20. The Optimality Theory — Harmonic Grammar Connection

Géraldine Legendre, Antonella Sorace & Paul Smolensky

Harmonic Grammar plays a key role in the overall ICS architecture: it mediates between the highest-level descriptions of fully symbolic Optimality Theory, and the lowest-level descriptions purely in terms of connectionist networks. What does Optimality Theory add to grammatical theory, beyond what Harmonic Grammar provides? And how exactly are the two connected in a consistent overall theory? The first question is addressed via a number of general sub-issues concerning the similarities and differences, real and apparent, between OT and HG. The first question is addressed by revisiting the problem at the syntax/semantics interface studied with Harmonic Grammar in Chapter 11, split intransitivity; a new OT analysis of this phenomenon is developed and compared to the HG account (contributing to ② of Figure 6 in Chapter 2's ICS Map). Finally, how exactly does HG enable a link between an OT grammar and a connectionist network? (⑤ and ⑦ of Figure 5, and (27), Chapter 2). An OT grammar (for syllabification) is explicitly reduced to a Harmonic Grammar, which in turn is reduced to a local connectionist network. Simulations show that the dynamics of this network allows it to build correct syllabifications, sometimes by quite indirect routes. Reducing local connectionist networks to still lower-level distributed networks is discussed in general in Chapter 11.

Contents

1. The OT/HG connection: General remarks....905
1.1. Universality....906
1.2. Numerically weighted constraints vs. strict domination....906
1.2.1. Empirical considerations....906
1.2.2. Computational considerations....910
1.2.3. Local conjunction and strict domination....912
1.3. Relative vs. absolute Harmony; graded acceptability....917
1.4. Interpretation- vs. production-orientation....918
2. A case study at the syntax/semantics interface: Revisiting split intransitivity 919
2.1. Symbolic approaches to the problem of cross-linguistic unaccusativity mismatches....919
2.2. Modeling cross-linguistic gradience in auxiliary selection as a semantic/aspectual hierarchy....922
2.3. A featural analysis of auxiliary selection in French and Italian....926
2.4. OT analysis....929
2.4.1. Harmonic Alignment....929
2.4.2. Formulating a $*1$ Power Hierarchy....930
2.4.3. Establishing the language-particular ranking of $*2$....932
2.4.4. Verifying predictions in French, Italian, and beyond....933
2.4.5. Partial constraint ranking....935
2.5. Concluding remarks: the larger perspective....936
3. A case study in phonology: Syllabification in Berber....937
3.1. Introduction to Berber....937
3.1.1. Sonority: Nuclear Harmony, HNUC....938
3.1.2. Syllable structure....939
3.2. The Dell-Elmedlaoui syllabification algorithm: DE....940
3.3. Positive OT....942
3.4. Positive HG....944
3.5. Negative OT....947
3.6. Negative HG....949
3.7. Connectionist implementation....950
3.7.1. BrbrNet....950
Figure 1. BrbrNet: Syllabifier for Imdlawn Tashlhiyt Berber....950
3.7.2. $HG_+ \leftrightarrow$ Net architecture....951
Figure 2. BrbrNet′: Equivalent syllabifier for Berber....953
3.7.3. BrbrNet dynamics....953
3.7.4. Competence/performance divergence....954
3.7.5. Simulations....956
Figure 3. Parsing /txznt/ → tx́zńt....958
Figure 4. Parsing sonority profile 8 1 2 1 3 4 5 7 8_7....959
Figure 5. Parsing sonority profile 1 2 7 8....960
Figure 6. Parsing sonority profile 1 2 3 7 8....961
Figure 7. Parsing sonority profile 8 7 2 1 2....962
3.8. Summary....962
References....963

The ICS architecture incorporates two optimization-based grammatical frameworks, Harmonic Grammar (Chapters 6, 10, and 11) and Optimality Theory (Chapters 4, 12, and 13–19). HG provides a level of description intermediate between the highest, fully-symbolic level of OT and the lowest, fully-connectionist level. Like OT, the representations employed in HG are symbolic. As in connectionist computation, the interaction of constraints is via numerical weighting. In this chapter we explore the HG-bridge between the connectionist and the symbolic. Primarily, this is done through two extensive case studies. The first considers how an HG account changes when it is pulled up the level of OT; the second examines how an OT account can be pushed down to an HG account which can then pushed be down to a connectionist network. The first study concerns syntax, the second, phonology.

The first study (Section 2) starts with the HG account of split intransitivity in Chapter 11, and recasts the analysis in OT. This highlights how OT provides a theory of universal typology: while the HG account is fully language-particular, the OT analysis focuses on cross-linguistic variation. The contrast with OT's rigidly restrictive constraint interaction by strict domination also puts in relief how HG's numerical interaction allows it to describe more complex language-internal patterns. The new theory of split intransitivity developed in this section uses OT to bring into sharp focus the phenomenon's highly challenging mix of strong universal tendencies together with dramatic cross-linguistic differences.

The second case study (Section 3) proceeds in the other direction. It starts with an OT analysis of the unusual syllabification system of Berber (from Prince and Smolensky 1993/2004). This is then recast in HG terms. An HG grammar for Berber syllabification can be readily realized as a local connectionist network — *BrbrNet*.[1] The computational properties of this network are, like the linguistic data it models, quite remarkable. While the connectionist level of ICS is assumed to involve distributed representations, Chapter 11 shows how a local network can itself be regarded as a kind of higher-level description — of a still lower-level, distributed network. Thus BrbrNet has a place in the overall ICS picture linking high-level OT grammars to low-level distributed connectionist networks.

Before launching into these two case studies, we begin in Section 1 by considering a number of general conceptual issues concerning the HG-OT relationship. HG and OT are both relatively new frameworks, and most of the questions we will address concerning their relationship are still open; much of the discussion will therefore be somewhat speculative.

1. THE OT/HG CONNECTION: GENERAL REMARKS

The fundamental similarities between Harmonic Grammar and Optimality Theory are fairly evident: the output of the mappings specified by the grammar — the structures declared grammatical — are those that optimally satisfy a set of constraints

[1] Chapter 10 shows how HGs for formal languages can be realized in local connectionist networks.

which apply in parallel to evaluate alternatives; the constraints are simple — that is, general — and therefore typically in conflict. Conflicts are adjudicated by the differential strength the grammar assigns to constraints. Intricate grammatical patterns emerge from the complexity of interaction of fundamentally simple constraints.

But there are also a number of differences between the two formalisms. In some ways, OT imposes further restrictions on a Harmonic Grammar. In other ways, the two theories appear to be in fundamental conflict. In certain respects, the kinds of grammatical questions the two theories have focused on differ. And the levels of description the theories adopt differ in their distance from a lower connectionist level. (For a direct comparison of the fundamental principles of OT with those of connectionism, see Chapter 22 (6).)

1.1. Universality

A principle absolutely central to Optimality Theory is the universality of grammatical structures and constraints. This is one respect in which OT is more restricted than HG. The constraints employed in the HG analysis of split intransitivity discussed in Chapter 11 were intended to embody universal tendencies, but the emphasis was on capturing difficult interactions within a single language. While language-particular analysis plays an important role in OT as in all grammatical theory, OT places a strong emphasis on explaining cross-linguistic patterns via the reranking of a fixed set of hypothesized universal constraints.

1.2. Numerically weighted constraints vs. strict domination

The most obvious difference between Harmonic Grammar and Optimality Theory is also one that underlies many of the other differences: both theories resolve constraint conflict by differentiating the strengths (or 'weights') of constraints, but in HG this is formalized with numerical strengths and in OT with a strict priority ranking. We consider two facets to this issue: empirical and computational.

1.2.1. Empirical considerations

The empirical question at issue is whether in fact grammatical constraints interact in accord with the principle of strict domination. The body of empirical work in OT to date seems to say: yes — with qualification. Without attempting to justify it here, our judgment is that, with respect to those constraints that seem implicated in broad cross-linguistic patterns, much empirical evidence supports the hypothesis that strict domination captures the core of grammatical constraint interaction, in at least most of phonology and much of syntax.

1.2.1.1. Grammars can't count

A central component of the empirical basis for strict domination is synopsized in the oft-repeated adage, *grammars can't count*. A concrete and simple illustration is provided by the interaction between two constraints from stress theory. The constraint STRESSHEAVY (more properly, WSP, the 'weight-to-stress principle': Prince 1990; Prince and Smolensky 1993/2004) requires that heavy syllables be stressed. (Essentially, a syllable is heavy if it has a long vowel or, in many languages, if it ends in a consonant.) The other constraint, MAINSTRESSRIGHT, requires a word's primary stress to be on the rightmost syllable; it is *gradiently violable*, violated once for each syllable that separates the main stress from the right edge of the word. Let's compare the possible patterns arising from the interaction of these two constraints under OT and HG.

In OT, there are only two possibilities: either STRESSHEAVY dominates MAINSTRESSRIGHT or the opposite ranking holds. These two situations are considered in (1)–(2). These are constraint tableaux introduced in Chapter 4. A light syllable is denoted simply by σ, a heavy syllable by σ_H. The acute accent marks the syllable bearing main stress: σ́. The input is a word with a heavy syllable σ_H followed by *n* light syllables σ (such as hypothetical *beetata⋯ta*). The question is: where is the stress? On the heavy (leftmost) syllable, satisfying STRESSHEAVY, or on the rightmost syllable, satisfying MAINSTRESSRIGHT? As always, this conflict between constraints is resolved in OT by ranking, with the higher-ranked constraint taking priority.

(1)

		Candidates	MAINSTRESSRIGHT	STRESSHEAVY
a.		$\acute{\sigma}_H\underbrace{\sigma\sigma\cdots\sigma}_{n}$	$\underbrace{**\ldots*}_{n}$	
b.	☞	$\sigma_H\underbrace{\sigma\cdots\sigma\acute{\sigma}}_{n}$		*

The tableau in (1) illustrates a language where MAINSTRESSRIGHT is dominant. The optimal candidate, (1*b*), has final stress, in satisfaction of the higher-ranked constraint. For our purposes, the interesting case is the reverse ranking, illustrated in (2).

(2)

		Candidates	STRESSHEAVY	MAINSTRESSRIGHT
a.	☞	$\acute{\sigma}_H\underbrace{\sigma\sigma\cdots\sigma}_{n}$		$\underbrace{**\ldots*}_{n}$
b.		$\sigma_H\underbrace{\sigma\cdots\sigma\acute{\sigma}}_{n}$	*	

Tableau (2) depicts a language where STRESSHEAVY dominates; now the optimal candidate has stress on the initial syllable, (2*a*). The point is this: the force of strict domination is to ensure that initial stress is optimal, *no matter how much this violates* the lower-ranked constraint favoring final stress, i.e., no matter how many light syllables σ follow the initial heavy syllable. That is, a single violation of the top-ranked constraint must outweigh *any number* of violations of the lower-ranked constraint.

So under OT, the typology of possible languages, with respect to the interaction of these two constraints and this type of input, contains only two language types. In one (1), stress will always fall on a heavy syllable (as in Hindi); in the other, stress will always be final (2) (as in French).

Under HG, however, the typology implicit in these two constraints exhibits *infinitely many* possibilities. For suppose the numerical strengths of the two constraints are $w_{\text{StressHeavy}}$ and $w_{\text{MainStressRight}}$: each violation of a constraint lowers the Harmony by a quantity equal to its strength. Then the Harmony of the two candidates are:

(3)

	Candidates	StressHeavy	MainStressRight	**Harmony**
a.	$\acute{\sigma}_H\underbrace{\sigma\sigma\cdots\sigma}_{n}$		$\underbrace{**\ldots*}_{n}$	$-n(w_{\text{MainStressRight}})$
b.	$\sigma_H\underbrace{\sigma\cdots\sigma\acute{\sigma}}_{n}$	*		$-w_{\text{StressHeavy}}$

Now suppose StressHeavy is the stronger constraint; under strict domination, this implies initial stress (3*a*) is optimal. Under HG, initial stress maximizes Harmony iff

$$-n(w_{\text{MainStressRight}}) > -w_{\text{StressHeavy}}$$

that is, iff

(4) $n < w_{\text{StressHeavy}}/w_{\text{MainStressRight}} \equiv r$

StressHeavy being dominant means its weight is greater, so the ratio r in (4) is greater than one — how much greater presumably varies typologically; in HG, r can inprinciple be any number. So if the word is short enough, if n is less than r, then initial stress maximizes Harmony, as with strict domination. But if the word is long enough, if n exceeds r, then stress flips to the final syllable: the cumulative effect of many violations of MainStressRight — the cost of stress falling n syllables from the right edge of the word — "gang up" and overcome the single violation of the higher-ranked constraint StressHeavy that mars the final-stress alternative (3*a*). The word length at which the stress flips from initial to final, set by the ratio of constraint weights r, can be any number. Thus the HG typology predicts that in some languages, in words with a single, initial, heavy syllable, stress will be initial for words shorter than four syllables, but final for longer words; in other languages, stress will be initial for words shorter than five syllables, otherwise final; etc., ad infinitum.

It is exactly this sort of typology that is banned as empirically impossible by the high-level generalization *grammars can't count.*

And the example we've developed here is only the very simplest typological prediction of HG-style, numerically weighted, constraint interaction: even more empirically suspect predictions are readily obtained by simply considering more than two constraints. If the weights of three conflicting constraints are, say, $w_A = 11$, $w_B = 3$, $w_C = 2$, then the strongest constraint A will over-ride the preferences of the two weaker constraints B and C, if each of them is violated at most once; but the weaker constraints can 'gang up' and over-rule A if: B is violated four or more times, or C is violated six or more times, or B is violated three times and C twice, or … . Once a realistically-sized set of conflicting constraints is considered, the 'counting' required of such grammars is staggering, and, it would appear, utterly without empirical basis in the extensive body of cross-linguistic studies of grammars of the world. (For a concrete example highlighting strictness of domination in an actual phonological analysis, see Prince and Smolensky 1993/2004: Sec. 7.4n.)

1.2.1.2. Conjunctive constraint interaction

A wide variety of empirical studies employing OT have confirmed strict domination as the basic mode of constraint interaction. It has nonetheless turned out in a number of language-specific analyses, and some typological studies, that deviations from this basic mode can be observed. Even in this regard, however, strict domination serves two crucial roles: it identifies the basic interactive mode providing the baseline for identifying these deviations, and it provides a formal device which can be deployed to develop a theory of the new type of interactions.

This theory asserts that in addition to strict domination, universal constraints sometimes interact via *local conjunction* (Chapter 12 (43); Chapter 14). That is, when two constraints are simultaneously violated within the same local domain, their joint effect can be stronger than their 'linear' sum. In particular, if an isolated violation of $\mathbb{C}_1$ is weaker than a violation of $\mathbb{C}_0$, and an isolated violation of $\mathbb{C}_2$ is also weaker than a violation of $\mathbb{C}_0$, it can sometimes happen that simultaneous violations of $\mathbb{C}_1$ and $\mathbb{C}_2$ 'at the same place' are *stronger* than a violation of $\mathbb{C}_0$.

The formalism developed to capture strict domination can actually be called into the service of a theory of such 'conjunctive' interactions. Given two constraints $\mathbb{C}_1$ and $\mathbb{C}_2$, their *local conjunction with respect to a domain* $\mathcal{D}$, written $\mathbb{C}_1 \,\&_{\mathcal{D}}\, \mathbb{C}_2$, is violated whenever $\mathbb{C}_1$ and $\mathbb{C}_2$ are both violated in a common domain $\mathcal{D}$.[2] This new constraint $\mathbb{C}_1 \,\&_{\mathcal{D}}\, \mathbb{C}_2$ is then ranked with all the other constraints in the hierarchy defining a language's grammar.

Obviously, allowing such conjunctive interactions makes OT considerably less restrictive, and there are at this point many open questions about whether the empirical facts require lessening the restrictiveness of the theory in this way. Regardless of the ultimate fate of the theory of local conjunction, there seems to be considerable consensus that OT has made a major contribution in identifying simple strict domination as the fundamental mode of grammatical constraint interaction.

1.2.1.3. Grammaticization

The HG study of split intransitivity in French discussed in Chapter 11 provides some evidence of numerical constraint interaction. How does this relate to OT's assumption of strict domination? One possibility is this. Knowledge relevant to language processing may combine (i) a system of constraints one might consider more strictly 'grammatical', interacting exclusively or primarily via strict domination, and (ii) a set of more pragmatically-based constraints, reflecting more directly, perhaps, statistical characteristics of experience, and interacting in a less restricted manner, via arbitrarily weighted constraints. The process of *grammaticization* may be one in which con-

[2] For example, if $\mathbb{C}_1$ = NOCODA and $\mathbb{C}_2$ = NOVOICEDSTOPS, and $\mathcal{D}$ = a segment, then the conjunction is violated by a voiced stop in a syllable coda.

straints effectively move from the latter category to the former. The constraints interacting in the HG analysis may constitute a mixture of both types of constraints, while the constraints focused upon in OT studies may be more completely contained in the 'grammatical' class. Computationally, it is possible to combine into one analysis, as perhaps the HG analysis of Chapter 11 does, both types of constraints because, as we now discuss, the strict domination interaction required by OT can be implemented by numerical weighting, putting the grammatical constraints into a computational arena in which they can interact with more pragmatic constraints with arbitrary numerical weights. The OT analysis of split intransitivity in Section 2 explores the implications of strict-domination interaction of certain of the types of constraints included in the HG study of French.

1.2.2. Computational considerations

In a sense, strict domination is one respect in which OT adds further restrictions to HG: strict domination can be seen as a special relation among numerical weights (Prince and Smolensky 1993/2004: Ch. 10). Several of the relevant considerations have already made their appearance in Section 1.2.1.1's discussion of 'counting' by grammars.

1.2.2.1. Exponential constraint weighting

Suppose first that each constraint can be violated only once per candidate. What strengths will yield the strict domination hierarchy $\mathbb{C}_n \gg \ldots \gg \mathbb{C}_2 \gg \mathbb{C}_1 \gg \mathbb{C}_0$? To set the arbitrary origin of weights, let the strength of $\mathbb{C}_0$ be 1. Then obviously the weight of $\mathbb{C}_1$ must be greater than 1; to set the arbitrary scale of weights, let the weight of $\mathbb{C}_1$ be 2. Now in order that $\mathbb{C}_2$ have strict priority over both lower-ranked constraints, the cost of a violation of $\mathbb{C}_2$ must be higher than the cost of violating *both* $\mathbb{C}_1$ and $\mathbb{C}_0$, i.e., greater than 2+1=3; we can keep integer weights by taking the strength of $\mathbb{C}_2$ to be 4. Continuing this logic, we see that the strict domination can be achieved by setting the weight of $\mathbb{C}_k$ to be 2^k. We will refer to this as *exponential weighting*: the strengths of constraints must grow exponentially as the hierarchy is mounted (see Section 21:6.4).

Typically, constraints can be violated more than once in a candidate linguistic structure (for example, the constraint NoCoda from Chapter 4 is violated once for each syllable that ends in a consonant; the constraint MainStressRight above, once for each syllable separating the main-stressed syllable from the word's right edge). Suppose each constraint can be violated anywhere from 0 to 9 times in a single candidate. To strictly dominate $\mathbb{C}_0$, the cost of a single violation of $\mathbb{C}_1$ must exceed that of any number of violations of $\mathbb{C}_0$, so the weight of $\mathbb{C}_1$ must exceed 9; say, 10. Then iterating the logic, we see that the strength of $\mathbb{C}_k$ must be 10^k. This is exponential weighting, but the base of exponentiation has grown from 2 to 10.

So in one sense, OT's strict domination is a special case of HG's numerical weighting: exponential weighting. But in another sense, strict domination is an idealization that no set of numerical weights can actually achieve. For strict domination puts no limit to the number of lower-ranked violations that a higher-ranking constraint overrides: thus no finite base of exponentiation — whether it be 2, or 10, or 10^6 — can *truly* implement strict domination. Thus it is most accurate to state that, as an idealized competence theory, OT constitutes a certain *limiting* case of HG: exponential weighting, in the limit as the base of exponentiation goes to infinity.[3]

Exponential weighting of constraints poses a number of unanswered questions about an underlying connectionist realization. Even for a small base of exponentiation, such weights quickly grow to be enormous as the number of crucially-ordered constraints increases. The challenge this presents to a connectionist realization is the large *dynamic range* required by computation with exponential weighting: no matter how the actual range of weights involved might be compressed by multiplying by a small scale factor, the number of *distinct values* of Harmony that must be accurately compared is large. It is unknown whether there is a way of embedding constraints in a network such that effective weighting of constraints will naturally be exponential.

1.2.2.2. Functional speculations

A quite different question is. *why* should constraint weighting be exponential? From the perspective of familiar connectionist networks, exponential weighting is entirely unexpected; is there some consideration from which it might actually follow? One possible functional motivation may come from the demands that grammar be a *sharable* knowledge system. The optimization involved in, say, planning a reaching movement can produce a different result each time — there's no strong requirement for reproducibility; any trajectory that basically meets the constraints is good enough. But it won't do if each hearer's optimization yields a different structure for a sentence, and the speaker's optimization yields yet another structure. One suggestion, due to David Rumelhart (personal communication), and independently to James McClelland (personal communication), is that exponential weighting of constraints may enable quick-and-dirty optimization algorithms of the sort embodied in connectionist computation to consistently find a single global Harmony optimum, whereas arbitrarily weighted constraints typically lead such algorithms to produce widely varying solutions, each only a local optimum. (See Section 6:2.4 for the relation of the competence/performance distinction to the global/local optimum distinction.) Some

[3] It is worth noting that it is not possible to have strict domination without finite, discrete levels of constraint violation. Thus continuous constraint penalties typical in phonetics, such as c([actual value] – [target value])^2, contrast with the discrete penalties of OT phonology (Flemming 2001). Suppose the weights of $\mathbb{C}_2 \gg \mathbb{C}_1$ are non-standard, infinite real numbers $w_2 > w_1$. For strict domination, any degree of violation ε_2 of $\mathbb{C}_2$ must overpower any finite degree of violation, e.g., 1, of $\mathbb{C}_1$ — i.e., it must be that $\varepsilon_2 w_2 > w_1$. Thus ε_2 must not be an infinitesimal with $\varepsilon_2 < w_1/w_2$. So continuous degrees of violation of $\mathbb{C}_2$ are incompatible with it strictly dominating $\mathbb{C}_1$. (For a less transparent argument, replace the infinite penalties with large finite ones.)

If phonological representations are discrete, it follows that degree of violation will be discrete, as required for strict domination. If, contrary to standard assumptions, phonological representations are continuous (Kirchner 1998), then strict domination requires that each individual constraint discretize its relevant representational continuum, assessing violations in a discrete, discontinuous fashion.

experimental evidence in support of this suggestion is discussed in Section 3.

Another possibility is that demands of *learnability* provide a pressure for strict domination among constraints of unknown strength. Rather strong formal results have been obtained concerning the efficient learnability of strictly ranked constraints (see Section 12:3) but it remains an open question problem to formally characterize exactly what is essential about strict domination to guarantee efficient learning.

In either case — whether due to pressures of reliable, rapid optimization, or efficient learnability — the idea is that while arbitrarily weighted constraints may be the typical case in connectionist networks, evolutionary pressures on the language system may have led to the development of a special architecture — as yet unknown — in which strict domination obtains, at least to the degree needed to ensure that an OT-like grammar would provide a good competence-theoretic idealization.

1.2.3. Local conjunction and strict domination

Consider for the sake of argument the following oft-repeated line of reasoning.

(5) The ganging-up argument: General form

The central premise of OT is strict domination, that no number of violations of lower-ranked constraints can overpower a single violation of a higher-ranked constraint. In particular, if $\mathbb{C} \gg \mathbb{A}$ and $\mathbb{C} \gg \mathbb{B}$, then no degree of violation of the lower-ranked constraints $\mathbb{A}$ and $\mathbb{B}$ can overrule a single violation of $\mathbb{C}$. The conjunction $\mathbb{A}$ & $\mathbb{B}$ is designed specifically so that in this situation, the ranking $\mathbb{A}$ & $\mathbb{B} \gg \mathbb{C}$ *does* entail that a violation of both $\mathbb{A}$ and $\mathbb{B}$ overrules a violation of $\mathbb{C}$. Thus admitting local conjunction into OT nullifies its most basic principle.

In fact, the need for local conjunction shows that strict domination is simply empirically incorrect. Strict domination in OT serves to eliminate the potentially much richer possibilities for constraint interaction that are provided in a theory like OT except that constraints have numerical strengths, and each constraint violation incurs a numerical penalty equal to that constraint's strength. Such a theory is in fact exactly Harmonic Grammar, HG, developed in Chapters 6, 10, and 11. Strict domination corresponds to a very special property of constraint strengths: stronger constraints are numerically *so much* stronger that it is simply numerically impossible for weaker constraints to "gang up on" and overpower a stronger constraint; as explained in Section 1.2.2.1, constraint strengths essentially need to grow exponentially as the ranking is mounted. What the need for local conjunction is telling us is that empirically, HG is correct and OT is not, that weaker constraints *can* gang up in the way that is allowed by HG but not by OT (without local conjunction).

The basic arguments developed in Chapter 14 for the introduction of local conjunction into OT center on the *BOWOW* pattern: inventories that Ban Only the Worst-Of-the-Worst. In this context, the argument in (5) can be made more concrete.

(6) The ganging-up argument: BOWOW inventories

An element that violates the markedness constraint $*\alpha$ is acceptable (i.e., in the inventory); an element that violates the markedness constraint $*\beta$ is acceptable; but an element that violates both is unacceptable (banned from the inventory). This BOWOW inventory requires the local conjunction $*\alpha$ & $*\beta$ in an OT analysis:

$$*\alpha \;\&\; *\beta \gg \mathrm{F} \gg *\alpha, *\beta$$

(F is the faithfulness constraint opposing elimination of the marked structures.)

Numerically, the BOWOW inventory calls for numerical strengths for these constraints with the property that, individually, the weight of the markedness constraint $*\alpha$, $w_{*\alpha\ *\alpha}$, is less than that of F, w_{F}; similarly, individually, $w_{*\beta}$ is less than w_{F}; but together, the *sum* $w_{*\alpha} + w_{*\beta}$ is *greater* than $w_{\mathrm{F},}$, so the markedness constraints "gang up and overpower" F when violated together. This sort of ganging-up is exactly what is barred by strict domination, which requires that $w_{*\alpha} + w_{*\beta}$ never be greater than w_{F}; under strict domination, w_{F} would be so great that no summation of weights of lower constraints could be larger. Thus BOWOW inventories show that strict domination is fundamentally incorrect, that weights for constraints must be allowed to "gang up". Using local conjunction amounts to admitting this failure of strict domination.

For understanding the implications of local conjunction for OT, it is important to see that *this argument is a fallacy*, in either its more general (5) or more concrete (6) form. Whether the weights satisfy the strict domination condition — whether the weights of the weaker constraints together, $w_{*\alpha} + w_{*\beta}$, exceeds that of the stronger constraint, w_{F} — *has nothing to do with* generating a BOWOW inventory.

The reason is extremely simple. To ban the worst of the worst, it is not sufficient that $w_{*\alpha} + w_{*\beta}$ exceeds w_{F}; what is required is that $w_{*\alpha} + w_{*\beta}$ exceed *twice* w_{F}. This is because in the 'worst of the worst' input, competition pits two markedness violations ($w_{*\alpha} + w_{*\beta}$) against *two* faithfulness violations ($2w_{\mathrm{F}}$). And it is simply impossible for $w_{*\alpha} + w_{*\beta}$ to exceed $2w_{\mathrm{F}}$ if, individually, neither $w_{*\alpha}$ nor $w_{*\beta}$ exceed w_{F} — no matter what the numbers, irrespective of whether the weights satisfy the strict domination condition $w_{\mathrm{F}} > w_{*\alpha} + w_{*\beta}$. Eliminating strict domination could allow two weaker violations to exceed *one* stronger violation, but *nothing* can allow two weaker violations to exceed *two* stronger violations.

To see why the competition is in general between two lower-ranked markedness violations and *two* higher-ranked faithfulness violations, consider the simple subset of the English obstruent inventory where place markedness and manner markedness interact: {*t*, *k*, *s*, **x*}; *x*, the velar fricative, is banned from the English inventory.

(7) BOWOW subset of English obstruent inventory

	Place	*[vel]
Manner	***t***	***k***
*[cont]	***s***	***x***

The unmarked segment *t* is in the inventory; and violating either *[vel(ar)] (place markedness) or *[cont(inuant)] (manner markedness) *individually* is allowed, admitting *k* and *s* into the inventory. But the segment violating *both* markedness constraints, *x*, is banned. This is a classic BOWOW inventory. Its analysis under numerically-weighted constraint interaction is simple. In order that the input /k/ have as its optimal output [k] (*[velar]) as opposed to [t], faithfulness to place, F[vel], must outweigh place markedness, *[vel]:

(8) *k* admitted $\Rightarrow$ $w_{F[vel]} > w_{*[vel]}$

By identical reasoning, for the segment *s* marked by *[cont] to be admitted, we must have

(9) *s* admitted $\Rightarrow$ $w_{F[cont]} > w_{*[cont]}$

Together, (8) and (9) entail that faithfulness to place and manner combined outweigh the markedness of place and manner together — so *x* must be in the inventory (10).

(10) $\therefore\ w_{F[vel]} + w_{F[cont]} > w_{*[vel]} + w_{*[cont]} \Rightarrow$ *x* admitted

This is displayed in the *HG tableaux* of (11); for concreteness actual numerical weights have been employed. To emphasize the irrelevance of strict domination, the weights chosen for faithfulness constraints exceed those of markedness constraints by less than 3%; the strict domination condition is not even close to being satisfied, as a single stronger violation is much less than the sum of two weaker violations.

(11) HG tableaux: Impossibility of BOWOW in HG

		F[vel] 5.02	F[cont] 5.01	*[vel] 5.0	*[cont] 4.9	*constraint weight*	
/k/						**Harmony**	
☞	[k]			* –5.0		– 5.00	✓ *k*
	[t]	* –5.02				– 5.02	
/s/							
☞	[s]				* –4.9	– 4.90	✓ *s*
	[t]		* –5.01			– 5.01	
/x/							
☞	[x]			* –5.0	* –4.9	– 9.90	✓ *x*
	[s]	* –5.02			* –4.9	– 9.92	
	[k]		* –5.01	* –5.0		–10.01	
	[t]	* –5.02	* –5.01			–10.03	

In (11), numbers show the Harmony of candidates, the optimal output being that of highest Harmony; the Harmony is the sum of negative contributions from each violation, weighted by the weight of the constraint violated.

The situation is depicted in another way in (12). For input /k/, the move from faithful [k] to unmarked [t] must be Harmony-decreasing, since *k* is in the inventory; call the Harmony difference –ε (this is just $-(w_{F[vel]}-w_{*[vel]})$ in the notation above; in the example (11), –ε = –0.02). Similarly, for input /s/, a move to [t] must produce a Harmony decrease, call it –δ. Now consider input /x/. The move from [x] to [s] incurs the same Harmony penalty as that from [k] to [t] $-(w_{F[vel]}-w_{*[vel]})$; this is –ε, so /x/→[x] has higher Harmony than /x/→[s]. Similarly, /x/→[x] has higher Harmony than /x/→[k] (by the amount δ). And finally, eliminating both marks via /x/→[t] incurs the Harmony penalty which is the sum of the other two, –ε–δ; this too must be less harmonic than the faithful mapping. Since the faithful mapping /x/→[x] has highest Harmony, *x* must be in the inventory.

(12) Net harmony penalties for unfaithful mappings

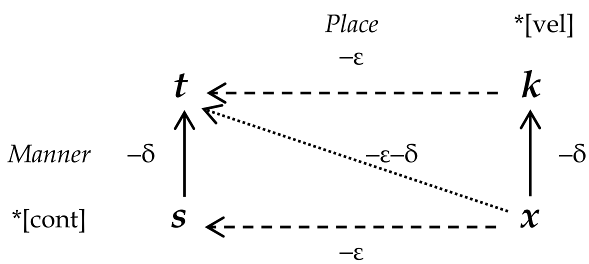

The difficulty of banning only the worst of the worst cannot be overcome by manipulating the strengths of constraints, which is why strict domination is not a relevant issue. The problem resides in the *means of combination* of markedness dimensions: simple summation. The above analysis shows that what is needed for a BOWOW inventory is that *the penalty for violating two markedness constraints is more than the sum of their individual penalties.*[4] We could say that markedness combination must be super-additive or *super-linear*. It is the additive or linear character of HG that prevents BOWOW; as long as markedness combination remains additive, it makes no difference whether strict domination is required of the weights.

What local conjunction centrally achieves is super-linear interaction, not nonstrict domination. For the numerical example in (11), the sum of the penalties for *[vel] and *[cont] is 9.9; if super-linear interaction yields instead a sufficiently higher penalty, e.g., 9.93, then the faithful mapping is no longer the most harmonic (/x/→[s] is), and *x* is now banned (13).

[4] Alternatively, FAITHFULNESS combination could be *sub*-linear, with the Harmony penalty of two faithfulness violations *less* than the sum of the individual penalties. (This would be a version of Burzio 2002's *gradient attraction*, where the strength of FAITHFULNESS between *A* and *B* diminishes as *A* and *B* become less similar.) Sub-linear faithfulness is instantiated in segmental-MAX: if this replaces the individual-feature faithfulness constraints F[φ] in the text, then deleting the entire segment avoids *both* marks ∗[vel] and ∗[cont] given /x/, while incurring the same *single* mark *MAX which is incurred if either one alone is avoided by whole-segment deletion given /k/ or /s/. In general, however, it is not sub-linear FAITHFULNESS but super-linear MARKEDNESS that is the numerical implementation of local conjunction. Given how optimality is computed in OT, making [α, β] more marked than [α] + [β] is straightforwardly achieved by a higher-ranked conjoined constraint, which adds a third mark to ∗α and ∗β; but making two FAITHFULNESS violations *less* costly than the two individual violations cannot be achieved by adding an additional mark: it requires *removing* at least one of the individual marks when two violations occur. Thus a general means of weakening FAITHFULNESS would require considerably more alteration to the basic theory than does strengthening MARKEDNESS.

(13) Possibility of BOWOW with super-linear constraint combination

	[vel]&[cont] 0.03 [+5.0+4.9=9.93]	F[vel] 5.02	F[cont] 5.01	*[vel] 5.0	*[cont] 4.9	*constraint weight*
/x/→						**Harmony**
[x]	* –0.03			* –5.0	* –4.9	–9.93
☞ [s]		* –5.02			* –4.9	–9.92
[k]			* –5.01	* –5.0		–10.01
[t]		* –5.02	* –5.01			–10.03

The argument developed above is quite general; indeed it applies directly to the general formalization of basic inventories developed in Chapter 14 (Appendix), so the theorems derived there all pertain to a general setting in which it is super-linearity, not non-strict domination, that is required to achieve the effects which can be achieved with local conjunction.

The situation we have seen above is covered by a general theorem of Alan Prince.

(14) The 'Anything Goes' Theorem (Prince 2002)

Let $\mathcal{H}$ be a strict domination hierarchy and let A be an optimal output for some input, and B a suboptimal output for that input. Suppose the following holds: For every constraint $\mathbb{C}$ in $\mathcal{H}$, the total number of violations by A of the constraints higher-ranked than $\mathbb{C}$ is less than the total number of violations by B of the same constraints. Then A has higher numerical harmony than B when computed by summing weighted constraint violations, for any set of numerical weights which order the constraints in accord with the ranking $\mathcal{H}$.

This theorem gives a sufficient condition for numerical weights to produce the same optimal outputs as OT harmonic evaluation with no restriction on the rate at which weights grow as the hierarchy is mounted (exponential growth is *not* required).[5] (Prince 2002 shows that this is a necessary condition as well.)

The condition of the Anything Goes Theorem is met in the little example just dis-

[5] *Proof sketch.* Since $A \gg B$, by the Cancellation/Domination Lemma (Prince and Smolensky 1993/2004), every mark of A is either cancelled or dominated by a mark of B. In both OT and the numerical theory, cancellation eliminates irrelevant marks. The highest uncancelled mark of A, call it $*\mathbb{W}$, is dominated by a mark of B, $*\mathbb{L}$. In the numerical theory, this condition is equivalent to $w_{\mathbb{W}} > w_{\mathbb{L}}$, with no restriction on the magnitude of the weights or their difference. Now throw out these two marks, noting that their respective contributions to the Harmonies of A and B favor A (lower penalty) and proceed to the new highest mark of A and repeat. When all marks of A have been discarded in this way, it is proved that the Harmony of A is higher than that of B (there may be remaining marks of B, but these just make it even lower in Harmony). This procedure can be executed whenever the condition of the theorem is met: for every uncancelled mark of A, there is guaranteed to be a higher-ranked remaining mark of B, even after some B-marks have been discarded. Each mark of A is cancelled or dominated by *its own distinct* mark of B. Thus the numerical Harmony of A is greater. But if the conditions of the theorem are not satisfied, all we know from the Cancellation/Domination Lemma is that each uncancelled mark of A is dominated by *some* uncancelled mark of B: but a *single* higher B mark can dominate any number of lower-ranked A marks — this is strict domination.

cussed, and in fact, it is met in general for all 'Basic Inventories' as defined in Chapter 14 (Appendix).[6] Thus the conclusion we have reached is quite general: For *any* numerical weighting of the constraints, the numerically optimal candidate is the same as that in plain OT. Thus, the inability of OT without local conjunction to generate BOWOW inventories applies equally to a theory employing unrestricted numerical weights. Since neither system can derive BOWOW inventories, as discussed at length in Chapter 14, neither is Strongly Harmonically Complete (SHarC). To achieve SHarC typologies, local conjunction is needed in OT, and super-linear constraint combination is needed in HG. Strict domination is irrelevant.

1.3. Relative vs. absolute Harmony; graded acceptability

In the HG analysis of split intransitivity, the absolute level of Harmony of a structure was taken to model its acceptability. Since multiple Harmony values are possible, this means multiple levels of acceptability can be modeled, and indeed the French data analyzed in that study consisted of a graded acceptability judgment, on a 5-point scale, for each sentence. In OT, however, only relative, not absolute, levels of Harmony matter: the highest-Harmony member of a candidate set is the only grammatical one, all other candidates being equally ungrammatical. Indeed it is readily seen upon inspecting many OT analyses that use of relative not absolute Harmony is crucial: in many cases, the absolute Harmony value of the optimal (and indeed correct) candidate in one candidate set is lower than the absolute Harmony of sub-optimal (and indeed incorrect) candidates in a different candidate set. An essential feature of competition-based explanation in OT is that a structure can be grammatical *even though it violates constraints that are fatal in other competitions,* because it is the best option in the candidate set. According to an OT grammar, the best of a bad batch is grammatical, while the second-best of a good batch is not. For this reason, use of absolute Harmony values as a model of graded acceptability, or graded processing accuracy or difficulty, is at odds with the fundamental structure of OT.

And indeed, there is no natural sense in which a connectionist network has access to its own absolute Harmony level: that is just a convenient measure for the analyst looking at the system from outside. Implicitly, a network has access to relative Harmony levels, because its processing algorithm takes it from a state at one moment to a state at the next moment that has higher relative Harmony; and this can be computed entirely locally, by spreading activation. But the absolute Harmony value of the state as a whole is a global quantity which would seem inaccessible to the net-

[6] *Proof* (assuming the definitions and notation of the Appendix of Chapter 14). Given an input *I* specified with feature values φ_i, the optimal output A has features ψ_i, where $\psi_i \equiv \varphi_i$, except $\psi_i \equiv -\varphi_i$ if φ_i is the marked value and $M[\varphi_i] \gg F[\varphi_i]$. That is, for A, for all *i*, either φ_i is unmarked, generating no mark $M[\varphi_i]$ and no mark $F[\varphi_i]$ as ($\psi_i \equiv \varphi_i$), or φ_i is marked, and A violates the lower-ranked constraint of $\{M[\varphi_i], F[\varphi_i]\}$. Let B be a suboptimal output for *I*. Then consider any mark of A; it is for some *i* the lower-ranked of $\{M[\varphi_i], F[\varphi_i]\}$. Either this mark is cancelled in B because B has the same value for ψ_i, or it is dominated because B is marked by the other, higher-ranked, constraint among $\{M[\varphi_i], F[\varphi_i]\}$. Every mark of A has *its own* canceling or dominating B mark. This ensures the condition of the theorem; indeed it's just what the condition of the theorem is designed to ensure: see the proof in note 5.

work itself. It would be surprising, then, if informants reporting acceptability judgments are in effect reporting absolute Harmony levels. Indeed we suspect that it is OT rather than HG that is on the right track here: only relative, not absolute, Harmony levels should be cognitively relevant.

How, then, can HG be reconceived without appealing to absolute Harmony? Simply by making it consistent with OT. To illustrate, we take the relatively clean case of Harmonic Grammars of formal languages, discussed in Section 6:2.3 and Chapter 10. It will turn out our proposal will involve the final difference between HG and OT we consider here, that between a theoretical bias towards grammatical mappings in the interpretation vs. production directions.

1.4. Interpretation- vs. production-orientation

Section 6:2.3 and Chapter 10 describe how, given a context-free formal language $\mathcal{L}$, a Harmonic Grammar can be designed so that a string of terminal symbols μ is in $\mathcal{L}$ iff $H(\mu) \geq 0$, that is, if and only if the maximum-Harmony parse tree of μ has non-negative (in fact, zero) Harmony. There, following what we will now refer to as the *absolute-H* interpretation of HG, we took this criterion to be *definitive* of grammaticality. This is essentially an *interpretation-oriented* criterion: what is given is an 'overt' string of symbols μ (idealizing a sequence of words a hearer might receive), and this is mapped onto the parse tree that maximizes Harmony. The absolute Harmony of this tree is then consulted: if it is negative, μ is not in the language $\mathcal{L}$; otherwise, it is.

There is, however, another view of this HG, the *relative-H* interpretation, which relies on the same *production-oriented* criterion for grammaticality that is used in OT. To generate, rather than recognize, strings in a formal language $\mathcal{L}$, we start with the grammar's start symbol, S, which we think of as denoting the category of grammatical sentence. We then use the grammar to generate all legal trees with S at the root and terminal symbols at the leaves of the tree. In the HG context, that means we take S as the 'input' and assign it to the maximum-Harmony tree with S at the root and terminal symbols at the leaves. What tree is this? There are many such trees — infinitely many, typically. For every grammatical string μ in $\mathcal{L}$, there is a zero-Harmony tree with μ on the leaves and S at the root; all other trees have negative Harmony. So the maximum-Harmony trees with S at the root are exactly the legal parse trees of grammatical strings in $\mathcal{L}$. Thus, simply by reinterpreting the HG using a production- rather than interpretation-oriented definition of grammaticality, we get a relative- rather than an absolute-Harmony characterization of the formal language.

How might the HG analysis of split intransitivity (Chapter 11) be redone in a production-oriented, relative-Harmony framework? In such a framework, to determine the grammaticality of a sentence, it does not suffice to simply evaluate the maximum-Harmony parse for that sentence; we must compare this Harmony to alternative parses of the same input, employing other constructions. A given syntactic expression, with a particular verb and set of arguments, is thus ungrammatical if the corresponding input — the proposition it would express — is optimally expressed via another expression.[7] Note that, as with OT, in this revised conception of grammaticality under HG, it is unclear how to model graded judgments of acceptability.[8]

For reasons discussed at some length in Chapter 16, there are some subtle conceptual and technical issues in OT syntax as regards propositions that have *no* grammatical expression in a language. Most of the unaccusativity tests discussed in Chapter 11 are of the form, 'if a verb V is grammatical in construction *X*, then V is unaccusative'. Thus for a verb that fails the test, there is no grammatical output for the input which corresponds to the construction *X*. It is most straightforward, then, to consider a different type of unaccusativity test. In this type, unaccusative verbs are grammatically expressed one way, and unergative verbs a different way: in every case there is a well-formed expression for the grammar to select. An extremely well-known test of this type is *auxiliary selection*: in certain constructions where a main verb requires an auxiliary verb, unaccusative main verbs 'select' (employ) one auxiliary verb ('be'), while unergative main verbs select a different auxiliary ('have'). A detailed analysis of this phenomenon is a primary topic of Section 2, a case study to which we now turn. It shows how a production-oriented, relative-Harmony approach can be developed to study split intransitivity. The analysis is in OT, but in principle the same approach could be adopted with numerically-weighted constraints in the style of HG.

2. A CASE STUDY AT THE SYNTAX/SEMANTICS INTERFACE: REVISITING SPLIT INTRANSITIVITY

In this section we inspect the Harmonic Grammar–Optimality Theory relation from the perspective of empirical linguistics, taking an HG account discussed earlier in the book and recasting it in OT terms to compare the views of a single phenomena which are revealed by the two different theoretical lenses.

2.1. Symbolic approaches to the problem of cross-linguistic unaccusativity mismatches

Recall from Section 6:2.5 and Box 11:1 that the Unaccusative Hypothesis formalized

[7] In some cases, the optimal output for a given input may be semantically unfaithful to that input, so that it does not in fact express the input proposition, which is only a semantic *target*. In this case, no grammatical sentence expresses the input semantics. See the discussion of ineffability in Chapter 16.

[8] It is a common misconception that OT is naturally more suitable than previous generative theories to modeling graded judgments because it employs violable rather than inviolable constraints. "The concept of violable constraint has still other potential advantages over derivational approaches that have yet to be explored. A prime example is the way such a model easily lends itself to the task of capturing the graded nature of alternative surface forms so characteristic of many variable processes" (Nagy and Reynolds 1997: 39). Ironically, exactly the opposite is the case. With inviolable constraints, it is simple enough to stipulate that, while the normal consequence of violating a constraint is a '*' (ungrammaticality), for certain constraints, a violation only induces a '?' (questionably grammatical). But in OT, for the reasons discussed in the text, the characterization of grammaticality in terms of relative Harmony makes a theory of partial grammaticality more difficult. It may be possible to construct an adequate theory of graded grammaticality by equating 'less grammatical' with 'less frequent' in the theory of variation provided by partial constraint ranking, developed in fact by Nagy and Reynolds 1997 among others (see Chapter 18).

in Perlmutter 1978 and Burzio 1986 states that the class of intransitive verbs divides into two subsets — *unaccusatives* and *unergatives* — which have distinct syntactic properties. The *Grammatical Relation (**GR**)* of the argument of an unaccusative verb is that of an underlying or *deep direct object*, so this argument displays many of the syntactic properties of the direct object of a transitive verb. In contrast, the single argument of an unergative verb is a *subject* at all levels of representation, and thus consistently displays the same syntactic behavior as the subject of a transitive verb. This syntactic difference, discussed at length in Chapter 11, can be represented as in (15).

(15) a. Unergative: NP [$_{VP}$ V] deep subject e.g., *He works hard*
b. Unaccusative: ___ [$_{VP}$ V NP] deep object e.g., *He died recently*

The earliest formulations of the Unaccusative Hypothesis noted that the distinction is also systematically related to certain lexical-semantic characteristics of the predicate: 'agentivity' tends to correlate with unergativity and 'patienthood' with unaccusativity (Perlmutter 1978; Dowty 1991). Much subsequent cross-linguistic research has shown, however, that the alignment between syntactic and semantic properties is not as consistent as originally predicted: some verbs with similar lexical semantics have different syntactic behavior across languages (for example, 'blush' is unaccusative in Italian and unergative in Dutch: Rosen 1984), and some verbs are classified as both unaccusative and unergative by the same diagnostic.[9] For example, *continuare* 'continue' can take both auxiliary *essere* 'be' (like an unaccusative) and auxiliary *avere* 'have' (like an unergative) in the Italian past tense (for the relevance of auxiliary verbs, see (11:3)).

Nevertheless, a substantial body of research has shown that these 'unaccusative mismatches' are problematic only to the extent that one expects unaccusative and unergative verbs to represent syntactically *and* semantically homogeneous classes. Most of the syntactic diagnostics of unaccusativity/unergativity reported in the literature tend to identify semantically coherent subsets of verbs (Levin and Rappaport Hovav 1995). The case of French, however, is a particularly challenging one, as summarized in Chapter 11, based on Legendre 1989 and Legendre, Miyata and Smolensky 1990a, b, 1991. (See Legendre and Sorace in press for a detailed discussion of the relevant empirical facts in French.)

The challenge has long been to identify the syntactically relevant components of meaning in different languages and to develop a theory that could account for the reciprocal syntax/semantics interaction. The principle underlying this endeavor is that neither a verb's ability to be found in the unaccusative or unergative syntactic configuration, nor the verb's particular semantic characteristics, are by themselves sufficient conditions to satisfy particular diagnostics, as explored in the HG account discussed in Chapter 11. A syntactic characterization of unaccusativity in terms of deep Grammatical Relations is necessary to account for phenomena not easily reducible to

[9] See Rosen 1984 for an early discussion of the absence of complete cross-linguistic overlap between a given semantic class and a given syntactic class. See also Legendre and Rood 1992 for a detailed illustration in the Siouan language Lakhóta.

purely semantic explanations, such as the similarity between unaccusatives and passives, the resultative construction in English, auxiliary selection in Italian,[10] and the complex facts of French (see Chapter 11; Legendre 1989).[11] The identification of syntactic constraints, however, is not sufficient; it is also crucial to explain how semantic characteristics (e.g., agentivity) and aspectual properties (e.g., telicity[12]) of individual verbs map to the binary syntactic representations underlying split intransitivity.[13]

Two main proposals have been made in strictly symbolic terms in the last decade or so. One, known as the 'projectionist' approach (see Levin and Rappaport Hovav in press for discussion) maintains that the lexical semantics of a verb *deterministically* specifies the hierarchical classification of its arguments, and that this in turn produces the syntactic behavior associated with unaccusativity or unergativity (Hale and Keyser 1986, 1993; Levin and Rappaport Hovav 1992, 1994, 1995, in press, among others). The most comprehensive account of this type is Levin and Rappaport Hovav 1995's model based on English, in which a small number of linking rules map semantic components of verb meaning (such as 'immediate cause', 'directed change' and 'existence') onto positions at argument structure. Within this approach, verbs with variable behavior have multiple meanings, and therefore multiple semantic representations, each with its own regular argument structure realization.

Alternatives to the projectionist view that have gained ground in recent years are the 'constructional' approaches (see Arad 1998; Borer 1994, 1998; McClure 1995; van Hout 1996, 2000). These models regard unaccusativity and unergativity not as lexical properties of verbs, but rather as clusters of properties derived from the syntactic configurations in which verbs appear, which in turn determine their aspectual interpretation. Since the lexical entry of verbs does not contain any specification of the

[10] In Italian, split intransitivity also manifests itself in participial absolute and reduced relative constructions (both illustrated in Table 11:1 for French), as well as use of the pronoun *ne*. *Ne* is a weak pronoun which stands for the equivalent of 'some' in English. Sorace 1995b, a shows that whether a verb allows *ne* or not systematically varies with whether the verb selects the auxiliary *essere* or *avere*. (Examples from Rosen 1984: 50.)

(i) a. Ne sonno (*essere*) morte tre
of-them were died three
'Three of them died.'

b. *Ne hanno (*avere*) risposto tre
of-them have replied three
'Three of them replied.'

The corresponding French pronoun *en* does not distinguish unaccusatives from unergatives (despite claims to the contrary in the literature). See Legendre and Sorace in press for thorough discussion.

[11] For example, the resultative construction in English is subject to a 'Direct Object Restriction' (Levin and Rappaport Hovav 1995) — it can be predicated only of a direct object NP governed by the verb:

(i) a. John licked his finger clean. (transitive)
b. The bottle broke open. (unaccusative)
c. *John shouted hoarse. (unergative)

[12] Telicity is a temporal property of an event (Box 11:1). A telic event includes achievement of a location, goal, or state as part of its meaning (e.g., *go to the train station:* one hasn't gone there until one is there). An atelic event is unbounded (*draw*, *cry*, etc.). The standard test is to add the temporal phrases *in an hour* and *for an hour. In an hour* is felicitous with telic verbs only while *for an hour* is felicitous with atelic verbs only. A given verb may describe both telic and atelic events, e.g., *draw for an hour* vs. *draw a picture in an hour*. See Dowty 1979 for further discussion.

[13] Various recent theories of argument structure (focused on the syntactically relevant properties of verb arguments) and event structure (focused on the temporal and aspectual organization of the event described by a verb) have set out to pursue this goal (Grimshaw 1990; Pustejovsky and Busa 1995; van Hout 1996; Rappaport Hovav and Levin 1998, among others).

GRs of its arguments, any verb is free to enter into more than one syntactic configuration and consequently to receive multiple aspectual interpretations. Unlike the projectionist model, this approach predicts flexibility in the syntactic realization of arguments, but at the price of insufficient restrictiveness. Constraints preventing overgeneration must thus be present at other levels (e.g., Cummins 1996; van Hout 1996).

2.2. Modeling cross-linguistic gradience in auxiliary selection as a semantic/ aspectual hierarchy

A challenge to both the projectionist and the constructional views has come from two fronts. One is the HG account discussed in Chapter 11, as it demonstrates the need for soft constraints to handle the kind of variation displayed by the split intransitivity phenomenon. The other is a series of empirical studies by Sorace and her collaborators (Sorace 1993a, b, 1995b, a; Keller and Sorace 2000; Sorace and Cennamo 2000; Sorace and Shomura in press); together these establish variation of a particular type as the norm cross-linguistically. The starting point of the latter studies are a set of facts which characterize split intransitivity in a number of Western European languages in two ways: (a) across languages, certain verbs tend to show consistent unaccusative/unergative behavior, whereas others do not; (b) within languages, certain verbs are invariably unaccusative/unergative regardless of context, whereas others exhibit variation. Sorace and colleagues' studies provide supporting evidence for these generalizations, mostly based on experiments testing native speakers' intuitions about auxiliary selection (perhaps the best known diagnostic of unaccusativity) in various languages that have a choice of past tense auxiliaries (such as Dutch, German, Standard Italian, and the Paduan dialect of Italian). In all these languages — and even to some extent in French — unaccusative verbs tend to 'select' the counterpart of English *be* (henceforth 'E' for Italian *essere* — *bE)* while unergative verbs select the counterpart of *have* ('A' for *avere* — *hAve*).[14] However, native intuitions on auxiliaries are categorical and consistent for certain types of verb, but much less determinate for other types. For example, native speakers have a very strong preference for auxiliary E with change-of-location verbs (e.g., 'arrive'), but express a weaker preference for the same auxiliary (or have no preference at all) with stative verbs (e.g., 'exist'). Further discussion of the experimental evidence appears below.

Sorace 2000's account of these systematic differences within the syntactic classes of unaccusative and unergative verbs is that there exist gradient dimensions or hierarchies which distinguish *core* unaccusative and unergative verbs from progressively more *peripheral* verbs. As we now see, these hierarchies, which are based on (potentially universal) event parameters, identify the aspectual notion of *telic dynamic change* as the core of unaccusativity and the semantic property *agentive non-motional activity* as the core of unergativity. The extremes of the hierarchies thus consist of maximally

[14] Legendre 1989, on which the HG account is based, rejected E selection as a *productive* diagnostic test for unaccusativity in French, compared to other tests. In the present chapter it is shown that the non-productivity of E selection in some languages is actually predicted by the derived typology.

distinct core verbs — verbs of change of location (e.g., *arrive*) and verbs of agentive non-motional activity (e.g., *work*) — which reliably display the greatest degree of consistency in auxiliary selection. In contrast, peripheral verb types between the extremes are susceptible to variation.[15]

The overall hierarchy is represented in (16).

(16) **Auxiliary Selection Hierarchy** (*ASH*)

change of location	Selects *be* (least variation)
change of state	
continuation of a pre-existing state	
existence of state	
uncontrolled process	
controlled processes (motional)	
controlled process (non-motional)	Selects *have* (least variation)

Peripheral unaccusative verb types include (arranged in order of closeness to the core): verbs denoting indefinite change in a particular direction (e.g., *rise*), change of condition (e.g., *wilt*), appearance (*appear*), continuation of a pre-existing condition (*stay*) and verbs denoting states (*exist, suffice*). Peripheral verbs closer to the unergative core include verbs denoting motional processes (e.g., *swim*), and various kinds of uncontrolled processes such as bodily functions (*sweat*), involuntary reaction (*tremble*) and emission (*rattle*).[16] The hierarchy in (16) embodies the claim that non-core verbs may receive multiple argument realizations, depending on how they are conceptualized. Thus, these verb classes do not display stable syntactic behavior across languages: they may be unaccusative in some languages and unergative in another. They may also show variable behavior within individual languages, for example by displaying syntactic characteristics of both unaccusative and unergative verbs.

The generalization that has emerged from these studies is that as soon as one moves away from a core one finds substantial but predictable indeterminacy in syntax-semantics mapping with intransitive verbs. This indeterminacy is difficult to accommodate insightfully within a projectionist model of the lexicon-syntax interface, since it would require multiple lexical semantic classifications for a great number of verbs (see van Hout 1996; Rappaport Hovav and Levin 1998 for discussion). It is also challenging for a constructional model, since core verbs display categorical behavior and the other verbs are variable, but to different degrees. For example, several verb classes in Italian allow both auxiliaries, as indicated below by E*/A* in Table 1 (p.

[15] In the proposed analysis, verbs are arrayed along a single scale extending from 'most unaccusative' to 'most unergative' (16). Thus there is potential confusion lurking in the terminology here, since verbs that are 'core' unaccusatives and unergatives lie at the *extremes* of the hierarchy, while verbs that are 'peripheral' unaccusatives and unergatives fall in the *center* of the hierarchy.

[16] The hierarchy does not include intransitive verbs alternating with transitive causative variants (e.g., *break, increase*), which are weakly unaccusative, and in some languages display unergative behavior (see Sorace and Shomura in press; Legendre 1989; Labelle 1992 on French; Haegeman 1994 on English). Nor does it include intransitive verbs which appear with reflexive morphology *se, s'* (e.g., French *s'évanouir* 'to faint', *s'évaporer* 'to evaporate'). See Legendre and Sorace in press for an analysis of reflexive unaccusatives which extends the present OT analysis.

927), but most indeterminate are the stative verbs in the center of the ASH (including verbs of physical and abstract existence, as well as psychological verbs). For verb classes closer to the unaccusative and unergative cores, there is usually a preference for E or A which follows the ASH: verbs closer to the unaccusative core allow E and A but prefer E, while verbs closer to the unergative core allow E and A but prefer A.

The hierarchy in (16) makes it possible to advance some specific typological predictions. Note that it does *not* predict that all languages differentiate among all verb classes, but only that there should not be complete reversals of the hierarchical order of verb types (e.g., languages in which stative verbs are core unaccusatives, or verbs denoting involuntary processes are core unergative). The data on auxiliary selection suggest that within any given language there is a cut-off point between the verbs that select auxiliary *be* and those that select *have*. The cut-off point cannot be identical in all languages, since if it were, all languages with a choice of auxiliaries would have exactly the same system of auxiliary selection. Thus, the locus of variation must be in the mapping governing the interface between the lexicon and the syntax. Mapping must be language-specific because the location of the cut-off points along the hierarchy may be different. However, variation in the location of the cut-off point most affects the verbs in the middle of the hierarchy: crucially, rarely the core.

As mentioned earlier, evidence for gradient variation can be found in a variety of experimental studies in Italian (Sorace 1993a, b, 1995b, a, Bard, Robertson and Sorace 1996) and in Germanic languages. Experiments on Dutch (Sorace and Vonk 1998) show orderly gradience in the judgments of native speakers on *zijn* 'be' and *hebben* 'have', largely corresponding to the intransitive hierarchies identified for Italian. In addition, they show that the acceptability of impersonal passives (a construction traditionally regarded as a diagnostic of unergativity) is affected by semantic factors, particularly agentivity, which cut across the unaccusative-unergative distinction (Zaenen 1993). For German, Keller and Sorace 2000 provide similar findings for native judgments on *sein* 'be' and *haben* 'have', and also show that inter-dialectal variation in auxiliary usage between Northern and Southern varieties is mostly found with peripheral (but not with core) verbs. In French, proper experiments are yet to be conducted, but the results of a pilot study of speaker variation in judgments of unaccusative and unergative verbs (collected by Legendre from five speakers in the same geographical region and age group) suggests that gradient variation is the norm in French as well. In acquisition of Italian as a non-native language, data show that syntactic properties such as auxiliary selection and use of *ne* (see note 10, p. 920) are acquired earliest with core verbs and then gradually extended to more peripheral verb types (Sorace 1993a, 1995b). Moreover, Italian learners of French find it more difficult to acquire *avoir* 'have' as the auxiliary for verbs closer to the core than for more peripheral verbs (Sorace 1993b, 1995a).

A preliminary look at the early acquisition of French verbs by young Grégoire, one of the children studied in Chapter 18, confirms the general findings. In his earliest four files (age 1;9-1;10) the only intransitive verbs Grégoire uses are unaccusative; he produces past tense forms with the correct auxiliary (E) with verbs of location first

— specifically *tomber* 'fall', *monter* 'go up', *partir* 'leave', in that order. The first unergative verbs to show up in the past tense with auxiliary A are controlled motional processes *bouger* 'move (for a person)' (age 2;0) and *rouler* 'move (for a car)' (age 2;3).

To sum up, auxiliary selection displays a gradient sensitivity to the aspectual and semantic properties of individual verbs; this gradience is captured by the Auxiliary Selection Hierarchy. While the ASH is a generalization and not a theory, it challenges existing theories of the syntax-lexicon interface: it cannot be accommodated within a projectionist account because it would entail too much duplication in the lexicon, and it does not fit a constructional account because the amount of variation is related to specific verb types. At the same time, it has features of both accounts: like the projectionist approach, it assumes a systematic relation between the syntax of auxiliary selection and the semantics of individual verbs; like the constructional approach, it allows some (but not all) verbs to have multiple syntactic projections.

The main idea of the HG account is that the overall acceptability (or Harmony) of a given verb in the syntactic context of a given unaccusativity test is the result of the complex interaction of a set of mapping constraints (or linking rules) pertaining to semantic and aspectual properties of the verb, the semantic properties of its argument, and properties of the diagnostic test itself. The rules themselves are very similar in content to Levin and Rappaport Hovav 1995. One important difference, however, is that the mapping constraints in HG are formalized as soft constraints rather than hard ones: they express numerically-weighted *preferences* of a given feature, say telicity, for a mapping to underlying direct object. The numerical values themselves are extracted from the French data by a connectionist learning algorithm in a computational model which encodes the strength of a preference for a particular mapping in the connection between abstract units, some of which represent features like telicity and animacy, others underlying GRs, yet others the individual verbs, and the diagnostic test itself. Finally, the model incorporates the binary deep GR distinction, underlying subject vs. direct object, in the form of hidden units which are assigned values automatically by the connectionist processing algorithm. There is explicit competition between unaccusative and unergative syntactic configurations for each individual combination of verb, argument, and construction. (See Chapter 11 for details.)

Another notable difference with Levin and Rappaport Hovav's linking rule approach is that the set of possible mappings is far richer in the HG account, allowing for a very fine-grained analysis. Finally, the HG analysis establishes the importance of telicity in characterizing unaccusativity (contra Levin and Rappaport Hovav).

However, the main import of the HG account is that it establishes a model of gradiency as syntactic competition driven by violable mapping constraints.[17] It also documents the necessity of a featural semantic and aspectual description of verbs in terms of which the constraints themselves are stated. The constraints do not refer to

[17] Sorace and Keller 2004 investigate the notion of gradience in syntax and elaborate on the difference between hard and soft constraints, providing experimental data from several domains.

pre-defined verb classes per se (unlike say Levin and Rappaport Hovav's linking rules for directed change verbs vs. existence verbs); the mapping preferences pertain to individual features for specific syntactic configurations. In other words, classes of similarly-behaving verbs are *emergent* properties, not analytic premises.

This suggests that the ASH too should be seen as the outcome of grammatical competition. This is indeed what we establish below by demonstrating that the empirical generalization embodied in the ASH is best explained in terms of an existing construct of OT, *harmonic alignment of (prominence) scales* (Prince and Smolensky 1993/2004; Chapter 12 (42)).[18] In fact, the OT account builds on the earlier HG account, specifically, on its featural descriptions. The main difference is that HG exploits numerical optimization (or constraint weighting) while OT exploits non-numerical optimization (based on a domination hierarchy).

The ASH suggests that there are two types of gradience to capture. One is the gradience across languages whereby a different cut-off point on the hierarchy determines which classes are unaccusative in a given language. This is the focus of the OT analysis presented in Section 2.4. The other type of gradience pertains to the fact that verb classes in the middle of the ASH are likely to exhibit some variation in auxiliary choice. We sketch an analysis of such variation in terms of *partial constraint ranking*, an OT construct well established in studies of variation, including synchronic (Nagy and Reynolds 1997), dialectal (Anttila 1997), diachronic (Slade 2003), and developmental (Legendre, Vainikka, Hagstrom and Todorova 2002; Davidson and Goldrick 2003; Davidson and Legendre 2003; Legendre, Hagstrom, Chen-Main, Tao and Smolensky in press; Chapter 18).

2.3. A featural analysis of auxiliary selection in French and Italian

The lexico-semantic and aspectual properties widely implicated in split intransitivity phenomena are telicity, directed change, change of state, motion, displacement (e.g., Van Valin 1990; Dowty 1991; Zaenen 1993; Levin and Rappaport Hovav 1995; Sorace 2000) and homogeneity (McClure 1995). As we shall see, (17) represents the smallest set of features necessary to exhaustively characterize auxiliary selection in Romance.

(17) Event features: INHERENT **DIS**PLACEMENT, HOMOGENEITY, INHERENT **TE**LICITY, **DIR**ECTION, **ST**ATE, INHERENT **VO**LITIONALITY, INTERNAL **MO**TION

Table 1 provides a featural description of each verb class; (17) shows the abbreviations. (A '+' under '–HOM' indicates the feature value [–HOM].)

[18] This was originally proposed in the context of the role of the sonority scale in syllable structure, formally stating that vowels are less marked syllable peaks than non-vowels, sonorant consonants are less marked peaks than obstruents, etc. Prominence alignment has been applied to other phenomena in syntax/semantics: see Artstein 1999; Asudeh 1999 and Chapter 15.

Table 1. Featural composition of French and Italian intransitive verbs

Aux. Sel. Fr.	It.	Features of event → / ↓ Emergent verb classes	DIS	–HOM	TE	DIR	ST	–VO	MO
E	E	Change of location: *arrive*	+	+	+	+	+	+	–
		Change of state							
E	E	a. Change of condition: *die*	–	+	+	+	+	+	–
E*	E	b. Appearance: (*dis*)*appear*	–	–	+	+	+	+	–
		c. Indefinite change in a particular direction:							
E*	E	*go up*	–	–	+	+	+	+	–
A	E	*rot*	–	–	+	+	+	+	–
A	E*	*worsen*	–	–	–	+	+	+	–
		States							
A	E*	a. Continuation of a pre-existing state: *last*	–	–	–	–	+	+	–
A	E*	b. Existence of state: *exist*	–	–	–	–	+	+	–
A	E	*be*	–	–	–	–	+	+	–
		Uncontrolled processes							
A	A*	a. Involuntary actions: *shiver*	–	–	–	–	–	+	–
A	A*	b. Emission: *resound*	–	–	–	–	–	+	–
A	A	c. Bodily functions: *sweat*	–	–	–	–	–	+	–
		Controlled processes							
A	A*	a. Motional: *swim*	–	–	–	–	–	–	+
A	A	b. Non-motional: *work*	–	–	–	–	–	–	–

* = Both auxiliaries are possible (see Section 2.4.5 for further discussion)

With the understanding that the roles of most features have been established in the relevant literature, we proceed with some selective discussion of this choice of features. Among these is INHERENT **TE**LICITY determined on the basis of diagnostic tests like occurrence with the adverbial phrase *for an hour/in an hour* (see footnote 12). On the basis of these tests change of location verbs like French *arriver* 'arrive' are telic (18) while controlled processes like French *travailler* 'work' are atelic (19).

(18) a. Pierre est arrivé chez lui en 1 h. TELIC
'Peter arrived home in one hour'

b. Pierre a pris/mis 1 h pour arriver chez lui.
'Peter took one hour to arrive home'

c. *Pierre est arrivé pendant 1 h.
Peter arrived for one hour'

(19) a. *Pierre a travaillé en 3 h. ATELIC
‘Peter worked in three hours’

b. *Pierre a pris/mis 3 h pour travailler.
‘Peter took three hours to work’

c. Pierre a travaillé pendant 3 h.
‘Peter worked for three hours’

A less-traveled feature is that of *HOMOGENEITY* or the extent to which all subintervals of an event are identical (McClure 1995). Events as different as controlled activities *travailler* ‘work’ and change of state *pourrir* ‘rot’ are homogenous: each sub-event of an on-going event of working or rotting entails that the referent has worked or rotted at least a bit (20). In contrast, an on-going event of dying (since it canonically takes time to die) is not made of homogeneous sub-events: at any point of getting closer to the endpoint one cannot say that the individual has died (even a bit).

(20) Event homogeneity (French examples)

a. *travailler* ‘work’

Jean est en train de travailler	⇒	il a travaillé
'John is (in the process of) working'	⇒	‘he has worked’

b. *pourrir* ‘rot’

Ta pomme est en train de pourrir	⇒	elle a pourri un peu
‘Your apple is rotting’	⇒	‘it has rotted a bit’

c. *mourir* ‘die’

Jean est en train de mourir	⇏	il est mort (un peu)
‘John is dying’	⇏	‘he has died (a bit)’

We need to distinguish change of state verbs like *die* from *rot* by at least one feature value because they select different auxiliaries in French. In the absence of the feature HOMOGENEITY, the featural profile of both verbs is identical, as shown in Table 1. We propose that the difference in auxiliary choice within this traditional verb class is due to the fact that *die* is non-homogeneous while *rot* is. Processes of rotting, dying, appearing or disappearing are all gradual processes. What distinguishes them is whether the gradual process is incrementally homogeneous or not.

Another important distinction to be made is between verbs which connote inherent displacement from point A to point B (e.g., *aller à* ‘go to’ and more generally change of location verbs) and verbs which connote internal motion (e.g., *nager* ‘swim’, *courir* ‘run’). Note that these types of motion can occur without displacement as revealed in the common expression *nager/courir sur place* ‘swim/run in place’. In Table 1 the feature INTERNAL MOTION distinguishes among controlled (volitional) processes e.g., *swim* vs. *yell*.

Finally INHERENT **VO**LITIONALITY distinguishes uncontrolled from controlled processes at the bottom of the hierarchy. Change of location and condition verbs (e.g., *venir* ‘come’, *mourir* ‘die’) are not inherently volitional. Volitionality, when present, is a property of their argument.

The remaining features in Table 1 are **DIR**ECTION and **ST**ATE. The former distinguishes change of location and state verbs from the rest because an important component of their meaning is 'directed change' and the latter characterizes verbs whose meaning includes being in or reaching a state (including location at some point).

Using these features and their appropriate values we obtain a set-inclusion hierarchy, as represented in Table 1. What is crucial here is that verb classes that select a different auxiliary be distinguished by at least one feature value and that these feature values express implicational relations (+DIS implies –HOM; –HOM implies +TE; +TE implies +DIR, etc.) This is true of the first six features, counting from the left. This distribution enables us to propose next an OT analysis grounded in a Power Hierarchy (Section 2.4.2) whose universal scope does not rely on any further stipulation.

2.4. OT analysis

2.4.1. Harmonic Alignment

One important outcome of much typological-functional research is the recognized existence of markedness relations (e.g., Jakobson 1965/1995; Croft 1990) which express favored associations in languages of the world.

Scales including 'animacy' (Local person > Pronoun 3rd > Human 3rd > etc.) and thematic properties (Agent > Patient) have been associated with the well-known GR hierarchy in (21)[19] to express markedness relations (e.g., Silvertein 1976; Keenan and Comrie 1977; Perlmutter 1983; Aissen 2001). In similar vein, we may formulate event scales for the features relevant to the A/E auxiliary distinction. (In fact, the telicity scale, atelic > telic (22c), is adopted by Grimshaw 1990.)

(21) GR scale: 1 (Subject) > 2 (Object)

(22) Event feature scales

a. displacement: –DIS > +DIS
b. homogeneity: +HOM > –HOM
c. telicity: –TE > +TE
d. directed change: –DIR > +DIR
e. state: –ST > +ST
f. inherent volitionality: +VO > –VO
g. internal motion: –MO > +MO

By aligning two scales at a time we come up with a set of relations which express the markedness of the mapping of a certain feature — say TELIC — to a certain GR — say 1. Such *harmonic alignments*, as they are defined in OT, formalize markedness rela-

[19] For example, languages which allow relativization of a direct object and languages that allow null objects are a subset of the languages which respectively allow relativization of a subject and null subjects (Keenan and Comrie 1977).

tions for mappings between certain properties *across* scales (12:42).[20] Note the change of symbol from '>' (higher on a scale) to '≻' (more harmonic — less marked) in the sample harmonic alignments in (23).

(23) Harmonic alignments

- a. 2/telic ≻ 1/telic — The mapping of [+telic] onto an unaccusative configuration (underlying 2) is less marked than the mapping of [+telic] onto an unergative configuration (underlying 1)
- b. 1/atelic ≻ 2/atelic
- c. 2/telic ≻ 2/atelic
- d. 1/atelic ≻ 1/telic
- e. etc.

Such alignments correspond to a hierarchy of constraints with polarity reversed (note again the change in symbol from '≻' (more harmonic) to '≫' (more dominant): (24).

(24) Constraint alignments:

- a. *1/telic ≫ *2/telic — 'don't map [+telic] onto an unergative configuration' outranks 'don't map [+telic] onto an unaccusative configuration'
- b. *2/atelic ≫ *1/atelic
- c. *2/atelic ≫ *2/telic
- d. *1/telic ≫ *1/atelic
- e. etc.

For the present we drop the lower-ranked constraints in each of these alignments, which target the unmarked mappings (see Chapter 12, note 27). Focusing first on those remaining constraints pertaining to GR1, we have {*1/telic ≡ *1/[+TE], *1/[+DIR], *1/[–VO], …}. We now encapsulate all these constraints.

2.4.2. Formulating a *1 Power Hierarchy

Consider again Table 1. Down to feature [–VOLITIONAL] Table 1 expresses implicational relations among feature values. We can thus define a set $\mathcal{C}$ of '2-preferring' feature values (25a) and state an encapsulated constraint *1/$\mathcal{C}$ which is violated each time a constraint in {*1/f | f ∈ $\mathcal{C}$} is violated.

(25)
- a. $\mathcal{C}$ ≡ {+DIS, –HOM, +TE, +DIR, +ST, –VO} ('2-preferring' feature values)
- b. F ≡ *1/$\mathcal{C}$ — An event with a $\mathcal{C}$-feature is not mapped to an unergative configuration

The fact that a candidate violating F = *1/$\mathcal{C}$ six time s is more marked than one violating it twice is implemented via a standard OT *Power Hierarchy* (Smolensky 1995;

[20] Alignment in OT in fact formalizes the concept of "mirroring" in statements like "the deep syntactic encoding mirrors the thematic hierarchy in markedness". See the Universal Alignment Hypothesis (UAH, Perlmutter 1978; Rosen 1984), also known as UTAH (Baker 1988) — a well-known principle governing the mapping between thematic roles and their (underlying) syntactic instantiation.

Legendre, Smolensky and Wilson 1998; the case here is formally identical to that of Section 14:6). The power hierarchy is given in (26a); the constraint F^k is violated whenever F is violated *k* (or more) times.

(26) Universal Mapping Constraint Hierarchy

a. $F^6 \gg \cdots \gg F^2 \gg F^1$ (GR/event semantics mapping)

b. *1/+DIS ≫ *1/−HOM ≫ *1/+TE ≫ *1/+DIR ≫ *1 /+ST ≫ *1/−VO

Given the implicational relations among features in Table 1., the Power Hierarchy can be written in the equivalent form (26b). For example, any candidate violating F at least 3 times must violate the constraints {*1/+DIR, *1/+ST, *1/−VO}: the feature sets yielding at least 3 violations are exactly {±DIS, ±HOM, ±TE; +DIR, +ST, −VO}. Indeed, any feature set violating *1/+DIR will necessarily also violate {*1/+ST, *1/−VO} so it will violate F at least 3 times. F^3 can thus be equated with *1/+DIR.[21]

As shown in Table 2, it is more marked for *arriver* 'arrive' to be assigned an unergative configuration than for *suer* 'sweat' because *arriver* 'arrive' violates constraint F 6 times while *suer* violates it only once.

Table 2: Markedness as determined by the *1 Power Hierarchy

Power Hierarchy:	$(*1/\mathcal{C})^6$	$(*1/\mathcal{C})^5$	$(*1/\mathcal{C})^4$	$(*1/\mathcal{C})^3$	$(*1/\mathcal{C})^2$	$(*1/\mathcal{C})$
Constraint ranking:	*1/+DIS	*1/−HOM	*1/+TE	*1/+DIR	*1/+ST	*1/−VO
1/*arriver* 'arrive' +DIS, −HOM, +TE, +DIR, +ST, −VO	*	*	*	*	*	*
1/*suer* 'sweat' −DIS, +HOM, −TE, −DIR, −ST, −VO						*
1/*nager* 'swim' −DIS, +HOM, −TE, −DIR, −ST, +VO						

Recall that an additional feature is needed to distinguish motional from non-motional controlled processes: *nager* 'swim' vs. *travailler* 'work'. The feature proposed is INTERNAL **MO**TION. Because all verb classes above controlled motional processes have the value –MO as does the lowest class on the ASH that includes *travailler*, the feature MO does not stand in an implicational relation with all other features. The constraint *1/+MO is therefore not part of the Power Hierarchy proper. In fact, Legendre to appear demonstrates that German auxiliary selection provides independent empirical

[21] It is worth entertaining the idea that the constraint hierarchy in (26b) derives in fact from the alignment of the scale –VO > +ST > +DIR > +TE > –HOM > +DIS with the GF scale 1 > 2 yielding two harmonic alignments: (i) 1/–VO ≻ 1/+ST ≻ 1/+DIR ≻ 1/+TE ≻ 1/–HOM ≻ 1/+DIS and (ii) 2/+DIS ≻ 2/–HOM ≻ 2/+TE ≻ 2/+DIR ≻ 2/+ST ≻ 2/–VO. In (26b) we are referring to only one of the two constraint hierarchies this alignment yields, namely the *1/*feature* hierarchy.

evidence for placing MO outside of the Power Hierarchy proper.

In Romance the constraint *1/+MO must be ranked below the lowest constraint in the Power Hierarchy to express the relative markedness of assigning GR2 to the *nager* vs. *travailler* classes. Diachronic evidence of auxiliary change in Spanish discussed in Legendre to appear supports the conclusion that 'work' is "more unergative" than 'swim' ; moreover selecting E is not completely rejected by native speakers in the case of 'swim'. With GR1, 'swim' violates *1/+MO, while 'work' violates no *1 constraint. Combining the *1 Power Hierarchy and *1/MO gives the *1 constraint hierarchy.

(27) *1 constraint hierarchy

*1/+DIS ≫ *1/–HOM ≫ *1/+TE ≫ *1/+DIR ≫ *1/+ST ≫ *1/–VO ≫ *1/+MO

Obviously, if no constraint against mapping onto GR2 ever entered the picture, selection of 2 (auxiliary E) would always be optimal, and we would have no cross-linguistic mismatches and no unergative verbs. While the harmonic alignments stated in (24b) result in a hierarchy of *2 constraints anti-parallel to that of *1 constraints (24a), for the purposes of this discussion and the formal results we seek, it suffices to state the *2 constraint in the encapsulated version given in (28).[22] Further refinements may be called for in future work.[23]

(28) *2: "Don't map onto an unaccusative configuration"

We propose that across languages this *2 constraint slides along the hierarchy of the harmonic alignment constraints in (27), resulting in a cut-off point which is cross-linguistically variable.

2.4.3. Establishing the language-particular ranking of *2

It should be clear by now that the cross-linguistic choice of auxiliary for a given featural profile will be determined by the relative ranking of *2 —just where *2 is interposed into the fixed *1 hierarchy.

We take the input to optimization to be the featural description of individual predicates or predicate subclasses, as specified in Table 1. The candidate set simply consists of two candidates corresponding to the two assigned configurations — underlying 2 and 1 — assumed to be directly mapped to E and A, respectively.[24]

A sample of tableaux is provided next which highlights the crucial optimizations responsible for the language-particular rankings of *2 stated to be derived in (33). (For further exemplification, see Legendre and Sorace 2003.)

[22] This formulation is independent of which specific structural terms this configuration might be stated in: AgrOP, etc.
[23] A fuller analysis would consider whether the constraint in (28) — *2 — should be replaced with a harmonic alignment constraint *2/E, derived from the alignment of the (underlying) GF prominence scale 1 > 2 (21) with A > E (*have* > *be*). Harmonic alignment would entail that 1/A is more harmonic than (≻) 1/E and 2/E ≻ 2/A. In turn the following constraint rankings obtain: *2/A ≫ *2/E and *1/E ≫ *1/A. As a result, the optimization would involve four candidate pairings of GF and auxiliary. But the formal result is the same with the simpler constraint in (28).
[24] This is the consequence of assuming that auxiliary selection is one reflex of the larger unaccusative/unergative distinction to be stated in syntactic terms.

(29) *pourrir* 'rot'

Input: –DIS, +HOM, +TE, +DIR, +ST, –VO, –MO

French	*1/DIS	*1/–HOM	***2**	*1/TE	*1/DIR	*1/ST	*1/–VO	*1/MO
a. 2/E			*!					
b. ☞ 1/A				*	*	*	*	

Despite the fact that it is telic, change-of-state verb *pourrir* selects A. This motivates the French sub-ranking *2 ≫ *1/TE, shown in (29). But *mourir* selects E, despite the fact that it is also a telic change-of-state verb. This is where the constraint on mapping the feature HOMOGENEITY comes in: *1/–HOM ≫ *2 in French, as shown in (30).

(30) *mourir* 'die'

Input: –DIS, –HOM, +TE, +DIR, +ST, –VO, –MO

French	*1/DIS	*1/–HOM	***2**	*1/TE	*1/DIR	*1/ST	*1/–VO	*1/MO
a. ☞ 2/E			*					
b. 1/A		*!		*	*	*	*	

In Italian, verbs of state select E. This entails that *2 must be outranked by *1/ST in Italian: see (31). In contrast, [–ST] verbs select A. Thus *2 ≫ *1/–VO: see (32).

(31) *esistere* 'exist'

Input: –DIS, +HOM, –TE, –DIR, +ST, –VO, –MO

Italian	*1/DIS	*1/–HOM	*1/TE	*1/DIR	*1/ST	***2**	*1/–VO	*1/MO
a. ☞ 2/E						*		
b. 1/A					*!		*	

(32) *sudare* 'sweat'

Input: –DIS, +HOM, –TE, –DIR, –ST, –VO, –MO

Italian	*1/DIS	*1/–HOM	*1/TE	*1/DIR	*1/ST	***2**	*1/–VO	*1/MO
a. 2/E						*!		
b. ☞ 1/A							*	

These crucial cases result in the language-particular constraint rankings stated in (33).

(33) a. French: *2 ↓ (between *1/–HOM and *1/+TE)

*1/DIS ≫ *1/–HOM ≫ *1/+TE ≫ *1/+DIR ≫ *1/+ST ≫ *1/–VO ≫ *1/MO

b. Italian: *2 ↑ (between *1/+ST and *1/–VO)

2.4.4. Verifying predictions in French, Italian, and beyond

For Italian the position of *2 in the *1 hierarchy entails that verbs that have any of the feature values [–DIS, –HOM, +TE, +DIR] select E. For example, *peggiorare* 'worsen' selects

E in Italian, as expected; see (34). Its French counterpart, however, is predicted to select A because *2 outranks *1/DIR in French; see (35).

(34) *peggiorare* 'worsen'

Input: –DIS, +HOM, –TE, +DIR, +ST, –VO, –MO

Italian	*1/DIS	*1/–HOM	*1/TE	*1/DIR	*1/ST	***2**	*1/–VO	*1/MO
a. ☞ 2/E						*		
b. 1/A				*!	*		*	

(35) *empirer* 'worsen'

Input: –DIS, +HOM, –TE, +DIR, +ST, –VO, –MO

French	*1/DIS	*1/–HOM	***2**	*1/TE	*1/DIR	*1/ST	*1/–VO	*1/MO
a. 2/E			*!					
b. ☞ 1/A					*	*	*	

In both languages, controlled processes (motional and non-motional) are predicted to select A. That is because only *2 and *1/MO are activated in the case of motional processes (36) and only *2 in the case of non-motional processes (e.g. *travailler/lavorare* 'work'). Given that *2 ≫ *1/MO in both languages, A is the preferred option.

(36) *nager* 'swim'

Input: –DIS, +HOM, –TE, –DIR, –ST, +VO, +MO

French	*1/DIS	*1/–HOM	***2**	*1/TE	*1/DIR	*1/ST	*1/–VO	*1/MO
a. 2/E			*!					
b. ☞ 1/A								*

Summing up: Auxiliary selection results from the competition of two Mapping Hierarchies: Power Hierarchies of mapping constraints themselves derived from harmonic alignment of simple scales referring to lexico-semantic and aspectual features and syntactic configuration. This entails that mapping rules cannot be stated in terms of verb classes, contra Levin and Rappaport Hovav 1995.

Significantly, the proposed general OT analysis does *not* predict total, unconstrained variation in auxiliary selection. Rather, it predicts a very specific typology of languages including languages in which all verb classes are syntactically unaccusative, languages in which all verb classes are unergative, and languages which each display one of a tightly limited set of splits.

First of all, it predicts that some languages do not show any split effects in auxiliary selection. In other words, languages in which all verb classes select E and languages in which all verb classes select A are predicted to exist. The latter formally result from *2 dominating all *1 constraints and the former from *1 constraints outranking *2. Both are found within Romance languages: Spanish only has A and the central Italian dialect Terracinese (Tuttle 1986) only has E.

Another correct prediction made by the present analysis is that further languages have different cut-offs along the universal hierarchy. Besides Standard Italian and its low cut-off point and Standard French with its high cut-off, we find Dutch and German with a threshold somewhere in between. In both these Germanic languages, change of location and change of state verb classes select E while the remaining verb classes — continuation of a pre-existing state, existence of state, uncontrolled processes, as well as controlled processes — select A.

Diachronically, there is evidence from studies on the historical development of auxiliaries in Romance (e.g., Benzing 1931; Tuttle 1986) showing that core verb types tend to be the last to be affected by the replacement of *esse*-reflexes by *habere*-reflexes, whereas peripheral verb types are the most vulnerable to the change. A remarkable example is provided by Spanish through the course of its history, as described in Aranovich 2003. In Modern Spanish all verb classes take A. In Old Spanish, only core unergative verbs like *trabajar* 'work' and *pecar* 'sin' occurred with A. Change from E to A started with the peripheral classes as predicted by our analysis. The first to go were verbs of manner of motion like *errar* 'wander' and verbs of existence of state *rastar* 'remain' (XIV century). Next to change were 'dynamic verbs of existence and appearance' (*aparecer* 'appear', *desaparecer* 'disappear', etc.) in the XV century. *Morir* 'die' and *ir* 'go' were the last ones to give up E (XVII century).

Our OT analysis also predicts some languages to be impossible. For example, there couldn't be a language where existence-of-state verbs select E but change-of-state verbs select A. As far as we know, this prediction is correct.

2.4.5. Partial constraint ranking

The OT analysis above captures gradience across languages whereby a different cut-off point on the hierarchy determines which classes of verbs are unaccusative in each language. As we saw in Section 2.2, there is another type of gradience whereby verbs in the middle of the ASH are peripherally unaccusative or unergative in the sense that they are likely to display a certain amount of fluctuation in their choice of auxiliary. For example, verbs of state in Italian and verbs of appearance in French may select either E or A (see Table 1). Can this type of gradience be derived from the present model? The answer is yes, provided the analysis is supplemented by *partial ranking*, i.e., some indeterminacy in the relative ranking of *2 and the *1/f constraints. By definition, a partial constraint ranking yields a set of rankings (e.g., Anttila 1997; Boersma 1997; Legendre et al. 2002; Slade 2003). This set of rankings yields potentially different optimal outputs, hence variation in outputs. See Chapter 18 for an illustration in the domain of early acquisition of syntax.

In the interest of space we do not provide an actual analysis of such gradience in auxiliary selection. Rather we provide an illustration of the mechanism in question, leaving a full analysis for future work. Consider, for example, the consequences if *2 were to float over four positions at the top of the *1/f hierarchy (27) we would obtain a set of four rankings for the class of [+DIS, –HOM, +TE] verbs, as represented in (38).

(For convenience, the lowest portion of the hierarchy, *1/−VO ≫ *1/+MO, is omitted.)

(37) Partial ranking with floating *2 constraint

Fixed: *1/DIS ≫ *1/−HOM ≫ *1/TE ≫ *1/DIR ≫ *1/ST

Floating: ⟵ ***2** ⟶

(38) Corresponding total rankings

a. ***2** ≫ *1/DIS ≫ *1/−HOM ≫ *1/TE ≫ *1/DIR ≫ *1/ST
b. *1/DIS ≫ ***2** ≫ *1/−HOM ≫ *1/TE ≫ *1/DIR ≫ *1/ST
c. *1/DIS ≫ *1/−HOM ≫ ***2** ≫ *1/TE ≫ *1/DIR ≫ *1/ST
d. *1/DIS ≫ *1/−HOM ≫ *1/TE ≫ ***2** ≫ *1/DIR ≫ *1/ST

Different rankings yield different proportions of verbs selecting a particular auxiliary. Specifically, verbs that are [+DIS] (hence also [−HOM, +TE]) are unaccusative in rankings (38b, c, d) so 75% of the time, assuming all four total rankings to be equally probable (see Chapter 18). Verbs that are [−DIS, −HOM] (hence [+TE]) are unaccusative for rankings (38c, d), so 50%; verbs that are [−DIS, +HOM, +TE] are unaccusative only 25% of the time — ranking (38d) only. And verbs that are [−TE] (hence [−DIS, +HOM]) are unergative 100% of the time. So the more extreme features +DIS or −TE are more homogenously unaccusative or unergative; the middling feature −HOM (hence [−DIS, +TE]) waffles 50/50 between the two. Put another way, there is more indeterminacy in the middle of the range than at the extremes.

2.5. Concluding remarks: the larger perspective

Summing up, we have proposed that the ASH derives from harmonic alignment of simple scales referring to semantic and aspectual features and syntactic configuration. In other words, verb classes like 'change of state', etc. (see the vertical axis in Table 1) have no theoretical status in our OT analysis. They are emergent classes. Yet, they serve the important function of making explicit how, given a set of constraints stated on relatively fine-grained semantic and aspectual features, and given a (typically binary) choice between two auxiliaries, E or A are alternatively selected, albeit differently in the languages forming the focus of this case study, Italian and French.

Another property of the OT analysis worth emphasizing is that the constraints are the same in different languages. What varies is the position of a single constraint (*2) relative to all others in the hierarchy. Thus, variation results from different interactions of the same set of mapping constraints. This stands in contrast with an OT analysis like Bentley and Eyrthorsson 2002 which is grounded in the ASH but posits different mapping rules in different languages.

While further confirmation of the predictions of the analysis awaits much further study, we hope these initial results provide new evidence from syntax/semantics for the utility of OT as a formal theory of typology; more specifically, we hope to have shown that the typological properties of split intransitivity (e.g., the existence of systematic cross-linguistic mismatches) receive a straightforward and enlightening account in terms of existing explanatory mechanisms in OT.

In Chapter 11, Harmonic Grammar was used to analyze the complex French pattern of language-internal unaccusativity mismatches as delicate quantitative interactions of general and cross-linguistically-attested constraints. But HG is not and was never intended to be a theory of typology. HG and OT provide deeply-related but complementary instruments for characterizing at different scales human knowledge at the syntax/semantics interface. They both find their place in the overall ICS cognitive architecture. If — as we believe a half-century of work in generative linguistics has demonstrated — typologies are real and robust, then theories of the mind must, among other things, wrestle with this fundamental reality. OT is the component of ICS grappling with this challenge, and the early results reported here give some significant basis for optimism.

3. A CASE STUDY IN PHONOLOGY: SYLLABIFICATION IN BERBER

In this section a remarkable phonological system is first analyzed within Optimality Theory; next, this OT account is transformed into an account in Harmonic Grammar; and finally the HG account is implemented in a local connectionist network. Thus this case study illustrates the general ICS theoretical reduction leading from OT to neural computation, a reduction in which HG plays a crucial bridging role. There are a number of ways in which this demonstration is incomplete or unsatisfactory. The OT account involves only two constraints, although one of them is arguably an encapsulation of an entire hierarchy of eight constraints. The structure imposed by the grammar during parsing is very simple: certain segments are marked as syllable nuclei, and that's it; there is no epenthesis or deletion, nor elaborate structure-building. The connectionist network employs local not distributed representations. Its temporal behavior is not (as yet) formally analyzed, so a proof of correctness, and calculation of processing times, are not (currently) available; it is only by computer simulation that the network's computational behavior is studied. Nonetheless, means of partly addressing all these shortcomings have been developed within this book.

3.1. Introduction to Berber

Dell and Elmedlaoui 1985 present a spectacular analysis of syllabification in the Imdlawn Tashlhiyt dialect of the language Berber (spoken in the Atlas mountains of North Africa; henceforth 'Berber' will refer to this particular dialect). Just as with the ***m*** and ***r*** of English *pris**m*** and *B**er**ber*, in this dialect of Berber, consonants can be syllable nuclei. But whereas English allows only the most vowel-like — the most *sonorous* — consonants to be nuclei, in Berber *any* consonant can be a syllable nucleus. Berber contains words like *tftkt* and *txẓnt* in which all segments have low sonority— yet Berber words are sequences of syllables just as in any other language, and the syllables are governed by universal principles. Since any segment can be parsed into any syllable position, the number of possible syllabifications of Berber words grows very quickly (exponentially) as the words increase in size.[25]

Analysis of syllabification in Berber is the starting point of Prince and Smolensky 1993/2004, and a theme returned to throughout that book as an illustration of many of the central concepts of Optimality Theory. As we will see in this section, Berber also provides an excellent case study for illustrating the relationship between OT and Harmonic Grammar. For this purpose we accept the approximation to Berber syllabification adopted in Prince and Smolensky 1993/2004, which abstracts away from certain complications, primarily at phrase edges, that are beside the present point. (See Clements 1997 for a fuller OT analysis of the complexity of Berber syllabification.)

3.1.1. Sonority: Nuclear Harmony, HNUC

The key idea is that the universal ideal for syllables is a certain *sonority profile*: starting at a low level, sonority rises quickly in the onset to a peak in the nucleus, and then drops gently in the coda (e.g., Clements 1990; see Chapter 14, Section 6). Sonority is an abstract grammaticization of the inherent phonetic prominence of segments: the low vowel *a* has the greatest sonority, while voiceless stops like *t* and *k* have the lowest sonority. The sonority of a segment ρ will be written *son*(ρ). While many aspects of the sonority scale are clearly universal, there are (at least apparent) minor variations across languages; these will not be dealt with here. Dell and Elmedlaoui argue that the sonority scale in Berber is as indicated in (39).

(39) Sonority scale

Segment Class	Ex. Segment ρ	Sonority *son*(ρ)
voiceless stops	*t, k*	1
voiced stops	*d, b, g*	2
voiceless fricatives	*s, f, x*	3
voiced fricatives	*z, γ*	4
nasals	*n, m*	5
liquids	*l, r*	6
high vocoids	*i/y, u/w*	7
low vowel	*a*	8

The high vocoids are *i/y* and *u/w*: when syllabified as a nucleus, such a segment is pronounced as a vowel (*i* or *u*), and when syllabified as a syllable margin — onset or coda — these same segments are pronounced as a glide (*y* or *w*). To aid the reader, in examples, the sonority level of each segment will often be indicated as a subscript.

The central point is that, despite their oddity, syllables in Berber respect the uni-

[25] Consider 6 segments; they can be fully syllabified in two ways: .CV.CV.CV. or .CVC.CVC. (see text below). So a string of $6k$ segments has at least 2^k syllabifications (roughly, then, a string of n segments has at least $2^{n/6}$ syllabifications). This seriously undercounts, since there are many syllabifications that cross the boundaries of the *k* substrings of 6 segments each, and we have counted only those that don't. Further, the syllabifications counted all satisfy ONSET and *COMPLEX (or *VV and *CCC) so if these constraints are external to *Gen*, then many more syllabifications are also candidates.

versal ideal for a sonority profile. While in languages like English only the most sonorous segments can be in the nucleus, in Berber syllable nuclei contain *the most sonorous segments possible*, given the segments provided by the word. In the word $t_1r_6g_2l_6t_1$ 'you locked,' the most sonorous segments — the liquids r_6 and l_6 — are the nuclei: the correct syllabification is $.t_1\acute{r}_6.g_2\acute{l}_6t_1.$, where, as in Chapter 13, periods mark syllable edges, and the acute accent ´ marks syllable nuclei. In another convenient notation, the syllable structure of $.t\acute{r}.g\acute{l}t.$ will be written .CV.CVC., using 'V' to denote any segment that has been syllabified into a nucleus (not necessarily a vowel) and 'C' any segment that has been syllabified into a syllable margin. The CV-string .CV.CVC. will then be called a *parse* of the segment string *trglt*.

The universal constraint demanding that syllable nuclei be sonorous will be formulated below: it is called HNUC (Harmony of the NUCleus). Interacting with HNUC is the ONSET constraint requiring syllables to have onsets, familiar from Chapter 13.

It is evident that ONSET has priority over HNUC in words like $t_1x_3z_4n_5t_1$ 'you stored', which is syllabified $.t_1\acute{x}_3.z_4\acute{n}_5t_1.$ Conforming to HNUC, the most sonorous segment in the word, n_5, is parsed as a nucleus. But the next-most sonorous segment, z_4 (a voiced coronal fricative), is passed over, and a lower-sonority segment, x_3 (a voiceless velar fricative), is selected as the other nucleus. As Dell and Elmedlaoui explain, this is because the syllable headed by *n* must have an onset; this forces *z* to be parsed as an onset, which takes it out of the running for the second nucleus. Of the remaining segments, *x* is the most sonorous, and it is selected.

3.1.2. Syllable structure

Syllables in Berber consist of an obligatory onset containing a single segment, an obligatory one-segment nucleus, and an optional one-segment coda. (In the notation of Chapter 13, Berber belongs to the typological class $\Sigma^{CV(C)}$.) The requirement that syllables must have a nucleus is suspended for the first syllable of the word (or phrase); because this subtlety has no bearing on the topic of this section — the connection between OT and HG — we will simply pretend that the ONSET constraint includes within it the waiver for initial syllables:

(40) (quasi-)ONSET (*VV)

A non-initial syllable must have an onset.

If a non-initial V is preceded by a C, that C can be parsed as an onset to satisfy ONSET. So the force of this constraint is to ban a non-initial V (i.e., nucleus) preceded by another V; that is, it bans *hiatus*: two adjacent syllable nuclei. The banned structure is simply VV, so it will sometimes be convenient to call this constraint '*VV'.

Note that, given the syllable structure requirements of Berber, it is always possible to locate syllable boundaries from a string of Cs and Vs: each non-initial V must be preceded by a single onset C. So before each CV there must be a syllable boundary, '.'. Thus, for example, VCCVCV must denote the syllabification

(41) .VC.**CV**.**CV**.

In this parse, each non-initial V, in bold, determines the '.' before the preceding C; and of course the parse must also begin and end with a syllable edge '.'. Because the location of syllable boundaries is determined by the CV-string, the '.' markings will often be omitted below.

That onsets and codas in Berber are limited to a single segment can be seen as the effect of a universal constraint introduced in Chapter 13: *COMPLEX, which prohibits syllable positions (onset, nucleus, or coda) which are 'complex' in the sense of containing multiple segments. That *COMPLEX is unviolated in Berber entails that a legal parse will never contain a sequence CCC: the middle C is not syllabifiable.

(42) *COMPLEX (*CCC)

No more than one segment may occupy a single syllable position.

Note that any string of Cs and Vs satisfying both the constraints *VV and *CCC is a legal syllabification in Berber. Such a CV string will be called a *legal* parse; for any given input string of segments, the optimal syllabification will be a legal parse. Any string of Cs and Vs at all — not necessarily satisfying *VV or *CCC — will be called a *potential* parse.

3.2. The Dell-Elmedlaoui syllabification algorithm: DE

Given an input like $t_1x_3z_4n_5t_1$, how can a syllabification be found that meets the demand of HNUC to have maximally sonorous nuclei, while respecting the overriding demand of ONSET? Dell and Elmedlaoui propose a syllabification algorithm that works as follows.

Start with the highest sonority level: 8, the sonority of *a*. Scan the string from left to right, looking for unsyllabified segments with this sonority level. Suppose one is found: call the segment ρ. Then look to see if there is an unsyllabified segment to its left (this is waived for the initial segment). If so, ρ is a *trigger*. Then, parse the trigger ρ as a syllable nucleus V and parse the preceding segment as a syllable onset C. Continue to the end of the string.

Then drop down one level in the sonority hierarchy: 7, the sonority of the high vocoids *i*/*y* and *u*/*w*. Follow the same procedure, scanning from left to right for an unsyllabified segment with sonority 7, parsing it as a nucleus if there is an unsyllabified segment to its left to parse as an onset.

Then continue down the sonority hierarchy until the lowest level has been scanned: 1, the sonority of voiceless stops like *t* and *k*.

At this point there may remain unsyllabified segments; parse each as a coda. The result is a complete parse, the output of the parsing algorithm DE.

As an illustration, $t_1x_3z_4n_5t_1$ is parsed as in (43). We see that at step (43e), parsing the most sonorous segment in the word, *n*, as a nucleus commits the parser to syllabify the preceding segment *z* as an onset, even though it is more sonorous than the segment to its left, *x*, which will then be syllabified as a nucleus in step (43f).

(43) Dell-Elmedlaoui algorithm (DE) parsing $t_1x_3z_4n_5t_1$

a.	$t_1\ x_3\ z_4\ n_5\ t_1$	(input)
b.	$t_1\ x_3\ z_4\ n_5\ t_1$	*no segments at sonority level 8*
c.	$t_1\ x_3\ z_4\ n_5\ t_1$	*no segments at sonority level 7*
d.	$t_1\ x_3\ z_4\ n_5\ t_1$	*no segments at sonority level 6*
e.	$t_1\ x_3$ .z$_4$ **ń**$_5$ t_1 .C V	on scan at sonority level 5 *current syllabification*
f.	.t$_1$ **x́**$_3$ .z$_4$ **ń**$_5$ t_1 .C V .C V	on scan at sonority level 3 *current syllabification*
g.	.t$_1$ **x́**$_3$ .z$_4$ **ń**$_5$ t$_1$ .C V .C V C.	on final step: parse codas (output)

The left-to-right scanning employed in DE is used to disambiguate parsing when two adjacent segments have the same sonority, e.g., *mn*. When that level of sonority is scanned, both segments are triggers: both are potentially available to become syllable nuclei — but both cannot be, since that would create a VV configuration in violation of ONSET = *VV. The left-to-right scan disambiguates the situation by selecting the leftmost segment, *m*, for the nucleus. If the rightmost segment *n* were instead selected, the preceding *m* would serve as the onset of a syllable .m**ń.** This has a poor sonority profile, as there is no sonority rise at all between onset and nucleus. Dell and Elmedlaoui themselves describe the left-to-right scanning as an indirect means of achieving a goal that is not actually concerned with directionality: maximizing the sonority difference between onset and nucleus (Dell and Elmedlaoui 1985: 127 n. 22). Since, again, this complication is beside the present point, we will restrict attention to inputs that *have no sonority plateaux*. That is, the sonority values of adjacent segments are never identical. For such inputs, the are no ambiguities of the sort the left-to-right scanning of the DE was designed to resolve, and directional scanning is unnecessary. When a given sonority level *s* is reached by the algorithm, all unparsed segments meeting the criterion (sonority = *s* with an unparsed preceding segment) can be parsed as nuclei, without conflict.

It will prove useful below to modify the DE algorithm slightly. Consider the point in DE when a segment α is parsed as a V, and the one preceding it as a C. Now consider the segment β following this new nucleus V. If β has been parsed already, it must necessarily have been parsed as a C: for had it previously been parsed as a V, the segment preceding it — α — would necessarily have been parsed as a C at the same time, which is impossible because α is now being parsed as a V. So if β has already been parsed, it has been parsed as a C. If β has not yet been parsed, we can say with confidence that it will later be parsed as a C: for to be parsed as a V, the segment preceding it, α, would have to be available to serve as an onset, which it is not since it has just been parsed as a nucleus. So in any event, β will end up parsed as a C.

In the modified algorithm, when α is parsed as a V, the segment β following it is also parsed as a C. Instead of inserting CV into the parse, the algorithm now inserts

CVC, the final C perhaps already having been parsed. (The second C is of course omitted if α happens to be the final segment, and no following segment β exists.) Henceforth, 'DE' will refer to this modified algorithm, which necessarily produces the same output as the original Dell-Elmedlaoui algorithm. The advantage is that in the modified algorithm, to qualify as a trigger at sonority level *s*, it suffices that a segment of that sonority simply be unparsed; there is no need to check its left neighbor to see if it too is unparsed and therefore free to serve as an onset. For if each step of parsing enters a sequence CVC into the parse, any unparsed segment γ must have to its left either a C or an unparsed segment; either way, γ qualifies as a trigger.

Since the DE algorithm constructs its output by inserting CVC strings, it follows that at most two Cs can end up adjacent to one another. And of course, by construction, DE never generates two adjacent Vs. Thus:

(44) The output of DE satisfies ∗CCC = ∗COMPLEX (42) and ∗VV = ONSET (40).

3.3. Positive OT

Several Optimality Theoretic analyses of Berber syllabification were developed in Prince and Smolensky 1993/2004. The one most intuitively close to the DE algorithm will be called the *positive* analysis, OT_+. In this formulation, $HNUC_+$ *rewards* good syllable *nuclei*; in the negative analysis discussed below ($OT_–$), $HNUC_+$ is replaced by constraints that *penalize* bad syllable *margins*.

(45) $HNUC_+$ (Nuclear Harmony, positive formulation)

If *son*(v) > *son*(τ), the nucleus $\acute{v}$ has higher Harmony than the nucleus $\acute{\tau}$.

When comparing two parses, compare first the highest-Harmony nucleus of each. If they have equal Harmony, discard them and continue to compare the highest-Harmony remaining nucleus. Continue until a nucleus in one parse has higher Harmony than its counterpart in the other parse; then declare the parse with the higher Harmony nucleus to be the one best satisfying $HNUC_+$. If all nuclei of one parse are completely exhausted before those of another parse, declare the latter to better satisfy $HNUC_+$.

The positive OT analysis is now simply:

(46) OT_+: ONSET ≫ $HNUC_+$

We posit that PARSE, FILL, IDENT are undominated, and so consider only candidates that satisfy these constraints.

(47) OT_+ tableau for *txznt*

$/t_1x_3z_4n_5t_1/$		ONSET	$HNUC_+$	*comment*
☞ a.	$.t_1\acute{x}_3.z_4\acute{n}_5t_1.$		$\acute{n}_5$ $\acute{x}_3$	*DE output*
b.	$.\acute{t}_1.x_3\acute{z}_4.\acute{n}_5t_1.$	∗!	$\acute{n}_5$ $\acute{z}_4$ $\acute{t}_1$	$\acute{z}_4 \succ x_3$ *but too late*
c.	$.\acute{t}_1.x_3\acute{z}_4.n_5\acute{t}_1.$		$\acute{z}_4$! $\acute{t}_1$ $\acute{t}_1$	$\acute{n}_5 \succ \acute{z}_4$
d.	$.\acute{t}_1x_3.z_4\acute{n}_5t_1.$		$\acute{n}_5$ $\acute{t}_1$!	$\acute{x}_3 \succ \acute{t}_1$

As an illustration of this analysis, (47) shows a constraint tableau for *txznt*. This tableau will be discussed in the context of theorem (48), one of the earliest results derived in OT.

(48) *Theorem: opt*(OT₊) = *out*(DE) (Prince and Smolensky 1993/2004)

The optimal parse as defined by OT_+ equals the output of the DE algorithm, for any input with no sonority plateaux.

Note: Subsequent theorems of this type — (49), (58), (59), (60), (63), (66) — all pertain, like (48), to any input with no sonority plateaux.

In the HNUC₊ column of tableau (47), $\acute{z}_4$ denotes the violation of $HNUC_+$ by the syllable nucleus containing z; this is a less serious violation of $HNUC_+$ than $\acute{x}_3$ because *son*(z) = 4 > *son*(x) = 3; i.e., $\acute{z}$ is a more harmonic nucleus than $\acute{x}$: $\acute{z}_4 \succ \acute{x}_3$.

Proof: Consider an input (say $t_1x_3z_4n_5t_1$) and let P be the corresponding output of DE. Is there any competing parse Q that has higher Harmony, on the OT analysis? Since P satisfies ONSET, any competitor Q that does not satisfy ONSET will lose to P when the higher constraint, ONSET, is evaluated (47b). So we need only consider competitors that satisfy ONSET, and consider evaluation by $HNUC_+$. Now suppose s is the highest sonority level present in the input (s = 5 for $t_1x_3z_4n_5t_1$). Segments of sonority s (n_5) will be parsed as nuclei in P (on the first scanning step of DE that finds any triggers). Consider a competing parse Q that does not have these segments parsed as nuclei. Then Q will lose to P when the highest-Harmony nuclei of P and Q are compared, because the highest-Harmony nuclei of P have higher sonority hence higher Harmony than any alternative nuclei (47c). So we need only consider competitors that have the same highest-sonority nuclei as P. Continuing this same logic, consider the next sonority level at which DE finds triggers (3: x_3). Any competitor Q that does not, like P, parse these triggers as nuclei will lose when the second-highest-Harmony nuclei in P and Q are compared by $HNUC_+$ (47d). So we need consider only competitors with the same second-highest-sonority nuclei as P. And so on to the bottom of the sonority hierarchy. When the lowest-sonority nuclei of P are reached, the logic entails that the only parse Q that does not have lower Harmony than P is a parse in which all nuclei are located at the same segments as the nuclei of P; that is, Q = P itself.

Note that the method employed within the $HNUC_+$ constraint for handling multiple syllables itself employs the strict domination characteristic of OT: if the highest-sonority nuclei of two parses differ, this will determine which is more harmonic, preempting all consideration of lower-sonority segments. The general strict domination mechanism of OT will be explicitly employed to achieve this in the second OT analysis presented below in Section 3.6.

(49) *Corollary.* OT_+ outputs satisfy *CCC (*COMPLEX) (42) and *VV (ONSET) (40).

Proof. By (49), the maximal-Harmony parse w.r.t. OT_+ is the output of the DE algorithm; such outputs always satisfy *CCC and *VV by (44).

Note that the ranking ONSET ≫ $HNUC_+$ already achieves the effects of *COMPLEX, without requiring an explicit *COMPLEX constraint: maximal satisfaction of $HNUC_+$ is achieved by having as many nuclei as possible, and ONSET restricts those possibilities to those satisfying *VV. This means the sub-parse CCC can never appear in an optimal parse, because HNUC$_+$ would be better-satisfied by replacing CCC with CVC, and such a move cannot run afoul of ONSET. This has the convenient consequence that the Harmonic Grammar developed next need not explicitly incorporate rules implementing *COMPLEX; it is sufficient to implement ONSET ≫ HNUC$_+$.

3.4. Positive HG

We are finally ready to develop a Harmonic Grammar — HG$_+$ — comparable to the OT$_+$ analysis. It is easy enough to state ONSET and HNUC$_+$ in HG terms.

(50) HG$_+$: Harmonic Grammar for Berber syllabification

a. ONSET A non-initial syllable must have an onset. (Negative)
Strength: $W_{ONS} = 2^M$

b. HNUC$_+$ Segment ρ must be a syllable nucleus. (Positive)
Strength: $W_\rho = 2^{son(\rho)} - 1$

ONSET *penalizes* a parse for each V that is preceded by another V; the Harmony penalty equals the strength W_{ONS}. HNUC$_+$ *rewards* a parse for each nucleus; the quantity of Harmony added for each nucleus depends on the sonority of the nuclear segment: higher sonority yields higher Harmony contribution.

The remaining question is: what should the strengths be? The answer is given in (50): the HNUC$_+$ contribution from a nuclear segment ρ is $2^{son(\rho)} - 1 \equiv W_\rho$, where *son*(ρ) is the sonority level of ρ on the numerical scale given in (39). Letting *M* denote the maximal sonority value (8 = *son*(*a*)), the largest possible contribution from a single nucleus is $2^M - 1 = W_a$. The strength of ONSET is larger, by one: $2^M \equiv W_{ONS}$; the HG strength relation $W_{ONS} > W_a$ corresponds to the OT$_+$ ranking ONSET ≫ HNUC$_+$. The task now is to derive these values.

The modified Dell-Elmedlaoui algorithm DE successively finds maximal-sonority unsyllabified segments, and parses them as a V flanked by a C on either side: CVC gets added into the parse, with the V at the trigger, contributing positive Harmony via HNUC$_+$. The *cost* of this move is that the two flanking segments, being assigned C, are no longer available as potential nuclei, and therefore no longer able to add positive Nuclear Harmony by being parsed as a V. In the worst case, the two flanking segments parsed C are just one level of sonority lower than the trigger. The Harmony contributions must be arranged so that a single positive contribution from parsing the trigger as V must be greater than the highest possible cost for the move, which is twice the Harmony contribution from one level down the sonority scale:

(51) Exponential growth of strength with sonority condition
$W_s > 2W_{s-1}$

Let us arbitrarily set the lowest HNUC$_+$ reward — for sonority level 1 — at 1:

(52) $W_1 \equiv 1$

Then (51) requires:

(53) $W_2 > 2W_1 = 2$

so let us choose

(54) $W_2 = 3$

Repeating the logic, we get $W_3 = 7$, $W_4 = 15$, etc. The result is:

(55) $W_s = 2^s - 1$.

The Harmony function for HNUC₊ defined by HG₊ (50) is:

(56) $$H_{\text{NUC+}}(\text{P}) = \sum_i nuc(\rho^i, \text{P})[2^{son(\rho^i)} - 1]$$

where $\{\rho^i\}$ are the segments in the parse $\text{P} = \rho^1\rho^2\cdots$, and $nuc(\rho^i, \text{P}) = 1$ if P's i^{th} segment ρ^i is parsed as a nucleus in P, and 0 otherwise. Because this is a simple sum of contributions from individual segments, H_{NUC} obeys (57).

(57) If the CV-parse P = LR is the concatenation of two sub-parses L and R, then

$$H_{\text{NUC}}(\text{LR}) = H_{\text{NUC}}(\text{L}) + H_{\text{NUC}}(\text{R})$$

This is useful in proving the next result:

(58) *Theorem*. $opt(\text{HG}_+) = out(\text{DE})$

For a given input string I, consisting of N segments with no sonority plateaux, let the candidate set $Gen(I)$ be the set of all potential parses of I by a CV-string of length N. Within $Gen(I)$, the maximum-Harmony parse according to HG_+ (50) is the same as the output of the DE algorithm given the input I.

For the proof of this result, it is useful to first establish (59).

(59) *Lemma*. A maximum-Harmony parse of HG_+ cannot violate ONSET = *VV.

Proof. Consider any parse Q that violates *VV. Intuitively: re-parse the first V as C, getting a competing parse P. Changing V to C eliminates one nucleus but also the ONSET violation; the Harmony increases by the strength of ONSET and decreases by the strength of HNUC for the re-parsed segment. But since the former is defined to be greater than the latter (corresponding to OT_+'s ONSET ≫ HNUC), there is an overall gain in Harmony; P is more harmonic than Q. Therefore Q can't be a Harmony maximum. This is a *harmonic bounding* argument: P harmonically bounds Q (Section 12:1.6).

More formally: Express Q as $\text{Q} = \text{LV}^1\text{V}^2\text{R}$ where L and R are respectively the portions of the string Q to the left and right of VV. Now let $\text{P} = \text{LC}^1\text{V}^2\text{R}$; this parse eliminates the violation at V^2 of *VV in Q. It may even eliminate a second violation of *VV in Q, if there is a violation at V^1 as well (i.e., if L ends in V); any other violations of *VV that there may happen to be within L or R are shared by P and Q, and do not differentiate their Harmony with respect to ONSET, denoted by H_{ONS}. Thus, P has at least one fewer violation of *VV than Q, so:

$$H_{\text{ONS}}(\text{P}) - H_{\text{ONS}}(\text{Q}) \geq W_{\text{ONS}}$$

Call the segments corresponding to V^1 and V^2 in the parse Q ρ^1 and ρ^2. Let V/ρ^1 denote the segment ρ^1 parsed as a nucleus (as in parse Q); let C/ρ^1 be ρ^1 parsed as a margin (as in P). Using (57), evaluate the difference between P and Q with respect to $H_{\text{HNUC+}}$, the Harmony arising from HNUC₊.

$$H_{\text{HNUC+}}(\text{P}) - H_{\text{HNUC+}}(\text{Q}) = [H_{\text{HNUC+}}(\text{L}) + H_{\text{HNUC+}}(\text{C}/\rho^1) + H_{\text{HNUC+}}(\text{V}/\rho^2) + H_{\text{HNUC+}}(\text{R})] - [H_{\text{HNUC+}}(\text{L}) + H_{\text{HNUC+}}(\text{V}/\rho^1) + H_{\text{HNUC+}}(\text{V}/\rho^2) + H_{\text{HNUC+}}(\text{R})] = -H_{\text{HNUC+}}(\text{V}^1) = -\text{W}_\rho$$

So P has higher Harmony than Q w.r.t. ONSET, by an amount at least W_{ONS}, while P has lower Harmony than Q w.r.t. HNUC_+, by an amount $\text{W}_{\text{\`{}}\text{,}1}$ that depends on the sonority of the segment ρ^1 which is nuclear in Q but not P. Now W_{ONS}, the strength of ONSET in HG_+ (50), is defined to be greater than all HNUC_+ strengths W_ρ, even that of the most harmonic nucleus, W_a. Thus P has higher total Harmony. In sum, for any parse Q of *I* that violates *VV, there is another parse P of *I* with higher Harmony that satisfies *VV. Thus the maximum-Harmony parse cannot violate *VV.

Now we can return to theorem (58).

Proof. Lemma (59) establishes that any maximal-Harmony parse for HG_+ satisfies ONSET, just like any output of the DE algorithm, so we now only need consider $H_{\text{HNUC+}}$, the Harmony due to HNUC_+. Given any input *I* with no sonority plateaux, let that parse of *I* that maximizes $H_{\text{HNUC+}}$ be denoted P†. It is now useful to view the DE algorithm as successively winnowing a candidate set: the initial set consists of all potential syllabic parses of *I*, and each time the algorithm parses a segment of *I*, the candidate set is reduced by discarding all syllabifications not consistent with the newly assigned structure. With no sonority plateaux, the DE algorithm is deterministic and terminates on a unique complete parse, all other candidates having been eliminated.

Let's jump into the operation of the DE algorithm anywhere during its operation. DE has just found a trigger ν at some sonority level *s*, and assigns CVC to the substring μνρ consisting of ν and its immediate neighbors μ and ρ. This winnows the candidate set; is it possible that the maximal-$H_{\text{HNUC+}}$ parse P† has been erroneously eliminated? There are a few different cases to consider. Exceptions associated with the left and right edges of the string *I* are obvious and will not be stated explicitly.

Suppose the segment μ in the critical substring μνρ has already been parsed; if so, it must be parsed as C since its neighbor ν is a trigger and hence previously unparsed: no unparsed segment can be adjacent to V since the (modified) DE algorithm imposes CVC whenever it parses a segment. If μ is already parsed, then the remaining candidates all agree that μ is parsed as a C and that part of the 'assignment' of CVC on μνρ has no effect at all.

Only if μ has not been previously parsed does imposing C on μ (as part of imposing CVC on μνρ) winnow the candidate set. But if μ has not been previously parsed, its sonority must be lower than the current level *s*; for in our version of the DE algorithm, having sonority higher than *s* and being unparsed would necessarily have qualified μ as a trigger at an earlier step of the algorithm. Now in the worst case, the sonority of μ is just one level lower than *s*. Imposing C on μ eliminates parses that assign V to μ and thereby earn an $H_{\text{HNUC+}}$ reward of W_{1ss-11}.

Inverting left and right in the previous argument shows that the segment ρ on the right of the trigger ν is either already parsed C and irrelevant, or it is unparsed, has a lower sonority than *s*, and assigning it C winnows candidates which parse ρ as V and receive an $H_{\text{HNUC+}}$ reward no more than W_{s-1}.

In the worst case, when *both* μ and ρ are unparsed and have sonority $s - 1$, imposing CVC on μνρ eliminates candidates with an HNUC_+ reward from μνρ of at most 2W_{s-1}. The candidates retained receive a benefit of W_s from μνρ under the parse CVC, since the sonority of ν is *s*. Now the strengths have been specifically designed so that for all *s*, $\text{W}_s > 2\text{W}_{s-1}$, as demanded by the exponential growth condition (51); thus the winnowed candidates have lower $H_{\text{HNUC+}}$ values than the retained candidates. Win-

> nowing in this fashion can never exclude the $H_{\text{HNUC}+}$-maximizing parse P†. And since the algorithm must terminate on a single output, this parse must be P†.
> P† maximizes $H_{\text{HNUC}+}$ and satisfies ONSET, so it maximizes total HG_+-Harmony H.

Combining (48) and (58) gives:

(60) *Corollary*. $opt(HG_+) = opt(OT_+)$

Intuitively, the connection between the HG and OT accounts here is quite transparent. It is notable however that the strengths needed to implement the OT account in HG do *not* exhibit the full exponential growth that is expected in the general (or perhaps the worst) case discussed in Section 1.2.2.1. To achieve the effect ONSET ≫ HNUC_+, it is not necessary to make W_{ONS} exponentially larger than W_{HNUC}: it suffices for the former to be even slightly greater than the latter. Within HNUC_+ itself, the strict domination character of cross-sonority-level interactions *does* lead to exponentially growing constraint strengths W_s. But here again the growth is not as severe as predicted by the general analysis of Section 1.2.2.1. There, the strengths grew as C^k, where k is the rank of the constraint and C is the maximum number of possible violations of a single constraint in a candidate. In HG_+, the base of exponentiation is fixed at only 2, even though the number of possible constraint violations grows without bound as the length of the input grows. Thus, if parse P of I has a *single* nucleus v that is more sonorous than all the nuclei in parse Q of I, then Q cannot maximize HNUC_+, even if *all* the nuclei in Q have the highest possible sonority value less than v. Since the number of such nuclei can grow without bound as the length of the input I grows, it might have been expected that the strength W_s for sonority s would have to exceed that of W_{s-1} by an unboundedly large factor.[26]

3.5. Negative OT

The 'positive' OT_+ analysis introduced in Section 3.3 is the one first introduced in Prince and Smolensky 1993/2004: Ch. 2, but the central constraint HNUC_+ does not fit the mold of 'standard' OT in which constraints *penalize* violation rather than *reward* satisfaction. The following 'negative' reformulation of HNUC as a *universal constraint subhierarchy* is developed in Prince and Smolensky 1993/2004: Ch. 8 (this hierarchy is introduced in Box 13:1 and derived in Section 14:6). Switching from positive to negative forces the focus of the constraint to shift from nucleus to margin, as we now see.

(61) HNUC_-: The universal margin subhierarchy

$$*C/v_1 \gg *C/v_2 \gg \cdots \gg *C/v_7 \gg *C/v_8$$

universally, where the constraint $*C/v_s$ is defined by:

[26] It is not necessarily the case that the relative Harmonies of all *sub-optimal* parses are the same in OT_+ and HG_+. Consider the hypothetical input *imrtrtrt*··· *rt* (with n repetitions of *rt*) and the two parses P = *.ím.rt́ .rt́ .rt́ .···.rt́* . and Q = *.yḿr.tŕ .tŕ .···.tŕ t.* In OT_+, HNUC_+ prefers P because its most sonorous nucleus, $í_7$, is more sonorous than Q's most sonorous nucleus $ḿ_5$, even though on the remainder of the string P has n miserable nuclei $t́_1$ while Q has $n-1$ respectable nuclei $ŕ_6$. In OT_+, P ≻ Q, but in HG_+, $H(P) = W_7 + nW_1 < H(Q) = W_5 + (n-1)W_6$ provided $n > 3$. But this does not contradict Theorem (60), because neither P nor Q is optimal; the OT_+-optimal parse is ☞P = *.í.mŕ .tŕ .tŕ .···.tŕ t.* and in HG_+, H(☞P) = $W_7 + nW_6$ which is greater than both $H(P)$ and $H(Q)$, for any m.

*C/v_s: A segment v_s of sonority *s* is not parsed as a syllable margin.

This version of HNUC is best understood via an example. The tableau (62) is the OT_- counterpart to the OT_+ tableau (47).

(62) OT_- tableau for *txznt*

		HNUC₋					
/$t_1x_3z_4n_5t_1$/	ONSET	…	*C/v_5	*C/v_4	*C/v_3	*C/v_2	*C/v_1
☞ a. .$t_1\acute{x}_3$.$z_4\acute{n}_5t_1$.				*			**
b. .$\acute{t}_1$.$x_3\acute{z}_4$.$\acute{n}_5t_1$.	*!				*		*
c. .$\acute{t}_1$.$x_3\acute{z}_4$.$n_5\acute{t}_1$.			*!		*		
d. .$\acute{t}_1x_3$.$z_4\acute{n}_5t_1$.				*	*!		*

Here, the strict domination structure that had been buried within the positive formulation of HNUC is made explicit by the hierarchy. As in OT_- (and DE), what must be attended to first are the highest-sonority segments. In the positive formulation, this meant that the demand that high-sonority segments be nuclei had to be evaluated before consideration is given to lower-sonority segments. In the negative formulation, the dominant consideration — that which must be demanded of high-sonority segments — is that they *not be margins* (i.e., they must be nuclei). The margin subhierarchy places highest the prohibition of putting sonority-8 segments into the margin; next-highest, sonority-7; and so on down to the lowest-ranked constraint, which bans sonority-1 segments from margins. Higher-ranked ONSET will of course prevail in requiring that *some* segments must be parsed into syllable onsets; the ranking in HNUC ensures that these segments will be those of lowest sonority. And because this is the negative formulation, the margin hierarchy is not demanding good margins — it is banning bad ones. All of this depends on the fact that in Berber, there is no epenthesis or deletion; every segment is parsed as either C or V.

(63) *Theorem*. *opt*(OT_-) = *opt*(OT_+).

The relative Harmonies of any two parses of a given input as evaluated by OT_+ and OT_- are identical. Thus these grammars define the same optimal outputs.

Proof. (It may be helpful to refer to the corresponding OT_+ and OT_- tableaux for *txznt,* (47) and (62), while following this argument.) ONSET is identical in the two analyses, so it suffices to consider evaluation by HNUC. Given an input *I*, consider two candidate parses P and Q. Consider the highest-sonority segments in *I*, which have sonority value *s*. Either they are all parsed identically in P and Q, in which case HNUC does not differentiate P and Q w.r.t these segments, or at least one such segment μ is parsed differently; say, as a V in P, and as a C in Q. In OT_+, P is then declared more harmonic than Q because there is a best nucleus $\acute{\mu}$ in P that is lacking in Q. In OT_-, P is also declared more harmonic than Q because the highest-ranked constraint in the margin hierarchy that is relevant to this input, *C/v_s, rejects Q for having at least one more violation (at μ) than P. This establishes the equivalence of OT_+ and OT_- unless all segments in *I* of maximum sonority (level *s*) are parsed identically in P and Q. In this case, proceed to consider the next-most-sonorous segments in ψ. By repeating exactly the same argument, OT_+ and OT_- will declare the same candidate to be more harmonic unless all these segments too are parsed identically. Continuing this reasoning as long as necessary, the conclusion is that the relative Harmony of P and Q will be decided identically by OT_+ and OT_- unless *all* segments are parsed identically in P and Q, in which case P = Q and both OT_+ and OT_- declare their Harmonies equal.

3.6. Negative HG

In the positive HG account, positive Harmony is awarded for satisfaction of HNUC+. In the negative account, negative Harmony is assessed for violations of HNUC–. The following Harmonic Grammar implements the negative OT analysis OT_- (61); here all constraints are negative: the punish violation but do not reward satisfaction.

(64) HG_-: Negative Harmonic Grammar for Berber syllabification

a. ONSET: A non-initial syllable must have an onset.
Strength: $s_{ONS} = 2^M$

b. HNUC–: Segment ρ must not be a syllable margin.
Strength: $s_\rho = 2^{son(\rho)} - 1$

The Harmony function evaluating ONSET is as before. But now the Harmony function evaluating HNUC is:

(65) $$H_{NUC-}(\mathrm{P}) = -\sum_i mar(\rho^i, \mathrm{P})[2^{son(\rho^i)} - 1]$$

where $mar(\rho_i, \mathrm{P})$ is 1 if the i^{th} segment, ρ_i, is parsed as a margin in P, and 0 otherwise.

(66) *Theorem*. $H_{HG-} = H_{HG+} + k$

For any given input *I*, the Harmony assigned by HG_- is equal to the Harmony assigned by HG_+ up to a constant *k* (which depends on *I*). Thus, for any given input, the two grammars declare the same parse to be optimal.

Proof. It suffices to consider only HNUC, since HG_+ and HG_- have the same other constraint, ONSET, with the same strength. Now every segment ρ^i is either parsed V or C in a parse P, so $mar(\rho^i, \mathrm{P}) = 1 - nuc(\rho^i, \mathrm{P})$. Thus, recalling (56), we have:

$$\begin{aligned} H_{NUC-}(\mathrm{P}) &= -\sum_i mar(\rho^i, \mathrm{P})[2^{son(\rho^i)} - 1] \\ &= -\sum_i [1 - nuc(\rho^i, \mathrm{P})][2^{son(\rho^i)} - 1] \\ &= -\sum_i [2^{son(\rho^i)} - 1] + \sum_i nuc(\rho^i, \mathrm{P})[2^{son(\rho^i)} - 1] \\ &= k + H_{NUC+}(\mathrm{P}) \end{aligned}$$

where $k \equiv -\sum_i [2^{son(\rho^i)} - 1]$; this is a constant for any given input.

3.7. Connectionist implementation

3.7.1. BrbrNet

The HG_{-} analysis can readily be translated into a local connectionist network: *BrbrNet* is shown in Figure 1. In this network each input segment is encoded in the activity of one input unit. The leftmost stack of input units — the lower, smaller disks — represents the first segment of the input. Since Berber syllabification depends only on sonority level, we can simplify the picture by having only one unit for each sonority level. For the example input /tbia/ shown in Figure 1, in the first stack of small disks the lowest unit, labeled *t*, is active while the remaining units in the stack are inactive. 'Active' and 'inactive' refer to activation values of 1 and 0 respectively; in the Figure 1, these are indicated by black or white shading, respectively. This same input unit would be active if the first segment of the input were *p* or any other segment with the same sonority level as that of its label *t*. Thus in the second stack of input units the active unit is labeled *d* although the input segment is *b*. The network is given an input by clamping on the input units the pattern of activity representing the input string.

The network computes a syllabification which is represented by a pattern of activity across the output units — the higher, larger disks. The first output unit encodes the syllabic role of the first segment: it is active (1) if the segment is parsed as a syllable peak (nucleus, 'V'), inactive (0) otherwise (margin, 'C'). Each output unit represents the syllabic position assigned to the corresponding input segment (which is represented in the stack of input units directly below the given output unit).

Among the segments parsed as syllable margins, the onset/coda distinction is not explicitly represented in the output: both types of margins are encoded by activity 0. The output representation could be enriched to explicitly make this distinction, but there is no point in doing so here. As pointed out in Section 3.1.2, given the positions of the peaks in a legal Berber syllabification, it is trivial to deduce the onset/coda status of each non-nuclear segment *X*: if *X* immediately precedes a peak, it is an onset, otherwise it is a coda. It would be trivial to add a second layer of output units encoding the onset/coda distinction, with activity 1 encoding onset and 0 coda; a negative bias on each unit would ensure that a unit is inactive unless it receives positive activation from the following peak unit. Berber is a language in which to study decisions not between onset and coda, but between peak and margin, because of its unique characteristic that any segment can, in the appropriate context, be a peak.

Ultimately, a version of BrbrNet employing distributed rather than local representations is needed to instantiate the principles of ICS, but that requires further progress, as noted in Section 2:8.

The following analysis will address Berber inputs with no sonority plateaux: no two adjacent input segments have the same sonority. This is because for such inputs, there is a uniquely defined correct syllabification. With sonority plateaux, ambiguities arise about which of two equally sonorous segments should be parsed as a nucleus, and the resulting complexities will not be taken up here. (So, with n segments, instead of all 8^n possible input sonority sequences, this analysis addresses only $8 \cdot 7^{n-1}$ inputs; e.g., for $n = 5$, only 19,208.)

Figure 1. BrbrNet: Syllabifier for Imdlawn Tashlhiyt Berber

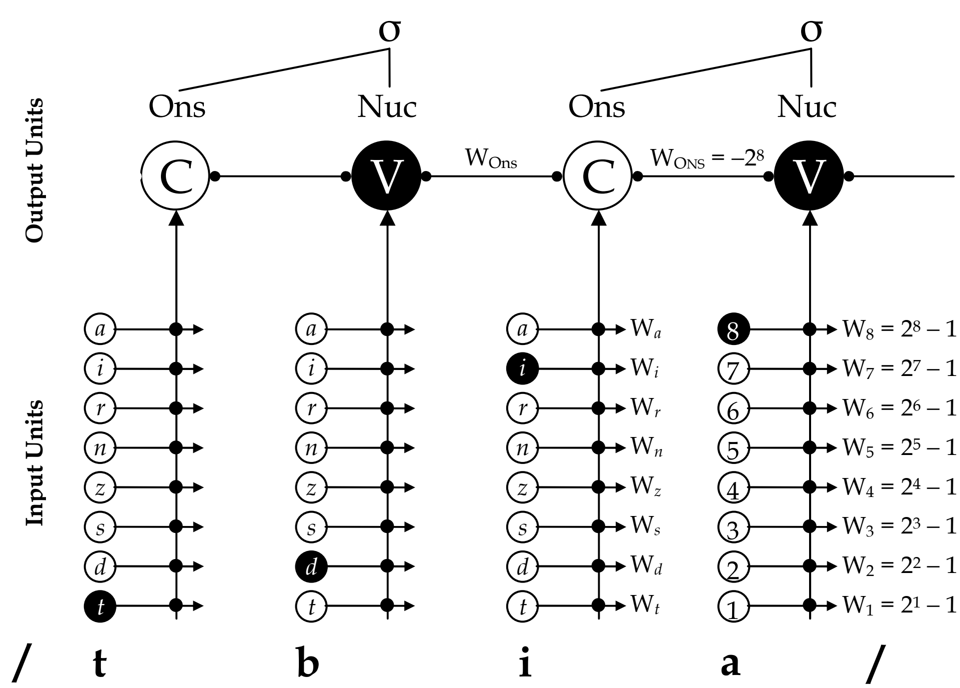

3.7.2. HG₊ ↔ Net architecture

The connections in BrbrNet realize the Harmonic Grammar for Berber called HG_+ (Section 3.4). This grammar is repeated in (67).

(67) HG_+: Harmonic Grammar for Berber syllabification

a. ONSET A non-initial syllable must have an onset. (Negative)
Strength: $W_{ONS} = 2^M$

b. HNUC₊ Segment v must be a syllable nucleus. (Positive)
Strength: $W_v = 2^{son(v)} - 1$

The ONSET constraint is implemented by inhibitory connections of strength 2^M between adjacent output units. When two adjacent output units α and β are active, there are two adjacent syllable peaks, which incurs an ONSET violation, changing the

Harmony by $\Delta H = a_\alpha W_{\alpha\beta} a_\beta = (1)(-2^M)(1) = -2^M$, just as required by (67a). (Recall that the Harmony contribution made by a connection is the weight of the connection times the activation values of the two units it connects: Section 6:1.)

The HNUC$_+$ constraint is realized by excitatory connections from input to output units. An input unit labeled with a segment v has a connection only to its corresponding output unit (directly above it). The strength of this connection is $W_v = 2^{son(v)} - 1$, *son*(v) being the sonority level of v, on a scale 1, 2, …, 8. In Figure 1 the rightmost stack of units is labeled by sonority values, and the corresponding weight values are shown explicitly. The weight pattern is the same for each stack of input units.

When an output unit α is active, a peak has been placed at that segment. In conformity with (67b), the resulting Harmony contributed by the input/output connections is $\Delta H = \Sigma_\beta\, a_\alpha W_{\alpha\beta} a_\beta = (1)(W_v)(1) = 2^{son(v)} - 1$, where v is the input segment in the position of α. That is: all input connections to α come from the input stack beneath it; no other segments are relevant. Suppose the input segment at the position of α is v. Then in the input stack below α, every unit β is inactive except β = v. The inactive units are all connected to α, but they contribute zero Harmony since their activation level is 0. When α is active, the only non-zero Harmony from the input units is $a_\alpha W_{\alpha v} a_v = (1)(W_v)(1)$, which is $2^{son(v)} - 1$ by the design of the connection weights.

In short, the output-output inhibitory connections implement the negative constraint ONSET while the input-output excitatory connections implement the positive constraint HNUC$_+$: the connection weights are the strengths of these constraints. The Harmony of a pattern of activity **a** as computed by the usual connectionist formula $\Sigma_{\alpha\beta}\, a_\alpha W_{\alpha\beta} a_\beta$ is the same as the Harmony of the syllabification represented by **a**, as computed by the Harmonic Grammar (67). At least, when **a** is a pattern of activity that *does* represent a syllabification, this equality holds. These are the activity patterns in which at most one input unit per stack has non-zero activity, and all network units have activity either 0 or 1. These are special network states, the ones with a symbolic semantic interpretation as a string of segments, each parsed as syllable nucleus or margin. For all other activity patterns, the connectionist formula $\Sigma_{\alpha\beta}\, a_\alpha W_{\alpha\beta} a_\beta$ gives the correct Harmony value, but since this pattern corresponds to no symbolic structure, it makes no sense to ask for the Harmony assigned by the Harmonic Grammar.

BrbrNet bears strong resemblance to the syllabification networks of Goldsmith 1992, but there are critical differences with major implications for general theoretical properties, and the correctness of the network's performance. Following the terminology of Prince 1993, where the behavior of these networks is solved analytically, Goldsmith's proposal will be called the *Dynamical Linear Model*, DLM. DLM's output units represent basically the same information as BrbrNet's, and in both networks each output unit is connected only to its neighbors. But unlike DLM, in BrbrNet these connections are necessarily symmetric: as strong from right to left as the reverse. This is necessary for the network to perform Harmony optimization (Chapter 9). The input units of DLM also correspond to those of BrbrNet, after a simple transformation which turns BrbrNet into BrbrNet′, shown in Figure 2. In BrbrNet, the first input segment v is represented by a unit of activity 1, connected to its corresponding out-

Figure 1

put unit with weight $2^{son(v)}$–1. From the point of view of Harmony computation, a completely equivalent arrangement would have a single input unit for the first segment, connected to its corresponding output unit with weight 1, and activity value $2^{son(v)}$–1; since the activation value gets multiplied by the weight in determining the Harmony value, the results are just the same. This new scheme is used in BrbrNet′. In BrbrNet′, there is a single input unit for the first segment, and we encode the sonority of this segment in the activity level of this input unit: the higher the sonority, the higher the activity. This is the way the input sonority values are provided to DLM, with an important difference. In DLM, the activation value for sonority level *s* is simply *s*, whereas in BrbrNet′, it is 2^s–1. The sonority/activation relation in DLM is linear, with sonority values 1, 2, 3, …, 8 being encoded in activity levels 1, 2, 3, …, 8. In BrbrNet′, the sonority/activation relation is exponential: sonority values 1, 2, 3, …, 8 are encoded as activity levels 1, 3, 7, …, 255. This exponential growth is essential for capturing the strict domination structure of the Optimality-theoretic Berber analysis OT_+ which is realized in HG_+ which is realized in BrbrNet′. And this exponential growth is indeed necessary for this sort of network — with the architecture common to DLM and BrbrNet′ — to compute the correct syllabifications.

Figure 2. BrbrNet′: Equivalent syllabifier for Berber

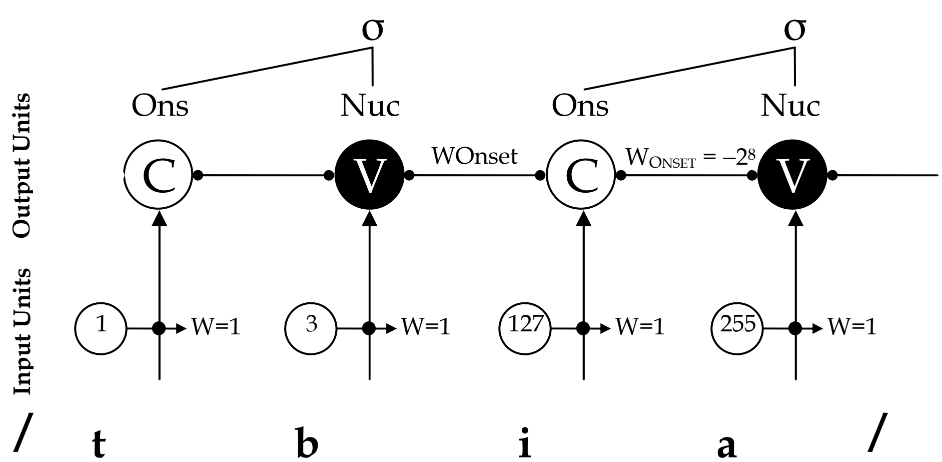

3.7.3. BrbrNet dynamics

To maximize Harmony, BrbrNet uses Brain-State-in-a-Box (BSB) dynamics (Section 9:3.2.3.4). As usual, the *input* to output unit β is ι_β, the sum of the activations of units connected to it, weighted by the connection strengths. An output unit in BrbrNet has two inhibitory connections, one each from its left and right neighbor (end units have only one neighbor). If we call the left neighbor β–1 and the right neighbor β+1, the input from these neighbors to β is $-W_{ONS}(a_{\beta-1} + a_{\beta+1})$, since the weight of these inhibitory connections is $-W_{ONS}$. The input from the input units is just W_v where v is the input segment at the position of β: the activation value of this input unit is 1, and the activations of all other input units connected to β — all the other input units in the stack beneath it — are 0. Thus

(68) $\iota_\beta = -W_{ONS}(a_{\beta-1} + a_{\beta+1}) + W_v$ — v = input segment corresponding to β

If the current activation of unit β is a_β, then Δt units of time later — one tick of the simulated clock — its activation is $a_\beta + \Delta a_\beta$, where the activation change Δa_β is simply proportional to the input to β, ι_β. If this change would raise a_β higher than 1, then the activation is simply set to 1; if the change would lower a_β beneath 0, the activation is set to 0. Thus the activations are always in the interval from 0 to 1: the space of activation states is the *N*-dimensional 'box' in which every coordinate — every activation value — lies in this interval. (*N* is the number of output units.) The activation rule just described is stated formally in (69).

(69) Brain-state-in-a-box dynamics

$$a_\beta + \Delta a_\beta = \begin{cases} a_\beta + k\Delta t\, \iota_\beta & \text{if this result } r \text{ is in the interval } [0,1] \\ 1 & \text{if } r > 1 \\ 0 & \text{if } r < 0 \end{cases}$$

3.7.4. Competence/performance divergence

From the formal analysis of HG_+ above, we know that the Harmony maxima with respect to that grammar are the correct Berber syllabifications. From the preceding analysis we know that, among the representations of syllabification as activity patterns in BrbrNet, the maximal-Harmony states are the same as the Harmony maxima defined by HG_+, hence correct syllabifications. From the formal analysis of the BSB dynamics in Chapter 9, we know BrbrNet will maximize Harmony. So it follows that BrbrNet will compute correct Berber syllabifications.

Well, not quite. There are two gaps in the logic. First, we know the Harmony function of BrbrNet is equivalent to the Harmony function defined by HG_+, but only as concerns activation patterns that are symbolically interpretable as syllabifications. It could be that among the network states that are *not* so interpretable, higher Harmony can be achieved. If so, maximizing Harmony in the network should lead to uninterpretable states, not realizations of correct syllabifications.

As concerns a BSB net, however, it turns out that this possibility cannot actually occur. Essentially, the reason is this. If a unit is getting positive input, it increases Harmony to raise its activation: its Harmony contribution is the product of its input and its activation. So it's always possible to increase Harmony by increasing activation on this unit. Except when that activation level reaches 1, at which point no further increase is allowed by the dynamics: the state cannot leave the box of legal states

Figure 2

in which all activations are between 0 and 1 (inclusive). If a unit is getting negative input, the corresponding reasoning leads to the conclusion that unless the unit's activation is 0, Harmony can be raised by decreasing the unit's activation. If a unit is getting input exactly equal to zero, its activation level does not affect the Harmony, so any value between 0 and 1 yields the same Harmony, and there is no true maximum. So a real maximum must be one of the 'corners of the box' where all output units' activations are either 0 or 1, and nowhere in between. All such states are symbolically interpretable as syllabifications, or more correctly, as parses of the string into syllable peaks and margins.

The second potential problem is that the BSB dynamics finds a *local* Harmony maximum, and the correct syllabifications are *global* maxima. If there are local Harmony maxima that are not global maxima (necessarily a symbolically interpretable activation pattern) then BrbrNet may produce a sub-optimal output, albeit one that has higher Harmony than any 'nearby' states —any states that can be achieved by changing the activation values by an arbitrarily small amount. Are there in fact non-global local maxima, and does BrbrNet erroneously produce them as output?

It is easy to show that there are indeed many non-global local Harmony maxima. First, since every local maximum is a string of 0s and 1s, interpretable as margin/peak assignments, we can notate any local maximum by a more linguistically-evocative string of Cs and Vs, where C denotes 0 and V denotes 1. (Recall that in Berber any segment can be a peak, and any except the most sonorous vowel, *a*, can be a margin. So for syllabification it's not useful to attempt to divide the segments into consonants and vowels. So C represents syllable margin, not 'consonant' per se.)

Now if a syllabification has three consecutive Cs, $C_1C_2C_3$, it cannot be a local Harmony maximum. (Recall (49).) Raising the middle activation C_2 must increase Harmony. This is because the C_2 unit gets no inhibition from its neighbors, both of which have zero activation. But this output unit does get excitation from the input units — more excitation the greater the sonority of the corresponding input segment. Hence increasing this output unit's activation will raise Harmony.

It is also true that a syllabification with two consecutive Vs, V_1V_2, cannot be a local maximum. (Recall (59).) Lowering the activation of either of these two output units from 1 must raise Harmony. Consider V_1; exactly the same reasoning applies to V_2. Since V_1 has a V-neighbor, it gets inhibition equal to $(1)W_{\text{ONS}}$ from this neighbor; it may also get even more inhibition from its other neighbor, but that doesn't matter. The excitation V_1 receives depends on the sonority level of the corresponding input segment: the maximal excitation possible is that for the most sonorous segment, W_a. But by design W_{ONS} is greater than all excitation weights, even the largest one, W_a. Thus the inhibition V_1 receives must exceed the excitation it receives; its input must be negative, so lowering its activation from 1 will increase Harmony.

So any local maximum must satisfy both *CCC and *VV: *CCC is violated by every C that has no V neighbor, and *VV is violated by every V that has, say, a V to its right. Violating *VV means violating ONSET and violating *CCC means gratuitously violating H_{NUC+}. And any CV string satisfying both *CCC and *VV is a string of legal Berber syllables. Every V must be preceded by a C (except an initial V); thus every V can be the nucleus of a syllable which has an onset (except initially, where onsets are not required in Berber). Any remaining Cs (not already serving as onsets) can be syllabified as a coda of the syllable to its left. This C must have a V to its left, since it must have at least one V neighbor to satisfy *CCC, and if its right neighbor were V, it would already be an onset. Thus a local Harmony maximum must be a candidate which is a string of syllables of shape CV, CVC, #V, or #VC, where # marks the edge of the word. These are exactly the legal syllable structures of Berber.

It is remarkable that every local maximum is sequence of legal syllables. Even more remarkable is that the converse is also true. Every string of legal syllables — every CV string satisfying *CCC and *VV — is a local Harmony maximum! Consider any V in the output. It cannot have a V to its left or right, so its neighbors must be either a C unit (activation 0) or an edge of the word. Either way, it gets no inhibition from neighbors. But it necessarily gets some excitation from the input units, so its input is positive, and therefore lowering its activation will lower Harmony. As for an output C, it must have a V either to its left or its right (or both) to satisfy *CCC. That V unit will send inhibition of strength W_{ONS}, and the excitatory input C receives cannot exceed this inhibition, as observed previously. So this input to this output C unit must be negative, and hence raising its activation from 0 will lower Harmony. In other words, if an output satisfies *CCC and *VV, every C unit is receiving negative input and every V is receiving positive input, so changing the activations can only lower Harmony. Since every small activation change lowers Harmony, the state is a local Harmony maximum. The network is pinned into its corner of the box.

(70) BrbrNet local Harmony maxima

An output pattern in BrbrNet is a local Harmony maximum if and only if it realizes a sequence of legal Berber syllables. That is, every activation value is 0 or 1, and the sequence of values is that given by a sequence of substrings taken from the inventory {CV, CVC, #V, #VC}, where C denotes 0, V denotes 1 and # denotes a word edge.

Thus the problem of local Harmony maxima in BrbrNet is far from hypothetical. The output space is rife with local maxima, only one of which is the global Harmony maximum, the *correct* output syllabification.

3.7.5. Simulations

Whether BrbrNet will succeed in finding the one global Harmony maximum among the multitude of local maxima cannot, as far as we know, be determined analytically. We therefore performed computer simulations, giving the network 10,000 randomly-generated inputs, with length up to 10 segments. The initial values of all output units was zero.[27]

[27] We are extremely grateful to Yoshiro Miyata, who performed these simulations with the PlaNet

Of these 10,000 inputs, all were correctly parsed, except for 103 cases. Of these, all but one contained a substring of sonority values meeting the template (71).

(71) Problematic input pattern (sonority values)

[Z] Y X 7 8 with [Z <] < Y < X

or the same sequence with left and right reversed. Instead of the correct parse [V].CV.CV, the erroneous output is [C]VC.CV. The bracketed segment may be present or absent.

The one remaining error instantiated the same pattern, with 8 and 7 replaced by 7 and 6, respectively.

Subsequent simulation employed a very small Δt — time interval between 'clock ticks' or steps of the simulation. With this more accurate simulation, every one of the 103 previously problematic inputs were parsed correctly.

Figure 3 shows a trace of the activation trajectory for a rather simple case, our example /txznt/ 'you$_{sing}$ stored', with sonority profile 13451. The correct syllabification is *.tx́.zńt.* or .CV.CVC.. With the time step set so that the activation rate coefficient in equation (69) is $k\Delta t$ = .00008, after 500 simulation steps, the activation values are 0, 0.27, 0, 1, 0. At this point the network is guaranteed to ultimately converge to the correct sequence 0, 1, 0, 1, 0 because, while the first peak *x́* has only reached activation value 0.27, its neighbors are both 0 so it receives no inhibition, only excitation. The non-peak segments initially increase in activation because no units receive inhibition until their neighbors have reached a significant activation level. The margin segments *t z t* respectively reach their maximal activation levels of .00034, 0.013, .00011 at time steps 11, 24, 4 and have become completely deactivated by time steps 21, 47, 7. As the lower graph shows, the Harmony monotonically rises throughout the entire computation. The Harmony of the correct parse is $(1)(W_x)(1) + (1)(W_n)(1) = (2^3–1) + (2^5–1) = 7 + 31 = 38$; after 500 steps, it reaches 32.89 = (1)(7)(.27) + (1)(31)(1): because the *x* unit has only reached activation level 0.27, it is not yet contributing the 7 units of Harmony it will provide when it ultimately reaches activation 1.

As might be expected, the highest sonority segment *n* quickly drives its output unit to 1, inhibiting its neighbors (enforcing *VV); the next-most-sonorous segment not thereby inhibited, *x*, drives its output unit up to 1 somewhat less quickly, inhibiting its neighbors in the process. And for longer inputs, this continues, with successively lower-sonority peaks being assigned. The process 'works down the sonority hierarchy' just like the Dell-Elmedlaoui algorithm does, but not because this sequence has been stipulated: the emergent seriality of processing arises from the widely disparate weights corresponding to different sonority levels.

One might say that stipulating the exponentially-growing weights is just a connectionist means of programming seriality into the network's search algorithm. While this is certainly true to a large extent, the story is not so simple. The Dell-Elmedlaoui algorithm *must* monotonically descend the sonority scale, for this is hard-

neural-network simulation environment (Miyata 1991).

coded in the algorithm. BrbrNet, on the other hand, exhibits much more complex decision-making in cases which are not so clear-cut as that shown in Figure 3.

Figure 3. Parsing /txznt/ → tx́zńt

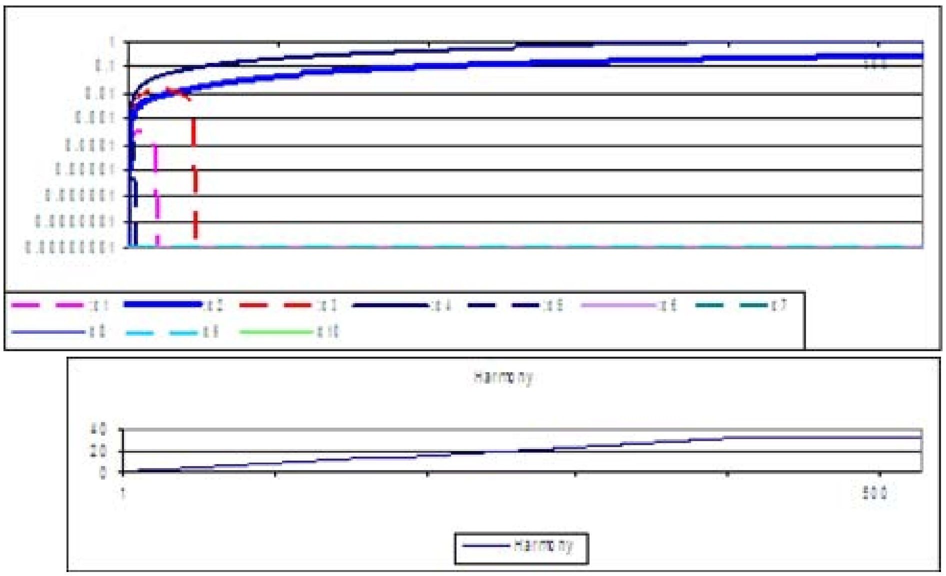

In fact, the time course is quite tricky exactly in the cases that were problematic for the less accurate simulation. And observing the network processing these challenging inputs explains why the difficult pattern (71) has the particular shape it does.

Figure 4 shows the activation state trajectory of BrbrNet parsing one of the challenging inputs not correctly parsed in the first, insufficiently fine-grained, simulation. The input sonority sequence is 8 1 2 1 <u>3 4 5 7 8</u> 7 as in hypothetical /apbtxznuai/ → *.á.pb́.tx́.zń.wáy.* Note that this instantiates the problematic pattern (71) (underlined). While some of the segments exhibit the simple behavior shown in the previous example, others do not. The fifth and seventh segments — hypothetical *x n*, with sonority values 3 5 — must be syllabified as a peak in the correct parse (as in actual *txznt* above). These activation values rise initially, as all do, but are then inhibited to zero. Later, these units rise phoenix-like out of activational oblivion to ultimately assume their proper places atop their respective syllables.

To understand the complex behavior arising in the problematic cases, it is helpful to consider a simpler input also incorrectly parsed in the coarser simulation: 1 2 7 8 as in /txia/ → *.tx́.yá.*; the relevant incorrect parse is **.t́x.yá*. This trajectory of BrbrNet while parsing this input is shown in Figure 5.

Figure 3

Figure 4. Parsing sonority profile 8 1 2 1 3 4 5 7 8_7

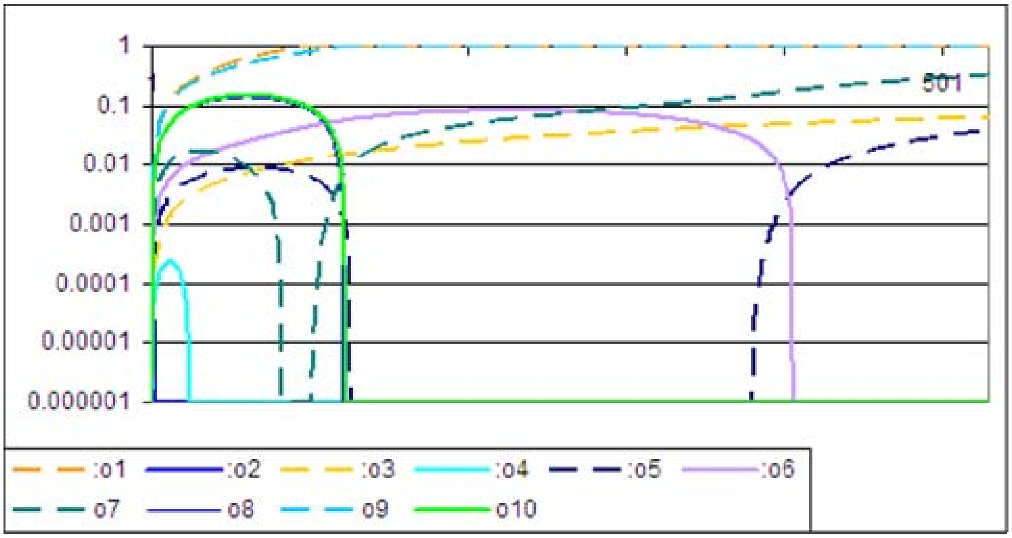

Examination of this plot reveals why the problematic input pattern is difficult. The sonority-8 final segment quickly rises to its correct value 1: this the final peak *á*. At the same time, the sonority-7 segment to its left rises to a considerable activation level before it becomes inhibited sufficiently by final *á* to overcome its strong excitatory input $2^7–1$ = 127. While the unit for the sonority-7 segment /i/ is active, it inhibits its left neighbor, which, with a mere sonority of 2, receives only excitatory input $2^2–1$ = 3 and is easily intimidated by its potent neighbor. It is driven to zero activation very quickly, and pinned there as long as 7 remains active, which is a considerable amount of time. During that time, when 2 is pinned at zero, it fails to inhibit its left neighbor; with no inhibition, even with its paltry sonority level of 1, /t/ rises appreciably, adding to the inhibition of 2. At this point (around 100 time steps) the net is proceeding directly toward the erroneous parse VCCV, **.t́x.yá*. But the action is not over because it is not yet the case that every active unit has only zero-activation neighbors. In particular, 7 does not. So its impressive rise will soon become a precipitous fall. The critical point is when 7 finally is quashed to activity 0. At this point, 2 is still at zero but its weaker neighbor 1 has a considerable head start. Released from the inhibition by 7, 2 is now inhibited only by 1. If 1 has such a big head start that it is already highly active, it is possible that its inhibition of 2 can exceed 2's excitation by the input: the weights realizing ONSET ≫ HNUC_{+} ensure that if 1 has reached full activation, its inhibition of 2 is unbeatable. So if 7 takes long enough to give up its ambitions of peak-hood, 1 can conceivably have become so close to complete success that 2 cannot even get off the ground. If however 1 has not risen sufficiently by the time 2 is released by 7, 2 will get less inhibition than excitation and it will start to rise. Eventually, it must overtake 1 because of its greater excitation level (2^2–1 = 3, vs. 2^1–1 = 1 for 1). But eventuality may never arrive. It will take time for 2 to overtake 1, and during that time 1 will continue to rise for a time. If that suffices to get 1 over the finish line of activation 1, the race is over, because as just observed the power of ONSET given a fully-active peak is dominant, and drives 2 back down to zero.

Figure 5. Parsing sonority profile 1 2 7 8

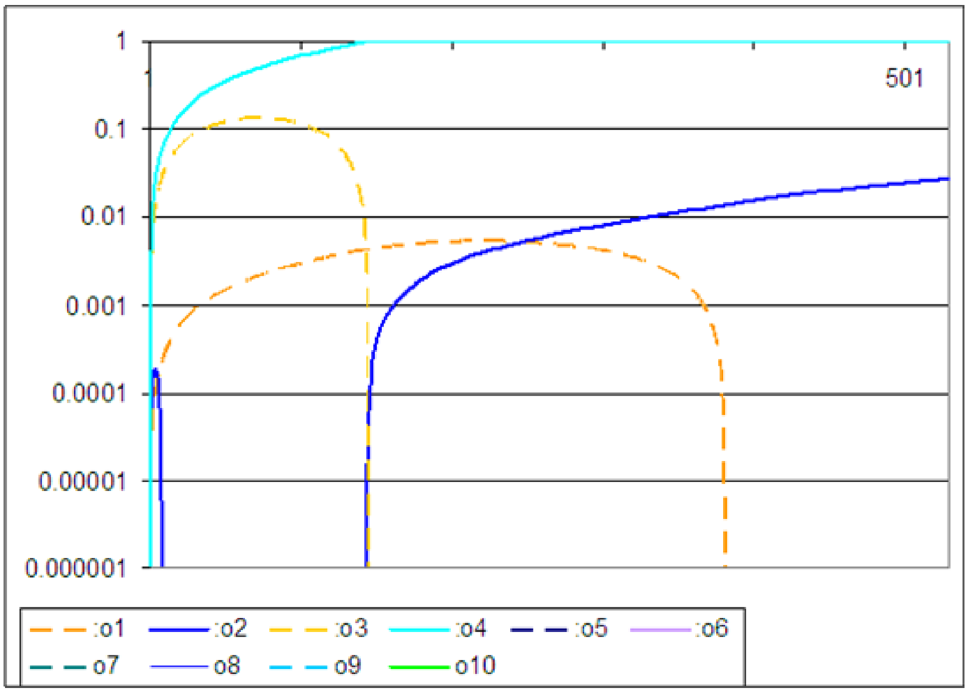

This scenario is responsible for making the difficult inputs difficult. But as Figure 3 shows, in actuality 2 wins the peakness race. Until the dynamics is sufficiently well-analyzed to quantify all the timing-dependent interactions played out in this scenario, it is not clear whether the just outcome results from skill or luck. That is, there is as yet no theorem guaranteeing success from BrbrNet. It is however true that in every one of the 103 problematic examples from the first simulation, the correct parse is selected. The network succeeds in avoiding the manifold local Harmony maxima, arriving ultimately at the global maximum.

The case shown in Figure 5B is in some sense the worst of the worst, and success on it is very encouraging. The reason it is the most difficult of the difficult cases is this. To maximize the opportunity for 1 to beat 2, yielding the incorrect parse, we must maximize the time during which 1 is rising and 2 is pinned at zero. And since 2 is pinned by 7, this means maximizing the time required for 7 to be driven to zero by

Figure 5

8. If the sonority difference between 7 and 8 is increased, the lower-sonority segment will be inhibited more quickly, so to maximize opportunity for error, we need to keep 7 as close in sonority as possible while still less than 8. So if 7 8 is to be modified to make the parsing harder, the replacement must be 6 7 or 5 6 etc. Of these *k* *k*+1 pairs, the larger *k*, the longer the inhibition time of the less sonorous segment so this time is maximized by the 7 8 pair.

Now to get the most out of the time during which 7 keeps 2 pinned, 1 should be as sonorous as possible: this allows it to rise fastest during that period. But raising 1 to say 3 requires raising 2 to 4, which means that after 7 releases its neighbor, now 4, then 4 will overtake 3 more quickly. Since the speed of rise is determined by the weights which grow exponentially up the sonority scale, the speed difference is much greater for the 3 4 pair than it is for the 1 2 pair. Thus it appears that with 7 8 maximizing the time for 7 to be inhibited, and 1 2 minimizing the speed at which 2 can overtake 1, the pattern 1 2 7 8 affords the best chance of error.

Figures 6–7 show a few more examples of BrbrNet correctly parsing problematic sonority patterns.

Figure 6. Parsing sonority profile 1 2 3 7 8

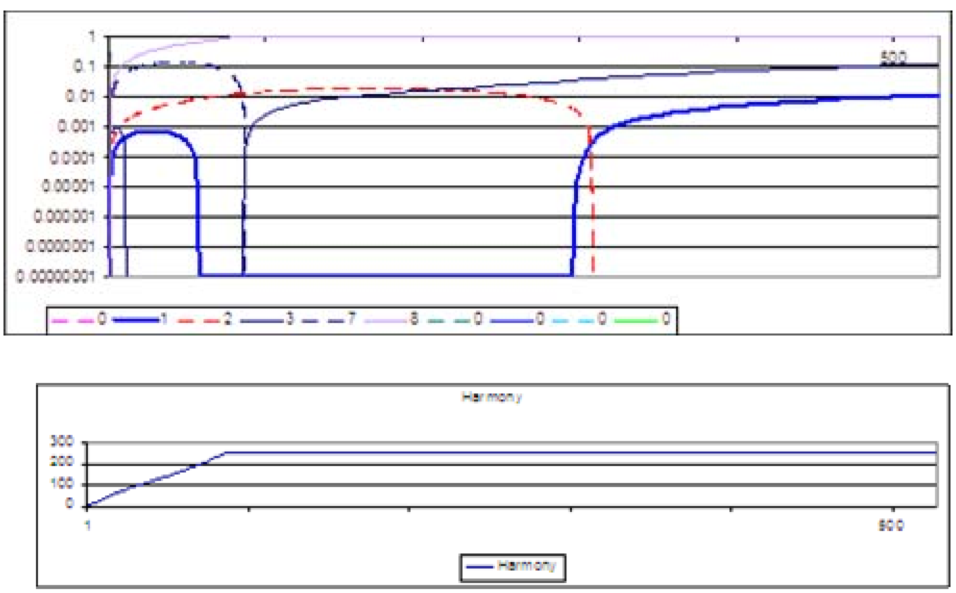

Figure 7. Parsing sonority profile 8 7 2 1 2

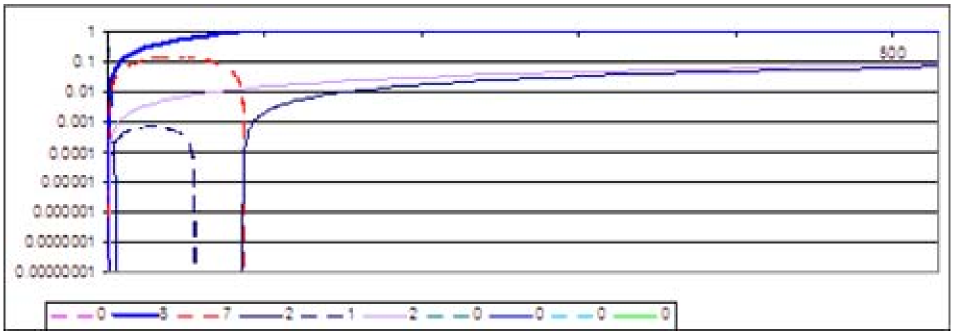

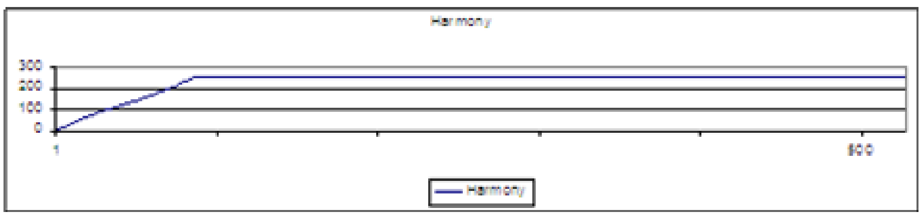

3.8. Summary

In this section we have examined the core of the complex syllabification system of a dialect of Berber. We have seen how the original Prince and Smolensky 1993/2004 analysis of this system can be directly realized in a Harmonic Grammar, where exponentially-weighted constraints implement strict domination. While exponential weighting is crucial for a core part of the constraint interaction, the need for a large range of weights is not nearly as extreme as expected for the general case. This OT analysis is rather unusual in its use of a positive constraint, so we next examined another OT Berber analysis from Prince and Smolensky 1993/2004 which employs only the standard, negative type of constraint — penalizing violation rather than rewarding satisfaction. This too is readily transformed to a Harmonic Grammar. The positive OT analysis was then realized in a local connectionist network, BrbrNet, which appears to correctly find the global Harmony maximum — the correct syllabification — amid myriad local maxima, despite using a simple, local Harmony-maximizing activation dynamics that is guaranteed only to find local maxima. For easy inputs, the resulting temporal trajectory implements the sequential parsing algorithm of Dell and Elmedlaoui. But for difficult inputs, the road to the correct parse includes major detours directed toward incorrect competitors.

References

Aissen, Judith. 2001. Markedness and subject choice in Optimality Theory. In *Optimality-Theoretic Syntax*, eds. Géraldine Legendre, Sten Vikner and Jane Grimshaw, 61-96. Cambridge, MA: MIT Press.

Anttila, Arto. 1997. *Variation in Finnish Phonology and Morphology.* Doctoral dissertation, Stanford University.

Arad, Maya. 1998. VP-Structure and the syntax-lexicon interface. MIT Occasional Paper in Linguistics 16.

Aranovich, Raúl. 2003. The semantics of auxiliary selection in Old Spanish. *Studies in Language* 27:1-37.

Artstein, Ron. 1999. Person, animacy and null subjects. Ms. Rutgers University.

Asudeh, Ash. 1999. Linking, optionality, and ambiguity in Marathi: An Optimality Theory analysis. Ms. Stanford University.

Baker, Mark C. 1988. *Incorporation: A Theory of Grammatical Function Changing*. Chicago: University of Chicago Press.

Bard, Ellen Gurman, Robertson, Daniel, and Sorace, Antonella. 1996. Magnitude estimation of linguistic acceptability. *Language* 71:32-68.

Bentley, Delia, and Eyrthorsson, Thorallur. 2002. Auxiliary selection and the semantics of unaccusatives. Ms. University of Manchester.

Benzing, Joseph. 1931. Zur Geschichte von ser als Hilfszeitwort bei den Intransitiven Verben im Spanischen. *Zeitschrift für rmoanische Philologie* 51:385-460.

Boersma, Paul. 1997. How we learn variation, optionality, and probability. *Proceedings of the Institute of Phonetic Sciences of the University of Amsterdam* 21:43-58.

Borer, Hagit. 1994. The projection of arguments. In *Functional Projections*, eds. E. Benedicto and J. Runner. Amherst: University of Massachusetts.

Borer, Hagit. 1998. Deriving passive without theta roles. In *Morphology and its Relation to Phonology and Syntax*, eds. Steven Lapoint, Diane Brentari and Patrick Farrell, 60-99. Stanford: CSLI Publications.

Burzio, Luigi. 1986. *Italian Syntax: A Government-Binding Approach*. Dordrecht: Reidel.

Burzio, Luigi. 2002. Surface-to-surface morphology: When your representations turn into constraints. In *Many Morphologies*, ed. Paul Boucher, 142-177. Cascadilla Press. Rutgers Optimality Archive 341.

Clements, George N. 1990. The role of the sonority cycle in core syllabification. In *Papers in Laboratory Phonology I: Between the grammar and the physics of speech*, eds. John Kingston and Mary Beckman, 283–333. Cambridge, MA: Cambridge University Press.

Clements, George N. 1997. Berber syllabification: Derivations or constraints? In *Derivations and Constraints in Phonology*, ed. Iggy Roca, 289-330. Oxford: Oxford University Press.

Croft, William. 1990. *Typology and Universals*. Cambridge: Cambridge University Press.

Cummins, Sarah. 1996. *Meaning and Mapping*. Doctoral dissertation, University of Toronto.

Davidson, Lisa, and Goldrick, Matthew. 2003. Tense, agreement and defaults in child Catalan: An Optimality Theoretic analysis. In *Linguistic theory and language development in Hispanic languages*, eds. S. Montrul and F. Ordóñez, 193-211. Cambridge, MA: Cascadilla Press.

Davidson, Lisa, and Legendre, Géraldine. 2003. Defaults and competition in the acquisition of functional categories in Catalan and French. In *A Romance Perspective on Language Knowledge and Use: Selected Papers from the 2001 Linguistic Symposium on Romance Languages (LSRL)*, eds. Rafael Nuñez-Cedeño, Luis López and Richard Cameron, 273-290. John Benjamins.

Dell, François, and Elmedlaoui, Mohamed. 1985. Syllabic consonants and syllabification in Imdlawn Tashlhiyt Berber. *Journal of African Languages and Linguistics* 7:105–130.

Dowty, David. 1979. *Word Meaning and Montague Grammar*. Dordrecht, Holland: Reidel.

Dowty, David. 1991. Thematic proto-roles and argument selection. *Language* 67:547–619.

Flemming, Edward. 2001. Scalar and categorical phenomena in a unified model of phonetics and phonology. *Phonology* 18:7-44.

Goldsmith, John A. 1992. Local modeling in phonology. In *Connectionism: Theory and Practice*, ed. S. Davis. Oxford, England: Oxford University Press.

Grimshaw, Jane. 1990. *Argument Structure*. Cambridge, MA: MIT Press.

Haegeman, Liliane. 1994. *Introduction to Government & Binding Theory*. Oxford: Blackwell.

Hale, Kenneth, and Keyser, Samuel Jay. 1986. Some transitivity alterations in English. Center for Cognitive Science, MIT. Cambridge.

Hale, Kenneth, and Keyser, Samuel Jay. 1993. On argument structure and the lexical representation of syntactic relations. In *The View from Building 20*, eds. Kenneth Hale and Samuel Jay Keyser, 53-109. Cambridge, MA: MIT Press.

Jakobson, Roman. 1965/1995. Quest for the essence of language. In *On Language: Roman Jacobson*, eds. Linda R. Waugh and Monique Monville-Burston, 407-421. Cambridge, MA: Harvard University Press.

Keenan, Edward L., and Comrie, Bernard. 1977. Noun Phrase accessibility and Universal Grammar. *Linguistic Inquiry* 8:63-100.

Keller, Frank, and Sorace, Antonella. 2000. Gradient auxiliary selection in German: an experimental investigation. Manuscript. University of Edinburgh.

Kirchner, Robert. 1998. *An effort-based approach to consonant lenition*. Doctoral, Linguistics Department, UCLA, Los Angeles. Rutgers Optimality Archive 276.

Labelle, Marie. 1992. Change of state and valency. *Journal of Linguistics* 28:375-414.

Legendre, Géraldine. 1989. Unaccusativity in French. *Lingua* 79:95–164.

Legendre, Géraldine. to appear. Optimizing auxiliary selection in Romance.

Legendre, Géraldine, Hagstrom, Paul, Chen-Main, Joan, Tao, Liang, and Smolensky, Paul. in press. Deriving output probabilities in child Mandarin from a dual-optimization grammar. *Lingua*.

Legendre, Géraldine, Miyata, Yoshiro, and Smolensky, Paul. 1990a. Can connectionism contribute to syntax? Harmonic Grammar, with an application. *Proceedings of the 26th Meeting of the Chicago Linguistic Society*, Chicago, IL, 237-252.

Legendre, Géraldine, Miyata, Yoshiro, and Smolensky, Paul. 1990b. Harmonic Grammar-A formal multi-level connectionist theory of linguistic well-formedness: An application. *Proceedings of the Twelfth Annual Conference of the Cognitive Science Society*, 884–891. Lawrence Erlbaum.

Legendre, Géraldine, Miyata, Yoshiro, and Smolensky, Paul. 1991. Unifying syntactic and semantic approaches to unaccusativity: A connectionist approach. *Proceedings of the Seventh Annual Meeting of the Berkeley Linguistics Society*, Berkeley, CA, 156-167.

Legendre, Géraldine, and Rood, David S. 1992. On the interaction of Grammar Components in Lakhóta: Evidence from Split Intransitivity. *Berkeley Linguistics Society* 18.

Legendre, Géraldine, Smolensky, Paul, and Wilson, Colin. 1998. When is less more? Faithfulness and minimal links in *wh*-chains. In *Is the Best Good Enough? Optimality and Competition in Syntax*, eds. Pilar Barbosa, Danny Fox, Paul Hagstrom, Martha McGinnis and David Pesetsky, 249–289. Cambridge, MA: MIT Press and MITWPL. Rutgers Optimality Archive 117.

Legendre, Géraldine, and Sorace, Antonella. 2003. Mapping lexical semantics onto syntactic structure: The problem of unaccusative mismatches in Romance languages. *Korean Journal of Cognitive Science* 4:43-78.

Legendre, Géraldine, and Sorace, Antonella. in press. Split intransitivity in French: An Optimality -theoretic perspective. In *Les langues romanes: problèmes de la phrase simple*, ed. Danièle Godard. Paris: CNRS éditions.

Legendre, Géraldine, Vainikka, Anne, Hagstrom, Paul, and Todorova, Marina. 2002. Partial constraint ordering in child French syntax. *Language Acquisition* 10:189-227.

Levin, Beth, and Rappaport Hovav, Malka. 1992. The lexical semantics of verbs of motion: the perspective from unaccusativity. In *Thematic Structure: Its Role in Grammar*, ed. Iggy Roca. Dordrecht: Foris.

Levin, Beth, and Rappaport Hovav, Malka. 1994. A preliminary analysis of causative verbs in English. *Lingua* 92:35-77.

Levin, Beth, and Rappaport Hovav, Malka. 1995. *Unaccusativity: At the Syntax-Semantics Interface*. Cambridge, MA: MIT Press.

Levin, Beth, and Rappaport Hovav, Malka. in press. *From Lexical Semantics to Argument Realization*. Cambridge: Cambridge University Press.

McClure, William. 1995. *Syntactic Projections of the Semantic of Aspect*. Tokyo: Hitsujishobo.

Miyata, Yoshiro. 1991. A user's guide to PlaNet version 5.6: A tool for constructing, running, and looking into a PDP network. Ms. Computer Science Department, University of Colorado, Boulder.

Nagy, Naomi, and Reynolds, William. 1997. Optimality Theory and variable word-final deletion in Fætar. *Language Variation and Change* 9:37-55.

Perlmutter, David M. 1978. Impersonal passives and the Unaccusative Hypothesis. *Proceedings of the Proceedings of the Fourth Annual Meeting of the Berkeley Linguistic Society*, Berkeley. University of California.

Perlmutter, David M. ed. 1983. *Studies in Relational Grammar I*. Chicago, IL: University of Chicago Press.

Prince, Alan. 1990. Quantitative Consequences of Rhythmic Organization. In *CLS 26-II: Papers from the Parasession on the Syllable in Phonetics and Phonology*, eds. Karen Deaton, Manuela Noske and Michael Ziolkowski. Chicago: Chicago Linguistic Society.

Prince, Alan. 1993. In defense of the number *i*: Anatomy of a linear dynamic model of linguistic generalizations. Technical Report RuCCS-TR-1. Rutgers Center for Cognitive Science, Rutgers University. New Brunswick, NJ.

Prince, Alan. 2002. Anything goes. Rutgers University, New Brunswick, NJ. Rutgers Optimality Archive 536.

Prince, Alan, and Smolensky, Paul. 1993/2004. *Optimality Theory: Constraint Interaction in Generative Grammar*. Malden, MA: Blackwell. Technical Report, Linguistics Department, Rutgers University, New Brunswick, NJ and Department of Computer Science, University of Colorado at Boulder (1993). Rutgers Optimality Archive 537 (2002).

Pustejovsky, James, and Busa, Frederica. 1995. Unaccusativity and event composition. In *Temporal Reference, Aspect and Actionality*, eds. P.M. Bertinetto, V. Bianchi, J. Higginbotham and M. Squartini. Torino: Roserberg & Sellier.

Rappaport Hovav, Malka, and Levin, Beth. 1998. Building verb meanings. In *The Projection of Arguments: Lexical and Compositional Factors*, eds. Miriam Butt and Wilhelm Geuder, 97-134. Stanford: CSLI Publications.

Rosen, Carol. 1984. The interface between semantic roles and initial grammatical relations. In *Studies in Relational Grammar 2*, eds. David Perlmutter and Carol Rosen, 38-77. Chicago, IL: University of Chicago Press.

Silvertein, Michael. 1976. Hierarchy of features and ergativity. In *Grammatical Categories in Australian Languages*, ed. R. M. W. Dixon, 112-171. Canberra: Australian Institute of Aboriginal Studies.

Slade, Benjamin. 2003. How to rank constraints: Constraint conflict, grammatical competition, and the rise of periphrastic *do*. In *Optimality Theory and Language Change*, ed. D. Eric Holt, 337-385. Dordrecht: Kluwer.

Smolensky, Paul. 1995. On the internal structure of the constraint component Con of UG. Talk presented at the UCLA Linguistics Department. April 7. Rutgers Optimality Archive 86.

Sorace, Antonella. 1993a. Incomplete vs. divergent representations of unaccusativity in non-native grammars of Italian. *Second Language Research* 9:22-47.

Sorace, Antonella. 1993b. Unaccusativity and auxiliary choice in non-native grammars of Italian and French: asymmetries and predictable indeterminancy. *Journal of French Language Studies* 3:71-93.

Sorace, Antonella. 1995a. Acquiring argument structures in a second language: the unaccusative/unergative distinction. In *The Current State of Interlanguage*, eds. Lynn Eubank, L. Selinker and Michael Sharwood Smith, 153-175. Amsterdam: John Benjamins.

Sorace, Antonella. 1995b. Contraintes sémantiques sur la syntaxe: l'acquisition de l'inaccusativité en italien L2. *Acquisition et Intéraction en Langue Etrangère* 5:79-113.

Sorace, Antonella. 2000. Gradients in auxiliary selection with intransitive verbs. *Language* 76:859-890.

Sorace, Antonella, and Cennamo, Michela. 2000. Aspectual constraints on auxiliary choice in Paduan. Manuscript. University of Naples and University of Edinburgh.

Sorace, Antonella, and Keller, Frank. 2004. Gradience in linguistic data. *Lingua*.

Sorace, Antonella, and Shomura, Y. in press. Lexical Constraints on the acquisition of split intransitivity: evidence from L2 Japanese. *Studies in Second Language Acquisition* 23.

Sorace, Antonella, and Vonk, Wietske. 1998. Gradient effects of unaccusativity in Dutch. University of Edinburgh and Max Planck Institute for Psycholinguistics, Nijmegen.

Tuttle, Edward. 1986. The spread of esse as a universal auxiliary in central Italo-Romance. *Medioevo Romanzo* XI.

van Hout, Angeliek. 1996. *Event Semantics of Verb Frame Alterations. A Case Study of Dutch and Its Aquisition*. Tilburg: University of Tilburg.

van Hout, Angeliek. 2000. Event semantics in the lexicon-syntac interface: Verb frame alterations in Dutch and their aquisition. In *Events as grammatical objects*, eds. C. Tenny and J. Pustejovsky, 239-282. Stanford: CSLI, Stanford.

Van Valin, R.D. 1990. Semantic parameters of split intransitivity. *Language* 66:221–260.

Zaenen, Annie. 1993. Integrating syntax and lexical semantics, ed. James Pustejovsky, 129-161. Dordrecht: Kluwer.