Most modern generative theories of phonology assume an overall analysis resembling the following three-part formal structure:

/pɝpl/ underlying representation (UR)

phonemic	consists only of contrastive sounds ('phonemes'); e.g. we know that /f/ and /v/ are phonemes of English because of the minimal pair *Eph* /if/ and *eve* /iv/
unpredictable	if information about a word must be memorized, it belongs in the UR; we must memorize that *Eph* ends with the phoneme /f/
abstract	not directly measurable; phonemes and URs are inferred from minimal pairs and patterns of pronunciation; we can't open up the brain to see phonemes and URs, so the true cognitive status of phonemes and URs is unclear

rules (e.g. syllabification, aspiration, /l/-velarization, . . .)

derivational	this is the 'generative' part of generative phonology; rules derive SRs from URs and describe allophonic pronunciation of phonemes as simply and naturally as possible; direction matters: aspiration is /p/ → [pʰ], and de-aspiration is /pʰ/ → [p]
abstract	like phonemes and URs, rules are theoretical devices that we use to model language, but we have no direct evidence that they exist in the minds of speakers

[pʰɝ.pɬ] surface representation (SR)

allophonic	pronunciation of phonemes in particular environments, as dictated by the rules; e.g., English aspirates voiceless stops word-initially, so /kaʊ/ is pronounced [kʰaʊ]
predictable	all differences between UR and SR are derived by rules, and thus can be predicted based on the UR and the set of rules available
concrete	represents actual, directly measurable pronunciation, depending on how narrow the phonetic transcription is

A phonological analysis should satisfy **simplicity** and **naturalness** whenever possible. A simple analysis collapses multiple related phenomena having recurrent patterns into a single rule. For example, it is simpler to write one rule to aspirate all word-initial voiceless stops than to write a separate aspiration rule for each one, "coincidentally" with the same change in the same environment.

A natural rule is one that obeys phonetic principles (acoustics, articulation, etc.). **Assimilation** is one of the most common kinds of natural rules, in which two adjacent sounds become more alike. **Dissimilation** is also common, when two adjacent sounds that are already very similar become more distinct. Many other rules have particular phonetic explanations (word-final devoicing, cluster simplification, etc.), and as we explore more data, we'll see many of these natural rules.

The classic standard formalism for a phonological rule is a rewrite rule in the following format:

$$\underset{\text{target}}{\text{A}} \quad \rightarrow \quad \underset{\text{change}}{\text{B}} \quad / \quad \underset{\text{environment}}{\text{X}\_\text{Y}}$$

where A, B, X, and Y can be any string of phonological material (features, segments, boundaries, etc.). The target is the set of distinctive features needed to specify all and only the segments affected by the rule. The change is the set of features that are changed or added to the target. The environment is the description needed to specify when the rule applies, with the underbar representing the position of the target with respect to X and Y, which are the descriptions of the triggering material on the left and right respectively (either or both of which could be nothing!). A sample formal rule is given below, which causes intervocalic voiced obstruents /b d ɡ/ to become fricatives [β ð ɣ]:

$$\begin{bmatrix} -\text{son} \\ +\text{voi} \end{bmatrix} \rightarrow [+\text{cont}] \,/\, \text{V}\_\text{V}$$

If the language has nasals /m n ŋ/, the target must be specified as either [−son] or [−nas] to prevent the rule from applying to them. Note that we do not need to specify both [−son] and [−nas] in the target, because nasals will be ruled out by either one. Having both would still create a working rule, but it would not be *minimal*.

Further, assuming the language has voiceless obstruents like /p t k/, the rule must include the feature [+voi] in the target description to prevent the rule from wrongly applying to voiceless obstruents.

However, we do not need the feature [−cont] in the target, even if the language has voiced fricatives like /z/. The reason is that if the rule 'applies' to voiced fricatives, they do not change! This is called **vacuous rule application**; the rule vacuously applies to voiced fricatives because the target description includes them, but the change does not make any actual change. Again, adding [−cont] to the target description would create a working rule, but then the target description would not be minimal. You may find vacuous rule application to be a bit unsettling, but it allows us to simplify many rules, and in the end, doesn't change the outcome. Occam's Razor at work!

The fricatives that result from the rule are all voiced obstruents, so you might be tempted to include [+voi] and/or [−son] in the change. But this is not necessary, since the targeted segments are already voiced obstruents, and by our assumption of **faithfulness**, no feature values are changed by a rule unless they are explicitly mentioned.

The environment of this rule tells us that the targeted segments must have a vowel on the left and a vowel on the right in order to be affected by the rule. This means that voiced obstruents next to consonants or at the beginning or end of a word will not become fricatives.

There are a variety of **special symbols** used in rules when features aren't appropriate:

∅	null symbol	A → ∅ / X__	delete A after X
		∅ → A / X__Y	insert A between X and Y
#	word boundary	A → B / #__	A becomes B at the beginning of a word
		A → B / __#	A becomes B at the end of a word
X_0	zero or more of X (usually C or V)	A → B / __C_0#	A becomes B at the end of a word, or before any number of word-final consonants
$[_\sigma$ $]_\sigma$	syllable boundaries	A → B / $[_\sigma$ __	A becomes B at the beginning of a syllable
		A → B / __ $]_\sigma$	A becomes B at the end of a syllable

The IPA symbol ˌ may also be used for a syllable boundary, but $[_\sigma$ and $]_\sigma$ are easier to see, so they are preferred.