



PROJECT MUSE®

Phoneme inventory size and population size

Jennifer Hay, Laurie Bauer

Language, Volume 83, Number 2, June 2007, pp. 388-400 (Article)

Published by Linguistic Society of America

DOI: <https://doi.org/10.1353/lan.2007.0071>



➔ *For additional information about this article*

<https://muse.jhu.edu/article/218148>

SHORT REPORT

Phoneme inventory size and population size

JENNIFER HAY

University of Canterbury

LAURIE BAUER

Victoria University of Wellington

This short report investigates the relationship between population size and phoneme inventory size, and finds a surprisingly robust correlation between the two. The more speakers a language has, the bigger its phoneme inventory is likely to be. We show that this holds for both vowel inventories and consonant inventories. It is not an artifact of language family.*

1. INTRODUCTION. Some researchers have previously speculated that there might be a link between the number of speakers of a language and how many phonemes that language is likely to have (Haudricourt 1961, Trudgill 2002). However, to our knowledge, no one has ever reported a statistical correlation between these two things. We don't find it particularly surprising that such a link has not been systematically investigated, as this is certainly not an association one would necessarily expect. However, in the process of proofreading a manuscript that contained information about a series of languages, including their populations and vowel inventories, it struck us that there appeared to be some connection between them. We couldn't resist checking this apparent link more systematically. In this short report we show that there is, indeed, an association between phoneme inventory and population size. We do not have well-developed arguments to offer about why this should be. The correlation seemed intriguing enough, however, that it was worth simply presenting the result, and leaving it up to readers to draw their own conclusions.

2. MATERIALS. Bauer 2007 is a handbook designed to provide a range of useful information for linguistics students. One part of the handbook is a list of some 250 languages with summary information about each, including its language family, where it is spoken, how many speakers it has, and typological features such as the relative order of subject, verb, and object; of noun and adjective; and so on. One piece of information that was collected for these purposes was the number of vowels each language is said to have.

When we noticed an apparent association between population size and vowel inventory size, we decided to supplement this information with information about the consonant inventory as well. Since not all of the original materials were still readily available, it proved difficult to collect all the relevant data here, and we were not able to collect information on consonants for as many languages as we had for vowels. We also decided to exclude from our analysis any language that did not have any living speakers. We therefore ended up with full population, vowel, and consonant information for a set of 216 languages.

While this is not a random sample, we think it is a reasonably representative one. An ideal approach to such a sample might be to do random selection from some list such as the *Ethnologue* (Grimes 1988); however, this would prove impractical in terms of gathering the required additional information, and would not have been appropriate for the original purpose of the selection in Bauer 2007. For the purposes of the book,

* We are grateful to Harald Baayen, Ann Bradlow, Brian Joseph, Christian Langstrof, Peter Trudgill, Paul Warren, and the *Language* referees for their comments and suggestions, and to Vladimir Pericliev for generously sharing his data with us.

languages were chosen to provide a geographical and genetic spread of languages, while including languages that students might be expected to or want to know something about.

Thus, 'big' languages like English, Hindi, and Mandarin were chosen, and 'small' but linguistically well-known languages like Basque, Diyari, and Hixkaryana were also selected. Languages that were not well described in works easily available in accessible libraries stood very little chance of being selected. Thus, the set of languages selected is somewhat biased toward 'big' languages and toward Indo-European and Pacific languages (because we are in New Zealand), but covers a range of languages from around the world. While the selection was not random, we cannot identify anything in the selection process that would have introduced an artifactual correlation between population size and phoneme inventory.

The phoneme inventory counts are taken directly from other linguists' analyses. No additional analysis of languages' phoneme inventories has been conducted by us. Thus, when we described the vowel information as representing how many vowels a language is 'said to have', the wording was deliberately careful. However much we may believe in the nonuniqueness of phonemic solutions (Chao 1957 [1934], Port & Leary 2005), it is surprising to see just how different two descriptions of the same language can be. For this reason, we have taken care to consider various subsets of the phoneme inventories separately, so that we could apply appropriate caution to any effect carried by parts of the inventory that are particularly prone to variation across analysts.

Thus, we distinguished between 'basic monophthongs', which differ in quality only, and 'extra monophthongs', which consisted of nonquality distinctions, such as length and nasalization. The basic monophthong counts are likely to be much more consistent across analysts. As Maddieson (2005) points out, the languages for which length and nasalized forms are listed as separate phonemes cannot necessarily be relied upon, since when analysts are considering whether such a distinction is phonemic 'the considerations which would lead to making one choice or the other are often finely balanced and lead different scholars to different conclusions' (Maddieson 2005:14). Maddieson excludes such forms from his analysis. We list them separately, and regard the counts with appropriate caution. Diphthongs were also listed separately, and here the analysis is even more open to interpretation, since diphthongs may be analyzed as independent phonemes, as sequences of vowel and glide, or as sequences of nonidentical vowels.

In terms of consonants, a distinction was made between obstruents (including plosives, affricates, implosives, ejectives, clicks, and fricatives) and sonorants (including nasals, liquids, and glides), but again some caution is required. The numbers here were also often different from one analysis to another.

3. ANALYSIS. Before conducting any statistical analysis, we inspected the consonant and monophthong inventory sizes for outliers and removed two extreme outlier languages (i.e. ones showing values more than four standard deviations above the mean). These were !Xu (for total consonants) and Acooli (for total monophthongs). We also took the log of the population size, in order to minimize the effect of outliers.

The left panel of Figure 1 shows the positive correlation between the log population of speakers of a language and their basic monophthong inventory, that is, including quality distinctions only. The right panel repeats the correlation, this time also including additional phonemic vowel differences such as length and nasalization. The correlation involving basic monophthongs is much tighter; if there turns out to be some causal relationship between population size and number of vowels, then this tighter correlation

for basic monophthongs might be attributed to the greater consistency here across analysts.

Any correlation regarding diphthong inventory may be very unreliable, due to the inherent role of the interpretation of the analyst. We note in passing, however, that population size is also well correlated with the number of diphthongs listed by analysts ($\rho = .28, p < 0.0001$).




FIGURE 1. Association between population size and vowel inventory. Each point represents a language. Left panel shows the inventory of basic monophthongs, distinguished by quality. Right panel shows the inventory including other distinctions analyzed as phonemic, such as nasalization and length. The line shows a nonparametric scatterplot smoother fit through the points (Cleveland 1979).

Figure 2 shows the correlation of population size with the size of a language's obstruent inventory, sonorant inventory, overall consonant inventory, and overall phoneme inventory. All of these return significant correlations. As described above, we have removed two languages from the sample because they fall more than four standard deviations from the mean. This was because we were worried that they would exert undue influence on the statistics. We have also checked all of the above correlations with the two languages included. All correlations remain significant.

It is important to note that what we are seeing here is an overall statistical tendency. The graphs in Figs. 1 and 2 all show a reasonable amount of scatter, reflecting the fact that there are individual languages that go against this general trend. Faroese, for example, has just forty-five thousand speakers, but twenty-one obstruents, eighteen sonorants, six monophthongal vowel qualities plus a length distinction, and eight diphthongs, many of which also show a long/short distinction (see Thráinsson et al. 2004).

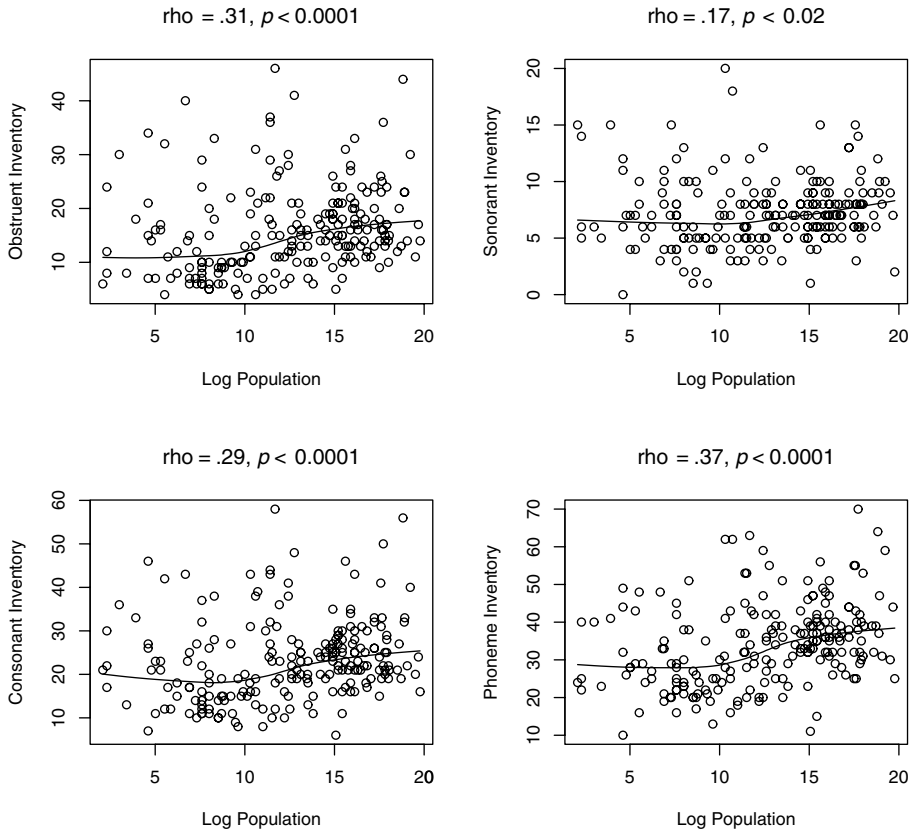


FIGURE 2. Association between population size and inventories of obstruents (top left), sonorants (top right), all consonants (bottom left), and all phonemes (bottom right). Each point represents a language. The line shows a nonparametric scatterplot smoother fit through the points (Cleveland 1979).

That vowel inventory and consonant inventory are both correlated with population size is quite remarkable. This is especially so because consonant inventory and vowel inventory do not correlate with one another at all in this data set ($\rho = -.01, p = 0.86$). Maddieson (2005) also reports that there is no correlation between vowel and consonant inventory size in his sample of 559 languages. Despite the fact that there is no link between vowel inventory size and consonant inventory size, both are significantly correlated with the size of the population of speakers.

We were suspicious that some of this trend might be carried by an association between language families and population size. Australian languages, for example, tend to have small vowel inventories and small populations. Indo-European languages tend to have larger vowel inventories and large populations. The fact that vowel and consonant inventory are not correlated at all, but that both correlate with population, suggests that language family is not the whole story here. That is, if the relationship between population size and both the consonant and vowel inventories were solely an artifact of language family, we might expect the consonant and vowel inventories to correlate with one another, and they do not. Nonetheless, it seemed important to assess the role that language family may be playing.

Each language in our database has been coded for the language family it belongs to. Given how controversial some language families are, it might seem that our analysis here could be skewed by the particular choice of labels for language families. The classification of the various languages was based on the data available in the sources for Bauer 2007. Since some of these sources were much older than others, there was some variation in the names of language families, and in the attribution of individual languages to particular families. However, since we were not concerned with the minutiae of the classifications, but with the top-level classifications, these variations were relatively easily eliminated on inspection and we believe the classification we used to be fairly robust.

In order to assess the significance of population while factoring in language family, we fit an ordinary least-squares linear-regression model, predicting the total phoneme inventory. We included language family as an independent variable, identifying all families for which we had seven or more languages represented, and classifying all other families as ‘other’. The threshold of seven languages was chosen so as to represent a reasonable number of language families in the model (seven language families met this threshold), while still retaining an appropriate number of degrees of freedom in our model. The model statistics are shown in Tables 1 and 2.

FACTOR	D.F.	PARTIAL SS	MS	F	p
Family group	6	3534.326	589.0543	6.4	< 0.0001
Log population	1	699.9648	699.9648	7.61	0.0063
REGRESSION	7	6415.775	916.5392	9.96	< 0.0001
ERROR	206	18952.06	92.00028		

TABLE 1. Analysis of variance for ordinary least-squares model predicting phoneme inventory size.

	VALUE	STD. ERROR	t	Pr(> t)
Intercept	32.0872	4.1607	7.712	0.0000
family = Altaic	−4.5173	4.6882	−0.9635	0.3364
family = Austronesian	−13.4739	3.5936	−3.7494	0.0002
family = Indo-European	1.689	3.2953	0.5125	0.6088
family = Niger-Congo	−0.8249	3.7932	−0.2175	0.8281
family = other	−5.0686	3.3043	−1.534	0.1266
family = Penutian	−5.2322	4.6232	−1.1317	0.2591
Log population	0.4718	0.1854	2.544	0.0117

TABLE 2. Coefficients for ordinary least-squares model predicting phoneme inventory size. (Residual standard error: 9.512 on 206 degrees of freedom; adjusted R²: .24)

Figure 3 plots the predictions of the model (overall model $r^2 = .24$). Language family does indeed have a significant influence, with Indo-European languages having the largest phoneme inventories, and Austronesian languages having smaller phoneme inventories (top panel). In addition to language family, however, the log population of speakers is a separate, significant predictor. This is shown in the bottom panel of Fig. 3, which plots the predicted effect of population size, while holding language family constant. Our sample of languages is, of course, not random. It is a function of the particular languages for which the relevant information happened to be readily available to us. It would be very difficult to create a fully random sample. This would involve taking all known languages, randomly sampling them, and then setting out to find population size and phoneme inventory for each selected language. In some cases this would require extensive research, including fieldwork. While this is not an impossible agenda, it is an extremely ambitious one. We chose to use a statistical method in its place.




FIGURE 3. Predictions of model predicting phoneme inventory size: the effects of language family (top panel) and population size (bottom panel). Dashed lines show 95% confidence intervals.

In order to assess the degree to which the specific selection of languages under investigation is responsible for the significance of the regression model, we conducted bootstrap validation of the model, using the ‘validate’ function from Harrell’s Design library in R (as described in Harrell 2001). The validation technique involves refitting the model over random subsamples of our data. About 63% of the languages are included in each random sample, together with replacement. That is, the random samples are all the same size as the original sample. They contain fewer languages, but some of these languages are represented more than once.

An automatic backwards step-down variable selection procedure is employed. This is repeated two hundred times. In all two hundred iterations, language family was retained as a significant predictor. Log population was retained in 180. This provides good evidence that the significance of the model is not due to the inclusion of any specific languages in our sample. The average R-squared value across these 200 models is .22.

This seems to provide some reassuring evidence that the observed correlation with population size is not due to the specific collection of languages included in our overall sample, nor is it due to an influence of language family. However, there is still a worrying element regarding the role of language family in the model reported above: 44% of our languages fall into the ‘other’ group. That is, 44% of the languages belong to a language family that is represented by fewer than seven languages in our sample. The model presented in Tables 1 and 2 certainly demonstrates that the observed difference is not due to differences between large family groups in the sample. For example,

there is a difference between Indo-European (many speakers, many phonemes) and Austronesian languages (few speakers, few phonemes), but this is not carrying the effect. But the potential effect of other, smaller, family groups may still be influencing the correlation.

As a final check, then, we decided to investigate how this correlation holds across family groups. That is, we removed the potential for undue influence by specific language families, by reducing each language family to one data point.¹ We did this by calculating the mean population size for each language family and the mean number of phonemes. There are forty-two language families in the sample.

Table 3 shows the Spearman's correlation between mean population size and mean inventory size for different parts of the phoneme inventory. There is a significant effect in most parts of the phoneme inventory. One exception is the sonorants—this is the subset of the consonant inventory that showed the weakest effect in Fig. 2. The other is the full monophthong count including the 'extra' monophthongs. This parallels the weaker effect within this subset of phonemes already observed in Fig. 1 and can be ascribed to the high degree of variability across analysts in terms of what might count as an 'extra' monophthong.

CORRELATION WITH	RHO	<i>p</i> <
All phonemes	.46	0.003
Basic monophthongs	.47	0.002
All monophthongs	.2	0.2
Diphthongs	.53	0.001
Plosives	.33	0.05
Fricatives	.53	0.001
Sonorants	.24	0.13
Obstruents	.43	0.005
All consonants	.45	0.005

TABLE 3. Spearman's correlation between mean language family population and mean inventory size.

Figure 4 shows plots for two of these correlations: the correlation for the basic monophthong inventory (top panel), and the correlation for the full consonant inventory (bottom panel). The points are plotted with the names of the language families, to give some sense of how the different language families are distributed across the space.

In these correlations, each language family is reduced to a single point. This analysis therefore eliminates the possibility that the observed correlation between population size and inventory is being carried by one or two overrepresented language families in our sample. Again, within this sample there is absolutely no correlation between the mean number of vowels a family has and the mean number of consonants ($\rho = -.03$, $p = 0.87$). From Fig. 4 we can see that there are some language families that are low on both consonants and monophthongs (e.g. Chibchan, Papuan), some that have many monophthongs and few consonants (Je), and vice versa (Caucasian). That there is no statistical relationship between vowel inventory and consonant inventory makes the fact that both correlate with mean population even more remarkable.

A referee suggested that perhaps the number of 'phonemes' in a language tends to increase as the language is studied, and that languages spoken by more speakers tend to receive more attention. This is an intriguing suggestion, which would provide a

¹ We owe thanks to Harald Baayen for suggesting this approach.




FIGURE 4. Mean log population and mean basic monophthong inventory (top panel) and consonant inventory (bottom panel). Each point represents a language family.

sociological explanation for our observed effect. However, this portion of our analysis has reduced each language family to one data point—investigating the mean population size and phoneme inventory by language family. This eliminates the possibility that a few highly studied language families (such as Indo-European) might be driving the effect in this way. The fact that the correlation is robust across both languages and language families suggests strongly that there is something here that requires explanation.

4. DISCUSSION. There is a surprisingly strong relationship between the size of a language's phoneme inventory and the number of speakers of that language. While such a correlation has not (to our knowledge) been reported before, some researchers have, in fact, speculated that there may be a link between the size of a community and the phoneme inventory.

For example, Haudricourt (1961, as cited in Trudgill 2002) has argued that small inventories are the result of 'impoverishment', which occurs in situations characterized by monolingualism, isolation, and/or by 'non-egalitarian bilingualism'. Haudricourt suggests that in certain environments a particular group may be sufficiently dominant that they have no motivation to articulate clearly. Such people are able to confuse two phonemes, or omit a phoneme without fear of 'mocking'. 'This is why we find fewer consonants in the language of the Iroquois who terrorized their neighbors, or in the languages of the people of Tahiti and Hawaii who combine island isolation with significant demographic development as compared to other less favoured archipelagos' (Haudricourt 1961:10, as cited in Trudgill 2002:720; Trudgill's translation). Trudgill appears relatively unconvinced by Haudricourt's interpretation, but does agree with the general prediction that isolated communities may have smaller inventories. This would be so because 'initial small community size . . . would have led in turn to tight social networks, which would have implied large amounts of shared background information—a situation in which communication with relatively low level of phonological redundancy would have been relatively tolerable' (Trudgill 2002:720). The factors Trudgill (2002) believes may lead to small phoneme inventory size include isolation from contact with other languages, initial small community size, tight social networks, and large amounts of shared background information. This hypothesis suggests that, in a fuller investigation, one should perhaps also attempt to take into account the degree to which each language is isolated from contact with other languages. This may well be partially correlated with population size and responsible for some of what we have found. However, this would be a much more major undertaking, and beyond our much more modest goals.

In addition to small community size potentially leading to small phoneme inventories, Trudgill claims that small communities can also lead to very large inventories. This is because of 'the ability of such communities to encourage continued adherence to norms from one generation to another, however complex they may be' (Trudgill 2004a:317). Thus, Trudgill claims, the combined effects of isolation, network structure, and language contact should lead languages with small populations to have either very small or very large inventories, and languages with larger populations to favor 'medium-sized inventories' (2004a:317). While we have clearly found some evidence of smaller populations favoring smaller inventories, there is no evidence in this data set that they also favor larger inventories.

In a commentary on Trudgill's (2004a) paper, Pericliev (2004) investigates the relationship between consonant inventory and population size in a set of 417 languages.

He examines the data in various ways and concludes that 'there is no correlation of the kind suggested by Trudgill between the size of a community speaking a language and the size of the consonantal inventory of that language' (2004:382). As his purpose is to provide counterevidence to Trudgill's theory that larger populations should lead to medium-sized inventories, Pericliev does not actually test for a straightforward positive correlation between population size and inventory. However, he has kindly shared his data with us, and we have tested this correlation. It is highly significant (Spearman's $\rho = .21, p < 0.0001$). That the overall correlation between phoneme inventory and population size is significant in this larger sample provides strong evidence that the observed correlation is not an artifact of our sampling procedure. Pericliev's sample was collected entirely independently and includes a different (larger) sample of languages. Yet it also contains the same correlation that we have observed.

In defending his thesis against Pericliev (and other commentaries), Trudgill argues that the effects of population size, network structure, and language-contact situation need to be considered together, and so there would be 'no reason at all to expect to find a simple correlation between the numbers of speakers in a language and the number of phonemes in that language' (2004b:386). While we agree with Trudgill that there is no obvious reason to expect such a correlation, the data we discuss in this short report certainly suggests that such a correlation exists.

One possible explanation for the correlation may come from issues relating to learnability. It would be a large leap to assume that speakers of languages with smaller populations are exposed to a narrower range of speakers (and/or dialects) than speakers of languages with larger populations. After all, each individual speaker of a language certainly does not necessarily interact with every other speaker. It is a tempting leap, though, because if this were true, it would suggest an explanation in terms of the robustness and learnability of categories based on this different exposure.

Experiments designed to teach nonnative phoneme distinctions show better learning, and considerably better long-term retention, if multiple voices are used in training (Lively et al. 1993, 1994, Logan et al. 1991). Results on listener adaptation to foreign-accented English also demonstrate that 'exposure to talker variability also facilitates rapid, talker-independent perceptual learning of a foreign accent which involves a wide range of acoustic-phonetic features' (Bradlow & Bent 2003:2884).

Such results tempt one to speculate that exposure to less variability would lead to less robustness of phonemic categories. Exposure to variability is important, as 'variability causes the need for abstraction' (Pierrehumbert et al. 2001).

The learning of phonemes involves abstraction over learned distributions of speech sounds (see e.g. Pierrehumbert 2000). The more exposure to more different speakers, the denser these distributions presumably are. Work on the acquisition of phoneme categories (Maye & Gerken 2000, Maye et al. 2002, Maye & Weiss 2003) shows that infants use distributional information in the signal to discern phoneme boundaries. That is, when an infant (or adult) is exposed to tokens from a particular phonetic space in a unimodal distribution, they tend to learn this as a single category. When a distribution over the same phonetic space is bimodal, it is learned as two categories. Increased exposure to a large number of speakers would lead to denser distributions and so (presumably) make learning of this kind more robust. With sufficient exposure, categories could be easily learned that would be difficult with more limited, less varied, exposure.

Variability facilitates the learning of categories, and repeated prolonged exposure also sharpens the boundaries of these categories. Lee and colleagues (1999) show that

phoneme boundaries sharpen considerably right up until late childhood. Pierrehumbert (2001) builds an effect of 'entrenchment' into her exemplar-theoretic model in order to simulate this kind of effect: categories become sharper with repeated exposure.

The effects of the size of a population of speakers are also revealed by multiagent modeling work. For example, Bart de Boer (2000, 2001) has done work constructing computer-based models that simulate the emergence and transmission of vowel systems. The models work by simulating a population of speakers that can produce and perceive vowels. After a series of iterations in which the agents attempt to imitate one another's productions, vowel systems that resemble human vowel systems emerge in the population. While investigating the properties of the parameters of his model, de Boer (2000) manipulates the size of the population, finding that in his modeling 'the success of all population sizes is comparable, but the vowel system size of small populations is smaller than that of large ones, reflecting the lower stability' (459). The effect of interaction among a large number of speakers is to increase the stability of systems with many vowels. Of course, to suggest that this is responsible for our correlation supposes that the number of speakers of a language is somehow correlated with the number of different speakers an individual is exposed to over the course of his or her lifetime. While this may be true, it is not necessarily so. And the degree to which de Boer's prediction would carry over to consonants is not so clear; most of this kind of work has focused on vowel systems.

Regardless of the explanation for the correlation, its existence raises a number of further empirical questions. One is the question of whether the relationship between population size and phoneme inventory is also relevant across dialects of individual languages. Are dialects spoken by fewer speakers also likely to have fewer phonemic categories? In addition, if there is a causal relationship between population size and phoneme inventory size, then we might also expect to see covariation of population size and phoneme inventories over time.² Neither the population of speakers of a language nor its phoneme inventory is constant. Some languages undergo drastic changes in their population of speakers due to catastrophic factors such as disease or genocide. The results presented here raise the question of whether such changes in population have a tendency to lead to changes in phoneme inventory size. Such changes need not be cotemporal, of course—changes in phoneme inventory size may lag considerably behind population trends. It would, of course, be a surprising finding if fluctuations in population size were indeed paralleled in the phonemic system, and we would not necessarily want to commit to such a prediction. But our results point to this as an intriguing question.

5. CONCLUSION. We have reported a positive correlation between how many phonemes a language has and how many speakers it has. This correlation exists both within the vowel inventory and within the consonant inventory. This is not an artifact of language family. We do not know what the underlying causes of this correlation are. But it is certainly intriguing, and we hope that this report will generate some discussion of the possible causes of such a relationship.

APPENDIX: LANGUAGES IN DATABASE.

!Xu, Abkhaz, Acooli, Afrikaans, Akan, Albanian, Amele, Amharic, Amoy, Apalai, Arabana-Wangganguru, Arabic, Armenian, Arrernte, Basque, Bengali, Berber, Blackfoot, Breton, Bulgarian, Burmese, Burushaski, Canela-Kraho, Cantonese, Cashinahua, Catalan, Cherokee, Cheyenne, Chipewyan, Chukchee, Cree, Croatian,

² Our thanks to Brian Joseph for this observation.

Crow, Czech, Dakota, Dan, Dani, Danish, Daur, Dinka, Diyari, Dutch, Dyirbal, Efik, English, Erromangan, Estonian, Evenki, Ewe, Faroese, Farsi, Fijian, Finnish, Fore, French, Friesian, Fula, Gaelic, Georgian, German, Gilyak, Greek (modern), Guarani, Gujarati, Haida, Hausa, Hawaiian, Hebrew, Hindi, Hixkaryana, Hopi, Hungarian, Icelandic, Igbo, Ijo, Illocano, Indonesian, Irish, Italian, Japanese, Javanese, Kabardian, Kamba, Kambera, Kannada, Kanuri, Karen, Kashmiri, Ket, Khmer, Kilivila, Kiowa, Kirghiz, Klamath, Kobon, Koiari, Korean, Kota, Kpelle, Kurdish, Kwakwaka, Laotian, Latvian, Lenakel, Lithuanian, Luiseño, Maasai, Madi, Maidu, Malagasy, Malay, Malayalam, Maltese, Mam, Mandarin, Maori, Marathi, Margi, Mari, Mazateco, Meithei, Mende, Miwok, Mixtec, Mongolian, Nahuatl, Nama, Navajo, Nez Perce, Ngiti, Nootka, Norwegian, Ojibwa, Oneida, Oromo, Ostyak, Panjabi, Papiamentu, Pashto, Pima, Pirahã, Pitjantjatjara, Polish, Pomo, Portuguese, Provençal, Quechua, Quiche, Quileute, Rapanui, Romanian, Romany, Rotuman, Russian, Saami, Samoan, Sanskrit, Sanuma, Seneca, Serbian, Shona, Shoshone, Sindhi, Sinhalese, Slovakian, Slovenian, Somali, Sorbian, Sotho, Spanish, Sundanese, Swahili, Swedish, Tagalog, Tahitian, Tamil, Telugu, Tetun, Thai, Tibetan, Tigrinya, Tiwa, Tiwi, Tlingit, Toba Batak, Tokelauan, Tol, Tongan, Totonaco, Trukese, Tswana, Tukang Besi, Turkish, Tuvaluan, Tzeltal, Ukrainian, Ulithian, Urubu-Kaapor, Vietnamese, Wai Wai, Warekena, Wari, Warlpiri, Washoe, Welsh, West Greenlandic, Wintu, Wolof, Yagua, Yiddish, Yimas, Yoruba, Zapotec, Zoque, Zulu, Zuni

REFERENCES

- BAUER, LAURIE. 2007. *The linguistics student's handbook*. Edinburgh: Edinburgh University Press, to appear.
- BRADLOW, ANN R., and TESSA BENT. 2003. Listener adaptation to foreign accented English. *Proceedings of the 15th International Congress of Phonetic Sciences*, Barcelona, Spain, ed. by Maria Josep Solé, Daniel Recasens, and Joaquín Romero, 2881–84. Barcelona: Universitat Autònoma de Barcelona.
- CHAO, YUEN-REN. 1957 [1934]. The non-uniqueness of phonemic solutions of phonetic systems. *Readings in linguistics 1*, ed. by Martin Joos, 38–54. Chicago: Chicago University Press.
- CLEVELAND, WILLIAM S. 1979. Robust locally weighted regression and smoothing scatterplots. *Journal of the American Statistical Association* 74.829–36.
- DE BOER, BART. 2000. Self-organization in vowel systems. *Journal of Phonetics* 28.441–65.
- DE BOER, BART. 2001. *The origins of vowel systems*. Oxford: Oxford University Press.
- GRIMES, BARBARA F. 1988. *Ethnologue*. Dallas: SIL International.
- HARRELL, F. E. 2001. *Regression modelling strategies*. Berlin: Springer.
- HAUDRICOURT, ANDRÉ. 1961. Richesse en phonèmes et richesse en locuteurs. *L'Homme* 1.5–10.
- LEE, SUNGBOK; ALEXANDROS POTAMIANOS; and SHRIKANTH NARAYANAN. 1999. Acoustics of children's speech: Developmental changes in temporal and spectral parameters. *Journal of the Acoustical Society of America* 105.1455–68.
- LIVELY, SCOTT E.; JOHN S. LOGAN; and DAVID B. PISONI. 1993. Training Japanese listeners to identify English /r/ and /l/. 2: The role of phonetic environment and talker variability in learning new perceptual categories. *Journal of the Acoustical Society of America* 94.1242–55.
- LIVELY, SCOTT E.; DAVID B. PISONI; REIKO A. YAMADA; YOH'ICHI TOHKURA; and TSUNEO YAMADA. 1994. Training Japanese listeners to identify English /r/ and /l/. 3: Long-term retention of new phonetic categories. *Journal of the Acoustical Society of America* 96.2076–87.
- LOGAN, JOHN S.; SCOTT E. LIVELY; and DAVID B. PISONI. 1991. Training Japanese listeners to identify English /r/ and /l/: A first report. *Journal of the Acoustical Society of America* 89.874–86.
- MADDIESON, IAN. 2005. Vowel inventories. *The world atlas of language structures*, ed. by Martin Haspelmath, Matthew S. Dryer, David Gil, and Bernard Comrie, 14–15. Oxford: Oxford University Press.
- MAYE, JESSICA, and LOUANN GERKEN. 2000. Learning phoneme categories without minimal pairs. *Proceedings of the 24th annual Boston University Conference on Language Development*, ed. by S. Catherine Howard, Sarah A. Fish, and Thea Keith-Lucas, 522–33. Somerville, MA: Cascadilla.

- MAYE, JESSICA, and DANIEL WEISS. 2003. Statistical cues facilitate infants' discrimination of difficult phonetic contrasts. *Proceedings of the 27th annual Boston University Conference on Language Development*, ed. by Barbara Beachley, Amanda Brown, and Frances Conlin, 508–18. Somerville, MA: Cascadilla.
- MAYE, JESSICA; JANET WERKER; and LOUANN GERKEN. 2002. Infant sensitivity to distributional information can affect phonetic discrimination. *Cognition* 82.3.B101–B111.
- PERICLIEV, VLADIMIR. 2004. There is no correlation between the size of a community speaking a language and the size of the phonological inventory of that language. *Linguistic Typology* 8.376–83.
- PIERREHUMBERT, JANET. 2000. What people know about sounds of language. *Studies in the Linguistic Sciences* 29.2.111–20.
- PIERREHUMBERT, JANET. 2001. Exemplar dynamics: Word frequency, lenition, and contrast. *Frequency and the emergence of linguistic structure*, ed. by Joan Bybee and Paul Hopper, 137–57. Amsterdam: John Benjamins.
- PIERREHUMBERT, JANET; MARY E. BECKMAN; and D. ROBERT LADD. 2001. Conceptual foundations of phonology as a laboratory science. *Phonological knowledge*, ed. by Noel Burton-Roberts, Philip Carr, and Gerard Docherty, 273–304. Oxford: Oxford University Press.
- PORT, ROBERT F., and ADAM P. LEARY. 2005. Against formal phonology. *Language* 81.927–64.
- THRÁINSSON, HÖSKULDUR; HJALMAR P. PETERSEN; JÓGVAN Í LON JACOBSEN; and ZAKARIS SVABO HANSEN. 2004. *Faroese*. Tórshavn: Føroya Fróðskaparfelag.
- TRUDGILL, PETER. 2002. Linguistic and social typology. *The handbook of language variation and change*, ed. by J. K. Chambers, Peter Trudgill, and Natalie Schilling-Estes, 707–28. Oxford: Blackwell.
- TRUDGILL, PETER. 2004a. Linguistic and social typology: The Austronesian migrations and phoneme inventories. *Linguistic Typology* 8.305–20.
- TRUDGILL, PETER. 2004b. On the complexity of simplification. *Linguistic Typology* 8.384–88.

Hay
 Department of Linguistics
 University of Canterbury
 Private Bag 4800
 Christchurch
 New Zealand
 [jen.hay@canterbury.ac.nz]

[Received 3 January 2006;
 revision invited 14 June 2006;
 revision received 14 July 2006;
 accepted 22 September 2006]

Bauer
 School of Linguistics and Applied Language Studies
 Victoria University of Wellington
 P.O. Box 600
 Wellington
 New Zealand
 [laurie.bauer@vuw.ac.nz]