No More than Necessary

Beyond the 'Four Rules', and a Bug Report

Alan Prince
Rutgers University
November 16, 2006. Updated 5/2/07. Typos Corrected 2/29/08

Abstract. After proposing four 'rules of inference' to be used in the program OTSoft for simplifying collections of ranking arguments, Hayes 1997 implicitly raises the question of whether these rules suffice. In this note, the simplification goal is spelled out within the analytical framework of Prince 2002a and Brasoveanu & Prince 2005, in prep. and the question is settled (negatively). A broader generalization subsuming two of Hayes's rules is offered (ex. (18), p. 9), and it is shown to provide a complete solution to the simplification problem as formulated. A tighter characterization of the role of disjunctive ranking relations then follows.

Having rules of inference in hand is not the same as having a procedure that uses them effectively. As of this writing, the only known algorithm that produces a fully simplified set of ranking conditions guaranteed to be both individually necessary and jointly sufficient is Fusional Reduction (FRed), presented in Brasoveanu & Prince (op. cit.), which relies on different inferential assumptions.

Consequences for the use of OTSoft are noted. The typological calculations and the stratified hierarchies produced in response to a ranking request are unchallenged. Those subparts of the program which deal with the *necessity* of ranking conditions (and concomitantly, necessary/non-necessary presence of constraints) must however be used with circumspection and should be supplemented with other methods. Some bugs in this part of the program are reported.

Acknowledgments. Thanks to the Rutgers Optimality Research Group for discussion and to Jane Grimshaw and Bruce Tesar for additional helpful comments.

0. Introduction

OTSoft, the invaluable software package developed by Bruce Hayes and his collaborators, accomplishes several distinct tasks (Hayes, Tesar & Zuraw 2003).[1] Among these are the generation a full factorial typology based on violation data supplied by the user. It can also produce a ranking, should one exist, that derives a single language, given a selection of desired optima. The ranking is determined by one of several algorithms: most neutrally, Recursive Constraint Demotion (RCD: Tesar & Smolensky 1993, Samek-Lodovici & Prince 1999, Prince 2002a, Brasoveanu & Prince 2005). RCD is a greedy algorithm that produces a stratified hierarchy in which each constraint is placed in the highest stratum (collection of nonconflicting constraints) that it can possibly reside in. A linearly-ordered ranking that generates the language can be derived from a stratified hierarchy by choosing linearizations (any will do) for each of the strata and hooking them together in a way that respects the stratal order.

[1] For a full description, see the detailed manual *OTSoft: Constraint Ranking Software: Version 2.1*, available at Hayes's UCLA website along with the program itself.

In addition to these utilities, OTSoft also provides a description of the ranking conditions *necessary* for each language.[2] This is a quite different task from supplying an RCD-based ranking which is sufficient to generate the language. The reason for this discrepancy is that, typically, many rankings (linear orderings on the constraint set) will generate the same language, and in general they need not be reducible to any one stratified hierarchy. To see this, consider the following example, which illustrates the commonly-encountered situation. Suppose we have a grammar of three constraints and a single candidate set of two members. This gives us precisely one optimum-suboptimum comparison to worry about: suppose it looks like this, where constraints are named by bold-face numerals; no prior ranking order is presupposed.

(1) **Typical disjunctive ranking condition**

Input	Winner	Loser	**1**	**2**	**3**
/u/	q	z	**W**	**W**	**L**

An Elementary Ranking Condition (ERC) is associated with each pairwise comparison (over the entire constraint set) of a desired optimum with a competitor. The ERC on display here declares : "**1** or **2** dominates **3**" and in any grammar in which ⟨u→q⟩ is better than ⟨u→z⟩, this condition must hold. Greedy RCD imposes the following stratification:

(2) **RCD stratification** {**1**,**2**}||{**3**}

Input	Winner	Loser	**1**	**2**	**3**
/u/	q	z	**W**	**W**	**L**

By stratal linearization, this admits any total orders in which **1** and **2** precede **3**, namely ***123*, *213***. But the original condition (1) — what the data says — also allows for **13*2*** and **23*1***. RCD's restriction to the first two linear rankings is an artifact of greediness, but it does absolutely no harm to the typological enterprise, which asks for the languages that are generated and does not need to know if they are generated by more than one ranking, or, if they are, how the collection of successful rankings is internally structured. Whenever you are asking for a ranking that works, RCD will give you (at least) one, if one exists, and it will let you know if there is none. The cost is loss of information about the details of necessity: which rankings the data absolutely requires.

To see how information is lost, consider the following system with two ERCs:

(3) **A different language**

Input	Winner	Loser	**4**	**5**	**6**
/u_1/	q_1	z_1	**W**	e	**L**
/u_2/	q_2	z_2	e	**W**	**L**

RCD gives the stratification {**4**,**5**} || {**6**}, which is entirely parallel to that derived in the first example. But in this case, the data tells us that each top-stratified constraint *must* dominate the bottom-ranked constraint.

In many contexts, we may not be interested in this kind of detail, and the cost of going greedy (particularly coupled with the alarming efficiency of RCD, which means that it can often be done by hand or by cut-and-paste) will be well worth paying. Nevertheless, there are situations where we must understand both necessity and sufficiency. Here's one example. Implications between one optimum and another — if *p* occurs in a language, then so must *q* —

[2] 'Language' is here understood to mean the collection of desired optima from the user-supplied candidates.

occur because the *necessary* conditions for the optimality of candidate *p* are *sufficient* for the optimality of *q* (Prince & Smolensky 1993/2004; Prince 2006). Here's another example: the disjunctive situation in (1) indicates that the degree of failure on constraint **3** can be explained as the result of *either* of the two other constraints: since they overlap in explanatory force, perhaps their content relationship can be fruitfully re-examined. But in a conjunctive situation like the one in ex. (3), each constraint is doing irreplaceable work. The active analyst will want to be aware of such things.

How do we know whether a given ranking condition is necessary? The answer must be dug out from the welter of disjunctive conditions that data usually provides. In 'Four rules of inference for ranking argumentation' (Hayes 1997), the principles that OTSoft uses are laid out and justified. Advances in our understanding of ranking make it possible to clarify the goals of the enterprise and answer (negatively) the question he implicitly poses about the sufficiency of the four cited principles for accomplishing those goals. The shortfall is not limited to clever constructions, but shows up in familiar cases. More positively, by finding and generalizing the essence of two of the proposed inference rules, we will uncover a single principle that includes the missed case as well, and provably achieves sufficiency.

Having sufficient principles on hand to solve a problem is by no means equivalent to having an algorithm that uses them effectively. The Fusional Reduction Algorithm (FRed) of Brasoveanu & Prince 2005 uses a rather different approach but is guaranteed to produce a maximally concise and informative representation of the individually necessary and jointly sufficient ranking conditions for any set of data. The analytical framework in which it is couched also provides the tools for demonstrating and refuting claims about necessity, and they will be used here.

A further, implementational question arises — are the principles used to maximal effect in the current state of the program? We will display a couple of examples which show that the program's results in matters of necessity should be double-checked, even within the domain where the four rules are efficacious. We also find some apparent flaws in the implementation.

The goal here is dual: to advance toward better understanding of ranking conditions, and to distinguish between the sound (and therefore comfortably usable) and the less sound (and therefore deserving of circumspection) in the output of what has proved to be an essential tool for the conduct of linguistic investigations.

1. Reducing Arguments

Given a set of ranking arguments derived from data, we may ask what they amount to. One way of answering is the provide a reduced set of arguments, equivalent to the original, which displays the information in nonredundant form. Brasoveanu and Prince (2005, in prep.: henceforth RN-I, II) have provided the basic framework of analysis, set within fusional ERC theory (Prince 2002ab, Prince 2006).

Assume that the ranking arguments are given in the form of Elementary Ranking Conditions (Prince 2002a). These are statements of the form 'at least one constraint of type 1 must dominate all constraints of type 2'. They are uniquely representable as lists (or 'vectors') of the comparative values W,L,*e*, where each constraint is assigned a fixed (if arbitrary) position in the list. The type 1 constraints are those whose position in the list is filled with W; together they form the 'W-set' of the list. The type 2 constraints are those marked by L, comprising its 'L-set'. The Elementary Ranking Condition (ERC) arises in the analysis of data from the comparison of a desired optimum with a single competitor. The W-constraints are those that favor the desired optimum, in accord with the wished-for outcome; the L-constraints perversely favor the competitor, which is globally desired to lose but may not be able to win without proper ranking; and the *e*-constraints are neutral between them. As a logical form — "some W dominates all L's" — the ERC stands on its own. We can manipulate and combine data-derived ERCs to produce yet further ERCs that express conditions on ranking, even though they may not directly represent the head-to-head contest between two pieces of data. ERCs can also be used to represent conditions of the type familiar from universally fixed hierarchies, or indeed any other condition one wishes to impose.

Suppose we have a set A of ERCs, from whatever sources. Our goal now becomes: find a potentially more useful set G of ERCs, such that every ERC in A is entailed by G, and every ERC in G is entailed by A. This means that in every hierarchy meeting the conditions of G, the conditions of A will also be satisfied. The set G is thus a perfectly valid representation of A: any consequence of A is also a consequence of G, and vice versa.

Let us call any such G a 'generating set' for A. (Such a thing must exist: A is a generating set for itself.) We are interested in finding generating sets that have desirable properties of conciseness and informativeness. Brasoveanu & Prince term a generating set of minimal cardinality a 'basis', and distinguish among various (equi-cardinal) bases in terms of their W/L/e structure. To appreciate the kind of distinctions that arise, consider the following example, which covers all the bases for the simple linear order on three constraints (RN-I, p.9)

(4) **Equivalent ERC Sets: Bases for 1□2□3**

B_1	**1**	**2**	**3**
α	**W**	**L**	**L**
β	e	**W**	**L**

B_2	**1**	**2**	**3**
α′	**W**	**L**	e
β	e	**W**	**L**

B_3	**1**	**2**	**3**
α″	**W**	**L**	**W**
β	e	**W**	**L**

The first of these is L-rich. The L-set in each ERC gives the complete list of (subordinated) constraints that the (possible dominator) constraints in the W-set are related to, either directly or by transitivity. For this reason, it is called the 'Most Informative Basis' (MIB). Basis B_2 is *e*-rich; it contains *none* of the consequences of transitivity; it is the 'Skeletal Basis' and is also quite informative (even in some ways arguably more so than the MIB) since it spells out exactly those constraints with which the W-set is involved in unmediated domination relations. Basis B_3 is W-rich: it is the 'Least Informative Basis', since it spells out the maximum number of local disjunctions that can be packed into the ERCs of a basis. In the general case, there will also be many other possible mixtures of W,L,*e*-content in the bases for a given target set of ranking conditions; these three are distinguished by their maximality along a certain dimension.

Observe also that the ranking at hand is also described by many other generating sets that do not qualify as bases. For example, if we form the union of all the ERCs cited in (4), we obviously still the generate the same ranking while adding no further consequences. The *basis* is distinguished by its conciseness.

OTSoft works with a kind of ranking condition in which only one constraint is subordinated, a subtype of ERC in which only one L appears. Let us call this a 'Primitive Ranking Condition' or PRC (*cf.* Prince 2002a:16). Specifically, let's require that a PRC contains *exactly* one L and at least one W.[3] To characterize the kind of PRC sets that OTSoft is implicitly looking for, we need

[3] This forbids ERCs lacking L or lacking W from being PRCs. W-free or L-free or all-*e* ERCs are 'trivial' in the sense that they either cannot be satisfied or are always satisfied. Since calculations with PRCs quickly lead outside

to find a parallel among PRC sets for the 'basis' idea that has been stated in terms of (general) ERC sets, whose members may include any number of L's.

Given a set of ranking conditions A (which we construe without loss of generality as being given in ERCs), we can define a Primitive Generating Set or Primitive Generator (PG) as a generating set for A that consists entirely of PRCs. We are of course interested in those PGs of minimal cardinality. As with the ERC-based basis, this minimality restriction immediately gives us a desirable property: logical independence. Any PG that happens to contain a PRC entailed by others in the PG can be reduced in size by simple elimination of the entailed PRC, producing a smaller PG without loss of any information.

Within the realm of general ERCs, as noted, the internal logical expressiveness of the ERC typically gives rise to more than one minimal-cardinality generating set. The same is true in the PRC world, because the PRC only restricts L-content, often leaving many equivalent options among the W's. Just as with bases, there can easily be more than one PG for a given set of ranking conditions. Consider this case:

(5) **A PRC set**. P = {π, ρ}
π = (W, W, W, L)
ρ = (e, e, L, W)

Observe that this set is 'logically independent': neither PRC entails the other, and it cannot be simplified by merely ejecting one of the PRCs. But it is overly disjunctive, locally, as may be seen by noting that it is equivalent to this set:

(6) **A Less Disjunctive Version**. P′ = {π′, ρ}
π′ = (W, W, e, L)
ρ = (e, e, L, W)

P′ is in fact the Skeletal Basis. Back in the original target set P, the PRC π declares that constraint **3** (among others) is a possible dominator for constraint **4**. But we may observe that ρ *requires* to the contrary that **4** dominate **3**. Therefore, when the whole system is considered, it becomes apparent that **3** is *not* a possible dominator for **4**. We can eliminate the subclause of π that declares it to be so, giving us π′.

This result may be derived in the following way using the Fusional Reduction Algorithm of RN-I.[4]

- π∘ρ = (W, W, L, L)
- π∘ρ is not entailed by ρ, and so is therefore part of the MIB
- If we erase from π∘ρ all L's shared by ρ and π∘ρ, replacing them with *e*, we obtain an element of the Skeletal Basis
- ρ is itself not entailed, so it must be retained.
- The Skeletal Basis is therefore {π′, ρ}. In this case, it consists of PRCs and we're done.

the set of PRCs into the greater ERC world that surrounds them, where 'triviality' can be expressed and detected, there appears to be little reason not to adopt the narrower characterization of the PRC.

[4] The operation of fusion, shown as α∘β (and *ƒ*X for a set X) combines ERC vectors coordinate-by-coordinate, according to the scheme X∘L=L∘X=L, e∘X=X∘e=X, W∘W=W. Any fusion is entailed by the *set* of its component fusands. In some cases, like this one, the fusion itself has desirable entailment properties. Any ERC entailed by a set of ERCs A is entailed by the fusion of some subset of A (Prince 2002a:14, Prop. 2.5). For introductory discussion, see Prince 2002b, Brasoveanu & Prince 2005, Prince 2006. For detailed analysis, see Prince 2002a.

The kind of PRC set we are seeking, then, not only has logical independence, but is also maximally free of local disjunction, i.e. has the minimal number of W's (equivalently construed as locally in each PRC and globally for the whole set, since we have fixed cardinality for the set).

Given the limitation to one L that defines the PRC, the minimal-W condition is alone sufficient for the notion we want: in a minimal-W PG, there can be no logical dependencies among the members. (Removing a logically dependent PRC from the set will decrease the overall W-count.) We therefore denominate the object of our search the Minimal Primitive Generator (MPG) and define it as follows:

(7) **Minimal Primitive Generator (MPG) for A**. Let A be a set of ranking conditions. Let M be a set of PRCs, and let M be a generating set for A. M is the MPG for A iff there is no other Primitive Generating Set for A with fewer W's.

The MPG is unique because the Skeletal Basis is unique (RN-II), and an MPG is simply a Skeletal Basis that has been 'exploded' so that any ERC with more than one L is replaced by a set of PRCs, each with the same W-set the original ERC, and each presenting one L of the original. For example, (W,L,L) is not an PRC, but it could be a member of a Skeletal Basis. It explodes to the equivalent PRC set { (W,L,e), (W,e,L)}.

In sum: with respect to the display of ranking arguments, we characterize the goal of OTSoft as being the production of the MPG for the ranking conditions implicit in some user-defined set of data. This is not quite all of what OtSoft does, because it displays arguments derived from transitivity as well, but since it marks those off, we can legitimately focus on the MPG-creating aspects of the program.

2. Four Rules of Inference

OTSoft uses four rules of inference to reduce PRCs derived from data. The last of these, which will be first here, is a basic defining property of partial orders like strict domination and has nothing to do with PRCs *per se*: transitivity.

(8) **Transitivity of Strict Domination.** For constraints A,B,C, $(A \gg B) \& (B \gg C) \Rightarrow (A \gg C)$.

To see its application, consider a case like this:

π = (W, L, e) **1≫2**
ρ = (e, W, L) **2≫3**
σ = (W, e, L) **1≫3**

This is a PG for **1≫2≫3**, but it's not minimal because the third PRC σ is entailed. Transitivity is a fundamental property, but it can relate constraints via an unlimited number of intermediary links, and thus greatly enriches the problem of reducing ERC or PRC sets.

The other three rules focus on relations between PRCs. One deals with the entailment relation between two individual PRCs. Hayes puts it this way (FRI:3, ex. (4)).

(9) **Superset Redundancy**

If, in two ranking arguments, the dominees are the same, and the dominator set of one is a superset of the dominator set of the other, discard the ranking argument that includes the superset.

The force of this is easy to grasp. The 'dominator set' of a PRC (its 'W-set', in our nomenclature) is involved in a disjunctive relationship with the 'dominee' (its 'L-set', which by definition contains only one element). As an example of the superset situation, consider two PRCs structured along these lines:

π = (W,W, e ,L) **1**≫**4** or **2**≫**4**

ρ = (W,W,W,L) **1**≫**4** or **2**≫**4** or **3**≫**4**

The first clearly entails the second: all the rankings in which **1** *or* **2** dominates **4** are also rankings of which it is true to say, more weakly, that **1**or **2** ***or 3*** dominates **4**. As the propositional calculus teaches (and we see no reason to object): $P \Rightarrow P \vee Q$.

In ERC theory, this effect is known as W-extension. If you take an ERC and replace any coordinate with W, the ERC thusly derived will be entailed by the original. We can therefore re-state 'Superset Redundancy' in terms of its logical core, as follows. We write $W(\alpha)$ for the W-set of ERC α, the set of coordinates (constraints) occupied by W in the ERC.

(10) **PRC Entailment**. For PRCs π and ρ, $\pi \Rightarrow \rho$ iff $W(\pi) \subseteq W(\rho)$.

The 'iff' means that the relation of W-set subsetting ('W-extension') is completely co-extensive with entailment among nontrivial PRCs; this follows from Prince 2002a, p.6, Prop. 1.1.

The 'discard' part of Superset Redundancy follows because the MPG is our ultimate goal. No MPG will contain entailed PRCs.

The third and fourth rules involve more elaborate relations in which two PRCs together entail a third that is less W-ful than one of the original pair, which may be replaced without loss of information. Recall our original example (the eliminable W is bolded):

(11) **A PRC set.** P = {π, ρ}

π = (W, W, **W**, L)

ρ = (e, e, L, W)

(12) **MPG(P)** = {π', ρ}

π' = (W, W, **e**, L)

ρ = (e, e , L, W)

The guiding fact is that π' is entailed by {π,ρ} and, at the same time, π is entailed by {π',ρ}. Ergo, these sets are equivalent, and, seeking fewer W's, we go with the second.

Hayes formulates the relevant condition as follows:

(13) **Reverse-Ranking Elimination**

Eliminate from the dominator set of a set argument any constraint C for which a pairwise ranking argument exists which subordinates C to the dominee of the set argument.

To rephrase this condition, let us use the notation $L(\alpha)$ to refer to the L-set of any ERC α, which consists of a single constraint among the PRCs. We are concerned with a PRC π and with a constraint C which is necessarily subordinated to $L(\pi)$, i.e. a constraint C for which $L(\pi) \gg C$, according to our target PRC set P. We construct a new PRC π' by retaining $L(\pi)$ and replacing $W(\pi)$ with $W(\pi)-\{C\}$. The conclusion is that we may replace π with π' in P.

It is important to note that the PRC stating this domination relation need not literally be resident in the PRC set under consideration . It need merely be *entailed* by it. Consider the following more articulated case, which expands on the one just reviewed:

(14) **A PRC set.** Q = {κ ,λ ,μ }

κ = (W, W, **W**, **L**, e)
λ = (e, e , L, e , W)
μ = (e, e, e, W, L)

The set Q does not contain any PRC declaring that **4**≫**3**, i.e. that $L(\kappa) \gg L(\lambda)$. It's a fact nonetheless, and motivates Reverse Ranking Elimination. This is a hint that applying the rules of inference effectively will require some labor, as Hayes himself notes (FRI:5, §4).[5]

The fourth and final rule of inference runs along similar lines, but it is concerned only with relations internal to the W-set. If, within the W-set of some PRC α, we know that some C is necessarily subordinated to some other W-set member, then that C is telling us nothing that the other constraints don't already tell us. Here's an example, modified from the first one:

(15) **A PRC set.** P= {π, ρ}

π = (W, W, **W**, L)
ρ = (e, W, L, e)

(16) **MPG(P)** = {π′, ρ}

π′ = (W, W, **e**, L)
ρ = (e, W, L, e)

The W at **3** in π (bolded) is uninformative. Since, by ρ, **2**≫**3**, it follows that in any hierarchy in which **3**≫**4**, it will also be the case that **2**≫**4.** Therefore **2** will always dominate when **3** does. So it is quite enough to remark, as in π′, that **1** or **2** must dominate **4**. This effect is given the name Part-for-Whole Selection (FRI:4, ex. (9)) and the rule of inference is stated as follows:

(17) **Part-for-Whole Selection**

Let S be a set of constraints, s be a proper subset of S, C a constraint included in S but not s, and S≫D be an ordinary set argument ("at least one constraint in S must dominate D"). If s >> C is also a ranking argument, then replace S≫D with the simpler ranking argument S′ ≫D, where S′ is the set obtained by removing C from S.

Paraphrasing in roughly the same style as before, we note that we are interested here in a constraint C in the W-set which is necessarily subordinated to some other element of the W-set. Suppose we have a π in as PRC set P in which this condition holds. Then there's some subset of $W(\pi)$, call it $V(\pi) \subseteq W(\pi)$, such that $V(\pi)$ is the W-set of another PRC ω in which C = $L(\omega)$. The PRC ω literally says: some constraint of $V(\pi)$ dominates C. Assume that the original PRC set P entails ω. In this case, we may replace $W(\pi)$ by $W(\alpha)-\{C\}$, retaining $L(\pi)$, to obtain another PRC π′ which will have fewer W's than π, while doing the same work in P.

[5] Hayes's attitude toward the project differs from ours, however, and it is worth reading his original pages.

Stepping back, we see that both Reverse-Ranking Elimination and Part-for-Whole Selection are really instances of a single more general principle that relates to the entire collection of constraints under discussion, regardless of the W-set/L-set distinction.

Any C that carries W in π can be removed from π's W-set if (and, as we'll see, only if) it is necessarily subordinate to a constraint mentioned as *either* W or L in π. This result is of sufficient importance to deserve a name of its own: Generalized Removal of W.

More explicitly: instead of considering $W(\pi)$ and $L(\pi)$ separately, as in Hayes's two rules of inference, let us consider $WL(\pi) = W(\pi) \cup L(\pi)$, the entire set of constraints that are 'polar' in π, i.e. that assess either W or L (but not *e*). In these terms, the truly basic principle of W-elimination is this: if C is necessarily subordinated to *any* constraint in $WL(\pi)$, then we can substitute for π a new PRC in which C bears *e*, i.e. a PRC π′ obtained by replacing $W(\pi)$ in π with $W(\pi)-\{C\}$.

(18) **Generalized Removal of W** (qualitatively put).

> For any Primitive Ranking Condition π in a set P of such conditions, a W may be removed from π — converted to *e* —, without affecting the force of P, if (and only if) the constraint assessing that W is necessarily subordinated to some polar constraint in π.

This verbal formulation is meant to include the disjunctive case, in which we only know that for some $X \subseteq WL(\pi)$, C must be subordinate to at least one of the constraints in X.

In fact, there's no need to restrict to a particular *sub*set of the polar constraints. All we need to license removal is the loosest possible requirement: that $WL(\pi)-\{C\}$ is the W-set for a PRC ω with $C=L(\omega)$, and that the original set P entails ω.

(19) **Generalized Removal of W**. (GRW)

> Let P be a set of PRCs. Let $\pi \in P$ and $C \in W(\pi)$. Let ω be a PRC with $WL(\omega)-\{C\}$ as its W-set and $C=L(\omega)$.
>
> Suppose P entails ω. Then P is equivalent to the PRC set P′, obtained by replacing π with π′, where $L(\pi')=L(\pi)$, and $W(\pi')=W(\pi)-\{C\}$.
>
> Conversely, if π is replaceable in P by such an p′, then $P \vDash \omega$.

It is important to note (and so we note it once again) that the notion of 'necessarily dominated' does not mean that there is already sitting in P a PRC which explicitly announces this relation. The overall consequences of a PRC set will typically far exceed its explicit contents. What's required is that the necessary domination relation be required by P; that is, logically entailed by it. In order to use this rule of inference, we must deal with *all* the consequences of P.

To see how the generalized criterion works out, let's apply it to the examples already discussed.

(20) **The PRC set from example (11).** P = {π, ρ}

π = (W, W, **W**, L)
ρ = (e, e, L, W)

To convince ourselves that bolded W, in constraint **3** of π, is indeed removable, we need simply note that the relevant ω is (W,W,L,W) and it is clearly entailed by P (via ρ). This ω is obtained by turning the targeted **3** into L, and carrying over the rest of $WL(\pi)$ as W.

Let's continue with the more elaborate PRC set from example (14).

(21) **A PRC set.** Q = {κ ,λ ,μ }

κ = (W, W, **W**, L, e)
λ = (e, e , L, e , W)
μ = (e, e, e, W, L)

Here we're targeting **3** in κ, and the relevant ω is (W,W,**L**,W,e). A quick calculation shows that it is entailed by λ∘μ = (e,e,L,W,L) and therefore by P.

Let's now move to the case of 'Part for Whole Selection', as in example (15)

(22) **A PRC set.** P= {π, ρ}

π = (W, W, **W**, L)
ρ = (e, W, L, e)

Here we target **3** in π and note that the relevant ω = (W,W,LW), which is entailed by ρ.

Generalized Removal of W (GRW) does more than join the two rules of inference just examined: it goes beyond them, because it allows for cases in which the set of dominators for the removable W spans the W-set/L-set divide. Each of the relevant two original rules of inference respects this distinction, with concomitant loss of coverage. To see the value of generality, and the cost of its loss, consider the following example, which falls under neither Reverse Ranking Elimination or Part-for-Whole selection.

(23) **Beyond the Four Rules**. P = {π, ρ}

π = (W, W, **W**, L)
ρ = (e , W, **L**, W)

Constraint **3** in π is the target, and the conditions of GRW are met. The relevant ω is (W,W,L,W) and it is entailed by ρ.

Looking into the particular details, we see that, by virtue of ρ, constraint **3** must be subordinated to either **2** or **4**, both of which are polar in π (and therefore form a subset of π's polar set.) Therefore, constraint **3** is doing no work in π and its W can be dispensed with. The MPG for the set {π,ρ}of example (23) looks like this:

(24) **MPG**(P)

π = (W, W, **e**, L)
ρ = (e , W, L,W)

To confirm that this case really doesn't fall under Reverse Ranking Elimination (13), recall that the inference rule speaks only of a *pairwise ranking condition* relating *L*(π), namely **4**, to C, namely **3** in this example. No such thing exists here: P does not require that **4**≫**3**.

And Part-for-Whole Selection speaks only of a proper subset of *W*(π). Since *W*(ρ) spans the W-set and the L-set of π, it does not fall under that condition either. But it is manifest that **3** is necessarily dominated by *some* polar constraint of π.

Lest it be thought that the region beyond the four rules is populated only by complex multiply disjunctive entities, which no working linguist will ever encounter, consider this case, which lies half-hidden within ex. (23).

(25) **Beyond the Four Rules Again**. P = {π, ρ}
π = (W, **W**, L)
ρ = (W, L, **W**)

There's an interesting mutuality here: the bolded W at **2** in π is removable because of ρ. And the bolded W at **3** in ρ is removable because of π.[6]

The resulting MPG therefore looks like this:

(26) **MPG**(P)
π′ = (W, **e**, L)
ρ′ = (W, L, **e**)

The striking fact is that we're not just simplifying a big disjunction to a smaller one: we're revealing *nondisjunctive* ranking requirements hidden among disjunctions. GRW is thus absolutely necessary if we're to have any hope of extracting the ranking information from perfectly ordinary disjunctive data.[7]

We now show that GRW is the *only* principle, beyond the basic logic of ERC entailment, that is needed to govern the contents of a Minimal Primitive Generator. To establish this, we want to establish first of all that it works as promised on the simplification side — if a member of the W-set of a PRC is necessarily dominated by a polar colleague, then simplifying the PRC by removing that W is indeed valid. This doesn't place much strain on intuition, and is already reasonably clear from the kind of cases examined above.

Second, going in the other direction, we want to show that whenever removing a W gives valid simplification of a PRC set, it must be that the constraint assessing the removed W is necessarily subordinated to some polar constraint in the simplified PRC. This claim is not entirely obvious, and it is sensible to demand explicit proof. Putting the two results together completely settles the question of what principles are needed to simplify a PRC set.

We begin by defining some key notions, go on to cite as a lemma the fundamental result on entailment between ERCs from Prince 2002a. We then prove the theorem. A qualitative discussion follows, which aims to lay out the content of the formal analysis.

(27) **Dfn. ERC Entailment.** For A, a set of ERCs, and β an ERC, A entails β, written $A \vDash \beta$, iff every ranking that satisfies the conditions of A also satisfies the conditions of β.

(28) **Dfn**. **Logically Independent**. A set of ERCs is logically independent iff for every $\alpha \in A$, it is not the case that $A - \{\alpha\} \vDash \alpha$. (RN-I:12, §2.3)

(29) **Dfn**. **Equivalence of ERC sets**. ERC sets A and B are equivalent, written $A \equiv B$, iff for every $\beta \in B$, $A \vDash \beta$ and for every $\alpha \in A$, $B \vDash \alpha$.

---

[6] In terms of the GRW, we observe that ρ provides the relevant ω for π, with respect to constraint **2**, and that π provides the relevant ω for **3** in ρ.

[7] OTSoft doesn't simplify example (25), as would be predicted from its adherence to the four rules.

(30) **Lemma**. **Entailment between two ERCs**. Let α and β be nontrivial ERCs. Then $\alpha \vDash \beta$ iff $W(\alpha) \subseteq W(\beta)$ and $L(\beta) \subseteq L(\alpha)$.

Proof. See Prince 2002a:6, Prop. 1.1.

Let us now state and prove the theorem.

(31) **Theorem**. **General Removal of W.** Let P be a logically independent set of PRCs over a constraint set Σ. Suppose $\pi \in P$. Let π′ be a PRC such that $L(\pi')=L(\pi)$ and $W(\pi')=W(\pi)-\{C\}$ for some $C \in W(\pi)$. Let P′ be a set of PRCs exactly like P except that π is replaced with π′, i.e. $P' = (P- \{\pi\}) \cup \{\pi'\}$. Let ω be a PRC, not necessarily included in P, such that $W(\omega) = WL(\pi) - \{C\}$ and $L(\omega)=\{C\}$.

Then, $P \equiv P'$ iff $P \vDash \omega$.

Proof of the theorem.

● **Right to left direction**. Suppose $P \vDash \omega$. Now consider $\pi \circ \omega$. We have $P \vDash \pi \circ \omega$ from the properties of fusion, because $\pi \in P$. By construction, $WL(\pi) = WL(\omega)$. Observe that $\pi \circ \omega$ has an L at $L(\pi)$ and an L at $L(\omega)=\{C\}$, while π′ shares the L at $L(\pi)$ and has an *e* at C. Furthermore, $W(\pi \circ \omega)=W(\pi')$, because $W(\pi \circ \omega)= WL(\pi)-\{L(\pi),C\}=W(\pi')$. Therefore, by the Lemma, $\pi \circ \psi \vDash \pi'$. Thus $P \cup \{\pi'\} \equiv P$. Since $\pi' \vDash \pi$ by W-extension, we have, for $P' = (P- \{\pi\}) \cup \{\pi'\}$, $P' \equiv P \cup \{\pi'\}$. This gives us the desired result, $P' \equiv P$.

● **Left to right direction**. Now assume $P \equiv P'$, where P′ is obtained from P by substituting π′ for some $\pi \in P$, where $L(\pi')=L(\pi)$ and $W(\pi')= W(\pi)-\{C\}$, as defined in the theorem statement.

We need to show that $P \vDash \omega$, where by definition $L(\omega)=\{C\}$ and $W(\omega)= WL(\pi)-\{C\}$.

Note first that we have $P \vDash \pi'$, since by assumption $P \equiv P'$ and $P' \vDash \pi'$ trivially. By Prop 2.5 of Prince 2002a, we are guaranteed the existence of a subset $X \subseteq P$ such that $fX \vDash \pi'$, writing fX for the ERC that is the fusion of all the elements in X.

Now let us inquire into the structure of fX, in particular into the value that it has at coordinate C, which we will write $fX[C]$.

- **Claim**: $fX[C] \neq W$. This cannot be, because we know $fX \vDash \pi'$ and $\pi'[C]=e$. A W-coordinate in the antedecent of entailment requires a corresponding W coordinate in the consequent, as the Lemma states.

- **Claim**: $fX[C] \neq e$. This is the interesting case.

Suppose for purposes of reductio, that $fX[C] = e$. We have $fX \vDash \pi'$, and $\pi' \vDash \pi$ by W-extension, since π′ is just π with one of its W coordinates changed to *e*. But then we have $fX \vDash \pi$ by transitivity of entailment. Notice that π itself cannot belong to fX, since π has W at C and fX has, by our reductio assumption, *e*. Therefore $X \subseteq P-\{\pi\}$ is a subset of P excluding π which entails π, contradicting the assumed logical independence of P. This contradiction establishes the claim.

-**Claim**: $fX[C]=L$. This is the only value left for it to have.

Now, since $fX \vDash \pi'$, $W(fX) \subseteq W(\pi') = W(\pi)-\{C\} \subseteq WL(\pi)-\{C\}= W(\omega)$, by the Lemma.. Since $L(\omega)=C \in L(fX)$, we have it that $L(\omega) \subseteq L(fX)$. These two conditions, one on the W-set ('W-extension'), the other on the L-set ('L-retraction), guarantee entailment.

Hence, $fX \vDash \omega$. This gives $P \vDash \omega$. QED.

Discussion. The second part of the proof asks this question: where could the new *e*-coordinate in a W-removed PRC come from, entailment-wise? Imagine some other ERC that entails it. It can't have an W there, and if it had an *e*, it would have to entail the original, pre-simplification PRC as well, violating one of the conditions of the theorem, namely that we're looking at a logically independent set of PRCs. So, the entailer must have an L, indicating subordination, and this is the only way to simplify, given logical independence. The proof cashes out this line of argument in a fully general context.

Any proof involves sorting through various details, but from an even broader perspective, the result should not come as a surprise. An MPG is nothing more than an exploded version of the Skeletal Basis. The Skeletal Basis replaces predictable L's with *e*'s. The way to arrive at the Skeletal Basis is through L-retraction on the L-rich MIB. Thus, it is inherently plausible that the new *e* in the replacement PRC should be permitted to occupy its coordinate only because a stronger ERC, which has an L there, is also entailed by the original set.

Finally, note that theorem only justifies simplification in a single coordinate. But any multiple simplifications can be arrived at by chaining together one-step moves.

The assurance of completeness in the GRW theorem immediately leads to a better characterization of the global role of constraints belonging to the W-set of a PRC in the MPG. Locally, in terms of a single generic ERC, a W-set member is laxly portrayed as a possible dominator of its entire accompanying L-set (and indeed of all the other polar constraints), with the caveat that other ERCs in the system may take away this tantalizing privilege. In an MPG, as indeed in the MIB and the Skeletal Basis, this local promise is made good globally: if a constraint belongs to the W-set of an ERC in any of these privileged collections, then we are guaranteed the existence of a grammar generating the target language (*qua* linear order on the constraints), in which that constraint actual does dominate its polar confrères.

(32) **Corollary to GRW Theorem**. **Local-to-Global Transport of Disjunctivity**. Let M be the MPG for a set of ERCs A over a set of constraints Σ.. For any $\pi \in$M and C$\in\Sigma$, if C$\in W(\pi)$, then there exists a linear order R on Σ, such that R satisfies M and C dominates every constraint in $WL(\pi)-\{\text{C}\}$.

Proof. Suppose we have an $\pi \in$M and C$\in W(\pi)$ such that (for purposes of reductio) there is no linear order on Σ satisfying M in which C dominates all the members of $WL(\pi)-\{\text{C}\}$. This means, equivalently, that in *every* order satisfying M some constraint from $WL(\pi)$−C dominates C. But if this is the case, then C and π satisfy the conditions of the GRW theorem (31). Therefore, by the Theorem, π cannot be a member of the MPG. Contradiction.

The same result must also extend, as claimed, to the Skeletal Basis and the MIB. The Skeletal Basis is exactly the MPG compressed by fusing PRCs with the same W-set. It therefore has exactly the same W-sets as the MPG, and result holds, because C belongs to the W-set of a PRC in the MPG iff it belongs to the W-set of an ERC in the Skeletal Basis, and they characterize exactly the same set of grammars. Similarly, the W-set structure of the MIB is identical to that of the Skeletal Basis, and the same follows for it. Thus the structure of these particularly 'informative' generative sets is such that the local guarantee of disjunctivity holds good in the global environment of the whole set.

3. The treatment of necessity in OTSoft

Since the underlying inferential structure of OTSoft is incomplete, it must necessarily fall short of producing the MPG in every case. In addition, while pursuing the issues discussed in the paper, it was found that the inferential system failed to perform up to its maximal capacity in certain cases. These issue affect three parts of the program's report: the list of ranking arguments, the diagrams (which depend directly on the ranking arguments), and the section labeled "Status of Proposed Constraints: Necessary or Unnecessary". With the exception of the last, there is no reason to believe that the assertions made by the program are ever false, in the sense of providing a statement that is contradicted by data. But the assertions can be *incomplete*; they only become false when global closure is assumed. Users should therefore avoid assuming closure and augment any claims based on necessity with analysis of their own. In many cases, it is entirely practical to check the assertions by hand using the ERC calculus, either rigidly following the steps of FRed or simply hunting down entailments by free-lancing among the ERCs in an spreadsheet or tabular environment where it is easy to shift rows and columns.

To provide a sense of where things can go wrong, let's look at some issues that came up in a study of the simplest typology of QI stress in Alber 2005. In this system, a LR iambic system with binary feet and no distinction of main stress has the following MPG (checked by hand). Constraints are numbered in the order in which they were presented to the program (left to right, occupying columns D-J in the spreadsheet). Exemplificatory winner-loser pairs are shown that lead to the cited PRCs.

(33) **MPG** for xX-xX-(x) in Alber 2005

PRC#	Winner	Loser	**6**:NoClash	**2**:Rh=I	**3**:FtBin	**7**:NoLapse	**4**:Parse-σ	**5**:AFL	**1**:Rh=T
a	xX-xX	Xx-Xx		**W**					**L**
b	xX-xX-x	X-xX-xX.			**W**		**L**	**W**	
c.i	xX-xX-x	xX. xxx				**W**	**W**		**L**
c.ii	xX-xX-x	xX. xxx				**W**	**W**	**L**	

Method: the data was presented in two candidate sets, 9 five-σ examples, followed by 5 four-σ examples. NB: this amounts to more examples than are needed for the MPG, and with different row and column order than is shown above. The constraints were arbitrarily listed as follows:

1:Rh=T 2:Rh=I 3:FtBin 4:Parse-σ 5:AFL 7:NoLapse

The program reported the following three ranking arguments (re-notated with '|' for disjunction).

3|5 ≫ **4**
2|3| 5 ≫ **4**
4|7 ≫ **1**

Notable here:

● Missed entirely is the ranking argument c.ii, **4|7**≫ **5**;

● Superset Redundancy has *not* been invoked to eliminate **2|3|5** in favor of **3|5** as a dominator set for **4**.

In addition, the program offered the following classification of the constraints:

(34) **Claims of necessity** from OTSoft

Constraint	**Status**
2:Rh=I	Necessary
3:FtBin	Necessary
7:NoLapse	Necessary
1: Rh=T	Not necessary
4:Parse-σ	Not necessary
5:AFL	Not necessary
6:NoClash	Not necessary

The MPG (33), however, reveals that **7**:NoLapse and **4**:Parse-σ are entirely comparable as far as necessity goes. Any work done by one can be done by the other, although **4**:Parse-σ must be subordinated and **7**:NoLapse need not be. (Their work is to dominate **5**:AFL and **1**:Rh=T.)

Similarly, **3**:FtBin is regarded as necessary and **5**:AFL unnecessary; but once again consultation of the MPG shows that they completely shadow each other. Each is as good as the other for covering **4**:Parse-σ. In this case, too, one must be subordinated and the other not. Tesar conjectures (p.c.) that it is the subordinated status of the "not necessary" member of each pair that determines its classification. Pursuing this observation, it appears that the choice of designation may be determined by a stratum-sensitive operation based on the RCD output. However, this cannot yield reliable information about necessity.

At bottom, the binary modal distinction itself is too coarse to portray the actual situation when disjunctions are involved. The logical structure of the MPG simply doesn't reduce to a two-way classification of this type. As we are instructed by the Corollary (32) to the GRW theorem, the global disjunctive structure of an MPG directly reflects the local structure of its component PRCs: *some* constraint in the W-set must dominate the L, and will certainly do so in some admitted grammar. Thus, any constraint appearing as a W among W's in some PRC of the MPG will be as necessary and as unnecessary as its fellow W's. We cannot do better than to describe it as one among the possible dominators in the W-set of the PRC where it sits.

In the course of experiments aiming to narrow down the source of the problem that leads to the overlooked ranking argument, the following subcollection of constraints and candidates was found to cause the same behavior:

(35) **Smallest known example with a missed ranking argument**

ERC#	Winner	Loser	**1**:Rh=T	2:Rh=I	**4**:Parse-σ	**5**:AFL	**7**:NoLapse
a	xX-xX-x	<x>.Xx.<xx>	**L**	**W**	**W**	**L**	**W**
b	xX-xX-x	xX.<xxx>	**L**		**W**	**L**	**W**
c	xX-xX	Xx-Xx	**L**	**W**			
d	xX-xX	xX.<xx>	**L**		**W**	**L**	**W**

The program finds **2**≫**1** and **4**|**7**≫**1**, but misses **4**|**7**≫**5**, just as it does with the full set. However, no superset redundancy was in evidence in the report. Removing row (c) and re-running the ranking algorithm leads to the discovery of both **4**|**7**≫**1** and **4**|**7**≫**5**, but the redundant (a) **2**|**4**|**7**≫**5** is also included; The equally redundant sub-PRC of (a), **2**|**4**|**7**≫**1**, is not mentioned.

This behavior was replicated with a generic, constructed example that follows the structure of the ERCs shown in (35), using both one candidate set and two.[8]

Completely missing a ranking argument raises the spectre of disaster: if **4**|**7**≫**5** goes unrecognized, then inserting **5**≫**4** and **5**≫**7** in blatant contradiction to the missed requirement ought to be allowed. However, the calculation of a sufficient ranking by RCD shows no sign of error in the program, and, in an experimental run in which these contradicting conditions were added, OTSoft had no problem detecting the inconsistency, reporting that no satisfactory ranking existed, a discernment well within the powers of RCD. (Tracking down the source of inconsistency is another matter again.)

We note finally that both order of presentation and structure of the data set had a significant effect on the program's performance. Presenting the full data set with rows and columns re-ordered in accord with the program's RCD output resulted in perfect performance. Various other orders and various inclusions of redundant data showed various effects on the appearance and nonappearance of ranking arguments and redundancies.

4. Summary and Conclusions

Issues with OTSoft:

1. The incompleteness of the inferential system behind OTSoft means that it can fail to give a fully reduced report of the ranking system. This can include omission of entailed nondisjunctive arguments, as in (25).

2. The classification of constraints as 'necessary' and 'not necessary' must be rejected because it imposes an unjustifiable dichotomy on the constraints.

3. Certain implementational problems were noted, where the inferential system did not perform up to its capacity. Given certain data, presented in certain orders, nonentailed arguments can completely missed and redundancies can go unremoved.

4. Problems discussed here are limited to the ones just cited. Nothing has been found that casts doubt on the correctness of the use of RCD to produce typologies and individual stratified hierarchies.

Computing a fully reduced representation of ranking content requires methods beyond those currently incorporated into OTSoft. The program may be used to serve as a guide to necessity, but its assertions must be checked.

The ideas behind OTSoft, as incorporated in the four rules of inference, submit to a productive clarification and generalization when viewed from the perspective of ERC theory. The sought-for collection of reduced arguments can be defined as a Minimal Primitive Generator (MPG), which echoes the Skeletal Basis, and the one condition that's needed to circumscribe it, beyond general commitment to the theory, emerges as a generalization of the essential content of two of the four rules.

[8] OTSoft does not, of course, handle ERCs natively, but they can be faked easily enough. For the desired optimum, enter a candidate with all 1's in its violation cells. For W in a competitor, enter 2; for L, enter 0. This establishes a 1:1 correspondence between integer violation values and W,L,e comparative values: 0 ↔L, 1↔*e*, 2 ↔ W.

5. References

ROA= http://roa.rutgers.edu

Brasoveanu, A. and A. Prince. 2005. Ranking and Necessity. Part I. ROA-794.

Brasoveanu, A. and A. Prince. In prep. Ranking and Necessity. Part II.

Hayes, B.1997. Four Rules of Inference for Ranking Argumentation. http://www.linguistics.ucla.edu/people/hayes/otsoft/argument.pdf.

Hayes, B., B. Tesar, and K. Zuraw. 2003. OTSoft 2.1. Software package. http://www.linguistics.ucla.edu/people/hayes/otsoft/

Prince, A. 2002a. *Entailed Ranking Arguments*. ROA-500.

Prince, A. 2002.b. Arguing Optimality. ROA-562.

Prince, A. 2006. Implication and Impossibility in Grammatical Systems: What it is & How to find it. ROA-880

Prince, A. & P. Smolensky 2004 (1993). *Optimality Theory: Constraint Interaction in Generative Grammar*. Blackwell: Malden-Oxford-Carlton. Also available as ROA-537.

Samek-Lodovici, V. and A. Prince. 2005. Fundamental Properties of Harmonic Bounding. ROA-785.

Tesar, Bruce, and Paul Smolensky. 1993. The Learnability of Optimality Theory: an Algorithm and some Basic Complexity Results. ROA-2.