tableaux not forming a σ set.

(84) A minimal proper representative S^{Min} is always a proper set of rankings.

Proof of 84. Suppose towards a contradiction there is some S^{Min} set which is not proper. Fix the set of tableaux in which S^{Min} is true as τ. It is clear that τ cannot be the σ^M set of tableaux for any M, for otherwise S^{Min} would have been proper.

We gather all the rows present in at least one normal form tableau in τ into a single tableau M_τ (which, in the general case, would not be a normal form tableau). Clearly S^{Min} is true in M_τ, and M_τ itself is in τ.

The set S^{Tot} of all total refinements of S^{Min} is by construction of S^{Min} the set of all total rankings true in M_τ. Therefore S^{Min} is M_τ-maximal. But then τ is after all the set σ^{M_τ}, contrary to assumption.

Therefore for any equivalence class of sets of rankings, the set of tableaux where the sets of rankings from the class are true is always a σ set, that is, a set generated by an entailment-closed single tableau. There are thus as many equivalence classes of sets of rankings as there are equivalence classes of tableaux, and each S^{Min} set is maximal for the tableau generating the σ set corresponding to S^{Min}. $\dashv$

(85) *Corollary to 84.*
There is a one-one correspondence between equivalence classes of sets of rankings and equivalence classes of tableaux.

Since 84 holds, the word *proper* in the term "minimal proper representative" has a double meaning: the set S^{Min} itself is a proper set, in addition to the fact we noted earlier but have not shown that it is not a truly minimal representative of its equivalence class.

But to show that S^{Min} is not necessarily "minimal" in the sense that there might be other T in the same equivalence class $\mathcal{C}$ such that $T \subset S^{Min}$, we need to further investigate the internal structure of equivalence classes of sets of rankings, and we do that in the next section.

5.3.3 The internal structure of an equivalence class of sets of rankings

Up to this point, we still have only a trivial characterization of what it means for two sets of rankings to be equivalent: they are when they are true in exactly the same tableaux. Despite the fact that we have already learned a great deal about equivalence classes of sets of rankings — for instance, we now know they are in one-one correspondence with equivalence classes of tableaux — we still cannot tell, given an arbitrary pair of sets of rankings, whether they are in the same equivalence class or not without comparing directly the sets of tableaux they are true in. And given that such sets are infinite, this brute force method of checking does not look too attractive. But the following fact will help us find the means to test for equivalency effectively.

(86) For each equivalence class $\mathcal{C}$ of sets of rankings, there exists a unique minimal proper extension S^{P+} set of rankings within that class.

Proof of 86. Each equivalence class of sets of rankings has its minimal proper representative S^{Min}. As S^{Min} is proper by 84, it is equal to its own minimal proper extension by 76. Hence there always exists one S^{P+} per class.

For uniqueness, suppose there are two distinct minimal proper extensions in some $\mathcal{C}$. By 76, both must be proper. Then by 83, both are equal to the S^{Min} representative of class $\mathcal{C}$. But then they are equal, contrary to assumption. ⊣

The last fact we have to prove before we can test for equivalency easily is the following one:

(87) For any set of rankings S and its minimal proper extension S^{P+}, S and S^{P+} are equivalent sets of rankings.

Proof of 87. Suppose towards contradiction that were not so. Fix some S and S^{P+} which are not equivalent.

By construction of S^{P+}, any $\phi \in S$ must be a refinement of a ranking from S^{P+}. Therefore the set of total refinements of S is a subset of the set of total refinements of S^{P+}. From that, it follows that $S^{P+} \models S$.

Thus if $S \not\models S^{P+}$, it must be witnessed by some tableau M where S is true, but S^{P+} is not. We fix a particular row r of M where S, but not S^{P+}, is true. W.l.o.g., we assume r has a single L in Ci.

r cannot be true in S by virtue of being accounted for by S^{Co}, for S^{P+} has the same core.

Suppose r is accounted for by the non-core part of S. Then in each $\phi \in S$, there is an atomic ranking accounting for the L in Ci in r. Each ranking $\psi \in S^{P+}$ contains the part ϕ_j^{Ci} from some $\phi_j \in S$. Therefore each ψ has to account for r just as well as any ϕ does.

Thus it is impossible to find any row at which S is true, but S^{P+} is not. Hence the two are compatible with the same set of tableaux, and thus S and S^{P+} are equivalent. ⊣

These results allow us to derive the following simple corollary:

(88) Sets of rankings S_1 and S_2 are equivalent iff $S_1^{P+} = S_2^{P+}$.

As we can (for finite Con) effectively compute the minimal proper extension of any set of rankings, and as those are also finite and thus easy to compare, we have an effective test for equivalency of two arbitrary sets of rankings.

Moreover, from 87 we learn exactly which sets are in a given equivalence class: those whose proper minimal extension is its S^{Min} representative. This, in turn, gives us a way to generate the members of the equivalence class smaller than S^{Min}: we can subtract rankings from S^{Min} without offending OT-compatibility as long as the result still has the same minimal proper extension.

In order to preserve what the minimal proper extension is, we need to make sure that the full range of Ci-alternatives present is still in the set. It is not necessary to have the same alternative present in S^{Min} more than once, so as long as each has at least one instance, the set remains equivalent to S^{Min}. This can be illustrated with the following example:

(89)

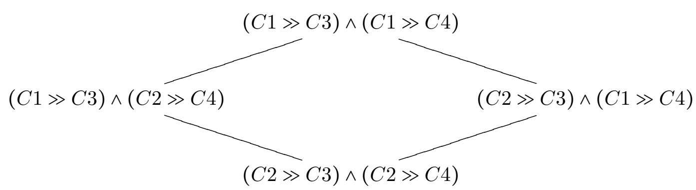

(90)

$C1$	$C2$	$C3$	$C4$
W	W	L	e
W	W	e	L

The set of rankings consisting of the four partial rankings in 89 is proper. It corresponds to the tableau in 90, and can be decomposed into S_{C3} and S_{C4}, both of which contain 2 alternatives:

(91) $S_{C3} = \{(C1 \gg C3),\ (C2 \gg C3)\}$
$S_{C4} = \{(C1 \gg C4),\ (C2 \gg C4)\}$

In the picture 89, each pair of rankings connected by a line can be transformed into each other by exchanging a particular Ci-alternative in one of them for the other possible Ci-alternative.

It is easy to see that the pairs of rankings opposite from each other — those pairs not connected by a line, and maximally different from each other — are each sufficient to recover the sets in 91. All four rankings are included in the S^{Min} representative of the equivalence class in question, but as long as the opposite rankings in 89 are kept, we can subtract one or two of the rankings maximal in the corresponding tableau 90 preserving OT-equivalency of the set.

Thus the set S^{Min}, which records all possible combinations of Ci-alternatives, and not just a subset of those sufficient for the recovery of S_{Ci} sets, is not truly a minimal representative. Yet the same example of the set in 89 shows that though there may be smaller sets $T \subset S^{Min}$ still in the same equivalence class, those T will not in the general case be unique. Therefore even though S^{Min} is bulkier than needed, it can serve as the unique representative of the class, but smaller equivalent T may not. Still, in practice it is of course better to have smaller representations. But even more efficient than any equivalent $T \subset S^{Min}$ for recording the equivalence class would be direct writing down of the core ranking S^{Co} and non-empty sets S_{Ci} into which S^{Min} of the equivalence class may be decomposed.

5.3.4 Behavior of non-proper sets of rankings under $\Cup$

From 55, we know that when we take the pairwise ranking-union of two proper sets of rankings, the result is always a proper set corresponding to the tableau which is the union of the corresponding tableaux for the arguments of the operation. $\Cup$ is thus a very useful operation on sets of rankings with an important OT meaning.

But what if we apply the same operation of $\Cup$ to non-proper sets? We give an example demonstrating that it need not be that $S \Cup T = S^{Min} \Cup T^{Min}$. In other words, if we simply take the pairwise ranking-union of two non-proper sets, the result may be in a totally different equivalence class of sets of rankings than $S^{Min} \Cup T^{Min}$. That means the members of an equivalence class do not behave uniformly with respect to $\Cup$, so basically we better keep clear of using non-proper sets if we want to stay out of trouble.

Consider two sets of rankings $\{\phi_1, ..., \phi_n\}$ and $\{\phi_1\}$, where all ϕ_i rankings are total. Suppose the set $\{\phi_2, ..., \phi_n\}$ is in the same equivalence class with $\{\phi_1, ..., \phi_n\}$. We observe that $\{\phi_1, ..., \phi_n\} \Cup \{\phi_1\} = \{\phi_1\}$, but $\{\phi_2, ..., \phi_n\} \Cup \{\phi_1\} = \{\varnothing\}$: indeed, for any two total rankings ϕ and ψ, the ranking $\phi \cup_r \psi$ is either $\phi = \psi$ if the two are equal, or not defined if $\phi \neq \psi$. But $\{\phi_1\}$ and $\{\varnothing\}$ are not equivalent. They belong to different equivalence classes. This shows that in the general case, applying $\Cup$ to non-proper sets may not have the same OT meaning as in the case of proper sets. So it is not that pairwise ranking-union as such which has an important OT interpretation, but rather $\Cup$ in the domain restricted to proper sets.

But what is this domain like? We will finish our discussion of sets of rankings with an analysis of that domain in the next section.

5.4 The domain of proper sets of rankings as an algebraic structure

The final set of questions we will address concerns the relations between different proper sets (and sets of their total refinements). Such questions include, for instance, whether there can be S^{Tot} sets such that one is a subset of the other, or whether the intersection of two S^{Tot} set can be, or must be, a proper set itself, and what kind of a proper set it will be then. In essence, all such questions are covered by one big question: what does the domain of proper sets of rankings looks like if we view it as an algebraic structure? Which operations relate which sets in it, and how are its members related to each other?

It is natural to start the analysis with identifying what are minimal, atomic proper sets of rankings. What could those be, aside from the trivial set of rankings $\{\Lambda\}$? The building blocks which we used when decomposing proper sets of rankings were S^{Ci} sets. We usually dealt with non-singleton S_{Ci} sets, but in fact singleton sets satisfying the conditions in 63 can be viewed as a special case of S^{Ci}, too.

The correspondence between S_{Ci} and O_{Ci} equivalence classes of tableaux guarantees that each S_{Ci} set is genuinely distinct from all others, and being proper, it defines an equivalence class of sets of rankings as its minimal proper representative S^{Min}.

Now, by 55, all other proper sets can be viewed as the pairwise ranking-union of several atomic sets.

The actual result in 55 guarantees decomposition into S_{Ci} *and* the core ranking. It is easy to show that we can "import" each atomic ranking in S^{Co} into the respective S_{Ci} set, conjoining it with all rankings in that S_{Ci}. Thus we can indeed view an arbitrary set as the pairwise ranking-union of several S_{Ci} sets (if we include into S_{Ci} sets singleton sets containing just one Ci-alternative.)

Thus the domain of proper sets is essentially the $\Cup$-closure of the set of all possible S_{Ci} sets.

Given the correspondence between proper sets and tableaux, this fact has a correlate in the realm of tableaux: the set of all possible tableaux is generated by a set of rows closed under row merger (that is, any tableau is simply a collection of rows.) As tableau merger is essentially the interpretation of $\Cup$ for sets of rankings in the domain of tableaux (and vice versa, $\Cup$ is the interpretation of tableau merger in the domain of proper sets), this parallelism is in fact inevitable.

But the parallelism is not precise so far: S_{Ci} sets correspond to O_{Ci} tableaux, not to individual rows. This suggests that among S_{Ci} sets, not all are equally atomic: there are true atoms which correspond to individual rows, and then there are also more complex S_{Ci} sets which can be generated by $\Cup$ from those atoms. It is easy to see, using the correspondence with individual rows, why the following is the characterization of the truly atomic part of the S_{Ci} realm:

(92) An S_{Ci} set $\{\phi_1, ..., \phi_n\}$ is **atomic** iff each ϕ_j contains exactly one meaningful atomic ranking.

This, of course, is hardly surprising: it is only natural that the building blocks of all proper sets are very close to individual atomic rankings. What is non-trivial, and what constitutes a major achievement of the current paper, is the understanding of how exactly those elementary building blocks are combined into proper sets — namely, the understanding that it is the operation of pairwise ranking-union $\Cup$ that generates proper sets from the atoms.

We have noted above that in general, $\Cup$ does not have its convenient OT interpretation when applied to non-proper sets of rankings. However, on S^{Tot} representatives, $\Cup$ behaves well enough. In fact, the following fact is useful for visualization of what happens when we take $\Cup$ of two proper sets:

(93) The largest total representative of the class defined by $S^{Min} \Cup T^{Min}$ is $S^{Tot} \cap T^{Tot} = S^{Tot} \Cup T^{Tot}$, where $\cap$ has the standard interpretation of set intersection.

Proof of 93. Let $\phi \in S^{Min}$, $\psi \in T^{Min}$. Consider $\phi \cup_r \psi$. Every refinement of $\phi \cup_r \psi$ is a refinement of both ϕ and ψ. Thus any total refinement of $S^{Min} \Cup T^{Min}$ was in both sets S^{Tot} and T^{Tot}.

Conversely, if some ranking was in both S^{Tot} and T^{Tot}, it was a refinement of some $\phi \in S^{Min}$, $\psi \in T^{Min}$, and thus it will also be a refinement of $S^{Min} \Cup T^{Min}$. So indeed the set of total refinements of $S^{Min} \Cup T^{Min}$ is $S^{Tot} \cap T^{Tot}$.

Finally, that $S^{Tot} \cap T^{Tot} = S^{Tot} \Cup T^{Tot}$ follows from the fact that the ranking-union $\cup_r$ of two total rankings is either equal to both, if they were equal, or not defined. ⊣

So whenever we take $S \Cup T$, the resulting set of rankings has fewer (or at least not more) total refinements than either of the original sets. This, again, has a parallel in the realm of tableaux: whenever we merge two tableaux, we put additional restrictions on what a ranking must have in order to be true in the tableau. We never cancel previously enforced restrictions.

It is important to note that while for S^{Tot} sets, the operations of $\Cup$ and $\cap$ coincide, this is not so in the general case: for instance, if we simply take the usual set intersection of two S^{Min} sets, the result is likely to be rubbish, in the sense that there is no useful OT characterization of the result.

This underscores the fact that while $\Cup$ is not the only conceivable operation on sets of rankings, including proper sets, unless there is a special reason to study them (such as, for instance, their significance for OT), there is not much merit in a simple enumerating of possibilities. This is why up to this moment we did not even mention taking the usual set union or intersection of proper sets: those do not have a natural OT interpretation. But it does not mean that no other useful operations may be defined for proper sets. We conclude this section illustrating it on a real example of such an operation.

Consider the intersection of σ_S and σ_T. It contains all and only tableaux where both S and T are individually true, so it may be an interesting set. But is there an operation on S and T which would produce a set of rankings true in exactly the tableaux in $\sigma_S \cap \sigma_T$?

We can answer such a question in two steps. First, we need to make sure that the set of tableaux in question, $\sigma_S \cap \sigma_T$, is necessarily a proper σ set. Otherwise there is simply no set of rankings such that it is true in precisely the tableaux in $\sigma_S \cap \sigma_T$. And second, if we can in the general case define the set of rankings corresponding to $\sigma_S \cap \sigma_T$ in terms of the sets S and T, then we essentially have discovered the sought operation.

So first, is $\sigma_S \cap \sigma_T$ a proper σ set? Any tableau M which belongs to both σ_S and σ_T can only contain rows entailed by both M_S and M_T, that is, rows included into $M_S^{En} \cap M_T^{En}$. In particular, the maximal collection of such rows, the tableau $M_S^{En} \cap M_T^{En}$, belongs at the same time to both σ_S and σ_T. All other tableaux in $\sigma_S \cap \sigma_T$ are subsets of $M_S^{En} \cap M_T^{En}$. Thus if $M_S^{En} \cap M_T^{En}$ is itself closed under row entailment, then $\sigma_S \cap \sigma_T$ is a normal σ set.

It is easy to show that $M_S^{En} \cap M_T^{En}$ is indeed a proper entailment-closed M^{En} tableau. For suppose there is a row r entailed by $M_S^{En} \cap M_T^{En}$ but not belonging to it. As our tableau in question is a subset of both M_S^{En} and M_T^{En}, the row r must be also entailed by M_S^{En} and M_T^{En}. But then $r \in M_S^{En} \cap M_T^{En}$, contrary to assumption. So $M_S^{En} \cap M_T^{En}$ is closed under row entailment, and thus is a proper σ set.

Now we need to define the set of rankings maximal for $M_S^{En} \cap M_T^{En}$ in terms of S and T. First we give an illustrative example, and then discuss the general case. Consider proper sets $S := \{C1 \gg C3 \gg C4\}$, and $T := \{(C2 \gg C3) \wedge (C2 \gg C4)\}$. Their corresponding M^{En} tableaux are as follows (the trivial rows are omitted, and the normal form subtableau is given above the double line):

(94) $= M^{En}_{(C1 \gg C3) \wedge (C3 \gg C4)}$

$C1$	$C2$	$C3$	$C4$
W	e	L	e
e	e	W	L
W	W	L	e
W	e	L	W
W	W	L	W
W	e	W	L
e	W	W	L
W	W	W	L
W	e	e	L
W	W	e	L
W	e	W	L
W	W	W	L
W	e	L	L
W	W	L	L

(95) $= M^{En}_{(C2 \gg C3) \wedge (C2 \gg C4)}$

$C1$	$C2$	$C3$	$C4$
e	W	L	e
e	W	e	L
W	W	L	e
e	W	L	W
W	W	L	W
W	W	e	L
e	W	W	L
e	W	L	L
W	W	L	L

The intersection of the two tableaux $N^{En} := M^{En}_{(C1 \gg C3) \wedge (C3 \gg C4)} \cap M^{En}_{(C2 \gg C3) \wedge (C2 \gg C4)}$ is this:

(96)

$C1$	$C2$	$C3$	$C4$
W	W	L	e
W	W	L	W
W	W	e	L
W	W	W	L
W	W	L	L

$= M^{En}_{(C1 \gg C3) \wedge (C3 \gg C4)} \cap M^{En}_{(C2 \gg C3) \wedge (C2 \gg C4)}$

The normal form of the tableau in 96 is the tableau in 97, and its corresponding maximal set of rankings U is given in 98:

(97)

$C1$	$C2$	$C3$	$C4$
W	W	L	e
W	W	e	L

(98) $$U = \left\{ \begin{matrix} (C1 \gg C3) \wedge (C1 \gg C4) \\ (C1 \gg C3) \wedge (C2 \gg C4) \\ (C2 \gg C3) \wedge (C1 \gg C4) \\ (C2 \gg C3) \wedge (C2 \gg C4) \end{matrix} \right\}$$

What is the relation between this set U and the original sets S and T? The set U can be decomposed into a pairwise ranking-union of two Ci-alternative sets S_{C3} and S_{C4}:

(99) $$\left\{ \begin{matrix} (C1 \gg C3) \wedge (C1 \gg C4) \\ (C1 \gg C3) \wedge (C2 \gg C4) \\ (C2 \gg C3) \wedge (C1 \gg C4) \\ (C2 \gg C3) \wedge (C2 \gg C4) \end{matrix} \right\} =$$

$$= \{(C1 \gg C3),\ (C2 \gg C3)\} \uplus \{(C1 \gg C4),\ (C2 \gg C4)\}$$

Now the relation between U and the original S and T becomes clear: the S_{Ci} sets of U are the set unions of the S_{Ci} sets of S and T![28] For instance, the S_{C3} set of S contains only the alternative $C1 \gg C3$; the S_{C3} set of T also contains a single $C3$-alternative $C2 \gg C3$. The S_{C3} set of U contains both of those alternatives.

This is the key to the understanding of the general case, as we will now show. Let S and T be arbitrary proper sets, and σ_S and σ_T their correspondent sets of tableaux. How can a ranking ϕ true in every tableau in $\sigma_S \cap \sigma_T$ look like?

First, obviously any refinement of a ranking in either S or T is bound to be true in all tableaux in $\sigma_S \cap \sigma_T$: that set of tableaux is a subset of both σ_S and σ_T, and any refinement

[28] Note that in order for that to hold, it is necessary to *not* factor out the core ranking of S and T before decomposition into S_{Ci} sets.

of S or T is true in all tableaux in the corresponding σ set. But are there other options? How would a ϕ which is not a refinement of any ranking in S and T look like if it is true in $\sigma_S \cap \sigma_T$?

Take some row r with a single L in Ci that can be found in a tableau from $\sigma_S \cap \sigma_T$. The alternatives from either S and T can equally well account for that row. Of course, for tableaux simply in σ_S the Ci-alternatives from T may be not suitable, of course. But for the subset of σ_S where all rankings in T are true, that is, in $\sigma_S \cap \sigma_T$, the alternatives from T are just as good as the ones from S. Below we formally define the operation of **alternative-addition** $\Cap$ on sets of rankings and prove that it is interpreted in the dual domain of σ sets as set intersection.[29]

(100) Let $S := \{\phi_1, ..., \phi_m\}$ and $T := \{\phi_{m+1}, ..., \phi_n\}$ be proper sets of rankings.
Let $\phi_k^{Ci} := \bigwedge_{(Cj \gg Ci) \in \phi_k} (Cj \gg Ci)$.
Then the **alternative-addition** of S and T is defined as follows:

$$S \Cap T := \Cup_{Ci \in \mathbf{Con}} \{\phi_1^{Ci}, ..., \phi_m^{Ci}, \phi_{m+1}^{Ci}, ..., \phi_n^{Ci}\}$$

NB: Alternative-addition $\Cap$ is not defined for non-proper sets. But the result of the operation is not guaranteed to be a proper set.[30]

(101)
$$\sigma_{S \Cap T} = \sigma_S \cap \sigma_T$$

That is, for proper sets S and T, set $S \Cap T$ is the set of rankings true in precisely those tableaux where both S and T are true.

Proof of 101.

Take an arbitrary row r with a single L in Ci appearing in $\sigma_S \cap \sigma_T$, and an arbitrary ψ_j in $S \Cap T$. ψ_j necessarily contains a Ci-alternative coming from either S or T. As both S and T are true in r, ψ_j is also true there. As ψ_j was arbitrary, $S \Cap T$ is true in r. As r was arbitrary, $S \Cap T$ is true in every tableau in $\sigma_S \cap \sigma_T$.

In the other direction, take an arbitrary row r with a single L in Ci in which $S \Cap T$ is true. The only part relevant for any set's truth in r is the set's S_{Ci} set of alternatives. In particular, the S_{Ci}

[29] The mnemonics for the chosen symbol $\Cap$ are as follows: the base part $\cap$ serves as a reminder that the operation has the interpretation of set intersection for σ sets (and M^{En} tableaux); the + part concerns the fact that the S_{Ci} sets of the result of the operation are a “sum” (more accurately, a set union) of the S_{Ci} sets of the arguments.

Note that while $\Cup$ and $\Cap$ are not directly connected, the choice of opposite base symbols for those two is well justified: $\Cup$ only diminishes the number of total refinements of the resulting set as compared to the argument sets, while $\Cap$ only enlarges that number.

[30] Consider $\{C1 \gg C2\} \Cap \{C3 \gg C4\}$. It is the set $\{(C1 \gg C2), (C3 \gg C4)\}$, which can only be true in a trivial L-less tableau. Any such set of rankings belongs to the equivalence class whose S^{Min} representative is $\{\Lambda\}$.

set of $S \Cap T$ is true at r. Therefore any its subset, including the S_{Ci} sets of S and T, is also true in r. Thus any such r appears in $\sigma_S \cap \sigma_T$. ⊣

It is easy to see that the atomic S_{Ci} sets can be generated by taking a $\Cap$-closure of the most elementary sets of rankings of the form $\{Ci \gg Cj\}$, containing only a single atomic ranking.

We finish the section by visualizing a part of the domain of sets of rankings and showing the structure that pairwise ranking-union $\Cup$ and alternative-addition $\Cap$ impose on it. We concentrate on proper sets, though we also provide a couple of examples of non-proper sets. Proper sets are given together with their representative tableaux, so the effects of applying $\Cup$ can be observed in parallel on the sets of rankings and the corresponding tableaux. Double arrows leading from some S and T to another set signify that that set is $S \Cup T$. Single arrows lead to $S \Cap T$, which is not always proper. Such links are not shown for all $S \Cup T$ and $S \Cap T$ in the picture, to avoid overcrowding.

(102)

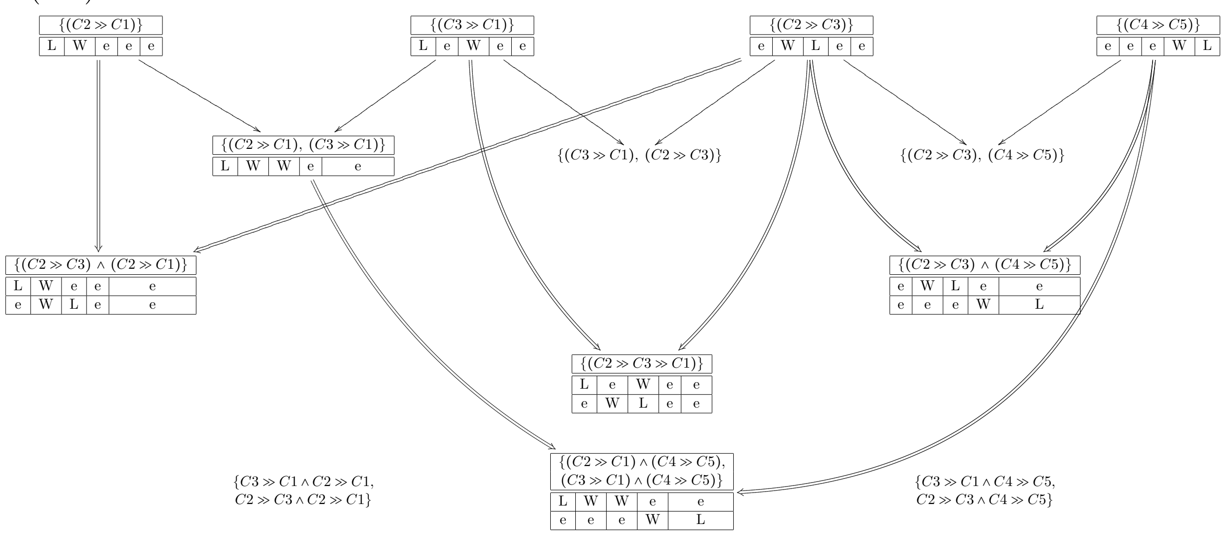

5.5 Summary of Section 5

- Truth for partial ranking can be naturally extended for truth for sets of (partial) rankings. A set S is true at tableau M iff every ranking in S is true at M, 44.
- Each equivalence class of sets of rankings has unique largest total representative S^{Tot} 48 and minimal proper representative S^{Min} 50.
- Each S^{Min} representative is the maximal set of rankings for some tableau M (as defined in 51), by 84, and each set of rankings maximal for some tableau (each proper set) is the S^{Min} representative of its equivalence class, 83.

- Each set of rankings which only contains meaningful rankings with one dominated constraint Ci (that is, each set of Ci-alternatives, see 63) is proper, 64.
- A set is proper iff it is equal to its minimal proper extension (74), by 76. Thus a proper set is the pairwise ranking-union of its S_{Ci} sets.
- There are procedures, effective in the finite case, which compute both directions of the one-one correspondence between arbitrary tableaux and their maximal sets of rankings (see 57 and 77). The procedure in 57 is essentially the solution for the general form of the OT Ranking problem.
- As S^{Min} representatives and proper sets of rankings are the same thing, there is a one-one correspondence between equivalence classes of sets of rankings and equivalence classes of tableaux.
- Two sets of rankings are equivalent iff they have the same minimal proper extension, 88.
- All proper sets are generated from special atomic S_{Ci} sets 92 by applying $\Cup$-closure.

6 Conclusion

We have concluded our investigation of the domain of sets of rankings as relevant for OT. We now know exactly which sets of ranking are proper, and thus represent possible OT grammars. In the course of our work, we have overcome the dichotomy between (sets of) rankings and tableaux, having identified a tableau and its corresponding maximal set of rankings as equivalent objects. The perspective on the problem of finding a ranking compatible with a given tableau is now entirely different from what it used to be: we have precise methods for computing all such rankings for any tableau.

The present work creates a framework in which previous research we often built on may be conveniently placed, and possibly seen in new light. For instance, in Appendix A we did not devise any new algorithms for tableau transformations, but instead analyzed the space of all such possible transformations in which the algorithms proposed earlier could be situated. Similarly for the main part of the paper and OT learning algorithms: specific algorithms such as RCD may now be viewed as efficient, heuristic ways to obtain at least some answer to the narrow form of the OT Ranking problem quickly, and their usefulness can now be measured by how much they allow us to gain from giving up the computation of the full corresponding ranking set. Besides, the existence of a precise answer to the general version of the Ranking problem allows for creation of new metrics for evaluating results produced by heuristic algorithms cutting corners.

Besides such contributions to the existing research areas, there are also four new areas of formal research opened by the analysis just presented.

First, our analysis has been conducted for the general case of Gen and Con, under the assumption that any sequence of W-s, L-s and e-s of the right length is a legitimate comparative row. Of course, practical choices of Gen and Con will almost always depart from that analytical ideal, at the very least because in realistic phonological theories different constraints are not necessarily independent. So realistic Gen and Con will collapse some distinctions which our ideal unrestricted Gen and Con make. Our analysis of the unrestricted case provides a baseline for the study of the spaces of grammars generated by actual Gen and Con components, and a toolkit of methods for such study. In particular, it becomes reasonable to ask how exactly the space created by a Con with non-independent constraints can be different from the domain of proper sets generated by the idealized Con without dependencies.

The second area for future research involves working with OT in the infinite. For some kinds of infinity in OT — e.g., the infinite nature of the set of candidates generated by Gen — it has been already established that those are quite harmless. But there may also exist not so harmless sources of infinity. Most importantly, what if the Con set is infinite? The harmlessness of infinite Gen is determined precisely by the finite nature of Con, which guarantees there can be only a finite number of harmonically unbound candidates for any input. Once Con also becomes infinite, this becomes false. But after having developed a logical analysis of OT rankings, we can use some guidance from logic when doing precise work with infinite Con.[31]

A related area involves heuristic, probabilistic analysis of very large tableaux, be they literally infinite, or simply very big. Even finite constraint sets may be huge, but at the same time phonological analytical practice seems to suggest that actual grammars only use relatively small portions of the very big universal space. It then becomes important to develop methods which would allow to drastically, and with a great degree of certainty, cut the space actually relevant for the grammar. The research in this area should inform, and be informed by, the research on infinite tableaux.

Finally, the precise characterization of the space of classical OT grammars allows us to, first, ask similar questions about other existing grammatical frameworks, and compare them explicitly with the classical OT, and secondly, to potentially formulate new modifications for OT, not wandering in the dark trying to come by a useful mutation, but rather formulating what exactly we want from our grammatical framework, and then tweaking some existing system in the right way.

[31] Is this, however, a realistic concern, or just an ivory tower type question? A skeptical reader may argue that we should not worry about infinite Con at all, for in practice we will never have to have such a constraint set. Besides, wouldn't it make the tasks the language user has to perform impossible to complete? Not necessarily so. Think, for instance, of infinite families of constraints based on some continual phonetic phenomenon, such as "articulate the feature F for x milliseconds". Of course, intuitively we feel such families are quite harmless, even though they are infinite. But the study of OT in the infinite is needed precisely because without it, we cannot really tell what kinds of infinities are good, and what kinds are disastrous.

References

[Anttila, 1997] Anttila, A. (1997). Deriving variation from grammar. In Hinskens, F., van Hout, R., and Wetzels, L., editors, *Variation, Change and Phonological Theory*, pages 35–68. John Benjamins.

[Anttila and Cho, 1998] Anttila, A. and Cho, Y. Y. (1998). Variation and change in Optimality Theory. *Lingua*, 104(1):31–56.

[Boersma, 1997] Boersma, P. (1997). How we learn variation, optionality, and probability. In *Proceedings of the Institute of Phonetic Sciences of the University of Amsterdam*, volume 21, pages 43–58. ROA 221.

[Brasoveanu and Prince, 2005] Brasoveanu, A. and Prince, A. (2005). Ranking and necessity, Part I. ROA 794.

[Brasoveanu and Prince, 2011] Brasoveanu, A. and Prince, A. (2011). Ranking and necessity: the Fusional Reduction Algorithm. *Natural Language and Linguistic Theory*, 29(1):3–70. *A revised version of ROA 794.*

[Hayes, 1997] Hayes, B. (1997). Four rules of inference for ranking argumentation. http://www.linguistics.ucla.edu/people/hayes/otsoft/argument.pdf.

[Jarosz, 2006] Jarosz, G. (2006). *Rich lexicons and restrictive grammars - maximum likelihood learning in Optimality Theory.* PhD thesis, Johns Hopkins University.

[Magri, 2009] Magri, G. (2009). *A Theory of Individual-Level Predicates Based on Blind Mandatory Implicatures. Constraint Promotion for Optimality Theory.* PhD thesis, MIT.

[Moreton, 1996] Moreton, E. (1996). Non-computable functions in Optimality Theory. *(A version revised in 1999 exists as ROA 364.).*

[Prince, 2000] Prince, A. (2000). Comparative tableaux. ROA 376.

[Prince, 2002] Prince, A. (2002). Entailed ranking arguments. ROA 500.

[Prince, 2006] Prince, A. (2006). No more than necessary: beyond the four rules, and a bug report. ROA 882.

[Prince, 2010] Prince, A. (2010). OT without ranking. *Handout for Harvard Linguistics Colloquium, February 2010.*

[Prince and Brasoveanu, 2010] Prince, A. and Brasoveanu, A. (2010). The formal structure of ranking arguments in ot. Rutgers University and UC Santa Cruz.

[Prince and Smolensky, 1993] Prince, A. and Smolensky, P. (1993). Optimality theory: Constraint interaction in generative grammar. ROA 537.

[Riggle, 2008] Riggle, J. (2008). Counting rankings. *Paper presented at LSA 2008*.

[Samek-Lodovici and Prince, 1999] Samek-Lodovici, V. and Prince, A. (1999). Optima. ROA 363.

[Tesar, 1995] Tesar, B. (1995). *Computational Optimality Theory*. PhD thesis, University of Colorado. ROA 90.

[Tesar, 1997] Tesar, B. (1997). Multi-recursive constraint demotion. ROA 197.

[Tesar and Smolensky, 1996] Tesar, B. and Smolensky, P. (1996). Learnability in Optimality Theory (longer version). Technical report, John Hopkins University, Department of Cognitive Science.

A Equivalence classes of OT tableaux

It is a well-known fact that some OT tableaux are equivalent, meaning that exactly the same sets of rankings are compatible with them.[32] As the simplest example, the order of rows in a tableau does not have any impact on which rankings the tableau is compatible with — this fact is so obvious that it is hardly perceived as worthy of any discussion. But in addition to row order changes, there are also less trivial transformations of tableaux which produce what is formally a different object, but yet remains just the "semantically"-same tableau in a different guise, in the sense that any ranking compatible with the first tableau will be compatible with the second, and vice versa.

It is useful to get a formal handle on this intuition of "tableaux which are essentially the same thing", describing when exactly two tableaux are such, and which relations hold between equivalent tableaux. Formally, we can partition all tableaux into equivalence classes, where each tableau in a class will be compatible with exactly the same set of rankings. The partition will be defined by the specific OT-compatibility relation — we will choose the classical OT compatibility, but other relations can also be used if the theory analyzed is different from standard OT — and the set of rankings which are used to distinguish tableaux — we will use the whole set of non-contradictory rankings for that purpose, and again, choosing a different set may yield different equivalence classes.[33]

Thus formally, two tableaux $T1$ and $T2$ belong to the same equivalence class iff they are compatible, under the relevant notion of compatibility, with exactly the same set of

[32] As in the rest of the paper, by "tableaux" we mean comparative tableaux simply representing data, and not including any representation of a ranking.

[33] For instance, if contradictory rankings are allowed, then what is a huge single class of inconsistent tableaux in classical OT will, on most possible definitions of what compatibility with a contradictory ranking may mean, be subdivided into narrower equivalence classes.

rankings out of the fixed set of legitimate rankings. In logical terms, that means $T1$ and $T2$ have the same theory. From the OT point of view, any two tableaux from the same equivalence class behave indistinguishably.

But the definition of what equivalence classes are is trivial, and what is interesting is the question of which tableaux fall into the same equivalence class. The rest of this section provides a long answer to that question. The plan is as follows: we will first define five syntactic transformations on tableaux which preserve OT-compatibility, and their inverses; so if we apply any sequence of these transformations to a tableau, we will always get a tableau from the same equivalence class. Then we will define a representative of the equivalence class: a certain normal form for tableaux. We will prove the existence and uniqueness of the normal form for every equivalence class, and thus establish that we can use the normal form as a proper representative of a class as a whole. And this will make things easier in the main part of the paper. We will finish by discussing several nice consequences from that main normal form result. In particular, it can be easily shown that any two equivalent tableaux can be transformed into each other by an effective procedure consisting of a number of applications of our equivalence-preserving operations. This essentially provides a syntactic counterpart of the semantic definition of equivalence classes of tableaux: two tableaux can be transformed into each other by a sequence of our five compatibility-preserving operations iff they are in the same equivalence class.

A.1 Compatibility-preserving transformations

There are five such operations, which we will now introduce one by one. For each one, we will first demonstrate how it works using simple examples, and then prove formally that it preserves OT-equivalency.

A.1.1 Row swap

It is obvious that if we simply change the order of rows in a tableau, that does not change the set of rankings which are compatible with it. So the operation of **row swap**, whose inverse is also a row swap, is always compatibility preserving. In fact, as the order of rows in a tableau is not used in OT at all, we can sloppily go from the tableau talk to the set-of-rows talk, which we will do throughout this section. So sometimes we will write that some row r is in a tableau T as $r \in T$; $T \setminus r$ for the result of deleting r from T; and $T \cup r$ for the result of adding r to T. But when we will define a normal form for tableaux later on, it will be more convenient to work with ordered sets of rows, so we do not reject order altogether.

To make formal definitions more "linear", we use a tuple $\langle r_1, r_2, ..., r_n \rangle$ to represent a tableau, with r_1 being its first row, r_2 its second row, and so forth. We will always assume that the order of columns in all rows is the same.[34]

[34]It is, of course, easy to allow to swap columns, which would also be an equivalency-preserving operation

(103) *Row swap* **Sw**:

For natural numbers k, l, and a tableau $T = \langle r_1, ..., r_k, ..., r_l, ...r_n \rangle$,

$\mathbf{Sw}(T)(k)(l) = \langle r_1, ..., r_l, ..., r_k, ...r_n \rangle$

It is trivial that row swap preserves equivalency. $\mathbf{Sw}(T)(k)(l)$ is its own inverse.

A.1.2 Row merger and splitting

The second, and more interesting, transformation that we can apply to a tableau without affecting the set of rankings which are compatible with it is the merging of rows with the same W-s into a single row, as well as its inverse, splitting of a row into several rows with the same W-s as the original one. In what follows, we call those two **row mergers** and **row splittings**.

We first give an example of two equivalent tableaux which can be turned into each other by row merger and its inverse:

(104)

$C1$	$C2$	$C3$
W	L	e
W	e	L

(105)

$C1$	$C2$	$C3$
W	L	L

Any ranking ϕ compatible with 104 has to include atomic rankings $C1 \gg C2$ and $C1 \gg C3$. But then ϕ is also compatible with 105. Conversely, any ranking ψ compatible with 105 has to have the same two atomic rankings $C1 \gg C2$ and $C1 \gg C3$, and thus is compatible with 104 as well. It is not too hard to notice that this is not an accident, and that more generally, we can always merge two rows from the same tableau without affecting the set of compatible rankings, if they share the same W-s and have different sets of L-s, into a row which has W-s in the same places, and L-s in every constraint which had an L in at least one of the original rows. Similarly we can split a row with multiple L-s into several rows just as well.[35]

To keep the definitions readable, I chose to state them in procedural terms; it is possible to state them declaratively as well, but that would require more technical notation, and could hinder the intuitive understanding of what is going on.

---

if done properly, but then we would have to keep track of the names of the columns. So to keep things simpler, we assume a fixed order.

[35] The first time such operations of splitting and merging were explicitly suggested as useful OT-equivalence-preserving operations that I know of is in [Magri, 2009, p. 145], but their correctness follows easily from the apparatus of [Prince, 2002].

(106) *Row merger* **Mrg**:

For a tableau T with n rows, and a set Ind of natural numbers such that, first, $i \leq n$, and second, for every $i, j \in Ind$, $W(r_i) = W(r_j)$,

$\mathbf{Mrg}(T)(Ind) = T'$,

where T' is the result of deleting from T the row r_i for each $i \in Ind$, and then adding to the resulting tableau a new row q s.t. $W(q) \coloneqq W(r_i)$ (any r_i will do, as they have the same W-sets) and $L(q) \coloneqq \bigcup_{i \in Ind} L(r_i)$.

(107) *Row splitting* **Spl**:

For a row r_k in a tableau T and a collection of sets Col s.t. for each $s_i \in Col$, $s_i \subseteq L(r_k)$, and $\bigcup_{s_i \in Col} s_i = L(r_k)$,

$\mathbf{Spl}(T)(k)(Col) = T'$,

where T' is the result of deleting r_k from T, and then for each $s_i \in Col$ adding a row q_i s.t. $W(q_i) = W(r_k)$, and $L(q_i) = s_i$.

It is easy to see that for each set of natural numbers Ind such that $\mathbf{Mrg}(T)(Ind) = T'$, there exist k and a collection of sets Col s.t. $\mathbf{Spl}(T')(k)(Col) = T$, and vice versa. In other words, each application of **Mrg** can be undone with a corresponding application of **Spl**.

(108) Row mergers **Mrg** and row splittings **Spl** preserve OT-compatibility.

Proof. Let A be an arbitrary tableau. Fix a row r and several rows $r_1, ..., r_n$ s.t. $W(r) = W(r_1) = ... = W(r_n)$, and $L(r) = \bigcup_i L(r_i)$ (in other words, r is merged from $r_1, ..., r_n$.) Consider $A \cup r$ and $A \cup \{r_1, ..., r_n\}$. If some ranking is true in $A \cup r$, it must put every L in r under one of the W-s. But then every L in $r_1, ..., r_n$ will also be covered under ϕ, as those rows have W-s in the same places as r, and can have an L only where r also has one. Conversely, if every row $r_1, ..., r_n$ is accounted for in B by some ranking ϕ, the same ϕ is compatible with r as well: for any L in r, there is some r_i where that L is covered by a W, and r has W-s in the same places. Similarly, it is easy to show that if a ranking is not compatible with r, then it will fail to be compatible with at least one of $r_1, ..., r_n$, and vice versa. So $A \cup r$ and $A \cup \{r_1, ..., r_n\}$ are compatible with exactly the same rankings. ⊣

Despite the fact **Mrg** and **Spl** are such simple operations, they often make proofs in the main part of the paper tremendously easier: the possibility to split any multiple-L row into single-L rows allows us to analyze without loss of generality the behavior of single-L rows instead of analyzing all arbitrary rows.

A.1.3 Inference elimination and introduction

The third operation which we need is **inference elimination** (and its inverse, **inference introduction**.) Inference elimination deletes from a tableau a *superfluous* row — a row that does not add any new restrictions on which rankings are compatible with the tableau. A tableau T with a superfluous row q is compatible with exactly the same rankings as $T \setminus q$.

But when is a row in a given tableau superfluous? There exists a simple syntactic answer to that question, discovered by [Prince, 2002], see his Prop. 2.5. Thus a reader not interested in the details may jump from here right to 118. In the rest of this section, we will provide another, semantic criterium for superfluousness.

Since OT rankings are transitively closed, we can omit rows which can be deduced by transitivity from a set of other rows: if a ranking says that $C1 \gg C2$ and $C2 \gg C3$, then it must also say that $C1 \gg C3$. Thus the bigger tableau in 109 is OT-equivalent to the smaller tableau in 110: the third row of 109 can be accounted for by the atomic ranking $C1 \gg C3$ which is required anyway if the ranking includes $C1 \gg C2$ and $C2 \gg C3$ needed to account for the first two rows. So every ranking compatible with 109 has to be compatible with 110, and vice versa.

(109)

$C1$	$C2$	$C3$
W	L	e
e	W	L
W	e	L

(110)

$C1$	$C2$	$C3$
W	L	e
e	W	L

But obviously detecting entailment by transitivity will not be as easy as it is in 109. Let's take an arbitrary row q with a single L and find out under what conditions that row is superfluous in a tableau $T = \langle r_1, ..., r_n \rangle$.

If q is superfluous, that means that however we can cover the L-s in the other rows with W-s, that covering will put at least one of q's W-s on top of its L. When does it happen? First of all, there should be some r_i with $L(q) \subseteq L(r_i)$, for otherwise the L-constraint of q will be undominated in some rankings compatible with T. It is useful to introduce the following notion:

(111) For a tableau T, a row $r_i \in T$, and a $Cj \in L(r_i)$, a **possible domination chain** is a sequence of constraints $\langle C_{k_1}, ..., C_{k_n}, Cj \rangle$ s.t. $C_{k_n} = Cj$, a single constraint never occurs twice in the chain, and for each C_{k_l}, $C_{k_{l+1}}$ there is a row $r_m \in T$ for which $C_{k_l} \in W(r_m)$, $C_{k_{l+1}} \in L(r_m)$.
A **maximal possible domination chain** is a possible domination chain for which there is no $r_m \in T$ s.t. $C_{k_1} \in L(r_m)$.

To illustrate how the notion works, let us find the domination chains for the L in q_2 in the following tableau:

(112)

	$C1$	$C2$	$C3$	$C4$
q_1	W	e	L	e
q_2	e	W	W	L

There are two maximal chains for $C4$ and q_2: one is $\langle C2, C4\rangle$, the other, $\langle C1, C3, C4\rangle$. What is the significance of those chains for entailment? Consider a row (W, e, e, L). This row would have been entailed by 112 if all rankings compatible with the tableau included the atomic ranking $C1 \gg C4$. But that is not the case: the only row requiring an atomic ranking of the form $Cx \gg C4$ is q_2; and it suffices to have $C2 \gg C4$ to account for q_2, with $C2$ being undominated by any other constraint, including $C1$. So there exists a ranking compatible with 112, but not compatible with (W, e, e, L). Now consider a row (W, W, e, L). Take some arbitrary ranking ϕ compatible with 112. Suppose ϕ accounts for the L in $C4$ in q_2 with the help of $C2 \gg C4$. This atomic ranking makes (W, W, e, L) true as well. Then suppose ϕ does not include $C2 \gg C4$, but includes $C3 \gg C4$ instead. If that would be all that ϕ includes, clearly it would not be true in (W, W, e, L). But ϕ has to account for the L in $C3$ in q_1, so it has to include some atomic ranking for that. Given what q_1 is, it can only be $C1 \gg C3$. But then from that and $C3 \gg C4$ we have $C1 \gg C4$ by transitivity, and thus ϕ will have to be compatible with (W, W, e, L) as well.

In the case of (W, e, e, L), there was a domination chain for the L in $C4$ in q_2 which did not include any of the W-constraints of our new row. And in the case of (W, W, e, L), every domination chain included a W-constraint from it. It is precisely in the latter situation that the new row was entailed.

If there are several rows in T which have an L in the L-constraint of q, it suffices for q to be superfluous that only one of those L-s in T is such that all its domination chains include at least one constraint for $W(q)$: the existence of such an L is enough to guarantee that the L in q will be covered by one of q's W-s in any ranking compatible with T.

Furthermore, if q has not a single, but multiple L-s, it suffices if for each of those L-s, there is a "buddy" L in some row in T such that all its domination chains include a constraint from $W(q)$.

(113) **Maximal chain lemma**. (If Con is finite, then) for any ranking ϕ compatible with T, row $r \in T$, and constraint $Ci \in L(r)$, there is a maximal domination chain $\langle C_{k_1}, ..., C_{k_n}, Ci\rangle$ for r and T such that $C_{k_1} \gg ... \gg C_{k_n} \gg Ci$ is in ϕ.

Proof. As ϕ is compatible with r, there exists some constraint $C_{k_n} \in W(r)$ s.t. $C_{k_n} \gg Ci$ is in ϕ. If there is no $r_j \in T$ s.t. $C_{k_n} \in L(r_j)$, $\langle C_{k_n}, Ci\rangle$ is a maximal chain, and we are done. Suppose there is such an r_j. As ϕ is compatible with r_j, it must contain an atomic ranking $C_{k_{n-1}} \gg C_{k_n}$ s.t. $C_{k_{n-1}} \in W(r_j)$. If $\langle C_{k_{n-1}}, C_{k_n}, Ci\rangle$ is a maximal chain, we are done, if not, we continue. The induction will have to stop at some moment because since CON is finite, we will eventually exhaust all the constraints in it, and ϕ, being compatible with a tableau, cannot be contradictory. $\dashv$

(114) **Superfluous row theorem**.
A tableau $T = \langle r_1, ..., r_n \rangle$ entails a row q iff for each $Ci \in L(q)$, there exists a row $r \in T$ s.t. in every maximal domination chain for Ci, r, and T, there is a constraint C_{k_l} in it s.t. $C_{k_l} \in W(q)$.

Proof. ($\Leftarrow$) Take an arbitrary ϕ compatible with T. Fix an arbitrary $Ci \in L(q)$. By the Maximal chain lemma 113, there is a maximal chain $\langle C_{k_1}, ..., C_{k_n}, Ci \rangle$ for r and T s.t. ϕ includes $C_{k_1} \gg ... \gg C_{k_n} \gg Ci$. By assumption, some C_{k_l} from this chain is in $W(q)$, which guarantees that ϕ puts a W on top of the L in Ci in row q. As Ci was arbitrary, the same holds for all constraints in $L(q)$, so ϕ accounts for all of q's L-s. As ϕ was arbitrary, all rankings compatible with T are compatible with q as well.

($\Rightarrow$) Suppose T entails q, and fix an arbitrary ϕ compatible with T. Without loss of generality, assume that ϕ is T-maximal. By assumption, for each $Ci \in L(q)$, there is some $Cj \in W(q)$ s.t. $Cj \gg Ci$ is in ϕ. As ϕ is T-maximal, $Cj \gg Ci$ cannot be subtracted from ϕ without making it incompatible with T. Therefore there is a sequence of rows in T s.t. there is a chain $Cj \gg C_{k_1} \gg ... \gg C_{k_n} \gg Ci$ in ϕ, and for each pair of neighbors in it, there is an $r \in T$ which has an L in the dominee of the pair, and a W in the dominator. But then $\langle Cj, C_{k_1}, ..., C_{k_n}, Ci \rangle$ is a domination chain for some r in T. If we extend that chain to a maximal chain, which can always be done with a finite CON, we derive the conclusion. $\dashv$

The Superfluous row theorem 114 tells us when exactly it is OK to delete a row because it is superfluous.

(115) *Inference elimination* **Inf**:

For a tableau $T = \langle q_1, ..., q_m \rangle$ and a row $q_k \in T$,
if for all $Ci \in L(q_k)$, there exists $r \in T$ s.t. for every maximal domination chain for Ci and r, some $Cj \in W(q_k)$ is in that chain, then
$\mathbf{Inf}(T)(k) = \langle q_1, ..., q_{k-1}, q_{k+1}, ...q_m \rangle$, that is, the result of deleting q_k from T;
otherwise, $\mathbf{Inf}(T)(k)$ is not defined.

(116) *Inference introduction* **InfIntro**:

For a consistent tableau $T = \langle q_1, ..., q_m \rangle$, a natural number $k \leq m$, and a row q s.t. for all $Ci \in L(q_k)$, there exists $r \in T$ s.t. for every maximal domination chain for Ci and r, some $Cj \in W(q_k)$ is in that chain,
$\mathbf{InfIntro}(T)(k)(q)$ is defined to be $\langle q_1, ..., q_{k-1}, q, q_{k+1}, ...q_m \rangle$, that is, the result of adding q as the k-th row into T.

Our characterization of row superfluousness is not the only possible one. [Prince, 2002] gives another characterization based on the syntactic operation of fusion, written $\circ$, of comparative rows.

(117) *The operation of fusion* ([Prince, 2002]):

- $L \circ X = X \circ L = L$
  (L-containing constraints always fuse to L)
- $e \circ e = e$
  (all-e constraints fuse to e)
- $W \circ e = e \circ W = W \circ W = W$
  (L-less constraints with some W-s fuse to W)

The operation straightforwardly generalizes from two rows to arbitrary number of rows, so it makes sense to speak of the fusion of a set of rows/a tableau T, written fT.

One of the main results of [Prince, 2002] (Proposition 2.5, p. 14) is that a row q is superfluous in tableau T exactly when there is a subset U of T whose fusion fU is such that any ranking compatible with fU is also compatible with q — namely, the singleton set $\{fU\}$ makes q superfluous.[36]

So we now have two criteria for identifying the situation when q is superfluous in T:

(118) A row q in a consistent tableau T is superfluous

1. iff for all $Ci \in L(q)$, there exists $r \in T \setminus q$ s.t. for every maximal domination chain for Ci and r, some $Cj \in W(q)$ is in that chain;
2. iff there is a subset U of $T \setminus q$ s.t. fU entails q.

Our domination chain requirement is thus essentially the semantic counterpart of Prince's syntactic fusion requirement. The existence of subset U which fuses to a row entailing the superfluous row q guarantees that the domination chain requirement will be met, and vice versa.

A.1.4 False W elimination (a.k.a. Generalized W Removal) and introduction

The fourth equivalency-preserving operation which we should introduce is called **Generalized W removal** (the term of [Prince, 2006]), or **false W elimination** (our term). Consider the following tableau:

(119)

$C1$	$C2$	$C3$
W	W	L
e	L	W

[36]Prince considers only total rankings when proving his results about fusion and entailment, but as we show in the main part of the paper that equivalence for total rankings implies equivalence for partial rankings and sets of partial rankings, it is straightforward to generalize his result to those entities as well.

The first row of 119 can be accounted for by either $C1 \gg C3$ or $C2 \gg C3$. However, the latter contradicts any atomic ranking which could take care of the L in the second row, because the second row requires $C3 \gg C2$ in order to be explained. So in effect the W in $C2$ in the first row does not do any real work in this tableau: we can omit it, getting 120, and still the tableau will be compatible with exactly the same set of rankings (namely, only with the total ranking $C1 \gg C3 \gg C2$.)

(120)

$C1$	$C2$	$C3$
W	e	L
e	L	W

Consider a slightly more complex example:

(121)

$C1$	$C2$	$C3$	$C4$
W	W	L	L
e	L	W	e

In 121, the W in $C2$ in the first row is false. Here is why: the second row forces each ranking compatible with it to include the atomic ranking $C3 \gg C2$. This means that the L in $C3$ in the first row can only be covered by $C1$, therefore we get that every compatible ranking includes $C1 \gg C3 \gg C2$. But then even if we want to account for the L in $C4$ in the first row ranking $C2$ over it, by transitivity we will always get $C1 \gg C2 \gg C4$. Thus any ranking compatible with the tableau includes $C1 \gg C3$ and $C1 \gg C4$. So the W in $C2$ does not do any useful work: and we can just as well work with a simpler 122 instead of 121: they are compatible with the same set of rankings.

(122)

$C1$	$C2$	$C3$	$C4$
W	e	L	L
e	L	W	e

None of the transformations we defined so far connects 119 and 120, or 121 and 122, so we need a new operation, and its inverse, which would transform the tableaux from those pairs into one another.

In fact, the exact elimination operation we need was introduced in [Prince, 2006] under the name Generalized W Removal, or **GWR**, and it covers not only the cases above, but a variety of cases where a W in some cell is false, that is, can be replaced with an e without affecting the set of compatible rankings. For a comprehensive discussion of the various cases of false W-s, we refer the reader to Section 2 of [Prince, 2006].

False W elimination **GWR** is not an instance of inference elimination **Inf**, but it also depends on drawing inferences. Both with **Inf** and **GWR**, given a row, we try to make it "simpler" on the basis of inferences derived from the rest of the tableau. Only in inference elimination we delete the whole row, and in false W elimination, we erase one of the W-s in it. Therefore the methods for showing correctness of **GWR** are similar to the ones we used for **Inf**.

Here is how **GWR** will work: Given a row r with multiple W-s in some consistent T, for each $Ci \in W(r)$ we can create a new technical row r_i s.t. $L(r_i) := \{Ci\}$, and $W(r_i) := (W(r_k) \cup L(r_k)) \setminus \{Ci\}$. Then we can check for this r_i whether it is entailed by the rest of the original tableau. If yes, false W elimination **GWR** may apply: we can replace the W in C_i in row r with an e, and it will preserve equivalence. It is not (and should not be) immediately obvious from this description why this works: the proof of **GWR**'s correctness, given below, is not trivial.

As usual, it is easy to define the inverse operation, **GWRIntro**: for each e of row r in some Ci, we can build a technical row r_i the same way we did for **GWR**, namely, with $L(r_i) := \{Ci\}$, and $W(r_i) := (W(r_k) \cup L(r_k)) \setminus \{Ci\}$. Then we, again as for **GWR**, check if r_i is entailed by the rest of the tableau. If yes, we can turn the e in Ci into a W. So essentially **GWR** and **GWRIntro** are flip-flop operations: they share the same precondition which establishes that it does not matter whether you have an e or a W in Ci in row r, so you can freely change one into another.

(123) *False W elimination* **GWR**:

For a tableau $T = \langle r_1, ..., r_n \rangle$, and i and k s.t. $r_k \in T$, $Ci \in W(r_k)$,

if row q defined by $L(q) := \{Ci\}$, $W(q) := (W(r_k) \cup L(r_k)) \setminus \{Ci\}$ is such that q is entailed by $T \setminus r_k$, then

$\mathbf{GWR}(T)(k)(i) = \langle r_1, ..., r_{k-1}, r'_k, r_{k+1}, ...r_n \rangle$, where r'_k is a row with $W(r'_k) := W(r_k) \setminus Ci$ and $L(r'_k) := L(r_k)$.

(124) *False W introduction* **GWRIntro**:

For a tableau $T' = \langle r_1, ..., r'_k, ...r_n \rangle$, and i and k s.t. $r'_k \in T'$, $Ci \notin W(r'_k) \cup L(r'_k)$,

if row q defined by $L(q) := \{Ci\}$, $W(q) := W(r'_k) \cup L(r'_k)$ is such that q is entailed by $T \setminus r'_k$,

then $\mathbf{GWRIntro}(T')(k)(i) = \langle r_1, ..., r_{k-1}, r_k, r_{k+1}, ...r_n \rangle$, where r_k is a row defined by $W(r_k) := W(r'_k) \cup Ci$, $L(r_k) = L(r'_k)$.

(125) **GWR** and **GWRIntro** preserve OT-equivalence.
Moreover, if changing some W into an e in some $r_k \in T$ does not change the set of rankings T is compatible with, and r_k is not entailed by $T \setminus r_k$, then the precondition for **GWR** applies to that W. Similarly for **GWRIntro** and changing an e into a W.

There is an alternative, equivalent formulation of the two statements above, which suggest more vividly the structure of the proof to follow:

- If q as defined above is entailed by $T \setminus r_k$, then the relevant W is false, and T and T' are compatible with the same sets of rankings.
- If q is not entailed by $T \setminus r_k$, and r_k is not entailed by $T \setminus r_k$, then the sets of rankings compatible with T and with T' are different.

[Prince, 2006, p. 12] gives an elegant fusion-based proof of the statement in 125, which is in fact a slightly disguised form of Prince's GWR Theorem, his (13). The right-to-left direction of Prince's theorem is the first statement of 125, and the left-to-right direction is the second, "moreover" statement. The only substantial difference between Prince's version of the theorem and ours is that we slightly strengthen it, pinning down exactly where the assumption of independence is needed (namely, it is required to prove the second statement), and requiring only that r_k was independent from $T \setminus r_k$, not that T as a whole did not have superfluous rows. This strengthening could in fact have been done with respect to Prince's original proof as well, as the independence precondition is also used there once, and for the independence of r_k (Prince's row π in his proof) only. The practical significance of this strengthening is that false W elimination can be applied without checking for independence of T's rows.

Just as we did for inference elimination, we give a semantic proof for false W elimination. Comparing the syntactic proof of Prince's with the semantic proof below, we can get a better insight into why **GWR** as defined is bound to work.

"Semantic" proof of 125. Let r_k be a row with a W in Ci, r'_k, an identical row with an e in Ci, and q defined as in 123. Let T be a tableau with r_k, and T' the same tableau with the relevant W withdrawn — with r'_k replacing r_k.

We will first prove that **GWR** and **GWRIntro** preserve OT-equivalence. Namely, we will show if $T \setminus r_k$ entails q, then there cannot be a ranking witnessing that the W in Ci is not false, and thus deleting this W does not offend equivalency.

Assume without loss of generality that our tableau T is non-contradictory. (If T is contradictory, then any W in it is false anyway.) Then the rankings which are compatible with T are a superset of the rankings compatible with T', regardless of whether the W is false or not: a new W instead of an e can only create new ways to account for the L-s of r_k, it cannot rule out any possibilities that already existed without the new W. So if our W is *not* false, there is some ranking compatible with T, but not with T'.

Towards a contradiction, assume there is such a ranking ϕ even though q is entailed by $T \setminus r_k$. ϕ should be compatible with r_k, but not with r'_k. Without a loss of generality, assume that ϕ is T-maximal. Clearly for some $Cj \in L(r_k)$, ϕ has to include the atomic ranking $Ci \gg Cj$ which explains the L in Cj in T, but not in T'. By assumption, ϕ is true in q. But in q, $Ci \gg Cj$ puts an L on top of some W, so the L in Ci itself has to be explained by some other atomic ranking in ϕ, of the form $Ck \gg Ci$ for $Ck \in W(r_k) \cup L(r_k)$.

Suppose $Ck \in W(r_k)$. Then in ϕ, $Ck \gg Ci \gg Cj$. But Ck is in $W(r'_k)$ as well, so ϕ is compatible with r'_k, and thus with T', contrary to assumption.

Suppose $Ck \in L(r_k)$. Then there should be some other constraint $Cl \in W(r_k)$ s.t. $Cl \gg Ck$ is in ϕ. From that, we get that $Cl \gg Ck \gg Ci \gg Cj$, but Cl is in $W(r'_k)$, and thus ϕ is compatible with r'_k as well.

We derived a contradiction, so if there is a q as defined, the W is false.

To prove the second statement of 125, we need to show that if q is *not* entailed by $T \setminus r_k$ and r_k is not superfluous in T, then there is a ranking compatible with T, but not with T'.

Assume that q is not entailed by $T \setminus r_k$, and fix a ranking ϕ witnessing it, so that ϕ is true in $T \setminus r_k$, but not in q. W.l.o.g., assume that ϕ is $T \setminus r_k$-maximal.

Then for every Cj in $(W(r_k) \cup L(r_k)) \setminus Ci$, the atomic ranking $Cj \gg Ci$ is not in ϕ — otherwise ϕ would have been compatible with q.

Now consider ϕ and T. As r_k is not superfluous in T, ϕ cannot be true in r_k, and there is some Cl in $L(r_k)$ s.t. for all $Cm \in W(r_k)$, the atomic ranking $Cm \gg Cl$ is not in ϕ. W.l.o.g., assume there is only one such Cl. We need to show that we can extend ϕ to a ranking compatible with r_k, but not r'_k.

Consider $\phi \wedge (Ci \gg Cl)$. If it is non-contradictory, clearly it is compatible with r_k, but not r'_k. Suppose $\phi \wedge (Ci \gg Cl)$ is contradictory. Then there ϕ includes the atomic ranking $Cl \gg Ci$. But then ϕ is compatible with q, contrary to assumption. Hence $\phi \wedge (Ci \gg Cl)$ is non-contradictory. ⊣

Note that iterating false W elimination alone does not guarantee we get the most streamlined version of the tableau. If there are multiple L-s in a row, some W may be half-false and half-genuine: it can be false for some L-s, and non-false for the others. Consider the following pair of tableaux:

(126)

$C1$	$C2$	$C3$	$C4$
W	W	L	L
e	L	W	e

(127)

$C1$	$C2$	$C3$	$C4$
W	W	L	e
W	W	e	L
e	L	W	e

Even though from the second row in 126 we see that the W in $C2$ in the first row cannot account for the L in $C3$, that W is not fully false because it can account for the L in $C4$. However, if we split the first row into two, each having one L, using **Spl**, as in 127, we can eliminate the W in $C2$ from one of the resulting rows, getting an equivalent tableau with more rows, but no even half-false W-s:

(128)

$C1$	$C2$	$C3$	$C4$
W	e	L	e
W	W	e	L
e	L	W	e

A.1.5 Contradictory jump and backward contradictory jump

So far, we have not discussed trivial, or degenerate rows: rows with only W-s and e-s, and rows with only L-s and e-s. But as we aim to provide a toolkit for working with all tableaux, including those with such degenerate rows, we should do so.

As for only-W rows ("happy rows"), they are entailed by any tableau whatsoever, including the empty one. Such rows can be eliminated by **Inf** and added by **InfIntro**: as they have no L-s, the precondition on those operations is vacuously met.

But for unhappy rows with only L-s and e-s, inference elimination does not work, and for a good reason. Any tableau with such a row is inconsistent, because no ranking is compatible with an unhappy row, and our inference elimination and introduction operations work in the realm of consistent tableaux. So we need two new operations if we want to be able to add and subtract L-only rows from tableaux preserving OT-compatibility.

On the adding side, any row added to an already contradictory tableau, including an L-only row, does not make it consistent. On the subtracting side, we can only subtract an L-only row from a tableau if the result remains contradictory; moreover, any row whatsoever may be subtracted as long as the result is still contradictory. For consistent tableaux, there exist a lot of equivalency classes, and one can "travel" from one member of a class to another using very restricted equivalency-preserving transformations; but all inconsistent tableaux form a huge single equivalence class, and even wild jumps preserve OT-equivalency.[37] So it is an easy exercise to show that the following two operations are equivalency-preserving:

(129) *Contradictory jump* **Jump**:

For an inconsistent tableau $T = \langle q_1, ..., q_n \rangle$, an arbitrary row r, and $1 \leq k \leq n$,

$\mathbf{Jump}(T)(r)(k) = \langle q_1, ..., q_{k-1}, r, q_{k+1}, ...q_n \rangle$

(130) *Backward contradictory jump* **JumpBackward**:

If $T \smallsetminus r$ is inconsistent, then

$\mathbf{JumpBackward}(T)(r) = T \smallsetminus r$

Obviously, **Jump** and **JumpBackward** are inverses.

[37] Note that the equivalence classes into which the domain of tableaux is divided are induced not only by the notion of OT-compatibility, but also by a particular choice of the set of legitimate rankings. The way we have defined rankings, excluding contradictory ones, forces us to say that all contradictory tableaux are equivalent with respect to each other. But had we included contradictory rankings and defined non-trivial conditions for their truth at tableaux, those contradictory tableaux could be distinguished from one another: different rankings having contradictory atomic rankings would be true in them. The equivalence classes of the subdomain of non-contradictory tableaux would still have remained the same, but a lot more equivalence classes would have arisen. Also, there still would be truly contradictory tableaux, those with an L-only row, and for them, we would still have needed a special operation to transform one into another.

A.2 Normal form for OT tableaux

Now we are ready to prove the fundamental result of this section. We know that the five operations we have defined and their inverses (row swaps, row mergers and splittings, inference eliminations and introductions, false W eliminations and introductions, and forward and backward contradictory jumps), when applied to a tableau, always preserve OT-equivalence. When we apply one of them to some T, the result T' is always in the same OT-equivalence class as T was.

But that is not ambitious enough: there is a number of things which the mere existence of such operations by itself does not give us. In particular, we do not have a decision procedure for determining whether a tableau is a given equivalence class, other than a brute force check of whether there exists a ranking compatible with the tableaux in the class, but not with our tableau. Equivalence classes of tableaux themselves are bulky objects, it is not very clear how to represent them efficiently. Finally, it would be nice to be able to transform *any* tableau in an equivalence class into *any other* tableau in that class. By the end of this Appendix, we will achieve all of that, and some more.

The way to prove those results is via normalization of tableaux. Here is the plan for the proofs:

(131)
1. We define a certain class $\mathcal{C}$ of tableaux, to be later shown to be the class of normal forms for equivalence classes of tableaux.
2. We show that any (finite) tableau can be (effectively) transformed into an equivalent tableau in $\mathcal{C}$.
3. We show that each two different tableaux in $\mathcal{C}$ belong to different equivalence classes. Thus a tableau in $\mathcal{C}$ can be used as the name for its equivalence class, its unique representative.
4. From 2 and 3, it follows that any two tableaux in the same equivalence class can be transformed into the same normal form tableau from $\mathcal{C}$. Since transformation sequences can always be inverted, we can always transform any tableau into any equivalent tableau going through the normal form. This also gives the decision procedure for membership in an equivalence class.

A.2.1 Defining the class of normal form tableaux

Fix an enumeration of the constraint set, sending each constraint to a natural number and using the natural ordering. We will assume that the leftmost constraint in a tableau is sent to 1, the second from the left to 2, and so forth. If C_i is sent to a number less than the number C_j is sent to, we write $C_i \prec C_j$. In terms of familiar linear presentations of tableaux, it can be interpreted as "C_i is on the left of C_j". For rows, we similarly write $r \prec q$ when r is higher than q in the tableau.

Let the normal form be the following diagonal form:

(132) **Normal form for OT tableaux**:

1. The only contradictory tableau in the normal form is the tableau with a single L in the first constraint.
2. Each row has at most a single L.
3. There are no rows which can be inference-eliminated.
4. In multiple-W rows, there are no false W-s.
5. The rows are alphabetically ordered in the following manner:
   - The closer the W-constraints to the first constraint, the higher the row: $[\exists C_i : (C_i \in W(r)) \wedge (C_i \notin W(q)) \wedge \nexists C_j : (C_j \prec C_i) \wedge (C_j \in W(q)) \wedge (C_j \notin W(r))] \rightarrow r \prec q$
   - Among the rows with the same set of W-s (each such group is contiguous by the previous condition), ordering is by the position of the L: $W(r) = W(q) \rightarrow [[\exists C_i \in L(r) : \exists C_j \in L(q) : \wedge C_i \prec C_j] \rightarrow r \prec q]$

There is a degree of arbitrariness in how we choose what we use as a normal form. For example, instead of requiring all false W-s to be eliminated, we could have required all of them being present. Or, for instance, we could have used any other contradictory tableau instead of ⟨(L, e, e, ...)⟩ as the normal form contradictory tableau. As long as there is only one normal form tableau per equivalence class, and there is always one, any normal form will do. I chose the particular formulation in 132 because it is convenient for direct use in proofs in this section and in the main part of the paper as well.

It should be obvious that conditions 2-4 are satisfiable in the same tableau. As for 5, it is not too hard to show it is satisfiable if 2-4 are met, as all rows in a tableau satisfying 2-4 have to be distinct, and then it is always possible to order them according to 5.

At this point, it is useful to give several examples of tableaux in the normal form. One such tableau is 133. Conditions 2-4 are obviously fulfilled, and the ordering has to be $\langle \mathbf{a}, \mathbf{b}, \mathbf{c} \rangle$: **a** and **b** have a W in C_1, so they go before **c**; **a** is higher than **b** because **a**'s second W is more to the left than **c** second W.

(133)

	$C1$	$C2$	$C3$	$C4$
a	W	W	e	L
b	W	e	W	L
c	e	W	L	e

On the other hand, the following tableau is not in the normal form, because the W in $C1$ in **d** cannot be really put on top of the L in $C4$, and thus is false: $C1 \gg C4$ would make the ranking incompatible with the row **e**.

(134)

	$C1$	$C2$	$C3$	$C4$
d	W	W	e	L
e	L	e	e	W

However, there is a normal form tableau equivalent to 134 (in fact, we will soon establish that there is always a normal form tableau equivalent to a non-normal form one):

(135)

	$C1$	$C2$	$C3$	$C4$
f	e	W	e	L

The pair of 134 and 135 may create the impression that tableaux in the normal form are somewhat more streamlined than non-normal-form equivalent tableaux. Whether this is true in the general case depends on one's notion of streamlined: normal form tableaux may be much bulkier than non-normal form equivalents, as in the following pair, where 136 is not in a normal form, and the equivalent 137 is:

(136)

$C1$	$C2$	$C3$	$C4$	$C5$
W	e	L	L	L
e	W	e	L	e

(137)

$C1$	$C2$	$C3$	$C4$	$C5$
W	e	L	e	e
W	e	e	L	e
W	e	e	e	L
e	W	e	L	e

A.2.2 Transforming an arbitrary tableau into its normal form

The next step of the proof is to show that we can in fact build a normal form tableau from an arbitrary tableau using only our five equivalence-preserving operations and their inverses. Note that if a single sequence of transformations leading to the normal form exists for a tableau, an infinite number of such sequences exist, as we can do and undo one operation any number of times. More meaningfully, there may very well be different routes from the tableau to its normal form than the sequence we have found even without doing and undoing the same thing. What suffices for us is to show that for each tableau, there exists at least one such finite sequence of transformations leading to a tableau in the normal form.

(138) **Normal form existence theorem**.
An arbitrary (finite) tableau T can be transformed into an equivalent normal form tableau by a (finite) sequence of equivalence-preserving transformations including

row swaps, row splittings, inference eliminations, false W eliminations, and contradictory jumps and backward contradictory jumps.

Proof. We will prove 138 by explicitly providing a procedure performing the transformation.[38]

First we check if the tableau is contradictory. If it is, we add a special unhappy row (L, e, e,...) to it, and subtract all other rows.

At the second step, we apply row splittings so that each row in the tableau has at most a single L.

The hard work starts at the third step. We process rows one by one, eliminating them if they are entailed by the rest of the tableau, and deleting all false W-s if there are any. For a row r in tableau T, we do it as follows. We check if r is entailed by $T \setminus r$. If it is, we delete it, and proceed to the next row. If it is not, we build for every W in r a special row q to check if that W is false or not. For those which are, we apply **GWR**.

It is important to consider what it means in terms of rankings compatible with T that r cannot be eliminated, and its W-s are not false in T. Since r is not entailed by $T \setminus r$, there are rankings ϕ which are compatible with $T \setminus r$, but not with r, and have to be refined in order to become compatible with T as a whole. If Ci is the L-constraint of r, it means that there are rankings which for every $Cj \in W(r)$ do not include $Cj \gg Ci$. Let us fix some $Ck \in W(r)$. Since the W in Ck is not false, $T \setminus r$ does not entail the row q defined by $L(q) := \{Ck\}$, $W(q) = W(r) \cup \{Ci\} \setminus \{Ck\}$. This means that there are rankings compatible with $T \setminus r$ which for every $Cl \in W(q)$, do not include $Cl \gg Ck$. Finally, the most important fact is that there must exist some ranking ϕ compatible with $T \setminus r$, but not compatible with either r or q.

Suppose it is not so, and every ϕ compatible with $T \setminus r$ and not entailing r entails q. This means that if ϕ compatible with $T \setminus r$ does not include any $Cj \gg Ci$, it has to include at least one $Cl \gg Ck$ atomic ranking. Since $Ci \neq Ck$, this cannot be forced by some single row in $T \setminus r$. So there must be a transitivity chain: some row r_1 should have an L in Ci which needs to be accounted for; and when it is not accounted for by $Cj \gg Ci$ where Cj comes from $W(r)$, the atomic ranking $Cm \gg Ci$ needed to account for r_1 should always trigger a transitive domination chain which leads to $Cl \gg Ck$, of the form $Cl \gg ... \gg Cm \gg Ci \gg ... \gg Ck$. But for such a chain to exist, it had to be that $Ci \gg Ck$ is forced by $T \setminus r$ already. And since $Ci \in W(q)$, $Ck \in L(q)$, that means $T \setminus r$ entails q, contrary to the assumption.

We have just shown that if every ranking compatible with $T \setminus r$, but not with r, entails q, then every ranking whatsoever compatible with $T \setminus r$ entails q. Thus if r is not entailed by $T \setminus r$, and the W in Ck is not false in r, there exists a ranking ϕ compatible with $T \setminus r$

[38] Exactly as was the case with our definition of truth for our logic of partial rankings in Section 4.1, we do not guarantee that the way to find such a sequence that we will now present is the most efficient one. But our purpose here is not to build a fast algorithm, but to prove that a correct algorithm exists. Once that is done, one may think of optimizations.

which includes neither $Ck \gg Ci$ nor $Ci \gg Ck$. Thus adding r to $T \setminus r$ changes the set of rankings that the tableau as a whole is compatible with — it rules out ϕ.

OK, now we know how to treat a single row r. But can we guarantee that changing r is not going to affect what we do with the other rows in T? First, consider counterfeeding row elimination. Eliminating a superfluous r or deleting a false W preserve equivalence in T as a whole. So if some s was superfluous in T, it will still be so in a modified T with a modified r.

What about false W elimination? Suppose s has a W in Ci which is false in T, and we modify some other r in T, obtaining T'. Towards a contradiction, assume that W is not false in T'. That means that $T \setminus s$ entails the q built for Ci from s, but $T' \setminus s$ does not. This can only be if every ϕ compatible with $T \setminus s$ contained $Cj \gg Ci$ for some $Cj \in W(s) \cup L(s) \setminus \{Ci\}$, but there exists a ψ compatible with $T' \setminus s$ which does not. If T' was a result of deleting a false W in some r, it cannot be: deleting a W from a tableau cannot enlarge the set of rankings compatible with it. If T' was a result of deleting the whole r because it was superfluous in T, that can in principle make $T' \setminus s$ compatible with more rankings than $T \setminus s$. Suppose it is so. Then it must have been r which forced $Cj \gg Ci$ for some $Cj \in W(s) \cup L(s) \setminus \{Ci\}$ into any ranking compatible with $T \setminus s$. But as r itself was superfluous in T, every ranking compatible with $T \setminus r = T' \setminus r$ must be forcing atomic rankings accounting for r. Therefore $T \setminus r$ will force $Cj \gg Ci$ for some $Cj \in W(s) \cup L(s) \setminus \{Ci\}$ as well. We only need to show that this cannot be because of s itself. Suppose that $T' \setminus s$ does not force any $Cj \gg Ci$ for $Cj \in W(s) \cup L(s) \setminus \{Ci\}$, but T' does. Fix as Ck the L-constraint of s. We get that for any $Cl \in W(s)$, $Cl \gg Ck$ forces $Cj \gg Ci$. This can only be via a chain $Cj \gg ... \gg Cl \gg Ck \gg ... \gg Ci$, with $Ck \neq Ci$. But if $Cl \gg Ck$ can create such a chain, it means that $T' \setminus s$ forces $Ck \gg Ci$, and thus $T' \setminus s$ entails s. Thus elimination of another row can only counterfeed false W elimination in s if s itself is superfluous.

To sum up, the procedure for a single r does not affect what we can do with the other rows in the tableau. Therefore it suffices to make one pass through the tableau in the third step checking if inference elimination or false W elimination can apply, and the result does not depend on the order in which we pass through the tableau.

To finish the normal form transformation, it only remains to sort the remaining rows by row swaps. $\dashv$.

A.2.3 Each equivalence class has exactly one normal form tableau

We now turn to the next part of the big proof: we need to show that there is only one normal form tableau per equivalence class. In other words, we want to prove that any normal form tableau *defines* its equivalence class, and may be thought of as the unique name for it. To show that, we need to prove any two distinct tableaux in the normal form are not equivalent.

(139) **Normal form uniqueness theorem**.

In each equivalence class of OT tableaux, there is at most one normal form tableau.

Proof of 139. The core of the proof is to show that any two normal form tableaux belong to different equivalence classes. The regularity of the normal form will make the proof easy.

Take two arbitrary distinct normal form tableaux $T1$ and $T2$. We will now build a ranking which is compatible with one of them, but not the other.

Without loss of generality, pick a row r of $T1$ which is not present in $T2$ (since $T1 \neq T2$, there will be such a row.) Since $T1$ is in the normal form, r cannot be entailed by $T1 \setminus r$. If r is not entailed by $T2$, there is a ranking compatible with $T2$ witnessing it. That ranking is not compatible with r, and hence with $T1$, so $T1$ and $T2$ are in different equivalence classes.

The interesting case is when r is entailed by $T2$. We fix a minimal subset $M = \{r_1, ..., r_n\}$ of $T2$ entailing r. It cannot be that $T1 \setminus r$ entails M, for that would have made r superfluous in $T1$. Now we want to show that there is a ranking compatible with T, but not with M.

Suppose there is no such ranking. Then for any ranking ϕ^- compatible with $T1 \setminus r$, refining it to ϕ compatible with r results in ϕ being compatible with M as well. We fix an arbitrary $r_k \in M$ such that $T1 \setminus r$ does not entail r_k. Let Ci be the L-constraint of r. Then any ϕ as just described should include for some $Cj \in W(r)$ the atomic ranking $Cj \gg Ci$. Since adding $Cj \gg Ci$ makes the resulting ϕ compatible with r_k, it must be that $Cj \gg Ci$ creates a chain of the form $Ck \gg ... \gg Cj \gg Ci \gg ... \gg Cl$ for $Ck \in W(r_k)$, $Cl \in L(r_k)$. But since M entails r, it forces any ranking ψ compatible with M to have for some $Cj \in W(r)$ the atomic ranking $Cj \gg Ci$, and as M is minimal, r_k should play a crucial role in it, creating a chain $Cj \gg ... \gg Ck \gg Cl \gg ... \gg Ci$.

Any ranking compatible with M has a chain $Cj \gg ... \gg Ck \gg Cl \gg ... \gg Ci$, yet some of them, namely ϕ's which are also compatible with T, should have a chain $Ck \gg ... \gg Cj \gg Ci \gg ... \gg Cl$. If $Cj \neq Ck$ or $Cl \neq Ci$, this is a contradiction, so the assumption must have been wrong, and there is a ranking compatible with $T1$, but not with M. This ranking then is not compatible with the whole $T2$ as well, and witnesses that $T1$ and $T2$ are in different OT equivalence classes.

On the other hand, in the case when both $Cj = Ck$ and $Cl = Ci$, row r_k entails r alone, and as $r_k \neq r$ (row r was not present in $T2$), r should have some W which r_k does not have. As this W cannot be false because of the normal form condition on $T1$, it is possible to extend some ranking compatible with $T1 \setminus r$ to account for r as well by using that W to cover the L, but that will not make the ranking compatible with r_k.

Thus if r is entailed by some $M \subseteq T2$, there is always a ranking compatible with $T1$, but not with $T2$.

To sum up, in either case the sets of rankings that $T1$ and $T2$ are compatible with differ. Thus any normal form tableau defines its equivalence class. ⊣

The theorem 139 establishes that each equivalence class has at most one normal form tableau. The theorem 138 guarantees that for any tableau, its equivalence class contains one normal form tableau. Thus there is exactly one normal form tableau per equivalence class.

This is the main result of this Appendix: it allows us to work with normal form tableaux as proper representatives of equivalence classes of tableaux in the main part of the paper.

A.3 Capitalizing on the normal form results

There is a number of nice corollaries from the normal form results.

For instance, we can easily define an effective procedure for transforming an arbitrary tableau $T1$ into another arbitrary tableau $T2$ equivalent to $T1$: we just transform $T1$ and $T2$ into the normal form, which is the same for both, as they are from the same equivalence class. Since all transformations we use for that have inverses, we can invert the sequence of transformations leading from $T2$ to the normal form. Composing the inverted sequence with the sequence transforming $T1$ into the normal form, we obtain a sequence of equivalence-preserving operations transforming $T1$ into $T2$.

Another nice consequence is the decidability of equivalence class membership: to find out which equivalence class some T belongs to, we just build T's normal form, which defines its equivalence class.

With a little bit of extra work, we can also formulate a procedure enumerating all the tableaux in the same equivalence class, if need be. That is hardly a particularly useful result for a practical OT analyst who works only with finite tableaux — in the finite, enumeration is trivial. But if we want to work with tableaux of denumerable size, an enumeration procedure smartly interleaving outputting different members will allow us to at least generate bigger and bigger portions of more and more members of the equivalence class.

Since our toolkit of equivalence-preserving operations is complete, all equivalence-preserving transformations proposed in the literature can be analyzed as specific cases in the general framework developed. For some algorithms like Fusional Reduction of [Brasoveanu and Prince, 2005] it has been shown that they preserve OT equivalency, but since the proof was done in an idiosyncratic way, it was not guaranteed that the same proof techniques would work for other tableau transformation algorithms. With the development of our full toolkit, to show a new algorithm's correctness it suffices to demonstrate that the algorithm affects the input tableau as a sequence of transformations from our toolkit does. Of course, whether such a proof will be more straightforward than a proof from first principles is for those who will attempt to do that to tell, but at least the current framework allows for such a general strategy.

Perhaps more importantly, the full toolkit we developed may suggest new equivalence-friendly algorithms useful from the practical point of view. The groundbreaking research reported in [Brasoveanu and Prince, 2005] had to address two problems simultaneously: what would be a good, convenient form for an arbitrary tableau, and how to get to that form. Their techniques of building different kinds of such forms (bases, in their terms) have heavily influenced the work reported in this Appendix. But the generalization we achieved here makes the task easier for those who follow: the basic toolkit for transformations is provided, all that is needed is to figure out what sort of a tableau one wants to get in the end.

Finally, yet another job which our toolkit makes easier is proving negative results. In general, it is much easier to prove a positive result than a negative one. For the former, one simply has to demonstrate a specific way to achieve the goal. For the latter, showing that one particular way does not allow us to reach the goal is not enough: that does not guarantee that some other way could not. But if there is a small set of operations which contains all we have at our disposal, and we can show that no combination of those does the trick, we thus show that the goal cannot be reached. This is exactly what we can do with our compete toolkit of operations preserving OT-equivalency.