Uncovering Mandarin Speaker Knowledge with Language Game Experiments

by

Boer Fu

B.A., University of California, Los Angeles (2017)

Submitted to the Department of Linguistics and Philosophy in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2023

Author ....

Department of Linguistics and Philosophy

August 31, 2023

Certified by ....

Adam Albright

Professor of Linguistics, Thesis Supervisor

Accepted by ....

Danny Fox

Anshen-Chomsky Professor of Language & Thought, Department Head

Uncovering Mandarin Speaker Knowledge with Language Game Experiments

by

Boer Fu

B.A., University of California, Los Angeles (2017)

Submitted to the Department of Linguistics and Philosophy in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

Abstract

Mandarin Chinese offers many intriguing puzzles for linguists because it has a shortage of morphophonological alternations. This has resulted in indeterminacy in various aspects of its phonological grammar, triggering much debate on syllable structure and allophonic mapping. The ambiguity of the data is also a problem for children acquiring Mandarin since alternative grammars can account for the surface forms equally well.

In order to find out what Mandarin speakers have learned about the phonology of their language, I conducted two language game experiments based on fanqie secret languages. It was found that markedness and faithfulness constraints are psychologically real for Mandarin speakers. Furthermore, the interactions between markedness and faithfulness constraints are shown to have an effect on glide movement in the language game. In addition, much speaker variation was observed in the experiment. I demonstrate that it is the result of constraint ranking variation. Nevertheless, general population-level trends on constraint ranking could still be identified. These trends lead to insights on phonological learning beyond Mandarin, showing evidence for naturalness bias and lexicon optimization.

Thesis Supervisor: Adam Albright

Title: Professor of Linguistics

Acknowledgements

First of all, I wish to thank my thesis committee members. Thanks to Michael Kenstowicz for providing insights on Mandarin phonology during our discussion over the years. I would like to thank Edward Flemming for helping me navigate through the experiment design and statistical analysis process. Many thanks to Donca Steriade, for debating with me and getting me out of rabbit holes. Last but not least, I say thank you to Adam Albright, for always reminding me why my project is exciting.

Special thanks to Anton Kukhto for assisting me to conduct the experiment. This thesis also benefitted from the audience feedback at AMP2022, CLS59, TEAL13, and MIT Phonology Circle. I appreciate the insights and suggestions shared by conference attendees. It goes without saying that all remaining mistakes are my own.

I am also grateful for the friendship and support from my cohort mates, as well as many others at the department. I will refrain from compiling a list of names, for fear of missing anyone. You have made my life brighter for the past six years in Stata.

Thanks to my parents, for always encouraging me to pursue my interest. Thanks to Patrick, for always believing in me.

A final thank you to the campus of MIT, for providing the secret tunnels for me to walk and think, and for keeping so many interesting science equipment in little-known corners to let me indulge in my curiosity.

Table of Contents

List of Tables

Chapter 1 Introduction

The study of phonology relies on a key source of data -- morphophonological alternations. For instance, we know that there is final devoicing in German, because the same noun stem ends with a voiceless consonant in singular form, but with a voiced consonant once a plural suffix is added. Works on vowel harmony often examine the vowel quality in a suffix, in order to investigate long-distance effects. Another example is trisyllabic laxing in English. The vowel sound change could only have been identified with the observation of morphologically related words. The list goes on and on. The different surface forms of the same morpheme in different phonological environments can tell us much about phonological rules, underlying representations, input-output mappings, and constraint ranking. Morphophonological alternations have contributed to much progress in the field of phonology.

But what happens if a language does not provide phonologists with enough information on morphophonological alternations? Can we still identify rules or constraints using surface forms that appear to have no relations to each other? Furthermore, what do language learners do when they are confronted with such a poverty of the stimulus? Do they acquire a consistent phonological grammar? Do speakers of the same language acquire the same grammar? These are the questions I aim to explore in this study.

1.1 Mandarin Chinese

Mandarin Chinese, with a shortage of morphophonological alternations, is a classic test case. Granted, the language is famous for its tone sandhi processes, which sees the same morpheme alternating in its tonal specification depending on the context. But when it comes to segmental phonology, there is only one suffix, the diminutive r-suffix, that can trigger alternations. All other word-formation processes in Mandarin are best described as the concatenation of two morphemes, involving no segmental change. For example, the plural morpheme [men] can be productively added to any animate noun, without changing segments in the noun or in itself.

For a linguist attempting to write a phonological grammar of Mandarin, the rarity of segmental morphophonological processes introduces indeterminacy in many aspects of its phonology. This has led to heated debates in the field. One such debate concerns the phonemic status of the palatal sibilants. They are in complementary distribution with three other groups of consonants, denti-alveolar sibilants, retroflexes, and velars. The palatal sibilants only appear before palatal glides ([ɕja] vs. *[ɕa]), while the other three groups never do ([sa] vs. *[sja], [ʂa] vs. *[ʂja], [xa] vs. *[xja]). There is no morphological process in the language that could have placed the palatal sibilant in an environment that does not contain the palatal glide. Therefore, the Mandarin lexicon does not provide any evidence for alternation between the palatals and one of the other three types of consonants. The indeterminacy problem in mapping the palatals was well described by Chao (1934). But this has not stopped linguists from trying. Many proposals have been made to map the palatals to another group of consonants. Based on diachronic sound change evidence, some scholars map the palatals to the velars, while others link them with the denti-alveolars. Still others claim the palatals alternate with both the velars and denti-alveolars. At the same time, many phonologists believe that there is no place for palatal place alternation in the synchronic grammar of Mandarin. Scholars also disagree with each other on the syllable structure of the language, especially concerning the prenuclear glide, which is situated between the initial consonant and the nucleus vowel. Some suggest that the glide is part of the onset, while others argue it is part of the rhyme.

1.2 A phonological learning problem

While structuralists debate on how to write a grammar for Mandarin, children acquiring the language are also faced with a number of challenges as a result of the shortage of morphophonological alternations. These obstacles are best explained in Optimality Theory (OT) constraint terms.

First of all, it is difficult to establish phonotactic constraints for Mandarin. Children can observe that the language has a fixed set of attested syllables. They can probably infer from this list which syllables are possible but unattested. Various generalizations can be made to distinguish the two set. But the question is which generalization should be learned as a proper phonotactic constraint for their grammar, and which should be treated as noise in the data. It is possible that some syllables are unattested by accident, not due to systematic phonological markedness.

Even if Mandarin learners have no trouble identifying relevant phonotactic constraints, there is still the challenge in figuring out what to do with syllables that are ill-formed. The lexicon only informs them on which unattested syllables should never surface, yet there is no instruction on how to repair it if one is found in the input. Take the example of the puzzling palatal sibilants. If an ill-formed palatal-initial syllable like /ɕa/ is fed into the input, should the speaker map it to an output that changes the place of articulation, as in [sa], or to one that inserts a glide, as in [ɕja]? In terms of faithfulness constraints, should they prioritize consonant identity or make sure no new segment is added to the input? Without evidence for morphophonological alternations, learning input-output mapping and faithfulness constraint ranking is a tricky business.

The challenges faced by children acquiring Mandarin represent problems in phonological learning in general. For phonotactics, do learners draw a line between systematic gaps and accidental gaps (see S. Wang 1998; Myers 2015; Gong & J. Zhang 2021; Hayes & White 2013)? The former group point to unattested forms that are ruled out by genuine phonological constraints, whereas the latter group are not. When it comes to faithfulness constraints, what can speakers learn in the absence of morphophonological alternations? In particular, can speakers acquire underlying representations (UR) that are different from their surface representations (SR)? This is a recurring question in the study of phonological acquisition (see Yip 1996; Alderete & Tesar 2002; McCarthy 2005; Rasin & Shefi 2019; Richter 2021). A more general question to be asked is what do speakers learn when they are exposed to ambiguous phonological input?

These are phonological acquisition questions not unique to Mandarin. With an understanding of how Mandarin speakers learn phonology, we can gain insight into universal generalizations on phonological learning. As a first step, I set out to find out what adult speakers of Mandarin have learned about the phonology of their language, which can reveal biases in the learning process. This is done with the help of a series of language game experiments.

1.3 The language game experiments

I conducted two language game experiments to uncover what Mandarin speakers know about their language. The game is based on fanqie secret languages (Chao 1931), which slice a syllable into two halves, before using them as ingredients to construct new forms for the purpose of secret communication. In the experiment, speakers of Mandarin are invited to learn a new way of encoding secret messages, which takes in a disyllabic word and swaps its two initial consonants, deriving a new disyllabic codeword. What speakers choose to produce as their codeword response can inform us on their phonological knowledge.

A conventional use of fanqie secret languages is to investigate syllable structure. By observing whether the glide moves with the consonant or the vowel, one can identify the glide's position in syllable structure. However, it was pointed out by Yip (2003) that fanqie secret languages do not provide conclusive evidence on syllable structure. The movement of the glide is better predicted by phonotactics and constraint interactions in the language. Therefore, I have repurposed the fanqie method to gain insights on the constraint-based grammar acquired by Mandarin speakers. The test stimuli are controlled for the phonological environment around the prenuclear glide, in order to detect the effect of phonotactic markedness constraints and faithfulness constraints. The introduction of phonological control, as well as the redefined goal of understanding constraint interactions, mark the main difference between this study and similar previous work (H. Wang & Chang 2001; Barnes 2002), which were conducted to identify syllable structure and segmentation.

The speakers' responses from the two experiments show clear signs of phonological constraints at work. Markedness constraints and faithfulness constraints are psychologically real for Mandarin speakers. Their interactions can predict the glide movement in the language game without reference to syllable structure. In addition, speaker variation was observed in the codeword responses, which I argue stems from a speaker variation in the constraint ranking, as opposed to a variation in syllable structure. In spite of constraint ranking variation, general population-level trends can still be identified. Speakers are less likely to violate phonotactic markedness constraints grounded in phonological principles, which supports a naturalness bias in phonological learning. In the case of the palatal sibilants, it is shown that consonant identity faithfulness constraints are prioritized in the speakers' grammar. The palatal sibilants are rarely mapped to a different type of consonant in the output. The results indicate that without morphophonological alternations, speakers will learn an underlying representation identical to the surface representation, providing evidence for lexicon optimization.

1.4 Overview

The thesis is organized as follows. Chapter 2 is a review of two debates in the Mandarin phonology literature that concern the prenuclear glide. It begins with an introduction to the Mandarin syllable and why it is a significant unit of sound in the phonology of the language. I then provide a survey on the different camps in the syllable structure debate. This is followed by an outline of the common methods used by phonologists to determine the structural constituency of the glide, including fanqie secret languages and speech error data. I discuss their limitations, based on Yip's (2003) observations. I also offer my proposal for modifying and repurposing the fanqie method, which I implemented in the language game experiments. An overview of a similar glide debate in Bulgarian and the application of the language game method by Barnes (2002) complete the section on syllable structure. The attention is then turned towards the palatal series debate. I present different sides of the debate and demonstrate how the language game experiment can inform us on the mapping of the palatal series. The final section of this chapter is dedicated to the phonological learning problems in Mandarin. I discuss why phonotactic markedness constraints and faithfulness constraints are difficult to learn when there is a shortage of morphophonological alternations.

In Chapter 3, I provide a detailed account of the potential phonotactic constraints in Mandarin that capture the distribution of the prenuclear glide. I split the discussion into two halves. The first is concerned with the interaction between the glide and the following vowel. The second addresses the co-occurrence restrictions between the glide and the preceding consonant. Glide-vowel phenomena discussed include feature agreement between the glide and the mid vowel and low vowel raising triggered by the palatal glides. Consonant-glide phonotactics surveyed include the predictability of palatal glides after palatal consonants, the avoidance of labial glides after labial consonants, as well as a few relatively unnatural constraints restricting the distribution of the labiopalatal glide. I include a discussion on the prevalence of feature agreement constraints in Mandarin, relating it to a preference to lengthen the duration of perceptual cues. I conclude the chapter with a section on the potential faithfulness constraints in the grammar of Mandarin. I demonstrate that different constraint rankings make different predictions on input-output mapping.

Chapter 4 reports on the methods and findings of the two codeword language game experiments. The controls on phonological environment in the test stimuli are explained in full. The first experiment was conducted online during the covid pandemic. Speaker responses were collected as text input, which led to problems in data interpretation. These issues were subsequently addressed in the second experiment, which was conducted in person. Speakers were recorded producing their response tokens. The improved data collection method made the results clearer for interpretation. With a series of statistical analysis, I show that the behavior of the glide in the speakers' responses can be predicted by phonotactics, without reference to syllable structure. Furthermore, the effect of markedness and faithfulness constraints in response distribution indicates that they are psychologically real for Mandarin speakers.

A general discussion on the experiment results is included in Chapter 5. It begins with a section on the speaker variation observed in the codeword responses, which aims to identify the source of variation. I compare two accounts, syllable structure variation and constraint ranking variation. I argue against the former and in favor of the latter. I also demonstrate how different rankings of constraints can predict the speaker variation observed in the codeword responses. This is followed by a closer inspection of the unattested syllables produced by speakers as part of their responses. The goal is to find out which constraints are better learned by Mandarin speakers. I show that there is a naturalness bias in learning phonotactic constraints. Finally, I discuss speakers' treatment of the palatal sibilants and demonstrate that there is a preferred ranking of faithfulness constraints, despite the lack of evidence for alternations. I argue that the experiment results for the palatal sibilants show strong support for lexicon optimization. Chapter 6 concludes.

Chapter 2 Questions Concerning the Mandarin Glide

The study of Mandarin phonology is both hampered and propelled by the language's shortage of morphophonological processes. On the one hand, without morphological alternation, phonologists have a hard time looking for evidence for allophonic alternation. But at the same time, such an ambiguity in the sound system of the language has led to a proliferation of competing grammars for Mandarin on the market. This in turn poses intriguing questions about the acquisition of phonology. If trained grammarians and phonologists cannot wade through the multitude of alternative accounts in the literature to identify the definitive Mandarin grammar, surely Mandarin-acquiring children are overwhelmed by the same colossal challenge. In other words, how does a Mandarin learner acquire a consistent grammar from an ambiguous phonological input? Additionally, do all learners agree on the same grammar?

In this chapter, I introduce two longstanding debates in the Mandarin phonology literature, both concerning the prenuclear glide. The first one is a debate on the glide's position in syllable structure (Section 2.1). The second one concerns the phonemic status of the palatal sibilants, whose distribution is sensitive to the glide (Section 2.2). In Section 2.3, I show that the two debates framed in the tradition of structuralism can be translated into questions on the learning of phonological constraints in OT terms. Proposals for a language game experiment are interspersed in the discussion. The experiment makes use of the fanqie method favored by participants to the two debates, but I show that it can shed light on the role of markedness and faithfulness constraints in the grammar acquired by Mandarin speakers.

2.1 Syllable structure

There is a long-standing debate on the structural position of the prenuclear glide in a Mandarin syllable. First, I offer an introduction to the Mandarin syllable and its significance in Mandarin phonology in Section 2.1.1. The different camps in the syllable structure debate are surveyed in Section 2.1.2. Common methods used by phonologists, including fanqie secret languages (Section 2.1.3), speech errors (2.1.4), along with their limitations are discussed. Section 2.1.5 is an overview of the modifications that I have made to these conventional methods for my language game experiments. In Section 2.1.6, I summarize the application of the language game experiment method in a similar puzzle on Bulgarian palatalization (Barnes 2002).

2.1.1 The Mandarin syllable

The syllable is a phonological unit of special importance in Mandarin, as well as in many other Chinese languages. A syllable usually corresponds to a word or morpheme. It employs a fixed template: CGVX, where C is the initial consonant, G is the prenuclear glide, V is the nucleus vowel, and X stands for a coda that can either be a nasal (N) or an offglide (G). There are three glides in Mandarin: the palatal /j/, the labial /w/, and the labiopalatal /ɥ/. They are cognates with the high vowels /i/, /u/, and /y/. They can all appear in the prenuclear glide position. Out of the four components of the syllable, only V is obligatory. All other members, C, G, X are optional, as in (C)(G)V(X). Some examples of Mandarin syllables are provided in (1).

(1) *Mandarin syllables*

a.	[ljɑŋ35]	CGVN	'cold'	g.	[lɥe51]	CGV	'to omit'
b.	[xwej55]	CGVG	'gray'	h.	[en55]	VXN	'gratitude'
c.	[ɥɛn35]	GVN	'circle'	i.	[aw55]	VXG	'concave'
d.	[waj51]	GVG	'outside'	j.	[je51]	GV	'leaf'
e.	[san55]	CVN	'three'	k.	[ʈʂu55]	CV	'pig'
f.	[low35]	CVG	'building'	l.	[ə35]	V	'goose'

The Mandarin syllable is almost impermeable to any form of cross-syllable segmental phonological processes. For instance, there is no resyllabification between neighboring syllables. The only exception to the rigidity of the syllable is perhaps the r-suffixation process, in which the diminutive suffix -ɻ is added to the right edge of the syllable, capable of triggering coda deletion and changing vowel quality. The left edge of the syllable, namely the CG portion, is not affected by any morphophonological process, which leads to ambiguity in the interpretation of the CG sound sequence. Scholars working on Mandarin phonology disagree on the structural position and even the segment status of the prenuclear glide.

2.1.2 The prenuclear glide debate

The syllable structure debate seeks to address a question left open by the traditional description of Chinese that dates back centuries. In the traditional view, a Mandarin syllable is divided into three components: the initial (C), the medial (G), and the final (VX). There is no indication whether the medial is structurally closer to the initial or the final. Some scholars proposed a CG constituent, where the medial is closer to the initial (see Firth & Rogers 1937; Bao 1990; Y. Wu 1994; H. Wang 1999; Lin 2002). Others argued that the medial is structurally closer to the final, therefore claiming GVX to be a constituent (see R. Cheng 1966; Chao 1968; C. Cheng 1973; L. Wang 1980; Hsueh 1986; Ji 1988; Lin 1989; Baxter 1992; Kuo 1994).

A second question concerning the glide asks whether it is a segment or not in the UR. We can safely assume that everyone who participated in the syllable structure debate is in agreement that the glide is a segment, for otherwise it cannot be analyzed as a node in the syllable hierarchy. Head of the challenge against establishing the prenuclear glide as a segment is Duanmu (1990; 2007). He observed that the production of the Mandarin syllable [swan55] 'sour' is very different from that of the English word *sway*. In the former, the lip rounding begins at the onset of the initial consonant. In the latter, the gesture for labialization only starts after the sibilant fricative. This is taken as evidence by Duanmu that the prenuclear glide is merely the secondary articulation of the initial consonant, as in [sʷan55].

Another challenge to the segment status of the prenuclear glide comes from (Ladefoged & Maddieson 1996:150). They claimed that after palatal consonants, the palatal glide is not a segment, but a predictable consonant-vowel transition period. This is probably an analysis proposed to account for the interesting distribution of /j/. It is contrastive after non-palatal consonants, but obligatory after palatal ones. This is illustrated in (2). With the denti-alveolar /l/ as the initial consonant, one can find a minimal pair (2a & b) that contrasts in the presence or absence of /j/. After the palatal fricative /ɕ/, a palatal glide is obligatory. (2c) *[ɕa] is unattested and ill-formed, while (2d) [ɕja55] 'shrimp' is well-formed. The predictability of the palatal glide after a palatal sibilant has led Ladefoged & Maddieson to analyze it as a natural CV transition. They observed that in the syllable [ɕja55] 'shrimp', which they transcribed as [ɕa55], the formant transition between the palatal fricative /ɕ/ and the nucleus vowel /a/ is nothing more than what one would expect from a natural CV transition.

(2) *Distribution of the palatal glides /j/ and /ɥ/*

a.	[la51]	'wax'	c.	*[ɕa]	*Unattested*	e.	*[ɕe]	*Unattested*
b.	[lja214]	'pair'	d.	[ɕja55]	'shrimp'	f.	[ɕje35]	'shoe'
						g.	[ɕɥe214]	'snow'

Note that the obligatory palatal glide does not have to be /j/. A labiopalatal /ɥ/ can appear after palatal consonants as well, as in (2g) [ɕɥe214] 'snow'. It forms a minimal pair with (2f) [ɕje35] 'shoe'. Syllables like [ɕɥe214] are omitted in Ladefoged & Maddieson (1996) report, which was restricted to the non-labialized palatal glide /j/. Their natural CV transition account did not offer a solution for transcribing (2g) [ɕɥe214], but there is an easy fix. The labialization can be transcribed as a secondary articulation of the palatal fricative, as in [ɕʷe214].

The debate on the segment status of the glide can be viewed as an extension of the debate on its structural position in the syllable. Duanmu's (1990) treatment of the glide as a secondary articulation of the consonant indicates that he agreed with establishing a CG constituent. Ladefoged & Maddieson's (1996) analysis of the palatal glide as a natural CV transition is also in general alignment with the CG constituent camp since the palatal CV transition is always found after palatal consonants. In the next two section, I provide an overview of the common methods phonologists use to identify the structural constituency of the glide.

2.1.3 Common methods and their limitations: fanqie secret languages

Phonologists engaged in the prenuclear glide debate mainly draw evidence from two sources, fanqie secret languages and speech errors. They both involve carving a Mandarin syllable into two halves, which are used as ingredients for a new construction, either intentionally or accidentally. Scholars examine which half contains the prenuclear glide, in order to identify its structural constituency. Yet they reach divergent conclusions. Yip (2003) offered some diagnosis as to why secret language and speech error data are inconclusive. Analyses of such artificial segmental processes tend to gloss over the role of phonotactics and fail to take speaker variation into account. I discuss the two methods and their limitations in turn. Modifications are proposed in Section 2.1.5.

A main source of evidence for the prenuclear glide debate is Chinese secret languages based on the fanqie method. Fanqie is a way to notate the pronunciation of Chinese characters in traditional rhyme books such as *Qieyun* (601 CE). First, a syllable is split into two parts, the initial and the final. Then, two characters are selected to represent the syllable, because they share the initial and the final with it, respectively. For instance, the character 東 [tʊŋ55] is spelt with 德 [tə35] and 紅 [xʊŋ35] using the fanqie method.

Chao (1931) reported eight types of secret languages in various Chinese languages and dialects that make use of the fanqie method. Secret languages are often created by a small community for the purpose of secret communication. Even though they are not devised by phonologists, they usually adhere to phonological principles which are representative of the grammar of the native language the creators speak. In fanqie secret languages, a syllable is split into two by using a template. A classic example is the *May-ka* secret language created by children of Beijing at the time of Chao's writing. In May-ka, a syllable of the shape CVX is split into two with the template of Caj kVX. The first syllable is composed of the original syllable's initial and the final [aj], or *ay*. The second syllable has [k] as the initial consonant, followed by the original syllable's final. To give an example, Chao's name [ʈʂaw51] can be transformed into [ʈʂaj214 kaw51].

Many linguists track the prenuclear glide's movement in fanqie secret languages, in order to identify its structural position. If the glide follows the initial, then it is taken to indicate that C and G form one constituent. On the other hand, if the glide follows the final, then it means G and V form one constituent. This is a diagnostic used by Bao (1990). In his analysis of May-ka, a syllable is first reduplicated into two syllables, an o-syllable and an r-syllable (o stands for onset and r stands for rhyme). An illustration of this process is provided in (3). Then, materials in the two syllables are replaced by the template. In the o-syllable, the rhyme [aw] is replaced by the templatic rhyme [aj]. In the r-syllable, the onset [ʈʂ] is replaced by [k]. Thus, one gets [ʈʂaj kaw]. Portions of the syllables to be replaced in (3b) are crossed out. Materials kept are highlighted in bold.

(3) (Bao 1990)*'s analysis of May-ka secret language*

	Reduplication			*Replacement*		
a.	[tʂaw]	→	b.	[**tʂ**~~aw~~ ~~tʂ~~**aw**]	→	c. [**tʂ**aj k**aw**]
				o-syl. *r-syl.*		*o-syl.* *r-syl.*

Bao observed that (4a) [t^hwo] is split into (4c) [t^hwaj kwo], not (4f) *[t^haj kwo]. He argued that the form [t^hwaj kwo] is only possible if the glide is part of the onset. The clue is in the o-syllable. If the glide /w/ is indeed part of the onset, then it will not get replaced by the templatic rhyme [aj]. The reduplication and replacement process is illustrated in (4a-c). However, had the glide been part of the rhyme, as shown by (4d-f), then rhyme replacement would have taken it out of the o-syllable, deriving *[t^haj kwo] instead. Therefore, Bao concluded that /w/ is part of the onset.

(4) *Bao's diagnostic for glide constituency*

	Reduplication			*Replacement*		
	If glide is part of the onset					
a.	[t^hwo]	→	b.	[**t^hw**~~o~~ ~~t^h~~**wo**]	→	c. [**t^hw**aj k**wo**]
				o-syl. *r-syl.*		*o-syl.* *r-syl.*
	Alternatively, if glide is part of the rhyme					
d.	[t^hwo]	→	e.	[**t^h**~~wo~~ ~~t^h~~**wo**]	→	f. *[**t^h**aj k**wo**]
				o-syl. *r-syl.*		*o-syl.* *r-syl.*

Yip (2003) pointed out that an obvious problem with Bao's analysis is that the glide in the r-syllable [kwo] remains unaccounted for. If /w/ is indeed part of the onset, we should expect the whole onset [t^hw] to be replaced in the r-syllable, deriving [ko] instead. The double copy of the glide did not go unnoticed by Bao. He proposed a zero onset preceding the glide, which somehow shields it from replacement. Nevertheless, Yip pointed out that the nature and mechanism of this zero onset is unclear.

In addition, Yip noted that secret language data can lead to divergent conclusions on glide constituency. She compared the results of May-ka and another secret language of Mandarin, Mey-ka, and showed that the palatal glide /j/ behaves differently in the two secret languages. Mey-ka operates in a similar way as May-ka, with the difference that it uses the template Cej-kVX instead. In May-ka, (5a) [ljɑŋ] is transformed into (5b) [ljɛ tɕjɑŋ], which is a repair of the ill-formed (5c) *[ljaj kjɑŋ] (for more on the observed velar palatalization, see Section 2.2). Mey-ka takes the same syllable and turns it into (5e) [le kjɑŋ], and not the expected (5f) *[ljej kjɑŋ]. May-ka contains two copies of /j/, whereas Mey-ka includes only one /j/ alongside the rhyme. May-ka would lead one to posit a CG onset constituent, but Mey-ka would suggest otherwise.

(5) *May-ka and Mey-ka show different conclusions*

	May-ka			
a.	[ljɑŋ]	→	b.	[ljɛ tɕjɑŋ]
		Not	c.	*[ljaj kjɑŋ]
	Mey-ka			
d.	[ljɑŋ]	→	e.	[le kjɑŋ]
		Not	f.	*[ljej kjɑŋ]

Yip argued that the difference between May-ka and Mey-ka stems from phonotactics, not glide constituency. In Yip's (1982) analysis of fanqie secret languages, she agreed with Bao (1990) that they begin with the process of reduplication. However, the subsequent steps are a bit different. She proposed that the template [ay-k] is used to replace some materials from the reduplicated form, with the preference to realize as much of the original materials from [ljɑŋ] as possible. This means that the default output always involves two copies of the palatal glide. May-ka would produce (5c) *[ljaj kjɑŋ], whereas Mey-ka would output (5f) *[ljej kjɑŋ]. The initial syllables in both forms are ill-formed, violating a general phonotactic constraint in Mandarin against having identical glides in a GVG sequence. The preferred repair in May-ka involves changing the diphthong in *[ljaj] into a monophthong, as in [ljɛ]. Mey-ka, on the other hand, simply gets rid of the glide in *[ljej] to avoid any violation of the phonotactic constraint.

With evidence drawn from various other Chinese secret languages, Yip demonstrated that phonotactics concerning the glide is a much better predictor of glide behavior in secret languages than glide constituency. Therefore, she concluded that secret languages data do not provide conclusive evidence on whether the glide is structurally part of the onset or the rhyme.

2.1.4 Common methods and their limitations: speech errors

The behavior of the prenuclear glide in speech errors has also been used to elucidate its structural position. This is the method employed by Wan (1999). In her thesis, she inspected naturally occurring speech errors made by Mandarin speakers, focusing on whether the glide acts as a unit with the initial or with the final. Wan found that the behavior of both the palatal glide and the labial glide is sensitive to the place of articulation of the consonant. /j/ moves with the initial consonant if it is also palatal, but separately otherwise. /w/ acts as one unit with a preceding velar consonant, but not with consonants of other places of articulation. (6) includes some examples of speech errors drawn from Wan's survey. Copied materials in the speech errors are in bold.

(6) *Speech errors provide evidence for glide constituency* (Wan 1999)

	Intended			Produced	Type of error	CG combo	Glide constituency
a.	[pʰ**jɛn**55 … pej51]	→	e.	[pʰ**jɛn**55 …p**jɛn**51]	*Rhyme copy*	Labial+j	GVX: [jɛn]
b.	[**tɕj**a55 tow55]	→	f.	[**tɕj**a55 **tɕj**ow55]	*Onset copy*	Palatal+j	CG: [tɕj]
c.	[tsɑŋ55 **ʂ**wej21]	→	g.	[**ʂ**ɑŋ55 **ʂ**wej21]	*Onset copy*	Retroflex+w	GVX: [wej]
d.	[**f**ej51 xwa51]	→	h.	[**f**ej51 **f**a51]	*Onset copy*	Velar+w	CG: [xw]

In (6a), the intended syllable [pej51] has its rhyme replaced with that of an earlier syllable [pʰjɛn55], resulting in [pjɛn51] in (6e). The palatal glide forms a constituent with the final when it is preceded by a labial consonant. In contrast, when it follows a palatal consonant, as in (6b) [tɕja55], it moves with the consonant as a single unit, resulting in (6f) [tɕja55 tɕjow55]. The two speech errors involving the labial glide tell a similar story. In (6g), as the retroflex [ʂ] is copied into a neighboring syllable, [w] stays put as part of the rhyme [wej]. (6d) is an instance of onset copying, in which the original onset [xw], including the glide, is replaced in whole by the onset of the neighboring syllable, [f].

The speech error data show that the behavior of the glide is sensitive to the place of articulation of the consonant. Wan concluded that glide constituency is also sensitive to consonant place. This is quite a leap. Yip (2003) pointed out that much of the glide movement in speech errors can be accounted for using phonotactics, in a similar fashion as secret languages. For example, for the speech error of (6f) [tɕja55 tɕjow55], had the speaker not moved the palatal glide with the consonant, they would end up with an ill-formed syllable *[tɕow55]. As to the labial glide in (6d) [fej51 xwa51], had it stayed put, it would have resulted in an ill-formed *[fwa51] after onset copying. In her survey, Wan made the observation that Mandarin speakers almost always produce well-formed syllables in speech errors. This means that ill-formed syllables are repaired even as speakers are making mistakes with their speech production.

Another problem with using speech errors as evidence for glide constituency is the potential complication of speaker variation. Yip reasoned that a database of naturally occurring speech errors is usually pooled from many speakers. It is possible that given the poverty of the stimulus offered by Mandarin, learners of the language have reached different conclusions about what to do with the glide. If speech errors point to two different directions regarding the structural position of the glide, that might very well be the result of two speakers who disagree with each other on the issue. Similarly, speaker variation could be a potential complication for secret languages as well, since they are the invention of a community of speakers, each perhaps bringing their own phonological grammar to the production of coded words used for secret communication.

2.1.5 Proposal for a language game experiment

Yip's (2003) criticism against using secret language and speech error data as evidence for glide constituency can be summarized in two points. The first is that phonotactics might be the real predictor of glide behavior in secret languages and speech errors, instead of glide constituency. The second is that this type of data tend to be pooled across multiple speakers, where speaker variation could play a role in glide behavior. She argued that phonotactics and speaker variation are the real reason behind the glide's seemingly divergent behavior on different occasions. This might explain why the debate on syllable structure is not resolved.

I propose a language game experiment that tests Yip's hypothesis, repurposing the fanqie method in order to gain insights on Mandarin phonotactics. To tease apart the effect of syllable structure and phonotactics, I control for the phonological environment of the glide. The consonant on the left is controlled for its place of articulation. The nucleus vowel on the right is controlled for whether it shows potential alternation sensitive to the glide. As to speaker variation, I have decided to recruit more than 50 speakers to participate in the experiment. Every individual speaker's preference for glide treatment can be put under the microscope for inspection.

Conducting controlled fanqie language game experiments with a sizable participant pool is not a new idea. H. Wang & Chang (2001) ran two experiments that invited 97 Mandarin speakers to fuse two syllables into one and split one syllable into two, using the fanqie method. Their goal was to find out whether the glide was part of the onset or the rhyme. They also wanted to investigate if consonant place could have an effect in determining glide movement in the speakers' response, in a similar fashion to what Wan (1999) observed in speech errors. No such effect could be detected for /j/ or /w/. They also concluded that the glide is part of the rhyme based on the experiment data. While Wang and Chang controlled for consonant place, they did not consider vowel alternation as a potential predictor for speaker response. Since the extraction or addition of a glide has the potential to affect vowel quality (for more details, see Chapter 3), I include vowel alternation as a factor in my language game experiments.

The benefit of adding phonological controls in the language game is twofold. Firstly, we can find out whether the movement of the glide in the fanqie secret languages can be predicted by constraint interactions alone, without reference to syllable structure. As a preview of the experiment results, I found that both consonant place and vowel alternation were significant predictors of the glide's behavior in the fanqie language game. Secondly, by evaluating the effect of phonotactics, we can gain insights into the phonological grammar acquired Mandarin speakers. This added benefit is discussed in Section 2.3. But before I conclude this section on the Mandarin syllable structure debate, I append a brief description of a similar debate in Bulgarian concerning palatalization, where the language game method is also applied to clarify the position of the glide.

2.1.6 Bulgarian palatalization and the language game method

The prenuclear glide debate is not unique to Mandarin, similar questions have been asked about the segment status and structural position of glides in other languages as well (see Barnes 2002, Pritchard 2012, Suh & Hwang 2016). Most notably, the palatal glide in Bulgarian has invited as much debate on its segment status as the Mandarin counterpart. Standard Bulgarian has palatalized consonants, but it has been reported that some palatalized consonants have undergone depalatalization in some areas of Bulgaria. In this account, a palatalized consonant C^j has been depalatalized into C, while the feature of palatality is now encoded in a separate segment /j/. Thus, a single segment has been decomposed into two sequential segments. Proponents of singling out /j/ as an independent segment include Scatton (1975) and G. Choi (1998). Pritchard (2012) demonstrated that different methods lead to contradictory conclusions about the status of palatalized consonants. A production experiment showed that Bulgarian speakers' production of the sound is identical to that of English Cj, but not Russian C^j. This suggests that depalatalization has taken place in Bulgarian. One segment has been split into two. However, a perception experiment using the gating method, in which incremental short windows of the consonant are played to Bulgarian speakers, gave the results that listeners are good at identifying palatalized consonants with limited information, without the CV transition. Pritchard took this as evidence that the consonants are still palatalized.

To identify how the sound is represented in Bulgarian speakers' grammar, Barnes (2002) took a language game experimental approach, in the same vein as Chinese fanqie secret languages. In the game, speakers were given a two-word phrase and asked to reverse the first segment of each word. When a two-word phrase contains a palatalized consonant such as (7a) [d^jasna lapa], what the speakers choose to do with d^j can reflect its segment status. They can move d^j as a single unit, deriving (7b) [lasna d^japa], which would affirm that palatalized consonants are still palatalized in Bulgarian. They might also move *d* without taking any palatalization with it, as in (7c) [l^jasna dapa]. This would corroborate the report that Bulgarian has depalatalized its C^j consonants into two consecutive segments.

(7) *Bulgarian language game experiment* (Barnes 2002)

a. [d^jasna lapa] → b. Choice 1: [**l**asna **d**^{j}apa] C^j single unit

c. Choice 2: [**l**^{j}asna **d**apa] Cj cluster

Barnes canvassed Bulgarian speakers from three different regions. The first region selected was Kozičino and Golica villages in the northeast, whose Erkeč dialect shows no sign of depalatalization. The second location was Bankja, near Sofia, where depalatalization was reported. A third transitional region, Trjavna, located in central Bulgaria, was also selected for the experiment. Participants from these three regions did not receive explicit instruction on the rule of the language game. Instead, they were asked to figure it out based on a few examples like the ones in (8). CC cluster examples like (8b) [mlada kotka] were also included, in order to demonstrate to the speakers that only the first segment was supposed to move. In the experiment phase, test tokens include those containing C^j like (7a), control items with a single initial C such as (8a), as well as some with CC clusters like (8b). The goal is find out whether C^j patterns with simple C or CC cluster in the speaker's response.

(8) *Example tokens used in the Bulgarian language game experiment*

a. [malka tapa] → c. [**t**alka **m**apa] C

b. [mlada kotka] → d. [**k**lada **m**otka] CC cluster

Speakers had no trouble understanding the task from the examples. In particular, for CC cluster items, they only moved the first segment of the cluster. This indicates that they understood the language game task to be moving initial segments, as opposed to entire onsets. Speakers from all three regions showed a general preference for the Cj cluster response, which appeared to indicate that the sound is treated as a sequence of two sounds. However, Barnes reached the conclusion that palatalized consonants are still parsed as a single segment by Bulgarian speakers. He gave two reasons for this surprising conclusion. The first is that even though speakers preferred the cluster response for C^j items, they did not show a preference as strong as they did for genuine CC clusters. This was taken to indicate that C^j is not equivalent to CC clusters in the grammar of Bulgarian speakers. Secondly, there is a confound from the orthography of Bulgarian, in which the palatalization of C^j is transcribed in the following vowel, and not the consonant itself, in line with other Cyrillic scripts like Russian and Mongolian. During the experiment, some speakers were observed saying letters to themselves, or using their hands to move imaginary letters in thin air. Therefore it is possible that participants preferred the cluster response based on orthography. Barnes also pointed out that sometimes speakers chose to leave the palatalization in its original word to avoid creating an ill-formed word otherwise. Phonotactics seems to play a role in the Bulgarian version of the fanqie secret language as well. Unfortunately, Barnes did not prepare any control for phonotactics in his experiment to monitor its effect.

2.2 Phonemic status of the palatal sibilants

As I have mentioned in Section 2.1.2, the palatal sibilants can only appear before a palatal glide, /j/ or /ɥ/. There is more to the distribution of palatal sounds in Mandarin. The palatal sibilants /tɕ/, /tɕʰ/, /ɕ/ are in complementary distribution with three groups of consonants, the velars /k/, /kʰ/, /x/, the denti-alveolar sibilants /ts/, /tsʰ/, /s/, and the retroflex sibilants /ʈʂ/, /ʈʂʰ/, /ʂ/. The other three groups of consonants cannot co-occur with palatal glides. They are either immediately followed by a nucleus vowel or a labial glide /w/. The high vowels /i/, /y/, and /u/ behave in the same way as their glide counterparts /j/, /ɥ/, and /w/ regarding co-occurrence restrictions with the preceding consonant. The complementary distribution between the palatal series and the other series is demonstrated with examples in Table 2.1. Grayed cells indicate that the syllable of the shape is unattested.

The palatal series of Mandarin is a classic case of complementary distribution without alternation. This results in ambiguity in the phonemic status of the palatal series. Are they phonemes in their own right or allophones of another series of consonants? If it is the latter, then which of the three series should the palatals be identified with? Do they alternate with the velars, denti-alveolar sibilants, or the retroflex sibilants?

	CjVX	CɥVX	CwVX	CVX
Palatals: tɕ, tɕʰ, ɕ	[ɕje35] 'shoe' [tɕjɛn55] 'sharp' [tɕi55] 'chicken'	[ɕɥe214] 'snow' [tɕʰɥɛn35] 'spring' [tɕʰy214] 'music'		
Velars: k, kʰ, x			[kwo55] 'wok' [kʰwan55] 'wide' [xu35] 'lake'	[xə51] 'crane' [kan55] 'dry'
Denti-alveolar sibilants: ts, tsʰ, s			[tswo51] 'to do' [swan55] 'sour' [tsu55] 'to rent'	[sə51] 'color' [tsʰan55] 'meal'
Retroflex sibilants: tʂ, tʂʰ, ʂ			[tʂwo55] 'table' [tʂʰwan35] 'boat' [ʂu55] 'book'	[tʂʰə55] 'car' [ʂan55] 'hill'

Table 2.1 Complementary distribution between the palatal sibilants, velars, denti-alveolar sibilants, and retroflex sibilants

These questions were first raised by Chao (1934) in his seminal paper titled *The Non-uniqueness of Phonemic Solutions of Phonetic Systems*. The complementary distribution between the palatals and three distinct groups means that a one-to-one allophonic mapping is difficult to determine. Generations of phonologists studying Mandarin have made attempts to identify the palatal series with one of the three groups. Many claimed that the palatals ought to be grouped with the velars. They include Chao (1934, 1948, 1968), H. Wong (1953), M. Fu (1957), Martin (1957), R. Cheng (1966), Hsueh (1986), Lin (1989), and Y. Wu (1994). Others argued the palatals are associated with the denti-alveolar sibilants. Among them are Hartman (1944), Hockett (1947), and Duanmu (2007). Still others suggest that the palatals can be derived from both the velar group and the denti-alveolar group (see Dong 1964, C. Cheng 1968). According to W. Li (1999), no one has argued for grouping the palatals with the retroflexes.

The above scholars' preference for aligning the palatals with the velars and denti-alveolars, as opposed to the retroflexes, is a result of their understanding of diachronic sound change. Historically, the palatal sibilants originated in two waves of palatalization. The first wave of palatalization affected the velars, which started in the 13th century. Later, the second wave acted on denti-alveolars, around the 18th century (W. Li 1999). The retroflexes were never a source of modern palatal sibilants, therefore they are not considered to be a possible UR for the palatals.

While many phonologists based their analysis of the palatal series on diachronic sound change, C. Cheng (1973), among others, argued that such historical relationship should have no bearing on the synchronic grammar of Mandarin. Speakers of the language have no access to etymological information, therefore they are expected to posit the palatals as independent phonemes, since there is no morphophonological alternation to support any particular allophonic mapping with other consonants. This is a view taken by Tung (1954), Shi (1957), C. Cheng (1973), You, Qian & Gao (1980), Pulleyblank (1984), Ao (1992), Halliday (1992), Kuo (1994), W. Li (1999), etc.

Answers to this debate can also be found in secret languages and speech errors. Both are Frankenstein processes that patch together parts of two syllables, creating novel phonological environment not naturally available in Mandarin, which provide the opportunity for potential alternations. This means that data from the two sources might shed some light on the phonemic status of the palatal consonants.

In May-ka, the templatic [k] initial for the second syllable can sometimes find itself in contact with a palatal glide. This triggers palatalization, as seen in (5), in which [ljɑŋ] is transformed into [ljɛ tɕjɑŋ]. Had the template Caj kVX been placed on the original syllable (5a) [ljɑŋ] with brute force, then we might expect to see (5c) [ljaj kjɑŋ] as the result. However, this is not the case. Since both syllables are unattested, some phonotactic repair has been done, which resulted in (5b) [ljɛ tɕjɑŋ], where the velar consonant has been palatalized. Chao (1934) used the palatalization in May-ka as evidence for an allophonic relation between the palatals and the velars.

Similarly, speech errors offer examples of palatalization. Wan (1999) found 17 instances of palatalization in her survey of naturally occurring speech errors. Of these, 12 are originally denti-alveolars, five are retroflexes. There is no instance of velar palatalization. Wan also discovered 13 examples of palatal consonants changing their place of articulation with the removal of a palatal glide or high vowel. Of these, 12 errors involve palatals turning into denti-alveolars, and only one into a velar. The absence of palatal-to-retroflex place change is not unexpected, since the speech errors were produced by Mandarin speakers in Taiwan, where retroflexes are relatively marked and often neutralized with denti-alveolars. Given the overwhelming ratio of denti-alveolars in errors involving consonants changing into and out of palatals, Wan concluded that there is a "close link between denti-alveolars and palatals in Mandarin". She stopped short of claiming that denti-alveolars and palatals are a single set of phonemes.

In OT terms, the palatal series debate can be phrased as a question on input-output mapping. If Mandarin speakers are confronted with an ill-formed input, like *[ɕa] or *[sja], do they map them to a non-identical consonant in the output? Or do they apply other methods of phonotactic repair to derive a well-formed output?

In my study, the language game experiment occasionally created such an ill-formed syllable. Specifically, a palatal consonant might be introduced into a syllable without any palatal glide. What Mandarin speakers choose to do with the palatal consonants, whether they depalatalize them to another group, keep their place of articulation, or repair with glide insertion, can uncover how the palatals are mapped from the input to the output. We can also gain insights on the faithfulness constraint ranking in the grammar of Mandarin speakers.

2.3 Phonological constraint learning

Many questions in Mandarin phonotactics and phonological learning can be answered via the language game experiment. In Section 2.3.1, I demonstrate how the experiment can address questions on phonotactic constraint learning. In Section 2.3.2, I discuss what the experiment can tell us about faithfulness constraint learning.

2.3.1 Phonotactic constraint learning

For a start, a phonologist inspecting the syllables of Mandarin can identify many generalizations about the distribution of sounds. For instance, they might discover that palatal sibilants are always followed by a palatal glide /j/ or /ɥ/ and write a consonant-glide agreement constraint for the palatal feature. They might also find out that the Mandarin lexicon has plenty of words shaped like Cje, but never any like *Cjə. In other words, between the two mid vowels, only the front [e] is allowed after palatal glides, but not the central [ə]. This might lead the phonologist to posit a glide-vowel agreement constraint using the front feature.

The phonologist can write as many phonotactic constraints as they wish, but it is difficult to find out whether such constraints are a psychological reality for Mandarin speakers. The lexicon only provides positive evidence of surface forms that are legal in the language. For all the potential syllables that are not attested in the language, it is not clear whether they violate genuine markedness constraints stored in the speakers' grammar. In addition, it is also not known given such input, what kind of repairs speakers might make, in order to produce a legal output form. Another way of putting this question is to say, what faithfulness constraints speakers find acceptable to violate, when they prioritize satisfying markedness constraints?

There have been a few studies to gauge Mandarin speaker's knowledge on phonotactics using nonce words judgment experiments (see S. Wang 1998; Myers 2015; Gong & J. Zhang 2021). In such an experiment, a list of nonce words composed of Mandarin sounds are shown to the speakers, who are asked to judge whether each item is an acceptable word in Mandarin or not. Some are ill-formed according to the phonotactic constraints identified by phonologists, others are well-formed based on these constraints. The former is called a systematic gap, because such an attested syllable is ruled out due to underlying phonological principles. The latter is called an accidental gap, because they do not violate any obvious phonotactic constraints, but are often the byproduct of diachronic sound change. The results from these studies show that Mandarin speakers judge the two types of nonce words differently, preferring accidental gaps to systematic gaps. It indicates that Mandarin speakers have indeed internalized generalizations on attested and unattested syllables as phonotactic constraints.

My language game experiments follow the footsteps of these phonotactic judgement studies. The results support that both consonant-glide phonotactic constraints and glide-vowel phonotactic constraints are learned by Mandarin speakers as part of their grammar. It also offers speakers an opportunity to produce unattested syllables, by sometimes bringing a sound to an unfriendly phonological environment. Whether the participants produce the unattested syllable faithfully or repair it can inform us on how they have ranked the relevant phonotactic constraint. This is a different approach to the phonotactics question. Instead of playing nonce word tokens to the speakers and asking them to judge on its well-formedness, here, the speakers are prompted to produce unattested forms themselves. The criteria for acceptability used in these two tasks could be quite different.

2.3.2 Faithfulness constraint learning

The biggest learning obstacle that comes with a shortage of morphophonological alternations is that there is no transparent way to find out how faithfulness constraints are ranked in the language. Learners of Mandarin can draw conclusions on which syllables cannot surface, but there is no evidence that shows how they can be repaired to meet the standard of well-formedness. In other words, Mandarin learners do not have enough information to establish input-output mapping. Therefore, it becomes a challenge to know which faithfulness constraints are ranked higher, and which are to be ranked lower.

Furthermore, phonologists also have no access to the ranking of faithfulness constraints in the grammar acquired by Mandarin speakers. The opaque nature of faithfulness constraint ranking leads to the distinct possibility that there is no such ranking in the grammar of Mandarin, if speakers never have any reasons to map ill-formed syllables to well-formed ones.

The language game experiment can provide insights into this puzzle. During the game, Mandarin speakers might find themselves confronted with an ill-formed syllable, such as *[ɕa]. Whether they choose to repair it and how they choose to repair it can reveal much about the input-output mapping and the ranking of faithfulness constraints in their grammar.

In this section, I have outlined some learning problems concerning phonotactic markedness constraints and faithfulness constraints in Mandarin. In Chapter 3, I provide a review of the potential constraints that can account for the surface forms in Mandarin.

Chapter 3 Phonotactics of Mandarin

In Chapter 2, I mentioned that the language game experiments of this study can shed light on whether phonotactic constraints and faithfulness constraints are a psychological reality for Mandarin speakers. In this chapter, I provide an overview of some potential phonotactic constraints on the distribution of the prenuclear glide, as well as a survey of the potential faithfulness constraints that can direct input-output mapping. The goal is not to offer a full phonological grammar of the language, which is an ambitious endeavor that warrants a separate thesis of its own. Instead, I merely seek to lay out the potential constraints concerning the prenuclear glide and its interaction with neighboring sounds. The discussion is mostly restricted to constraints that can be used to account for the speaker response in the language game experiments, which I will demonstrate in Chapter 4 and Chapter 5.

The rest of the chapter is organized as follows. I introduce phonotactic constraints that govern the interaction between the glide and the following vowel in Section 3.1. I offer an account of the co-occurrence restrictions between the preceding consonant and the glide in Section 3.2. In Section 3.3, I show that many of the phonotactic constraints in Mandarin can be described as agreement constraints, in which a feature is required to be realized over the span of multiple neighboring sounds. This means that in many cases, a sound in Mandarin relies on its neighbor to share its features, rendering the neighbor indispensable in fanqie secret languages, which might lead to the illusion of glide constituency. Finally, in Section 3.4, I discuss the role of faithfulness constraints in mapping ill-formed input to well-formed output. I demonstrate that with a shortage of morphophonological alternations, input-output mapping remains undetermined, which leads to indeterminacy in the ranking of faithfulness constraints.

3.1 Glide-vowel phonotactics

The vowels of Mandarin can be grouped into high vowels and non-high vowels. They are shown in the vowel chart of Table 3.1, where solid lines indicate phoneme boundaries. Dashed lines point to the lack of phonemic contrast between two neighboring vowels in the table. The high vowels display more contrasts along the front-back dimension and roundness, with little variation in surface representation (SR). The non-high vowels, on the other hand, have their feature values determined by their phonological environment. The mid vowel and the low vowel can appear in many forms, depending on what glide precedes them and which nasal follows them. The non-high vowels of Mandarin can be said to form a “vertical” vowel inventory akin to Marshallese, where vowels only contrast in height and lack front-back specification (Bender 1968; J. Choi 1992).

In this section, I discuss two phenomena involving glide-vowel interactions, feature agreement between the glide and the mid vowel in an open syllable (Section 3.1.1), as well as the raising of the low vowel in a closed syllable (Section 3.1.2). Some notes on the high vowels and their relations to the glides are included in Section 3.1.3.

		FRONT [+front, -back]		CENTRAL [-front, -back]	BACK [-front, +back]
HIGH	[+high] [-low]	i ɪ	y	ɨ	u ʊ
MID	[-high] [-low]	e		ə	ɤ o
LOW	[-high] [+low]	ɛ a		ä	ɑ

Table 3.1 Vowel chart of Mandarin

3.1.1 The mid vowel in an open syllable

In an open syllable, the mid vowels [e], [ə], and [o] participate in co-occurrence restrictions with the preceding glide. The front [e] can only appear after a palatal glide [j] or a labiopalatal glide [ɥ]. The round [o] only shows up after a labial glide [w]. The central [ə][1] cannot appear after any glide. The distribution between the glide and the mid vowel is illustrated in Table 3.2. The top row lists Mandarin mid vowels by their specification for the features [front] and [round]. The leftmost

[1] This vowel is often transcribed as a back [ɤ].

column shows the possible glides that the vowels can combine with, including a glide-less condition. If a glide-vowel combination is possible, an example syllable is provided. A shaded cell means the glide-vowel combination is unattested.

	[-high, -low] vowels		
	[+front, -round] e	[-front, -round] ə	[-front, +round] o
C	*[se]	[ʂə35] 'snake'	*[mo]
Cj	[tʰje214] 'iron'	*[ljə]	*[pjo]
Cw	*[xwe]	*[twə]	[kwo55] 'wok'
Cɥ	[ɕɥe214] 'snow'	*[nɥə]	*[tɕɥo]

Table 3.2 Distribution of Mandarin glides and mid vowels in open syllables

I propose a pair of agreement constraints to account for the glide-vowel distribution, They are defined in (1) and (2) using contrastive features. The two constraints are abbreviated as AGREE[FRONT]-GV_{MID} and AGREE[ROUND]-GV_{MID}, which are listed in parentheses. I will refer to the constraints using the abbreviation for the rest of the discussion. All phonotactic constraints proposed follow this practice. They are first defined with contrastive features and then given a shorthand for easy reference. For the two AGREE constraints, the palatal glide /j/ and labiopalatal glide /ɥ/ are specified as [+front] since they are the glide equivalents of the front vowels /i/ and /y/. These constraints are capable of ruling out the ill-formed syllables in the shaded cells of Table 3.2.

(1) AGREE[front]- $\begin{bmatrix}-\text{consonantal}\\-\text{syllabic}\end{bmatrix}\begin{bmatrix}+\text{syllabic}\\-\text{high}\\-\text{low}\end{bmatrix}$ (AGREE[FRONT]-GV_{MID}):

A neighboring glide and [-high, -low] vowel must agree in the feature [front]. Assign a violation when a palatal or labiopalatal glide is followed by a [-front] vowel. Assign a violation when a [+front] vowel is not preceded by a palatal or labiopalatal glide.

(2) AGREE[round]- $\begin{bmatrix}-\text{consonantal}\\-\text{syllabic}\end{bmatrix}\begin{bmatrix}+\text{syllabic}\\-\text{high}\\-\text{low}\end{bmatrix}$ (AGREE[ROUND]-GV_{MID}):

A neighboring glide and [-high, -low] vowel must agree in the feature [round]. Assign a violation when a labial or a labiopalatal glide is followed by a [-round] vowel. Assign a violation when a [+round] vowel is not preceded by a labial or labiopalatal glide.

(3) *ø ≫ AGREE[FRONT]-GV_{MID} ≫ AGREE[ROUND]-GV_{MID}

(4) [ɥe] *can surface despite a mismatch in rounding*

	/ɥə/	*ø	AGREE[FRONT]-GV_{MID}	AGREE[ROUND]-GV_{MID}
a.	[ɥə]		*!	*
☞b.	[ɥe]			*
c.	[ɥø]	*!		
d.	[ɥo]		*!	

Note that AGREE[ROUND]-GV_{MID} is violated frequently in Mandarin. The glide-vowel combination [ɥe] can surface, even though there is a mismatch in the feature [round]. This is possible because AGREE[FRONT]-GV_{MID} takes priority over AGREE[ROUND]-GV_{MID}. Since there is no mid front rounded vowel in the Mandarin vowel inventory to fully agree in features with [ɥ], an unrounded front vowel [e] suffice. This can be captured with the ranking in (3), where the avoidance of [-high, +front, +round] vowels is expressed with a markedness constraint *ø. (4) is an Optimality Theory (OT) tableau that illustrates how the ranking selects [ɥe] as the SR. The tableau assumes that faithfulness constraints regarding the glide are ranked high. Candidates with a change of glide or deletion of the glide are omitted here. Due to the absence of morphophonological alternations in

Mandarin, the ranking of faithfulness constraints remains undetermined. More discussion on the role of faithfulness constraints is provided in Section 3.4. This tableau only aims to show that [ɥe] can surface from the input /ɥə/, because it is more optimal than *[ɥo] and *[ɥø].

The low vowel /a/ has the same surface value for [front] and [round] regardless of its left side neighbor in an open syllable. It can follow a palatal glide /j/ or a labial glide /w/. This is shown in (5). The two agreement constraints AGREE[FRONT]-GV_{MID} and AGREE[ROUND]-GV_{MID} apply only to mid vowels.

(5) *Low vowel in an open syllable*

a.	[ja55]	'duck'	d.	[ɕja55]	'shrimp'
b.	[wa35]	'baby'	e.	[kwa55]	'melon'
c.	[a51]	'Ah'	f.	[ʈʂʰa35]	'tea'

Interestingly, the requirement on a neighboring glide and mid vowel to agree in feature values is also valid when the two sounds are in the opposite sequence. The diphthongs in (6) show that the mid vowel agrees in [front] and [round] with the offglide. The low vowels are not subject to the requirement.

(6) *Diphthongs show mid-vowel-glide feature agreement*

a.	[ej35]	'eh'	e.	[lej35]	'thunder'
b.	[ow55]	'Europe'	f.	[ʈʂow55]	'week'
c.	[aj51]	'love'	g.	[baj35]	'white'
d.	[aw214]	'winter jacket'	h.	[tʰaw35]	'peach'

3.1.2 The low vowel in a closed syllable

The palatal glide has a raising effect on the low vowel in a closed syllable. To understand the interaction between /j/ and /a/, one need to know about Rhyme Harmony. It describes the phenomenon in which the nucleus vowel has to agree in frontness and backness with the nasal coda. Two nasals can occupy the coda position in Mandarin, the denti-alveolar /n/ and the velar

/ŋ/. The denti-alveolar nasal is specified as [+front], whereas the velar nasal is [+back]. Rhyme harmony dictates that when a vowel is followed by the front /n/, it also has to be front. When a vowel precedes a back /ŋ/, it ought to be back as well. Without a nasal coda, the mid and low vowels surface as central vowels, specified as [-front, -back]. The paradigm is shown with examples in (7). Note that in (7f), the central low vowel is transcribed as [ä]. The umlaut is added to highlight its difference from the front [a] in rhyme harmony. Elsewhere in this thesis, where the distinction is not paramount to the interpretation of the data, the umlaut is omitted for the central vowel to simplify the transcription.

(7) *Rhyme harmony*

a.	[pen214]	'notebook'	d.	[lan35]	'blue'
b.	[ʈʂɤŋ55]	'to steam'	e.	[kɑŋ55]	'steel'
c.	[kʰə35]	'shell'	f.	[ʈʂʰä35]	'tea'

Rhyme harmony can be captured using the constraints AGREE[FRONT]-VN and AGREE[BACK]-VN. They are defined in (8) and (9). They generally apply to high vowels as well, in the form of co-occurrence restrictions, as shown in (10). The only exception is the [iŋ] rhyme in (10b), which violates both constraints. One possible explanation to the well-formedness of [iŋ] is that there is no unrounded high back vowel in Mandarin to agree with the velar nasal in [back]. Therefore, if the input is /iŋ/, it gets to surface despite markedness constraint violations. In addition, it is reported that many Mandarin speakers insert an epenthetic vowel between [i] and [ŋ] to avoid the rhyme harmony violation (Carden 2016).

(8) AGREE[front]-[+syllabic][+nasal] (AGREE[FRONT]-VN):

A neighboring vowel and nasal coda must agree in the feature [front]. Assign a violation when a [+front] vowel is followed by a velar nasal, which is [-front]. Assign a violation when a denti-alveolar nasal, which is [+front], is preceded by a [-front] vowel.

(9) AGREE[back]-[+syllabic][+nasal] (AGREE[BACK]-VN):

A neighboring vowel and nasal coda must agree in the feature [back]. Assign a violation when a [+back] vowel is followed by a denti-alveolar nasal, which is [-back]. Assign a violation when a velar nasal, which is [+back], is preceded by a [-back] vowel.

(10) *Rhyme harmony for high vowels*

a.	[pin55]	'guest'	c.	[sʊŋ55]	'pine'
b.	[piŋ55]	'ice'			

When a palatal or labiopalatal glide is added before the vowel-nasal rhyme [an], the low vowel is raised to [ɛ], as shown in (11). The labial glide does not affect the vowel height in any way.

(11) *Low vowel raising*

a.	[pan55]	'class'	d.	[pjɛn55]	'side'	g.	*[pɛn]
b.	[an55]	'safe'	e.	[ɥɛn55]	'wronged'	h.	*[ɛn]
c.	[kan55]	'dry'	f.	[kwan55]	'to close'		

Low vowel raising in closed syllables can be attributed to *EFFORT, which punishes rapid movement of the articulator (see Flemming 2011). To produce /j/ or /ɥ/ with a narrow constriction, the jaw needs to be high. The jaw is also high for the production of /n/. The target between the glide and the coda, /a/, requires a lowered jaw. Therefore, if the speaker wishes to reach the targets /j/, /a/, and /n/, they need to raise the jaw, lower it, and then raise it back up in quick succession, which might prove to be effortful. To ease the articulator movement, the low vowel target is not reached in completion. A halfway vowel [ɛ] is produced instead. For the specific case of low vowel raising, I propose a constraint *EFFORT-PAL-an, which is defined in (12). An accompanying general markedness constraint *ɛ is added to reflect that [ɛ] is not observed anywhere else in the Mandarin lexicon. It can account for unattested syllables like (11g) *[pɛn], in which [ɛ] immediately follows a consonant. It is ranked below *EFFORT-PAL-an, as shown in (13). The workings of the two constraints are shown in the tableaux of (14) and (15).

(12) *EFFORT-$\begin{bmatrix}-\text{syllabic}\\+\text{high}\\+\text{front}\end{bmatrix}\begin{bmatrix}+\text{syllabic}\\+\text{low}\end{bmatrix}\begin{bmatrix}+\text{nasal}\\+\text{front}\end{bmatrix}$ (*EFFORT-PAL-an):

Do not fully realize the low vowel target [a] between a palatal or labiopalatal glide and an alveolar nasal. The slightly higher target [ɛ] is preferred. Assign a violation to the sequences [jan] and [ɥan].

(13) *EFFORT-PAL-an ≫ *ɛ

(14) *Vowel raising after palatal glides*

	/pjan/	*EFFORT-PAL-an	*ɛ
a.	[pjan]	*!	
☞b.	[pjɛn]		*

(15) *Vowel raising does not take place after consonants*

	/pan/	*EFFORT-[PAL]an	*ɛ
☞a.	[pan]		
b.	[pɛn]		*!

The mid vowels do not undergo similar effort minimizing processes when they are situated between a prenuclear glide and a nasal coda. Note that for the mid vowels, both their left-hand side neighbor and right-hand side neighbor have an assimilating effect on them. The glide on the left determines their frontness and roundness, while the nasal on the right controls their frontness and backness. One might wonder which neighbor gets the final say when there is a conflict. The answer is the nasal coda. The mid vowels prioritize rhyme harmony, sometimes at the cost of disagreeing in features with the preceding glide. This is evidenced by the examples in (16). The mid vowel in (16a) [t^hwen35] 'village' shares the feature [+front] with the denti-alveolar nasal, but it does not agree in [round] with the labial glide. (16b) [wɤŋ55] 'old man' also displays glide-vowel disagreement in [round] and [front]. This means that the pair of rhyme harmony constraints AGREE[FRONT]-VN and AGREE[BACK]-VN are ranked above the pair of glide-vowel assimilation constraints AGREE[FRONT]-GV$_{\text{MID}}$ and AGREE[ROUND]-GV$_{\text{MID}}$, as stated in (17). The tableau in (18) illustrates the derivation for [tʰwen35] ‘village’.

(16) *The mid vowel’s features are controlled by the nasal coda*

a [tʰwen35] ‘village’ b. [wɤŋ55] ‘old man’

(17) {AGREE[FRONT]-VN, AGREE[BACK]-VN} ≫ *ø ≫ AGREE[FRONT]-GV$_{\text{MID}}$ ≫ AGREE[ROUND]-GV$_{\text{MID}}$

(18) *Rhyme harmony is prioritized over glide-mid-vowel assimilation*

	/twən/	AGR[FRONT] -VN	AGR[BACK] -VN	*ø	AGR[FRONT] -GV$_{\text{MID}}$	AGR[ROUND] -GV$_{\text{MID}}$
a.	[twən]	*!				*
☞b.	[twen]				*	*
c.	[twøn]			*!	*	
d.	[twon]	*!	*			

3.1.3 High vowels and glides

The high vowels /i/, /y/, and /u/ have a special relationship with the glides /j/, /ɥ/, and /w/. There does not appear to be any substantive phonetic difference between the high vowels and the glides. /j/ is the glide equivalent of /i/, /ɥ/ that of /y/, and /w/ that of /u/. As some scholars have pointed out, there is no underlying distinction between the Mandarin high vowels and their glide counterparts (see C. Cheng 1973; Lin 1989; J. Wang 1993; Kuo 1994). Whether a sound surfaces as a high vowel or a glide depends on its syllable position. If it occupies the nucleus position, then it appears as a high vowel. When it shows up in the medial position between the consonant and the nucleus, then it is a glide. In other words, the high vowels and glides only differ in the feature [syllabic]. This is schematized in (19). (20) shows their distribution with some example words.

(19) *High vowels and glides are underlyingly the same*

	As the nucleus		*Not as the nucleus*	
	[+syllabic]		[-syllabic]	
/i~j/	[i]	*High front unrounded vowel*	[j]	*Palatal glide*
/y~ɥ/	[y]	*High front rounded vowel*	[ɥ]	*Labiopalatal glide*
/u~w/	[u]	*High back rounded vowel*	[w]	*Labial glide*

(20) *High vowels and glides occupy different syllable positions*

a.	[li214]	'mile'	d.	[ljɑŋ35]	'to measure'
b.	[tɕʰy214]	'music'	e.	[tɕʰɥɛn55]	'circle'
c.	[ku55]	'mushroom'	f.	[kwɑŋ55]	'light'

Further doubt on a substantive distinction between the high vowels and the glides was cast by Duanmu (2007). He pointed out that [ji] and [i] do not contrast with each other in Mandarin, unlike *yeast* and *east* in English. The same can be said for the pair [ɥy] and [y], as well as [wu] and [u]. In these cases, the glide is included in the transcription to signify that the onset position is filled.

I concur with the dominant view in the literature that the only distinction between the high vowels and the glides is their feature specification in [syllabic]. As I show in Section 3.2, the high vowels pattern in the same way as glides when it comes to consonant-glide phonotactics.

There are two lax high vowels [ɪ] and [ʊ]. They only appear in closed syllables with a nasal coda, as shown in the examples of (21). Note that it is not the case that all high vowels in closed syllables are lax, as evidenced by the tense vowel in [ɕin55] 'new' and [piŋ55] 'ice'. Therefore, it is not clear why the high vowels in the syllables in (21) are lax. Some articulatory constraints might be at play. The distribution of high vowels in a closed syllable is an interesting puzzle, but it does not have any bearing on the language game experiment data, therefore I will not speculate on the exact reasons, phonological or phonetic, as to why these vowels are lax.

(21) *Lax high vowels in closed syllables*

a.	[tɕʰɥɪn35]	'dress'	c.	[lʊŋ35]	'dragon'
b.	[ɕɥɪn55]	'smoked'	d.	[ɕɥʊŋ35]	'bear'

The central high vowel [ɨ], or the "apical vowel", only appears after denti-alveolar sibilants /ts/, /tsʰ/, /s/ and retroflexes /ʈʂ/, /ʈʂʰ/, /ʂ/, /ɻ/. It is sometimes transcribed as [z], [ʐ] (see Duanmu 2007), [ɿ], [ʅ] (see C. Cheng 1973) in the literature. [ɨ] is homorganic to the preceding sibilant, in that it shares the place of articulation of its consonantal neighbor. In a way, it is the vocalic continuation of the sibilant. Since the apical vowel can follow both denti-alveolar sibilants and retroflexes, its specific articulation can be slightly different depending on the place of the preceding consonant. The vowels in [sɨ55] 'silk' and [ʂɨ55] 'teacher' are not exactly the same, but the slight variation is often ignored in transcription. For instance, C. Cheng (1973) noted the variants [ɿ] after denti-alveolars and [ʅ] after retroflexes, but he decided to simplify the transcription for the two variants with a single symbol "ɨ". I follow Cheng's practice in my description of Mandarin vowels.

3.2 Consonant-glide phonotactics

In Section 3.1, I have introduced the various ways in which a Mandarin prenuclear glide can interact with its vowel neighbor to the right. In this section, I show that the glide also participates in phonotactics with its consonant neighbor to the left. The three glides and their interactions with the consonants are discussed in turn, in the order of the palatal glide /j/ (Section 3.2.1), the labial glide /w/ (Section 3.2.2), and the labiopalatal glide /ɥ/ (Section 3.2.3).

The consonant inventory of Mandarin is presented in Table 3.3. With the exception of the velar nasal /ŋ/, all other consonants can appear in syllable-initial position. The distribution of initial consonants relative to the glides are shown in Table 3.4. This serves as the basis for the following discussion on consonant-glide phonotactics. The top row lists Mandarin consonants by place of articulation. The leftmost column shows the possible glides that the consonants can combine with, including a glide-less condition. If a consonant-glide combination is possible, an example syllable is provided. A shaded cell means the consonant-glide combination is unattested. Note that the high vowels /i/, /u/, and /y/ behave in the same way as the glides when it comes to combinatorial possibility with a preceding consonant. This is illustrated in Table 3.5.

	Bilabial	Labiodental	Denti-alveolar	Retroflex	Palatal	Velar
Plosive	p pʰ		t tʰ			k kʰ
Nasal	m		n			ŋ
Fricative		f	s	ʂ	ɕ	x
Affricate			ts tsʰ	ʈʂ ʈʂʰ	tɕ tɕʰ	
Approximant			l	ɻ		

Table 3.3 Mandarin consonant inventory

	Labials	Denti-alveolars			Retroflexes	Palatals	Velars
		Stops	Sonorants	Sibilants			
	p, pʰ, m, f	t, tʰ	n, l	ts, tsʰ, s	ʈʂ, ʈʂʰ, ʂ, ɻ	tɕ, tɕʰ, ɕ	k, kʰ, x
C	[fɤŋ55] *'wind'*	[tʰow35] *'head'*	[la55] *'to pull'*	[san55] *'three'*	[ʈʂʰɑŋ51] *'to sing'*	*[tɕa]	[xej55] *'black'*
Cj	[mjɛn51] *'flour'*	[tʰjɛn55] *'sky'*	[njow35] *'ox'*	*[sjɑŋ]	*[ʂjaw]	[ɕja55] *'shrimp'*	*[kja]
Cw	[pwo55] *'wave'*	[twen55] *'ton'*	[nwan214] *'warm'*	[tswo51] *'to do'*	[ɻwan214] *'soft'*	*[tɕʰwo]	[kʰwaj51] *'fast'*
Cɥ	*[pʰɥe]	*[tɥɛn]	[lɥe51] *'to omit'*	*[tsɥe]	*[ʈʂɥɛn]	[tɕɥɛn55] *'silk'*	*[xɥe]

Table 3.4 Distribution of Mandarin initial consonants and glides

	Labials	Denti-alveolars			Retroflexes	Palatals	Velars
		Stops	Sonorants	Sibilants			
	p, pʰ, m, f	t, tʰ	n, l	ts, tsʰ, s	ʈʂ, ʈʂʰ, ʂ, ɻ	tɕ, tɕʰ, ɕ	k, kʰ, x
C	[fa214] *'law'*	[tʰə51] *'special'*	[la55] *'to pull'*	[sə51] *'color'*	[ʈʂʰa35] *'tea'*	*[tɕa]	[xə55] *'to drink'*
Ci	[mi214] *'rice'*	[tʰi55] *'to kick'*	[ni35] *'mud'*	*[si]	*[ʂi]	[ɕi55] *'west'*	*[ki]
Cu	[pu51] *'cloth'*	[tu55] *'capital'*	[nu51] *'anger'*	[tsʰu55] *'thick'*	[ɻu51] *'to enter'*	*[ɕu]	[kʰu214] *'bitter'*
Cy	*[my]	*[ty]	[ly51] *'green'*	*[tsy]	*[ʈʂy]	[tɕʰy214] *'music'*	*[xy]

Table 3.5 Distribution of Mandarin initial consonants and high vowels

3.2.1 Co-occurrence restrictions of the palatal glide

As shown in Table 3.4, the palatal glide /j/ can appear after labials, denti-alveolar stops and sonorants, and palatals. It cannot co-occur with denti-alveolar sibilants, retroflexes, and velars, as discussed in Section 2.2. Since the three types of consonants in question, denti-alveolar sibilants, retroflexes, and velars, do not form a natural class, they each have their own markedness constraint, as defined in (22-24). The palatal glide is expressed using the features [+high] and [+front]. This is because it shares its feature specifications and distribution with the high front vowel /i/. The constraints can thus account for the co-occurrence restrictions between these consonants and /i/ as well (see Table 3.5). The labiopalatal glide /ɥ/ and the high front rounded vowel /y/ are also not allowed after denti-alveolar sibilants, retroflexes, and velars. The constraints in (22-24) do not specify the [round] feature of the glide, which means they can account for the distribution of /ɥ/ and /y/ as well.

(22) $*\begin{bmatrix}+\text{coronal}\\-\text{dorsal}\\+\text{anterior}\\+\text{strident}\end{bmatrix}\begin{bmatrix}+\text{high}\\+\text{front}\end{bmatrix}$ (*DENTISIBILANT-PAL):

Assign a violation when [ts], [tsʰ], or [s] is followed by one of [j], [ɥ], [i], or [y].

(23) $*\begin{bmatrix}+\text{coronal}\\-\text{dorsal}\\-\text{anterior}\end{bmatrix}\begin{bmatrix}+\text{high}\\+\text{front}\end{bmatrix}$ (*RETRO-PAL):

Assign a violation when [ʈʂ], [ʈʂʰ], or [ʂ] is followed by one of [j], [ɥ], [i], or [y].

(24) $*\begin{bmatrix}-\text{coronal}\\+\text{dorsal}\end{bmatrix}\begin{bmatrix}+\text{high}\\+\text{front}\end{bmatrix}$ (*VELAR-PAL):

Assign a violation when [k], [kʰ], or [x] is followed by one of [j], [ɥ], [i], or [y].

Among the group of consonants that can precede /j/ and /i/, the palatals are especially interesting, because they obligatorily require the presence of a following [+high] and [+front] segment. This can come in the form of a palatal glide /j/ or a labiopalatal glide /ɥ/. In addition , the high front vowels /i/ and /y/ can also take this position. To capture this requirement, I make use of the markedness constraint defined in (25). It is abbreviated as *PAL[-HIGH]. It punishes sequences of a palatal sibilant immediately followed by a mid or low vowel.

(25) $*\begin{bmatrix}+\text{coronal}\\+\text{dorsal}\end{bmatrix}$ [-high] (*PAL[-HIGH]):

Assign a violation when [tɕ], [tɕʰ], or [ɕ] is followed by a [-high] vowel like [ə] and [a].

The four constraints listed above are all phonotactic constraints that are never violated in the Mandarin lexicon. In other words, they are undominated, making sure that illegal syllables never get to surface in the language. In (26), I use a pseudo-tableau to illustrate the workings of these four constraints. There is no input. Potential surface syllables are listed in the candidate column to show whether they get to surface in the language or not. They are not in competition with each other. Any violation of one of the five phonotactic constraints means they cannot surface.

(26) *Pseudo-tableau that shows how the phonotactics constraints rule out ill-formed syllables*

		*DENTISIBILANT -PAL	*RETROFLEX -PAL	*VELAR -PAL	*PAL[-HIGH]
a.	[sjɑŋ]	*!			
b.	[si]	*!			
c.	[ʂjaw]		*!		
d.	[ʈʂy]		*!		
e.	[kja]			*!	
f.	[xɥe]			*!	
g.	[tɕa]				*!
☞ j.	[san]				
☞ k.	[ʈʂʰɑŋ]				
☞ l.	[xə]				
☞ m.	[ɕja]				
☞ n.	[tɕʰy]				
☞ o.	[mjɛn]				
☞ p.	[tʰi]				

No input-output OT tableau in the conventional sense is provided, since the language does not provide enough data on what kind of repair strategy is preferred in the case of a phonotactically ill-formed input syllable. In other words, when the faithful candidate violates a phonotactic constraint, it is not clear which unfaithful candidate surfaces, because the ranking of faithfulness constraints is ambiguous and opaque. This problem is discussed in more detail in Section 3.3.

3.2.2 Co-occurrence restrictions of the labial glide

The labial glide /w/, as well as the high back vowel /u/, cannot appear after palatal consonants, as shown in Table 3.4 and Table 3.5. The co-occurrence restriction can be stated as the markedness constraint *PAL[+BACK], as defined in (27). The labial glide /w/ is assumed to have the same feature specifications as the high back vowel /u/. Therefore the constraint *PAL [+BACK] is capable of ruling out syllables like *[tɕ^{h}wo] and *[ɕu].

(27) $*\begin{bmatrix}+\text{coronal}\\ +\text{dorsal}\end{bmatrix}\begin{bmatrix}+\text{high}\\ +\text{back}\end{bmatrix}$ (*PAL[+BACK]):

Assign a violation when [tɕ], [tɕ^h], or [ɕ] is followed by a back vowel like [u] or a labial glide [w].

It is very often claimed that the labial glide is also not allowed after labial consonants (see Duanmu 2007). This is evidently not true. Both [pwo55] 'wave' and [pu51] 'cloth' are attested words in Mandarin. However, the proposal of a phonotactic constraint against labial-labial consonant-glide sequences is not without reason. Many syllables of the shape PwVX, in which the P stands for a labial consonant, are unattested in the lexicon. It appears to be the case that there is indeed a general phonotactic constraint against the syllable template PwVX. In fact, [pwo55] 'wave' and [pu51] 'cloth' can be considered the odd ones out. The limited distribution of PwVX syllables can be seen in stark contrast compared to those that begin with a velar consonant, KwVX, as demonstrated in (28). For example, velar-initial syllables such as (28m) [xwa55] 'flower', (28n) [k^{h}wan55] 'wide', and (28o) [kwɑŋ55] 'light' are all perfectly good lexical items in Mandarin, whereas their counterparts that begin with a labial consonant, (28d) *[fwa], (28e) *[mwan], and (28f) *[p^{h}wɑŋ] are all unattested syllables.

(28) *Labial vs. Velar consonants before /w/ and /u/*

a.	[pwo55]	'wave'	j.	[kwo55]	'wok'
b.	*[p^{h}wen]	*Unattested*	k.	[k^{h}wen51]	'sleepy'
c.	*[mwej]	*Unattested*	l.	[xwej55]	'gray'
d.	*[fwa]	*Unattested*	m.	[xwa55]	'flower'
e.	*[mwan]	*Unattested*	n.	[k^{h}wan55]	'wide'
f.	*[p^{h}wɑŋ]	*Unattested*	o.	[kwɑŋ55]	'light'
g.	*[pwaj]	*Unattested*	p.	[k^{h}waj51]	'fast'
h.	[pu51]	'cloth'	q.	[xu214]	'tiger'
i.	*[p^{h}ʊŋ]	*Unattested*	r.	[kʊŋ55]	'male'

It is the attempt to account for the general unattestedness of syllables of the shape PwVX that led phonologists to posit phonotactic constraints like *Pw. With such an analysis, the presence of syllables like [pwo55] 'wave' becomes a problem that needs fixing. Note that it is common practice to transcribe (29a) [pwo55] 'wave' as (29e) [po55] in the literature (see Duanmu 2007), as well as in alphabetical writing systems of Mandarin such as *pinyin* (used in mainland China) and *bopomofo* (used in Taiwan). The latter's very name is based on such a transcription of the first four labial-initial syllables used in reciting Mandarin initials. By transcribing 'wave' as [po55], it takes away the exception to the *Pw constraint.

(29) *Alternative transcription for Pwo syllables*

a.	[pwo55]	'wave'	e.	[po55]	'wave'
b.	[p^hwo55]	'slope'	f.	[p^ho55]	'slope'
c.	[mwo51]	'ink'	g.	[mo51]	'ink'
d.	[fwo35]	'Buddha'	h.	[fo214]	'Buddha'

Since there is an audible glide period between the initial consonant and the vowel of the word for 'wave', I have decided to include the labial glide in my transcription [pwo]. The word is not equivalent in pronunciation to the French *peau* [po] 'skin', which genuinely has no glide between the consonant and the vowel. With the syllables of (29) duly transcribed as Pwo, they serve as counterexamples to the constraint *Pw. Therefore, I conclude that the labial glide is free to co-occur with the labial consonants /p/, /p^h/, /m/, /f/.

I have demonstrated that there is no strict phonotactic constraint against Pw sequences, however, it is still very much the case that such syllables are a rarity in the language. Most syllables that have Pw on the left edge are ill-formed. Here, I only venture a guess at what factors might be at play. The verification of my hypothesis proposed here is left for future studies on coarticulatory phonetics.

The key to the difference between the labials and other consonants might lie in their ability to accommodate coarticulation of the labial glide. /w/ and /u/ are usually heavily coarticulated with the initial consonant in Mandarin. Based on my own observation, the gesture for rounding begins as early as the onset of the initial consonant. Thus (28j) [kwo55] 'wok' can be transcribed as [k^wwo55], if detailed coarticulation information is to be included. Labial consonants in Mandarin appear to be incompatible with the gesture of rounding. Impressionistic observation of Mandarin speakers producing the syllable (28a) [pwo55] 'wave' seems to indicate that the rounding gesture only begins at the offset of the consonant. The articulation of the labial stops appears to require the lips to be unrounded. Therefore, I venture to propose a constraint against rounded labial consonants in (30), with the shorthand *P^w.

(30) $* \begin{bmatrix} +\text{labial} \\ +\text{round} \end{bmatrix}$ (*P^w): Labial consonants cannot have a secondary articulation that is [+round].

Assign a violation for [p^w], [p^{hw}], [m^w], and [f^w].

Why does coarticulation matter to the distribution of the labial glide? My hypothesis is that /w/ requires a neighboring segment to help realize the [+round] feature, in order to maximize its perceptual cue for the feature. Next to a velar consonant, this can be easily done with coarticulation on the consonant itself, as in (31d) [x^wwa55] 'flower'. However, when /w/ follows a labial, the consonant cannot assist in extending the perceptual cue for [+round], due to the constraint *P^w. In this case, the labial glide can look in the other direction, the nucleus vowel, for help. This is what happens in (31a) [pwo55] 'wave', in which the rounded vowel makes sure the rounding is realized over two segments. In (31b) *[fwa], however, the low vowel cannot be rounded. This means that [+round] on the labial glide cannot spread to the left or the right, rendering the syllable ill-formed.

(31) [+round] *requires two segments*

a.	[pwo55]	'wave'	c.	[k^wwo55]	'wok'
b.	*[fwa]	*Unattested*	d.	[x^wwa55]	'flower'

One might wonder why the syllable [pu51] 'cloth' is allowed to surface, given that there is a single rounded vowel? This could be because in an open glide-less syllable CV, the vowel is long in duration (Duanmu 2007). Thus, 'cloth' can be transcribed as [puː51]. The perceptual cue for rounding is salient enough on the long vowel, without requiring the help from its neighbors.

I use the constraint [+ROUND]-2X to express the requirement that rounding needs two segments or a long vowel for perceptual saliency. This is defined in (32). The tableaux in (33-36) demonstrate how it works alongside the phonotactic constraint *P^w.

(32) [+ROUND]-2X: The feature [+round] needs to be realized over two timing slots. It can either be satisfied by two neighboring segments that are both [+round], or by a single long [+round] vowel. Assign a violation when there is only one short [+round] segment in a syllable.

(33) [pwo55] *satisfies* *[+ROUND]-2X *by vowel rounding*

	/pwo/	*P^w	[+ROUND]-2X
☞ a.	[pwo]		
b.	[p^wwo]	*!	

(34) *[fwa] *can never surface*

	/fwa/	*P^w	[+ROUND]-2X
a.	[fwa]		*!
b.	[f^wwa]	*!	

(35) [x^wwa55] *satisfies* *[+ROUND]-2X *by consonant coarticulation*

	/xwa/	*P^w	[+ROUND]-2X
a.	[xwa]		*!
☞ b.	[x^wwa]		

(36) [puː51] *satisfies* *[+ROUND]-2X *using a long vowel*

	/pu/	*$LABIAL^w$	[+ROUND]-2X
a.	[pu]		*!
b.	[p^wu]	*!	
☞ c.	[puː]		

3.2.3 Co-occurrence restrictions of the labiopalatal glide

The labiopalatal glide /ɥ/ has the most restricted distribution when it comes to combinatorial possibilities with the initial consonant. It can follow only two types of consonants, the palatals and the denti-alveolar sonorants. Consonants of all other places of articulation cannot appear next to it. Its co-occurrence restrictions regarding the denti-alveolar sibilants, retroflexes, and velars can be accounted for using the constraints in (22-24). Another group of consonants that cannot precede /ɥ/ or /y/ is the labials, which can be captured with the constraint *Pɥ in (37). This could also be the combined result of *P^w (see Section 3.2.2) and the general markedness of /ɥ/.

(37) $*\,[+\text{labial}]\begin{bmatrix}+\text{high}\\+\text{front}\\+\text{round}\end{bmatrix}$ (*Pɥ): Labial consonants cannot be followed by a labiopalatal glide or a high front rounded vowel. Assign a violation for [pɥ] and [py].

After palatal sibilants, the labiopalatal glide /ɥ/ satisfies *PAL[-HIGH] as much as the plain palatal glide /j/. As discussed in Section 3.2.1, a palatal sibilant cannot be immediately followed by a non-high vowel without a buffering palatal glide, /j/ or /ɥ/. The phonotactic constraint *PAL[-HIGH] in (25) captures the relationship between palatal consonants and the labiopalatal glide, as well as that between palatal consonants and the high front rounded vowel /y/.

The denti-alveolar sonorants /n/ and /l/ are a pair of more mysterious neighbors of the labiopalatal glide. There are altogether seven consonants that have denti-alveolar as their place of articulation in Mandarin, yet it is only the pair of sonorants that allow for the presence of /ɥ/ or /y/. The denti-alveolar stops /t/ and /t^h/, as well as the sibilants /ts/, /ts^h/, /s/, are averse to such a high front rounded neighbor. For the denti-alveolar sibilants, this is expected. The constraint *DENTISIBILANT-PAL in (22) can rule out ill-formed syllables like *[tsɥe] and *[tsy]. The denti-alveolar stops, on the other hand, require a separate phonotactic constraint to account for their aversion to /ɥ/ and /y/. This is stated in (38).

The phonotactic constraint *Tɥ captures a real generalization about unattested syllables in the Mandarin, but it is not clear what motivates it. I provide a possible explanation for the constraint in the following paragraphs. Note that this is speculation on my part. A rigorous study on the acoustics and perception of the labiopalatal glide is needed, in order to verify this account.

(38) $$*\begin{bmatrix}+\text{coronal}\\-\text{dorsal}\\-\text{sonorant}\\-\text{strident}\end{bmatrix}\begin{bmatrix}-\text{consonantal}\\+\text{high}\\+\text{front}\\+\text{round}\end{bmatrix}$$ (*Tɥ):

Assign a violation when [t] or [t^h] is followed by [ɥ] or [y].

I suspect that the denti-alveolar stops, unlike the sonorants, do not provide an amenable environment for the labiopalatal glide to be saliently perceived. For /ɥ/ to be correctly identified by a listener of Mandarin, it should require cues that can contrast it against /j/ and /w/. The contrast between /ɥ/ and /j/ relies on the difference in roundness. The contrast between /ɥ/ and /w/ is one of frontness. This means that a three-way contrast between /j/, /ɥ/, and /w/ is more difficult to maintain than a two-way contrast between /j/ and /w/, since the palatal glide and the labial glide contrast in both [front] and [round]. In the framework of Dispersion Theory (Flemming 1995), one could say that the perceptual distance between /j/ and /w/ is larger than that between /j/ and /ɥ/ or that between /ɥ/ and /w/.

(39) *Three-way contrast after denti-alveolar sonorants*

a.	[lje51]	'to crack'	d.	[ni214]	'you'
b.	[lɥe51]	'brief'	e.	[ny214]	'woman'
c.	[lwo51]	'to fall'	f.	[nu51]	'anger'

(40) *Two-way contrast after denti-alveolar stops*

a.	[tje35]	'dish'	d.	[t^hi55]	'to kick'
b.	*[tɥe]	*Unattested*	e.	*[t^hy]	*Unattested*
c.	[two55]	'many'	f.	[t^hu35]	'picture'

As shown in (39), the denti-alveolar sonorants allow for a three-way contrast in the nature of the following glide. All three high vowels are contrastive after /n/ and /l/ as well. The stops /t/ and /t^h/, however, only permit a two-way contrast between /j/ and /w/ (40a-c), or that between /i/ and /u/ (40d-f).

Steriade (1997) observed that certain phonological environments provide better perceptual cues for contrasts than others. This is the Licensing by Cue hypothesis. In the case of Mandarin, there is a distinct possibility that /n/ and /l/ could be more accommodating to the three-way contrast in glides than /t/ or /tʰ/. Specifically, the sonorants might be able to license a more difficult three-way contrast, because they are longer in duration than the stops. The labiopalatal glide /ɥ/, along with the high front rounded vowel /y/, might be heavily coarticulated with the initial consonant. The sonorants /n/ and /l/ have a lengthy voiced period, which could offer a window for the formants of the co-articulated glide and vowel to be audible to the listener. The same could be said for the palatal sibilants, in which a long frication period should also provide good noisy background for the formant values of glides to show up. The denti-alveolar stops /t/ and /tʰ/, both voiceless and short, perhaps leave no space for coarticulatory cues from the following vocalic segment. Therefore, a two-way contrast might be preferred. The expectation is that the labiopalatal glide would be the one to give way, since it is the most marked sound out of the three, while also being perceptually too close to each of /j/ and /w/. To verify this account, acoustic studies need to be conducted, in order to examine the coarticulation of the labiopalatal glide with its preceding consonant. Furthermore, a perceptual experiment will also help to clarify on the robustness of the three-way glide contrast in different phonological environments. I look forward to seeing future research, which might address these two points.

The denti-alveolar sonorants /n/ and /l/ are the only non-palatal consonants that can be followed by /ɥ/ and /y/. Nevertheless, they are more restricted in what type of syllable they can appear in, compared to the palatal sibilants. Specifically, [nɥ] and [lɥ] are only found in open syllables. In contrast, [ɕɥ] can appear in both open and closed syllables. This is shown in (41).

To capture the restriction on sonorant-initial closed syllables, I use a phonotactic constraint *LɥVN. This is defined in (42).

(41) *Sonorant-ɥ sequences only allowed in open syllables*

a.	[ny214]	‘woman’	f.	[tɕʰy214]	‘music’
b.	[lɥe51]	‘brief’	g.	[ɕɥe214]	‘snow’
c.	*[lɥɛn]	*Unattested*[3]	h.	[tɕɥɛn55]	‘to donate’
d.	*[nɥɪn]	*Unattested*	i.	[tɕʰɥɪn35]	‘skirt’
e.	*[lɥʊŋ]	*Unattested*	j.	[ɕɥʊŋ35]	‘bear’

(42) $\begin{bmatrix}+\text{anterior}\\ -\text{dorsal}\\ +\text{sonorant}\end{bmatrix}\begin{bmatrix}-\text{consonantal}\\ +\text{high}\\ +\text{front}\\ +\text{round}\end{bmatrix}[+\text{syllabic}][+\text{nasal}]$ (*LɥVN):

A closed syllable cannot begin with a denti-alveolar sonorant and labiopalatal glide. Assign a violation to any syllable in the shape of nɥVN or lɥVN.

The lengthy description of this phonotactic constraint reflects its unnaturalness. Since both Lɥ and ɥVN are well-formed sound sequences in the language, it is unclear why the concatenation of the two strings results in an unattested form. Again, I offer only speculation on what might be the motivation for the constraint *LɥVN.

One possible explanation to the unattestedness of syllables in the shape of nɥVN or lɥVN is a potential compression effect induced by the addition of the nasal coda. In the open syllables (41a) [ny214] and (41b) [lɥe51], there is plenty of time for the acoustic cues of [y] and [ɥ] to be realized. In a closed syllable with a nasal coda, there are four articulatory targets for speakers to reach. Even though Mandarin syllables do not have a fixed duration (see F. Wu & Kenstowicz 2015), it might still be possible that the orchestration of the four targets could put some time pressure on each of the targets. Therefore, the duration of the glide might need to be compressed. If there is indeed a shortened time window to realize the cues for frontness and roundness, the labiopalatal glide, which is already quite marked, could lose its perceptual saliency. It might be preferred to have a two-way contrast between /j/ and /w/ in a closed syllable, as shown in (43a-c), compared to a full

[3] [lɥɛn] is listed as a variant pronunciation of [lwan] ‘twin’ in many dictionaries. It is an archaic pronunciation that is only associated with a single lexical item. Therefore, it does not contradict the observation that [lɥ] cannot appear in closed syllables.

three-way contrast between all three glides. This might account for why we do not see any syllables shaped as nɥVN or lɥVN.

The palatal sibilants, on the other hand, are expected to have no issue with co-occurring with /ɥ/ in a closed syllable. This is because the maximum number of glides that can contrast after them is two to begin with, since the sequence *[ɕw] is banned. The two-way contrast between /j/ and /ɥ/ should be easily maintained after palatal sibilants in a closed syllable, as shown in (43d-f). To verify this account, one needs to measure and compare the duration of the glide in open and closed syllables, in order to find out if there is a compression effect in closed syllables. This is a potential area of research for future studies on Mandarin phonotactics.

(43) *Two-way glide contrast in closed syllables*

a.	[ljɛn51]	'practice'	d.	[ɕjɛn55]	'fresh'
b.	*[lɥɛn]	*Unattested*	e.	[ɕɥɛn214]	'to choose'
c.	[lwan51]	'mess'	f.	*[ɕwan]	*Unattested*

3.3 Co-dependence between neighboring sounds

Throughout the discussion on the phonological interaction between the Mandarin prenuclear glide and its consonant and vowel neighbors, a common theme has emerged. A feature is often associated with two neighboring segments in the surface form. The requirement for feature sharing can be understood in perception terms. A perceptual cue requires a lengthy duration, in order to be salient enough to contrast with other syllables.

We observe feature sharing in the interaction between a glide and a neighboring mid vowel, in which the two sounds agree in the features [front] and [round]. In rhyme harmony, a vowel agrees in the features [front] and [back] with its nasal coda neighbor. When it comes to co-occurrence restrictions between the initial consonant and the glide, we have seen that the palatal sibilants require the following vocalic segment to share its features, [+high] and [+front]. The labial glide, in a quest for neighbors to share its rounding feature, is often left empty-handed when it finds itself next to a labial consonant, which is averse to taking on a rounding secondary articulation. The labiopalatal glide, with its three-feature bundle, [+high], [+front], and [+round], has an even more difficult time locating hospitable neighbors. Most initial consonants are not able to share all three features with it.

The Mandarin glide is in a co-dependent relationship with both its neighbors in terms of perceptual cues. At times, it is indispensable to the consonant on the left. On other occasions, it is integral to the surface realization of the vowel to the right. In a fanqie secret language, the co-dependence between the glide and its neighbors is put at risk. The rules of the game dictate that a cut needs to be made near the left edge of the syllable. If the speaker cuts between the consonant and the glide, then whatever feature that is supposed to be shared between the two sounds, might find itself realized over a single segment, thus prone to misperception. On the other hand, if the speaker makes the cut between the glide and the vowel, then the vowel quality might be compromised without the aid of the glide. In certain syllables, the glide is simultaneously relied upon by the consonant and the vowel for cue enhancement. A player of a fanqie secret language would be hard pressed to find a solution that satisfies both the consonant and the vowel. From the lens of phonotactics, we can see that the results of fanqie secret language might be heavily influenced by phonological consequence of slicing near the glide, as opposed to the structural constituency of the glide itself. In the next chapter, with the aid of language game experiments, I demonstrate that phonotactic constraints concerning the glide indeed play a significant role in predicting speaker's treatment of the glide in fanqie secret languages.

3.4 Faithfulness constraints

The discussion so far has focused on markedness constraints and their role in ruling out unattested syllables in the Mandarin lexicon. However, I have not addressed what happens when an unattested form is fed to the grammar as the input. What phonotactic repairs are made to them? How are they mapped to well-formed syllables? Assuming Richness of the Base (ROTB), any possible Mandarin syllable should be generated as an input, which is then mapped to a well-formed output after repairs have been made. The mapping between ill-formed input and well-formed output depends on the ranking of faithfulness constraints. The shortage of morphophonological alternations in Mandarin means that there is no direct evidence of non-identical input-output mapping. Learners of

Mandarin not only have to figure out how to map ill-formed input to well-formed output, but also need to infer the faithfulness constraint ranking that can predict such a mapping.

This is a learning problem for Mandarin speakers, as well as a research puzzle for phonologists. We do not know how faithfulness constraints are ranked in the grammar of Mandarin speakers. In this section, I explore the way in which different rankings of faithfulness constraints can predict different output forms to the same input. In particular, I focus on three scenarios of unattested input form. The first case involves input forms that violate a glide-vowel markedness constraint (Section 3.4.1), while the second contains input forms that incur violations of consonant-glide markedness constraints (Section 3.4.2). Finally, in Section 3.4.3, I consider an unattested input that goes against a markedness constraint governing the entire syllable.

3.4.1 Mapping input forms that violate glide-vowel phonotactics

First, I discuss input forms that violate the glide-vowel markedness constraint AGREE[FRONT]-GV_{MID} and how they can be mapped to well-formed output forms. One such input form is /le/, which can be mapped to either [lə51] 'happiness' or [lje51] 'row', as shown in Table 3.6. Note that this is a slightly modified copy of the upper left corner of Table 3.2, which illustrates the distribution of glide and mid vowels. Gray cells indicate ill-formed syllables.

	[+front, -round] e		[-front, -round] ə
C	/le/	→	[lə51] 'happiness'
	↓		
Cj	[lje51] 'row'		*[ljə]

Table 3.6 Potential input-output mappings for /le/

To derive the output [lə] from /le/, one need to change the vowel quality by violating IDENT-V[FRONT]. To get [lje], on the other hand, a glide is inserted, which violates DEP-G. Whether /le/ is mapped to [lə] or [lje] depends on the relative ranking between the two faithfulness constraints, as demonstrated in the tableaux of (44) and (45). [lə] surfaces if DEP-G outranks IDENT-V[FRONT], whereas [lje] wins if the two constraints are swapped.

(44) /le/ *is mapped to* [lə] *if* DEP-G >> IDENT-V[FRONT]

	/le/	AGREE[FRONT] -GV$_{MID}$	MAX	DEP-G	IDENT-V[FRONT]
a.	[le]	*!			
☞ b.	[lə]				*
c.	[lje]			*!	

(45) /le/ *is mapped to* [lje] *if* IDENT-V[FRONT] >> DEP-G

	/le/	AGREE[FRONT] -GV$_{MID}$	MAX	IDENT-V[FRONT]	DEP-G
a.	[le]	*!			
b.	[lə]			*!	
☞ c.	[lje]				*

	[+front, -round] e		[-front, -round] ə
C	*[le]		[lə51] ‘happiness’
			↑
Cj	[lje51] ‘row’	←	/ljə/

Table 3.7 Potential input-output mappings for /ljə/

Similarly, the ill-formed input /ljə/ could be mapped to either [lə51] 'happiness' or [lje51] 'row', as shown in Table 3.7. The former involves the deletion of the glide, violating MAX, whereas the latter makes changes to vowel quality. To derive [lə] from /ljə/, IDENT-V[FRONT] must be ranked above MAX (see the tableau in (46)). If [lje] is to win, then the ranking must be flipped (see (47)). We do not know which ranking is representative of the grammar of Mandarin speakers, due to a shortage of morphophonological alternations.

(46) /ljə/ *is mapped to* [lə] *if* IDENT-V[FRONT] >> MAX

	/ljə/	AGREE[FRONT] -GV$_{MID}$	DEP-G	IDENT-V[FRONT]	MAX
a.	[ljə]	*!			
☞ b.	[lə]				*
c.	[lje]			*!	

(47) /ljə/ *is mapped to* [lje] *if* MAX >> IDENT-V[FRONT]

	/ljə/	AGREE[FRONT] -GV$_{MID}$	DEP-G	MAX	IDENT-V[FRONT]
a.	[ljə]	*!			
b.	[lə]			*!	
☞ c.	[lje]				*

3.4.2 Mapping input forms that violate consonant-glide phonotactics

A similar mapping problem exists for input forms that violate the consonant-glide constraint *PAL[-HIGH]. For instance, the ill-formed input /ɕa/ in Table 3.8 can be mapped to four distinct well-formed output forms. The first option [ɕja55] 'shrimp' involves glide insertion, which violates DEP-G. Alternatively, one can change the place of articulation of the palatal sibilant, deriving [sa55] 'three people', [ʂa55] 'sand' or [xa55] 'Ha'. The three output candidates each violate a set of IDENT-C constraints, including IDENT-C[CORONAL], IDENT-C[ANTERIOR], IDENT-C[DISTRIBUTED], and IDENT-C[DORSAL]. In tableaux of (48-51), I illustrate the rankings that can derive the four output forms. In (48), the four IDENT-C constraints are ranked above DEP-G, which leads to the glide-insertion candidate [ɕja] to win. In the other three tableaux, DEP-G is ranked above the four IDENT-C constraints. Therefore, consonant feature change is preferred over glide insertion. The exact surface consonant depends on the relative ranking of the four IDENT-C constraints, as shown in (49-51). In Table 3.9, I include the feature chart for the consonants involved, as a guide for understanding the violations to the various IDENT-C constraints.

	Palatals tɕ, tɕʰ, ɕ		Denti-alveolar sibilants ts, tsʰ, s		Retroflex sibilants ʈʂ, ʈʂʰ, ʂ		Velars k, kʰ, x
C	/ɕa/	→	[sa55] *'three people'*	OR	[ʂa55] *'sand'*	OR	[xa55] *'Ha'*
	↓						
Cj	[ɕja55] *'shrimp'*		*[sja]		*[ʂja]		*[xja]

Table 3.8 Potential input-output mappings for /ɕa/

	[coronal]	[anterior]	[distributed]	[dorsal]
Palatals tɕ, tɕʰ, ɕ	+	+	+	+
Denti-alveolar sibilants ts, tsʰ, s	+	+	-	-
Retroflex sibilants ʈʂ, ʈʂʰ, ʂ	+	-	+	-
Velars k, kʰ, x	-	-	-	+

Table 3.9 Feature chart for four types of consonants: palatals, denti-alveolar sibilants, retroflex sibilants, and velars

(48) /ɕa/ *is mapped to* [ɕja] *if* IDENT-C >> DEP-G

	/ɕa/	*PAL [-HIGH]	IDENT-C [CORONAL]	IDENT-C [ANTERIOR]	IDENT-C [DISTRIBUTED]	IDENT-C [DORSAL]	DEP-G
a.	[ɕa]	*!					
☞ b.	[ɕja]						*
c.	[sa]				*!	*	
d.	[ʂa]			*!		*	
e.	[xa]		*!	*	*		

(49) /ɕa/ *is mapped to* [sa] *if* DEP-G >> IDENT-C *and that* IDENT-C[DISTRIBUTED] *and* IDENT-C[DORSAL] *are ranked lowest*

	/ɕa/	*PAL [-HIGH]	DEP-G	IDENT-C [CORONAL]	IDENT-C [ANTERIOR]	IDENT-C [DISTRIBUTED]	IDENT-C [DORSAL]
a.	[ɕa]	*!					
b.	[ɕja]		*!				
☞ c.	[sa]					*	*
d.	[ʂa]				*!		*
e.	[xa]			*!	*	*	

(50) /ɕa/ *is mapped to* [ʂa] *if* DEP-G >> IDENT-C *and that* IDENT-C[ANTERIOR] *and* IDENT-C[DORSAL] *are ranked lowest*

	/ɕa/	*PAL [-HIGH]	DEP-G	IDENT-C [CORONAL]	IDENT-C [DISTRIBUTED]	IDENT-C [ANTERIOR]	IDENT-C [DORSAL]
a.	[ɕa]	*!					
b.	[ɕja]		*!				
c.	[sa]				*!		*
☞ d.	[ʂa]					*	*
e.	[xa]			*!	*	*	

(51) /ɕa/ *is mapped to* [xa] *if* DEP-G >> IDENT-C *and that* IDENT-C[DORSAL] *is ranked highest among the* IDENT-C *constraints*

	/ɕa/	*PAL [-HIGH]	DEP-G	IDENT-C [DORSAL]	IDENT-C [CORONAL]	IDENT-C [DISTRIBUTED]	IDENT-C [ANTERIOR]
a.	[ɕa]	*!					
b.	[ɕja]		*!				
c.	[sa]			*!		*	
d.	[ʂa]			*!			*
☞ e.	[xa]				*	*	*

As shown in (49-51), different relative rankings between the four IDENT-C constraints lead to different mappings for the ill-formed input /ɕa/. This is a concrete illustration of the palatal series debate using OT faithfulness constraints. Phonologists debate on which consonant group the palatals are mapped to in the synchronic grammar of Mandarin. Essentially, their debate can be described as an attempt to identify which faithfulness constraint ranking corresponds to the grammar acquired by Mandarin speakers.

	Palatals tɕ, tɕʰ, ɕ		Retroflexes ʈʂ, ʈʂʰ, ʂ
C	*[ɕa]		[ʂa55] *'sand'*
			↑
Cj	[ɕja55] *'shrimp'*	←	/ʂja/

Table 3.10 Potential input-output mappings for /ʂja/

On the flip side of the palatal mapping problem, is the treatment of ill-formed input forms like /ʂja/. This is shown in Table 3.10. It can be repaired with two alternative strategies. The speaker can either delete the glide, deriving [ʂa55] 'sand', or change the place of articulation of the retroflex sibilant into a palatal one, as in [ɕja] 'shrimp'. To get the former output, IDENT-C needs to be ranked above MAX (see (52)). For the latter output to win, the ranking is reversed, as seen in (53). Yet again, we do not know which ranking is acquired by Mandarin speakers. Other ill-formed input forms /sja/ and /xja/ face the same challenge as /ʂja/. I do not include tableaux for them, in the interest of brevity.

(52) /ʂja/ *is mapped to* [ʂa] *if* IDENT-C >> MAX

	/ʂja/	*RETROFLEX-PAL	DEP-G	IDENT-C [ANTERIOR]	IDENT-C [DORSAL]	MAX
a.	[ʂja]	*!				
☞ b.	[ʂa]					*
c.	[ɕja]			*!	*!	

(53) /ʂja/ *is mapped to* [ɕja] *if* MAX >> IDENT-C

	/ʂja/	*RETROFLEX-PAL	DEP-G	MAX	IDENT-C [ANTERIOR]	IDENT-C [DORSAL]
a.	[ʂja]	*!				
b.	[ʂa]			*!		
☞ c.	[ɕja]				*	*

3.4.3 Mapping input forms that violate syllable shape phonotactics

Finally, I discuss the potential input-output mappings for the input /lɥɛn/, which violates *LɥVN, a phonotactic constraint that is concerned about the syllable shape. In Table 3.11, I list three possible attested output forms that the unattested /lɥɛn/ could be mapped to. It might be repaired to [lwan51] 'messy', with a change in the [front] feature of the glide. Alternatively, it could be mapped to [ljɛn51] 'chain', violating IDENT-G[ROUND]. It is also possible to change into [lan51] 'rotten', incurring a violation of MAX through glide deletion. The constraint rankings of (54-56) illustrate how each of the three input-output mappings can be derived. There is no evidence in the Mandarin lexicon to tease these three rankings apart.

	Denti-alveolar sonorants n, l
C	[lan51] *'rotten'*
	OR
Cj	[ljɛn51] *'chain'*
	OR
Cw	[lwan51] *'messy'*
	↑
Cɥ	/lɥɛn/

Table 3.11 Potential input-output mappings for /lɥɛn/

(54) /lɥɛn/ *is mapped to* [lwan] *if* IDENT-G[FRONT] *and* IDENT-V[LOW] *are ranked the lowest*

	/lɥɛn/	*LɥVN	MAX	IDENT-G [ROUND]	IDENT-G [FRONT]	IDENT-V [LOW]
a.	[lɥɛn]	*!				
☞ b.	[lwan]				*	*
c.	[ljɛn]			*!		
d.	[lan]		*!			*

(55) /lɥɛn/ *is mapped to* [ljɛn] *if* IDENT-G[ROUND] *is ranked the lowest*

	/lɥɛn/	*LɥVN	MAX	IDENT-G [FRONT]	IDENT-V [LOW]	IDENT-G [ROUND]
a.	[lɥɛn]	*!				
b.	[lwan]			*!	*	
☞ c.	[ljɛn]					*
d.	[lan]		*!		*	

(56) /lɥɛn/ *is mapped to* [ljɛn] *if* MAX *and* IDENT-V[LOW] *are ranked the lowest*

	/lɥɛn/	*LɥVN	IDENT-G [FRONT]	IDENT-G [ROUND]	MAX	IDENT-V [LOW]
a.	[lɥɛn]	*!				
b.	[lwan]		*!			*
c.	[ljɛn]			*!		
☞ d.	[lan]				*	*

As illustrated by the three cases of ill-formed input above, different rankings of faithfulness constraints predict different input-output mappings in the grammar of Mandarin. There is not enough evidence to identify the input-output mapping that corresponds to the grammar acquired by Mandarin speakers, due to a shortage of morphophonological alternations in the language. In the next two chapters, I show that Mandarin speakers' behavior during the language game experiments can provide some inkling to the ranking of faithfulness constraints in their grammar.

Chapter 4 Codeword Language Game Experiments

In this chapter, I detail the design of the codeword language game experiments and report their findings. The goal of the experiments is to uncover what Mandarin speakers have learned about the sound system of their language, when there is a shortage of morphophonological alternations. The codeword language game creates an artificial phonological environment in which phonological alternations and phonotactic constraints, should they exist in the speaker's grammar, can be forced to surface. The experiments can not only lead to insights on the phonology of Mandarin, but also provide some answers to a longstanding question in phonological acquisition – what do speakers learn in the absence of evidence for phonological alternation?

The codeword language game is adapted from the fanqie secret languages (see Chao 1934) introduced in Chapter 2. In the experiments, native speakers of Mandarin were invited to disassemble and rearrange Mandarin syllables with the guise of encoding secret messages. The specific fanqie method used for the codeword language game is slightly different from the original game. In *May-ka*, as well as many other fanqie secret language, a Mandarin syllable is split into an initial and a final, in order to form a new disyllabic construction with a template. Here, the input is in the form of a disyllabic Mandarin word. Participants were asked to swap the initials of the two syllables, in order to form a new disyllabic construction, often a nonce word. What they have produced as codeword responses can inform us on various aspects of Mandarin phonology.

The conventional use of data from fanqie secret languages is to identify syllable structure, by inspecting where speakers pronounce the prenuclear glide. More specifically, whether it is part of the onset or the rhyme. However, as Yip (2003) pointed out, the movement of the glide in such language games can be predicted by a combination of markedness constraints and faithfulness constraints in the language. To test this hypothesis, I have included controls for the phonological environment of the glide in the experiment design. They are introduced to identify the role that consonant-glide and glide-vowel phonotactic constraints play in determining the surface position of the glide. To my knowledge, previous experiments based on the fanqie method did not make full use of phonological environment controls. This experiment is a first in this respect.

In addition, the codeword language game has the potential of forcing phonological alternations that are otherwise not seen in the Mandarin lexicon. At times, participants might be confronted with an ill-formed input, which they need to map to a well-formed output. Specifically, it might shed some light on the palatal series debate (see Chapter 2). If the palatal sibilants are indeed mapped to one of the three other series of consonants, then we expect such mappings to show up in the speakers' response.

The rest of the chapter is organized as follows. Section 4.1 reports on the first online experiment, in which test stimuli and participant responses were both in written format. Section 4.2 include results from the second experiment, which was conducted in a sound booth, where the test stimuli were presented in audio and participant responses were collected as verbal tokens. A more general discussion on the speaker variation observed in the experiments and constraint rankings can be found in Chapter 5.

4.1 Writing-based language game experiment

The first codeword language game experiment was conducted online, using PCIbex Farm (Zehr & Schwarz 2018). The stimuli were presented in written Simplified Chinese characters on a screen. Participant responses were collected as typed Chinese characters as well. This is in contrast to the second in-person experiment, which is based on audio stimuli and verbal responses. The focus was placed on the palatal glide /j/ and its behavior in the fanqie codeword language game. The labiopalatal glide /ɥ/ was also investigated. The data collection method in this experiment was approved by MIT's Committee on the Use of Humans as Experimental Subjects (COUHES), in accordance with protocol 0902003098.

4.1.1 Methods

The codeword language game invites Mandarin speakers to disassemble the syllable in an artificial setting that is akin to fanqie secret languages. In the language game, participants are instructed to try a novel way of encoding secret messages. With an input of a disyllabic word in Mandarin, the encoding technique involves swapping the initial consonants of the two syllables. For example, the word (1a) [**k^h**a55 **f**ej55] ‘coffee’ is encoded as (1b) [**f**a55 **k^h**ej55], a nonce word. The rhymes of the two syllables stay in their original positions, while the initials [k^h] and [f] (in bold) have switched places.

(1) *Codeword language game task*

	Original word				*Codeword*	
a.	**k^h**a55 **f**ej55	‘coffee’	→	b.	**f**a55 **k^h**ej55	*Nonce word*

With a disyllabic word that contains a prenuclear glide, whether Mandarin speakers move the glide with the initial consonant into the target syllable or leave it next to the vowel in the source syllable, can tell us much about the phonotactics of the language. In order to elicit spontaneous responses on how speakers treat the prenuclear glide in the language game, no instruction is provided on what to do with the glide. In fact, during the training phase, participants only witness the encoding of glide-less items like [k^ha55 fej55].

The ambiguity on glide movement is built into the experimental design, in order to find out how speakers of Mandarin treat the sound in their own grammar. This also opens up the range of possible codeword responses to a test item. Take the example test item [ta51 ljaw51] ‘star anise’ in (2), it can lead to three distinct types of codeword responses, depending on the speaker’s decision regarding the palatal glide. The speaker might choose to pronounce the glide in its source syllable, as in (2b) [la tjaw]. Only the consonants [t] and [l] are moved. I label this as the “GV response”, since the glide stays close to its vowel neighbor in the response, making sure that the GV sequence in the original word remains as such in the codeword. Tones are omitted for the codeword responses, since they often differ from the original word and are irrelevant to the task. Alternatively, the speaker might move the glide along with the consonant, introducing it to the target syllable. This is seen in the “CG response” of (2c) [lja taw], thus named due to the faithful replication of the original CG sequence in the codeword. A third possible response is for the speaker to pronounce the glide twice, as in (2d) [lja tjaw]. It is named the “GG response” to signify the double copies of G in both its source and target syllables in the codeword.

(2) *Three codeword response types*

	Original word				*Codeword responses*	
a.	ta51 ljaw51 ‘star anise’	→	b.	la tjaw	*GV response*
			c.	lja taw	*CG response*
			d.	lja tjaw	*GG response*

The conventional use of fanqie secret language data is to identify syllable structure. With such an approach, different hypotheses on syllable structure would make different predictions on the movement of the glide in the language game. If the glide is part of the rhyme, the prediction is that it should surface in its source syllable, resulting in the GV response. However, if the glide forms a constituency with the initial consonant, it should be expected to move into the target syllable with it, deriving the CG response. A GG response would be ambiguous between the two theories on glide constituency. The fact that the glide is pronounced next to the vowel in the source syllable could be indicative of its constituency with the final. On the other hand, the glide's movement into the target syllable with the consonant could be taken to demonstrate its constituency with the initial. In short, the GG response would not offer conclusive evidence on the structural position of the glide in a Mandarin syllable.

Yip (2003) claimed that the movement of the glide in fanqie secret languages can be captured with the phonotactics of the language, without reference to syllable structure. I set out to test this hypothesis by controlling the phonological environment of the glide in the experiment, investigating the role of markedness constraints and faithfulness constraints in the grammar acquired by Mandarin speakers. In Chapter 5, I give due consideration to the alternative syllable structure account for glide movement. I show that it fails to account for the speaker variation observed in the codeword responses.

In order to identify the effect of phonotactics on glide movement, I first set out the baseline case in which there can be no interference from such concerns. This is exemplified with the test item [ta51 ljaw51] ‘star anise’ in (2). Here, the denti-alveolar consonants [t] and [l] do not display co-occurrence restrictions with the palatal glide. The low vowel monophthong [a] and diphthong [aw] do not participate in potential alternation sensitive to the glide. This means that the glide's movement in the language game should not be affected by phonotactic concerns, which makes it a good baseline case.

In the tableau of (7) for the baseline case [ta51 ljaw51] 'star anise', I show that the combination of various faithfulness constraints predicts the GV response to be the default response type. The task of the language game is expressed via an anti-correspondence constraint ¬ANCHOR-L, which is defined in (3). It dictates that the left edge of each syllable in codeword must not be identical to the left edge in the original word. It therefore penalizes the faithful candidate (a) [ta ljaw], which does not serve the purpose of encoding secret messages. This constraint is ranked at the top, since violating it means not completing the codeword task at all.

The language game task also requires that as many components of the original word are realized faithfully in the codeword, starting from the right edge of each syllable. This is expressed with the constraint ANCHOR-R. Defined in (4), it demands the right edge of the syllables be kept intact. This is violated by candidate (f) [ljaw ta], which swaps the two syllables wholesale. Segments to the left of the right edge are glued in place via the two IDENT-TRANSITION constraints in (5) and (6). IDENT-TRANSITION-GV demands that the formant transition between the prenuclear glide and the nucleus vowel remain identical in the codeword. Therefore, it penalizes any move to take the glide out of its source syllable, as in the CG response candidate (c) [lja taw]. If a syllable in the original word does not contain any glide-vowel transition, then the addition of a glide also incurs a violation. This takes place in both the CG response candidate (c) and the GG response candidate (d) [lja tjaw]. Similarly, IDENT-TRANSITION-VX governs the identity of transition from the nucleus vowel to the coda, be it a nasal or an offglide. No candidate in (7) violate this constraint. The tableau in (8), however, contains a candidate that violates IDENT-TRANSITION-VX. It shows the response candidates for another baseline case [$ʈʂ^h$wan55 ts^haj51] 'Szechuan cuisine'. Candidate (h) [ts^hwen $ʈʂ^h$aj], where the nucleus vowel has been changed, incurs a violation of IDENT-TRANSITION-VX, as well as IDENT-TRANSITION-GV.

(3) ¬ANCHOR-L: Any element at the left edge of the syllable in the original word does not have a correspondent at the same edge in the codeword.

(4) ANCHOR-R: Any element at the right edge of the syllable in the original word has a correspondent at the same edge in the codeword.

(5) IDENT-TRANSITION-GV: The glide-vowel transition in the original word must be identical to the glide-vowel transition in the codeword.

(6) IDENT-TRANSITION-VX: The vowel-nasal or vowel-glide transition in the original word must be identical to the vowel-nasal or vowel-glide transition in the codeword.

Last but not least, MAX punishes glide deletion, which is found in candidate (e) [la taw] in (7). It also rules out a candidate that replaces [l] with a random consonant [p], as in candidate (g). DEP-G, penalizes glide insertion, as seen in the GG response [lja tjaw] in candidate (d).

Note that these are output-output faithfulness constraints, since the participants are basing their codeword response on the output form of the original test item. Nevertheless, they serve as a good approximation of input-output faithfulness constraints. The output-output faithfulness constraints used in the language game setting can be described as an application of the input-output faithfulness constraints built in the grammar acquired by Mandarin speakers.

(7) *GV response is the default response for the baseline case* [ta51 ljaw51] *'star anise'*

	Test item	[ta ljaw]	¬ANCHOR-L	ANCHOR-R	MAX	DEP-G	ID-TRANS-GV	ID-TRANS-VX
a.		[ta ljaw]	*!*					
☞b.	*GV*	[la tjaw]						
c.	*CG*	[lja taw]					*!*	
d.	*GG*	[lja tjaw]				*!	*	
e.		[la taw]			*!		*	
f.		[ljaw ta]		*!*				
g.		[pa tjaw]			*!			

(8) *GV response is the default response for the baseline case* [ʈʂʰwan55 tsʰaj51] *'Szechuan cuisine'*

	Test item	[ʈʂʰwan tsʰaj]	¬ANCHOR -L	ANCHOR -R	MAX	DEP -G	ID-TRANS -GV	ID-TRANS -VX
a.		[ʈʂʰwan tsʰaj]	*!*					
☞b.	*GV*	[tsʰwan ʈʂʰaj]						
c.	*CG*	[tsʰan ʈʂʰwaj]					*!*	
d.	*GG*	[tsʰwan ʈʂʰwaj]				*!	*	
e.		[tsʰan ʈʂʰaj]			*!		*	
f.		[ʈʂʰaj tsʰwan]		*!*				
g.		[tsʰwan kʰaj]			*!			
h.		[tsʰwen ʈʂʰaj]					*!	*

None of the candidates in (7) or (8) violate any markedness constraint. They are all made of well-formed syllables in Mandarin. Without risk of markedness constraint violation, the GV response always wins, since it violates no faithfulness constraint. Note that none of the faithfulness constraints in the two tableaux is defined in reference to syllable structure, which is in line with Yip's (2003) proposal.

The tableaux in (7) and (8) show that in the absence of any phonotactic concerns, faithfulness constraints predict the GV response to be the optimal candidate in the baseline case. I now turn my attention to cases in which phonotactic constraints are at risk of being violated. Both consonant-glide phonotactics and glide-vowel phonotactics might have a role in predicting the behavior of the glide in the language game (see Yip 2003). This is illustrated with the examples in (9) and (10). The test item (9a) [ta51 tɕʰjaw35] 'big bridge' shows consonant-glide phonotactics. As explained in Chapter 2 and 3, palatal consonants are always followed by a palatal glide in the lexicon of Mandarin. The question posed by this experiment is whether such distribution of two neighboring sounds really translates to a genuine co-occurrence restriction constraint in the phonotactic grammar of Mandarin. If it does, then Mandarin speakers are expected to avoid any surface construction that includes a palatal consonant immediately followed by a non-high vowel. Therefore, in the context of the codeword language game, participants are expected to prefer the

CG and GG responses (9c) [tɕʰja taw] and (9d) [tɕʰja tjaw] over the GV response (9b) [tɕʰa tjaw], which violates the consonant-glide phonotactic constraint.

(9) *Consonant-glide phonotactics influence response type*

	Original word				Codeword responses	
a.	ta51 tɕʰjaw35	'big bridge'	→	b.	*tɕʰa tjaw	*GV response*
				c.	tɕʰja taw	*CG response*
				d.	tɕʰja tjaw	*GG response*

(10) *Glide-vowel phonotactics influence response type*

	Original word				Codeword responses	
a.	ta21 lje51	'to hunt'	→	b.	la tje	*GV response*
				c.	lja *te/tə	*CG response*
				d.	lja tje	*GG response*

Glide-vowel phonotactics might have a similar effect in the codeword response. The test item in (10), [ta lje] 'to hunt', shows glide-vowel agreement in the second syllable. If such agreement is learned as a markedness constraint by Mandarin speakers, then they are expected to disfavor the CG response (10c) *[lja te] in the language game experiment. Such a response would either risk violating the markedness constraint with an ill-formed syllable *[te], or commit a faithfulness constraint violation by changing the vowel quality, as in [tə]. Therefore, participants are expected to prefer the GV or GG response in cases like this.

With a view of phonotactics, the experiment is designed to test two hypotheses. If the consonant-glide co-occurrence restriction constraint is part of the phonotactic grammar that Mandarin speakers have learned, then we expect to observe divergent patterns for palatal-initial test items and non-palatal-initial test items. Specifically, palatal-initial test items would lead to fewer GV responses compared to the non-palatal-initial ones. Secondly, if glide-vowel agreement is internalized as a markedness constraint by Mandarin speakers, then test items that contain a potential alternating vowel ought to lead to fewer CG responses compared to those that do not.

Finally, the codeword language game experiment is well-equipped to check the input-output mapping of the palatal sibilants. As discussed in Chapter 2, the palatals are in complementary distribution with three distinct groups of consonants, the denti-alveolars, the retroflexes, and the velars. There is a longstanding debate on whether the palatals can be mapped to one of the three groups, and if yes, which one. In the experiment, a palatal consonant might sometimes end up in a syllable that has no palatal glide, like the GV response in (9b), [tɕʰa tjaw]. If the palatals are indeed mapped to another group, then we expect to observe a change in consonant place in the codeword response. [tɕʰa] could be modified into [tsʰa], [ʈʂʰa], or [kʰa], as shown in Section 3.4.2. Such a place change is only optimal when DEP-G is ranked above IDENT-C constraints. However, if DEP-G is ranked below IDENT-C instead, we expect to see the glide insertion repair, as in [tɕʰja], which corresponds to the GG response (9d) [tɕʰja tjaw].

The various hypotheses tested in the codeword language game experiment, along with their predictions, are summarized in Table 4.1.

Baseline case hypothesis:
If no phonotactic constraint is at risk of violation, *then GV responses will dominate.*
Phonotactic constraint hypotheses:
If consonant-glide phonotactic constraints are learned, *then palatal-initial test items will elicit fewer GV responses than non-palatal-initial ones.*
If glide-vowel phonotactic constraints are learned, *then test items with an alternating vowel will elicit fewer CG responses than those without.*
Palatals mapping hypothesis:
If palatal consonants are mapped to another group of consonants (DEP-G >> IDENT-C), *then consonant place change will take place in response tokens.*

Table 4.1 Summary of writing-based codeword language game experiment hypotheses and predictions

4.1.2 Participants

16 native speakers of Mandarin participated in the online experiment. They were all born in China. Their age ranged from 21 to 58. Eight participants identified as male, and eight as female. They were randomly assigned to three test groups, each with a different set of experiment materials (see Section 4.1.3 *Materials*). Three participants were assigned to Group A, seven to Group B, and six to Group C. Consent was elicited electronically at the beginning of the online experiment. There was no monetary compensation provided for their participation.

4.1.3 Materials

The word list used for the experiment includes 48 disyllabic items containing a palatal glide /j/ in one of its syllables, 12 items containing a labiopalatal glide /ɥ/, and another 48 filler items with no glide. The labial glide /w/ is not included in the word list. In addition, there are 10 glide-less training items.

Out of concern for the duration of the online experiment, the 108 words are distributed into three test groups. Every participant was randomly assigned to one of the three test groups. Each test group includes 24 glide test items and 24 glide-less test items. Group A and Group B focus on the palatal glide, splitting the 48 /j/ items into two halves. Group C, on the other hand, compares the behavior of the palatal glide with that of the labiopalatal glide. It contains 12 /j/ items and 12 /ɥ/ items. The training items are the same for all three test groups.

In Groups A and B, the /j/ items are controlled for their phonological environment on both sides of the glide. The consonant on the left of the glide is controlled for its place of articulation, palatal vs. non-palatal. The consonants cannot be identical between the source and target syllable. The vowel following the glide is controlled for potential alternation. There is also a third control on the position of the glide, whether it is in the first or second syllable of the disyllabic test item. The three binary parameters, consonant place, source vowel alternation, and glide position, divide the

/j/ items in each test group evenly in half. They combine to create 2 × 3 × 2 = 8 types of test items, as summarized in Table 4.2, along with example test items.

PARAMETER			EXAMPLE TEST ITEM	
Consonant Place	Source Vowel Alternation	Glide Position		
Palatal	Yes	First syllable	[tɕʰjɛn55 la55]	‘to pull’
		Second syllable	[taw51 tɕʰje51]	‘to steal’
	No	First syllable	[ɕjaw21 pan55]	‘kindergarten grade’
		Second syllable	[law21 tɕjow51]	‘maternal uncle’
Non-palatal	Yes	First syllable	[ljɛn35 ta51]	‘UN General Assembly’
		Second syllable	[low35 mjɛn51]	‘floor area’
	No	First syllable	[ljow35 taw51]	‘to keep until’
		Second syllable	[tow51 mjaw35]	‘sprout’

Table 4.2 The three binary parameters for /j/ items in Groups A & B of the writing-based experiment

PARAMETER		EXAMPLE TEST ITEM	
Source Nucleus Vowel	Glide Position		
/ɛ/	First syllable	[ɕɥɛn35 lə]	‘unlikely’
	Second syllable	[tʰə51 tɕʰɥɛn35]	‘privilege’
/e/	First syllable	[ɕɥe55 lan35]	‘Xuelan (name)’
	Second syllable	[nan35 tɕɥe35]	‘baron’
/ʊ/	First syllable	[tɕʰɥʊŋ35 lan214]	‘to read a lot’
	Second syllable	[lan35 tɕʰɥʊŋ35]	‘blue dome (sky)’

Table 4.3 The two parameters for /ɥ/ items in Group C of the writing-based experiment

In Group C, all /ɥ/ items have palatal initial consonants. This is because the labiopalatal glide has a very limited distribution beyond palatal-initial syllables. In addition, vowel alternation is not a predictor for /ɥ/ items. This is because the vowel after /ɥ/ always alternates. Instead, the 12 /ɥ/ items are evenly divided into three groups, depending on the nucleus vowel: /ɛ/, /e/, or /ʊ/ in the source syllable. Glide position is also considered as a parameter. The two parameters combine to create 2 × 3 = 6 types of /ɥ/ items. They are demonstrated with examples in Table 4.3. For the purpose of consistency between the two types of test stimuli, the 12 /j/ items selected for Group C are also all palatal-initial and show vowel alternation.

Since the participants were asked to type their codeword response, the test items were selected on the basis that their corresponding codeword options are made of attested syllables in the Mandarin lexicon that can be written with Chinese characters. For non-palatal-initial items, all three responses ought to be writable, as shown for the example test item [ta ljaw] 'star anise' in Table 4.4. For palatal-initial items, only the CG and GG responses need to correspond to existing Chinese characters, since the GV response is usually unavailable, due to the presence of an unattested syllable. For example, in Table 4.4, the test item [ta tɕʰjaw] 'big bridge' might lead to a GV response that contains the unattested syllable *[tɕʰa]. It is possible that some speakers might repair this with consonant place change, resulting in a syllable that can be typed with a Chinese character.

	NON-PALATAL-INITIAL TEST ITEM			PALATAL-INITIAL TEST ITEM		
Test item	[ta51 ljaw51]	大料	'star anise'	[ta51 tɕʰjaw35]	大桥	'big bridge'
GV response	[la51 tjaw51]	辣掉	*available*	[tɕʰa tjaw51]	? 掉	*unavailable*
CG response	[lja21 taw51]	俩到	*available*	[tɕʰja51 taw51]	恰到	*available*
GG response	[lja21 tjaw51]	俩掉	*available*	[tɕʰja51 tjaw51]	恰掉	*available*

Table 4.4 Availability of codeword responses in the writing-based experiment

The stimuli were presented in written Simplified Chinese on a computer screen, along with audio pronunciation produced by a native Mandarin speaker (author). Audio stimuli were included to avoid ambiguity in character pronunciation, since there are Chinese characters that are mapped to more than one pronunciation. The full list of test stimuli used in the writing-based experiment is included in Appendix I.

4.1.4 Procedure

At the beginning of the codeword language game experiment, the participants were randomly assigned to one of the three test groups, Group A, B, or C. The participants were invited to learn a new method of encoding secret messages in Mandarin Chinese. They were given explicit instruction in Mandarin, written in Simplified Chinese, to swap the consonants (辅音[fu21 jin55]) or initials (声母[ʂɤŋ55 mu214]) of a disyllabic word. The two codeword examples in (11) were displayed on the screen in turn, written in both Simplified Chinese and pinyin. The participants were instructed that the tones in the original word do not have to remain the same in the codeword, as demonstrated by the example item in (11b) [xɑŋ35 ʈʂow55] 'Hangzhou'. They were informed that the codeword did not have to be a real word. No instruction regarding what to do with a prenuclear glide was provided.

(11) *Codeword language game instruction examples*

a.	xaj21 nan35	'Hainan'	→	naj21 xan35	*Tones remain the same*
b.	xɑŋ35 ʈʂow55	'Hangzhou'	→	ʈʂɑŋ21 xow51	*Tones have changed*

After reading the instructions, the participants were given the chance to try their hand on encoding a few words in the training phase. The 10 training items do not contain any prenuclear glide. Each training trial was presented as a forced-choice task. The disyllabic training item was displayed as Simplified Chinese characters on the screen. There was no pinyin, but the item's pronunciation was provided in audio form, which could disambiguate between variant pronunciations for the same character. Two choices for the corresponding codeword were displayed as Simplified Chinese characters on the screen, with no accompanying audio. One was the correct initial-swapping codeword. The other was the incorrect result of swapping entire syllables, which was a common mistake made by volunteers in a pilot experiment. Examples of training items and response options are shown in Table 4.5. The choice could be between two real words (row 1), between two nonce words (row 2), or between a nonce word and a real word (row 3). Both real word and nonce words options were included to reinforce the idea that the codeword could be real or not. After the participant made their selection by clicking on the option button, they were immediately given feedback on whether they had correctly encoded the word.

Training item		Correct response: Initials swapped		Incorrect response: syllables swapped	
[pan51 taw214]	半岛 ‘peninsula’	[tan55 paw214]	担保 ‘guarantee’	[taw51 pan214]	盗版 ‘pirated’
[kan55 tʂʰaj35]	干柴 ‘firewood’	[tʂʰan55 kaj55]	掺该 *nonce word*	[tʂʰaj55 kan214]	拆敢 *nonce word*
[fɑŋ51 ta51]	放大 ‘magnify’	[tɑŋ55 fa214]	当法 *nonce word*	[ta51 fɑŋ55]	大方 ‘generous’

Table 4.5 Training phase item examples in the writing-based experiment

Following the training phase was the experiment phase, broken into two sessions. During the break between sessions, participants were invited to read about a Chinese language fun fact unrelated to phonology. The order of the 48 experiment items (24 glide test items and 24 glide-less filler items) was randomized for each participant. In each trial, the test item appeared on the screen as Simplified Chinese characters, accompanied with pronunciation audio. Participants were then asked to type in their codeword in a text box, before clicking on the “next” button to submit their response. The experiment was self-paced.

There was an exit survey at the end of the experiment. Information on age, gender, and place of birth was collected. Participants were asked to rate the difficulty of the task and comment on which words they had found to be more difficult to encode. They were also asked if they wrote down any notes in pinyin to help them during the experiment. This question was added to check whether the explicit use of an alphabetical writing system could have had an effect on their response or not.

4.1.5 Results and discussion

The responses in Simplified Chinese text were converted into pinyin using the *Talking Chinese to Pinyin/Zhuyin Converter* (Purple Culture 2021), which was then translated into IPA. Some response tokens were deemed uninterpretable, including encoding mistakes such as entire syllable swap, wrong consonant, wrong rhyme, deletion of glide. These tokens were discarded. The remaining interpretable response tokens were labeled as GV, CG, and GG accordingly. The response tokens in Group A and Group B (10 speakers) were pooled, since they include the same type of test items. Results from Group C, which contains /ɥ/ items, are reported separately.

In Table 4.6, the distribution of the three response types is reported as percentages for four types of test items. The four test item types are created from crossing the two phonological environment factors, consonant place and vowel alternation. For every test item type, if a response type enjoys an overwhelming majority of 75% or above, it is highlighted in bold. Responses elicited for Group C, which had six participants, are reported in Table 4.7. This group contains both /j/ items and /ɥ/ items. The response type distribution is reported for the two glides separately, with no further division regarding phonological environment. This is because all test items in this group begin with a palatal consonant and contain an alternating vowel.

TEST ITEM TYPE	ONLINE EXPERIMENT RESULTS		
	GV	CG	GG
C: Palatal-initial V: Alternates	0%	40%	60%
C: Palatal-initial V: No alternation	0%	41%	59%
C: Non-palatal-initial V: Alternates	**76%**	14%	10%
C: Non-palatal-initial V: No alternation	**77%**	17%	6%

Table 4.6 Distribution of response types for /j/ items in Groups A & B of the writing-based experiment, listed by test item type according to consonant place and vowel alternation

Test Item Type		Online Experiment Results		
		GV	CG	GG
G: /ɥ/	C: Palatal-initial V: Alternates	24%	34%	42%
G: /j/	C: Palatal-initial V: Alternates	24%	52%	24%

Table 4.7 Distribution of response types for /j/ items and /ɥ/ items in Group C of the writing-based experiment, listed by test item type according to glide type and vowel alternation

Different response types were preferred depending on the initial consonant. As shown in Table 4.6, when faced with a non-palatal-initial test item, participants overwhelmingly preferred the GV response, as evidenced by its share of response at 76% and 77% percentage. On the other hand, palatal-initial test items have elicited no GV response at all. The response tokens to palatal-initial items are split between the CG response and the GG response. There appears to be a slight preference for the GG response, with a share at around 60%.

In Table 4.7, all three types of responses were elicited for /ɥ/ items, without a clear winner preferred by the participants. Most interestingly, GV responses were observed in Group C, in contrast to Group A and Group B. Upon closer inspection, it is found that they were largely provided by two individual participants. One participant consistently changed the place of articulation of the palatal sibilants in their GV responses, in order to avoid an unattested syllable like *[ɕa]. The palatals were sometimes changed into a retroflex, as in [ʂa], and sometimes to a denti-alveolar, as in [sa]. Another participant resorted to typing in pinyin for the unattested syllable *[ɕa], transcribing it as *xa* (*x* stands for a palatal sibilant).

The baseline case, in which no phonotactic constraints could affect the results, can be found in the non-palatal-initial /j/ items that show no vowel alternation (bottom row, Table 4.6). Recall that it is predicted that GV responses should dominate for the baseline group. This is indeed the case. 77% of the responses elicited for the baseline category are GV responses. Note that the tableau in

(7) predicts the GV response to be the outright winner, since it harmonically bounds all other response candidates. Therefore, the remaining 23% of CG and GG responses appear to be unaccounted for. This might be the result of noise from the language game task. During the experiments, participants are exposed to items that genuinely prompted a CG or GG response due to phonotactic concerns. Their experience with these items might have pushed them to produce CG and GG response tokens, even when they are not called for.

In Table 4.6, the divergent pattern of responses elicited for palatal-initial and non-palatal-initial test items is predicted by the consonant-glide phonotactic hypothesis. The absence of GV response tokens for palatal-initial test items indicates that the co-occurrence restriction between a palatal consonant-glide combination has indeed been learned by Mandarin speakers as a phonotactic constraint. The 10 participants consistently moved the palatal glide with the palatal consonant into the target syllable, deriving the CG or GG response. Such a phonotactic effect is absent in non-palatal-initial test items, since the consonants do not require the presence of the glide. When it comes to glide-vowel phonotactics, however, the picture is less clear-cut. Recall that a glide-vowel agreement constraint predicts there to be fewer CG responses to test items that contain an alternating vowel. The share of CG response tokens reported in Table 4.6 seems to marginally support this hypothesis. For the non-palatal-initial group, test items that contain an alternating vowel lead to 14% of CG response, compared to the slightly higher 17% elicited for those that do not contain an alternating vowel. The difference is even smaller for the palatal-initial group (40% versus 41%). The response type percentage results are inconclusive.

In order to accurately identify the factors predicting response type preferences, I tried to fit a multinomial logit model to the response data in Group A and Group B in RStudio (R Core Team 2023). I used the mblogit() function in the mclogit package (Elff 2022), which is especially designed for multinomial categorical responses with a baseline category. The GV response was chosen as the baseline response. Factors considered included consonant place, vowel alternation, glide position, and pinyin usage (five participants took notes in pinyin and five did not). Significance level was taken at $p = 0.05$. The attempt was unsuccessful, mostly likely due to the absolute zero number of GV response tokens elicited. This resulted in complete separation in the statistical analysis, in which the maximum likelihood estimation could not be found.

Recall that the language game experiment is also designed to provide a phonological environment for palatal sibilant to appear before [-high] vowels, which might prompt a change of consonant place. No such tokens were recorded for Group A or Group B. All palatal sibilants were replicated faithfully in the response tokens. In Group C, however, place change tokens are observed. They were produced by a single speaker, who sometimes change the palatal sibilants into denti-alveolar ones, and sometimes produce a retroflex. There does not appear to be a unique repair strategy for this speaker. In other words, the palatals are mapped to two distinct groups of consonants in their grammar. In faithfulness constraint terms, this speaker has the ranking DEP-G >> IDENT-C, while all other speakers surveyed have the ranking IDENT-C >> DEP-G instead.

As a reality check, I summarize the results of the online experiment as answers to the questions posed by the language game. Table 4.8 is a copy of Table 4.1, which lists the hypotheses and predictions of the experiment. For every hypothesis, I use a checkmark to indicate that its prediction is born out. A question mark means the results are inconclusive, requiring further tweaks to the experiment design.

Baseline case hypothesis:	
	If no phonotactic constraint is at risk of violation,
✓	*then GV responses will dominate.*
Phonotactic constraint hypotheses:	
	If consonant-glide phonotactic constraints are learned,
?	*then palatal-initial test items will elicit fewer GV responses than non-palatal-initial ones.*
	If glide-vowel phonotactic constraints are learned,
?	*then test items with an alternating vowel will elicit fewer CG responses than those without.*
Palatals mapping hypothesis:	
	If palatal consonants are mapped to another group of consonants (DEP-G >> IDENT-C),
?	*then consonant place change will take place in response tokens.*

Table 4.8 Summary of writing-based codeword language game experiment hypotheses, predictions, and results

4.1.6 Remaining problems

The design of the online experiment has three main problems, which are highlighted by the data collected from participants. They can be found in three areas, sample size, control for vowel alternation, and medium of response collection. I discuss them in turn.

Firstly, the sample size for the online experiment is too small. Any individual speaker could inordinately influence the pool of data collected, potentially leading to a skewed tendency that is not representative of Mandarin speakers as a whole. For instance, the GV response was completely avoided by the 10 participants of Group A and Group B, which led to complete separation in the statistical model. Yet the same type of response was favored by two out of the six participants assigned to Group C. It is for this reason that for the second experiment, I recruited more participants, doubling the sample size.

Another issue concerns the control for vowel alternation. Only the vowel in the source syllable was controlled for whether it might alternate or not. The vowel in the target syllable, which does not originally contain the glide, was not controlled. Vowel alternation in the glide-less target syllable can also play a role in determining the codeword response. This is because the introduction of a glide into this target syllable might lead to a markedness constraint violation or force a repair that changes vowel quality. The effects of vowel alternation in the two syllables might also interact with each other, creating confounds. In the second experiment, I added a control for target syllable vowel alternation, and consider it as a factor for response type distribution.

Finally, eliciting codeword responses as text input creates many confounds, even though it brings much convenience to data collection during covid. First and foremost is an attested syllable bias. Because participants could only key in responses that can be transcribed using a Chinese character, they were unable to record responses that contain unattested syllables in the language's lexicon. Therefore, spontaneous response from the participants could be lost in the process. In a pilot experiment, in which the participants were asked to record themselves playing the language game, many instances of unattested and marked syllables like *[ɕjan] were recorded. In the writing-based experiment reported here, it is very likely that when a speaker had the unattested syllable *[ɕjan] in mind, but was asked to transcribe it with a Chinese character, they would reach for the ones that sound most similar to *[ɕjan], namely a character pronounced as [ɕjɛn] (e.g. 先, 贤, 显, 县). Therefore when we see one of these characters, it is hard to know whether the participant had meant [ɕjɛn] or *[ɕjan] by their response. It is also possible for speakers to produce marked forms that are so distant from attested syllables that it would prove difficult for the speakers to find a real Chinese character to approximate their intended pronunciation. One participant in Group C indeed had to resort to typing marked syllables like *[ɕa] in pinyin letters.

It is not simply marked syllables that present a problem to the method of data collection that uses text input, certain well-formed syllables, or accidental gaps, can also slip through the net of available Chinese characters that can be used for transcription. If a participant were to come up with well-formed yet unattested syllable *[pow], they would have no way of transcribing it with a Chinese character, since it has no associated lexical entry. A variant form of accidental syllable gaps is accidental tonal gaps. If a participant wanted to produce a tonal gap syllable *[p^han214] as part of their response, they would have to alter the tonal specification of the syllable to find a Chinese character they could use to transcribe the intended syllable. This makes the task unnecessarily difficult for the participant. The second experiment, which is production-based, eliminates the attested syllable bias by collecting verbal responses instead.

A more serious problem with the written-based experiment is that some participants might bypass their phonological grammar altogether when performing the encoding task. More often than not, the participant would type Chinese characters using some form of romanization scheme, be it pinyin (used in mainland China) or bopomofo (used in Taiwan). This leaves the possibility that some participants would resort to simple string manipulation, without referring to their own phonological grammar. Since the goal of the language game experiments is to find out what Mandarin speakers have learned about their language, it is paramount that the participants actually get to listen to the test stimuli and pronounce their responses out loud. The second experiment, which took place in a sound booth, presented the stimuli in audio format, and collected responses in spoken form as well. This can ensure that participants are actively engaging with their grammar as they play the codeword language game.

4.2 Production-based language game experiment

The second experiment was conducted in-person in a sound booth at the MIT Phonetics Lab, in which responses were collected as audio recordings of the speakers' production. All three glides, /j/, /w/, and /ɥ/ were investigated. The data collection method in this experiment was approved by MIT's Committee on the Use of Humans as Experimental Subjects (COUHES), in accordance with protocol 0902003098.

4.2.1 Methods

The methods used in this experiment largely follow that of the first one, which is described in detail in Section 4.1.1. Improvements were made to address the issues that arose from the first experiment, which are discussed in full in Section 4.1.6. These measures included an increased sample size, an added control for target syllable vowel alternation, and a data collection method based on speaker production. The test stimuli were modified and expanded, in order to test two additional hypotheses concerning phonotactics.

(12) *Glide-vowel phonotactics in target syllable influence response type*

	Original word				*Codeword responses*	
a.	ʂə55 xwa35	'luxury'	→	b.	xə ʂwa	*GV response*
				c.	*xwə/xwo ʂa	*CG response*
				d.	*xwə/xwo ʂwa	*GG response*

The first new hypothesis tested concerns glide-vowel phonotactics. In the first experiment, only potential vowel alternation in the source syllable was considered as a factor for response type. The vowel in the target syllable might also play a role in predicting the codeword response. If it shows feature agreement with a preceding glide, then the introduction of the glide from the source syllable into the target syllable might trigger markedness constraint violations or force a change of vowel quality. This is exemplified by the test item in (12), [ʂə55 xwa35] 'luxury'. Its target syllable [ʂə] contains a vowel that potentially alternates with [o]. The CG response in (12c) and the GG response in (12d) both introduce the labial glide into the target syllable, which means the speaker has to choose between either a marked syllable *[xwə] or an unfaithful syllable [xwo]. Therefore, an alternating vowel in the target syllable is predicted to make GV responses more likely.

An additional hypothesis on consonant-glide phonotactics is tested as well, which involves the labial glide and velar consonants. In the results for /j/ items in the online experiment, we have seen that palatal consonants and non-palatal consonants lead to divergent response patterns. There is a possibility that velar consonants could have a similar effect on the response distribution for /w/ items. This is based on Wan's (1999) study of naturally occurring Mandarin speech error data. She made the parallel observation that the palatal glide moves and copies with palatal consonants, and that the labial glide moves and copies with velar consonants. In contrast, /j/ operates independently when it is preceded by a non-palatal consonant. So does /w/ when it follows a non-velar consonant. Given that divergent behavior of the palatal glide is seen in the language game experiment, it would be reasonable to expect that the labial glide should also lead to two different response patterns sensitive to consonant place in the language game.

However, I argue that it is far more likely that Mandarin speakers will not treat the labial glide differently in velar-initial test items and non-velar-initial test items in the language game. This is because unlike the palatal-palatal sequence, velar-labial combinations are not affected by consonant-glide phonotactics. Palatal consonants never appear without a palatal glide, but velar consonants can very easily appear on its own. The labial glide can also freely combine with many other types of initial consonants as well. Therefore, if consonant-glide phonotactic constraints are the driving force behind the divergent response pattern for /j/ items we observed in the first experiment, then we should expect no such dichotomy for /w/ items in the second. In other words, if Mandarin speakers have learned consonant-glide co-occurrence restrictions as phonotactic constraints, then in the codeword language game, only the behavior of the palatal glide should be sensitive to consonant place, but not that of the labial glide.

The additional phonotactic hypotheses in the production experiment are added to the original set. They are summarized in Table 4.9. New additions are highlighted with an arrow.

Baseline case hypothesis:
If no phonotactic constraint is at risk of violation, *then GV responses will dominate.*
Phonotactics hypotheses:
If consonant-glide phonotactic constraints are learned, *then palatal-initial test items will elicit fewer GV responses than non-palatal-initial ones;* ➢ *while velar-initial test items will not elicit fewer GV responses than non-velar-initial ones.*
If glide-vowel phonotactic constraints are learned, *then test items with an alternating vowel in the source syllable will elicit fewer CG responses than those without;* ➢ *while test items with an alternating vowel in the target syllable will elicit more GV responses than those without.*
Palatals mapping hypothesis:
If palatal consonants are mapped to another group of consonants (DEP-G >> IDENT-C), *then consonant place change will take place in response tokens.*

Table 4.9 Summary of production-based codeword language game experiment hypotheses and predictions

4.2.2 Participants

42 speakers of Mandarin with a diverse linguistic background participated in the in-person experiment. There were 33 native speakers, eight heritage speakers, and one speaker who self-identified as somewhere in between a native and a heritage speaker. 28 speakers were born in mainland China, five in Taiwan, one in Singapore, and eight in the United States. The participants signed consent forms prior to the experiment and were compensated for their time at the rate of $10/half-hour.

Out of the 42 participants, the responses produced by 33 speakers were analyzed. A participant's data was discarded if they misunderstood the codeword language game task, produced a high percentage of uninterpretable responses, or had experienced technical problems during the experiment.

4.2.3 Materials

There are 100 test stimuli and 20 training stimuli. The test stimuli include 64 disyllabic words containing a glide, made up of 24 /j/ items, 24 /w/ items, and 16 /ɥ/ items. There are 36 filler items with no glide. The 20 training items are also glide-less disyllabic words.

The 24 /j/ items are controlled for their consonant place, vowel alternation, and glide position. Unlike the first experiment, there are three levels for the vowel alternation parameter: source syllable vowel alternation, target syllable vowel alternation, and no vowel alternation. They are illustrated with example test items in Table 4.10. Note that only one of the two syllables in a test item can contain an alternating vowel. If a test item shows source syllable vowel alternation, then its target syllable must contain a vowel that does not alternate, and vice versa. The three parameters are crossed to create 2 × 3 × 2 = 12 types of /j/ items. Each type is represented by two test items.

LEVEL	EXAMPLE TEST ITEM	
Source syllable vowel alternation	[tjɛn51 paw51]	'telegraph'
Target syllable vowel alternation	[lan35 tɕʰjaw35]	'blue bridge'
No vowel alternation	[njow35 taw55]	'butcher knife'

Table 4.10 Three levels of the vowel alternation parameter in the production-based experiment

The 24 /w/ items are also evenly split using the same three parameters as the /j/ items, which include consonant place, vowel alternation, and glide position. For /w/ items, consonant place has two levels: velars and non-velars. In the non-velar group, the initial consonant can be denti-alveolar or retroflex. Labial initial consonants are excluded from the experiment materials because they generally do not co-occur with the labial glide (see Section 3.2.2). Note that in the target syllable, the initial consonant is always a denti-alveolar or retroflex. Velar consonants are avoided for the target syllable, in order to avoid a potential confounding velar effect.

The 16 /ɥ/ items are composed of 14 palatal-initial items and two non-palatal-initial items. Due to the very limited distribution of /ɥ/ after non-palatal consonants, such test items are hard to find. Only [lɥe] and [nɥe] are used to construct non-palatal-initial /ɥ/ items. Vowel alternation cannot be used as a parameter, since all vowels alternate after a labiopalatal glide. Glide position remains a binary parameter.

Audio stimuli were created for the 100 test items and 20 training items, as well as 20 codeword example responses for the training items. The stimuli were produced by a female native speaker of Mandarin from Shanghai, China, who had no knowledge of the purpose of the experiment. The recording took place in a sound booth at the MIT Phonetics Lab, using an Audio-Technica AT804 omnidirectional dynamic microphone, with a sampling rate of 44.1kHz, 32 bits. The speaker read off a word list written in Simplified Chinese, except for a few tonal gap syllables in the codeword example responses, which were transcribed in pinyin. The disyllabic test stimuli were produced in isolation. The full list of test stimuli used in the production-based experiment is included in Appendix II.

4.2.4 Procedure

The second experiment took place in a sound booth at the MIT Phonetics Lab. Each participant signed a consent form that permitted the recording of the entire experiment session as a method of data collection. They listened to audio stimuli using a set of Audio-Technica ATH-SR50BT Bluetooth wireless over-ear headphones. The response tokens were recorded via an Audio-Technica AT804 omnidirectional dynamic microphone, with a sampling rate of 44.1kHz, 32 bits.

Like the first experiment, participants were invited to learn a new method for encoding words in Mandarin. But unlike the first experiment, they were not given explicit instruction on how to encode words in the second experiment. There was no mention of terms such as consonant or initial. Instead, they were asked to figure out the encoding method by listening to pairs of original words and codewords on their own. They were also instructed to pay special attention to the pronunciation of the words. The instruction was written in English and displayed on a computer screen. This way, both native and heritage speakers of Mandarin could understand the instruction.

There were two demonstration phases and two training phases prior to the experiment phase. In the first demonstration phase, participants listened to five examples of word encoding. The demonstration items contain no glide. Their codewords retain the same tonal specifications. The participants could listen to the original word and its corresponding codeword by clicking on two buttons on the screen, for as many times as they wished, at their own pace. No written form of the stimuli was displayed.

The demonstration phase was followed by a training phase, in which the participants were invited to try encoding five glide-less words on their own. In each training trial, the participant first listened to the audio stimulus of the word to be encoded. They were asked to repeat the original word before saying out loud what they thought the codeword was. Afterwards, they could check whether their own codeword matched the correct answer or not, by clicking on a button that plays the correct codeword. The process was self-paced. Like the demonstration phase, no written form of the training stimuli was provided. There was a rinse-and-repeat of the demonstration and training phase, with another set of glide-less items. The second demonstration and training phases were added to make sure that participants who had difficulty figuring out the encoding method got another chance to learn it.

The participants then moved on to the experiment phase. They were told that for some of the words they would hear, there was no correct answer, and that they should say whatever came to their mind first. During each experiment trial, the participant listened to the test stimulus, repeated it out loud, and produced their own codeword response. It was self-paced. No feedback was provided. The participants were not allowed to take any notes during the experiment, either on paper or on a digital device. The 100 trials were split into four self-paced sessions of 25. Between sessions, fun facts about the Chinese language were displayed on the computer screen.

The order of the 100 test items in which they appeared was not entirely randomized. The first nine items were predetermined, because they were relatively easier to encode. Nevertheless, the order within this small set was randomized for each participant. In the set, the first three items were all glide-less filler items like the ones participants had heard in the demonstration and training phases. The next six items contained only non-palatal /j/ items and /w/ items. /ɥ/ items were excluded from this set because volunteers who participated in a pilot experiment found them harder to encode. The remaining 91 test items were completely randomized in order.

At the end of the experiment, participants were asked to fill out an exit survey, which collected information on their place of birth, age, and other Chinese dialect or language they speak. Heritage speakers were also asked questions on their level of familiarity with pinyin or bopomofo.

4.2.5 Results and discussion

The participants' responses in audio format were annotated using Praat (Boersma & Weenink 2022). For each experiment trial, I transcribed the stimulus repeat token and the codeword response token in pinyin and IPA. Sometimes participants produced more than one response tokens during a single trial. They either repeated the same codeword response multiple times or made corrections to a previous attempt. For consistency, the final response token that a participant produced was selected for analysis. Uninterpretable response tokens were discarded, which included entire syllable swap, consonant copying, rhyme copying, incomplete response, etc. A summary of uninterpretable response types is listed in Table 4.11, along with examples. If a participant produced more than 20 uninterpretable responses, their entire data was discarded. The data of five participants were discarded for this very reason. There were altogether 1966 interpretable responses tokens, which were labeled as GV, CG, and GG responses.

One objective of the production-based experiment was to reduce the attested syllable bias observed in the writing-based experiment. By asking participants to produce verbal responses, it was expected that some would produce unattested and marked syllables in Mandarin. The prediction was borne out. Some examples are listed in (13), where unattested and marked syllables are in bold. 16% of all interpretable codeword responses produced by participants contain unattested or marked syllables.

UNINTERPRETABLE RESPONSE TYPE	TEST STIMULUS	STIMULUS REPEAT TOKEN	RESPONSE TOKEN
Entire syllable swap	[tow mjaw]	[tow mjaw]	[mjaw tow]
Consonant copying	[tjɛn paw]	[tjɛn paw]	[tjɛn taw]
Rhyme copying	[tjɛn paw]	[tjɛn paw]	[pjaw taw]
Impossible consonant alternation	[tsʰwen ʈʂə]	[tsʰwen ʈʂə]	[tswen kʰə]
Impossible rhyme alternation	[paw ɕjɛn]	[paw ɕjɛn]	[ɕjaw paj]
Glide deletion	[kʰwaj ʈʂʰə]	[kʰwaj ʈʂʰə]	[ʈʂʰaj kʰə]
Glide item misheard as glide-less	[ɕɥɛn tə]	[ʂen tə]	[ten ʂə]
Item misheard as glide-initial	[lɥe nan]	[ɥe nan]	[nɥe jɛn]
Item misheard with identical consonants	[lɥe nan]	[nɥe nan]	[nɥe nan]
Item misheard with identical rhymes	[tɕjaw low]	[tɕjow low]	[ljow tsow]
Item misheard as containing two glides	[tjaw low]	[tjaw ljow]	[law tjow]
Glide misheard as a different glide	[tɕɥʊŋ tan]	[tɕjow tan]	[tjow ʈʂan]
Incomplete response	[nan tɕɥe]	[nan tɕɥe]	[tɕjan]

Table 4.11 Summary of uninterpretable response types excluded from analysis in the production-based experiment

(13) *Response tokens with unattested or marked syllables in Mandarin*

	RESPONSE TOKEN	TEST ITEM	
a.	pjaw **ɕjan**	ɕjaw21 pan55	'kindergarten grade'
b.	njɑŋ **tɕan**	tɕjɑŋ55 nan35	'Jiangnan (place name)'
c.	**te** tɕja	tɕje21 ta35	'solution'
d.	**swɑŋ** ʈʂwej	tsɑŋ55 ʂwej214	'dirty water; slander'
e.	ʈʂʰaj **kʰwə**	kʰwaj51 ʈʂʰə55	'express train'
f.	ɕjɛn **lɥʊŋ**	lan21 ɕɥʊŋ35	'lazy bear'
g.	**lɥɪn** ʈʂan	tɕɥɪn55 lan35	'Junlan (brand name)'
h.	lan **tɥe**	tan21 lɥe51	'courage and resourcefulness'
i.	twan **ɕə**	ɕɥɛn55 tə35	'Xuande (emperor)'

The response type distribution is summarized in Tables 4.12-14, for /j/ items, /w/ items, and /ɥ/ items respectively. In each table, test items are categorized by their consonant place and vowel alternation. For every type of test item, the share of GV, CG, and GG responses are reported as percentages. If a response type has a share of 75% or above, it is highlighted in bold.

/J/ TEST ITEM TYPE	PRODUCTION EXPERIMENT RESULTS		
	GV	CG	GG
C: Palatal-initial V: Source V alternates	12%	11%	**77%**
C: Palatal-initial V: Target V alternates	13%	22%	65%
C: Palatal-initial V: No V alternates	11%	20%	69%
C: Non-palatal-initial V: Source V alternates	**96%**	0%	4%
C: Non-palatal-initial V: Target V alternates	**95%**	3%	2%
C: Non-palatal-initial V: No V alternates	**88%**	2%	9%

Table 4.12 Distribution of response types for /j/ items in the production-based experiment, listed by test item type according to consonant place and vowel alternation

Table 4.12 shows that for /j/ items, participants generally prefer the GG response for palatal-initial /j/ items, while opting for the GV response for non-palatal-initial items. The dominance of GV response for non-palatal-initial items is stronger than the one enjoyed by GG response for palatal-initial items. This is a repeat of the findings of the online experiment, in which participants also demonstrated a strong preference for the GV response in non-palatal-initial items. In fact, the GV response preference, at around 90%, is stronger in the production-based experiment, compared to the 77% seen in the writing-based experiment.

/w/ Test Item Type	Production Experiment Results		
	GV	CG	GG
C: Velar-initial V: Source V alternates	**85%**	1%	14%
C: Velar-initial V: Target V alternates	**97%**	3%	1%
C: Velar-initial V: No V alternates	**79%**	6%	15%
C: Non-velar-initial V: Source V alternates	**88%**	0%	12%
C: Non-velar-initial V: Target V alternates	**98%**	1%	1%
C: Non-velar-initial V: No V alternates	**84%**	6%	9%

Table 4.13 Distribution of response types for /w/ items in the production-based experiment, listed by test item type according to consonant place and vowel alternation

/ɥ/ Test Item Type	Production Experiment Results		
	GV	CG	GG
C: Palatal-initial V: Both V alternate	20%	16%	64%
C: Non-palatal-initial V: Both V alternate	**90%**	5%	5%

Table 4.14 Distribution of response types for /ɥ/ items in the production-based experiment, listed by test item type according to consonant place

For the newly introduced /w/ items, participants overwhelmingly favor the GV response for every test item type, as shown in Table 4.13. There does not appear to be any divergent response patterns for /w/ items beginning with a velar consonant and those beginning with a non-velar consonants. This is in contrast to the role consonant place plays in determining the response for /j/ items.

In the three test groups that involve no potential phonotactic constraint violations, GV responses dominate. These include the non-palatal-initial /j/ items with no vowel alternation, the velar-initial /w/ items with no vowel alternation, and the non-velar-initial /w/ items with no vowel alternation. This confirms the baseline case hypothesis.

In Table 4.14, it can be observed that the response distribution for /ɥ/ items largely replicates that for /j/ items. Participants prefer the GG response for palatal-initial /ɥ/ items, while opting for the GV response for non-palatal-initial /ɥ/ items. The GV response enjoys a bigger majority in the non-palatal-initial group, compared to what the GG response has for the palatal-initial group. This is similar to the results for /j/ items as well.

The general response preference of participants is summarized in Table 4.15, for each type of glide test item. The GV response is the clear winner for non-palatal-initial /j/ items, /w/ items, and non-palatal-initial /ɥ/ items. It is only when the test item contains a palatal consonant and a palatal glide, that participants prefer the GG response instead. This generalization appears to indicate that the GV response is the default codeword response, and that the GG response is only adopted when consonant-glide phonotactics are at play, further confirming the baseline case hypothesis.

GLIDE TYPE	PARTICIPANT RESPONSE PREFERENCE
/j/ items	GV response preferred for non-palatal-initial items
	GG response preferred for palatal-initial items
/w/ items	GV response preferred for all items
/ɥ/ items	GV response preferred for non-palatal-initial items
	GG response preferred for palatal-initial items

Table 4.15 Summary of participant response preference for the production-based experiment, listed by the type of glide in the test item

In order to verify the various effects of phonotactics on response types, a series of statistical analysis was performed. First of all, multinomial logit models were fit to the response data using RStudio (R Core Team 2023). The mblogit() function in the mclogit package (Elff 2022) was used. It is designed to test multinomial categorical responses with a baseline category, which was chosen to be the GV response. The model compares each alternative response to the baseline category. In the case of the codeword language game, the likelihood of CG responses is compared to that of GV responses, to see which factors best predict the outcome. The same is done for the comparison between GG responses and GV responses. For /j/ items and /w/ items, factors considered include consonant place, vowel alternation, and glide position. For /ɥ/ items, only consonant place and glide position are considered. Subject is included as a random intercept. An example mblogit() function is shown in (14). The results are summarized in (15-17), for /j/ items, /w/ items, and /ɥ/ items respectively, in which *p*-values that have reached significance ($p = 0.05$) are highlighted in bold.

(14) *Example multinomial logit model*

```
mb1 <- mblogit(AnswerType ~ ConsonantPlace + VowelAlternation + GlidePosition, random = ~1|Subject, data = j_data)
```

(15) *Summary of multinomial logit model fitted to /j/ item responses*

Equation for CG vs GV:

	Estimate	Std. Error	z value	Pr(>|z|)	
(Intercept)	-4.33537	0.56109	-7.727	**1.1e-14**	***
ConsonantPlacePalatal	5.23241	0.52133	10.037	**< 2e-16**	***
VowelAlternationSourceV	-1.45254	0.49842	-2.914	**0.00356**	**
VowelAlternationTargetV	-0.36344	0.42579	-0.854	0.39334	
GlidePosition1	-0.07749	0.18652	-0.415	0.67781	

Equation for GG vs GV:

	Estimate	Std. Error	z value	Pr(>|z|)	
(Intercept)	-3.1043	0.3938	-7.883	**3.2e-15**	***
ConsonantPlacePalatal	5.6365	0.3587	15.716	**< 2e-16**	***
VowelAlternationSourceV	-0.3734	0.3489	-1.070	0.2845	
VowelAlternationTargetV	-0.7416	0.3555	-2.086	**0.0370**	*
GlidePosition1	0.2712	0.1441	1.881	0.0599	.

(Co-)Variances:

Grouping level: Subject	Estimate		Std. Error	
CG~1	2.4648		1.114	
GG~1	0.6078	1.9417	0.764	NaN

Number of observations: Groups by Subject: 33, individual observations: 749

(16) *Summary of multinomial logit model fitted to /w/ item responses*

Equation for CG vs GV:

	Estimate	Std. Error	z value	Pr(>|z|)	
(Intercept)	-2.92827	0.41805	-7.005	**2.48e-12**	***
ConsonantPlaceVelar	0.40059	0.47140	0.850	0.39544	
VowelAlternationSourceV	-2.84690	1.04166	-2.733	**0.00628**	**
VowelAlternationTargetV	-1.70639	0.58604	-2.912	**0.00359**	**
GlidePosition1	0.09156	0.23522	0.389	0.69710	

Equation for GG vs GV:

	Estimate	Std. Error	z value	Pr(>|z|)	
(Intercept)	-2.471205	0.351320	-7.034	**2.01e-12**	***
ConsonantPlaceVelar	0.424502	0.297049	1.429	0.153	
VowelAlternationSourceV	-0.009829	0.302381	-0.033	0.974	
VowelAlternationTargetV	-3.153334	0.757934	-4.160	**3.18e-05**	***
GlidePosition1	-0.060245	0.147518	-0.408	0.683	

(Co-)Variances:

Grouping level: Subject	Estimate		Std. Error	
CG~1	1.011		0.7619	
GG~1	1.012	1.492	0.9642	1.2175

Number of observations: Groups by Subject: 33, individual observations: 742

(17) *Summary of multinomial logit model fitted to /ɥ / item responses*

Equation for CG vs GV:

	Estimate	Std. Error	z value	Pr(>|z|)	
(Intercept)	-4.4216	0.8422	-5.250	**1.52e-07**	***
ConsonantPlacePalatal	3.9619	0.7920	5.002	**5.67e-07**	***
GlidePosition1	-0.6056	0.1991	-3.042	**0.00235**	**

Equation for GG vs GV:

	Estimate	Std. Error	z value	Pr(>|z|)	
(Intercept)	-3.2882	0.6597	-4.984	**6.22e-07**	***
ConsonantPlacePalatal	4.7626	0.6421	7.417	**1.20e-13**	***
GlidePosition1	-0.2298	0.1401	-1.640	0.101	

(Co-)Variances:

Grouping level: Subject	Estimate		Std. Error	
CG~1	4.017		2.6528	
GG~1	1.728	1.969	1.4548	0.5757

Number of observations: Groups by Subject: 33, individual observations: 475

In (15) and (17), it can be seen that consonant place is a significant predictor for response type elicited for /j/ items and /ɥ/ items. However, it does not reach significance for /w/ items (see (16)). This is in line with the observation based on the response type percentage report in Table 4.12-14.

When it comes to vowel alternation, it has reached significance for both /j/ items and /w/ items. (15) shows that for /j/ items, source syllable vowel alternation makes the CG response less likely than the baseline GV response. In addition, target syllable vowel alternation makes a GG response less likely than the GV response. This is expected if glide-vowel phonotactic constraints are at play. Recall that an alternating vowel in the source syllable makes the glide more likely to stay where it is (GV or GG response), while an alternating vowel in the target syllable makes the glide less likely to move, making GV the best response option. Therefore, both of these factors should favor the GV response. For /w/ items, both source and target syllable vowel alternation make the CG response less likely than the baseline GV response, as shown in (16). For the comparison between GG and GV responses, target syllable vowel alternation makes the former less likely than the latter.

Surprisingly, in (17), glide position returns a significant *p*-value for the comparison between the CG and GV response elicited for /ɥ/ items. A glide in the first syllable of the test item makes it more likely for a participant to produce a GV response, and less likely to produce a CG response. The effect is a marginal one. One possible explanation is that this is a working memory effect. Participants generally spent longer periods pondering over the /ɥ/ items compared to the rest of the experiment stimuli, due to their complexity. In the experiment recordings, there is often a pause between the first and second syllable of the codeword response for /ɥ/ items. It is plausible that these hesitating participants might had a clearer memory that there was a glide when they were producing the first syllable of the codeword, compared to a few seconds later, as they made an attempt to produce the second syllable. Thus, if the glide was in the first syllable, then the participant would be much more likely to produce it in its source syllable, as in a GV response. If the glide was in the second syllable, then the participant would be less likely to produce it there, resulting in a higher number of CG response.

A secondary statistical analysis was performed, in order to directly test the hypotheses concerning phonotactics. Recall that the three predictions are as follows. Palatal-initial consonants are expected to lead to fewer GV responses. Source syllable vowel alternation ought to render fewer CG responses. Target syllable vowel alternation should make the GV response more likely. Therefore, I fit two binomial logit models for each type of glide. One model pits the GV response to the sum of the other two types of response, with the aim of finding out whether palatal consonant place and target syllable vowel alternation are significant predictors of the GV response. The other model compares the likelihood of the CG response with that of the other two responses combined. Its goal is to check if source syllable vowel alternation makes the expected prediction. For /ɥ/ items, since vowel alternation is not considered as a factor, there is no need for a CG response model. I made use of the glmer() function for the mixed-effects binomial logit models. Fixed effects include consonant place, vowel alternation, and glide position. Subject is added as a random intercept. An example glmer() function is provided in (18). Significance level is taken at $p = 0.05$. The results are summarized in (19-23). Factors expected to be significant by the phonotactic constraint predictions for each model are underlined. *p*-values that have reached significance are highlighted in bold.

(18) *Example mixed effect binomial logit model*

```
gm1 <- glmer(GVResponse ~ (1|Subject) + ConsonantPlace + VowelAlternation +
GlidePosition, family = binomial(link = "logit"), data = j_data)
```

(19) *Summary of mixed-effects binomial logit model fitted to /j/ item GV responses*

Random effects:

Groups	Name	Variance	Std.Dev.
Subject	(Intercept)	2.268	1.506

Number of obs: 749, groups: Subject, 33

Fixed effects

	Estimate	Std. Error	z value	Pr(>|z|)	
(Intercept)	2.9470	0.4043	7.290	**3.1e-13**	***
<u>ConsonantPlacePalatal</u>	-6.1657	0.4444	-13.874	**< 2e-16**	***
VowelAlternationSourceV	0.6113	0.3501	1.746	0.0808	.
<u>VowelAlternationTargetV</u>	0.6713	0.3524	1.905	0.0568	.
GlidePosition1	-0.2022	0.1432	-1.412	0.1581	

(20) *Summary of mixed-effects binomial logit model fitted to /j/ item CG responses*

Random effects:

Groups	Name	Variance	Std.Dev.
Subject	(Intercept)	3.325	1.824

Number of obs: 749, groups: Subject, 33

Fixed effects

	Estimate	Std. Error	z value	Pr(>|z|)	
(Intercept)	-5.3819	0.6707	-8.025	**1.02e-15**	***
ConsonantPlacePalatal	3.3339	0.5143	6.482	**9.05e-11**	***
<u>VowelAlternationSourceV</u>	-1.2096	0.4348	-2.782	**0.0054**	**
VowelAlternationTargetV	0.1580	0.3586	0.441	0.6595	
GlidePosition1	-0.2818	0.1612	-1.748	0.0805	.

(21) *Summary of mixed-effects binomial logit model fitted to /w/ item GV responses*

Random effects:

Groups	Name	Variance	Std.Dev.
Subject	(Intercept)	1.807	1.344

Number of obs: 742, groups: Subject, 33

Fixed effects

	Estimate	Std. Error	z value	Pr(>|z|)	
(Intercept)	2.22788	0.35786	6.226	**4.80e-10**	***
ConsonantPlaceVelar	-0.45224	0.27224	-1.661	0.0967	.
VowelAlternationSourceV	0.45632	0.28585	1.596	0.1104	
<u>VowelAlternationTargetV</u>	2.57495	0.49079	5.247	**1.55e-07**	***
GlidePosition1	0.02378	0.13481	0.176	0.8600	

(22) *Summary of mixed-effects binomial logit model fitted to /w/ item CG responses*

Random effects:

Groups	Name	Variance	Std.Dev.
Subject	(Intercept)	1.219	1.104

Number of obs: 742, groups: Subject, 33

Fixed effects

	Estimate	Std. Error	z value	Pr(>|z|)	
(Intercept)	-3.39508	0.54570	-6.221	**4.92e-10**	***
ConsonantPlaceVelar	0.33031	0.47615	0.694	0.48787	
<u>VowelAlternationSourceV</u>	-2.90300	1.04467	-2.779	**0.00545**	**
VowelAlternationTargetV	-1.41093	0.58152	-2.426	**0.01525**	*
GlidePosition1	0.09915	0.23714	0.418	0.67586	

(23) *Summary of mixed-effects binomial logit model fitted to /ɥ/ item GV responses*

Random effects:

Groups	Name	Variance	Std.Dev.
Subject	(Intercept)	2.782	1.668

Number of obs: 475, groups: Subject, 33

Fixed effects

	Estimate	Std. Error	z value	Pr(>|z|)	
(Intercept)	3.1100	0.6080	5.115	**3.13e-07**	***
<u>ConsonantPlacePalatal</u>	-5.0568	0.6196	-8.162	**3.29e-16**	***
GlidePosition1	0.3116	0.1444	2.158	**0.0309**	*

The results in (19-23) largely support the hypotheses concerning phonotactic constraints. For consonant-glide phonotactics, it can be seen in (19) and (23) that palatal consonant place is a significant predictor for GV responses. Specifically, a palatal initial consonant makes a GV response less likely to appear. The same effect is not observed for velars, as evidenced by the report in (21), in which consonant place does not reach significance. The different results for palatal consonants and velar consonants support the hypothesis that only consonant-glide phonotactic constraints can play a role in determining glide movement. The affinity Wan (1999) observed for velar consonants and the labial glide was probably an artefact of the small sample size and lack of phonological control in the naturally occurring speech error data.

When it comes to glide-vowel phonotactics, it can be seen in (20) that source syllable vowel alternation leads to fewer CG responses for /j/ items, verifying the prediction made in Table 4.9. The same is true for /w/ items in (22). As to target syllable vowel alternation, it reaches significance for the /w/ item model in (21), but only marginally so for the /j/ item model in (19). Specifically, an alternating vowel in the target syllable is linked to more GV responses. This is exactly what was predicted in Table 4.9 as well.

There were also attempts to fit binomial logit models containing random slopes to the data. However, these models either failed to converge or received a boundary (singular) fit warning, which indicates that the random effects were too small. In addition, I tried to fit binomial logit models that included the interaction of consonant place and vowel alternation as a factor. None of the interaction factors in these models returned a significant *p*-value.

The results of statistical models presented in (19-23) strongly support the hypotheses that consonant-glide phonotactic constraints, as well as glide-vowel phonotactic constraints, are learned as part of the phonological grammar by Mandarin speakers. Such phonotactic concerns are seen to influence the decision on codeword response type as well. The palatal consonant effect especially helps explain the high share of GG responses elicited for palatal-initial test items.

Last but not least, the language game data can shed light on the mapping of palatal consonants. GV response tokens involving a change in the place of articulation of the initial palatal consonant have been observed for 10% of tokens elicited for palatal-initial test items. 18 participants have produced such a response. In particular, five speakers changed consonant place in their response quite frequently, each producing five or more such tokens. The other 13 speakers who modified the palatals produced fewer than five such tokens, usually only amounting to one or two instances. For the group of frequent consonant place change speakers, a summary of their consonant place change token counts and their choice of initial consonants is provided in Table 4.16, along with speaker origin information collected from the exit survey.

Speaker	Place change tokens	ɕ → s tokens	ɕ → ʂ tokens	ɕ → x tokens	Speaker origin
#15	19	0	19	0	Shanghai, China
#7	17	5	12	0	Heritage (USA)
#10	11	4	7	0	Singapore
#16	6	2	4	0	Heritage (USA)
#34	6	2	0	4	Fujian, China

Table 4.16 Summary of palatal sibilant place change tokens, as well as consonant place preference, listed for the top five speakers who frequently changed the palatals in the production-based experiment

Even though these participants have resorted to changing the place of articulation of palatal consonants, they do not appear to have a unique choice of repair. Take the example of speaker 7, they swapped a palatal for a denti-alveolar five times, for a retroflex 12 times. Their habit of consonant place change does not appear to be based on a rigorous allophonic alternation built in their grammar, but more likely a result of probabilistic faithfulness constraint ranking that maps the same ill-formed input to various output forms. This can be understood in a stochastic phonological grammar like Maximum Entropy (MaxEnt) (Goldwater & Johnson 2003; Hayes & Wilson 2008), in which different weights are assigned to the faithfulness constraints to represent their influence on predicting the likelihood of various possible output forms. For speaker 7, one can say that DEP-G has a lower weight compared to the IDENT-C constraints, leading the palatal consonants to frequently change place. In particular, IDENT-C[DISTRIBUTED] has a higher weight compared to IDENT-C[ANTERIOR], which means a retroflex output is more likely than a denti-alveolar output.

Dialectal or cross-linguistic influence might also have played a role. For example, in Taiwanese, or Southern Min, some velars are cognates of palatals in Mandarin. Speakers familiar with this correspondence might be more at ease swapping a palatal for a velar. Speaker 34, who hails from

Fujian, where Southern Min is spoken, might have produced velar-initial tokens exactly for this reason.

One might also wonder if the specific Mandarin dialect spoken by these participants might genuinely include a deterministic mapping for the palatal consonants. For instance, speaker 15 frequently produced a retroflex to replace the palatal consonant, never opting for another repair consonant. Is it possible that a rigorous palatal-retroflex mapping is part of the phonological grammar of Shanghai Mandarin? The answer appears to be no. There were two other speakers from Shanghai who participated in the experiment, as well as six speakers from the neighboring Jiangsu province. The Mandarin dialect spoken in Jiangsu is heavily influenced by Wu Chinese, very much like Shanghai. Out of these eight Shanghai and Jiangsu speakers, three participants produced a total of four tokens of palatal consonant place change (three to retroflex, one to velar). The remaining five speakers never changed the place of articulation of the palatal sibilant once. Speaker 15's response data is not representative of the Mandarin grammar in the region. Therefore, I conclude that the palatal consonant place change observed in the language game experiment is not indicative of any deterministic mapping for the palatals in the synchronic Mandarin grammar. The preferred repair strategy in the case of an ill-formed input like *[ɕa] appears to be glide insertion, given the popularity of the GG response for palatal-initial items.

The production-based language game experiment is much more successful at answering questions about Mandarin phonology. Table 4.17 summarizes the findings and how they relate to the hypotheses listed in Table 4.8. A checkmark indicates that a prediction is born out. A cross suggests that it is not. Verified hypotheses are highlighted in bold. More discussion on the experiment results, speaker variation, alternative accounts, phonological constraint learning can be found in Chapter 5.

Baseline case hypothesis:	
	If no phonotactic constraint is at risk of violation,
✓	*then GV responses will dominate.*
Phonotactic constraint hypotheses:	
	If ***consonant-glide phonotactic constraints are learned****,*
✓	*then palatal-initial test items will elicit fewer GV responses than non-palatal-initial ones;*
✓	*while velar-initial test items will not elicit fewer GV responses than non-velar-initial ones.*
	If ***glide-vowel phonotactic constraints are learned****,*
✓	*then test items with an alternating vowel in the source syllable will elicit fewer CG responses than those without;*
✓	*while test items with an alternating vowel in the target syllable will elicit more GV responses than those without.*
Palatals mapping hypothesis:	
	If palatal consonants are mapped to another group of consonants (DEP-G >> IDENT-C),
×	*then consonant place change will take place in response tokens.*

Table 4.17 Summary of production-based codeword language game experiment hypotheses, predictions, and results

Chapter 5 General Discussion

In Chapter 4, I have shown that Mandarin phonotactics can account for the behavior of the prenuclear glide in a fanqie language game experiment, without reference to syllable structure. Both markedness constraints and faithfulness constraints are observed to play a role in predicting speaker response type in the production-based language game experiment. The effect of palatal consonant place and vowel alternation on the choice of codeword response shows that Mandarin speakers have internalized surface distribution of consonant-glide and glide-vowel combinations as phonotactic constraints. Finally, the low frequency of responses involving a place of articulation change for the palatal sibilants, taken together with the indeterminacy of direction of change, strongly suggests that there is no consistent mapping between the palatal consonants and other groups of consonants in the synchronic grammar of Mandarin.

For every question the experiment is designed to address, the results provide a satisfactory answer. But what more can we learn from the codeword response corpus produced by Mandarin speakers in the language game? This chapter seeks to delve a bit deeper into the results, in order to search for more answers, as well as puzzling questions, on Mandarin phonology and phonological learning in general. It is organized as follows. In Section 5.1, I examine the speaker variation observed in the experiment data and try to identify its source. In the same section, an alternative interpretation of the language game data is entertained, one that attributes the movement of the glide to syllable structure. I argue that the speaker variation in the experiment cannot be the result from a variation in syllable structure. I also demonstrate that constraint ranking variation can capture the response type distribution well. Section 5.2 takes a closer look at the marked response tokens produced by speakers, in order to find out which phonotactic constraints are better learned by Mandarin speakers. In Section 5.3, I explore how Mandarin speakers establish input-output mappings based on their treatment of the palatal sibilants. I also discuss what insights on phonological learning can be gleaned from the bias and tendency of Mandarin speakers.

5.1 Source of speaker variation

Plenty of speaker variation in codeword response types is observed in both experiments. For any type of test item, there is never a single unique response type that all participants converged on. What is the source of this variation? I consider two possible sources, syllable structure variation and constraint ranking variation. The former follows the conventional interpretation of fanqie data. I show that the distribution of response types strongly supports the latter as the real source of variation. In Section 5.1.1, I explain why syllable structure variation fails to account for the data. Arguments in favor of constraint ranking variation are detailed in Section 5.1.2. Two puzzling response patterns are explored in Section 5.1.3.

5.1.1 Syllable structure variation account

First, I entertain the possibility that the glide's behavior in the codeword language game might be accounted for by its structural constituency in the syllable. This is the conventional approach to fanqie secret language data. Recall that in the syllable structure account, a GV response would be taken as evidence that the glide is part of the rhyme, because it stays with the nucleus vowel in the source syllable. A CG response, on the other hand, would lead to the conclusion that the glide is part of the onset, as it moves with the initial consonant into the target syllable. A GG response would be inconclusive since the glide patterns with both its neighbors.

Following the syllable structure account, the variation in speaker's codeword response in the language game would be analyzed as the result of a variation in syllable structure. In this account, the structural constituency of the prenuclear glide could be subject to interspeaker variation. One speaker might treat it as part of the onset, while another might parse it as part of the rhyme. In addition, there could be intraspeaker variation in syllable structure. The same speaker might parse the three glides differently. Within the grammar of one speaker, a glide's structural constituency might be determined by the phonological environment as well. For instance, one could point to the divergent response patterns for /j/ items in palatal-initial and non-palatal-initial environments and conclude that the palatal glide's structural constituency is sensitive to its preceding environment. Specifically, the majority GG response for palatal-initial /j/ items would be taken to indicate that the palatal glide forms a constituent with the palatal consonant, but it is independent in other types of syllables.

Indeed, it is reasonable to claim that syllable structure shows variation that is sensitive to speaker, type of glide, as well as the phonological environment. However, it would be quite a stretch to say that the same speaker could parse the same glide differently in the same phonological environment. Yet this is exactly what was observed in the production-based experiment. Table 5.1 shows the response pattern five participants produced for non-palatal-initial /j/ items that contain no alternating vowel. These five participants all show response variation for this group. It is unlikely that there is syllable structure variation in the same phonological environment, within the grammar of the same speaker.

SPEAKER	RESPONSE TOKENS		
	GV	CG	GG
#1	3	0	1
#6	3	0	1
#14	2	0	2
#21	3	1	0
#25	3	1	0

Table 5.1 Distribution of response types for non-palatal/j/ items produced by five participants in the production-based experiment

Nevertheless, the variation in response type in Table 5.1 needs to be accounted for. Since there is no consonant-glide phonotactic interaction or glide-vowel interaction within this group of test items, phonotactics could not have played a role in contributing to the variation. I believe the CG response and GG response tokens collected were noise from the task itself, given that they have a very low share of the total response tokens. During the experiment, a participant playing the codeword language game was exposed to many instances of palatal-initial items that genuinely warranted a CG or GG response for phonotactic reasons. Having produced some CG and GG response tokens, the participant might have readjusted their task strategy, believing that these two types of response were just as likely as the GV response. Thus, they might produce CG and GG responses even when there was no need for them.

The syllable structure variation account suffers a fatal blow when it comes to accounting for the target syllable vowel alternation effect, which is observed for /w/ items (see the statistical analysis report in Chapter 4). Participants are shown to produce more GV responses and fewer CG responses when there is an alternating vowel in the target syllable, compared to when there is no vowel alternation at all. Recall that the CG response could be taken as evidence that the glide is part of the onset. Therefore, the syllable structure variation account would conclude that the labial glide is less likely to form a constituent with the consonant if the vowel in the neighboring syllable displays alternation. This is a preposterous statement. It is one thing to say that the parsing of a glide is sensitive to its immediate phonological environment, but quite another to say that it is dictated by the phonological environment of another syllable.

For the above reasons, I conclude that the source of speaker variation in response type cannot possibly be due to a variation in syllable structure. Note that this refutes the claim I made in an earlier report on the language game experiment (B. Fu 2022), in which I attributed response variation to a speaker variation in glide segmentation. The earlier study only included a preliminary analysis of responses to the /j/ items in the writing-based and production-based experiments, without inspecting the responses for the /w/ items. The narrower scope of data analysis might have led to the erroneous claim on syllable structure variation.

5.1.2 Constraint ranking variation account

Alternatively, the speaker variation in response types could be the result of constraint ranking variation. This can be illustrated with the test item [tɕje21 ta35] 'solution', which includes both consonant-glide and glide-vowel phonotactics. Different constraint rankings select different response tokens, as shown in the tableaux of (1-5). Each tableau evaluates various response types as output candidates for the test item [tɕje21 ta35]. Tones are omitted for space. Note that these are output-output tableaux, in which the test item [tɕje21 ta35] is not an input, but an output form that codeword response candidates are compared to. This means that all faithfulness constraints included in these tableaux are output-output constraints. Nevertheless, they are a close approximation of input-output constraints.

I begin the discussion with the tableau in (1). Its constraint ranking predicts the GG response to be the optimal candidate. Candidate (a) is the original test item [tɕje ta]. It violates the constraint ¬ANCHOR-L, failing to encode a secret message. Candidate (b) is the faithful GV response [tje tɕa]. Its second syllable violates the phonotactic constraint *PAL[-HIGH], because the palatal affricate is immediately followed by a low vowel. To fix this, the speaker might change the place of articulation of the palatal consonant, for example, into a retroflex. They would end up with a GV-repair response [tje ʈʂa], as seen in candidate (c). It incurs a violation of the faithfulness constraint IDENT-C[ANTERIOR]. Other types of consonants can be used as a repair too, including the denti-alveolar [ts] and the velar [k]. These candidates behave in the same way as candidate (c). The only difference is that they violate IDENT-C constraints based on other features. I use the retroflex candidate as a representative for the other similar GV-repair candidates. The CG response candidate (d) [te tɕja] avoids a violation of *PAL[-HIGH] by moving the glide with the palatal consonant. However, this leaves the first syllable *[te] ill-formed, violating AGREE[FRONT]-GV_{MID}. Fortunately, there is a simple repair. The speaker could change the vowel quality in the first syllable, producing [tə] instead. Thus, they arrive at the CG-repair response [tə tɕja] in candidate (e). It avoids a violation of the markedness constraint, at the cost of violating the faithfulness constraint IDENT-V[FRONT]. The GG response [tje tɕja] in candidate (f), on the other hand, violates no markedness constraint and makes no change to the consonant or vowel in the original test item. However, it violates DEP-G, which penalizes the insertion of a glide in the output. The ranking in (3) has DEP-G as one of the lowest ranked constraint. It predicts that the GG response is the optimal candidate.

(1) *Constraint ranking that predicts the GG response*

	Test item	[tɕje ta] 'solution'	¬ANCHOR-L	ANCHOR-R	MAX	*PAL[-HI]	ID-C[ANT]	AGR[FR]-GV$_{MID}$	ID-V[FR]	DEP-G	ID-TRANS-GV
a.	*Original*	[tɕje ta]	*!*								
b.	*GV*	[tje tɕa]				*!					
c.	*GV-repair*	[tje ʈʂa]					*!				
d.	*CG*	[te tɕja]						*!			**
e.	*CG-repair*	[tə tɕja]							*!		**
☞ f.	*GG*	[tje tɕja]								*	*
g.	*G deleted*	[te tɕa]			*!	*		*			*
h.	*G deleted*	[tə ʈʂa]			*!		*		*		*
i.	*Syl. swap*	[ta tɕje]		*!*							
j.	*Random C*	[tje la]			*!						

Note that I also include several unlikely response candidates. In candidate (g) [te tɕa], the glide is deleted altogether. It is harmonically bound by the GV response candidate (b), which means it can never surface. Even after some repairing, the glide deletion response of candidate (h) [tə ʈʂa] still can never win, since it is harmonically bound by the GV-repair response candidate (c). Therefore, I conclude that a glide deletion response can never surface in the codeword language game, because it will always be harmonically bound by other candidates. Glide-deletion candidates are not considered for the rest of the discussion. Candidate (i) [ta tɕje] swaps syllables wholesale, which goes against the word encoding task. It violates ANCHOR-R. Finally, in candidate (j) [tje la], a random consonant is introduced, incurring a violation of MAX. To produce one of these responses means the speaker has not completed the language game task. Since the following discussion compares different response candidates that fulfill the task when phonotactic constraints are at risk, I exclude these non-completion candidates from consideration. The original word candidate (a) is also treated as a non-completion candidate. In addition, constraints that are only violated by these non-completion candidates are omitted from the following tableaux. The list includes ¬ANCHOR-L, ANCHOR-R, and MAX. The GV, CG, and GG responses, no matter their shape, all satisfy this set of constraints.

The tableau in (1) shows that with a low-ranked DEP-G, the GG response wins over other codeword candidates. In (2-5), I consider constraint rankings that rank DEP-G at the top, which means that the GG response is eliminated early on in the race. For example, (2) has a minimally different ranking compared to (1). DEP-G is moved to the top of the ranking, while all other constraints remain at the same position. It predicts that the CG-repair candidate (d) [tə tɕja] will surface. Here, the pair of constraints governing the consonant, *PAL[-HIGH] and IDENT-C[ANTERIOR] are both ranked high, which means the sequence [tɕj] is not to be altered in the response. Therefore, only CG response candidates have a chance to win. Between the two constraints on vowels, AGREE[FRONT]-GV$_{\text{MID}}$ is ranked above IDENT-V[FRONT]. This leads the vowel to change, resulting in the CG-repair response [tə tɕja].

The constraint ranking in (3) predicts the faithful CG response [te tɕja] to surface. It is minimally different from the ranking in (2). The pair of vowel constraints, IDENT-V[FRONT] and AGREE[FRONT]-GV$_{\text{MID}}$, have swapped places. A speaker with this grammar prioritize preserving vowel quality, at the expense of violating glide-vowel feature agreement.

(2) *Constraint ranking that predicts the CG-repair response*

	Test item	[tɕje ta] ‘solution’	DEP-G	*PAL[-HI]	ID-C [ANT]	AGR[FR]-GV$_{\text{MID}}$	ID-V [FR]	ID-TRANS -GV
a.	*GV*	[tje tɕa]		*!				
b.	*GV-repair*	[tje ʈʂa]			*!			
c.	*CG*	[te tɕja]				*!		**
☞d.	*CG-repair*	[tə tɕja]					*	**
e.	*GG*	[tje tɕja]	*!					*

(3) *Constraint ranking that predicts the CG response*

	Test item	[tɕje ta] ‘solution’	DEP-G	*PAL[-HI]	ID-C [ANT]	ID-V [FR]	AGR[FR]-GV_{MID}	ID-TRANS -GV
a.	*GV*	[tje tɕa]		*!				
b.	*GV-repair*	[tje t̪sa]			*!			
☞c.	*CG*	[te tɕja]					*	**
d.	*CG-repair*	[tə tɕja]				*!		**
e.	*GG*	[tje tɕja]	*!					*

(4) *Constraint ranking that predicts the GV-repair response*

	Test item	[tɕje ta] ‘solution’	DEP-G	AGR[FR]-GV_{MID}	ID-V [FR]	*PAL[-HI]	ID-C [ANT]	ID-TRANS -GV
a.	*GV*	[tje tɕa]				*!		
☞b.	*GV-repair*	[tje t̪sa]					*	
c.	*CG*	[te tɕja]		*!				**
d.	*CG-repair*	[tə tɕja]			*!			**
e.	*GG*	[tje tɕja]	*!					*

(5) *Constraint ranking that predicts the GV response*

	Test item	[tɕje ta] ‘solution’	DEP-G	AGR[FR]-GV_{MID}	ID-V [FR]	ID-C [ANT]	*PAL[-HI]	ID-TRANS -GV
☞a.	*GV*	[tje tɕa]					*	
b.	*GV-repair*	[tje t̪sa]				*!		
c.	*CG*	[te tɕja]		*!				**
d.	*CG-repair*	[tə tɕja]			*!			**
e.	*GG*	[tje tɕja]	*!					*

In (4-5), I show the constraint rankings that can predict GV responses. Here, the two vowel constraints are ranked higher than the consonant constraints. This means that the glide-vowel sequence [je] must remain the same in the codeword response, which leaves GV response candidates as the only viable options. In (4), the markedness constraint *PAL[-HIGH] is ranked above the faithfulness constraint IDENT-C[ANTERIOR]. It favors the GV-repair response candidate (b) [tje ʈʂa]. (5) has the opposite relative ranking between the two consonant constraints, which leads the marked GV response candidate (a) [tje tɕa] to surface.

Throughout (1-5), I have shown that five distinct constraint rankings can generate five types of codeword responses to the same test item [tɕje21 ta35] 'solution'. This is true for any palatal-initial test item, as shown by the unranked tableaux of (6-8). Each tableau evaluates the candidate responses for one type of palatal-initial /j/ items categorized by vowel alternation. (6) is a repeat tableau for the test item [tɕje ta] 'solution', which displays source syllable vowel alternation. (7) has an example test item that involves target syllable vowel alternation, [tə35 tɕja55] 'Degas'. The test item in (8), [tɕjaw21 low35] 'turret' shows no vowel alternation. One thing that unite all three unranked tableaux, no matter the nature of the vowel, is that every response candidate has violated some constraint. No candidate is harmonically bound by another. Therefore, different rankings of the constraints listed here will predict different optimal response candidates.

(6) *Palatal-initial test item with source syllable vowel alternation*

	Test item	[tɕje ta] 'solution'	*PAL[-HI]	ID-C [ANT]	AGR[FR]-GV$_{MID}$	ID-V [FR]	ID-TRANS -GV	DEP-G
a.	*GV*	[tje tɕa]	*					
b.	*GV-repair*	[tje ʈʂa]		*				
c.	*CG*	[te tɕja]			*		**	
d.	*CG-repair*	[tə tɕja]				*	**	
e.	*GG*	[tje tɕja]					*	*

(7) *Palatal-initial test item with target syllable vowel alternation*

	Test item	[tə tɕja] 'Degas'[4]	*PAL[-HI]	ID-C [ANT]	AGR[FR]-GV$_{MID}$	ID-V [FR]	ID-TRANS -GV	DEP-G
a.	*GV*	[tɕə tja]	*					
b.	*GV-repair*	[tʂə tja]		*				
c.	*CG*	[tɕjə ta]			*		**	
d.	*CG-repair*	[tɕje ta]				*	**	
e.	*GG*	[tɕjə tja]			*		*	*
f.	*GG-repair*	[tɕje tja]				*	*	*

(8) *Palatal-initial test item with no vowel alternation*

	Test item	[tɕjaw low] 'turret'	*PAL[-HI]	ID-C [ANT]	AGR[FR] -GV$_{MID}$	ID-V [FR]	ID-TRANS -GV	DEP-G
a.	*GV*	[ljaw tɕow]	*					
b.	*GV-repair*	[ljaw tʂow]		*				
c.	*CG*	[law tɕjow]					**	
d.	*GG*	[ljaw tɕjow]					*	*

However, it is not always the case that different rankings lead to distinct responses. For non-palatal-initial test items, no matter how the constraints are ranked, the output is always predicted to be the GV response, due to harmonic bounding between candidates. This can be seen in the tableaux of (9-11), each evaluating a non-palatal-initial example test item with a different configuration of vowel alternation. The constraints in these tableaux are unranked. Nevertheless, it is not hard to see that no matter the constraint ranking, the GV response will always win, because it incurs no constraint violation. All other response candidates are harmonically bound by it.

[4] [tə tɕja] 'Degas' is not an actual test item used in the experiments. It is listed as an example here to keep the discussion on various test item types simple, since they all involve the markedness constraint AGR[FR]-GV$_{MID}$.

(9) *Non-palatal-initial test item with source syllable vowel alternation*

	Test item	[ta lje] 'to hunt'	*PAL[-HI]	ID-C [ANT]	AGR[FR]-GV	ID-V [FR]	ID-TRANS -GV	DEP-G
☞a.	*GV*	[la tje]						
b.	*CG*	[lja te]			*		**	
c.	*CG-repair*	[lja tə]				*	**	
d.	*GG*	[lja tje]					*	*

(10) *Non-palatal-initial test item with target syllable vowel alternation*

	Test item	[ljow tə] 'to retain'	*PAL[-HI]	ID-C [ANT]	AGR[FR]-GV_{MID}	ID-V [FR]	ID-TRANS -GV	DEP-G
☞a.	*GV*	[tjow lə]						
b.	*CG*	[tow ljə]			*		**	
c.	*CG-repair*	[tow lje]				*	**	
d.	*GG*	[tjow ljə]			*		*	*
e.	*GG-repair*	[tjow lje]				*	*	*

(11) *Non-palatal-initial test item with no vowel alternation*

	Test item	[ta ljaw] 'star anise'	*PAL[-HI]	ID-C [ANT]	AGR[FR]-GV_{MID}	ID-V [FR]	ID-TRANS -GV	DEP-G
☞a.	*GV*	[la tjaw]						
b.	*CG*	[lja taw]					**	
c.	*GG*	[lja tjaw]					*	*

Comparing the palatal-initial test item tableaux of (6-8) and the non-palatal-initial ones of (9-11), we can see that the constraint ranking variation account makes different predictions on the response type for these two types of test items. For palatal-initial items, the three response types, GV, CG, and GG, ought to be all available, depending on the constraint ranking. For non-palatal items, only the GV response is optimal. It is expected to dominate. This is exactly what was observed in the production-based language game experiment data.

In Table 5.2, I compare side by side the prediction on response distribution made by the constraint ranking variation account and the actual response distribution elicited from the 33 participants in the production-based experiment. I list the test item types by phonological environment in the leftmost column. The middle three columns are predictions on whether a response type is available, based on constraint ranking variation. A checkmark indicates that the response type is expected to surface. A shaded cell means it is predicted to never be optimal, regardless of constraint ranking. There is a clear dichotomy based on consonant place. In the rightmost three columns, I copy the response type percentages from Table 4.12, as calculated from the response tokens elicited from the production-based experiment.

/j/ TEST ITEM TYPE	CONSTRAINT RANKING VARIATION ACCOUNT PREDICTION			PRODUCTION-BASED EXPERIMENT RESULTS		
	GV	CG	GG	GV	CG	GG
C: Palatal-initial V: Source V alternates	✓	✓	✓	12%	11%	**77%**
C: Palatal-initial V: Target V alternates	✓	✓	✓	13%	22%	65%
C: Palatal-initial V: No V alternates	✓	✓	✓	11%	20%	69%
C: Non-palatal-initial V: Source V alternates	✓			**96%**	0%	4%
C: Non-palatal-initial V: Target V alternates	✓			**95%**	3%	2%
C: Non-palatal-initial V: No V alternates	✓			**88%**	2%	9%

Table 5.2 Response type percentage for /j/ items in the production-based experiment, listed by test item type, compared to the prediction of the constraint ranking variation account

The actual response distribution roughly aligns with the prediction made by the constraint ranking variation account. Palatal-initial test items have elicited far more speaker variation in response type, compared to the non-palatal-initial test items. In the latter group, the GV response dominates. Even though there is a smattering of CG and GG responses for non-palatal-initial items, which are predicted to not be able to surface at all, this could be attributed to noise from the language game task. Participants' exposure to palatal-initial items that genuinely prompted a CG or GG response during the experiment might have made them more ready to produce such a token, even when it was not called for.

Table 5.2 shows that not only is the constraint ranking variation account able to predict the existence of speaker variation in response types, it can also predict when speaker variation is reduced to a minimum, namely in non-palatal-initial test items. Therefore, I conclude that the source of variation in the experiment data is a variation in the constraint ranking.

5.1.3 Puzzling response patterns

With the affirmation that constraint ranking variation can account for speaker variation in response types, I now turn my attention to two puzzling response patterns collected in the language game experiment. The first is the surprising frequency of GG responses for /w/ items. The other is the presence of [j] and [w] in responses for /ɥ/ items.

The response type distribution table for /w/ items is reproduced in Table 5.3. For four out of the six test item types, there is an unusual amount of GG responses. These are test items that either have a source syllable vowel alternation or no vowel alternation, regardless of consonant place. Recall that there is no consonant-glide phonotactics concerning the labial glide, therefore there is no reason for /w/ to move to the target syllable with the consonant. Even though there are glide-vowel phonotactics at stake, they are not expected to push the glide into the target syllable. What makes the amount of GG response more surprising is the fact they are predicted to not surface at all. In the tableaux for non-palatal-initial /j/ items in (9-11), the CG and GG response candidates are harmonically bound by the GV response. These tableaux are applicable for /w/ items since neither non-palatal-initial /j/ items nor /w/ items show consonant-glide phonotactics. Therefore, for the /w/ items, the CG and GG responses are also expected to be harmonically bound by the GV response. The low percentage of CG responses can be attributed to task noise, but the unusually high share of GG responses requires a separate explanation.

/w/ Test Item Type	Production Experiment Results		
	GV	CG	GG
C: Velar-initial V: Source V alternates	**85%**	1%	14%
C: Velar-initial V: Target V alternates	**97%**	3%	1%
C: Velar-initial V: No V alternates	**79%**	6%	15%
C: Non-velar-initial V: Source V alternates	**88%**	0%	12%
C: Non-velar-initial V: Target V alternates	**98%**	1%	1%
C: Non-velar-initial V: No V alternates	**84%**	6%	9%

Table 5.3 Distribution of response types for /w/ items in the production-based experiment, listed by test item type according to consonant place and vowel alternation.

One possible answer is that Mandarin speakers might be prioritizing the preservation of the consonant's acoustic features in the language game. The labial glide has a rounding coarticulation with the preceding consonant. For example, in [swan55] 'sour', the rounding gesture is reported to begin as early as the sibilant onset, unlike the English word *sway*, where rounding only takes place after the sibilant offset (Duanmu 2007). Therefore, it is possible that during the language game experiment, participants were trying to preserve the rounding coarticulation on the initial consonant. Sometimes, when they moved the rounded consonant into the target syllable, they brought the labial glide with it in order to license the rounding coarticulation. I use the faithfulness constraint MAX-C[ROUND] to capture this. Its working is shown in the tableaux in (12) and (13).

The consonants are transcribed with rounding coarticulation to highlight the effect of this newly introduced constraint. I do not consider candidates that contain a rounded consonant with no licensing glide, such as [ʂʷɑŋ tswej].

In (12), the test item [tsɑŋ55 ʂwej214] 'slander' has no vowel alternation. A highly ranked Max-C[Round] means the GG response will win. It is not harmonically bound by the GV response anymore. This accounts for the high percentage of GG responses for /w/ items. In (13), the test item [ʂə55 xwa35] 'luxury' shows target syllable vowel alternation. The GG response candidate violates the glide-vowel markedness constraint Agree[Round]-GV$_{MID}$, whereas the GG-repair candidate incurs a violation of Ident-V[Round]. This means that preserving the rounding coarticulation on the consonant has an unwelcomed effect. Therefore, speakers are expected to produce the GG response less frequently in this case, which is exactly what we observe in Table 5.4.

(12) *GG response can win for /w/ test item with no vowel alternation*

	Test item	[tsɑŋ ʂʷwej] 'slander'	Max-C [Rd]	Agr[Rd]-GV$_{MID}$	Id-V [Rd]	Id-Trans -GV	Dep-G
a.	*GV*	[ʂɑŋ tsʷwej]	*!				
b.	*CG*	[ʂʷwɑŋ tsej]				**!	
☞c.	*GG*	[ʂʷwɑŋ tsʷwej]				*	*

(13) *GG response dispreferred for /w/ test item with target syllable vowel alternation*

	Test item	[ʂə xʷwa] 'luxury'	Agr[Rd]-GV$_{MID}$	Id-V [Rd]	Id-Trans -GV	Max-C [Rd]	Dep-G
☞a.	*GV*	[xə ʂʷwa]				*	
b.	*CG*	[xʷwə ʂa]	*!		**		
c.	*CG-repair*	[xʷwo ʂa]		*!	**		
d.	*GG*	[xʷwə ʂʷwa]	*!		*		*
e.	*GG-repair*	[xʷwo ʂʷwa]		*!	*		*

Thus, the unusual amount of GG responses for some /w/ items are accounted for. Note that for the /j/ items, those with a non-palatal consonant and no vowel alternation also have led to an unexpected amount of GG responses (9%). The rounding coarticulation account cannot provide an explanation for the /j/ items. Upon closer inspection of the data, it is observed that two thirds of the GG response tokens collected for this group come from a single test item, [njow35 taw55] 'butcher knife'. It remains unclear to me why this particular test item has prompted so many GG response, but it does not seem to be representative of a wider pattern for non-palatal-initial /j/ items.

(14) *Variety of responses for* [lə51 tɕʰɥɛn35] *'Spring of Happiness'*

a.	[tɕʰə lɥɛn]	*GV response*	3 tokens	i.	[tɕʰɥe ljɛn]	*GG response*	2 tokens
b.	[tɕʰe lɥɛn]	*GV response*	1 token	j.	[tɕʰɥe lwan]	*GG response*	3 tokens
c.	[kʰə lɥɛn]	*GV response*	1 token	k.	[tɕʰɥo lwan]	*GG response*	1 token
d.	[tɕʰə lwan]	*GV response*	1 token	l.	[tɕʰjə lɥɛn]	*GG response*	1 token
e.	[tɕʰɥə lan]	*CG response*	3 tokens	m.	[tɕʰjə lwan]	*GG response*	2 tokens
f.	[tɕʰɥe lan]	*CG response*	5 tokens	n.	[tɕʰje lɥɛn]	*GG response*	3 tokens
g.	[tɕʰɥə ljɛn]	*GG response*	3 tokens	o.	[tɕʰje lwan]	*GG response*	2 tokens
h.	[tɕʰɥə lwan]	*GG response*	1 token	p.	[tɕʰje lwen]	*GG response*	1 token

(15) *Variety of responses for* [ɕɥɛn35 lə] *'unlikely'*

a.	[lɥɛn ɕə]	*GV response*	7 tokens	h.	[lɥɛn ɕɥe]	*GG response*	2 tokens
b.	[lɥɛn ʂə]	*GV response*	1 token	i.	[lɥɛn ɕjə]	*GG response*	3 tokens
c.	[lwan ɕə]	*GV response*	4 tokens	j.	[lɥɛn ɕje]	*GG response*	1 token
d.	[lwan ʂə]	*GV response*	1 token	k.	[ljɛn ɕɥə]	*GG response*	2 tokens
e.	[lan ɕɥə]	*CG response*	1 token	l.	[lwan ɕjə]	*GG response*	1 token
f.	[lan ɕɥe]	*CG response*	2 tokens	m.	[lwan ɕje]	*GG response*	3 tokens
g.	[lɥɛn ɕɥə]	*GG response*	1 token				

A second point of interest is the variety of responses elicited for /ɥ/ items, as shown by the example test items [lə51 tɕʰɥɛn35] 'Spring of Happiness' and [ɕɥɛn35 lə] 'unlikely' in (14) and (15). The former has led to 16 distinct responses, while the latter resulted in 13 different codewords. The labels GV, CG, and GG do not seem to do justice to the variety of shapes responses come in. Most interestingly, the palatal glide [j] and labial glide [w] make their appearance in these tokens.

The cause of the wide range of variation in response tokens for /ɥ/ items might be that there are simply too many markedness and faithfulness constraints at stake in the encoding process. Most /ɥ/ items begin with a palatal consonant, which means consonant-glide phonotactics need to be considered. They all contain vowel alternation in both syllables, due to the limited distribution of the labiopalatal glide itself. This means various glide-vowel phonotactic constraints are vying for attention. To repair these marked forms, multiple faithfulness constraints are violated. For instance, a palatal glide [j] can be inserted to fix an ill-formed syllable that violates *PAL[-HIGH]. This repair violates DEP-G. In the case of an ill-formed syllable *[lɥɛn], the glide might change into [j] or [w] as a repair, violating IDENT-G[ROUND] or IDENT-G[FRONT]. This has probably led to an explosion of response candidates for the /ɥ/ items. Naturally, it follows that many of the candidates are given an opportunity to surface, depending on the constraint ranking. The increased search space might also explain why participants spent a longer period of time figuring out the codeword response for /ɥ/ items during the experiment. I illustrate this point with an unranked tableau for the test item [ɕɥɛn35 lə] 'unlikely' in (16). I include genuine response tokens as candidates. Rounding coarticulation is not transcribed as part of the surface representation of the candidates, in the interest of space. However, any consonant preceding /w/ or /ɥ/ is assumed to be rounded. MAX-C[ROUND] assigns violations accordingly.

There is a new constraint in this tableau, DEP-ɥ. The insertion of a labiopalatal glide is against the principle of THE EMERGENCE OF THE UNMARKED (TETU). In derived phonological environment such as insertion, we usually expect the inserted segment to be unmarked in nature. The insertion of a palatal glide [j] is natural, since it is relatively unmarked. The labiopalatal glide [ɥ], on the other hand, with its rarity in the Mandarin lexicon, as well as the restrictions in combinatorial possibility associated with it, is relatively marked in the language. [ɥ] is a universally marked sound as well, appearing in fewer languages than [j]. Therefore, when the phonological environment calls for an insertion repair, it is the service of [j] that is usually solicited, as opposed to [ɥ]. The preference for j-insertion as opposed to ɥ-insertion is captured by the constraint DEP-ɥ, which is ranked above the more general DEP-G constraint.

(16) *Various response candidates allowed to surface for /ɥ/ items*

Token		Test item	[ɕɥɛn lə]	*Pal[-High]	Agr[Fr]-GV	*LɥVN	Id-V[fr]	Id-C[ant]	Max-C[rd]	Id-G[fr]	Id-G[rd]	Id-V[low]	Dep-ɥ	Dep-G	Id-Trans-GV
7	a.	*GV*	[lɥɛn ɕə]	*		*			*						
1	b.	*GV*	[lɥɛn ʂə]			*		*	*						
4	c.	*GV*	[lwan ɕə]	*					*	*		*			*
1	d.	*GV*	[lwan ʂə]					*	*	*		*			*
1	e.	*CG*	[lan ɕɥə]		*							*			**
2	f.	*CG*	[lan ɕɥe]				*					*			**
1	g.	*GG*	[lɥɛn ɕɥə]		*	*							*	*	*
2	h.	*GG*	[lɥɛn ɕɥe]			*	*						*	*	*
3	i.	*GG*	[lɥɛn ɕjə]		*	*			*					*	*
1	j.	*GG*	[lɥɛn ɕje]			*	*		*					*	*
2	k.	*GG*	[ljɛn ɕɥə]		*						*		*	*	**
1	l.	*GG*	[lwan ɕjə]		*				*	*		*		*	**
3	m.	*GG*	[lwan ɕje]				*		*	*		*		*	**

As shown in the tableaux of (12), (13), and (16), two puzzling response patterns in the experiment can be attributed to the interaction between various markedness and faithfulness constraints, providing further evidence for the constraint ranking variation account.

5.2 Lessons for phonotactic constraint learning

A key challenge to the study of Mandarin phonology is that even though there is plenty of surface co-occurrence restrictions, it is difficult to find concrete evidence for actual phonotactic constraints in the grammar of Mandarin speakers. The results from the codeword language game experiment have shown that certain phonotactic constraints are indeed learned by speakers as part of their grammar. What remains unclear is whether these constraints are learned to the same extent by speakers. In other words, what is the relative ranking between various phonotactic constraints in a speaker's grammar?

In the preceding section, I have shown that constraint ranking is subject to much speaker variation. Nevertheless, we might be able to detect trends at the population level. There could be constraints that are ranked relatively high by most speakers, and those that are ranked low by everyone. In order to find out if such trends exist, I surveyed the response tokens from the experiment, some of which violate phonotactic constraints. The goal is to identify which phonotactic constraints are more frequently violated by participants to the experiment, and which are more productive in forbidding marked syllable from surfacing. If speakers are shown to tolerate violations to a certain constraint in their response tokens, it is taken as evidence that the constraint is usually ranked low in Mandarin speakers' grammars. On the other hand, if speakers avoid violating a certain constraint to a high degree, then it means the constraint is ranked high for most speakers.

To calculate the frequency of violation for each phonotactic constraint in the language game response tokens, I took the following steps. First, I identified the test items in which a phonotactic constraint could potentially be violated. For example, the consonant-glide phonotactic constraint *PAL[-HIGH] could be violated by all palatal-initial test items, but never by any of the non-palatal-initial test items. Every phonotactic constraint is given a list of potential violator test items. Then, I counted the total number of response tokens elicited for the list of potential violators, as well as the number of marked tokens that have indeed violated the phonotactic constraint. As a final step, I divided the number of marked tokens by the total number of tokens that risk violation for a phonotactic constraint, deriving a violation frequency. A few phonotactic constraints introduced in Chapter 3 are never at risk of violation in the language game, therefore they are excluded from analysis. The results are summarized in Table 5.4, in which phonotactic constraints are in reverse order of their violation frequency. Constraints to the right are violated more frequently than those to the left.

	*PAL[-HIGH]	*EFFORT-PAL-an	AGR[FR]-GV	*LɥVN	*Tɥ
Potential violator test item example	[tɕje ta] *'solution'*	[tɕjɑŋ nan] *'Jiangnan'*	[tʰə tɕʰɥɛn] *'privilege'*	[lan tɕʰɥɪn] *'blue dress'*	[ɕɥɛn tə] *'Xuande (Emperor)'*
Number of potential violators	26	14	4	7	4
Total number of response tokens	789	417	85	212	118
Constraint-violating response tokens	55	59	30	117	69
Frequency of constraint violation	7%	14%	35%	55%	58%

Table 5.4 Frequency of phonotactic constraint violation by Mandarin speakers in the production-based codeword language game

Note that the order of constraints listed in Table 5.4 is not an actual ranking, but it is a good approximation of population-level trends. The constraint at the left end of the table, *PAL[-HIGH] is violated only 7% of the time by participants in the language game. This means it is very likely to be ranked very high in the grammar of most Mandarin speakers. At the other end of the table, we have the pair of phonotactic constraints regarding the labiopalatal glide, *Tɥ and *LɥVN. Participants violated these two constraints more than half of the time. This shows that violation to these constraints are frequently tolerated by the speakers, indicating that they are very likely to be ranked low in most speakers' grammar. The remaining phonotactic constraints, *PAL-an and AGREE[FRONT]-GV$_{MID}$, are situated somewhere in the middle. Speakers do not treat them as stringent constraints like *EFFORT-PAL-an, but they also do not violate them casually.

The frequency of constraint violation report in Table 5.4 shows that not all phonotactic constraints are learned to the same extent by Mandarin speakers. They appear to have learned *PAL[-HIGH] and *EFFORT-PAL-an better compared to *Tɥ and *LɥVN, applying it more productively in forming codeword responses. What might have contributed to the different degree of learning?

One possible answer is that the phonotactic constraints with lower violation frequency are more grounded in phonology and phonetics. The constraint *PAL[-HIGH] serves a key role in ensuring that the palatal sibilants remain contrastive to the denti-alveolar and retroflex sibilants. It demands that a palatal consonant always be followed by a high front glide or vowel, extending the duration of the palatal feature. For a speaker to violate this constraint, they risk losing the three-way contrast between the sibilants. In the case of *EFFORT-PAL-an, it is a *EFFORT constraint grounded in articulatory phonetics. Speakers have real difficulty reaching the three targets, /j/, /a/, and /n/ in quick succession. I suspect the reason this constraint is violated 14% of the time during the experiment is due to a syllable lengthening effect. Participants playing the language game usually pronounce the novel response tokens with a longer duration compared to natural speech. Therefore, they might have had more time to reach the low vowel target, deriving the sequence [jan] that is usually not found in the language.

In contrast, the frequently violated constraints *Tɥ and *LɥVN might appear to be more arbitrary to Mandarin learners. For *Tɥ, it is not clear why denti-alveolar sonorants are allowed to be next to /ɥ/, but not obstruents. As to *LɥVN, the fact that both LɥV and ɥVN are legal sequences in the language makes the unattestedness of LɥVN ever more perplexing. In a way, the discussion on these two constraints in Chapter 3 is telling. I could not identify any concrete motivation behind them, offering only speculation.

Another lesson to take from the results of Table 5.4 is that there is no clear boundary between systematic gaps and accidental gaps in Mandarin. Systematic gaps describe unattested forms in a language that can be attributed to underlying phonological principles, whereas accidental gaps point to those that do not violate any phonotactic constraints. Experiments on phonotactic judgement using wug tests have repeatedly shown speakers tend to judge accidental gaps to be better than systematic one. This is observed for Mandarin speakers (see S. Wang 1998; Myers 2015; Gong & J. Zhang 2021), as well as for speakers of other languages (see Frisch & Zawaydeh 2001; Hayes & White 2013). At the same time, many scholars have pointed out that there does not appear to be a clear categorical boundary between these two types of gaps (see Frisch, Large & Pisoni 2000; Boersma & Hayes 2001). The frequency at which participants produced unattested syllables in the codeword language experiment appears to support this view.

The violation frequency data show a gradient spectrum of phonotactic constraints. Some phonotactic constraints are stringent, like *PAL[-HIGH], which forbids systematic gaps like *[ɕa] from surfacing. Others are based on accidental gaps like *[faj], which is completely dormant when the English word *Wi-Fi* is borrowed faithfully into the language. Mandarin speakers who have almost no knowledge of English are observed to have no trouble producing the syllable [faj]. Yet many phonotactic constraints are situated somewhere in between these two extremes. *LɥVN is one such example. In the experiment, speakers did not resist *[lɥɛn] to the same degree in which they blocked the systematic gap *[ɕa]. But they were also not freely producing it on the fly like they usually do for the accidental gap [faj]. The unattested *[lɥɛn] is neither a systematic gap nor an accidental gap.

5.3 Lessons for faithfulness constraint learning

The experiment results also offer a glimpse into population-level trends in faithfulness constraint ranking. The clue is in the Mandarin speakers' treatment of the palatal sibilants in the language game. When they are confronted with a GV response candidate containing an ill-formed syllable like *[ɕa], the speakers' choice of repair can inform us on their faithfulness constraint ranking. Pooling the response tokens elicited for palatal-initial items, we can see that *PAL[-HIGH] is violated in 7% of the tokens, based on the violation frequency data in Table 5.4. To produce such ill-formed tokens, the faithfulness constraints IDENT-C and DEP-G must be ranked above the markedness constraint *PAL[-HIGH]. Another 10% of response tokens for this group involve a place of articulation change. The constraint ranking that can derive such a form must have *PAL[-HIGH] and DEP-G ranked higher than IDENT-C. Finally, the vast majority of response tokens (83%) for palatal-initial items use the repair of glide insertion, as in [ɕja]. This means that the preferred ranking of Mandarin speakers is {*PAL[-HIGH], IDENT-C} >> DEP-G. The frequency of the three types of response tokens is summarized in Table 5.5, along with the constraint ranking that can predict them.

Response type	Frequency	Constraint ranking
Ill-formed, no repair: *[ɕa]	7%	{IDENT-C, DEP-G} >> *PAL[-HIGH]
Change place of articulation: [sa/ʂa/xa]	10%	{*PAL[-HIGH], DEP-G}>> IDENT-C
Glide insertion: [ɕja]	83%	{*PAL[-HIGH], IDENT-C}>> DEP-G

Table 5.5 Frequency of three response types to palatal-initial test items in the production-based experiment, matched with a constraint ranking that derives the response type

The fact that the vast majority of Mandarin speakers have converged on the same faithfulness constraint ranking, namely IDENT-C >> DEP-G, shows that speakers can learn a consistent phonological grammar, without the assistance of morphophonological alternations. Whether this ranking reflects a universal learning bias, regardless of the phonological input, is an interesting question for future research.

The treatment of the palatal sibilants in the experiment also addresses a longstanding question in phonological learning. If a learner has access to only complementary distribution without morphophonological alternation, can they learn non-identical UR-SR mappings (see Yip 1996, Alderete & Tesar 2002; McCarthy 2005; Rasin & Shefi 2019; Richter 2021)? The answer seems to be a resounding no. The palatal sibilants are in complementary distribution with three other groups of consonants, the denti-alveolar sibilants, retroflexes, and velars. But there is no morphophonological alternation between any of the groups. The codeword language game experiment shows that Mandarin speakers rarely change the place of articulation of the palatal consonants, which indicates that the palatal sibilants are not mapped to another group of consonants. Therefore, I conclude that Mandarin speakers have not learned non-identical UR-SR mapping for the palatal consonants. This result adds support for lexicon optimization, which states that there is a learning bias towards faithful UR-SR mapping (see Prince & Smolensky 1993).

The participants' treatment of palatal sibilants can also shed some light on another important question in phonology, namely how is contrast maintained between groups of consonants with closely situated places of articulation. It is difficult to pack three groups of sibilants, denti-alveolars, palatals, and retroflexes in the tight area of the front oral cavity, while making sure there is salient contrast between them (see Flemming 2018). Many languages only have a two-way contrast for coronal sibilants. Sibilant place contrast is usually expressed with two perceptual cues, spectral center of gravity (CoG) during frication and sibilant-vowel F2 transition (Żygis & Padgett 2010). Mandarin is no different. C. Lee, Y. Zhang & X. Li (2014) demonstrated with a production experiment that Mandarin denti-alveolar sibilants have the highest CoG, whereas retroflex sibilants have the lowest CoG. Palatal sibilants are situated somewhere between these two groups. They also found that palatal sibilants have a much higher vowel onset F2 than the other two types of sibilants.

The pattern of response tokens produced for palatal-initial test items strongly suggests that transitional F2 is an indispensable perceptual cue for the palatal sibilants. In the codeword language game, speakers are tasked with moving the palatal consonant into the target syllable, which might result in a loss of transitional F2 cue. The most popular course of action is to bring the palatal glide with the palatal sibilant, as in a CG or GG response, thus recreating the high transitional F2 in the target syllable. 83% of all response tokens collected for palatal-initial items were CG and GG responses. The rest of tokens were GV responses. Some speakers changed the place of articulation of the palatal consonant. This practice is found in 10% of response tokens. To these speakers, transitional F2 is also an important perceptual cue for palatal sibilants. When the transitional F2 cue is removed in a GV response, all that is left is CoG. For these speakers, perhaps the CoG cue alone is not enough to maintain the three-way contrast, thus the palatal sibilants become indistinguishable from the denti-alveolars or the retroflexes. Finally, 7% of response tokens were GV responses with no place change, which means some speakers ended up producing a marked syllable like *[ɕa]. It is possible that these speakers can produce a three-way contrast with CoG as the sole acoustic cue. Transitional F2 is not as important to this group of speakers compared to the rest. The divergent behavior of Mandarin speakers in the experiment appears to show that there is speaker variation on the weight of different perceptual cues in realizing contrast. More production and perception experiments need to be conducted, in order to identify the exact nature and degree of this variation.

Chapter 6 Conclusion

In order to find out what Mandarin speakers know about their language, in the absence of morphophonological alternations, I conducted two codeword language game experiments. I repurposed the fanqie method, which is usually used for identifying syllable structure, for understanding phonological constraint interactions. To this end, tClick or tap here to enter text.he experiment stimuli are fully controlled for the phonological environment of the glide. The preceding consonant is controlled for its place of articulation and the following vowel is controlled for whether it displays potential alternation. This means that the effect of markedness and faithfulness constraints can be detected.

The experiment results show that there is much speaker variation in the shape of response tokens. Nevertheless, the distribution of response types shows clear signs of markedness and faithfulness constraint effects. A closer inspection of the data reveals population-level trends in constraint ranking as well.

In terms of phonotactic markedness constraints, the Mandarin speakers' habit for moving the palatal glide with the palatal consonant in spite of syllable structure is an indication that consonant-glide phonotactic constraints are psychologically real for them. The preference to pronounce the glide in its original position, when the source syllable contains an alternating vowel, coupled with the aversion to introducing the glide in the target syllable when the vowel there potentially alternates, are evidence that Mandarin speakers have learned glide-vowel phonotactic constraints. In addition, an inspection of the response tokens containing unattested syllables shows that not all phonotactic constraints have been learned to the same extent. In particular, constraints with phonological or phonetic grounding are learned more consistently than those without. This supports the view that there is a naturalness bias in phonological learning.

Lessons on faithfulness constraints are drawn from the speakers' treatment of the palatal sibilants in the experiment. The codeword language game provides a chance for palatal sibilants to appear in a hostile environment, which prompts mapping between ill-formed syllables to well-formed ones, thus revealing faithfulness constraint ranking. Instances of palatal sibilant place change are rare in the response tokens. Furthermore, even among participants who made place changes more frequently than others, the destination was not unique. In the response tokens of the same speaker, we find palatals changed into denti-alveolars, as well as palatals swapped for retroflexes. The two facts combined argue against the existence of consistent mapping between the palatal sibilants and another group of consonants in the synchronic grammar of Mandarin. The preferred repair strategy to such ill-formed syllables appears to be glide insertion, which improves the hostile phonological environment for the palatals. Speakers generally prioritize consonant identity over guarding the output against newcomer segments. The results also offer evidence for lexicon optimization, which states that in the absence of morphophonological alternations, speakers will posit underlying representations that are identical to surface representations.

The language game method has been demonstrated to reveal the phonological knowledge of adult Mandarin speakers. It offers a glimpse into how speakers have acquired markedness and faithfulness constraints in their early years. This in turn can inform us on some universal learning biases when learners are confronted with ambiguous phonological input. The language game method can be applied to other languages as well, providing more insight into the process of phonological acquisition.

Bibliography

Alderete, John & Bruce Tesar. 2002. *Learning covert phonological interaction: an analysis of the problem posed by the interaction of stress and epenthesis*. Piscataway.

Ao, Benjamin. 1992. Non-Uniqueness Condition and the Segmentation of the Chinese Syllable. *Ohio State University Working Papers in Linguistics No. 41*.

Bao, Zhiming. 1990. "Fanqie" languages and reduplication. *Linguistic Inquiry* 21. 317–350.

Barnes, Jonathan. 2002. Palatalization in Bulgarian Dialects: An Experiment in Phoneme Categorization. In Ronelle Alexander & Vladimir Zhobov (eds.), *Revitalizing Bulgarian Dialectology*, 24–43. University of California Press.

Baxter, William H. 1992. *A Handbook of Old Chinese Phonology*. Berlin, New York: De Gruyter Mouton. https://doi.org/doi:10.1515/9783110857085.

Bender, Byron. 1968. Marshallese phonology. *Oceanic Linguistics* 7(1). 16–35.

Boersma, Paul & Bruce Hayes. 2001. Empirical Tests of the Gradual Learning Algorithm. *Linguistic Inquiry* 32. 45–86.

Boersma, Paul & David Weenink. 2022. Praat: doing phonetics by computer. https://www.fon.hum.uva.nl/praat/.

Carden, Kelly Ann. 2016. *Vowel-consonant interaction in two dialects of Mandarin*. Iowa City: University of Iowa PhD dissertation.

Chao, Yuen Ren. 1934. The Non-Uniqueness of Phonemic Solutions of Phonetic Systems. *Bulletin of the Institute of History and Philology, Academia Sinica* 4. 363–397.

Chao, Yuen-Ren. 1931. 反切語八種 [Eight varieties of secret languages using Fan-ch'ieh]. *Bulletin of the Institute of History and Philology, Academia Sinica* 2(3). 312–354.

Chao, Yuen-Ren. 1948. The voiced velar fricative as an initial in Mandarin. *Le Maître Phonétique* 89. 2–3.

Chao, Yuen-Ren. 1968. *A Grammar of Spoken Chinese*. Berkeley and Los Angeles: University of California Press.

Cheng, Chin-chuan. 1968. *Mandarin Phonology*. University of Illinois.

Cheng, Chin-Chuan. 1973. *A Synchronic Phonology of Mandarin Chinese*. Berlin, New York: De Gruyter Mouton.

Cheng, Robert L. 1966. Mandarin phonological structure. *Journal of Linguistics* 2. 135–158.

Choi, Gwon-Jin. 1998. The phonological value of the feature [palatalness] in the contemporary Bulgarian language. In *Paper presented at 12th International Congress of Slavic Scholars, Poland.*

Choi, John-Dongwook. 1992. Phonetic Underspecification and Target Interpolation: An Acoustic Study of Marshallese Vowel Allophony. *UCLA Working Papers in Phonetics No. 82*.

Dong, Shao-wen. 1964. *语音常识 [General Knowledge of Phonetics]*. Beijing.

Duanmu, San. 1990. *A formal study of syllable, tone, stress and domain in Chinese languages*. Massachusetts Institute of Technology PhD dissertation.

Duanmu, San. 2007. *The Phonology of Standard Chinese*. 2nd edn. Oxford: Oxford University Press.

Elff, Martin. 2022. Multinomial Logit Models, with or without Random Effects or Overdispersion. https://cran.r-project.org/web/packages/mclogit/index.html.

Firth, J R & Benjamin Bickley Rogers. 1937. The Structure of the Chinese Monosyllable in a Hunanese Dialect (Changsha). *Bulletin of the School of Oriental and African Studies* 8. 1055–1074.

Flemming, Edward. 1995. *Auditory Representations in Phonology*. University of California, Los Angeles PhD dissertation.

Flemming, Edward. 2011. The grammar of coarticulation (English draft). In Christelle Dodane & Mohamed Embarki (eds.), *La Coarticulation: Des Indices à la Représentation*. Paris: Editions L'Harmattan.

Flemming, Edward. 2018. Systemic Markedness in Sibilant Inventories. In *Poster presented at AMP 2018*. San Diego.

Frisch, Stefan A, Nathan R Large & David B Pisoni. 2000. Perception of Wordlikeness: Effects of Segment Probability and Length on the Processing of Nonwords. *Journal of memory and language* 42 4. 481–496.

Frisch, Stefan A & Bushra Zawaydeh. 2001. The Psychological Reality of OCP-Place in Arabic. *Language* 77. 106–91.

Fu, Boer. 2022. The Segment Status of the Mandarin Glide: A Language Game Experiment. In *Proceedings of the 2022 Annual Meeting on Phonology*. Los Angeles.

Fu, Maoji. 1957. 北京話的音位和拼音字母 [The phonemes of Beijing Mandarin and the Pinyin letters]. *Zhongguo Yuwen* 5. 3–12.

Goldwater, Sharon & Mark Johnson. 2003. Learning OT constraint rankings using a maximum entropy model. In *Proceedings of the Stockholm workshop on variation within Optimality Theory*, 120.

Gong, Shuxiao & Jie Zhang. 2021. Modelling mandarin speakers' phonotactic knowledge. In *Phonology*, vol. 38, 241–275. Cambridge University Press. https://doi.org/10.1017/S0952675721000166.

Halliday, Michael Alexander Kirkwood. 1992. A systematic interpretation of Peking syllable finals. In Paul Tench (ed.), *Studies in Systematic Phonology*, 98–121. London: Pinter.

Hartman, Lawton M. 1944. The Segmental Phonemes of the Peiping Dialect. *Language* 20. 28.

Hayes, Bruce & James White. 2013. Phonological Naturalness and Phonotactic Learning. *Linguistic Inquiry* 44. 45–75.

Hayes, Bruce & Colin Wilson. 2008. A Maximum Entropy Model of Phonotactics and Phonotactic Learning. *Linguistic Inquiry* 39(3). 379–440. https://doi.org/10.1162/ling.2008.39.3.379.

Hockett, Charles Francis. 1947. Peiping phonology. *Journal of the American Oriental Society* 67. 253–267.

Hsueh, Fengsheng. 1986. *國語音系解析 [An Anatomy of the Pekingese Sound System]*. Taipei: Taiwan Student Book.

Ji, Xianlin (ed.). 1988. *中国大百科全书：语言文字 [Chinese Encyclopaedia: Language and Orthography]*. (Ed.) Xianlin Ji. Beijing and Shanghai: Encyclopedia of China Publishing House.

Kuo, Feng-Lan. 1994. *Aspects of segmental phonology and Chinese syllable structure*. University of Illinois at Urbana-Champaign PhD dissertation.

Ladefoged, Peter & Ian Maddieson. 1996. *The sounds of the world's languages*.

Lee, Chao-Yang, Yu Zhang & Ximing Li. 2014. Acoustic characteristics of voiceless fricatives in Mandarin Chinese. *Journal of Chinese Linguistics* 42(1). 150–171.

Li, Wen-chao. 1999. *A diachronically-motivated segmental phonology of Mandarin Chinese*. New York: Peter Lang.

Lin, Yen-Hwei. 1989. *Autosegmental treatment of segmental processes in Chinese phonology*. University of Texas at Austin PhD dissertation.

Lin, Yen-Hwei. 2002. Mid Vowel Assimilation Across Mandarin Dialects. *Journal of East Asian Linguistics* 11. 303–347.

Martin, Samuel Elmo. 1957. Problems of hierarchy and indeterminacy in Mandarin phonology. *Bulletin of the Institute of History and Philology, Academia Sinica* 29. 209–229.

McCarthy, John J. 2005. Taking a Free Ride in Morphophonemic Learning. *Catalan journal of linguistics* 4. 19–55.

Myers, James. 2015. Markedness and Lexical Typicality in Mandarin Acceptability Judgments. *Language and Linguistics* 16. 791–818.

Prince, Alan & Paul Smolensky. 1993. *Optimality Theory: Constraint Interaction in Generative Grammar*.

Pritchard, Sonia. 2012. *A Cross-Language Study of the Production and Perception of Palatalized Consonants*. University of Ottawa.

Pulleyblank, Edwin George. 1984. *Middle Chinese: A Study in Historical Phonology*. Vancouver: University of British Columbia Press.

Purple Culture. 2021. Talking Chinese to pinyin/zhuyin converter. https://www.purpleculture.net/chinese_pinyin_converter/.

R Core Team. 2023. R: A language and environment for statistical computing. Vienna: R Foundation for Statistical Computing. https://www.r-project.org/.

Rasin, Ezer, Itamar Shefi & Roni Katzir. 2020. A unified approach to several learning challenges in phonology. In *NELS 50: Proceedings of the Fiftieth Annual Meeting of the North East Lingusitic Society Volume Three*, 73–86.

Richter, Caitlin. 2021. *Alternation-sensitive Phoneme Learning: Implications for Children's Development and Language Change*. Philadelphia: University of Pennsylvania PhD dissertation.

Scatton, Ernest A. 1975. *Bulgarian phonology* . *Bulgarian phonology*. Cambridge, Mass: Slavica Publishers.

Shi, Cun-zhi. 1957. 北京話音位問題商較 [A discussion of Pekingnese phonemes]. *Zhongguo Yuwen* 56. 9–12.

Steriade, Donca. 1997. *Phonetics in phonology: The case of laryngeal neutralization*. Los Angeles.

Suh, Yunju & Jiwon Hwang. 2016. The Korean Prevocalic Palatal Glide: A Comparison with the Russian Glide and Palatalization. *Phonetica* 73. 85–100.

Tung, Tung-ho. 1954. *中國語音史 [A Phonological History of Chinese]*. Taipei.

Wan, I-Ping. 1999. *Mandarin phonology: Evidence from speech errors*. *ProQuest Dissertations and Theses*. United States -- New York: State University of New York at Buffalo.

Wang, H Samuel & Chih-ling Chang. 2001. On the Status of the Prenucleus Glide in Mandarin Chinese. *Language and Linguistics* 2(2). 243–260.

Wang, Hongjun. 1999. *汉语非线性音系学 [Chinese Non-linear Phonology]*. Beijing: Peking University Press.

Wang, Jenny Zhijie. 1993. *The geometry of segmental features in Beijing Mandarin*. University of Delaware PhD dissertation.

Wang, Li. 1980. *汉语音韵 [Chinese Phonology]*. 2nd edn. Beijing: Zhonghua Shoji.

Wang, Samuel. 1998. An experimental study on the phonotactic constraints of Mandarin Chinese. In Benjamin K. T'sou (ed.), *Studia linguistica serica*, 259–268. Hong Kong: Language Information Sciences Research Center, City University of Hong Kong.

Wong, Helen. 1953. Outline of the Mandarin Phonemic System. *WORD*. Routledge 9(3). 268–276. https://doi.org/10.1080/00437956.1953.11659474.

Wu, Fei & Michael Kenstowicz. 2015. Duration reflexes of syllable structure in Mandarin. *Lingua* 164. 87–99.

Wu, Yuwen. 1994. Mandarin segmental phonology. *Toronto Working Papers in Linguistics* (University of Toronto Working Papers in Linguistics).

Yip, Moira. 1982. Reduplication and C-V Skeleta in Chinese Secret Languages. *Linguistic Inquiry* 13(4). 637–661.

Yip, Moira. 1996. Lexicon Optimization in Languages without Alternations. In Jacques Durand & Bernard Laks (eds.), *Current Trends in Phonology: Models and methods* . Salford: University of Salford Publications.

Yip, Moira. 2003. Casting doubt on the Onset - Rime distinction. *Lingua* 113(8). 779–816. https://doi.org/10.1016/S0024-3841(02)00130-4.

You, Rujie, Nairong Qian & Zhengxi Gao. 1980. 论普通话的音位系统 [On the Phonemic System of Modern Standard Chinese]. *Zhongguo Yuwen* 5. 328–334.

Zehr, Jeremy & Florian Schwarz. 2018. PennController for Internet Based Experiments. https://doc.pcibex.net/.

Żygis, Marzena & Jaye Padgett. 2010. A perceptual study of Polish fricatives, and its implications for historical sound change. *J. Phonetics* 38. 207–226.

Appendix I: Test stimuli in the online writing-based language game experiment

Group A test items:

Test item	Simplified Chinese	English	Glide	Consonant place	Source vowel alternation	Glide position
tɕʰjɛn55 la55	牵拉	‘to pull’	j	Palatal	Yes	1
tɕʰjɛn21 low51	浅陋	‘shallow’	j	Palatal	Yes	1
ɕje51 low51	泄漏	‘leak’	j	Palatal	Yes	1
taw51 tɕʰje51	盗窃	‘theft’	j	Palatal	Yes	2
taw21 ɕjɛn51	导线	‘wire’	j	Palatal	Yes	2
paw21 tɕjɛn51	保健	‘healthcare’	j	Palatal	Yes	2
tɕjaw21 low35	角楼	‘turret’	j	Palatal	No	1
tɕʰjaw51 pʰan51	翘盼	‘to look forward to’	j	Palatal	No	1
ɕjaw21 pan55	小班	‘kindergarten grade’	j	Palatal	No	1
law21 tɕjow51	老舅	‘maternal uncle'	j	Palatal	No	2
tʰan35 ɕjaw51	谈笑	‘to talk and laugh’	j	Palatal	No	2
ta55 tɕjow51	搭救	‘to rescue’	j	Palatal	No	2
ljɛn35 ta51	联大	‘UN General Assembly’	j	Non-palatal	Yes	1
tʰjɛn35 paw214	填饱	‘to fill’	j	Non-palatal	Yes	1
ljɛn35 paw214	联保	‘warranty’	j	Non-palatal	Yes	1
low35 mjɛn51	楼面	‘floor area’	j	Non-palatal	Yes	2
ta21 lje51	打猎	‘to hunt’	j	Non-palatal	Yes	2
tʰaw35 pʰjɛn51	陶片	‘pottery’	j	Non-palatal	Yes	2
ljow35 taw51	留到	‘to keep until’	j	Non-palatal	No	1
tjaw55 low35	碉楼	‘watchtower’	j	Non-palatal	No	1
ljaw35 ta51	辽大	‘Liaoning University’	j	Non-palatal	No	1
tow51 mjaw35	豆苗	‘sprout’	j	Non-palatal	No	2
low51 tjaw51	漏掉	‘to omit’	j	Non-palatal	No	2
lan35 pjaw55	蓝标	‘blue label’	j	Non-palatal	No	2

Group A filler items:

Test item	Simplified Chinese	English	Glide
pʰa51 lɤŋ214	怕冷	‘to be afraid of cold’	None
lan35 ʈʂow55	兰州	‘Lanzhou (city)’	None
kʰaw51 ʈʂen214	靠枕	‘cushion’	None
tɑŋ55 ɻan35	当然	‘of course’	None
pan51 kʰa214	办卡	‘to apply for a card’	None
tʰə51 ʈʂɤŋ55	特征	‘characteristic’	None
pɑŋ55 ʈʂʰen51	帮衬	‘to assist’	None
ta51 xow214	大吼	‘to shout’	None
kan55 ʈʂə	甘蔗	‘sugarcane’	None
kow51 tsaw51	构造	‘structure’	None
fa55 lɤŋ51	发愣	‘to daze’	None
paj35 tʰa214	白塔	‘white tower’	None
xə51 tsʰaj214	喝彩	‘to cheer’	None
lə51 kaw55	乐高	‘Lego’	None
naj51 tsɑŋ55	耐脏	‘to not stain easily’	None
ʈʂʰaw55 tsan51	超赞	‘very good’	None
ma35 tsaj214	马仔	‘gang member’	None
xə55 ʈʂow55	喝粥	‘to eat congee’	None
mɤŋ21 ɻan35	猛然	‘suddenly’	None
taw21 kow51	导购	‘shop assistant’	None
ʈʂaw21 tsə35	沼泽	‘swamp’	None
ʈʂʰaj55 san51	拆散	‘to break apart'	None
mɑŋ35 ɻan35	茫然	‘clueless’	None
tow51 ʈʂa55	豆渣	‘tofu dreg’	None

Group B test items:

Test item	Simplified Chinese	English	Glide	Consonant place	Source vowel alternation	Glide position
tɕjɛn21 paw51	简报	'briefing'	j	Palatal	Yes	1
tɕjɛn21 taw55	剪刀	'scissors'	j	Palatal	Yes	1
tɕjɛn21 ta35	简答	'short answer'	j	Palatal	Yes	1
tʰaw55 tɕʰjɛn35	掏钱	'to pay'	j	Palatal	Yes	2
lan35 tɕje35	拦截	'to intercept'	j	Palatal	Yes	2
paw21 ɕjɛn55	保鲜	'to keep fresh'	j	Palatal	Yes	2
tɕjɑŋ55 nan35	江南	'Jiangnan (region)'	j	Palatal	No	1
tɕʰjaw35 tan55	乔丹	'Jordan'	j	Palatal	No	1
ɕjaw21 lan35	晓岚	'Xiaolan (name)'	j	Palatal	No	1
tow51 tɕjaw214	豆角	'green bean'	j	Palatal	No	2
lan35 tɕʰjaw35	蓝桥	'blue bridge'	j	Palatal	No	2
lan35 ɕjɑŋ35	蓝翔	'Lanxiang (school)'	j	Palatal	No	2
tʰje21 law35	铁牢	'jail'	j	Non-palatal	Yes	1
tʰje21 maw35	铁锚	'anchor'	j	Non-palatal	Yes	1
tjɛn51 paw51	电报	'telegraph'	j	Non-palatal	Yes	1
law21 tʰjɛn55	老天	'god'	j	Non-palatal	Yes	2
taw51 pjɛn55	道边	'curbside'	j	Non-palatal	Yes	2
lan51 tʰje214	烂铁	'metal scrap'	j	Non-palatal	Yes	2
njow35 taw55	牛刀	'butcher knife'	j	Non-palatal	No	1
tjaw55 pan214	雕版	'woodblock print'	j	Non-palatal	No	1
mjaw51 tʰan51	妙探	'marvelous detective'	j	Non-palatal	No	1
nan35 pʰjaw51	男票	'boyfriend'	j	Non-palatal	No	2
nan35 mjaw51	南庙	'south temple'	j	Non-palatal	No	2
pʰan35 tʰjaw35	盘条	'wire rod'	j	Non-palatal	No	2

Group B filler items:

Test item	Simplified Chinese	English	Glide
fen51 nej51	份内	‘one’s duty’	None
pan55 ʈʂɑŋ214	班长	‘class president’	None
tan55 tsʰɤŋ35	单层	‘single layer’	None
pʰa51 law35	帕劳	‘Palau’	None
xɤŋ35 saw214	横扫	‘clean sweep’	None
ma21 laj35	马来	‘Malay’	None
taw51 kʰow214	道口	‘crossroad’	None
xaj35 tsaw214	海藻	‘seaweed’	None
kʰə21 ɻan35	可燃	‘inflammable’	None
taj51 ʈʂʰan214	待产	‘expectant (mother)’	None
taj51 kʰə51	代课	‘to substitute to teach’	None
pɑŋ21 tan55	榜单	‘chart’	None
maj51 lɑŋ51	麦浪	‘wheat field’	None
kʰow21 ʈʂaw51	口罩	‘mask (medical)’	None
pʰaj35 xan51	排汗	‘to sweat’	None
paw21 ʈʂɑŋ51	保障	‘guarantee’	None
man21 tʰɑŋ35	满堂	‘full house’	None
lɤŋ21 ʈʂan51	冷战	‘Cold War’	None
tʰan51 ʈʂɑŋ214	探长	‘chief inspector’	None
kaw55 ʈʂʰan214	高产	‘prolific’	None
pʰaw21 tʰɑŋ35	跑堂	‘waiter’	None
xɤŋ55 ʈʂʰɑŋ51	哼唱	‘to hum’	None
pʰɑŋ35 ɻan35	庞然	‘gigantic’	None
fa35 tan55	罚单	‘fine ticket’	None

Group C test items:

Test item	Simplified Chinese	English	Glide	Consonant place	Source vowel alternation	Glide position
tɕʰjɛn55 la55	牵拉	'to pull'	j	Palatal	Yes	1
tɕʰjɛn21 low51	浅陋	'shallow'	j	Palatal	Yes	1
ɕje51 low51	泄漏	'leak'	j	Palatal	Yes	1
tɕjɛn21 paw51	简报	'briefing'	j	Palatal	Yes	1
tɕjɛn21 taw55	剪刀	'scissors'	j	Palatal	Yes	1
tɕjɛn21 ta35	简答	'short answer'	j	Palatal	Yes	1
taw51 tɕʰje51	盗窃	'theft'	j	Palatal	Yes	2
taw21 ɕjɛn51	导线	'wire'	j	Palatal	Yes	2
paw21 tɕjɛn51	保健	'healthcare'	j	Palatal	Yes	2
tʰaw55 tɕʰjɛn35	掏钱	'to pay'	j	Palatal	Yes	2
lan35 tɕje35	拦截	'to intercept'	j	Palatal	Yes	2
paw21 ɕjɛn55	保鲜	'to keep fresh'	j	Palatal	Yes	2
ɕɥe55 lan35	薛岚	'Xuelan (name)'	ɥ	Palatal	Yes	1
tɕʰɥe51 lan35	雀兰	'Quelan (brand)'	ɥ	Palatal	Yes	1
ɕɥɛn35 lə	悬了	'unlikely'	ɥ	Palatal	Yes	1
ɕɥɛn55 tə35	宣德	'Xuande (emperor)'	ɥ	Palatal	Yes	1
tɕʰɥʊŋ35 lan214	穷览	'to read exhaustively'	ɥ	Palatal	Yes	1
tɕɥʊŋ21 tan51	囧蛋	'awkward person'	ɥ	Palatal	Yes	1
nan35 tɕʰɥe51	难却	'difficult to let go'	ɥ	Palatal	Yes	2
tʰə51 tɕʰɥɛn35	特权	'privilege'	ɥ	Palatal	Yes	2
lan21 ɕɥʊŋ35	懒熊	'lazy bear'	ɥ	Palatal	Yes	2
lan35 tɕʰɥʊŋ35	蓝穹	'blue dome (sky)'	ɥ	Palatal	Yes	2
lə51 tɕʰɥɛn35	乐泉	'Spring of Happiness'	ɥ	Palatal	Yes	2
nan35 tɕɥe35	男爵	'baron'	ɥ	Palatal	Yes	2

Group C filler items:

Test item	Simplified Chinese	English	Glide
pʰa51 lɤŋ214	怕冷	‘to be afraid of cold’	None
fen51 nej51	份内	‘one’s duty’	None
lan35 ʈʂow55	兰州	‘Lanzhou (city)’	None
pan55 ʈʂɑŋ214	班长	‘class president’	None
kʰaw51 ʈʂen214	靠枕	‘cushion’	None
tan55 tsʰɤŋ35	单层	‘single layer’	None
tɑŋ55 ɻan35	当然	‘of course’	None
pʰa51 law35	帕劳	‘Palau’	None
pan51 kʰa214	办卡	‘to apply for a card’	None
xɤŋ35 saw214	横扫	‘clean sweep’	None
tʰə51 ʈʂɤŋ55	特征	‘characteristic’	None
ma21 laj35	马来	‘Malay’	None
pɑŋ55 ʈʂʰen51	帮衬	‘to assist’	None
taw51 kʰow214	道口	‘crossroad’	None
ta51 xow214	大吼	‘to shout’	None
xaj35 tsaw214	海藻	‘seaweed’	None
kan55 ʈʂə	甘蔗	‘sugarcane’	None
kʰə21 ɻan35	可燃	‘inflammable’	None
kow51 tsaw51	构造	‘structure’	None
taj51 ʈʂʰan214	待产	‘expectant (mother)’	None
fa55 lɤŋ51	发愣	‘to daze’	None
taj51 kʰə51	代课	‘to substitute to teach’	None
paj35 tʰa214	白塔	‘white tower’	None
pɑŋ21 tan55	榜单	‘chart’	None

Appendix II: Test stimuli in the production-based language game experiment

/j/ test items:

Test item	Simpl. Chinese	English	Glide	Consonant place	Vowel alternation	Glide position
tɕje21 ta35	解答	'solution'	j	Palatal	Source syl.	1
tɕʰjɛn21 low51	浅陋	'shallow'	j	Palatal	Source syl.	1
taw51 tɕʰje51	盗窃	'theft'	j	Palatal	Source syl.	2
paw21 ɕjɛn55	保鲜	'to keep fresh'	j	Palatal	Source syl.	2
tɕjɑŋ55 nan35	江南	'Jiangnan (region)'	j	Palatal	Target syl.	1
ɕjaw21 pan55	小班	'kindergarten grade'	j	Palatal	Target syl.	1
tan51 tɕjaw214	蛋饺	'egg dumpling'	j	Palatal	Target syl.	2
lan35 ɕjɑŋ35	蓝翔	'Lanxiang (school)'	j	Palatal	Target syl.	2
tɕjaw21 low35	角楼	'turret'	j	Palatal	No	1
tɕʰjow35 law35	囚牢	'jail'	j	Palatal	No	1
ta51 tɕʰjaw35	大桥	'big bridge'	j	Palatal	No	2
law21 tɕjow51	老舅	'maternal uncle'	j	Palatal	No	2
pje35 naw51	别闹	'knock it off'	j	Non-palatal	Source syl.	1
tjɛn51 paw51	电报	'telegraph'	j	Non-palatal	Source syl.	1
ta21 lje51	打猎	'to hunt'	j	Non-palatal	Source syl.	2
taw51 pjɛn55	道边	'curbside'	j	Non-palatal	Source syl.	2
mjaw51 tʰan51	妙探	'marvelous detective'	j	Non-palatal	Target syl.	1
tjaw55 pan214	雕板	'woodblock print'	j	Non-palatal	Target syl.	1
nan35 mjaw51	南庙	'south temple'	j	Non-palatal	Target syl.	2
man21 ljow35	满流	'full flow'	j	Non-palatal	Target syl.	2
njow35 taw55	牛刀	'butcher knife'	j	Non-palatal	No	1
tjaw55 low35	碉楼	'watchtower'	j	Non-palatal	No	1
tow51 mjaw35	豆苗	'sprout'	j	Non-palatal	No	2
ta51 ljaw51	大料	'star anise'	j	Non-palatal	No	2

/w/ *test items:*

Test item	Simpl. Chinese	English	Glide	Consonant place	Vowel alternation	Glide position
kʰwo51 san51	扩散	‘spread’	w	Velar	Source syl.	1
kwo21 ʈʂʰa35	果茶	‘fruit tea’	w	Velar	Source syl.	1
ʂan55 xwo214	山火	‘wildfire’	w	Velar	Source syl.	2
san55 kwo35	三国	‘Three Kingdoms’	w	Velar	Source syl.	2
kʰwaj51 ʈʂʰə55	快车	‘express train’	w	Velar	Target syl.	1
kwan55 tsʰə51	观测	‘to observe’	w	Velar	Target syl.	1
sə51 xwan35	色环	‘color ring’	w	Velar	Target syl.	2
ʂə55 xwa35	奢华	‘luxury’	w	Velar	Target syl.	2
xwɑŋ35 ʂan55	黄山	‘Yellow Mountain’	w	Velar	No	1
xwen55 tsʰaj51	荤菜	‘meat dish’	w	Velar	No	1
tsɑŋ55 kʰwan214	赃款	‘embezzled money’	w	Velar	No	2
ʂan21 xwej35	闪回	‘flashback’	w	Velar	No	2
tswo51 ʈʂʰan35	坐禅	‘seated meditation’	w	Non-velar	Source syl.	1
ʂwo55 san55	说三	‘to gossip’	w	Non-velar	Source syl.	1
ʂan35 two214	闪躲	‘to evade’	w	Non-velar	Source syl.	2
tsʰan55 ʈʂwo55	餐桌	‘dinner table’	w	Non-velar	Source syl.	2
tsʰwen35 ʈʂə35	存折	‘deposit book’	w	Non-velar	Target syl.	1
ʈʂʰwen55 sə51	春色	‘spring scenery’	w	Non-velar	Target syl.	1
ʂə51 tʰwan35	社团	‘student club’	w	Non-velar	Target syl.	2
tsʰə51 ʂwan55	侧栓	‘side bolt’	w	Non-velar	Target syl.	2
tsʰwej51 ʂan55	翠山	‘green hill’	w	Non-velar	No	1
ʈʂʰwan55 tsʰaj51	川菜	‘Szechuan cuisine’	w	Non-velar	No	1
ʂan55 swen214	山笋	‘bamboo shoot’	w	Non-velar	No	2
tsɑŋ55 ʂwej214	脏水	‘slander’	w	Non-velar	No	2

/ɥ/ *test items:*

Test item	Simpl. Chinese	English	Glide	Consonant place	Vowel alternation	Glide position
ɕɥe55 lan35	薛岚	‘Xuelan (name)’	y	Palatal	Both syl.	1
tɕʰɥe51 lan35	雀兰	‘Quelan (brand)’	y	Palatal	Both syl.	1
ɕɥɛn35 lə	悬了	‘unlikely’	y	Palatal	Both syl.	1
ɕɥɛn55 tə35	宣德	‘Xuande (emperor)’	y	Palatal	Both syl.	1
tɕʰɥʊŋ35 lan214	穷览	‘to read exhaustively’	y	Palatal	Both syl.	1
tɕɥʊŋ21 tan51	囧蛋	‘awkward person’	y	Palatal	Both syl.	1
tɕɥɪn55 lan35	君兰	‘orchid’	y	Palatal	Both syl.	1
tʰə51 tɕʰɥɛn35	特权	‘privilege’	y	Palatal	Both syl.	2
lan21 ɕɥʊŋ35	懒熊	‘lazy bear’	y	Palatal	Both syl.	2
lan35 tɕʰɥʊŋ35	蓝穹	‘blue dome (sky)’	y	Palatal	Both syl.	2
lə51 tɕʰɥɛn35	乐泉	‘Spring of Happiness’	y	Palatal	Both syl.	2
nan35 tɕʰɥe51	难却	‘difficult to let go’	y	Palatal	Both syl.	2
nan35 tɕɥe35	男爵	‘baron’	y	Palatal	Both syl.	2
lan35 tɕʰɥɪn35	蓝裙	‘blue dress’	y	Palatal	Both syl.	2
lɥe51 nan35	略难	‘slightly difficult’	y	Non-palatal	Both syl.	1
tan21 lɥe51	胆略	‘courage and resourcefulness'	y	Non-palatal	Both syl.	2

Filler items:

Test item	Simplified Chinese	English	Glide
xaj21 nan35	海南	‘Hainan’	None
kʰa55 fej55	咖啡	‘coffee’	None
tan51 kaw55	蛋糕	‘cake’	None
fen51 nej51	份内	‘one’s duty’	None
pan55 ʈʂɑŋ214	班长	‘class president’	None
tan55 tsʰɤŋ35	单层	‘single layer’	None
pʰa51 law35	帕劳	‘Palau’	None
xɤŋ35 saw214	横扫	‘clean sweep’	None
ma21 laj35	马来	‘Malay’	None
taw51 kʰow214	道口	‘crossroad’	None
xaj35 tsaw214	海藻	‘seaweed’	None
kʰə21 ɻan35	可燃	‘inflammable’	None
taj51 ʈʂʰan214	待产	‘expectant (mother)’	None
taj51 kʰə51	代课	‘to substitute to teach’	None
pɑŋ21 tan55	榜单	‘chart’	None
maj51 lɑŋ51	麦浪	‘wheat field’	None
kʰow21 ʈʂaw51	口罩	‘mask (medical)’	None
pʰaj35 xan51	排汗	‘to sweat’	None
paw21 ʈʂɑŋ51	保障	‘guarantee’	None
man21 tʰɑŋ35	满堂	‘full house’	None
lɤŋ21 ʈʂan51	冷战	‘Cold War’	None
tʰan51 ʈʂɑŋ214	探长	‘chief inspector’	None
kaw55 ʈʂʰan214	高产	‘prolific’	None
pʰaw21 tʰɑŋ35	跑堂	‘waiter’	None
xɤŋ55 ʈʂʰɑŋ51	哼唱	‘to hum’	None
pʰɑŋ35 ɻan35	庞然	‘gigantic’	None
fa35 tan55	罚单	‘fine ticket’	None
xaj21 paw51	海报	‘poster’	None
men35 kʰow214	门口	‘doorway’	None
taj51 kow55	代沟	‘generation gap’	None

Test item	Simplified Chinese	English	Glide
ʂan55 fɤŋ55	山峰	‘mountain peak’	None
ma21 fɤŋ55	马蜂	‘wasp’	None
p^hen35 tsaj55	盆栽	‘bonsai’	None
lan35 sə51	蓝色	‘blue’	None
ʈʂow55 san55	周三	‘Wednesday’	None
saj51 $ʈʂ^h$ə55	赛车	‘car racing’	None
ʂa55 sɤŋ55	沙僧	‘Sha Wujing’	None
sɑŋ55 ʂen51	桑葚	‘mulberry’	None