Nishant Suria
Professor Rachel Walker
LING 102
8 December 2021

ANALYSIS OF VOWEL DEVOICING IN STANDARD JAPANESE

1. INTRODUCTION

Japanese (endonym 日本語, [nihoŋgo]) is a commonly spoken language, spoken by about 128 million people, primarily concentrated in Japan. It is considered to be in the Japonic (which includes Japanese and a few Ryukyuan languages) language family, and while a connection to other language families has been posed several times, no such theory has received scholarly acceptance.

The phonology of Japanese is reasonably restricted, commonly described as only allowing morae /CV/, /V/, along with /N/, a placeless nasal that assimilates with the following segment, and /Q/, a placeless archiphoneme that assimilates with a following obstruent, forming a geminate (Labrune 2012). On the other hand, Japanese features many instances of allophony, that being variation in the articulation of a phoneme depending on context, and one such example is voicing in vowels, especially high vowels [i] and [ɯ]. While vowels are generally [+voi], in Japanese, vowels can devoice under certain circumstances.

2. OVERVIEW OF JAPANESE PITCH ACCENT

Japanese vowel devoicing has some interactions with Japanese pitch accent (低高アクセント), and so the analysis will be prefaced with a foundational explanation thereof, drawn from Labrune (2012). Japanese pitch accent distinguishes two 'heights', which are called low (低, L) and high (高, H), which I notate in this analysis with ↓ and ↑, respectively, written under the narrow transcription. Every mora in a Japanese word is attested to have one of these two pitch levels, and while not as prevalently as in tonal languages, can be used to contrast words in narrow cases.

(i) Minimal pair 箸 [ha↑.ɕi↓] and 橋 [ha↑.ɕi↑]

箸	[ha.ɕi] ↑ ↓	"chopsticks"
橋	[ha.ɕi] ↑ ↑	"bridge"

Japanese words have pitch assigned to every mora, and additionally, one that gets assigned to the following particle, if present. For example, the following forms have the particle の [no] following them, which marks the genitive case.

(ii) Minimal triplet, and differential pitch of attaching particle

箸の	[ha.ɕi no] ↑ ↓ ↓	"of the chopsticks"
橋の	[ha.ɕi no] ↑ ↑ ↓	"of the bridge"
端の	[ha.ɕi no] ↓ ↑ ↑	"of the edge"

This kind of three-way distinction is very rare in Japanese, but is clearly possible under Japanese phonetics. Without a particle present, however, the second two forms, 橋 "bridge" and 端 "edge" are neutralized, with no difference in pronunciation.

The contours of Japanese pitch accent follow one of two patterns, which reasonably limits the amount of different articulations possible for Japanese words. The first pattern occurs in words that contain a 'downstep', which is a location in a word between two morae where the pitch drops (called 頭高 [atamadaka] "head-high", 中高 [nakadaka] "middle-high", or 尾高 [odaka] "tail-high" in Japanese depending on where the downstep occurs). In this pattern, the pitch accent is high (↑) on all morae before the downstep, and low (↓) on all morae after the downstep, as well as the aforementioned pitch that is assigned to the attaching particle. While usually notated with a (ꜜ), for this analysis I will notate it with a (ˋ), to avoid ambiguity. The forms 箸 "chopsticks" and 橋 "bridge" follow this pattern.

(iii) Minimal pair 箸 [ha↑.ɕi↓] and 橋 [ha↑.ɕi↑], with underlying downstep

箸	はし haˋ.shi	[ha.ɕi] ↑ ↓ ↓	"chopsticks"
橋	はし ha.shiˋ	[ha.ɕi] ↑ ↑ ↓	"bridge"

As shown, the downstep of 箸 "chopsticks" occurs after the [ha] mora, and so the [ha] mora is articulated high (↑), and the rest low (↓). In contrast, the downstep of 橋 "bridge", occurs after the [ɕi] mora, and so both the [ha] and [ɕi] morae are articulated high (↑), and only the attaching particle's articulation is low (↓). The other pattern, *Heiban* (平板 [he:ban] "flat") is considered the default pattern, and occurs when a word has no downstep. In such words, the first mora is articulated low (↓), and then all subsequent morae, through the attaching particle, are articulated high (↑). The form 端 "edge" follows this pattern:

(iv) 端 (ha.shi), with Heiban contour

端	はし ha.shi	[ha.ɕi] ↓ ↑ ↑	"edge"

As shown, the first mora [ha] is articulated low (↓), and the rest high (↑). This is considered "flat" to many Japanese speakers because the most audible part of pitch accent articulation in Japanese is the downstep from a high to a low mora, and Heiban contours have no downstep.

Additionally, the downstep is attested to be what is underlying with regards to Japanese pitch accent, as the presence or absence, and location of a downstep can fully predict the surface pitch accent contour of a Japanese word.

This analysis concerns what is known as Standard Japanese (標準語 [ɕjoːd͡zɯŋgo]) specifically, as pitch accent patterns, as well as vowel devoicing tend to vary heavily depending on dialect. Standard Japanese is the de facto national language of Japan, and language of the media and education, and is based on the dialect of Tokyo. Pitch accent data for this analysis is taken from Daijirin (大辞林), 3rd edition (Matsumura 2006).

3. ANALYSIS OF VOWEL DEVOICING

Vowel devoicing (母音の無声化) in Japanese surfaces as allophonic variation between voiced and voiceless vowels. In essence, in Modern Japanese, there are no words where the presence or lack of a voiceless vowel causes a difference in the word, and its appearance can be predicted, for which an analysis will be outlined and argued in this paper.

The underlying cause for vowel devoicing is transparently assimilation, what feature specifically is being assimilated is unclear; the [voice] feature is understandably the most common notion, but it has been argued that [spread glottis] is being assimilated (Tsuchida 2001). In all environments of vowel devoicing, the devoiced vowel is in contact with at least one consonant with the feature [-voice]. In the form below, for example, the [ɯ̥] is preceded by [k], and followed by [s], both voiceless segments.

(v) /kusa/[1], with demonstration of assimilation

/kusa/[1] → [**kɯ̥.s**a]
| | |
[-voi]

Analysis of such allophony with Optimality Theory uses a context-sensitive markedness constraint, a context-free markedness constraint, and a faithfulness constraint (Kager 1999). Let these constraints be abstracted as the following:

(vi) Constraints, for abstract analysis of allophony based on Kager (1999)

*V̥/X (Context-Sensitive Markedness)	Assign a violation for each [+syll, +voi] segment in the output, in context X.
*V̥ (Context-Free Markedness)	Assign a violation for each [+syll, -voi] segment in the output.
Ident-IO(voi)	Assign a violation for each segment that differs in the feature [voice]

(Faithfulness)	between the input and output.

Via the schema devised by Kager (1999), these constraints demonstrate allophony, as in the following tableau, using the strict ranking Context-Sensitive Markedness >> Context-Free Markedness >> Faithfulness, which in this case would be *V̬/X >> *V̥ >> Ident-IO(voi).

(vii) Tableau /kusa/ → [kʊ̥.sa]

/kusa/	*V̬/X	*V̥	Ident-IO(voi)
→ kʊ̥.sa		*	
1 kɯ.sa	*!		*

In this tableau, the vowel [ʊ̥] is marked by *V̥, but because the alternative, form 1, violates the markedness constraint *V̬/X, the voiceless vowel is present in the surface form. Thus, in a form where the alternative does not violate the *V̬/X constraint, the *V̥ constraint is able to stop devoicing from occurring, as in the following tableau:

(viii) Tableau /kɯma/ → [kɯ.ma]

/kuma/	*V̬/X	*V̥	Ident-IO(voi)
→ kɯ.ma			
1 kʊ̥.ma		*!	*

As shown, a violation of the context-sensitive markedness constraint *V̬/X is required to overcome the context-free markedness constraint *V̥, and where that isn't the case, no devoicing occurs. The faithfulness constraint Ident-IO(voi) is specified because Optimality Theory appeals to Richness of the Base, which posits that any underlying form is possible, and an appropriate OT analysis of a language must be able to fix any potential input, not only those which are posited to be present in speakers' lexicons. Therefore, malformed inputs, such as /kʊ̥ma/ should also be considered, as in the following tableau:

(ix) Tableau /kʊ̥ma/ → [kɯ.ma]

/kʊ̥ma/	*V̬/X	*V̥	Ident-IO(voi)
→ kɯ.ma			*
1 kʊ̥.ma		*!	

As shown in this tableau, it is specifically because the faithfulness constraint Ident-IO(voi) is ranked lower than *V̥, forms such as /kʊ̥ma/ can be repaired, and true allophonic distribution can be demonstrated.

The rest of this analysis will concern the nature of *V̥/X, in essence, which environments will pressure a vowel to devoice, each having a new context-sensitive markedness constraint, and where appropriate, additional context-free markedness constraints which are dominated by higher ranked context-sensitive markedness constraints. As such, the constraint relationship theorized by Kager that was explained here will still apply even as more constraints are added. Therefore, the Ident-IO(voi) will be omitted from further tableaux, and malformed inputs such as /kʉma/ will not be considered, with the understanding that they are assumed.

3.1. WORD-MEDIAL HIGH VOWEL DEVOICING

The most observable phenomenon in vowel devoicing in Japanese is that high vowels /i, ɯ/ become devoiced between two voiceless sounds. This analysis posits that the [-voice] is being assimilated to these high vowels, which are specifically susceptible to phonological change, as they are the vowels with by far the shortest articulation (citation needed).

This phenomena is explained by a context-sensitive markedness constraint which pressures a vowel to devoice between two voiceless sounds, and two context-free markedness constraints, one that marks all voiceless vowels, and one that specifically marks voiceless [-high] vowels.

(x) Constraints, for word-medial high vowel devoicing

*[-voi]V̥[-voi]	Assign a violation for each [+syll, +voi] segment between a [-voi] segment and a [-voi] segment in the output.
*V̥	Assign a violation for each [+syll, -voi] segment in the output.
*Ḁ	Assign a violation for each [+syll, -high, -voi] segment in the output.

These constraints show strict ranking *Ḁ >> *[-voi]V̥[-voi] >> *V̥, as the context-sensitive markedness constraint *[-voi]V̥[-voi] is strong enough to overcome the markedness of voiceless vowels (caused by *V̥), but not of [-high] vowels (caused by *Ḁ).

(xi) Tableau /kusa/ → [kʉ̥.sa]

/kusa/	*Ḁ	*[-voi]V̥[-voi]	*V̥
→ kʉ̥.sa			*
1 kɯ.sa		*!	

As mentioned, the environment for this pattern of devoicing is specifically when a [+high] vowel is between two voiceless segments. This is demonstrated in Tableau (xi), where the constraint *[-voi]V̥[-voi] overcomes the markedness of voiceless vowels, which results in the /u/ vowel, being devoiced to [ʉ̥].

(xii) Tableau /kuma/ → [kɯ.ma]

/kuma/	*A̰	*[-voi]V̥[-voi]	*V̥
→ kɯ.ma			
1 kɯ̥.ma			*!

(xiii) Tableau /kesa/ → [ke.sa]

/kesa/	*A̰	*[-voi]V̥[-voi]	*V̥
→ ke.sa		*	
1 kɐ̥.sa	*!		*

However, if this specific environment is not met, such both consonants not being [-voi] as in Tableau (xii), or if the vowel is [-high], as in Tableau (xiii), the interconsonantal vowel does not devoice.

3.2. WORD-FINAL HIGH VOWEL DEVOICING

To a lesser extent, high vowels at the end of a word devoice when preceded by a voiceless sound. Unlike word-medial devoicing, which happens almost universally when the environment occurs, word-final devoicing only occurs in certain words. Let the following examples be considered:

(xiv) Minimal pair /ha˘chi/ and /hachi/

/ha˘chi/ "eight"	[ha.t͡ɕi̥] ↑ ↓↓
/hachi/ "bee"	[ha.t͡ɕi] ↓ ↑↑

These words show a minimal pair in which the final devoicing may or may not occur solely due to a difference in pitch accent. Saliently, the [t͡ɕi] mora is variably high or low pitch, and this is correlated with whether or not the [i] vowel can be devoiced. This is established by the following additional constraints.

(xv) Constraints, for word-final high vowel devoicing

*[-voi]V̥#	Assign a violation for each [+syll, +voi] segment following a [-voi] segment at the end of word in the output.
*V̥↑	Assign a violation for each [+syll, -voi] segment in a high pitch mora in the output.

These constraints have strict ranking *V̥↑ >> *[-voi]V̥#, as devoicing occurs word-finally unless *V̥↑ blocks it.

(xvi) Tableau /kusa/ → [kɯ̥.sa]

/kusa↘/	*Ḁ	*[-voi]V̥[-voi]	*V̥↑	*V̥
→ kɯ̥ɪ.sa ↑ ↑↓			*	*
1 kɯɪ.sa ↑ ↑↓		*!		

This tableau is a restatement of Tableau (vii), with pitch accent and the *V̥↑ constraint added. From this, it is clear that *[-voi]V̥[-voi] must strictly dominate *V̥↑, as the output would be incorrect otherwise.

(xvii) Tableau /ha↘chi/ → [ha.tɕi̥]

/ha↘chi/	*Ḁ	*[-voi]V̥[-voi][2]	*V̥↑	*[-voi]V̥#	*V̥
→ ha.t͡ɕi̥ ↑ ↓↓		*			*
1 ha.t͡ɕi ↑ ↓↓		*		*!	

(xviii) Tableau /hachi/ → [ha.tɕi]

/hachi/	*Ḁ	*[-voi]V̥[-voi][2]	*V̥↑	*[-voi]V̥#	*V̥
→ ha.t͡ɕi ↓ ↑ ↑		*		*	
1 ha.t͡ɕi̥ ↓ ↑ ↑		*	*!		*

In these tableaux which show the minimal pair /hachi/ and /ha↘chi/, show differing outputs of [ha.t͡ɕi̥], and [ha.t͡ɕi], respectively. For both words, the final /i/ violates *[-voi]V̥# where it surfaces as [i], and for /ha↘chi/, this pressure causes the surface form to be [ha.t͡ɕi̥]. However, in /hachi/, because the final [i] becomes high pitch, this incurs a violation of *V̥↑, which blocks devoicing and causes the surface form to be [ha.t͡ɕi]

3.3. INTERACTIONS OF MEDIAL AND FINAL HIGH VOWEL DEVOICING

In words wherein multiple vowels match environments which prompt devoicing, there are more than two likely candidates considered from GEN. For a variety of words, the current set of constraints is sufficient to output the surface form.

(xix) Tableau /futsuka/ → [ɸɯ̥ɪ.t͡sɯ̥ɪ.ka]

/futsuka/	*Ḁ	*[-voi]V̥[-voi]	*V̥↑	*[-voi]V̥#[3]	*V̥
→ ɸɯ̥ɪ.t͡sɯ̥ɪ.ka			*	*	**

↓ ↑ ↑ ↑					
1 ɸɯ.t͡su̥ɯ.ka ↓ ↑ ↑ ↑		*!	*	*	*
2 ɸu̥ɯ.t͡sɯ.ka ↓ ↑ ↑ ↑		*!		*	*
3 ɸɯ.t͡sɯ.ka ↓ ↑ ↑ ↑		*!*		*	

In this form, both morae /fu/ and /tsu/ are in the devoicing environment for high-vowel word medial devoicing, while the mora /ka/ is unable to devoice in this context. The surface form, quite logically, contains both vowels devoiced, as [ɸu̥ɯ.t͡su̥ɯ.ka].

(xx) Tableau /kishitsu/ → [ki̥.ɕi̥.t͡sɯ]

/kishitsu/	*Ḁ	*[-voi]V̥[-voi]	*V̥↑	*[-voi]V̥#	*V̥
→ [ki̥.ɕi̥.t͡sɯ] ↓ ↑ ↑ ↑				*	**
1 [ki̥.ɕi̥.t͡su̥ɯ] ↓ ↑ ↑ ↑			*!		***
2 [ki.ɕi̥.t͡sɯ] ↓ ↑ ↑ ↑		*!		*	*
3 [ki.ɕi̥.t͡su̥ɯ] ↓ ↑ ↑ ↑			*!		**
4 [ki̥.ɕi.t͡su̥ɯ] ↓ ↑ ↑ ↑		*!	*		**
5 [ki̥.ɕi.t͡sɯ] ↓ ↑ ↑ ↑		*!		*	*
6 [ki.ɕi.t͡sɯ] ↓ ↑ ↑ ↑		*!*		*	

In this form, all three morae /ki/, /shi/, and /tsu/ are in devoicing environments, but as we see in the surface form, only the first two morae are devoiced, as [ki̥.ɕi̥.t͡sɯ]. However, this occurrence is due to the aforementioned *V̥↑ context-free markedness constraint, which marks all [-voi] vowels of high pitch, and is critically ranked over *[-voi]V̥#, which would ordinarily cause a vowel like the one in the [tsu] mora to devoice.

However, for some forms, an additional constraint is necessary, which will block *[-voi]V̥[-voi] in certain circumstances. In forms like /tsuku˴su/ [t͡su̥ɯ.kɯ.su̥ɯ], the [kɯ] mora governs the pitch

accent of the word (as the downstep occurs between [kɯ] and [sɯ̟]) so it is specifically not devoiced, even where the final mora is devoiced.

(xxi) Constraint *V̟μ↘

*V̟μ↘	Let 'pre-downstep mora' be a mora of high pitch after which all morae are low pitch. (notated with ↘) Assign a violation for every [+syll, -voi] segment in a pre-downstep mora in the output.

As this constraint blocks word-medial high vowel devoicing, it must strictly dominate *[-voi]V̟[-voi], and no case of it being violated is observed in the data.

(xxii) /tsuku↘su/ → [t͡sɯ̟.kɯ.sɯ̟]

/tsuku↘su/	*A̟	*V̟μ↘	*[-voi]V̟[-voi]	*V̟↑	*[-voi]V̟#	*V̟
→ [t͡sɯ̟.kɯ.sɯ̟] ↑ ↑ ↓ ↓			*	*		**
1 [t͡sɯ̟.kɯ̟.sɯ] ↑ ↑ ↓ ↓		*!		**	*	**
2 [t͡sɯ.kɯ̟.sɯ̟] ↑ ↑ ↓ ↓		*!	*	*		**
3 [t͡sɯ.kɯ̟.sɯ] ↑ ↑ ↓ ↓		*!	*	*	*	*
4 [t͡sɯ̟.kɯ̟.sɯ̟] ↑ ↑ ↓ ↓		*!		**		***
5 [t͡sɯ.kɯ.sɯ̟] ↑ ↑ ↓ ↓			**!			*
6 [t͡sɯ.kɯ.sɯ] ↑ ↑ ↓ ↓			**!		*	

With the addition of this constraint, the lack of devoicing in the [.kɯ] mora can be accounted for, where otherwise form 4, [t͡sɯ̟.kɯ̟.sɯ̟] would be preferred. Although it may seem arbitrary, the blocking of devoicing by being a pre-downstep mora actually is quite reasonable. As mentioned in [IV], it can be reasonably hypothesized that the phonemic [word] of pitch accent is specifically the location of the downstep (or lack thereof), which is also the most [word] part of the articulation. Because voiceless sounds, lacking vibration of the vocal cords, cannot carry a tone, devoicing the downstep vowel would rid the word of pitch accent entirely, and thus never occur in surface forms[4].

3.4 NON-HIGH VOWEL DEVOICING

While not as prominent as with high vowels, other vowels also devoice slightly, in specific words where the same vowel is repeated. Let the following forms be compared:

(xxiii) Devoicing in /haka/, /hako/

はか˪ ha.ka˪	[hḁ.ka] ↑ ↑ ↓
はこ ha.ko	[ha.ko] ↓ ↑ ↑

Due to having the same vowel twice in a row, the first one can devoice. However, this must also occur with the restrictions that are present on high-vowel medial devoicing.

(xxiv) Devoicing in /hata/, /hara/

はた ha.ta˪	[hḁ.ta] ↑ ↑ ↓
はら ha.ra˪	[ha.ɾa] ↑ ↑ ↓

Specifically, the devoicing cannot occur where a high vowel wouldn't devoice, in essence, it only devoices between two [-voi] segments. Therefore, this analysis posits the following constraints:

(xxv) Constraints, for non-high vowel devoicing

*$\mu_{\alpha\underset{\circ}{A}}\mu_{\alpha\underset{\circ}{A}}$	Assign a violation for each pair of morae in a row containing the same [+voi, +syll] segment.
*[-voi]V̥[-voi]	Assign a violation for each [+syll, +voi] segment between a [-voi] segment and a [-voi] segment in the output.
*$\mu_{\alpha\underset{\circ}{A}}\mu_{\alpha\underset{\circ}{A}}$ & *[-voi]V̥[-voi]	Assign a violation iff *$\mu_{\alpha\underset{\circ}{A}}\mu_{\alpha\underset{\circ}{A}}$ and *$\alpha\underset{\circ}{A}C_0\alpha A$ are violated in the output.

This analysis posits that while neither of the constraints *$\mu_{\alpha\underset{\circ}{A}}\mu_{\alpha\underset{\circ}{A}}$ nor *[-voi]V̥[-voi] rank above *$\underset{\circ}{A}$, their conjunction, *$\mu_{\alpha\underset{\circ}{A}}\mu_{\alpha\underset{\circ}{A}}$ & *[-voi]V̥[-voi] does, allowing [-high] vowels to devoice in those specific circumstances. This is demonstrated in the following tableaux.

(xxvi) Tableau /haka˪/ → [hḁ.ka]

/ha.ka˪/	*$\mu_{\alpha\underset{\circ}{A}}\mu_{\alpha\underset{\circ}{A}}$ & *[-voi]V̥[-voi]	*V̥μ˪	*$\underset{\circ}{A}$	*[-voi]V̥ [-voi]	*V̥↑	*[-voi] V̥#	*V̥
→ [hḁ.ka] ↑ ↑ ↓			*		*	*	*
1 [ha.ka] ↑ ↑ ↓	*!			*		*	

2 [ha.kạ̥] ↑ ↑ ↓		*(!)	*	*(!)			*
3 [hạ̥.kạ̥] ↑ ↑ ↓		*(!)	**(!)		*		**

This tableau demonstrates the occasion for which a [-high] vowel, [a] in this case, devoices. In this instance, the underlying form, form 1, violates the constraint *$\mu_{\alpha\underset{\circ}{A}}\mu_{\alpha\underset{\circ}{A}}$ & *[-voi]V̥[-voi], because it contains both two [+syll, +voi] segments in a row, [a], and the first [a] is in the environment of being preceded and followed by a voiceless consonant. Forms 2 and 3 show why it is only the first [a], which is in the position that violates *[-voi]V̥[-voi] that must devoice.

(xxvii) Tableau /hako/ → [ha.ko]

/ha.ko/	*$\mu_{\alpha\underset{\circ}{A}}\mu_{\alpha\underset{\circ}{A}}$ & *[-voi]V̥[-voi]	*V̥μ↘	*Ḁ	*[-voi]V̥[-voi]	*V̥↑	*[-voi]V̥#	*V̥
→ [ha.ko] ↓ ↑ ↑				*		*	
1 [hạ̥.ko] ↓ ↑ ↑			*!			*	*

(xxviii) Tableau /hara↘/ → [ha.ɾa]

/ha.ra↘/	*$\mu_{\alpha\underset{\circ}{A}}\mu_{\alpha\underset{\circ}{A}}$ & *[-voi]V̥[-voi]	*V̥μ↘	*Ḁ	*[-voi]V̥[-voi]	*V̥↑	*[-voi]V̥#	*V̥
→ [ha.ɾa] ↑ ↑ ↓							
1 [hạ̥.ɾa] ↑ ↑ ↓			*!		*		*

As shown in these two tableaux, however, a specific violation of the conjunction constraint *$\mu_{\alpha\underset{\circ}{A}}\mu_{\alpha\underset{\circ}{A}}$ & *[-voi]V̥[-voi] is necessary to overcome the context-free markedness constraint *Ḁ, which otherwise causes all [-high, -voi] vowels to be disallowed. In both cases, *Ḁ is the constraint that blocks devoicing.

3.5 SUMMARY

In this analysis, several phenomena have been introduced which are governed by several constraints, which are depicted below in this Hasse diagram:

(xxxix) Hasse diagram

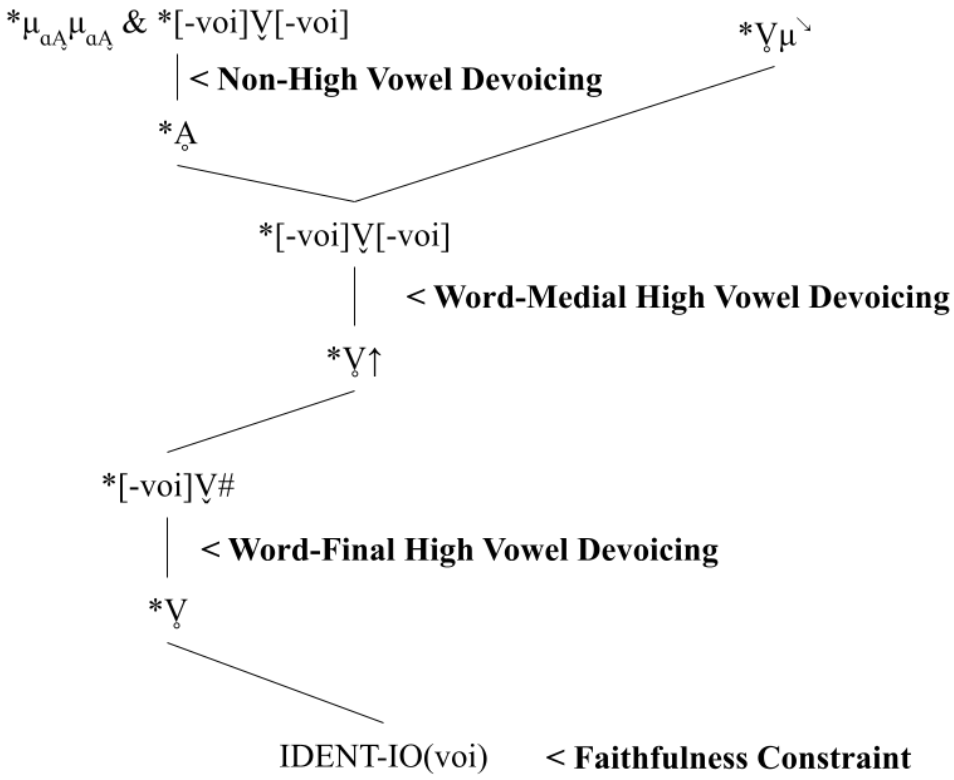

Although the relationship between the eight constraints shown is relatively linear (with the exception of *V̥μ↘) it can be seen as being separated between three different separate phenomena, which all correspond to an allophonic variation as described by Kager (1999). In Word-Final High Vowel Devoicing, *V̥ is overcome by *[-voi]V̥#, meaning that vowels are pressured to be devoiced at the end of word if preceded by a voiceless consonant. However, the higher ranked *V̥↑, which marks all voiceless vowels of high pitch, blocks this phenomenon. In the same way, Word-Medial High Vowel Devoicing (devoicing vowels between two [-voi] segments) can devoice vowels of high pitch, but not [-high] vowels. Furthermore, Non-High Vowel Devoicing acts on all vowels. In each of these cases, the phenomenon can be understood as a context-sensitive markedness constraint being ranked over a context-free markedness constraint, which is ultimately ranked over the bottom ranked IDENT-IO(voi), the common faithfulness constraint. In this system, *V̥μ↘ acts on its own, blocking devoicing specifically when the devoicing mora is before the downstep, and is not dominated by any known constraint.

4. OPEN QUESTIONS

The data used for this analysis was collected and compiled from research papers and dictionaries, and was confirmed with a native Japanese speaker (Sakae Fujita, personal interview, 2021). However, this analysis was done by ear, and not through modern research methods.

Therefore, the degree and prevalence of devoicing still remains further explained. The following nuances were observed, but ignored by this analysis:

(1) Vowels that could be devoiced word-finally showed variation in whether or not the vowel actually surfaced devoiced depending on speaker and rate of speech. What was described as a devoiced vowel in this analysis (such as in the word /ha↘chi/ [ha.t͡ɕi̥]) could more accurately be described as a vowel permitted to devoice, and the voiced counterpart (such as in the word /hachi/ [ha.t͡ɕi]), a vowel that devoices under no circumstances. In contrast, medial high-vowel devoicing, such as in the form /ku.sa↘/ [kɯ̥ɹ.sa], was present in the surface form regardless of rate of speech, unless hyper-articulated.

(2) Non-high vowel devoicing was significantly more subtle than high vowel devoicing, and was observed as a slightly quieter vowel sound. Some works, such as Tsuchida 2001, claim that no such devoicing occurs at all. It is also possible that this phenomena may be 'breathy voice', rather than true devoicing.

(3) In accordance with the variation outlined in Shaw 2018, the articulation of the voiceless counterpart of vowel [ɯ] can range from just devoicing to vowel shortening and, in some cases, complete deletion.

The exact nature of these phenomena, and any phonological predictability thereof, remains to be studied, and was not acknowledged in this analysis.

Additionally, it was observed that some morae, notably [ki] and [ɸɯ], showed resistance to word final devoicing, as demonstrated by the following forms.

(xxx) Forms ending in [ki] and [ɸɯ]

あき a.ki	[a.ki] ↑ ↓ ↓
ふうふ fu↘.u.fu	[ɸɯː.ɸɯ] ↑↓ ↓ ↓

These exceptions were not able to be explained on any phonological basis, and thus were not mentioned as part of this analysis. Whether any specific pattern can be gleaned from these morae, and/or if further morae seem to be similarly resistant to devoicing remains a topic of further study.

Finally, this analysis has posited that Japanese pitch accent is fully phonemic, and via the constraints and methods demonstrated above, devoicing can be fully predicted from underlying forms. This is, however, an assumption. While I have argued above that devoicing is restricted by pitch accent, the opposite may also be the case, that devoicing changes and restricts pitch accent, and that is the reason for their connection. Forms like /tsu.ku↘.su/ [t͡sɯ̥ɹ.kɯ.sɯ̥], which I have above explained as a case of pitch accent blocking devoicing, may in fact be a case of a pressure to create voiced-voiceless contours, as posited Tsuchida 2001, and the pitch accent was shifted to allow for the downstep to be on a voiced mora.

5. CONCLUSION

In conclusion, in this paper the variation between voiced and voiceless vowels has been argued to be allophonic variation, and this variation has been shown to be fully predictable given the underlying pitch accent. In accordance with Kondo (2005), this analysis posits that high vowels are devoiced between voiceless segments, and between a voiceless segment and the end of word. This analysis also claims that word-final devoicing can be blocked by that mora being of high-pitch, and that word-medial devoicing can be blocked by that mora being right before the downstep in pitch. Additionally, this analysis also explained the occurrence of non-high vowel devoicing, which is confined to word-medially between voiceless consonants, with the additional condition of the same vowel occurring twice in a row in two morae. This analysis also prompts further research into variance in the amount of voicelessness, as well as a few forms which seem to not abide with the proposed framework, specifically those ending in /ki/ or /fu/. Hopefully, more rigorous testing with modern scientific methods will be able to give more insight into these processes.

6. REFERENCES

Kager, René. 1999. "Conflicts in grammars." Optimality Theory, pp. 43-48. Cambridge: Cambridge University Press.

Kondo, Mariko. 2005. "Syllable structure and its acoustic effects on vowels in devoicing environments." *Voicing in Japanese,* Walter de Gruyter GmbH & Co. KG.

Labrune, Laurence. 2012. *The Phonology of Japanese*. Oxford University Press.

Matsumura, Akira; Sanseido Henshujo. 2006. 大辞林 / Daijirin. 3rd ed., 三省堂.

Shaw, Jason A.; and Shigeo Kawahara. 2018. The lingual articulation of devoiced /u/ in Tokyo Japanese. *Journal of Phonetics*, 66, 100-119.

Tsuchida, Ayako. 2001. Japanese Vowel Devoicing: Cases of Consecutive Devoicing Environments. *Journal of East Asian Linguistics*, *10*(3), 225–245. http://www.jstor.org/stable/20100803

7. NOTES

[1] Underlying representations are noted here in the Hepburn romanization system, which was devised for both ease-of-use for speakers of Western languages, especially English, and for its

correspondence with Japanese kana. Despite orthography often not being perfectly correspondent with a speaker's theoretical mental representation of their lexicon, in the case of Japanese, a language often described as moraic, the moraic writing system can be postulated to be reasonably close to Japanese speaker's mental representation, especially because Japanese orthography (in kana) has reasonably minimal variation from pronunciation in many dialects.

[2] In this tableau, there are violations of *[-voi]V̥[-voi], which are due to the word medial /a/. However, any form containing a devoiced /a/, rendering [hḁ...], would incur a violation of the highly ranked constraint *Ḁ, thus blocking it.

[3] In this tableau, there is consistently a violation of *[-voi]V̥# for every form, due to the final mora /ka/. However, any form containing a devoiced /a/, rendering [...kḁ], would incur a violation of the highly ranked constraint *Ḁ, thus blocking it.

[4] While the idea that voiceless segments cannot carry pitch and thus neutralize, or otherwise alter pitch contours is consistent with the classical understanding of Japanese phonology and with the data I personally collected, but as noted in Laburne 2012, this isn't necessarily the case. Words with a downstep on a devoiced mora in specific scenarios have been attested, but not observed in my findings.

8. DATA

The following is the entire dataset that was compiled and verified during research.

Japanese	**Romanization**	**Transcription**	**Glossary**
0: Forms used to exemplify pitch-accent. (devoicing not notated)			
箸	はし ha↘.shi	[ha.ɕi] ↑ ↓ ↓	"chopsticks"
橋	はし ha.shi↘	[ha.ɕi] ↑ ↑ ↓	"bridge"
端	はし ha.shi	[ha.ɕi] ↓ ↑ ↑	"edge"
1: Forms following word-medial high vowel devoicing patterns.			
草	くさ ku.sa↘	[kɯ̥.sa] ↑ ↑ ↓	'grass'
熊	くま ku↘.ma	[kɯ.ma] ↑ ↓ ↓	'bear'

今朝	かさ ke↘.sa	[ke.sa] ↑ ↓ ↓	'this morning'
下	した shi.ta	[ɕi̥.ta] ↓ ↑ ↑	'below'
島	しま shi.ma↘	[ɕi.ma] ↑ ↑ ↓	'island'
2: Forms following word-final high vowel devoicing patterns.			
出す	だす da↘.su	[da.sɯ̥] ↑ ↓ ↓	'to take out'
残す	のこす no.ko↘.su	[no.ko.sɯ̥] ↑ ↑ ↓ ↓	'to leave behind'
椅子	いす i.su	[i.sɯ] ↓ ↑ ↑	'chair'
バルサミコ酢	ばるさみこす ba.ru.sa.mi.ko↘.su	[ba.ɾu.sa.mi.ko.sɯ̥] ↑ ↑ ↑ ↑ ↑ ↓ ↓	'balsamic vinegar'
行く	いく i.ku	[i.kɯ] ↓ ↑ ↑	'to go'
行くか	いく か i.ku ka	[i.kɯ̥.ka] ↓ ↑ ↑	'shall we go?'
八	はち ha↘.chi	[ha.t͡ɕi̥] ↑ ↓ ↓	'eight'
蜂	はち ha.chi	[ha.t͡ɕi] ↓ ↑ ↑	'bee'
3: Forms with more than one possible high vowel devoicing environment			
2日	ふつか fu.tsu.ka	[ɸɯ̥.t͡sɯ̥.ka] ↓ ↑ ↑ ↑	'two days'
靴	くつ ku.tsu↘	[kɯ̥.t͡sɯ] ↑ ↑ ↓	'shoes'
靴下	くつした ku.tsu↘.shi.ta	[kɯ̥.t͡sɯ.ɕi̥.ta] ↑ ↑ ↓ ↓ ↓	'socks'
尽くす	つくす tsu.ku↘.su	[t͡sɯ̥.kɯ.sɯ̥] ↑ ↑ ↓	'to use up'

気質	きしつ ki.shi.tsu	[kɨ̥.ɕɨ̥.t͡sɯ] ↓ ↑ ↑ ↑	'temperament'
4: Forms with [-high] vowel devoicing			
墓	はか ha.ka↘	[hɑ̥.ka] ↑ ↑ ↓	'grave'
箱	はこ ha.ko	[ha.ko] ↓ ↑ ↑	'box'
心	こころ ko.ko.ro	[ko̥.ko.ro] ↓ ↑ ↑ ↑	'heart'
旗	はた ha.ta↘	[hɑ̥.ta] ↑ ↑ ↓	'flag'
腹	はら ha.ra↘	[ha.ɾa] ↑ ↑ ↓	'stomach'
5: Anomalous data			
秋	あき a.ki	[a.ki] ↑ ↓ ↓	'autumn'
夫婦	ふうふ fu↘.u.fu	[ɸɯː.ɸɯ] ↑↓ ↓ ↓	'husband and wife'