Variation and gradience in non-standard Turkish affix ordering

Eva Neu*

Abstract. Semantically vacuous affix order variation is a cross-linguistically rare phenomenon that promises to shed light on the principles underlying affix order in general. This paper discusses a case of free affix order in Turkish: if a verb contains more than one TAM morpheme, any or all of them can be followed by an agreement morpheme. I seek to determine which factors influence the acceptability of a given order for a given set of morphemes by collecting new data from 19 native speaker consultants. The results reveal a sharp bifurcation between some orderings that are perceived as categorically grammatical by all speakers and others that result in gradient and variable judgments without any discernible patterns. I discuss these findings against the background of previous approaches to affix order variation, highlighting the extent to which analyses are influenced by methodological choices about data collection.

Keywords. affix order; variation; agreement; morphology; Turkish

1. Background.

1.1. PREVIOUS APPROACHES TO AFFIX ORDER VARIATION. A long-standing debate in morphology concerns the principles that determine the relative order of affixes on a word (see Manova & Aronoff 2010; Rice 2011 for an overview). A particularly promising, yet understudied window into the nature of affix order is offered by cases of ordering variation. This paper focuses on a cross-linguistically rare subtype of variable affix order in which all possible orderings are assigned the same interpretation. In previous work, two main approaches to semantically vacuous affix order variation have emerged, which I briefly outline in the following.

In Chintang (Sino-Tibetan), prefixes can appear in any order, shown for two prefixes in (1):

(1) a. a-mai-ep-t-e.
2-NEG-get.up-NEG-PST
b. mai-a-ep-t-e.
NEG-2-get.up-NEG-PST
'You didn't get up.' (Bickel et al. 2007:57)

Bickel et al. (2007) have argued that Chintang prefixes select not for a morphologically defined constituent but for a prosodic word, and that both the stem and the prefixes constitute independent prosodic words. The prosodic structure of (1) is shown in (2); note that vowel-initial prosodic words surface with a glottal stop to fulfill an onset requirement.

(2) a. (ω ʔa) (ω mai) (ω ʔepte)
b. (ω mai) (ω ʔa) (ω ʔepte) (Bickel et al. 2007:57)

A prefix that selects for a prosodic word can variably combine with the stem or any other prefix. Under this analysis, free affix order in Chintang is thus the result of prosodic subcategorization.

A different approach has been developed by Ryan (2010) for variable affix order in Tagalog, which has an aspectual marker realized as a reduplicant prefix that can occupy various positions

* Thanks to Faruk Akkuş, Joe Pater, Kyle Johnson, the audience at the 2025 LSA, Betülay Aras and my informants.
Authors: Eva Neu, University of Massachusetts Amherst (eneu@umass.edu).

in the word without any changes to the semantics. Different orderings are attested with different frequencies. To analyze these data, Ryan uses weighted morphotactic bigram constraints that encode the strength of coherence of a pair of morphemes compared to other bigrams in the language. A high weight assigned to such a bigram constraint corresponds to a high probability that the two morphemes appear consecutively in a word. By way of example, the bigram constraint X-Y (3) is satisfied by the strings in (3a) and (3b) but violated by those in (3c) and (3d):

(3) X-Y
  a. X-Y-Z
  b. Z-X-Y
  c. *X-Z-Y
  d. *Y-X-Z

These bigram constraints are put to work in a MaxEnt grammar, a framework in the broader OT tradition that generates a probabilistic output for every input. A simplified example tableau is given in (4). Each constraint is assigned a weight; here, X-Y has a weight of 2, Y-X of 1. For each input – here, the collection of morphemes {X, Y, Z} –, all possible candidates, that is, possible orderings, are generated. For simplicity's sake, (4) lists only three possible orderings: X-Y-Z, Z-Y-X and X-Z-Y.

(4)

		X-Y	Y-X	H	eH	Prob.
		2	1			
{X, Y, Z}	X-Y-Z	0	1	-1	0.0368	0.67
	Z-Y-X	1	0	-2	0.1353	0.24
	X-Z-Y	1	1	-3	0.0498	0.09

The probability of a candidate is computed based on the constraints it violates. In (4), the candidate X-Y-Z complies with the constraint X-Y but incurs one violation of Y-X, and vice versa for Z-Y-X. The candidate X-Z-Y violates both constraints. For every candidate, each violation is multiplied by the weight of the constraint violated and summed up. The negative sum is this candidate's harmony score, abbreviated here as H. From this, the eHarmony (eH) value is derived by raising the Harmony score to e. Finally, the probability of each candidate is obtained by dividing its eHarmony score by the sum of all eHarmony scores.

As a result, the weight assigned to a bigram constraint maps onto the probability that the pair of morphemes in question will appear adjacent to each other in the specified order. Ryan shows that a MaxEnt grammar equipped with such bigram constraints can closely match the frequencies of different orderings attested in Tagalog corpus data. Under this analysis, speakers' knowledge about the acceptability of different affix orderings hence simply consists in knowledge of surface probabilities of bigrams.

Following up on Bickel et al. (2007), Mansfield (2015) and Dąbkowski (2022) have proposed that free affix order in Murrinhpatha and Paraguayan Guaraní, respectively, is equally due to prosodic subcategorization. Newbold (2013), on the other hand, has developed an account of variable affix order in Kuna along the lines of Ryan (2010). Against this background, Dąbkowski (2022) has argued that variable affix order cross-linguistically falls into two distinct typolological classes, one being prosodically, the other morphotactically conditioned.

1.2. VARIABLE AFFIX ORDER IN TURKISH. The present paper adds to the debate on semantically vacuous affix order variation by presenting new empirical data from Turkish, a heavily agglutinative language. In verbs with more than one TAM (tense/aspect/mood) morpheme, the agreement morpheme can either surface in word-final position (5a) or between the TAM morphemes (5b). Alternatively, the verb can contain two agreement morphemes, each of which follows a TAM morpheme (5c) (Good & Yu 1999, 2005; Güneş 2020, 2021).[1]

(5) a. gel-di-yse-*k*
come-PST-COND-1PL
root-TAM-TAM-Agr *Final agreement*
'we would come'

b. gel-di-*k*-se
come-PST-1PL-COND
root-TAM-Agr-TAM *Medial agreement*

c. gel-di-*k*-se-*k*
come-PST-1PL-COND-1PL
root-TAM-Agr-TAM-Agr *Double agreement*

Following Güneş (2020, 2021), I will refer to the three possible ordering patterns as final (5a), medial (5b) and double (5c) agreement. No difference in interpretation between the three orderings has been reported. For verbs with three TAM morphemes, agreement can equally surface after any or all TAM morphemes; for simplicity's sake, I omit the relevant data. Agreement morphemes are not licensed directly after the root or any non-TAM affix (but see footnote 2).

To complicate the picture, agreement morphology in the Turkish verbal domain surfaces in three different paradigms, summarized in Table 1. The choice of paradigm depends on the preceding TAM morpheme; e.g., past tense *-DI* and conditional *-sE* in (5) both must be followed by an agreement morpheme from the *k*-paradigm. In the following, I will gloss agreement morphemes from the *k*-, *z*- and *rz*-paradigm as Agr_k, Agr_z and Agr_{rz}, respectively, and I will use the labels TAM_k, TAM_z and TAM_{rz} for TAM morphemes that have been reported by Güneş (2020, 2021) to precede the corresponding agreement paradigm.

	k-paradigm (Agr_k)	*z*-paradigm (Agr_z)	reduced *z*-paradigm (Agr_{rz})
1SG	*-m*	*-(y)Im*	*-m*
2SG	*-n*	*-sIn*	*-n*
3SG	∅	∅	∅
1PL	*-k*	*-(y)Iz*	*-z*
2PL	*-nIz*	*-sInIz*	*-nIz*
3PL	*-lEr*	*-lEr*	*-lEr*

Table 1. Agreement paradigms in the Turkish verbal domain

[1] Examples follow the Leipzig glossing conventions. Note that some verbs cited in this paper feature a glide (*-y*) between the TAM markers, commonly analyzed as a copula (Güneş 2021) or auxiliary (Sezer 2001) that is always present underlyingly between two TAM morphemes but only spelled out in certain phonologically defined contexts. Since it is orthogonal to our purposes, I do not gloss the copula separately.

So far, the data presented have suggested that the three possible orderings – final, medial and double agreement – are always available for any set of morphemes and for any speaker. But this is arguably not the case. Previous work has argued that the position(s) in which an agreement morpheme can appear depend on its paradigm, specifically, that the *z*-paradigm can only surface word-finally (Good & Yu 1999, 2005; Güneş 2020, 2021). The examples in (6) illustrate this claim with two TAM morphemes that need to be followed by the *z*-paradigm, progressive *-Iyor* and evidential *-mIş*.

(6) a. gel-iyor-muş-*sun*
come-PROG-EVID-2SG
root-TAM$_z$-TAM$_z$-Agr$_z$ *Final agreement*
'you (sg.) are apparently coming'

b. *gel-iyor-*sun*-muş
come-PROG-2SG-EVID
root-TAM$_z$-Agr$_z$-TAM$_z$ **Medial agreement*

c. *gel-iyor-*sun*-muş-*sun*
come-PROG-2SG-EVID-2SG
root-TAM$_z$-Agr$_z$-TAM$_z$-Agr$_z$ **Double agreement*

The reported contrast between (5) and (6) relates to a broader discussion about the syntax of Turkish agreement. Specifically, the data in (6) have been taken as evidence for the claim that the three paradigms, while superficially similar, differ in their underlying syntactic structure (see also Bobaljik 2000; Kelepir 2001; Kornfilt 1996, but see Neu 2024). I briefly return to this issue in Section 3, but it will turn out to be tangential to our purposes.

No other restrictions on the availability of the three different orderings have been reported in previous research. Prescriptively, final agreement is the norm, with the exception of the 3PL morpheme *-lEr*, which is standardly realized in medial position. The question remains whether non-3PL medial agreement and double agreement are categorically acceptable, whether they are subject to a constraint against non-final Agr$_z$ as outlined above, whether they show inter-speaker variation, whether they are licensed only for specific TAM morphemes, roots or ϕ-feature combinations, and so on. Clarifying this empirical picture is fundamental to any analysis of Turkish affix order variation.

Hence, the present paper seeks to establish what determines which of the three ordering patterns are licensed for a given set of morphemes. In particular, I ask whether there is indeed a restriction against non-final Agr$_z$ and whether any other patterns are systematically unavailable. To this end, I collected new data, as described in Section 2. The empirical findings are summarized in Section 3. In Section 4, I then discuss the results against the background of previous work on affix order variation, in particular Bickel et al. (2007) and Ryan (2010). Section 5 concludes.

2. Methods.

2.1. PROCEDURE AND MATERIALS. Collecting data on variable affix order posed a methodological challenge. Medial and double agreement – with the exception of medial 3PL agreement – are considered colloquial and substandard. As a result, they are largely confined to spoken language and thus cannot be investigated using corpus data. What is more, non-3PL medial agreement and double agreement might be routinely rejected on prescriptive grounds by participants in a standard acceptability judgment experiment. At the same time, Güneş (2020, 2021) reports

dialectal variation in the data, especially with respect to double agreement; hence, working with only one or two speakers would give us a very incomplete picture.

The solution I settled on was to conduct one-hour fieldwork interviews with 21 native speakers of Turkish, which made it possible to collect data from a reasonably large number of informants while still being able to engage with them directly and in more depth than in an experimental setting. The goal was to elicit grammaticality intuitions about verbs with final, medial and double agreement, respectively. To make the task more natural, all verbs were used in the context of a complete sentence, and each sentence was presented three times, once for each ordering pattern. An example of such a set of items is given in (7):

(7) a. Buraya kadar *gel-di-yse-k* bir çayınızı içeriz.
here so-far come-PST-COND-1PL a tea drink
'Now that we have come here, let us have a cup of your tea.'

b. Buraya kadar *gel-di-k-se* bir çayınızı içeriz.
here so-far come-PST-1PL-COND a tea drink

c. Buraya kadar *gel-di-k-se-k* bir çayınızı içeriz.
here so-far come-PST-1PL-COND-1PL a tea drink

While most items contained monosyllabic high-frequency roots such as *gel* 'come,' *gör* 'see,' *bul* 'find' or *koş* 'run,' we also tested with some speakers a few longer and lower-frequency roots such as *deneyimle* 'experience,' *güncelle* 'update' or *barış* 'reconcile.' The results were not noticeably different from those for short high-frequency roots; hence, I do not report them separately in the following.

The interviews took place remotely over Zoom and were conducted in Turkish with the help of a Turkish-speaking research assistant unless the informant was fully fluent and comfortable in English. All items were pronounced by the RA, who was a native speaker. Given the considerable amount of inter-speaker variation with respect to affix order variation, we found it more helpful to ask informants specifically whether they would use a given verb form themselves and/or whether others used it, rather than whether they found a form grammatical. In the same interviews, we also collected intuitions on a number of other data points, not reported in the following.

To reduce noise in the data, we had prepared a small set of clearly ungrammatical sanity check items which we used for two speakers who had indiscriminately accepted all sentences they were presented. One of them accepted all, the other several of the sanity check items; hence, their data were discarded, leaving us with a sample size of 19. No such issues emerged with the other participants.

It bears highlighting that the data thus collected are strictly qualitative in nature; they are fieldwork results, not experimental findings. The number of intuitions collected for each individual item is low overall and no efforts were made to counterbalance or randomize the items presented.

2.2. PARTICIPANTS. We recruited a diverse set of informants in terms of age, gender, socioeconomic and geographic background with the aim of collecting judgments from as many different varieties of Turkish as possible. Eight informants (42%) were male, eleven (58%) female, and ages ranged from 18 years to 62 years, with an average of 37 years. Four speakers were bi- or multilingual. Four of the informants were professional linguists who grew up in Turkish but lived

in the US by the time of the interview while the others were naive speakers who lived, and had always lived, in Turkey. No systematic differences between the judgments of naive and linguistically trained speakers emerged. Additional demographic data can be found in the appendix. In the following, informants will be referred to by pseudonyms.

3. Findings. The main finding that emerged from the data was a clear bifurcation between judgments on final agreement and 3PL medial agreement on the one hand side – forms that are prescriptively considered the norm – and judgments on non-3PL medial agreement and double agreement on the other. The former examples were invariably and without qualifications found perfectly grammatical. While the 3PL morpheme *-lEr* was considered more natural in medial position, it was not rejected in final position either. In contrast, speakers' intuitions about non-3PL medial agreement and double agreement were considerably more nuanced. Almost all speakers accepted a subset of such forms, but both the size and the extension of this subset varied radically. Acceptability of medial and of double agreement for a given form were correlated but neither entailed the other. Moreover, judgments could often not be clearly categorized as either grammatical or ungrammatical. Speakers frequently gave intermediate judgments, or reported that they themselves did not use a given form but knew that others did, or vice versa. Occasionally, informants pointed out that they would use a given form as a joke, or volunteered sociolinguistic and prescriptive judgments. This overall picture of pervasive variation held up for all informants, including linguistically trained speakers.

Unexpectedly, I did not replicate the pattern reported by Güneş (2021, 2020) and Good & Yu (2005, 1999) that Agr_z can only appear in final position. Two examples which were accepted by some speakers are given below; (8a) was explicitly called 'normal' by informant E8, (8b) 'beautiful' by informant S9:

(8) a. gel-ecek-siniz-di
come-FUT-2PL-PST
root-TAM_z-Agr_z-TAM_z
'you (pl.) will have come'

b. bul-uyor-sun-muş
find-PROG-2SG-EVID
root-TAM_z-Agr_z-TAM_z
'you (sg.) are apparently finding'

Acceptance rates for Agr_z in non-final position varied considerably between speakers. For medial agreement, informant E8 judged 5/6 forms to be acceptable or reported that others would use them, informant S9 7/8, informant T16 6/10 and informant B20 5/10. While double agreement was rejected more often, informant E8 accepted 3/6 and informant M19 4/12 forms. Other informants accepted Agr_z in non-final position only in one or two instances, or not at all. It is correct that in some cases, changing an Agr_z morpheme in non-final position to Agr_{rz} improved acceptability (9):

(9) a. *bul-*uyor-uz*-muş-uz
find-PROG-1PL-EVID-1PL
root-TAM_z-Agr_z-TAM_z-Agr_z
'we are apparently finding'

b. bul-*uyo-z*-muş-uz
find-PROG-1PL-EVID-1PL
root-TAM_{rz}-Agr_{rz}-TAM_z-Agr_z
'we are apparently finding'
(Informant: T16)

However, judgments could also be reversed (10):

(10) a. *koş-*uyo*-*z*-muş
run-PROG-1PL-EVID
root-TAM$_{rz}$-Agr$_{rz}$-TAM$_z$
'we are apparently running'

b. koş-*uyor*-*uz*-muş
run-PROG-1PL-EVID
root-TAM$_z$-Agr$_z$-TAM$_z$
'we are apparently running'
(Informant: M19)

Overall, medial and double agreement with Agr$_z$ in non-final position is subject to considerable inter-speaker variation but far from being unattested, in contrast to what has been reported previously. As noted earlier, the claim that the three agreement paradigms differ from each other in terms of their ordering properties has been taken to corroborate the claim that they equally differ in their underlying syntax (Good & Yu 1999, 2005; Güneş 2020, 2021, see also Bobaljik 2000; Kelepir 2001; Kornfilt 1996). The finding reported here does not constitute evidence against this analysis but it does cast doubt on one of the arguments in its favor. I have outlined an alternative approach in Neu (2024); the present paper is not the space to discuss it further.

More in general, I could not identify any categorical necessary or sufficient conditions on the acceptability of medial and double agreement, with the sole exception noted already that medial agreement is systematically grammatical with 3PL morphemes. Speakers' judgments were affected by small changes in the verb form. For instance, changing the person/number features of the agreement morpheme (11) or the root (12) could result in different judgments:

(11) a. *gid-iyo-*z*-du
leave-PROG-1PL-PST
'we were leaving'

b. gid-iyo-*n*-du
leave-PROG-2SG-PST
'you (sg.) were leaving'
(Informant: E8)

(12) a. **bul*-acağ-ız-dı
find-FUT-1PL-PST
'we will have found'

b. *gel*-eceğ-iz-di
come-FUT-1PL-PST
'we will have come'
(Informant: T16)

However, these effects would not generalize: informant E8 did not systematically prefer medial agreement forms with 2SG over those with 1PL, or informant T16 with *gel* over *bul*. For other forms, the contrast could even be reversed.

A few factors, while still not clear-cut, were more systematic. The 3PL morpheme, always accepted in medial position, also tended to make double agreement more acceptable. While far from being exceptionless, this pattern was robustly attested for many informants, exemplified by the minimal pairs in (13) and (14):

(13) a. gel-ecek-*ler*-di-*ler*
come-FUT-3PL-PST-3PL
'they will have come'

b. *gel-ecek-*siniz*-di-*niz*
come-FUT-2PL-PST-2PL
'you (pl.) will have come'
(Informant: S9)

(14) a. gör-dü-*ler*-se-*ler*
see-PST-3PL-COND-3PL
'if they saw'

b. *gör-dü-*nüz*-se-*niz*
see-PST-2PL-COND-2PL
'if you (pl.) saw'
(Informant: O2)

Similarly, medial agreement seemed to be accepted more readily after some TAM markers than others: by and large, speakers preferred medial agreement after *-DI*, *-Iyo* and *-Iyor* over medial

agreement after *-sE*, *-EcE* and *-EcEk*. Example (15) gives such a minimal pair for the contrast between the morpheme sequence *-DI-sE* and *-sE-DI*, example (16) between *-Iyo-mIş* and *-EcE-sE*:

(15) a. gel-*di*-k-*se*
come-PST-1PL-COND
'if we came'
b. *gel-*se*-k-*ti*
come-COND-1PL-PST
'if we had come'
(Informant: E14)

(16) a. bul-*uyo*-z-*muş*
find-PROG-1PL-EVID
'we are apparently finding'
b. *bul-*aca*-z-*sa*
find-FUT-1PL-COND
'if we will find'
(Informant: M19)

However, such observations are only anecdotal in nature. As noted above, the data collected here are purely qualitative and do not allow for probabilistic inference.

Finally, there is one more odd finding to mention. In previous literature, the three possible orderings have always been described as perfectly synonymous, and for most of the items I tested I have no reason to doubt this claim. In a handful of cases, however, informants pointed out that two forms received a slightly different interpretation, but in none of these cases could they fully articulate the meaning difference in question. These observations were made by different speakers for different items and, to the extent that they became clear, about different semantic contrasts. All attempts to replicate these contrasts with similar items – e.g., by using a different root – failed; the meaning differences seemed highly item-specific. I have nothing to say about why these effects would arise; I am reporting them here merely for completeness' sake.

To sum up, final and 3PL medial agreement is categorically grammatical, while judgments on the remaining forms vary, without any categorical effects. The data suggest the presence of probabilistic patterns, which would need to be backed up by more robust quantitative evidence. In the next section, I discuss whether the Turkish data are amenable to one of the two previous approaches to affix order variation discussed earlier, Bickel et al. (2007) and Ryan (2010), starting with the former.

4. Discussion.

4.1. PROSODIC SUBCATEGORIZATION. To recap briefly, Bickel et al. argue that in Chintang, prefixes select for a prosodically defined constituent, specifically, a prosodic word. Since both the stem itself and each prefix are independent prosodic words, prefixes can appear in any order. For this analysis to apply to Turkish, both the root together with the first TAM morpheme and the second TAM morpheme in isolation would thus have to constitute independent prosodic units of the same size (17):

(17) a. ($_\omega$ root-TAM) ($_\omega$ TAM) ($_\omega$ Agr)
b. ($_\omega$ root-TAM) ($_\omega$ Agr) ($_\omega$ TAM)
c. ($_\omega$ root-TAM) ($_\omega$ Agr) ($_\omega$ TAM) ($_\omega$ Agr)

But this is not the case. In the following, I only sketch out the relevant facts; see Güneş (2021) for a more thorough analysis of the prosodic structure of Turkish verbs. Relying on stress as the main diagnostic for prosodic wordhood, Güneş argues that verbs with more than one TAM morpheme are parsed as two prosodic words, such that primary stress is assigned at the right edge of the first prosodic word and secondary stress at the right edge of the second. In the case of final

agreement, the first prosodic word is formed by the root and the first TAM morpheme, and the second prosodic word by the remaining affixes (18). A BINMAX constraint ensures that a syntactic word maps onto a maximum of two prosodic words.

(18) gel-iyor-muş-sun → (ω gel-iyor) (ω muş-sun)
come-PROG-EVID-2SG
'you (sg.) are apparently coming'

In the case of medial and double agreement, the non-final agreement morpheme is realized as part of the preceding prosodic word if it is a single segment (19a). If it is syllabic, it is variably either integrated into the preceding or the following prosodic word (19b).

(19) a. gel-di-k-se → (ω gel-di-k) (ω se)
come-PST-1PL-COND
'we would come'

b. gel-di-niz-se → (ω gel-di-niz) (ω se) *or* (ω gel-di) (ω niz-se)
come-PST-2PL-COND
'you (pl.) would come)

In short, the agreement morpheme does not consistently combine with a constituent of a certain prosodic size. Overall, I am not aware of any evidence that the second TAM morpheme constitutes a prosodic constituent of the same type – e.g., a prosodic word – as the root together with the first TAM morpheme, such that the agreement morpheme could select for this prosodic constituent.

While Turkish affix order variation thus does not appear to be prosodically conditioned, there is a very basic sense in which it can be analyzed along the same lines as Chintang, namely, as the result of a rule that yields more than one output. In Chintang, Bickel et al. argue, prefixes select for a prosodic word, and if there is more than one prosodic word, the prefix can be freely ordered. In Turkish, one might argue, agreement affixes must follow a TAM marker, and if there is more than one TAM marker, the affix can be freely ordered.[2] In a sense, the logic is straightforward, regardless of whether it is implemented by locating the variation in the position occupied by the agreement projection in the clausal spine or the merging site of a dissociated agreement morpheme (Güneş 2021), by appealing to head movement, a purely morphotactic copy-paste operation (Güneş 2020) or paradigm-based word formation rules.

But there is a more general problem, which is that such an account has no handle on the gradient judgments and inter-subject variation observed earlier for medial and double agreement. By itself, it predicts that final, medial and double agreement are invariably licensed under all circumstances, but this is not what we observed. Bickel et al.'s proposal is a story about why affix order variation occurs in the first place, not why some patterns are more felicitous than others. Hence, I now turn to Ryan (2010), who explicitly tries to model different frequencies of different orderings.

[2] I gloss over an exception to this generalization, which is that the polar question marker *-mI* in some forms surfaces between the TAM and the agreement morpheme (i):

(i) gel-iyor-mu-sunuz
come-PROG-Q-2PL
'are you (pl.) coming?'

4.2. BIGRAM CONSTRAINTS. As outlined earlier, Ryan (2010) analyzes free affix ordering in Tagalog by means of bigram constraints that encode the relative strength of coherence of a pair of morphemes. These bigram constraints are implemented in a MaxEnt grammar, which generates a probability distribution over the candidates. Unlike Bickel et al. (2007), Ryan is thus not concerned with explaining why variable affix order is possible in certain cases, but with deriving the different frequencies of the possible patterns. This makes it, at first sight, particularly suited for analyzing the Turkish data, which are equally subject to gradience and fine-grained variation.

However, MaxEnt is a framework that deals in frequencies, not acceptability judgments, as collected here for Turkish. The probability distributions generated do not straightforwardly translate into intuitions about grammaticality or acceptability. What is more, the present data are only qualitative, too small-scale to confirm or disconfirm predictions about probability distributions. Hence, I will not attempt to implement a full MaxEnt analysis for Turkish. Nonetheless, I will make some general remarks to the effect that the Turkish data raise questions for a bigram approach to affix order variation.[3]

To begin with, Ryan makes the assumption that all possible permutations of relevant morphemes serve as input to the constraint set. That is, if the word in question consists of three morphemes, six different candidates will be evaluated by the bigram constraints. For Turkish, however, this raises problems. Some pairs of TAM morphemes such as past *-DI* and conditional *-sE* can occur in either order, with the different orderings giving rise to different interpretations (20):

(20) a. gel-*di*-*yse*-k
come-PAST-COND-1PL
'if we came'

b. gel-*se*-*ydi*-k
come-COND-PAST-1PL
'if we had come'

Under Ryan's approach, examples (20a) and (20b) would compete with each other as different outputs for the same input. Since the constraints as assumed so far cannot see the semantics of the candidates, one of the two orders in (20) would emerge victorious for both interpretations, contrary to fact. We might respond to this problem by making the constraint set in some way sensitive to the semantics of the morpheme sequence evaluated. But more generally, affix order variation in Turkish is highly restricted, with most orderings being fully ungrammatical and never attested. It is questionable whether the grammar truly generates fully ungrammatical orderings such as (21) before filtering them out by means of constraints:

(21) *iyor-sunuz-gel-di
PROG-2PL-come-PST
TAM-Agr-root-TAM

Hence, another possible strategy would be to restrict the set of candidates submitted for evaluation. In constraint-based architectures such as OT and MaxEnt, the two central components of the grammar are GEN, a function that generates candidates, and EVAL, a function that selects the winning candidate based on a set of constraints. By and large, the explanatory burden for grammatical patterns has been placed on EVAL, with Prince & Smolensky (1993/2004) denying that GEN is subject to any restrictions and instead arguing that it blindly applies all possible operations to the input set of lexical items. However, others have proposed that the computational

[3] One of the issues I need to sidestep here is that fact that the Turkish data have two peculiar properties not present in Tagalog, namely, multiple exponence of a single morpheme and allomorphy. How bigram constraints would handle these patterns is an open question.

power of GEN must be limited (see e.g., Blaho et al. 2007; Uffmann 2007). In particular, work in OT syntax typically relies on well-defined, rule-based syntactic operations in GEN in conjunction with the constraint set CON (e.g., Arregi & Nevins 2012; Grimshaw 1997; Wolf 2008; Wurmbrand 2008, 2018). For Turkish, we might thus restrict GEN such that it produces only candidates in which the agreement morphemes surface after TAM markers.

Finally, a related problem that remains is that the data reported above have revealed a sharp bifurcation between categorical judgments for final and 3PL medial agreement on the one hand and variable and gradient judgments for the remaining forms on the other. While a MaxEnt grammar might be able to assign the highest probability to the final agreement forms, it cannot model the contrast between categorical and gradient judgments since it is simply not a model of grammaticality judgments but of probabilities. In other words, the results summarized above suggest that the two sets of forms have a different grammatical status, but MaxEnt can only treat them as more or less frequent. In consequence, even if a bigram model could successfully match the different frequencies of the different affix orderings in Turkish, a fundamental aspect of the data would still remain unaccounted for.

4.3. TOWARDS A SYNTHESIS. To give an intermediate summary, neither the prosodically driven approach to affix order variation in Bickel et al. (2007) nor the morphotactic bigram model in Ryan (2010) extends straightforwardly to the Turkish data. Bickel et al.'s analysis, besides other challenges, has no handle on the inter-speaker variation and gradience that I documented for medial and double agreement forms. Ryan, on the other hand, has nothing to say about the categorical patterns in the Turkish data – the fact that some logically possible morpheme orderings are strictly ungrammatical and never surface, but also that final and 3PL medial agreement are invariably grammatical.

All of this is not necessarily due to an intrinsic difference between Chintang, Tagalog and Turkish affix order variation, but rather to different research questions and different kinds of data collected. Bickel et al. set out to solve a typological mystery – why is free affix ordering, while cross-linguistically rare, attested in Chintang? –, and they do so by drawing on small-scale fieldwork data from a low-resource language. In fact, they do observe variation in the judgments: 'not all informants accept all logically possible orderings [...], but judgments vary widely and, as far as we can tell, not systematically' (Bickel et al. 2007:45) – just as reported above for Turkish. But the goal of Bickel et al.'s paper is not to model this variation but to explain why affix order in Chintang would be free in the first place.

In contrast, while Ryan's (2010) bigram model closely matches the corpus frequencies of different affix orderings, he never answers the very basic question of why the Tagalog reduplicant morpheme can surface in different positions at all – which is a question worth asking, given that cross-linguistically, and even within Tagalog, the vast majority of affixes are fixed in their orderings. To frame matters differently, since bigrams are not considered particularly successful models of linguistic knowledge in other domains, the question arises what makes them perform so well for Tagalog affix order. Furthermore, the contrast between categorical and gradient judgments observed above for Turkish might obtain in Tagalog as well; since Ryan relies on frequency counts and not on grammaticality judgments, it is hard to say.

The Turkish data I collected for the present paper, in a somewhat unusual medium-scale format, sit in between Bickel et al.'s fieldwork findings and Ryan's corpus counts. They are sufficiently detailed to clearly document inter-speaker variation and contrasts between categorical

and gradient judgments, and even to suggest probabilistic patterns, yet not detailed enough to allow for quantitative methods. The general picture in all three languages may ultimately be very similar (contra Dąbkowski 2022), and to the extent that Bickel et al.'s and Ryan's analyses fail to account for Turkish, they might equally do so for the languages for which they were developed. A complete analysis would need to explain why affix order variation arises, why some of the logically possible variants are more acceptable and more frequent than others, and why speakers' intuitions are gradient and variable for some variants but categorical and stable for others.

In the discussion above, I have begun to sketch out a possible synthesis of the two accounts that could deal with the Turkish data. One component of the grammar might generate all forms in which an agreement morpheme follows a TAM marker, and another component might evaluate these forms using a set of violable constraints that yield gradient results. In addition, a substantive bias of sorts might ensure that final agreement is always grammatical, in line with the cross-linguistic tendency for Agr to be realized syntactically higher than, and morphologically outside of, TAM heads. Overall, affix order in the Turkish verbal domain appears to be subject to a multitude of pressures, some general and others item-specific, some violable and others categorical: to place an agreement affix after a TAM marker, to place an agreement affix in word-final position, to reproduce frequently observed sequences of morphemes, and so forth. Grammatical knowledge about licit orderings may be fluid even for adult speakers, who continue to look for patterns in the input, extrapolate rules and form analogies.

5. Conclusion. This paper has been concerned with affix order variation in Turkish, asking what determines which of the three possible orderings – final, medial and double agreement – are available for a given set of morphemes. I found that final agreement as well as 3PL medial agreement are categorically grammatical, whereas non-3PL medial agreement and double agreement give rise to gradient judgments and bother inter- and intra-speaker variation with no discernible patterns. I could not replicate the claim made in earlier work that agreement morphemes from the *z*-paradigm are confined to word-final position (Good & Yu 1999, 2005; Güneş 2020, 2021).

I have then discussed these findings against the background of two previous approaches to variable affix order, Bickel et al. (2007) and Ryan (2010). I have argued that Turkish is not amenable to the prosodic analysis of Bickel et al., but also presents challenges for Ryan's bigram model. In particular, an account along the lines of Bickel et al. predicts that all orderings should be equally attested and acceptable and has nothing to say about the gradient and variable judgments for Turkish medial and double agreement forms. On the other hand, Ryan's analysis might be able to match the frequency distribution of the different orders but cannot explain why some forms but not others are categorically grammatical for all speakers.

Finally, perhaps the key upshot of this paper has been a methodological one. In my discussion of Bickel et al. and Ryan, I have highlighted that the difference between their analyses might stem not so much from a difference between Chintang and Tagalog affix order variation as such, but rather between the types of data they draw on. The data I collected for this project straddle the line between in-depth judgments elicited from a few consultants and quantitative experimental or corpus data, and in consequence, it has revealed discontents with both previous proposals. Without corpus frequencies or other quantitative data, linguistic variation without clear patterns must remain unexplainable noise, but frequency counts do not always reveal important differences in judgments. Intuitions without numbers are lame, numbers without intuitions are blind, and how data are collected often already limits the space of possible analyses. For future work,

relying on a broad range of different kinds of data might give us a more complete picture of affix order variation.

References

Arregi, Karlos & Andrew Nevins. 2012. *Morphotactics: Basque auxiliaries and the structure of spellout*, vol. 86. Springer Science & Business Media. https://doi.org/10.1007/978-94-007-3889-8.

Bickel, Balthasar, Goma Banjade, Martin Gaenszle, Elena Lieven, Netra Prasad Paudyal, Ichchha Purna Rai, Manoj Rai, Novel Kishore Rai & Sabine Stoll. 2007. Free prefix ordering in Chintang. *Language* 83. 43–73. https://doi.org/10.1353/lan.2007.0002.

Blaho, Sylvia, Patrik Bye & Martin Krämer. 2007. Freedom of analysis? In Sylvia Blaho, Patrik Bye & Martin Krämer (eds.), *Freedom of analysis?*, Berlin/New York: De Gruyter Mouton. https://doi.org/10.1515/9783110198591.1.

Bobaljik, Jonathan David. 2000. The ins and outs of contextual allomorphy. *University of Maryland working papers in Linguistics* 10. 35–71.

Dąbkowski, Maksymilian. 2022. Paraguayan Guaraní and the typology of free affix order. *Proceedings of the Linguistic Society of America* 7. 5159. https://doi.org/10.3765/plsa.v7i1.5159.

Good, Jeff & Alan Yu. 1999. Affix-placement variation in Turkish. In *Annual meeting of the Berkeley linguistics society*, vol. 25, 63–74. https://doi.org/10.3765/bls.v25i2.1209.

Good, Jeff & Alan Yu. 2005. Morphosyntax of two Turkish subject pronominal paradigms. In Lorie Heggie & Francisco Ordóez (eds.), *Clitic and affix combinations: Theoretical perspectives*, 315–341. John Benjamins. https://doi.org/10.1075/la.74.13goo.

Grimshaw, Jane. 1997. Projection, heads, and optimality. *Linguistic Inquiry* 28. 373–422.

Güneş, Güliz. 2020. Variability in the realization of agreement in Turkish: A morphotactic account. In A. Gürer, D. Uygun Gökmen & B. Öztürk Başaran (eds.), *Morphological complexity within and across boundaries: Essays in honour of Aslı Göksel*, 236–261. Amsterdam: John Benjamins. https://doi.org/10.1075/slcs.215.09gun.

Güneş, Güliz. 2021. Morphosyntax and phonology of agreement in Turkish. *Syntax* 24. 143–190. https://doi.org/10.1111/synt.12210.

Kelepir, Meltem. 2001. *Topics in Turkish syntax: Clausal structure and scope*. Cambridge, MA: MIT dissertation.

Kornfilt, Jaklin. 1996. On copular clitic forms in Turkish. In A. Alexiadou, N. Fuhrhop, P. Law & S. Löhken (eds.), *ZAS papers in Linguistics*, vol. 6, 96–114. Berlin.

Manova, Stela & Mark Aronoff. 2010. Modeling affix order. *Morphology* 20. 109–131. https://doi.org/10.1007/s11525-010-9153-6.

Mansfield, John. 2015. Morphotactic variation, prosodic domains and the changing structure of the Murrinhpatha verb. *Asia-Pacific Language Variation* 1(2). 163–189. https://doi.org/10.1075/aplv.1.2.03man.

Neu, Eva. 2024. Hybridity and change in Turkish inflectional morphology. Ms., UMass Amherst.

Newbold, Lindsey. 2013. Variable affix ordering in Kuna. In John Sylak-Glassman & Justin Spence (eds.), *Structure and contact in languages of the Americas*, .

Prince, Alan & Paul Smolensky. 1993/2004. *Optimality Theory: Constraint interaction in generative grammar*. Oxford: Wiley-Blackwell. https://doi.org/10.1002/9780470759400.

Rice, Keren. 2011. Principles of affix ordering: An overview. *Word Structure* 4(2). 169–200. https://doi.org/10.3366/word.2011.0009.
Ryan, Kevin M. 2010. Variable affix order: Grammar and learning. *Language* 86. 758–791. https://doi.org/10.1353/lan.2010.0032.
Sezer, Engin. 2001. Finite inflection in Turkish. In Eser Erguvanlı Taylan (ed.), *The verb in Turkish*, 1–45. Amsterdam/Philadelphia: John Benjamins. https://doi.org/10.1075/la.44.02sez.
Uffmann, Christian. 2007. Restricting GEN. In Sylvia Blaho, Patrik Bye & Martin Krämer (eds.), *Freedom of analysis?*, De Gruyter Mouton. https://doi.org/10.1515/9783110198591.281.
Wolf, Matthew Adam. 2008. *Optimal interleaving: Serial morphology-phonology interaction in a contraint-based model*. Amherst, MA: University of Massachusetts Amherst dissertation.
Wurmbrand, Susanne. 2018. The cost of raising quantifiers. *Glossa: a journal of general linguistics* 3. 19. https://doi.org/10.5334/gjgl.329.
Wurmbrand, Susi. 2008. Word order and scope in German. *GAGL: Groninger Arbeiten zur germanistischen Linguistik* (46). 89–110.

Appendix: Demographic data

Housewife (2x)
Babysitter
Driver
Social media manager
Freelance content creator
Nurse
Building constructor
Financial specialist
Publisher
Publisher/writer
Student
Student/activist
Student/journalist
Graduate student (3x)
Professor (2x)

Table 2. Current (or, if retired, previous) occupation

Bitlis, Istanbul, USA
Bursa, Istanbul, USA
Denizli, Ankara, Istanbul, USA
Çorum, Istanbul
Erzurum
Erzurum, Istanbul
Hopa, Istanbul
Istanbul (5x)
Istanbul, Çanakkale
Istanbul, Trabzon, Balıkesir, Bursa
Mersin, Istanbul
Rize, Samsun, Istanbul
Sivas, Bolu, Erzurum, Erzincan, Giresun, Istanbul
Tokat, Istanbul
Trabzon, Istanbul, USA

Table 3. Current and previous places of residence
(at least 3 consecutive years, ordered chronologically)

Albanian, Macedonian
Arabic, Zazaki
Georgian
Laz
None (15x)

Table 4. Native languages other than Turkish