How to Represent Natural Sound Patterns*

John J. Ohala

Phonology is the study of sound patterns in language, whether these 'sound patterns' are patterns in the acquisition or borrowing of speech sounds, sound changes (or their present-day result, sound alternations between words), etc. The first job in studying sound patterns is to notice them--no easy task--then to develop a system for classifying them, which includes a method for naming them. Like any science, though, phonology must go beyond this taxonomic task and must eventually provide an explanation for the observed sound patterns.

Among the many sound patterns phonologists discover and seek to explain, they find that some are more expected or intuitively more natural than others. For example, the alternation or change of 'k' to 's', that is, velar softening, before a front vowel, is more natural than a change of 'k' to a pharyngeal fricative before a back vowel. However there is nothing in the conventional distinctive feature representation of these processes that indicates this. The notion of markedness was introduced to solve this. Specifically, it was said that the expected or natural value for a phonological feature in a given environment would henceforth be the unmarked value for that feature and the value that would surprise us if we found it, would be the marked value. But this still puts us no closer to a solution. Before markedness was proposed we were given no reason why a given pattern was more expected or natural than another. Now we are told it is due to one being unmarked and the other marked. But these are only labels; we are given no explanation as to what determines the marking of values for features. So, in fact, this relabelling has yielded no net gain in our knowledge: we are still unable to represent or model sound patterns in a way which will make evident the naturalness or lack of it of a given sound pattern.[1] The solution to this problem can be reached--or at least approached--by a re-examination of what is meant by 'natural' with respect to sound patterns and what is meant by 'explaining' a sound pattern.

---

* Paper delivered at winter meeting of the Linguistics Society of America, St. Louis, Mo., Dec. 28-30, 1971. My sincere thanks to Robert Krones for many useful discussions relating to this paper.

[1] Schachter (1969) proposes that the marking conventions be supplemented by a set of naturalness conventions which would specify the natural value, N, of a feature, f, in assimilation rules as follows:

Nf ---> <u>m</u>f / <u>m</u>f (a)

(Continued on the next page)

It should be possible to make explicit the basis for our intuitive judgment of the naturalness of given sound patterns. A sound pattern is natural because it is expected or non-surprising when we discover it in a language. Why is it expected? Because it is encountered so frequently in our studies of diverse languages. But how can it be that

----

That is, "...the natural value of a feature is the marked value of that feature ([$\underline{m}$f]) when an adjoining segment shows the marked value of the feature..." The marking conventions would then apply to specify the actual + or - value of the feature. This is just a further relabelling of the problem of representing the naturalness of sound patterns, and it suffers from the same inadequacies as the marking conventions. We are still led to ask why (a) should be true and not either (b) or (c), below:

$$Nf \dashrightarrow \underline{m}f \;/\; \underline{u}f \qquad (b)$$

$$Nf \dashrightarrow \underline{u}f \;/\; \underline{m}f \qquad (c)$$

Schachter bases (a) on a meta-theoretic constraint (d):

> Unmarked feature values assimilate to adjacent marked feature values rather than conversely. (d)

But it is not obvious that (d) is justified. Schachter posited it in part on the basis of the extreme commonness of assimilatory nasalization, e.g., $V \dashrightarrow \tilde{V} / N$, and the presumed absence (or extreme rarity) of assimilatory denasalization, e.g., m --> b / [oral vowel], the former representing the assimilation of the unmarked value for nasal, [-nasal], to the marked value, [+nasal], in the environment of the marked value for nasal, and the latter the assimilation of the marked value to the unmarked. But assimilatory denasalization does occur often enough to suggest caution before labelling it "extremely uncommon." The assimilatory process represented in (e) occurs in Maxakalí

$$\begin{bmatrix} C \\ +nasal \end{bmatrix} \dashrightarrow [-nasal] \;/\; \underline{\qquad} \begin{bmatrix} V \\ -nasal \end{bmatrix} \qquad (e)$$

(Gudschinsky, Popovich, and Popovich 1970:84), and is well known in many Chinese dialects, especially some of the Min dialects, including Xiamen (Chen 1972). Data from Mohawk (provided by W. Chafe, personal communication) and Hindi (Beames 1970:254) suggest that it may have occurred in these languages. Moreover, the common appearance of epenthetic stops between nasals and following obstruents (see below), e.g., "warmth" [wɔrmpθ], may be viewed as a partial denasalization of the nasal segment in the environment of the following (and necessarily non-nasal) obstruent.

certain sound patterns are encountered so frequently in languages that are unrelated to each other and are distant from each other both geographically and chronologically? The answer must be that these sound patterns are actuated by factors which are universal to all human speakers, namely the universal facts of human anatomy, physiology, and perception. Natural sound patterns, then, are physically-actuated sound patterns. Velar softening, nasalization of vowels before nasals, and all of the other classic examples of natural sound patterns fulfill this definition.

Unnatural sound patterns also occur, of course, but these are rare and less expected and are accordingly actuated by non-universal factors, i.e., by language- and culture-specific factors. Spelling pronunciations, various types of analogy, including hypercorrection, and so on, typically give rise to highly unnatural sound patterns.

Now that we know roughly what it is that we seek to explain, let us consider what it means to explain a sound pattern. Actually, it is not necessary to take up the rather major philosophical problem of what constitutes an explanation. Rather, it is sufficient to point out that there is more than one way to explain something: it all depends on what the question is. We can ask at least two kinds of questions about sound patterns: 1. How did they originate? 2. How are they maintained in the language? It has been pointed out above that one can explain the origin of given sound patterns by referring either to the physical or to the socio-cultural factors which actuated them. In this task then, the question of whether these are natural or non-natural sound patterns is relevant. But to explain how a sound pattern is maintained in a language, one could say perhaps that it is a productive pattern in the language, that is, it is known by the speakers, e.g., the regular English plural, or one could perhaps say that it is simply present in the language due to the sheer weight of usage, but is in no way productive, e.g., the irregular English plurals of 'goose'--'geese', 'foot--feet', etc. The assignment of psychological reality or the lack of it to a given sound pattern in order to explain its retention in a language does not involve the question of naturalness at all.

Most phonologists claim to be interested in describing the psychologically real phonological grammar that is known by the native speaker and seem to assume that natural sound patterns are preferred in this grammar. But it has never been proved that ordinary language users--by which I mean non-linguists--can differentiate, let alone prefer natural sound patterns over non-natural ones. Linguists may know that the alternation between the words smelt/melt, snare/nerve (residue of the famous s-movable in Indo-European) represents an unnatural sound pattern, probably induced by a fact of Indo-European grammar (Hoenigswald 1952), and linguists may know that the alternations due to "velar softening" such as cool/chill, are natural sound patterns, but do the people on the street know this--even tacitly? There is no evidence that they do. Therefore, in our phonological grammars the need to represent natural sound patterns in a way that renders their naturalness self-evident is unnecessary.[2]

For all anybody knows the appropriate representation of such grammars may require the conventional binary distinctive features, feature matrices, etc. Of course, no one does know this for sure, although many people pretend that they do. Only careful psychological tests will reveal the actual form of grammars. However this highly interesting matter of finding out what a speaker knows about the sound patterns in his language and how to represent what he knows is beyond the scope of this paper. In what follows I propose to pursue instead the equally interesting problem of how to represent natural sound patterns, that is, those which are physically-actuated. It is not difficult to show that in this case, where we are describing physical not psychological phenomena, binary features, 2-dimensional feature matrices and the accompanying rule formalism are hopelessly inadequate for the task.

To see why they are inadequate we might briefly review some of the properties of a good model. And let it be clear that it is a model of sound patterns that we are after, not merely a mode of naming or denoting entities. We need a representation, it has truthfully been said, which takes into consideration the intrinsic content of the sounds manifesting the patterns, that is, a system which incorporates the inherent properties of the sounds. A model which accomplishes this must include two things: 1. primitive entities or substantive elements, that is, the basic ingredients of the system studied, and 2. some framework or rules that express the way the primitive entities interact, that is, the so-called formal elements of the model. The latter are typically given in the form of rules, equations, graphs, schematic figures (circuit diagrams), flow charts, etc.

Primitive entities must be independent elements not analyzable in terms of any of the other selected primitives, or, as the mathematicians say, they must be orthogonal with respect to each other. Note that primitive entities need not be primitive in any absolute sense--they may very well be capable of being analysed into something even more elementary. Rather, all that is required is that all the primitives chosen at a given level of analysis be independent relative to each other. Together, these substantive and formal elements of the model should enable us to derive or predict other facts or relations which we observe to be true of our system. When we find, however, certain facts cannot be derived from our model the usual step is to revamp the model--and to continue to revamp it until finally only the observed facts can be derived from it and no non-observed relations can be. Of course, if we desire a model that explains only selected relations we need not include irrelevant details in it. We obviously mention only what is essential to the particular issue at hand. The models used in chemistry provide some convenient

---

[2] The point has also been made by Bach and Harms (1969).

examples of these points. Consider the models (empirical formulas) of ethyl alcohol and that of dimethyl ether, at the bottom of figure 1. The letters stand for atoms of carbon, hydrogen, and oxygen and the subscripts tell how many atoms of each are bonded together to form one molecule of these substances. Here the atoms are the primitive entities or the substantive part of the model and the fact that they can group together in varying (whole) numbers is the formal part of the model. Together these allow us to derive the fact that these substances combine in exactly these proportions no matter how much ethyl alcohol one deals with, quite a significant and useful fact. If the proportion of these elements in these substances is all one wishes to explain then this is an adequate model. But if one wants to explain also the different properties of these two substances--one is a poison, the other is not--then the model is inadequate since the empirical formulas for these two different substances are identical. So it is necessary to escalate and include more detail in the model such as is incorporated in the structural formulas at the top of figure 1. These now show not only the component atoms of the molecule, but also something about how these atoms are joined together. These models, the structural formulas, now help to explain the different properties of the substances. The primitive entities in these models, the atoms, are, of course, not absolutely primitive because for some other purpose or at some other level of analysis we might want to take the atom as a derived entity and consider protons, electrons, and neutrons as our primitive entities, whereas for still another purpose we may want to take methyl alcohol itself as a primitive, say, if we were specifying the recipe for bathtub gin.

Let us turn to phonological models. The first criticism that can be leveled against the conventional distinctive feature matrix as a representation of sounds is that too often the primitive entities, i.e., the small set of distinctive features themselves, are not truly primitive; rather they are frequently dependent on each other. This is most obvious in the case of features such as [±high] and [±low] for tongue body position since a plus value for either requires a minus value for the other. This is also the case with the features [±sonorant] and [±voice], since it is well known that in certain environments non-sonorants, that is, obstruents, tend to induce voicelessness.[3] But

---

[3] In fact in any phonological rule which changes or fills in the value for a given feature F_1 in a feature matrix, given the environment provided by the values of the surrounding features F_2, F_3, etc., F_1 is necessarily non-orthogonal with respect to F_2, F_3, that is, it is dependent on or derived from F_2, F_3. There is nothing wrong with deriving some features (parameters) from others as long as one clearly recognizes what one is doing. Vennemann and Ladefoged (1971) advocate the use of derived features or "cover features" in phonological statements and thus are making explicit the common though unacknowledged practice in phonology for decades. However they go beyond traditional practice by correctly insisting that primitive features and derived features be clearly differentiated.

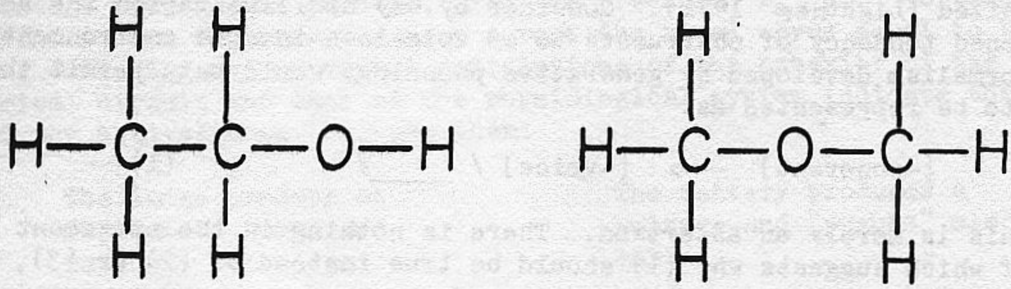

C_2H_6O ETHYL ALCOHOL

C_2H_6O DIMETHYL ETHER

FIGURE I

selecting non-primitive primitives is actually the least damaging criticism of this system since these distinctive feature inventories can be revised and replaced with truly independent features, as indeed, fortunately, they are being revised.

Rather, the truly defective aspect of the distinctive feature system is the extremely limited capacity of the framework to show how these features or parameters interact. The 2-dimensional distinctive feature matrix is scarcely more than a list--no better than the empirical formulas in chemistry. Like them, the distinctive feature matrix and the extremely simplistic rule formalism that goes with it are fine as far as they go, but they just do not go far enough--as many linguists have recognized (Lightner 1971). Consider by way of illustration the above-mentioned tendency of obstruents to be voiceless in some environments. The formalism developed by generative phonology would only permit this fact to be represented as:

[-sonorant] ---> [-voice] / ____# (1)

But this is merely an assertion. There is nothing in the statement itself which suggests why (1) should be true instead of (2) or (3), below.

*[-sonorant] ---> [+voice] / ____# (2)

*[-sonorant] ---> [-voice] / V____V (3)

As shown above, the addition of a marking convention such as (4) is still no help since it still leaves us in the dark as to why that should be true instead of the opposite marking convention, (5).

[u voice] ---> [-voice] / [$\overline{\text{-sonorant}}$]# (4)

*[u voice] ---> [+voice] / [$\overline{\text{-sonorant}}$]# (5)

In fact the probable reasons for the truth of (1) and (4) and the falsity of (2), (3), and (5) are known; the formalism developed by generative phonology just cannot incorporate these reasons in any convenient way: The closure of the vocal tract traps air in the oral cavity causing the pressure above the glottis to approach that below the vocal cords; as the pressure drop across the vocal cords decreases the air flow through the vocal cords falls and thus voicing is inhibited. What is needed, is a way of satisfactorily representing the simple primitive entities of vocal cord configuration, oral air capacity, obstruent closure, and expiratory force and the interactions between these elements, in order to derive the resultant variations in air pressure, air flow, and voicing.

For a first approximation we could take the verbal description mentioned above. This would be difficult to work with, however, since the meanings of words are often not carefully defined and even if they were, the resulting description would be so long and unwieldy that it would be difficult to grasp conceptually. Fortunately there already exists a "language" which is suited to our needs, that is, one which employs carefully defined terms and which is compact. This is the language of electrical circuits and the mathematical expressions associated with them. The relations we need to express are represented in an equivalent way in the electrical circuit depicted schematically in figure 2 (cf. similar analogues in Fant 1960, Flanagan 1965). Here we equate expiratory force with the voltage developed by the battery, air flow with the electrical current (the flow of electrons), glottal resistance with the electrical resistance, the air capacity of the oral cavity with the capacitor (which is capable of storing a certain amount of electrons), and the potential oral constriction with the switch.[4] Further, subglottal air pressure, P_s, is equated with the voltage measured at point A, oral air pressure, P_o, with voltage at point B, and the pressure drop across the vocal cords equal to $V_A - V_B$. The following paired verbal descriptions of the behavior of the electrical circuit and that of the physiological system indicate the nature of the equivalences between them:

The lungs produce an expiratory force and push air out.	The battery produces a voltage and "pushes" electrons out.
The air forced out meets with resistance at the glottis (assuming the vocal cords are adducted for voicing) and thus there is a buildup of subglottal air pressure, P_s.	The electrons are impeded in their passage by the resistance, R, and so some electrons accumulate to the left of R, thus creating a voltage offset at point A.
If there is no obstruction in the oral cavity the air that does pass through the vocal cords simply flows out of the mouth.	If there is no further impedance to the electrons (if the switch is closed), they travel freely past the switch.
However, if the air flow is blocked by an oral constriction, oral air pressure, P_o, begins to build up and in time approaches P_s.	However, if the current is blocked by an open switch, electrons begin to accumulate on the upper plate of the capacitor, C, and soon the accumulation of electrons (the voltage) at point B approaches that at point A.

----

[4] A better analogue of the oral constriction during fricatives would require replacing the switch by a variable resistance.

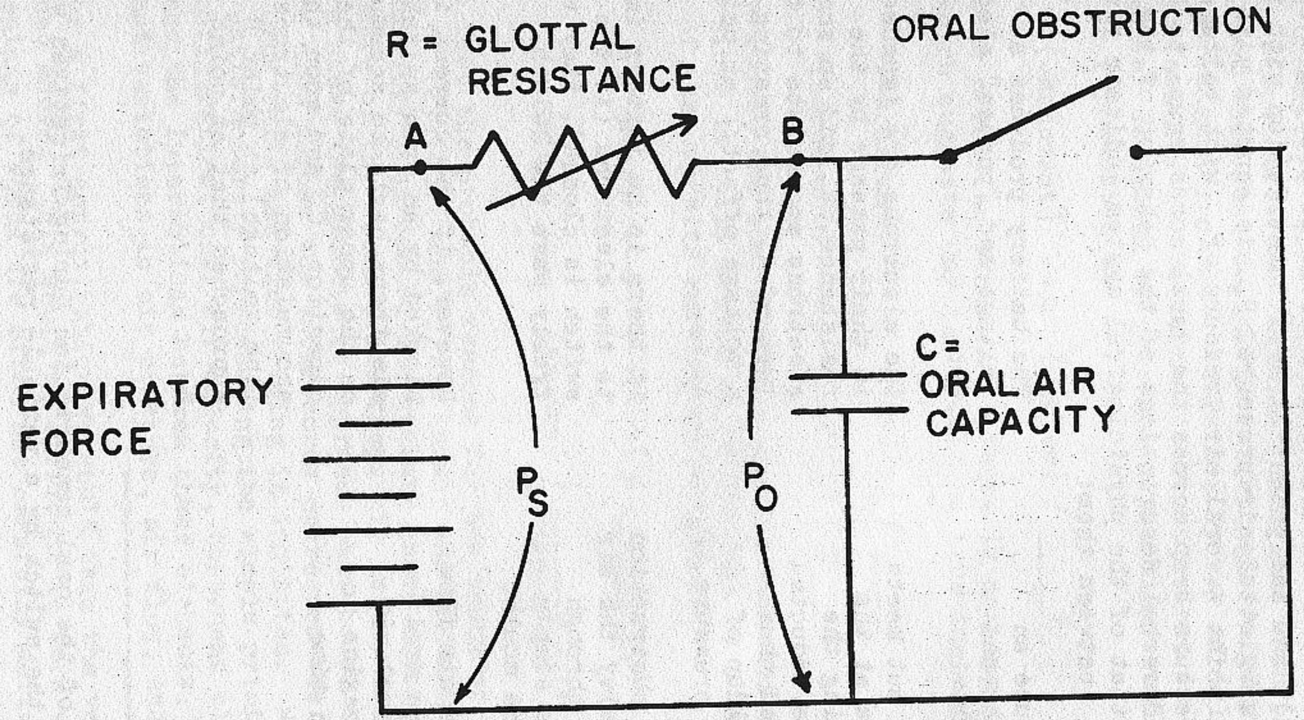
R = GLOTTAL
RESISTANCE
ORAL OBSTRUCTION
A
B
EXPIRATORY
FORCE
P_S
P_O
C =
ORAL AIR
CAPACITY

FIGURE 2

When the air pressure on both sides of the vocal cords are equal, air flow ceases and voicing stops, since voicing requires air flow through the cords. (Actually voicing would cease even before the air flow reached zero.)

When the voltage at point A equals that at point B, current ceases. (A slight embellishment on this model would be required to show how decreasing current stops oscillation; cf. Flanagan and Cherry 1969).

Electrical circuit theory provides equations that summarize the above quantitatively. One such equation would be that in (6) below which predicts how the voltage at B would vary as a function of time, the values of the resistance R, the capacitance C, and the voltage at A, once the switch is open.

$$V_B = V_A\left(1 - \epsilon^{-\frac{t}{RC}}\right) \qquad (6)$$

$(\epsilon = 2.7183)$

The curved solid line labelled $V_B(=P_o)$ in figure 3 presents a graph of equation (6). Plotted on the same graph is $V_A(=P_s)$ and (as the broken curved line) the derived function $\Delta V = V_A - V_B (= \Delta P = P_s - P_o)$, that is, the pressure drop across the vocal cords. If the straight broken line labeled P_{min} represent the minimum pressure drop capable of maintaining voicing, then it follows that the oral obstruction must last at least as long as the time interval T_{min} in order for voicing to be inhibited. Thus the duration of obstruents in final position must generally exceed T_{min}. However it appears that intervocallically, after stress, the closure of obstruents is very short and is probably less than T_{min}, which is why voicing is not inhibited in that environment. Thus, unlike the formalism in (1) and (4), above, the formalism of this model provides an explanation as to why voicing is inhibited in certain environments and not others. This is not to say that this model provides the complete explanation for the devoicing of obstruents finally, because we would still want to find out why obstruents tend to be longer in duration finally than medially. However it should be clear the partial explanation is a genuine one and is not merely a relabelling of the problem.

Let us consider another example: the well-known case of the epenthetic stops that occur between nasal consonants and following obstruents (and vice-versa, Varma 1961, p.123) such as the /p/ in "warmth" [wɔrmpθ] or "something" [sʌmpθIŋ], the /t/ in "pence" [pɛnts] or sense" [sɛnts], the /k/ in "length" [lɛŋkθ] and "strength" [strɛŋkθ], and so on. The underlying physiological cause of these stops can be demonstrated using a device called a nasograph (Ohala 1971a).

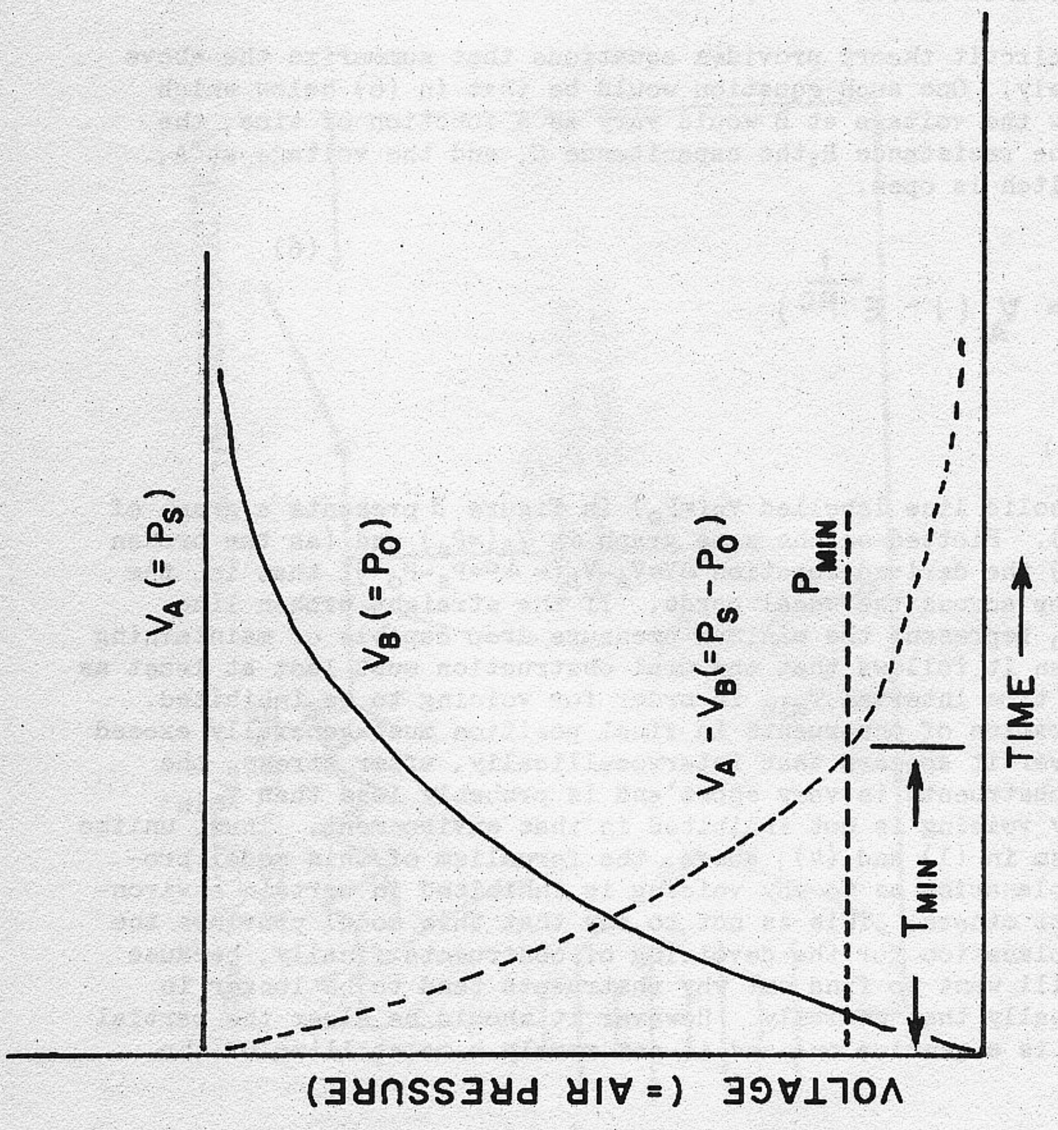

FIGURE 3

(The nasograph (see figure 4) detects movements of the soft palate. It consists of a flexible transparent plastic tube which contains a tiny light and a tiny light sensor. The whole thing is inserted into the subject's nasal cavity and pharynx such that the light is in the pharynx and the light sensor in the nasal cavity. When the velum is lowered, light shines on the light sensor causing it to develop a given voltage; when the velum is elevated the light is blocked and a different voltage is produced by the light sensor, which voltage changes can be fed to an oscillograph or FM tape recorder and related to other speech events. The voltage produced by the light sensor is taken to be approximately proportional to the elevation of the soft palate.)

Figure 5 shows two pronunciations of the name "Samson" first without the epenthetic /p/ and then with the epenthetic /p/. Time progresses from left to right. The traces from the top are: oral pressure sampled by a tube just inside the lips, the microphone signal, and the nasograph signal, with a more elevated velum being indicated by the nasograph signal being near the top of its range. The cause of the epenthetic stop is simple: in uttering a nasal consonant followed by an obstruent requiring a change in articulatory configuration the speaker has to synchronize the elevation of the soft palate with the release of the first oral closure and the formation of the second. The epenthetic stop apparently results from the velum closing too soon thus partially denasalizing the nasal consonant. In figure 5 the dotted lines indicate the approximate duration of the labial closures, and the arrows the approximate moments of complete velic closure. In the first utterance where the closure is synchronized with the labial release, no epenthetic stop results. In the second utterance the velic closure is earlier, thus forming an obstruent (revealed by the buildup of oral air pressure). This is a very common phenomenon and surely counts as a natural sound pattern. A model which incorporates the relevant facts can be devised which is only a slight revision of the model shown earlier.

Figure 6 is essentially the same circuit as that in figure 2, except that now there are two switches representing the two oral constrictions and an extra switch for the velopharyngeal opening. The voltage buildup (the analogue of the air pressure buildup for stops) occurs when all switches are open. This model can also explain the appearance of the epenthetic stops between laterals and following homorganic fricatives, as in "false," [fɔlts], "pulse," [pʌlts], etc., if the two switches are arranged in parallel instead of in series.

Again, this model does not explain everything since we still do not know why the soft palate movement is not always exactly synchronized with the oral articulations, however, except for this, the model renders self-evident the appearance of the epenthetic stops in the given environments as well as their non-appearance in other environments. This is something the marking conventions cannot do.

None of the above models are offered as being particularly original,

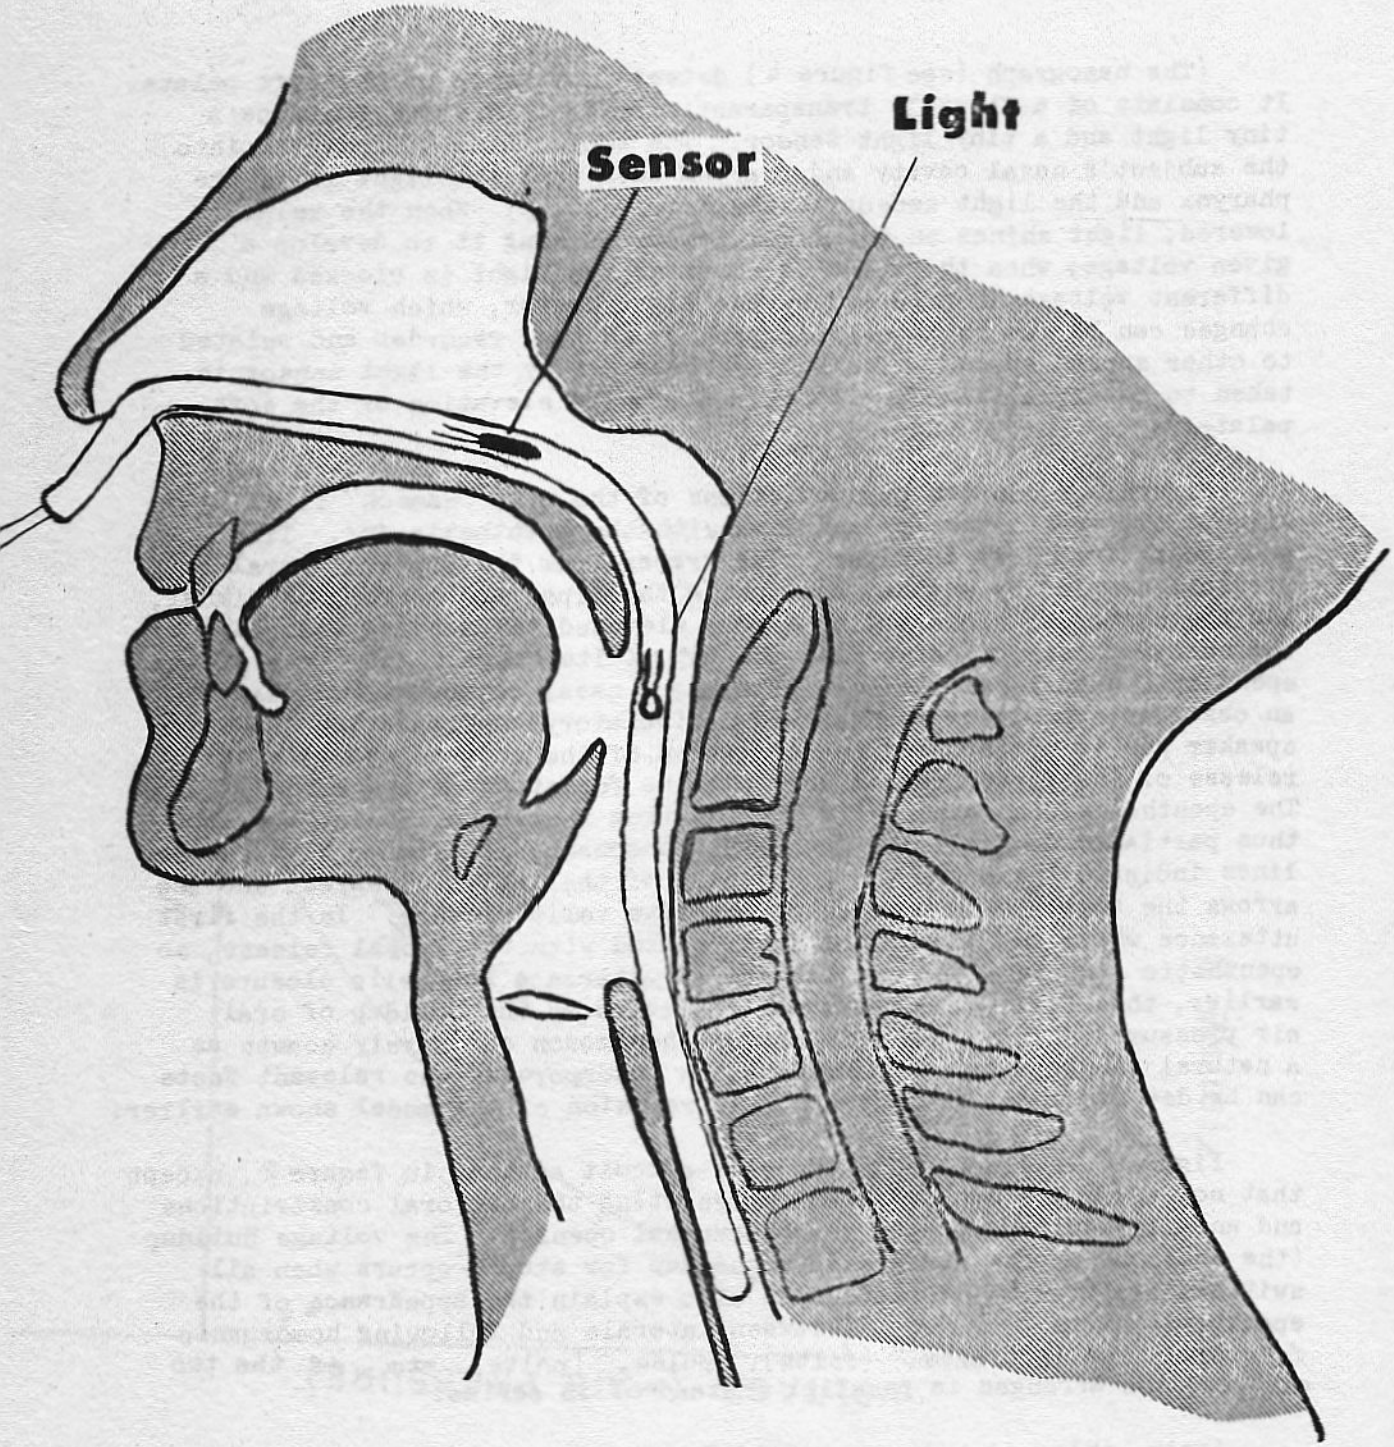

FIGURE 4

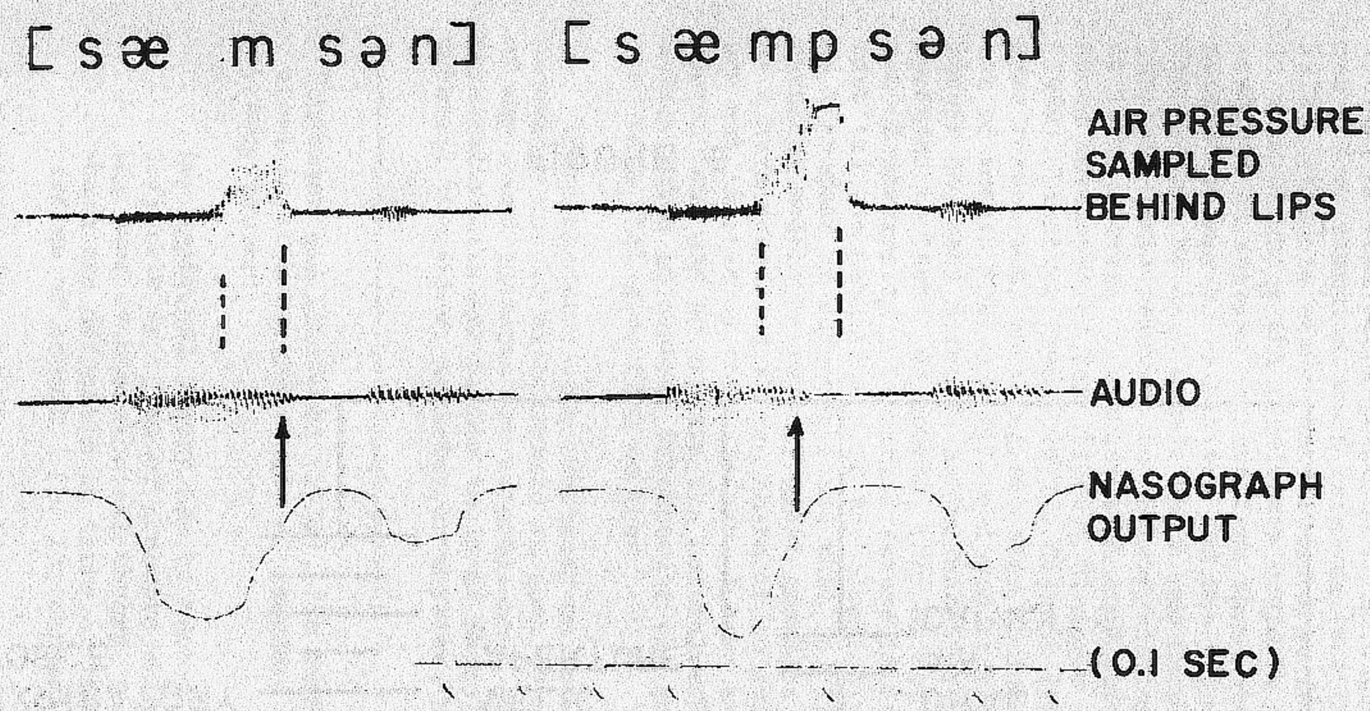
[sæ m sən] [s æ mp s ə n]
AIR PRESSURE SAMPLED BEHIND LIPS
AUDIO
NASOGRAPH OUTPUT
(0.1 SEC)

FIGURE 5

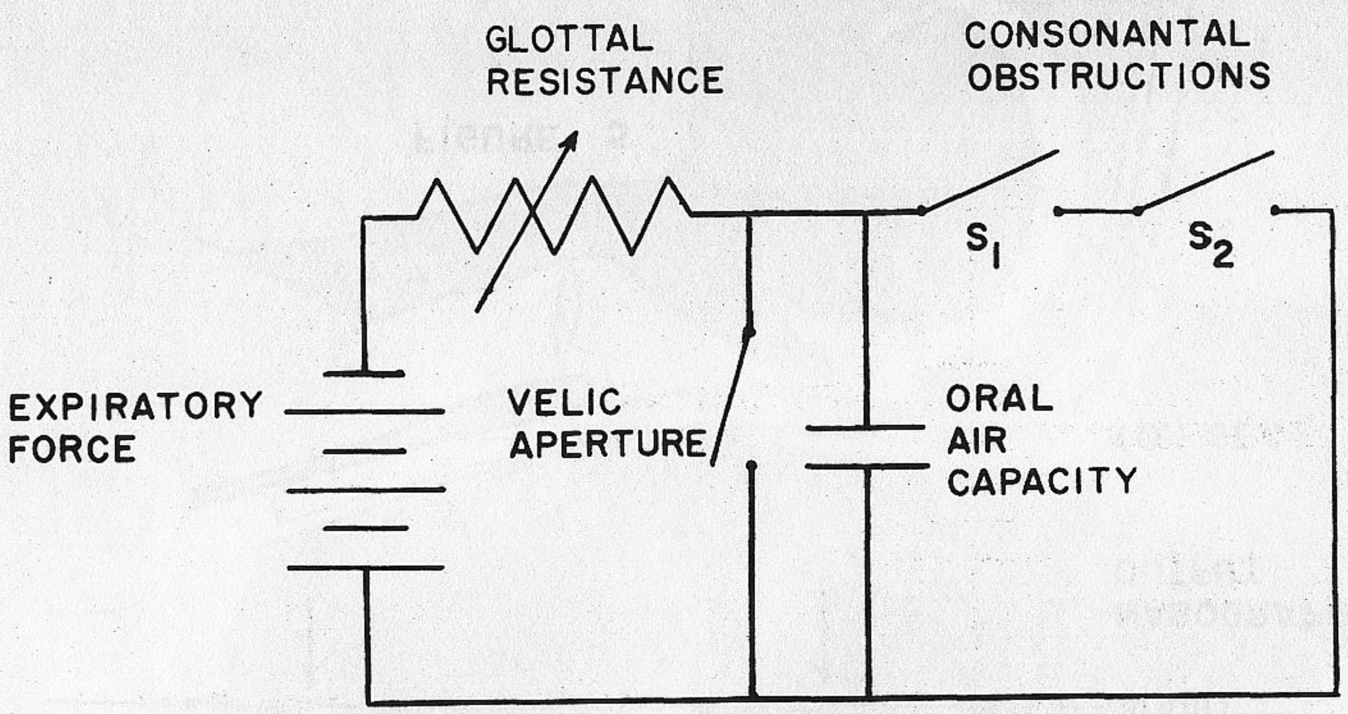
GLOTTAL
RESISTANCE
CONSONANTAL
OBSTRUCTIONS
S_1
S_2
EXPIRATORY
FORCE
VELIC
APERTURE
ORAL
AIR
CAPACITY

FIGURE 6

nor do they reveal any previously unknown facts about speech or sound change. They are presented simply to illustrate the obvious utility of various physiological, acoustic, and perceptual models in representing natural sound patterns in an explanatory way. And, of course, many other examples could be given (see Ohala 1971a, b, c; Chen 1971, 1972). What is needed in fact is a series or collection of models, in some of which the primitives would be physiological entities with acoustic parameters derived from them, whereas in others the acoustic parameters may be primitives and the perceptual parameters derived from them. We need models which can incorporate all and only the physical facts relevant to the explanation of a given sound pattern. It is no secret, of course, that many such models already exist, e.g., those developed by Fant (1960), Flanagan (1965), Flanagan and Cherry (1969), Lindblom (1963, 1971), Stevens (1968), Öhman (1967). But for some reason most phonologists seem to insist on forcing explanations out of models that are ill-equipped to provide such explanations, no matter how many marking, naturalness, or redundancy conventions are appended to them. Of course, they do not reject these better models out of perversity. Rather, I assume it is due to a sincere belief that (a) the naturalness of phonological processes have to be reflected in a grammar of competence and that (b) it is inappropriate to use physical models in psychological grammars. No doubt (b) is true, but, as I have suggested above, there is no evidence that (a) is, and so we are not limited by (b). Rather, I have tried to show that the appropriate way to explain the naturalness of sound patterns is to use models which adequately incorporate the known universal physical processes that give rise to them.

There may also be a hesitancy on the part of some to admit that phonology should be concerned in such detail with answering the question of why sound patterns are natural--or a reluctance to admit that such physical models are proper tools of phonology--correct, perhaps, for physical phonetics, medicine,or physics, but not for phonology.[5] But no

---

[5] Vennemann and Ladefoged (1971), for example, seem to claim that the physical aspects of the sounds of language are not part of phonological theory proper:

> Any empirical theory has to have a number of primitives which are definable in terms of concepts which belong outside the theory. In the case of phonological theory, these are the prime features which are definable in terms of the acoustic or physiological properties of sounds. (p. 13; emphasis added)

I obviously do not think it is necessary or profitable to take such a highly restricted view of the domain of phonology. If phonologists cannot refer directly to the physical properties of the mechanisms which transmit and receive speech, then how will common, natural sound patterns be explained? Certainly not by the empty formalism that has characterized generative phonology and its derivatives to date.

scientific discipline can arbitrarily limit the tools it will use in the pursuit of its goals. If I may paraphrase and embellish a quotation of Karl Popper's: As scientists, we are not students of particular disciplines or students of particular methodologies, we are students of problems (Popper 1963). Our problem is explaining sound patterns and we are compelled to pursue this goal using whatever tools do the best job.

References

Bach, E. and Harms, R. (1969) How do languages get crazy rules? Paper presented at UCLA Conference on Historical Linguistics in the Perspective of Transformational Theory.

Beames, J. (1970) A comparative grammar of the modern Aryan Languages of India. New Delhi: Munshiram Manoharlal.

Chen, M. (1971) Metarules and universal constraints in phonological theory. Project on Linguistic Analysis Reports, 2nd Series (Berkeley) 13:MC1-MC56.

Chen, M. (1972) Nasals and nasalization in Chinese: explorations in phonological universals. Unpublished Doctoral dissertation, University of California, Berkeley.

Fant, G. (1960) Acoustic theory of speech production. The Hague: Mouton.

Flanagan, J. L. (1965) Speech analysis synthesis and perception. New York: Academic Press.

Flanagan, J. L. and Cherry, L. (1969) Excitation of vocal-tract synthesizers. JASA 45:764-69.

Gudschinsky, S. C., Popovich, H., and Popovich, E. (1970) Native reaction and phonetic similarity in Maxakalí phonology. Language 46:77-88.

Hoenigswald, H. (1952) Laryngeals and s movable. Language 28:182-85.

Lightner, T. (1971) Generative phonology. In: Dingwall, W. O. (ed.) A survey of linguistic science. University of Maryland, pp. 498-574.

Lindblom, B. (1963) Spectrographic study of vowel reduction. JASA 35: 1773-1781.

Lindblom, B. (1971) Numerical models in the study of speech production and speech perception: some phonological implications. Paper delivered at the 7th International Congress of Phonetic Sciences, Montreal, Plenary session 4.

Ohala, J. (1971a) Monitoring soft palate movements in speech. JASA 50:140. Printed in full in: Project on Linguistic Analysis Reports, 2nd Series (Berkeley) 13:J01-J015.

Ohala, J. (1971b) The role of physiological and acoustic models in explaining the direction of sound change. Project on Linguistic Analysis Reports, 2nd Series (Berkeley) 15:25-40.

Ohala, J. (1971) Physical models in phonology. Reports of the 7th International Congress of Phonetic Sciences, Montreal, Canada, 22-28 August 1971 (in press). Reprinted in: Project on Linguistic Analysis Reports, 2nd Series (Berkeley) 15:41-45.

Öhman, S. (1967) Numerical model of coarticulation. JASA 41:310-320.

Popper, K. (1963) Conjectures and refutations. Routledge and Kegan Paul.

Schachter, P. (1969) Natural assimilation rules in Akan. IJAL 35: 342-355.

Stevens, K. N. (1968) Acoustic correlates of place of articulation for stop and fricative consonants. Q. P. R. (M.I.T.) No. 89:199-205.

Varma, S. (1961) Critical studies in the phonetic observations of Indian grammarians. Delhi: Munshiram Manoharlal.

Vennemann, T. and Ladefoged, P. (1971) Phonetic features and phonological features. Working Papers in Phonetics, (UCLA) 21:13-24.