

Predicting vowel inventories from a Dispersion-Focalization model: New results*

Nathan Sanders and Jaye Padgett
Williams College and UC Santa Cruz

1 Overview

Since the seminal work of Liljencrants & Lindblom (1972), a key testing ground for functional, evolutionary, or emergentist approaches to sound systems has been the typology of vowel inventories (for example, Lindblom 1986, Schwartz et al. 1997a, de Boer 2000). An important innovation of Schwartz et al.'s Dispersion-Focalization Theory (DFT) was calculating the optimality (“energy”) of a vowel system as a weighted combination of *two* separate auditory parameters:

- (1) **dispersion:** maximization of the auditory distance between vowels (as in Liljencrants & Lindblom 1972)
- focalization:** maximization of the importance of “focal” vowels such as [i] and [y].

We report results of new vowel system simulations following the original DFT formulas of Schwartz et al. However, our means of selecting candidate systems for comparison explores the search space more effectively, allowing for more thorough and accurate computation of DFT’s predictions. In particular, we find a greater number of optimal systems than originally reported for DFT by Schwartz et al. throughout the entire range of possible parameter settings in DFT. This seems to be a good result, since many of the newly discovered optimal systems have analogues in the UCLA Phonological Segment Inventory Database (UPSID; Maddieson 1984, Maddieson & Precoda 1989).

2 How Dispersion-Focalization Theory works

In DFT, vowel systems are compared according to their total “energy” according to the distribution and types of vowels in each system. The lower a vowel system’s total energy is, the more optimal it is. For a given vowel system $\{V_1, \dots, V_N\}$, each vowel V_i is characterized by its first four formants $\langle F1_i, F2_i, F3_i, F4_i \rangle$, measured in Bark.¹ A system’s total energy E_{DF} (2) is just the simple sum of its dispersion energy E_D (§2.1) and its focalization energy E_F (§2.2):

$$(2) \quad E_{DF} = E_D + E_F$$

*Our appreciation goes to audiences at Stanford, UCSC, and MIT, for their helpful feedback on various stages of this work. All responsibility for the content of this paper lies solely with the authors.

¹We follow Schwartz et al. in calculating Bark with the formula $f_{Bk} = 7 \cdot \sinh^{-1}(f_{Hz}/650)$.

2.1 Dispersion energy

The dispersion energy of a vowel system is a measure of the overall auditory distance between the vowels in the system. The basic auditory distance between two vowels V_i and V_j is the euclidean distance d_{euc} between them in the auditory space based on their values of F1 and *effective* F2 (a.k.a. “F2 prime”, a hypothesized perceptual integration of F2, F3, and F4, symbolized as F2’; Carlson et al. 1970, Carlson et al. 1975), as demonstrated in Fig. 1. However, because F2’ spans a significantly larger range (about 10–11 Bk) than F1 does (only about 4–5 Bk), this




Figure 1: Ordinary euclidean auditory distance between vowels

simple euclidean measure of auditory distance overgenerates color (F2’) contrasts in comparison to height (F1) contrasts. In order to generate more realistic predictions about color vs. height contrasts, phonetic models of vowel dispersion must compress the color space (Fig. 2). Furthermore, there is independent acoustic and




Figure 2: Sample F2’ compression

perceptual support for weighting F1 more heavily than F2’. For example, F1 is known to be louder than higher formants, and louder formants will weight more heavily in perceptibility (see Lindblom 1986, Schwartz et al. 1997a, Benkí 2003). The amount of weighting F2’ receives is represented in DFT by the parameter λ , which falls between 0 (for which dispersion is determined solely by F1) and 1 (for which F1 and F2’ contribute equally to dispersion).

In DFT, the total dispersion energy E_D (3) of a vowel system with N vowels is the sum of the inverse squares of the λ -weighted distances d_{ij} between every pair

of vowels V_i and V_j in the system:

$$(3) \quad E_D = \sum_{\substack{i=1, \dots, N-1 \\ j=i+1, \dots, N}} \frac{1}{d_{ij}^2} \quad \text{where } d_{ij} = \sqrt{(F1_i - F1_j)^2 + \lambda^2(F2'_i - F2'_j)^2}$$

From this equation, it is clear that the more vowels in a system that are at the periphery of the auditory space, the lower the system's E_D will be. Thus, optimal vowel systems will tend to be more peripheral in structure, giving a preference to front unrounded vowels, back round vowels, high vowels, and low vowels, and a strong avoidance of mid central vowels. In addition, optimal systems will have vowels spread apart from each other, rather than clustered together.

2.2 Focalization energy

DFT additionally assumes that some vowels, so-called “focal vowels,” are preferred in vowel systems due to their own inherent acoustic qualities, irrespective of the relational role they play in the system as a whole. Specifically, a focal vowel in DFT has one or more pairs of adjacent formants that are close together, causing the formants to enhance each other, and making the vowel more perceptually robust overall (Schwartz & Escudier 1987, 1989; cf. Stevens 1972). The focalization energy E_F of a vowel system is the sum of the focalization energies for each vowel in the system. Each individual vowel's focalization energy is the negative sum of the inverse squares of the differences between adjacent formants:

$$(4) \quad E_F = \alpha \sum_{i=1}^N \left(\frac{-1}{(F1_i - F2_i)^2} + \frac{-1}{(F2_i - F3_i)^2} + \frac{-1}{(F3_i - F4_i)^2} \right)$$

The value of the parameter α determines the relative influence of *focalization* vis-à-vis *dispersion*. Like λ , α falls between 0 and 1.²

Note that the focalization energy is always negative. For a vowel like [i], which has F3 and F4 very close together, $(F3 - F4)^2$ will be small, which means its inverse will be large. This inverse is given a negative sign, so the overall focalization energy for [i] will have a large magnitude negative component that reduces the total energy of the vowel system by a larger amount than if its F3 and F4 were farther apart. Thus, optimal vowel systems will have a larger number of focal vowels. The most focal vowels in DFT by far are [i] and [y], with others ranked roughly as in (5):

- | | |
|---|--|
| (5) low E_F
(high magnitude negative) | high E_F
(low magnitude negative) |
| [i y] < [ɪ] < [e] < [ɘ] < [ɛ] < [æ a ɑ] < [u ʊ ø œ o ɔ ɒ] < [ə ʌ] < [ɨ ʉ ɤ] | < [ɨ ʉ ɤ] |
| most focal | least focal |

²Unlike with λ , there seems to be no reason why α cannot be greater than 1. It may be worthwhile in future work to explore variations of DFT for higher values of α .

2.3 Prototypes

To limit the amount of computation time required to find optimal vowel systems, Schwartz et al. utilize a finite, predetermined set of 37 vowel “prototypes” (Fig. 3). These prototypes are based primarily upon the vowel system from UPSID, plus two extra vowels in the gap in the acoustic vowel space between the back round vowels and the back unrounded vowels. For each prototype vowel, Schwartz et al. set fixed values for F1–F4 that are typical of an adult male speaker, with F2' calculated from F2–F4 by Mantakas et al.'s (1986) computation.




Figure 3: Vowel prototypes in DFT

3 What Dispersion-Focalization Theory predicts

3.1 Phase spaces

With these equations and prototypes in hand, the task is to find, for any number of vowels N , the optimal vowel system of size N , which has the lowest total energy $E_{DF}(2)$ of all systems of size N . The result depends heavily on the particular values of the two weighting parameters λ and α , so there will potentially be different optimal vowel systems of size N for different choices of $\langle \lambda, \alpha \rangle$ pairs. To more easily visualize the optimal vowel systems that are found for various choices of

$\langle \lambda, \alpha \rangle$, Schwartz et al. plot the optimal vowel systems in the $\lambda \times \alpha$ space by means of a “phase space” diagram, which divides the $\lambda \times \alpha$ space into regions where particular vowel systems are found to be optimal. For example, the phase space




Figure 4: Schwartz et al.’s phase space for $N = 5$

diagram in Fig. 4 shows that the vowel system [i e a o u] is found to be optimal in the region to the far left, where $\lambda \leq 0.2$ (and for some values $0.2 \leq \lambda \leq 0.3$ for $\alpha > 0$), while the vowel systems [i ‘e’ a ‘o’ u], [i y a ‘o’ u], and [i ε a u u] are each optimal in other regions of the $\lambda \times \alpha$ space.

3.2 Search algorithm

Searching the total set of possible systems to find the one single optimal system is not a computationally trivial task, even with the limitation of only having 37 vowel prototypes to choose from. Thus, some search algorithm must be used which picks out only certain vowel systems for consideration of being the most optimal. Schwartz et al. utilize the following algorithm to make the search procedure more manageable:

- (6) a. For each value of $N = 3, \dots, 9$, pre-select a number of systems of size N to be candidates for being *the* optimal vowel system for some $\langle \lambda, \alpha \rangle$ in DFT.³
- b. For various $\langle \lambda, \alpha \rangle$ pairs, compute the energy of every candidate system, according to equations (2), (3), and (4).⁴
- c. For each $\langle \lambda, \alpha \rangle$ pair, select as optimal the candidate system with the lowest energy.

³Schwartz et al. do not specify precisely how these candidate systems are chosen, but presumably, they are hand-picked based on inferences about what should be optimal in various regions of the $\lambda \times \alpha$ space, as well as typological and phonetic plausibility. For example, [i e a o u] would most certainly have been a candidate 5-vowel system, but there would be no apparent need to even consider a system such as [ø i ə ɾ u], so it would likely have been completely ignored.

⁴Again, Schwartz et al. do not specify precisely how these $\langle \lambda, \alpha \rangle$ pairs are chosen. Given the locations of their key plotted points in their phase spaces, $\langle \lambda, \alpha \rangle$ seems to have been chosen largely at fixed intervals, though this appearance could just be an artifact of their graphing method.

However, there are at least two ways this search algorithm can go wrong:

- (7) a. The initial posited candidate systems in step (6a) might not include the system that is actually the most optimal. Some truly optimal systems may be overlooked because they incorrectly appear to a human to be “obviously” less optimal than some other system, when in fact, the actual calculations prove otherwise.
- b. The choice of $\langle \lambda, \alpha \rangle$ pairs in step (6b) may not be fine-grained enough. This would result in gaps in the phase space, which could cause some regions of the phase space to be presumed to be larger than they actually are and potentially miss out on finding some optimal systems. For example, had Schwartz et al. not considered any value of $\alpha < 0.1$, then the vowel system [i ε a u u] would not have been found to be optimal, as can be seen in the phase space diagram in (4).

3.3 Our revised search algorithm

Using the same prototypes and same functions as Schwartz et al., we used a different search algorithm (described in (8), graphically schematized in Fig. 5) for finding optimal systems and mapping out the phase space, to alleviate or avoid the problems in (7).

- (8) a. For each value of $N = 3, \dots, 9$, initialize a catalog K_N of all vowel systems of size N already shown to be optimal by Schwartz et al. anywhere in the $\lambda \times \alpha$ space. For example, as per (4):

$$K_5 = \left\{ \begin{array}{l} [i \ e \ a \ o \ u], [i \ y \ a \ 'o' \ u], \\ [i \ \varepsilon \ a \ u \ u], [i \ 'e' \ a \ 'o' \ u] \end{array} \right\}$$

This initialized catalog only needs to be created once for each N .

- b. For each value of N , randomly sample 5000 $\langle \lambda, \alpha \rangle$ pairs drawn from $[0, 1] \times [0, 1]$.⁵
- c. For each $\langle \lambda, \alpha \rangle$ pair, randomly sample enough candidate vowel systems of size N drawn from the 37 vowel prototypes to have a 99% chance of finding a system that is in the top 0.1% of all possible systems in terms of optimality (lowest energy).⁶ Add to this set of candidates all of the known optimal systems from K_N .
- d. For each $\langle \lambda, \alpha \rangle$ pair and its set of candidate vowel systems, compute the energy of every candidate system, including those from K_N , according to equations (2, 3, 4).
- e. For each $\langle \lambda, \alpha \rangle$ pair and its set of candidate vowel systems, select as optimal the candidate system with the lowest energy. If this optimal system is not yet in K_N , add it. Otherwise, make no change to K_N .

⁵Random sampling in this search algorithm was done using the `runif()` and `sample()` functions in the R programming language (Ihaka & Gentleman 1996).

⁶See Appendix A for mathematical discussion of what constitutes “enough” candidates.

- f. Repeat steps (b)–(e) five times, and then continue repeating them until K_N no longer changes.




Figure 5: Schematic diagram for (8)

3.4 Comparison of results

Fig. 6 shows the phase space diagrams obtained by Schwartz et al.’s original search algorithm (6) and by our improved algorithm (8). The two vowel systems found to




Figure 6: Phase space diagrams for $N = 3$

be optimal by Schwartz et al., $[i \ a \ 'o']$ and $[i \ a \ u]$, are both found to be optimal by

our algorithm as well (indicated by ✓ in the phase space diagram), but in a smaller region of the $\lambda \times \alpha$ space (the volcano-shaped region in the middle, centered around $\lambda = 0.3$). Our algorithm also finds five new vowel systems to be optimal in the remaining regions of the $\lambda \times \alpha$ space. Of the five new optimal systems, [i a o] is optimal only in a tiny sliver of the phase space and is very similar to already predicted [i a ‘o’], so this system is only marginally “new.” The remaining four are not attested in UPSID and seem to be phonetically implausible. Thus, while we do not find any definitively new optimal systems for $N = 3$ that are attested in UPSID, our results show that the parameters λ and α must be limited to the “magic volcano” region, since values outside that region result in unattested, implausible optimal vowel systems.

There are still some remaining 3-vowel systems attested in UPSID that do not yet appear to be optimal in DFT, even with our new search algorithm. For example:

[i æ u]	Shilha
[e a o]	Alabama and Amuesha
[‘ə’ a ‘o’]	Qawasqar

* * *

Interlude: Evaluating the results using UPSID

UPSID is an invaluable resource, but using it to evaluate predictions raises significant challenges. First, there is the problem of **insufficient precision**. When a researcher notates ⟨e⟩, does this represent [e], [‘e’], or [ɛ]? UPSID often uses ⟨‘e’⟩ when the source description doesn’t make the precise height clear, conflating the difference between truly mid and unknown mid. Analogous questions apply to other vowels, such as ⟨u⟩ (does this represent [u], [ʊ], or even [ɯ] in a less careful source?). Furthermore, if a phoneme has multiple different allophones, which one is actually recorded in UPSID?

On the flip side, UPSID also seems to have **too much precision**. There are so many symbols used by UPSID that many vowel systems appear only once. This makes it hard to talk about the frequency of vowel systems. Schwartz et al. (1997b) address this problem by collapsing various kinds of distinctions in UPSID and in DFT’s outputs, using a representative vowel for a set of similar vowels when only one member is present, according to the following scheme:

[i I] → [i]	[ɯ ʉ u ʊ] → [u]	[ø ‘ø’ œ] → [ø]
[y Y] → [y]	[i ɨ ɥ ɰ] → [ɨ]	[ɣ ‘ɣ’ ʌ] → [ɣ]

Both of these problems make it difficult to use UPSID more rigorously. However, even with these problems, UPSID is the best available resource of its kind, so we use it here as our measurement of how well DFT predictions match attested reality, with the understanding that the attestations themselves may be somewhat inaccurate.

* * *

As with $N = 3$, our new search algorithm finds all of Schwartz et al.’s optimal systems for $N = 4$, but not in the exact same regions of the $\lambda \times \alpha$ space (Fig. 7). Three of the four matching vowel systems, which happen to also be attested in




Figure 7: Phase space diagrams for $N = 4$

UPSID, are found in the same magic volcano region, while the fourth matching system [i y a u], lies outside this region and is unattested in UPSID. In addition, our new search algorithm finds six new vowel systems to be optimal in the $\lambda \times \alpha$ space, listed in Fig. 8, labeled with italicized letters that correspond to the phase space regions with matching letters in Fig. 7(b). Two of these new optimal systems,




Figure 8: Newly discovered optimal systems for $N = 4$

^a[y e ε a] and ^b[i e ε a], are unattested in UPSID, which is not surprising given their phonetic implausibility. The third new optimal system, ^c[i e a ə], is not attested directly in UPSID, but it is similar to attested systems, such as Klamath and Tacana’s [i ‘e’ a ‘o’]. The three remaining new systems, ^d[i ε a u], ^e[i ε a u], and ^f[i ə a u], are each attested in UPSID, either exactly (in Murinhapatha, Moxo, and Jebero, respectively) or very similarly (Cayapa’s [i ε a u], Lushootseed’s [ɪ ‘ə’ a u], and Ivatan, Paiwan, and Yupik’s [i ‘ə’ a u]). Thus, our revised search algorithm shows that DFT’s ability to predict extant systems is better than originally thought.

The system ^f[i ə a u] is also interesting because it contains [ə], which calls into question the need for Schwartz et al. (1997b) “transparency hypothesis” for [ə], in which neither the presence nor absence of [ə] plays a role in the acoustic dispersion of a vowel system. Boë et al. (1994) argue that [ə] is difficult to obtain in DFT, but as our results show, there is indeed a region of the phase space (albeit small!)

around $\langle 0.55, 0.01 \rangle$ where an optimal 4-vowel system containing $[\text{ə}]$ is predicted by DFT without modification.

There are still numerous remaining 4-vowel systems attested in UPSID that do not yet appear to be optimal in DFT. For example:

$[\text{i a } \text{ɔ} \text{u}]$	Tiwi	$[\text{i } \text{ɛ a 'o'}]$	Margi
$[\text{i } \text{æ a u}]$	Quileute	$[\text{i } \text{ɜ a } \text{ɹ}]$	Yessan-Mayo
$[\text{ɪ } \text{ɐ a } \text{ʉ}]$	Nunggubuyu	$[\text{'e' 'ə' a 'o'}]$	Upper Chehalis
$[\text{i } \text{i a 'o'}]$	Maranao		

The pattern established with $N = 3$ and $N = 4$ continues with $N = 5$, with a minor twist (Fig. 9). Our search algorithm finds three of Schwartz et al.'s optimal




Figure 9: Phase space diagrams for $N = 5$

systems, again inside the magic volcano, but not $[\text{i y a 'o' u}]$. Since this system lies outside the magic volcano and is not attested in UPSID, it does not seem to be problematic that we do not find it. Again, we also find new optimal systems not found by Schwartz et al.'s search algorithm (Fig. 10). Of the eight new optimal




Figure 10: Newly discovered optimal systems for $N = 5$

systems, the two systems $^a[\text{i y 'e' a } \text{ɹ}]$ and $^b[\text{i } \text{ɪ 'e' a } \text{ɹ}]$ are unattested in UPSID and seem phonetically implausible. Three other systems, $^c[\text{i 'e' a } \text{ɹ u}]$, $^g[\text{i y } \text{ɛ a u}]$, and $^h[\text{i y } \text{æ } \text{ɹ u}]$, while potentially plausible, are also unattested in UPSID. Furthermore, all five of these systems are optimal only in phase space regions outside the magic volcano. The system $^d[\text{i } \text{æ a 'o' u}]$ is not attested directly in UPSID, but it is similar to attested systems, such as Taishan's $[\text{i } \text{æ a } \text{ɔ} \text{u}]$, and with a bit of a stretch, possibly

Nez Perce’s [i æ a ɔ u]. The remaining two systems ^e[i ε a o u] and ^f[i ε a ‘o’ u] are both directly attested in UPSID (Jacaltec and Nasioi; and Batak, Hawaiian, Yucuna, Yagaria, and Baining), again showing that our new search algorithm confirms better predictions for DFT.

There are still numerous 5-vowel systems attested in UPSID that do not yet appear to be optimal in DFT. For example:

[i i a ɔ u]	Papago	[i e ɐ a u]	Koya
[i i ‘e’ a ‘o’]	Abipon	[i ø æ a ʏ]	Hopi
[i i ε a ‘o’]	Cofan	[i ø ε a ʊ]	Malakmalak
[i ε a ɔ o]	Tseshah	[‘e’ æ ɔ u u]	Hixkaryana
[i e a ʊ u]	Kunimaipa		

Full discussion of the results for $N = 6, \dots, 9$ is beyond the scope of this paper, but the raw results themselves can be found in Appendix B.

4 Wrap-up and future work

Our improved search algorithm leads to different results for DFT’s predictions. In particular, we find a greater number of optimal systems throughout the entire $\lambda \times \alpha$ space for $N = 3, \dots, 7$ (see Table 1; note that Schwartz et al. do not give details of their results for $N = 8$ or 9). Many of these newly discovered optimal systems are

N	3	4	5	6	7	8	9
Schwartz et al.	2	4	4	4	5	–	–
Sanders and Padgett	7	10	11	11	9	13	10

Table 1: Total number of optimal systems found for each N

attested in UPSID, which shows that the predictive power of DFT is greater than originally thought.

In addition, we have gathered preliminary statistics on the percentage of systems missing one of the “corner” vowels ([i], [u], and [a]). They show an interesting asymmetry, with [u] generally more likely to be missing and [a] less likely to be missing (Fig. 11; there is also an overall downward trend, which is the result of larger vowel systems being forced to make use of corner vowels more often in order to satisfy dispersion). An articulatory explanation may help explain the difference in behavior among the corner vowels, so an important future extension of this work is to add an articulatory energy component to the basic DFT equation, with more articulatorily extreme vowels having higher articulatory energy, and thus, creating less optimal systems. This would also increase the relative optimality of systems containing vowels like [ə], further obviating Schwartz et al.’s (1997b) transparency rule for [ə].

Finally, further thought needs to be given to precisely what it is that DFT is trying to predict. Are we only concerned with truly optimal systems, or should relative optimality of systems tie into relative *frequency* of those systems? Answering this question is necessary for deciding if and how DFT should be adjusted to also predict “crazy” vowel systems like Qawasqar’s [ə a ‘o’].




Figure 11: Percentage of UPSID systems missing a corner vowel

Appendix A: How many candidates do we need?

In order to have a 99% chance of randomly selecting one of the systems in the top 0.1% of all possible systems in terms of optimality, we need to randomly sample 4,603 candidate systems.

Proof: Let S be the total number of systems to be randomly sampled. The probability that not a single one of them is in the top 0.1% is also the probability that all of them are in the bottom 99.9%. Each one has a probability of 0.999 of being in the bottom 99.9%, so the probability that *all* of them are in the bottom is all S of those 0.999 probabilities multiplied together, which is 0.999^S .

Having all of our sampled systems be in the bottom 99.9% would of course be a bad result, since our hope is to find the one true optimal system, so we want the probability of this event occurring to be as small as possible. A 1% chance of failing to get any optimal system seems acceptable, so we set $0.999^S = 0.01$ and solve for S to see how many systems we need to sample:

$$\begin{aligned}
 0.999^S &= 0.01 \\
 S \cdot \log(0.999) &= \log(0.01) \\
 S &= \frac{\log(0.01)}{\log(0.999)} = 4,603
 \end{aligned}$$

Our actual failure rate should ultimately be smaller than 1%, due to the use of the known optimal systems from the catalogs (because some of these are likely to be the truly optimal systems we hope to find), due to the repetition in step (8f) (because we are sampling a potentially different set of 4,603 systems in each cycle), and due to the pseudo-genetic nature of the search algorithm (because this adds every newly discovered optimal system to the catalog, ensuring that all known strong contenders for optimality must be considered in every subsequent cycle).

Appendix B: Predictions for $N = 6, \dots, 9$




Figure 12: Phase space diagrams for $N = 6$; [i e ε a 'o' u] is missing in (b)




Figure 13: Newly discovered optimal systems for $N = 6$




Figure 14: Phase space diagrams for $N = 7$; [i e ε a 'o' u u] and [i y ε 'o' u u] are missing in (b)




Figure 15: Newly discovered optimal systems for $N = 7$




Figure 16: All predicted optimal systems for $N = 8$




Figure 17: All predicted optimal systems for $N = 9$

References

- Benkí, José R. 2003. Analysis of English nonsense syllable recognition in noise. *Phonetica* 60:129–157.
- Boë, Louis-Jean, Jean-Luc Schwartz, and Nathalie Vallée. 1994. The prediction of vowel systems: Perceptual contrast and stability. In *Fundamentals of speech synthesis and speech recognition*, ed. E. Keller, 185–213. Chichester, UK: John Wiley.
- de Boer, Bart. 2000. Self-organization in vowel systems. *Journal of Phonetics* 28:441–465.
- Carlson, Rolf, Gunnar Fant, and Björn Granström. 1975. Two-formant models, pitch and vowel perception. In *Auditory analysis and perception of speech*, ed. Gunnar Fant and Mark A. A. Tatham, 55–82. London: Academic Press.
- Carlson, Rolf, Björn Granström, and Gunnar Fant. 1970. Some studies concerning perception of isolated vowels. *STL-QPSR* 11:19–35.
- Ihaka, Ross, and Robert Gentleman. 1996. R: A language for data analysis and graphics. *Journal of Computational and Graphical Statistics* 5:299–314.
- Liljencrants, L., and Björn Lindblom. 1972. Numerical simulations of vowel quality systems: The role of perceptual contrast. *Language* 48:839–862.
- Lindblom, Björn. 1986. Phonetic universals in vowel systems. In *Experimental phonology*, ed. John J. Ohala and Jeri J. Jaeger, 13–44. Orlando: Academic Press.
- Maddieson, Ian. 1984. *Patterns of sounds*. Cambridge, MA: Cambridge University Press.
- Maddieson, Ian, and Karen Precoda. 1989. Updating UPSID. *UCLA Working Papers in Phonetics* 74:104–111.
- Mantakas, M., Jean-Luc Schwartz, and Pierre Escudier. 1986. Modèle de prédication du ‘deuxième formant effectif’ F_2^l —application à l’étude de la labialité des voyelles avant du français. In *Proceedings of the 15th journées d’étude sur la parole*, 157–161. Société Française d’Acoustique.
- Schwartz, Jean-Luc, Louis-Jean Boë, Nathalie Vallée, and Christian Abry. 1997a. The Dispersion-Focalization Theory of vowel systems. *Journal of Phonetics* 25:255–286.
- Schwartz, Jean-Luc, Louis-Jean Boë, Nathalie Vallée, and Christian Abry. 1997b. Major trends in vowel system inventories. *Journal of Phonetics* 25:233–253.
- Schwartz, Jean-Luc, and Pierre Escudier. 1987. Does the human auditory system include a large scale spectral integration? In *The psychophysics of speech perception*, ed. M. E. H. Schouten, 284–292. Dordrecht: Martinus Nijhoff Publishers.
- Schwartz, Jean-Luc, and Pierre Escudier. 1989. A strong evidence for the existence of a large-scale integrated spectral representation in vowel perception. *Speech Communication* 8:235–259.
- Stevens, Kenneth N. 1972. The quantal nature of speech: Evidence from articulatory-acoustic data. In *Human communication: A unified view*, ed. E. E. Davis Jr. and P. B. Denes, 51–66. New York: McGraw-Hill.