Yers, Logical Phonology, and Absence

Philippe Larkin - LING621

May 11th, 2026

1 Introduction/Presentation of subject

The question of the mapping between phonology and phonetics requires insights that are not always physically evident. Yers, unlike segments that exist in the usual mappings between phonology and phonetics, are not consistently predictable by the standard rules. They are clearly a special case, since for some given environments, you may in one instance see the presence of a yer and in the very same environment, you may not see the yer. For this reason, yers are a test case for analyzing the phonology-phonetics mapping, while continuing to show how different analyses of the same data can be. They present a challenge for how phonological forms should be stored and also how they decide to surface. When something exists phonologically but does not always surface phonetically , the question is if we have to review our perception of the phonology/phonetics mapping. Although this mapping is commonly thought to be reversible, since perception is to be taken into consideration, yers represent a special case.

Yer is a term that refers to vowel-zero alternation in Slavic languages. In Slavic, we observe the presence and absence of vowels in the phonetic output that do not seem to be fully predictable. Their distribution suggests that although phonetically similar (or identical) to their strong vowel counterpart, they are underlyingly different in some way.

Table 1: Yer data from Polish

GLOSS	NOM.SG	GEN.SG	INSTR.PL
sweater	sfɛt<u>ɛ̞</u> r-∅	sfɛtr-a	sfɛtr-ami
bicycle	rovɛr-∅	rovɛr-a	rovɛr-ami
peter	pʲotr-∅	pʲotr-u	pʲotr-ami

The underlined vowel in the table above is the yer. The term “ghost vowel” is sometimes used as well, referring to elements that surface or are absent in some predictable contexts (moreover, ghost aptly describes that the alternation with zero is sometimes visible, sometimes invisible) (Lindsey, 2019). Even when absent, the phonological presence of the yer is well justified, since its effects are observed on surrounding segments through things like palatalization. Palatalization is more of a diachronic change, since it comes from the slow and gradual removal of the weak yer, which affected the preceding consonant by palatalizing it. However, this can still count towards evidence of the presence of a yer.

For those unfamiliar with yers, the question becomes why this is not treated as the deletion of a segment in a specific ordering of rules, which would theoretically allow us to see the effects on

the surface of the vowel being there, but after deletion the vowel would be absent phonetically. In other, words, why does the yer deserve special status? In similar or even identical environments, a yer will often be absent (delete), while a "normal" vowel will not. This is well illustrated by Polish:

The three-way paradigm in Table 1 illustrates the behaviour of the yer, which is underlined for clarity. The yers appear in the final syllable of a root. They also appear if they are in a syllable that is followed by a syllable containing a yer. The distribution of the normal, or strong, vowels has overlap, as can be seen in Table 1. Although there are rules that govern keeping the yer to not produce illegal final clusters, there are exceptions, such as 'sweater' since it is a loanword. Why is this special? It is because yers are not so easily explained as just being normal vowels that delete or appear in certain contexts. They are phonological segments that sometimes surface phonetically. This paradigm should display the necessity to theorize yers as being something different from regular vowel behaviour.

The analytic problem of yers is their distribution. Their distribution points to a unique case of underlying segments that seem to neutralize at the surface. Many non-phonetically grounded frameworks have their own way of dealing with the yer case. This paper sets out to improve on a major issue of both the existing rule-based and autosegmental analyses of this classic problem.

2 Past analyses: autosegmental and SPE-style

This squib attempts to address the problem that exists in many analyses of yers and that is referring to the absence of something. Although it is still a debated topic, there are interesting reasons to try to restrict our system in this particular way. Since yers can be analyzed as something that is hard to hear phonetically, this would make them the perfect test case for this issue.

As we will see in subsequent sections, the autosegmental and SPE-style approaches to this problem actually share a lot in common. They both rely on an abstract approach to yers. The autosegmental analysis may, for example, target a segment which has no associated v-slot. The SPE approach may use a certain feature, like [tense] to posit the existence of the phoneme for the yer and have it absolutely neutralize at the surface. Both of these are abstract, in that they do not rely on phonetic differences in the signal to explain how the yer distribution is learned/occurs. My analysis falls along the same line of thinking, but takes a position on absence.

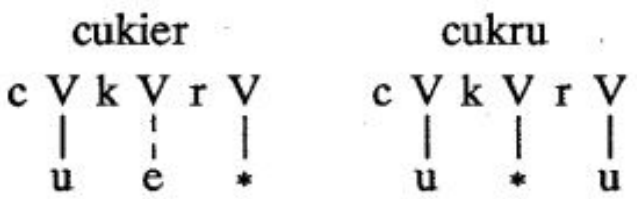
The position that a phonological theory should not make reference to the absence of something is explored in this squib, and possible, alternative ways to approach the yer problem are also discussed.

2.1 Autosegmental analyses

The autosegmental analyses for yers usually involve some form of differentiation for yers compared to normal vowels. One way to do this is to say the yer vowel is something completely empty (featurally) that is not tied to the segmental tier, which is then added by some set of rules. For example, Spencer (1986) has yers as unassociated V-slots on the skeletal CV tier. In this analysis, the [e] is an empty epenthetic vowel. It is the only vowel that has no featural specification, it is neither high, nor back, nor round. "Empty V-slots are filled with melodic material by a set of partly universal and partly language-specific redundancy rules, and surface as [e], with the exception of

the last occurrences of empty V, which are marked extrametrical, and are left uninterpreted at the end of the derivation."(Piotrowski, 1988) It looks something like this:

(8)



The problem of figuring out from the phonetic signal which is the yer and which is the normal vowel is usually the critique given to those who posit only one surface form, with absolutely no phonetic difference between a yer and a normal, 'phonetic' [e] (and it is not particular to Spencer's analysis). A complaint of this analysis found in (Piotrowski, 1988) is that the instances of phonetic [e] are indistinguishable from each other, whether deleting or not, since they are all empty slots on Spencer's analysis. As mentioned, the abstractness of this or Lightner's approach is not considered a problem by those who suggest it. However, one must at least address the learning issue of having two different underlying representations that surface as one sound if you cannot hear the difference. If we believe there is no phonetic difference, abstract representations are not so hard to accept. If a learner does not find a normal pattern that is predicting when the yer surfaces, this may lead the learner to posit a ghost segment, feature, or other abstract concept to help them learn the distribution.

Exception marking, like in the case of (Chociej & Kochetov, 2009), is no less abstract in its proposal than the original tense feature. Storing exceptions in a different way in the lexicon remains the same type of proposal, as long as one thinks the acoustic reality is that yers and non-yers sound the same on the surface.

Lightner's approach is historical in argumentation, since it's said that the yers come from the historically short Proto-Slavic and even further back PIE vowels, while the strong vowels come from the long variants. Those historically short vowels were assigned as [-tense] by Lightner.

2.2 Why not the traditional SPE-style approach?

The SPE approach makes a couple predictions. Firstly, it predicts that yer distribution can be accounted for by rules applied to underlying forms, as long as you use the [tense] feature right. It also predicts a certain level of opacity in rule application. The surface form's path from underlying form to surface form is not exactly clear under absolute neutralization. It makes the prediction that the learner must decide that there is another feature that is causing the special distribution of these vowels. [tense] is a possibility, but really any other theorized, unused feature, like [yer] is equally good. It assumes that the listener cannot hear from the signal the difference between a yer and a strong vowel, and must thus assume a special feature that will have the listener store both the yer and the strong vowel. The distribution and the two underlying forms then allow the speaker to produce the correct output.

I will briefly note a couple of reasons that the traditional system was left behind, trying to summarize the critiques. The rule-based approach to the yer problem is to treat the vowels you want to represent the yers as underlyingly having some feature. In Theodore Lightner's dissertation

(Lightner, 1965) it is [-tense], and then this marking with a feature allows them to be distinguished and then targeted by tthe lower rule or any kind of other process you want to do. The SPE approach seems to be a case of absolute neutralization. The reason absolute neutralization can sometimes be controversial to propose is that on the surface you have no direct evidence that there are two different underlying segments that surface as the same segment. Someone might contest that it is impossible to acquire the two underlying segments with this pattern. You can only theoretically posit its existence based on its necessity in a certain context. A certain phonological phenomenon that seems like it has no reason to behave a certain way may be solved by theorizing the existence of an underlying segment that never surfaces, causing effects on its surrounding environments (a la SPE stress rules for a word like "ellipse") You have underlying vowels that might or might not have this kind of marking on them but you will not hear it phonetically, they will surface the same. In this way, positing these features could be considered as too ad hoc. The autosegmental analysis is also technically guilty of this problem, however it seems the thought process is that it is mechanically more satisfying, since allowing empty slots and using that to your advantage seems more like a feature of autosegmental analyses and not an exception or stipulation on them.

It is important to be fair to both sides, and it is not uncommon to still see phonological phenomenon being analyzed as cases of absolute neutralization, like the potential case of the harmonically irregular vowels of Hungarian, found in (Reiss, 2021).

(29)	Feature-changing absolute neutralization of ϕ and ψ with Hungarian parallel
	
	• Rules: A subtraction rule for ψ-to-Δ and a unification rule for Δ-to-ϕ

Keeping in mind the ideas of abstractness, absolute neutralization, and now also referring to the absence of something, I will show that on the groundwork set by Logical Phonology, we can tie all of these up into one single conversation.

3 What is Logical Phonology and why do we want it?

Substance-Free Logical Phonology is a theory of phonological computation. by substance-free, what is meant is a theory of phonology that doesn't need to make reference to well-formedness, markedness, child phonology, constraints, phonotactics, articulatory phonetics, etc. It states that phonological theories should ideally not account for every true generalization about attestable phonological systems. LP uses set-theory as the foundation on top of which the phonological rules are built.

What can Logical Phonology buy us that other frameworks may be lacking? The answer lies in positively specifying targets. This seems like a small change, but being able to refer to the absence of something is actually undesirable for a few reasons.

A note on negation, nothingness

Why should we want a theory which restricts the ability to refer to nothing? This can be answered similarly with a question: why should we want to restrict anything at all? Pure negation and referring to the space of nothing are powerful tools. If they turn out to be absolutely necessary, than so be it. However, a theory which tries to push the possibility of being able to handle all possible phonological rules without referring to nothing has actually substantially restricted itself in power. Restricting yourself this way is minimalist in philosophy, and can be argued to be a very good thing. Full negation will raise the computational complexity of the logic you have decided dictates the phonological system. It can be paralleled by the early attempts to classify natural language in the Chomsky hierarchy. What point is there to start at the Turing machine level? It serves little purpose from a scientific perspective to just assume you have infinite computational power. What can we get then that can help us do some of the heavy lifting full negation used to do? 'Not equal to' is actually a lower computational complex operation. In any case, the point remains that restricting the system is one way to try to be a minimalist even in phonology.

Point 1: Breaking the natural class definition

The first is that being able to refer to the absence of something breaks the definition of natural classes. The natural class containing segments /t/ and /d/ should also contain /D/, which is the segment underspecified for voice, for example. This approach which uses archiphonemic-type underspecification is found in analyses of Turkish vowel harmony. See the following long but illustrative quote:

"if a rule be permitted in English phonology which distinguishes between the features of /p/ and /b/ on the one hand and on the other, the set of these same features with the exception that voice is unspecified (a set which we shall designate by means of the "archiphonemic" symbol / B/), then we could easily eliminate from English phonology, without knowing anything about English pronunciation, the otherwise relevant feature of Voice from all occurrences of either /p/ or /b/, or in fact from all occurrences of any voiced stop." This quote specifically calls for the inclusion of the underspecified /B/ in the natural class containing /p/ and /b/, lest we break the natural class system.

Point 2: Infinite regress? Or other problems with allowing absence

Imagine the raven paradox.

1. All ravens are black.
2. If something is not black, then it is not a raven.
3. This green apple is not black, and it is not a raven.

Th statement 3 is evidence that if something is not black then it is not a raven. The statement is also logically equivalent to 1 (all ravens are black), so it would follow that the sight of a green apple is evidence supporting the notion that all ravens are black. This would mean that by my looking at an apple, I've discovered information about ravens. This doesn't seem right. The reason I bring up this paradox is to illustrate how careful we have to be with logically equivalent statements. Stating things in terms of positively specified information is not just word play. Restricting the system to not allow reference to absence is actually restricting quite significantly the amount of logically equivalent ways we can express a same target or rule, and this is a good thing. Not only does this

restriction permit, as we mentioned above, to restrict the computational complexity of the overall system, we are also restricting the amount of ways we can state things.

Finally, the following should illustrate the issue brought up in this squib. Say we would want a rule for intervocalic voicing:

[-voice] → [+voice] / V__V

We could write this as: $\forall x[-voice(x) \wedge \exists y(Vowel(y) \wedge Precedes(y,x)) \wedge \exists z(Vowel(z) \wedge Precedes(x,z))] \rightarrow [+voice]$

This way I've written it using formal logic. Once I add the $\neg$ symbol, as mentioned previously, this is full negation and adds much complexity to my logic.

What is Logical Phonology?

It decomposes the traditional arrow operator ($\rightarrow$) into a couple different set-theoretic notions:

- subtraction, noted as the set minus ($\setminus$).
- unification ($\sqcup$)
- segment insertion and deletion via $\mapsto$ (called 'maps to') and ϵ

Segments are feature specifications linked to an X-slot. Underspecification is allowed and is utilized in LP to generate the right rules. A segment would look like:

$$/\text{i}/ = \left\{ \begin{array}{c} +syll \\ -back \\ -round \\ +high \\ -low \\ +ATR \end{array} \right\}$$

3.1 Subtraction

Since each segment is a set, a natural class is a set of sets of valued features. Subtraction is quite intuitive since it works in the same way that is learned in basic set theory:

$$\left\{ \begin{array}{c} +high \\ -back \\ -round \end{array} \right\} \setminus \left\{ \begin{array}{c} +high \\ +round \end{array} \right\} = \left\{ \begin{array}{c} -back \\ -round \end{array} \right\}$$

In a rule, the environment is specified using natural classes like:

[-back] \ {-high} / ___ [+nasal]

3.2 Unification

Unification is different from *union*. In set union and in particular with valued features, the union of two sets can accidentally cause two contradicting features to join into the same set, which is simply not possible. Unification, however, cannot do this. If it tries to union with a feature and value it already contains, we call this **vacuous unification**. If it tries to union with a feature and value directly opposite, we call this **unification failure**. For example, a vacuous unification looks like:

$$\left\{\begin{matrix} +high \\ -back \\ -round \end{matrix}\right\} \sqcup \{-back\} = \left\{\begin{matrix} +high \\ -back \\ -round \end{matrix}\right\}$$

And unification failure would look like:

$$\left\{\begin{matrix} +high \\ -back \\ -round \end{matrix}\right\} \sqcup \left\{\begin{matrix} +high \\ +round \end{matrix}\right\} = undefined$$

4 Logical Phonology Approach

The LP approach to yers would resemble any of the Lightner, Rubach, or Spencer approaches in what they are theoretically positing. To remind ourselves of those, they are:

- Rubach (1986): lacking an underlying X slot
- Spencer (1986): bare V slot
- Rubach (2016): lack an underlying mora
- Gussmann (2007): not linked to a nuclear X slot

Here is the data that is used as an example:

	nom.sg.	gen.sg.	gen.pl.		
a.	kub**eł**	kubł-a	kubł-ów	'bucket'	(masc.)
b.	klusk-a	klusk-i	klus**e**k	'noodle'	(fem.)
c.	żebr-o	żebr-a	żeb**e**r	'rib'	(neut.)
d.	kuf**e**r	kufr-a	kufr-ów	'chest, trunk'	(masc.)
e.	szof**e**r	szof**e**r-a	szof**e**r-ów	'driver'	(masc.)
f.	cyfr-a	cyfr-y	cyfr	'digit'	(fem.)

In LP's way of approaching this problem, which can be seen in (Gorman, 2025), the fleeting vowels (yers) are considered both the target and the trigger of the rules allowing this fleeting *e* to surface, for example. The high lax / i̯, ĭ / are used as the underlying forms that never actually surface as they are underlyingly. The basic generalization is that yers surface as *e* if there is a yer in the following syllable. They will delete otherwise. The two rules:

(9)

$$\begin{bmatrix} +high \\ -tense \end{bmatrix} \setminus \begin{Bmatrix} +high \\ +back \end{Bmatrix} /__C_0 \begin{bmatrix} +high \\ -tense \end{bmatrix}$$

(10)

$$[\,] \sqcup \begin{Bmatrix} -high \\ -back \end{Bmatrix}$$

Then there is a rule that deletes the remaining yers:

(11)

X
|
$\begin{Bmatrix} -high \\ -back \end{Bmatrix} \mapsto \epsilon$

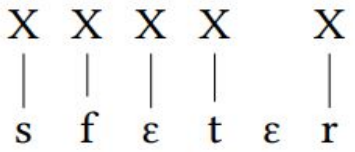
This combination of rules correctly predicts the data. In the end, however, this is quite similar to the original rule-based analysis, using restated tools of LP.

5 How to improve upon this problem

Here, I will add to the conversation surrounding the LP approach to yers in one important way. One aspect of Logical Phonology, mentioned in: "Recall that LP segmental rules cannot refer to the absence of a feature specification. In a similar spirit, I propose LP autosegmental rules cannot refer to the absence of an association line between some node n on tier T and any other node on tier T'. This proposal makes it impossible to target just the fleeting vowel". We have now seen why this objection may be well motivated.

This is a stipulation in LP, an assumption that is made. The obvious problem then is twofold. On the segmental tier, we cannot refer to an e that is not associated to an X-slot, so then we cannot target just the yer. On the X tier, we cannot refer to the absence of the X or even if there were an X, to the absence of an association line.

One way of seeing this is that LP only accepts positive evidence for things. Trying to find workarounds like "a slot that interrupts the timing tier" or something similar is still technically empty, and undesirable to target for a rule. However, if we refer back to the timing tier diagram for the Polish loan word:



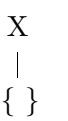
Option 1: Flanking consonants

With flanking consonants, in this case the segments t and r between our yer, the yer is between two consonants which are adjacent on the X tier. If you refer to the yer as lying between two consonants which are adjacent to each other on the X tier, you are only using something that is positively specified. If this is the case, the autosegmental analysis can be retained using positive evidence.

One important problem with this workaround is the issue of word-final yers, which are not surrounded by consonants. The orthodox approach for the x-tier does not allow for boundary symbols (say a #). Allowing the # breaks the division of labour, or modularity of the system. The x-tier is for timing and should possibly not allow this boundary symbol. There are a few simple things that seem to be more difficult under these restrictions we set out for ourselves, such as simply referring to word final segments. Depending whether you take the logical phonology approach or a different one, word-final environments can be described positively. "Target the segment within [] in which all segments come before it". By doing this, you refer to both the segment you're targeting and the stuff before it, without ever having to worry about referring to that empty space ahead of you.

Option 2: The second LP method

Since the word-final yers may present a problem, it is clear that the LP way may present to us a second, even simpler analysis. The empty set is not the same thing as having nothing. By virtue of calling it the empty set, you are admitting it is something, a structure, a set. It just contains no elements. Since I have previously described why LP phonologists may be against referring to the absence of something, unassociated x-slots are not allowed. However, you can have an x-slot associated with an empty set. Then, you still have something to refer to, but it is empty. It is a sort of combination of the earlier theories. Just like the following:



It's still important to mention, however, that this approach has its own set of issues. Although it is nice that these extra tiers exist in LP, putting an empty set symbol down there may restrict some of the fundamental properties of leaving the segmental tier slot empty. We have many rules in autosegmental phonology which are able to cross only if there is no blocker (aka empty slot). If we add an empty set there, we may very well have to find a workaround to this kind of issue.

6 Conclusion and why this is interesting

In the LP approach by Gorman, and also in the two separate approaches or modifications I propose here, there is a modification of the rules and approaches to fix a problem. If I can just as well provide a rule which gives me the correct distribution of yers but does not refer to the absence of something, then I have proposed an alternative analysis which is more minimal.

For the very specific reason that yers are usually treated as the absence of something, or being 'ghosts', this is precisely why the conversation surrounding absence is important. A theory which

restricts the system along this dimension, and a theory which uses only the basic structures laid out is a theory which has many positives. For this reason, LP, which does not simply state an opposition to referring to absence but argues it on logical grounds (and also its building blocks in set theory and definition of natural class), is one worth considering against a backdrop of dozens more. Whether the adjusted approach to the x-tier in terms of flanking consonants, the approach based on Lightner's in (Gorman, 2025), or the adjustment to the x-tier in the last point above is found to be the most elegant one, the point I make is hopefully clear by now: **as long as it is not referring to the absence of something**.

The question of yers is a test case for the mapping from phonology to phonetics. Although this mapping is commonly thought to be reversible, since perception is to be taken into consideration, yers represent a case where we have something phonologically present, very clearly, but for which we cannot find evidence phonetically ((Lindsey, 2019)). There is some evidence that yers are actually phonetically different in some sense that could be perceived by the learner. In this case, there wouldn't be anything particularly special about yers. This squib assumes there is no phonetic difference between the yers and the non-yers, as an exercise in trying to sketch out an abstract dilemma.

References

Chociej, J., & Kochetov, A. (2009). Patterns in polish vowel˜ zero alternations: An optimality theory analysis. In *Manuscript. university of toronto ma forum paper.*

Gorman, K. (2025). *A logical phonology of some "minor rules" of polish.*

Lightner, T. M. (1965). *Segmental phonology of modern standard russian.* (Unpublished doctoral dissertation). Massachusetts Institute of Technology.

Lindsey, K. L. (2019). *Ghost elements in ende phonology.* Stanford University.

Piotrowski, M. (1988). Polish yers in non-linear phonology. *Phonologica*, 215–227.

Reiss, C. (2021). Towards a complete logical phonology model of intrasegmental changes. *Glossa: a journal of general linguistics*, *6*(1).