The Typology of TSL Case Assignment

Kenneth Hanson
Department of Linguistics
Stony Brook University
mail@kennethhanson.net

Abstract

Long-distance syntactic dependencies such as movement and agreement have been argued to be *tier-based strictly 2-local* (TSL-2) over Minimalist Grammar (MG) dependency trees, placing tight constraints on the range of predicted patterns (Graf 2022b; Hanson 2025b). Here, I argue that the same is true of the syntax of case. Building on Vu et al. (2019) and Hanson (2023), I show that small variations on a TSL-2 tree grammar correspond to well-known phenomena such as case spreading, variation in alignment, and several types of differential argument marking. At the same time, many simple yet unnatural patterns are ruled out. This study therefore strengthens the TSL-2 hypothesis while providing a computational explanation for several aspects of the typology of case.

1 Introduction

The *tier-based strictly 2-local* (TSL-2) formal languages over strings and trees have been found to be a good model of many long-distance phonological and syntactic dependencies, respectively (Heinz 2018; Graf 2022a). In addition to providing a tight upper bound on predicted structural patterns, the patterns that can be expressed correspond closely to attested cross-linguistic variation in long-distance phonotactics (McMullin and Hansson 2016), syntactic movement (Graf 2022b), and agreement (Hanson 2025b). These results thus show how formal considerations provide insight into issues of locality and typology which are of broad interest within linguistic theory.

The syntax of case has also been studied from a TSL perspective by Vu et al. (2019) and Hanson (2023), but no such cross-linguistic study has yet been conducted. In addition, prior work assumed that case may utilize structure-sensitive tier projection, which increases expressive power, and did not emphasize the typological significance of limiting the constraint window on the tier to two elements. This work addresses all of the above limitations. It argues that many well-attested case patterns are understood as follows: case assigners (licensors) determine the cases of one or more nominals in their domain (clause, VP, etc.) according to a TSL-2 tree language. Furthermore, the natural parameters of the model—the set of visible elements and the structure of the bigram constraints—correspond to key points of attested variation, while many unnatural patterns are ruled out. Case therefore turns out to be less mysterious than is sometimes thought: it occupies a small and well-structured corner of the space of formal patterns and is computationally similar to other long-distance dependencies.

The resulting model, which makes full use of both dimensions of a tree tier, is highly reminiscent of existing configurational theories of case (Yip et al. 1987; Baker 2015), but unlike those theories, it uses only computational mechanisms which are independently necessary. The analysis does raise several formal questions, in particular, why certain TSL-2 patterns do not seem to have any correlate in the realm of case. While some of these gaps are presumably best explained by extra-grammatical factors, I propose two additional formal generalizations which further narrow down the space of possible case patterns.

The remainder of this paper proceeds as follows. First, I introduce the TSL-2 model of case assignment and discuss some important characteristics of the model (2). Next, I show how the model accounts for many classic case phenomena while ruling out unattested patterns (3). Finally, I address the overgeneration question, and discuss how case fits within the larger picture (4).

2 Background and model

This section lays out the background on case and subregular syntax, and presents a TSL-2 model of case which forms the basis for the cross-linguistic analysis which follows. It begins with a brief overview of subregular complexity (2.1), the syntax of case (2.2), and previous subregular analyses of case (2.3). Next, I introduce the MG dependency tree framework (2.4) and how it can be used to model case dependencies (2.5), with examples from English. Finally, I discuss some formal and linguistic issues related to the proposed model (2.6).

2.1 Subregular complexity

The Chomsky Hierarchy (Chomsky 1959) measures the complexity of linguistic patterns in terms of string languages (string sets); similar hierarchies can be described for tree languages, string-to-string maps, and so on. Within these classifications, studies to date have found that linguistic patterns are overwhelming *subregular*, meaning that they do not require the full power of even the *regular*, or finite-state class (see Graf 2022a and references therein). Of the various subregular classes, TSL appears to be the upper bound for most individual dependencies, and M[ulti]-TSL for a grammar containing many such dependencies. This is true even for syntax if we use an appropriate tree-based representation. Limited exceptions also exist; so far these have been modeled with a set of extensions known as *structure-sensitive TSL*, or SS-TSL.[1]

2.2 Case in syntax

This work is solely concerned with the syntactic distribution of case, not the details of morphological realization, nor the distribution of overt/covert nominals (‘Case’ in the sense of Chomsky 1981). Given a syntactic structure, we seek only to correctly determine the cases assigned to each nominal. Within this scope, I focus almost exclusively on *structural cases* such as nominative, accusative, dative, ergative, and genitive, which have a non-trivial distribution. *Lexical case*, which is assigned to arguments of specific predicates, will be treated briefly, as it is interesting insofar as it interacts with structural case. *Semantic cases* such as locative/instrumental/etc. which occur mainly on adjuncts, are essentially self-licensed, and will therefore be ignored.[2] See Blake (2001), Butt (2006), and Malchukov and Spencer (2008) for additional general background on case.

Note that I assume to nominal phrases be DPs, but nothing crucial hinges on this; some or all nominals must instead be NPs, KPs, etc. Next, note that a given case may have multiple modes of assignment; for our purposes these will be considered distinct cases. Finally, although I speak of case being *assigned* to DPs, the present account is, mathematically speaking, a pure *licensing* model; see Section 2.6 for explanation.

2.3 Previous work

Vu et al. (2019) made the initial argument that the syntax of case is TSL. Technically, their model is SS-TSL rather than TSL proper; however, the need for SS-TSL comes solely from their use of traditional MG derivation trees, as is true of Graf (2018), on which their model is based. With the introduction of MG dependency trees (a different subtype of derivation tree) the need for SS-TSL appears to be eliminated for movement (Graf 2022b), agreement (Hanson 2025b), and case (this work), though not for binding (Graf and Shafiei 2019).

The present analysis therefore starts with the model in Hanson (2023), which uses dependency trees to analyze a variety of case patterns in Japanese. Although not emphasized in that paper, the analysis is entirely TSL-2 except for a brief use of SS-TSL for reasons of conciseness; Hanson (2025a, ch. 5) revises this to achieve a pure TSL-2 analysis. I also retain the approach of gathering all DPs in some case domain as the daughters of the domain node on a tier, with the case of each being determined by the tier mother and sisters together, as introduced there.

2.4 MG dependency trees

Here, I describe the MG dependency tree model as used in recent work in subregular syntax.[3]

MGs (Stabler 1997, 2011) are a formalization of ideas from Minimalist syntax. They can generate a variety of structures, but for present purposes the relevant structure is a *derivation tree*, which encodes the sequence of operations (Merge, Move, etc.) in a syntactic derivation. MG *dependency trees* are a subtype of derivation tree in which every node represents a lexical item; the operations themselves are indicated only in the form of *feature diacritics* on the nodes themselves. See Graf and Kostyszyn (2021) for a formal definition.

[1] See De Santo and Graf (2019) for a formal definition and Mayer and Major (2018), Graf and Mayer (2018), and Graf and Shafiei (2019) for empirical studies.

[2] Some theories distinguish an additional category of *inherent case*, which is assigned to arguments with a specific thematic role. For our purposes, it can be rolled into lexical case, since it is similarly local in character.

[3] The dependency trees in Kobele (2012) are essentially identical. In contrast, Boston et al. (2010) generate surface dependency trees using MGs.

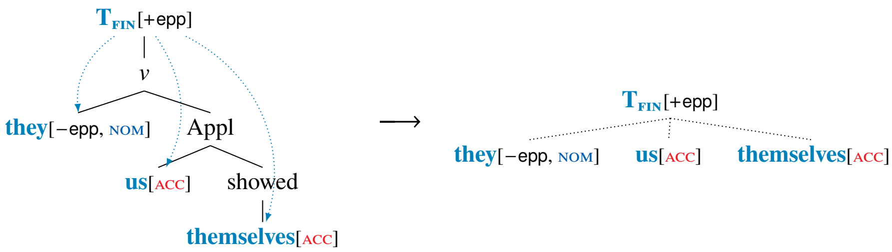

Figure 1: MG dependency tree and case tier for *They showed us themselves*. Left: full tree with case tier overlain. Right: case tier in isolation. Other dependencies such as EPP movement are controlled by a separate tier (not shown).

An example dependency tree for *They showed us themselves* is given in Figure 1, in which the subject moves to Spec-TP (call this EPP movement) and receives nominative case; both objects receive accusative case. Each daughter of a node is an argument, ordered from left to right in reverse order of first merge; this order also corresponds to asymmetric c-command among co-arguments. The features +epp and −epp indicate the landing site of EPP movement and the mover, respectively. Additionally, case features show which DPs receive nominative/accusative/genitive case. Following Hanson (2023), case assigners are not given a feature; instead, the label of the case assigner determines the string of cases of its dependents on a tree tier, as described in the next section. For simplicity, selectional features are omitted.[4]

As in traditional MG derivation trees, all elements appear in their base positions only. This works well for case, since scrambling and *wh*-movement, for example, preserve case. Even when movement is thought to alter case marking, the movement features themselves may be enough to determine the correct case; see Section 3.6 for an example. We diverge from standard MGs in one important way: the negative movement (and case) features are unordered, which is crucial if the tree language is to be TSL. For example, a DP bearing −epp and −wh simply moves to the closest landing site for each, and traces or copies are inserted as appropriate (cf. Graf 2022b, 2023).

2.5 TSL case assignment

I now explain the TSL model of case assignment as presented in Hanson (2023). I start with the definition of a TSL string language, which is then generalized to trees. I demonstrate the system with an analysis of some case patterns in English.

[4]Unlike in Minimalist syntax, MG features indicate operations that occur *in the present derivation* rather than the *potential* to undergo some operation.

A string language is *tier-based strictly k-local* (TSL-k) iff there is a set of salient symbols such that, when all others are deleted, licit/illicit strings can be distinguished with a *strictly k-local* (SL-k) grammar (Heinz et al. 2011). The set of salient symbols is the *tier alphabet*, notated T, and the string obtained by deleting all non-salient symbols is a *tier projection*. The SL-k grammar may be either a *negative grammar* (G^-), which is a set of forbidden k-grams, or a *positive grammar* (G^+), which is a set of permitted k-grams; for present purposes the positive form is more convenient. Also, we include the *edge markers* ⋊/⋉ in the tier projection so that we can refer to the beginning/end of the string. For example, consider a natural language whose high vowels harmonize in backness, such that every word contains only [i] or [u] but never both. This pattern is TSL-2, with $T = \{i, u\}$ and $G^+ = \{⋊i, ⋊u, ii, uu, i⋉, u⋉\}$. As a result, words like ‘kibili’ (⋊iii⋉) and ‘kubulu’ (⋊uuu⋉) are ruled in, but ‘kibulu’ (⋊iuu⋉) and ‘kubilu’ (⋊uiu⋉) are ruled out.

We generalize this idea to tree languages as follows. First, we ignore the non-salient elements and construct a *tree tier* which preserves dominance and precedence relations among those that remain. Continuing the preceding example, let us assume a single tier regulating case in English, containing all case assigners and assignees, which include at least T_{FIN} and all D heads (we consider other candidates below). The result is that all verbal arguments become daughters of T_{FIN}, as illustrated in Figure 1. Second, we associate each node on the tier with a TSL-k *daughter string language* (DSL), expressed in the form of a TSL-k *daughter string grammar* (DSG). In English finite clauses, the first (= highest) DP is nominative, and any others are accusative. Taking the position in the daughter string as a proxy for structural height (see Section 2.6), this pattern is described using the following grammar, which is

SL-2, and therefore also TSL-2.[5] A dot is added between elements for readability, and movement features are omitted for simplicity.

(1) Daughter string grammar for T_{FIN} (SL-2)
$G^+ = \{\rtimes \cdot \mathrm{D}[\text{NOM}],\ \mathrm{D}[\text{NOM}] \cdot \mathrm{D}[\text{ACC}],\ \mathrm{D}[\text{ACC}] \cdot \mathrm{D}[\text{ACC}],\ \mathrm{D}[\text{NOM}] \cdot \ltimes,\ \mathrm{D}[\text{ACC}] \cdot \ltimes\}$

This grammar allows daughter strings such as D[NOM], D[NOM] · D[ACC], D[NOM] · D[ACC] · D[ACC], etc., while disallowing an accusative subject followed by a nominative object, among many other configurations. Assuming that infinitive *to* does not appear on the tier, the analysis immediately extends to raising infinitives as well as ECM infinitives. These are illustrated below.

(2) She seems to like me (raising)

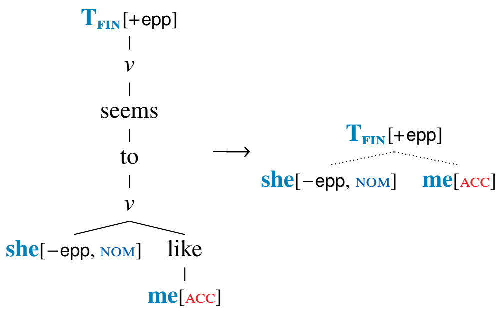

(3) We expect her to win (ECM)

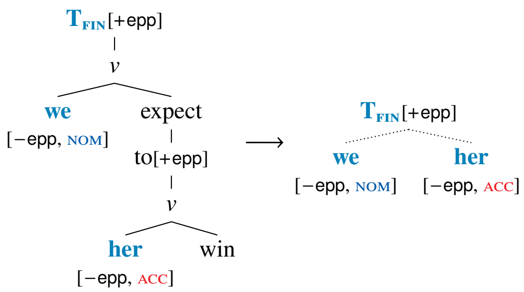

Several other aspects of English were previously treated by Vu et al. (2019) and could be translated into the current system with little difficultly. Since developing an in-depth analysis of English is not our goal, I will expand the grammar only slightly. (Movement features are henceforth omitted when not directly relevant.) I assume that DPs allow up to one genitive argument/possessor, possibly followed by one or more accusative arguments in the case of a *poss-ing* gerund (e.g. *his devouring the caviar*) and that other domains such as non-finite CPs (e.g. *for him to love caviar*) only allow accusative DPs. These structures are shown below. The case of the root is intentionally left unspecified, but would be accusative by default.

(4) his devouring the caviar (DP gerund)

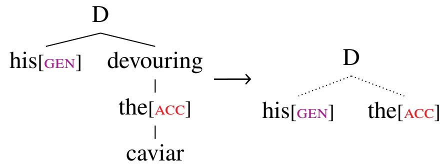

(5) *for him to love caviar* (non-finite CP)

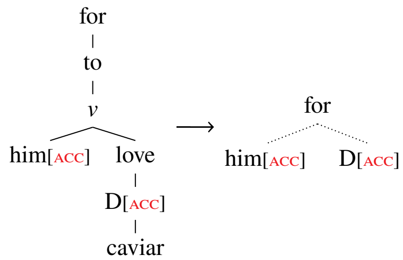

We must now deal with the fact that T/C heads also occur in the DSL. Let us assume that they *absorb* the case that would otherwise be assigned in that position. That is, they are marked for case, but show no morphological realization, just as lexical DPs do not. The following example demonstrates absorption of nominative by a clausal subject. Later, we will consider examples where certain elements are invisible rather than filling a case slot.

(6) That he owned a motorcycle surprised us

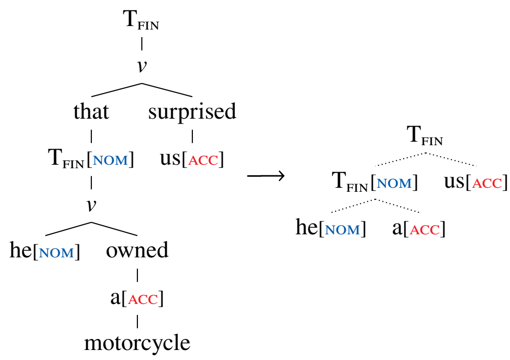

Putting all of this together, we arrive at the following TSL grammar.[6]

(7) Case tier for English:

- Project: $\{T_{\text{FIN}}, \mathrm{D}, C_{\text{INF}}\}$
- Daughter string grammars:

  T_{FIN}: $\{\rtimes \cdot \text{NOM}, \text{NOM} \cdot \text{ACC}, \text{ACC} \cdot \text{ACC}, \text{NOM} \cdot \ltimes, \text{ACC} \cdot \ltimes\}$

  D: $\{\rtimes \cdot \text{GEN}, \text{GEN} \cdot \text{ACC}, \text{ACC} \cdot \text{ACC}, \text{GEN} \cdot \ltimes, \text{ACC} \cdot \ltimes, \rtimes \cdot \ltimes\}$

  C_{INF}: $\{\rtimes \cdot \text{ACC}, \text{ACC} \cdot \text{ACC}, \text{ACC} \cdot \ltimes, \rtimes \cdot \ltimes\}$

This grammar incorporates several notational shortcuts. A lone category label stands in for any node of the given category *regardless of its MG*

[5] Although we do not make use of tier projection in the DSL at the moment, the definition of tree TSL allows it, and we do need it for case. See Sections 2.6 and 3.4.

[6] A more elegant analysis might use finite C for the domain of nominative case, mirroring C_{INF} for non-finite clauses.

features. Additionally, a label like NOM stands in for any nominative XP, regardless of its other features. The reader may confirm that each daughter string in the immediately preceding examples is well-formed according to the grammar.

To review, we have a TSL-2 model of case assignment which brings together all nominals in a case domain as the tier daughters of a case assigner, using the tier mother and sisters together in determining the case of each nominal. Compared to earlier work, the key innovation is the use of relative order among tier sisters. In existing TSL analyses of movement (Graf 2018, 2022b), all daughter string languages are closed under permutation, that is, the daughters can be shuffled with no change in wellformedness; here, this no longer holds. As shown in the next section, by making full use of both dimensions of the tree tier, we can account for a wide range of case phenomena in a straightforward manner.

2.6 Some consequences of the model

To close this section, I wish to briefly discuss a few formal and linguistic issues related to the model just presented.

Two chances for tier projection The definition of a TSL tree language allows for tier projection in the construction of the tree tier, and again in the daughter string languages. In the preceding example, there was no need for tier projection in the second step, but this is not always true. As discussed in Hanson (2023), adjuncts must be ignored for purposes of case assignment, but if they are ignored in the tree tier then any DPs inside them will be strewn among those in the containing clause. This problem can be avoided if we preserve adjuncts in the tree tier, then filter them out in the daughter string language. I make use of this technique in the treatment of lexical case in Section 3.4.

Configurational case assignment As mentioned in the introduction, the configurational perspective on case has ample precedent in the frameworks of *Case in Tiers* Yip et al. (1987) and *Dependent Case Theory* (Marantz 2000; Baker 2015), among others. While I do not conduct an in-depth comparison here, I do wish to highlight a few key points. Case in Tiers is a direct adaptation of autosegmental phonology (Goldsmith 1976), which itself inspired the TSL formal languages; unlike the current approach, it is framed in terms of linear surface order among DPs. The core idea of Dependent Case Theory is that case is assigned according to relative c-command relations among DPs in some domain; the present model is an implementation of the same idea.

As noted by Vu et al. (2019), tree tiers also subsume standard generative case theory, in which case is a pairwise relation between assigner and assignee. This is especially clear in the present model: such relations are just the special case in which the daughter string language is closed under permutation, as with movement. In this way, the model explains why case patterns which are described so naturally by configurational theories should exist—they are just examples of the more general capabilities of tree tiers.

Tier sisters and command Although we are using sister order as a proxy for structural height/prominence, it is possible for a node X to be the left sister of Y on the tree tier even though X does not c-command Y in the phrase structure tree. This could be useful in deriving certain c-command violations, such as the limited possibility for an NPI to be licensed by an element embedded in a left branch (e.g. [[*No one's*] directions] were *any* help). On the other hand, if c-command is in fact the relevant structural relation between DPs for purposes of case assignment, this could lead to incorrect predictions.

In practice, there is usually some independent reason to project the root of a complex left branch; this applies to adjuncts as just mentioned, but also to clausal subjects and PP subjects. That said, the issue points towards the need for additional constraints beyond those provided by tree tiers, as was already acknowledged by Graf (2022b). The other major approach to subregular syntax applies string constraints to paths which encode command relations (Graf and Shafiei 2019; Hanson 2025b). It may be possible to recast much of the present analysis in that system. However, given the naturalness of the tree tier analysis of case, it seems to me that the ideal solution for case would be a hybrid system in which tier sister order is guaranteed to respect c-command.[7]

Multi-tier patterns While a single case tier is sufficient for the above fragment of English, this might not be true for all languages. For example,

[7] While writing this paper, I considered redefining the tier sister relation in terms of the d[derivational]-command relation of Graf and Shafiei (2019), but this only created a new problem: a single node could end up with several discontinuous daughter strings. The complement spine model in Hanson (2025b), which filters out the unwanted command relations, may be a better starting point.

the analysis of Japanese in Hanson (2023) includes three case tiers, corresponding to case in the verbal domain (nominative, accusative, dative), case in the nominal domain (genitive), and lexical case; the entire case system is therefore MTSL-2.[8] In this paper, I focus on typological predictions of a single tier system, which by hypothesis is enough to handle the structural cases of a single domain, leaving a proper study of multi-tier patterns to future work. That said, I do include several examples of interactions across tiers in Section 3.

Assignment vs licensing Finally, a comment on the issue of case *assignment* versus *licensing*. A dependency tree is well-formed if it does not violate any constraints in an (M)TSL tree grammar. Conceptually, the presence of a case feature on a given node can be interpreted as saying that it will be assigned said case as part of the syntactic derivation. In a mathematical sense, however, the model is a pure licensing model; this applies not only to case, but to all syntactic dependencies. It may also be insightful to model case as a TSL transduction from the dependency tree to the corresponding phrase structure tree (see Graf 2023 regarding movement).

3 Cross-linguistic study

In this section, I demonstrate how the proposed model of case assignment accounts for many well-known and theoretically interesting patterns in the syntactic distribution of case. As in previous studies of this type, I emphasize how the natural parameters of the model—the tier projection and the constraints on the daughter strings—correspond closely to attested cross-linguistic variation (3.1). Specific phenomena of interest include case spreading (3.2), variation in alignment (accusative, ergative, etc.) (3.3), interactions between lexical and structural case (3.4), domain effects (3.5), and four types of differential argument marking (3.6). I also briefly discuss what we do *not* see, and why the model rules out such patterns (3.7).

3.1 Parameters of variation

I wish to highlight two aspects of the TSL-2 model of case which are the focus of the present study. First, the model provides two basic parameters for manipulation: i) the tier alphabet, which determines which items are visible, and ii) the bigrams of the daughter string grammars. Each of these is expected to vary across languages, and I show that the obvious ways that they may differ aligns closely with the variation that we see. For example, Graf (2022a) explains how the difference between single and multiple *wh*-movement boils down to presence or absence of a single bigram in the daughter string language. Similarly, many island effects arise from the presence of extra nodes on the tier which interrupt the relation between mother and daughter. Here, I show that the same sorts of variation can be found in the realm of case.

Second, we have considerable control over the relationship between a domain node and two adjacent DPs on the tier, both in the mother-daughter and left-sister relations. For example, we can refer to the properties of the the domain node to determine the case pattern for its daughters, and to the properties of the sisters in the derivation of differential argument marking. This also has precedent in previous work in TSL syntax (Graf 2022a; Hanson 2025b); as mentioned above, the main innovation here is the specific focus on tier sisters. In contrast, our ability to coordinate elements which do not stand in such a relationship is extremely limited. Accordingly, I show how many logically simple yet linguistically unnatural patterns are ruled out on the grounds that they are not TSL-2.

Of course, not every unattested pattern can be ruled out in this manner. Many other factors are involved, including substantive differences between different types of dependencies as well as constraints on diachronic development and language acquisition. Even so, we have a strong hypothesis for what constitutes a possible case pattern, which plausibly derives from some limitation of the underlying cognitive system (Graf 2022a). This in turn clarifies which gaps *must* be accounted for by other factors, computational or otherwise; I suggest some additional formal factors in Section 4.

3.2 Spreading

One of the major motivations of Yip et al. (1987) was to account for case spreading, that is, situations where a single case is assigned to multiple DPs in some domain. We saw spreading of accusative in the previous section, though as Yip et al. note, Swedish would be a better example, since accusative is plausibly the default case in English.

(8) Multiple accusative in Swedish (Yip et al. 1987)
Kungen gav honom henne (till maka).
the.king gave him.ACC her.ACC (for wife)

[8] An MTSL language is just the intersection of several TSL languages (De Santo and Graf 2019).

Similarly, Japanese allows multiple nominatives and genitives in several contexts.

(9) Multiple nominatives in Japanese

a. Stative object (Hiraiwa 2001)

Mary ga eigo ga yoku dekiru.
Mary NOM English NOM well can.do

'Mary can speak English well.'

b. Possessor raising (Kuno 1973)

Yama ga ki ga kirei desu.
mountain NOM trees NOM pretty are

'The mountains—their trees are pretty.'

(10) Multiple genitives in Japanese

a. Multiple possessors (Saito et al. 2008)

Taroo no Chomsky no hon
Taroo GEN Chomsky GEN book

'Taroo's book by Chomsky'

b. Nominal subject/object (Saito et al. 2008)

yabanzin no Rooma no hakai
barbarian GEN Rome GEN destruction

'the barbarians' destruction of Rome'

As discussed by Graf (2022b), multiple licensing is also a straightforward prediction of the TSL-2 model: all we need is a bigram in the relevant daughter string grammar which duplicates the same item twice. Thus, by including ACC · ACC, NOM · NOM, or GEN · GEN, etc., we enable spreading without the need to posit additional (covert) case assigners. We also predict that spreading will continue if additional DPs are present, since a daughter string with three or more repetitions of, say, accusative, has exactly the same bigrams: $\{\rtimes \cdot$ ACC, ACC · ACC, ACC $\cdot \ltimes\}$. Put another way, an SL-2 grammar cannot distinguish "exactly two" from "two or more".

While the literature rarely discusses such constructions, this appears to be true at least of possessor raising in Japanese (Kuno 1973).[9] Furthermore, restricting spreading to a *maximum* of two DPs requires at least TSL-3. This kind of counting is highly unusual in natural language, and TSL-3 can generate many other strange patterns that TSL-2 cannot (see Section 3.7). In absence of contradictory evidence, the TSL-2 analysis of spreading is therefore preferred.

Another superficially similar phenomenon, not treated by Yip et al. (1987), is *case concord*, in which the case of the DP is realized in multiple places, including the head noun, the determiner, adjectival modifiers, and perhaps other elements. An extreme example of this type from Panjima (Pama-Nyungan), with five concordial elements, is shown below.

(11) Ngatha wiya-rna ngunha-yu maripa-yu paka-lalha-ku nharniwalk-ku warrungkamu-la-ku
I.NOM see-PST that-ACC man-ACC come-PRF-ACC hither-ACC morning-LOC-ACC

'I saw that man who came this way this morning.'
(Blake 2001, ch. 4)

There are several ways this might be implemented. The direct approach involves duplicating the bigrams of the spreading grammar to allow iteration of each case. The trouble with strategy is that the case shared among the concordial elements is not arbitrary, but rather determined by the case of the entire DP. One possible solution is to *split* the DSG for D on the case tier into several, one for each case.

(12) Schematic grammar for case concord

D[NOM]: $\{\rtimes \cdot$ NOM, NOM · NOM, NOM $\cdot \ltimes$, $\rtimes \cdot \ltimes\}$
D[ACC]: $\{\rtimes \cdot$ ACC, ACC · ACC, ACC $\cdot \ltimes$, $\rtimes \cdot \ltimes\}$
D[DAT]: $\{\rtimes \cdot$ DAT, DAT · DAT, DAT $\cdot \ltimes$, $\rtimes \cdot \ltimes\}$
etc.

A better approach might be to treat case concord as a kind of agreement, since person/number/gender features are often shared as well. Furthermore, the class of items involved is somewhat different, including nouns and adjectives but excluding possessors and arguments (e.g. the possessor of an accusative DP does not become accusative). While the basic idea remains the same, the tier for case *assignment* would be completely separate from those for case *concord*; see Hanson (2024a) for a possible approach to the latter.

To briefly summarize, the TSL-2 model predicts the existence of iteration. The appears to be bourne out in several ways in the realm of case, mirroring its appearance in long-distance harmony, multiple *wh*-movement, and concord in the DP. However, we cannot reasonably expect every such pattern to occur in each of these domains. In particular, there are TSL-2 patterns which are highly characteristic of case but which are not nearly as robustly attested elsewhere. We consider some of these below.

3.3 Alignment

Case marking correlates imperfectly with both semantic role and grammatical function. In line with

[9] Numeral quantifiers and PP modifiers of nouns also take the *no* clitic, which would seem to imply an unlimited number of genitives. But *no* can also be an adnominal copula, and it is difficult to tell apart these two functions.

the configurational view adopted above, I show that many of the most common patterns correspond exactly to minor differences in the structure of the SL-2 daughter string language on a case tier. Note that many languages employ multiple alignments in the same domain, conditioned by various structural and lexical factors, which I abstract away from here; in Section 3.5 we consider case conditioned by aspect and finiteness.

Following common convention, I refer to the sole argument of an intransitive verb as S, the "more agent-like" argument of a transitive verb as A, and the "less agent-like" argument as O (cf. Bickel and Nichols 2008). The case of S is usually the (literally) unmarked case, which we call *nominative*. When A and S share a case and O receives a different case, we call the case of O *accusative*. This is known as accusative alignment, and is the dominant pattern in most Indo-European languages. It is also common for S and O share a case to the exclusion of A. We call the case of A *ergative*, and the overall pattern is known as ergative alignment. This is the dominant pattern in Shipibo (Panoan family).

(13) Shipibo (Baker 2015)

a. Maria-nin-ra ochiti noko-ke.
Maria-ERG-PRT dog find-PRF
'Maria found the dog.'

b. Maria-ra ka-ke.
Maria-PRT go-PRF
'Maria went.'

For the languages under discussion, I assume a consistent structural relationship such that A is higher than O. I further assume following Baker (2015) that ergative case in Shipibo is assigned to the higher of two DPs (the A argument) in the same clause.[10] The basic grammar for ergative alignment is therefore nearly identical to that for accusative. The key difference is that the 'new' case of the transitive structure is introduced to the left of nominative case.

(14) Grammar for ergative alignment
$G^+ = \{\rtimes \cdot \textsc{nom}, \rtimes \cdot \textsc{erg}, \textsc{erg} \cdot \textsc{nom}, \textsc{nom} \cdot \ltimes\}$

It is also possible for S, A, and O to each have a distinct case ('tripartite' alignment), or for them to all share the same case ('neutral' alignment). The neutral grammar is identical to what we called

[10]It is likely that not every language with a case called 'ergative' works like this, which Baker acknowledges. See Butt (2006, ch. 6) for additional discussion. In Section 3.5, I examine another type of 'ergative' system.

Alignment	Case of S	A	O
Accusative	NOM	NOM	ACC
Ergative	NOM	ERG	NOM
Tripartite	NOM	ERG	ACC
Neutral	NOM	NOM	NOM

Table 1: Common case alignments

'spreading' in 3.2, and the tripartite grammar is given below.

(15) Grammar for tripartite alignment
$G^+ = \{\rtimes \cdot \textsc{nom}, \rtimes \cdot \textsc{erg}, \textsc{erg} \cdot \textsc{acc}, \textsc{nom} \cdot \ltimes, \textsc{acc} \cdot \ltimes\}$

We see that every logical possibility for the cases of the canonical transitive verb is realized: either A or O may have a special case, or both, or neither. This is summarized in Table 1. Furthermore, each variant differs in only a few bigrams compared to the others. Note that some authors, such as Bittner and Hale (1996), distinguish even more systems. For example, internal and external arguments of intransitive verbs may behave differently, which is a major factor in certain dialects of Basque. Additionally, alignment may change according to tense or aspect, as in Hindi. From the present perspective, such "alignments" usually involve multiple domain nodes, and will be addressed in Section 3.5.

Now, we ask what happens when there are more than two arguments. Following Bickel and Nichols (2008), I refer to the two objects of a ditransitive verb as G (for the more goal-like argument) and T (for the more theme-like argument); ditransitive subjects are rarely distinguished from transitive subjects, so they are still A. It is common for either G or T to receive the same case as O, while the other receives a new case. For example, G can be *dative*, and T aligned with O. This occurs both in languages with accusative alignment, such as German and Japanese, as well as languages with ergative alignment, such as Ingush (Nakh-Daghestanian).

(16) Japanese ditransitive verb
Jin ga Yumi ni hon o ageta.
Jin NOM Yumi DAT book ACC gave
'Jin give Yumi a book.' (Hanson 2023)

(17) Ingush ditransitive verb
Aaz cynna mashen jelar
1SG.ERG 3SG.DAT car.NOM gave
'I gave him/her a car' (Bickel and Nichols 2008)

Maintaining our earlier assumption that the default linear order indicate c-command among base argument positions, the grammar for such systems are given below.

(18) Accusative alignment with dative G
$G^+ = \{\rtimes \cdot \text{NOM}, \text{NOM} \cdot \text{ACC}, \text{NOM} \cdot \text{DAT}, \text{DAT} \cdot \text{ACC}, \text{NOM} \cdot \ltimes, \text{ACC} \cdot \ltimes\}$

(19) Ergative alignment with dative G
$G^+ = \{\rtimes \cdot \text{NOM}, \rtimes \cdot \text{ERG}, \text{ERG} \cdot \text{NOM}, \text{ERG} \cdot \text{DAT}, \text{DAT} \cdot \text{NOM}, \text{NOM} \cdot \ltimes\}$

It is also possible for T to have a special case and G aligned with O, or for all three to be aligned, with various relationships to S/A. On top of this, some languages may employ low goals rather than high goal, or perhaps include both structures. While I cannot do justice to the full typology here, what is clear is that as with transitives, several well-attested systems correspond to simple manipulations of the TSL-2 grammar: is there a third basic case and does it occur before or after the case of O?

We could continue in this fashion by asking what happens when there are four arguments. Although simplex predicates with four DP arguments are rare (perhaps nonexistent), they can be constructed by valency increasing operations such as morphological causatives. Unfortunately, as with spreading, the literature on case rarely discusses such data. For instance, not a single example of a causative of a ditransitive could be found by searching all index entries for both “causative” and ”ditransitive” in the 900-page Oxford Handbook of Case (Malchukov and Spencer 2008). However, Hanson (2023) notes that such constructions are possible in Japanese, and the middle two DPs are both dative (both datives are plausibly structural, or at least ambiguous between structural and lexical dative).

(20) Causative of ditransitive

Ken	ga	Jin	ni	Yumi	ni	hon	o	agesaseta.
Ken	NOM	Jin	DAT	Yumi	DAT	book	ACC	gave.CAUS

‘Ken made/let Jin give Yumi a book.’

The grammar proposed there, shown below with extraneous details removed, involves spreading of dative.

(21) Verbal cases in Japanese
$G^+ = \{\rtimes \cdot \text{NOM}, \text{NOM} \cdot \text{ACC}, \text{NOM} \cdot \text{DAT}, \text{DAT} \cdot \text{DAT}, \text{DAT} \cdot \text{ACC}, \text{NOM} \cdot \ltimes, \text{ACC} \cdot \ltimes\}$

While further investigation is clearly needed, our survey of alignment clearly supports the idea that the various structural cases available in a given domain can be assigned according to a TSL-2 grammar. There are, of course, alternatives that could fit the readily available data. For example, we could employ a traditional generative analysis, in which each case is associated with a unique functional head: T, Aspect, Voice, Appl, etc. But then we would lose generalizations like the dative rule in Japanese.[11] In contrast, such patterns are a straightforward prediction of the TSL-2 model with very coarse case domains.

3.4 Lexical case and visibility

In the preceding sections, we focused mostly on the structure of the constraints. Now, we begin to consider the contents of the tier alphabet as well.

In Section 2.5, I showed how a clausal subject fills the nominative slot in English, which I called case absorption. When a lexically case-marked DP occurs in a particular case position, it is possible for the corresponding structural case to instead be *displaced* to the next available DP. This occurs with lexical datives in Icelandic.

(22) Case displacement in Icelandic (Yip et al. 1987)

a. Siggi leyndi konuna sannleikanum.
S.NOM concealed the.woman.ACC the.truth.LDAT
‘Siggi concealed the truth from the woman.’

b. Siggi sagði barninu söguna.
S.NOM told the.child.LDAT the.story.ACC
‘Siggi told the child the story.’

c. Barninu batnaði veikin.
the.child.LDAT recovered.from the.disease.NOM
‘The child recovered from the disease.’

When LDAT occurs in the lowest (oblique) position, the subject and first object are unaffected. But when it marks in the first object position, accusative is displaced to the second object, and when it marks the subject, nominative is displaced to the object. Icelandic is a particularly good example of the latter phenomenon, as it is clear that the lexically-case marked subject is in fact in the usual subject position (Zaenen et al. 1985).

The core of the analysis is to make lexical datives as invisible on the tier controlling structural case in the clause. That is, the tier alphabet contains T and any D marked for structural case (nominative,

[11] The idea that dative is the case of the middle of three DPs originates in Kuno (1973). In Baker (2015), dative applies to the upper DP in the domain of VP, but only when *v*P is a phase. Kuno’s generalization therefore becomes an accident.

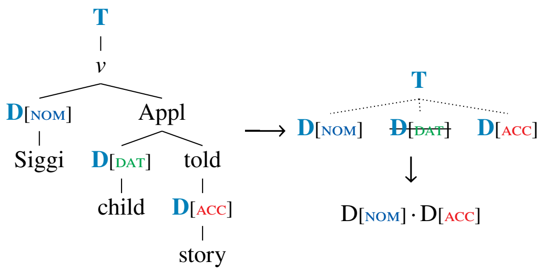

Figure 2: Derivation of case displacement. Lexical datives appear on the tree tier, but are ignored in the daughter string language of T. Nominative and accusative are adjacent as usual.

accusative, and perhaps structural dative) but not lexical dative. As a result, the tiers for the above examples look the same as they would for an ordinary transitive/intransitive clause.

As mentioned in Section 2.6, there is a complication, which is that any DPs inside the invisible lexical dative will be strewn among the arguments of the containing clause. The solution, borrowed from Hanson (2023), is to retain lexical datives in the tree tier, but ignore them in the daughter string language, which is now properly TSL-2. This is illustrated for (22b) in Figure 2.[12]

As a final remark, though I have not specified how lexical dative itself is assigned in Icelandic, lexical case is generally held to be assigned locally to arguments of specific predicates. This means that lexical case occupies the tier which contains everything; in other words, a simple SL grammar will do. But as we have seen, it is still formally interesting due to its interaction with cases assigned on non-trivial tiers.

3.5 Domain effects

Next, we consider effects arising from the the nature of the domain nodes which appear on the tier.

As mentioned in Section 2.6, it is normal to have several case domains, which might include TP, VP, DP, and perhaps others. Case may also correlate with clausal properties such as aspect, as in Hindi.

(23) Alignment split by aspect in Hindi (Mahajan 1990)

a. Raam rot̤ii khaataa thaa.
Raam.NOM bread.NOM eat.IPFV be.PST
'Raam ate bread (habitually).'

b. Raam-ne rot̤ii khaayii.
Raam-ERG bread.NOM eat.PFV
'Raam ate bread.'

Based on this data, imperfective clauses have neutral alignment while perfectives are ergative. (Hindi also has differential object marking, which is covered in 3.6.) Leaving aside the details, let us assume that the Asp(ect) head, realized by an auxiliary only in the imperfective, controls alignment according to the following grammar:

(24) Split alignment in Hindi

- Project: {Asp, D}
- Daughter string grammars:
  Ipfv: {⋊ · NOM, NOM · NOM, NOM · ⋉}
  Pfv: {⋊ · NOM, ⋊ · ERG, ERG · NOM, NOM · NOM, NOM · ⋉}

In reality, it is more likely that neutral and ergative alignment are associated with Asp and T or vice versa (Thomas McFadden, p.c.), but without delving into the details, it is difficult to distinguish between these alternatives. The key point is that major category alone need not be the distinguishing factor.

However, a clear example of this type can be found in Basque, in which finiteness is the distinguishing factor. For context, transitive clauses show ergative alignment, while the case of an intransitive subject varies. The data is complex, but in western/central dialects, the basic generalization is that a lone external argument receives ergative case, while a lone internal argument receives absolutive (=nominative) case, a pattern which is sometimes called 'Split S'.

(25) Split S in Basque (Berro and Etxepare 2017)

a. Jon-ek dantza-tu du.
Jon-ERG dance-PRF AUX
'Jon has danced.' (unergative)

b. Jon eror-i da.
Jon.NOM fall-PRF AUX
'Jon has fallen.' (unaccusative)

By testing a variety of constructions, it is possible to show that ergative case is only available in the domain of finite T (Rezac et al. 2014). For example, embedded finite clauses can have ergative subjects, whereas in gerunds, which are non-finite, ergative marking is lost.

(26) Ergative is lost in gerunds (Rezac et al. 2014)

a. [Katu-ek sagu-ak harrapa-tu dituzte-la] ikusi dut.
cat-PL.ERG mouse-PL.NOM caught AUX-that seen AUX

[12] Hanson (2023) uses this technique to factor out PP adjuncts—lexical datives do not require this treatment due to a quirk of the grammar of Japanese. We might also use it for PP arguments that induce case displacement, to the extent that such PPs can be distinguished from DPs with lexical case.

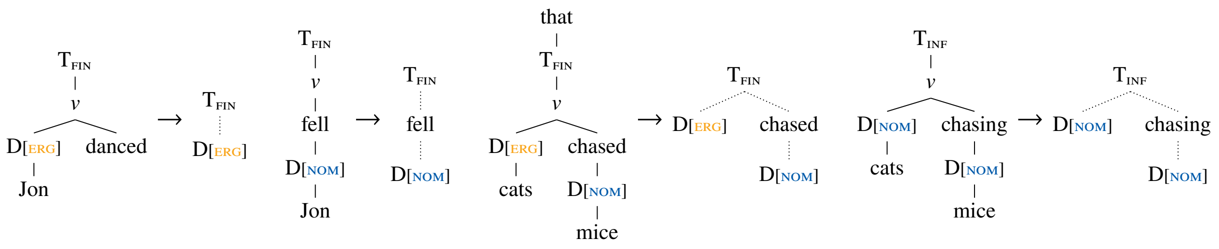

Figure 3: Ergative in Basque is available only in the domain of finite T. From left to right: unergative, unaccusative, transitive finite, transitive gerund.

'I saw that the cats caught the mice.'

b. [Katu-ak sagu-ak harrapa-tzen] ikusi ditut.
cat-PL.NOM mouse-PL.NOM catch-ing seen AUX
'I saw the cats catch the mice.'

We can derive this pattern if we assume that TP and VP are case domains and that only finite T can assign ergative; both V and non-finite T cannot. This is illustrated in Figure 3. I leave open whether the domain nodes should absorb case, or whether they should be ignored in the daughter string language as in Section 3.4.

It is also possible for a domain node that is usually present to be invisible in certain contexts. This is the basis of the analysis of long-distance (cross-clausal) case assignment in Hanson (2023), as exemplified by constructions such as the following.

(27) Finite ECM (adapted from Kishimoto 2018)

Ken ga [Eri {ga/o} kawaii to] omotteiru.
Ken NOM Eri {NOM/ACC} be.cute C think
'Ken thinks that Eri is cute.'

While I cannot go into detail here, the basic idea is that certain lexical predicates optionally select an embedded complementizer which is invisible on the relevant tier. Thus, all arguments become dependents of the upper case assigner, effectively merging two case domains, with the result that the embedded subject receives accusative case like a canonical object.[13] Though it is perhaps more common to analyze such phenomena as involving a truncated clause structure, the TSL approach allows us abstract away from the precise nature of the structural difference. This is helpful for Japanese since the particle *to* appears extremely high in the extended CP (Saito 2015), yet is still present under long-distance assignment of accusative.

To summarize what we have seen so far, we can derive many of the most common variations in case marking simply by manipulating the contents of the tier alphabet, the number of cases, and their positions in the bigrams of the daughter string grammars. In the next section, we consider the last obvious variable, which is to utilize features of DPs other than case in our grammars.

3.6 Differential argument marking

Differential argument marking (DAM) refers to phenomena in which some morphological marking of a DP, including but not limited to case, is only sometimes present; see Witzlack-Makarevich and Seržant (2018) for an overview. DAM subsumes both differential object marking (DOM) as well as differential subject marking (DSM), and tends to be linked to discourse and semantic factors, which themselves could involve the argument itself or a co-argument. Here, I show that the TSL-2 model derives this four-way typology.

We begin with a classic example of DOM, as illustrated with accusative case in Sakha (Turkic). In Sakha, accusative case marking is associated with a definitive/specific interpretation as well as movement out of the VP, as diagnosed by the position of manner adverbs.

(28) Sakha DOM (Baker and Vinokurova 2010)

a. Masha salamaat-*(y) türgennik sie-te.
Masha porridge-ACC quickly eat-PST.3SG
'Masha ate the porridge quickly.'

b. Masha türgennik salamaat-(#y) sie-te.
Masha quickly porridge-ACC eat-PST.3SG
'Masha ate porridge quickly.'
(ACC on 'porridge' only with contrastive focus)

Baker and Vinokurova (2010) posit that accusative is assigned to the lower of two DPs in the domain of T (as I have been assuming), and that

[13] In Japanese, the embedded subject need not vacate the clause for this to happen, though as discussed below, we can also handle case alternations triggered by movement in the present system.

	Locus of Controlling Property	
Subtype	Self	External
DSM	Kham, Folopa, Umpithamu	Eastern Ostyak, Ika, Kanuri
DOM	Hindi, Spanish, Turkish, Sakha	Ik

Table 2: Typology of self/externally driven DAM

movement of the object out of the VP is what causes it to receive accusative case. That movement is the trigger for DOM is not a foregone conclusion, as I discuss below, but let us adopt this view for now.

As discussed in Graf (2022b) and Hanson (2025b), we can handle certain interactions between syntactic dependencies by referencing the MG features for one dependency in the TSL grammar for another, as there is no *a priori* reason why such information should not be available. For example, case-sensitive agreement is modeled by projecting only D heads bearing an appropriate case on the tier regulating agreement. Here, we assign accusative only to DPs which bear the feature for object shift in addition to appearing after another another DP on the tier. This is illustrated in the grammar below.

(29) Accusative DOM triggered by object movement

- Project: $\mathrm{T}_{\textsc{fin}}$, all D
- DSG of $\mathrm{T}_{\textsc{fin}}$: $\{\rtimes \cdot \mathrm{D}[\textsc{nom}], \mathrm{D}[\textsc{nom}] \cdot \mathrm{D}[-\text{oshift}, \textsc{acc}], \mathrm{D}[\textsc{nom}] \cdot \mathrm{D}[\textsc{nom}], \mathrm{D}[\textsc{nom}] \cdot \ltimes, \mathrm{D}[-\text{oshift}, \textsc{acc}] \cdot \ltimes\}$

If such constraints are generally possible, we predict that movement of the object may also influence ergative case marking of the subject, a kind of DSM. Baker (2015, pp. 127–129) also gives several examples of this type. In addition, we predict that DSM could be triggered by an internal property of the subject, or conversely, DOM triggered externally by a property of the subject. As discussed by Daniel (2024), all four possibilities are realized. The set of examples she cites are shown in Table 2.

At least one cell in the typology is problematic for many approaches to case. For example, self-driven DSM is unexpected under Baker's assumptions, since subjects do not normally move for semantic/discourse reasons, while externally-driven DAM is generally unexpected under functionalist theories (Baker 2015, pp. 129–130). From the present perspective, all four cells are equally easy to derive. Furthermore, movement need not be the trigger; it could also be agreement, or the underlying semantic/discourse feature, if we make it available in the syntax. Indeed, Hindi also has differential object marking (not shown previously), but unlike in Sakha it is not thought to involve movement.

We can therefore provide a maximally simple explanation for the four-way typology: given that the subject and object appear in the same tier bigram, we can reference the properties of either in determining their cases. In future work it would be important to confirm whether global case splits (cf. Bárány 2017), in which the combined properties of *both* the subject and object influence case or agreement, can be handled in the same manner. In principle this should be possible, but the data is complex enough that a closer look is warranted.

As this marks the end of our final case study, the case phenomena which have been treated in this paper are summarized in Table 3. Just from this brief survey, we have identified clear correlates of visibility in both the tree tier and daughter string language as well as various manipulations of the mother-daughter and sister-sister relations on the tier. Conversely, many common and sometimes theoretically problematic case phenomena have a straightforward analysis in the TSL-2 model.

3.7 Non-existent case patterns

Having shown how many of the positive predictions of the TSL-2 model for case are bourne out, I close this section by examining some non-existent patterns which are ruled out purely on formal grounds. Graf (2022b) constructs a selection of logically simple yet unattested island constraints, and notes that none of them are TSL-2. Now, consider the following hypothetical case systems:

(30) **2x2 alignment:** the first two DPs are nominative; the next two DPs are accusative; repeat

(31) **Anti-local dependent case:** the subject is nominative if there are one or two arguments and ergative if there are three or more; objects are always accusative

(32) **Coordinated alignment:** when two clauses are coordinated, one must have accusative alignment and the other ergative, in either order

(33) **Rationed dative case:** up to two clauses may contain dative DPs; the rest use genitive instead

(34) **Rotating alignment:** the main clause is accusative; a singly embedded clause is ergative; a doubly embedded clause is accusative; etc.

All of these are straightforwardly ruled out due to requiring some non-TSL-2 mechanism, either in domain node identification or in the daughter string language:

Phenomenon	Section	One line summary
Spreading	3.2	Allow NOM · NOM, ACC · ACC, etc. in the DSL
Alignment	3.3	Does the marked case precede or follow the unmarked case?
Split alignment	3.5	PFV/IPFV, T_{FIN}/T_{INF}, etc., have different DSLs
Long-distance case	3.5	Certain C heads are invisible on the structural case tier
Case absorption	2.5	P/C/etc. are visible on the tree tier and in the DSL
Case displacement	3.4	Lexical case-marked D is invisible in the DSL
Diff. argument marking	3.6	DSL can reference non-case features of D

Table 3: Summary of case patterns and their TSL-2 analysis

- A window larger than two (30, 31)
- Global boolean logic (32)
- Threshold counting (33)
- Modulo counting (34)

These examples were selected in part because they are conceptually straightforward yet extremely unnatural from a linguistic perspective. They are also simple in computational terms, all being regular over dependency trees, and with the exception of (34), subregular. As such, we have good reason to believe it is TSL-2 specifically that characterizes the class of possible case patterns—it is powerful enough, but also not overpowered.

One possible objection to this line of reasoning is that all of the above systems are clearly dysfunctional as a means of mapping DPs to grammatical positions, and are therefore already ruled out for this reason, irrespective of computational complexity. But this mapping is tenuous at best even in real languages. Already, we have seen many examples of grammatical roles being conflated under the same case, and of the same role being split. The mapping is also subject to extensive but not unlimited variation. There remains ample room for formal factors to play a role in explaining the typology, and the fact that TSL-2 makes approximately the right cut is unlikely to be a coincidence. This is especially true when we consider the broader picture, since seemingly unrelated phenomena such as movement and long-distance phonotactics also seem to obey the same restriction.

4 The overgeneration question, or, what makes a possible case pattern?

In the preceding section, we saw how many robustly attested case patterns are TSL-2, and in contrast, how many simple yet unattested patterns are not. Now, we should also ask whether there are any instances of undergeneration and overgeneration. I am not aware of any attested case patterns which are clearly not TSL-2 (undergeneration), but we can construct certain unattested patterns which are TSL-2 (overgeneration). For instance, nominative and accusative could alternate according to the grammar below:

(35) $G^{+} = \{\rtimes \cdot \text{NOM},\ \text{NOM} \cdot \text{ACC},\ \text{ACC} \cdot \text{NOM},\ \text{ACC} \cdot \ltimes,\ \text{NOM} \cdot \ltimes\}$

This pattern is only slightly different from (30) in that it utilizes a window of size 2 rather than 3, but it is just as unnatural.

However, as discussed in Section 3.1, we cannot expect every (T)SL-2 pattern to be realized in every type of linguistic dependency. Some, like iteration, are shared across many dependency types, while others show up only here and there. Indeed, although alternating patterns like the above are unnatural for case, they *can* be found in the realm of category selection—D selects N, which selects P, which selects D, and so on—as well as syllable structure (Hanson 2024b). Another set of patterns is characteristic of in feature-matching operations such as long-distance harmony and syntactic agreement (Hanson 2025b). Case exhibits yet another slice of the overall space of TSL-2 patterns.

Can we say anything more specific about the subclass of TSL-2 which is representative of case? I believe that we can. As mentioned at the outset, I have assumed that case is fundamentally a relation between one or more DPs which indicates their relative position in some domain. If this is on the right track, it is reasonable to posit that the mapping from position to case must be *monotonic*. This would rule out the example in (35), in which the first and third DPs are mapped to nominative, with the second mapped to accusative, much like the *ABA constraint in morphology. Furthermore, this would be just one of several monotonicity requirements affecting case, along with morphological syncretism

(Graf 2019), visibility (Bobaljik 2008), and locality (Poole 2022; also see Graf 2020).

Another possible formal restriction relates to iteration/spreading.[14] In previous examples, only a single case in each daughter string language could iterate. I propose that this is a hard constraint. Suppose that we had a language which is like English except that *both* nominative and accusative iterate.

(36) $G^{+} = \{\rtimes \cdot \textsc{nom}, \underline{\textsc{nom} \cdot \textsc{nom}}, \textsc{nom} \cdot \textsc{acc}, \textsc{acc} \cdot \textsc{acc}, \textsc{nom} \cdot \ltimes, \textsc{acc} \cdot \ltimes\}$

Now, whenever there are three DPs in a clause, the middle may be either nominative or accusative, with no corresponding difference in structure or meaning; if there are four, the number of possibilities only increases. This seems extremely unnatural.

Intuitively, the order of cases assigned in a given domain should be *deterministic* in the sense that only one sequence is permitted for a string of a given length, modulo category and other features. Restricting iteration to a single case slot achieves this for the present example. More generally, it could be derived if the string-to-string mapping was found to be a *input tier-based strictly local* (ITSL) function in the sense of Burness et al. (2021). This provides further motivation for revisiting case as a transduction, as suggested in Section 2.6.

In summary, there may be some unattested case patterns that are TSL-2, but there are far more that are not. Once embedded in a larger theory, the unwanted patterns can be further whittled down by the other factors—computational or otherwise—for which we have independent evidence.

5 Conclusion

In this paper, I have argued that the typology of case dependencies in syntax is explained well by the assumption that they are TSL-2. Major points of cross-linguistic variation such as presence of iteration, the structural order among cases, the visibility of DPs and domain nodes, and interactions between DPs in the same domain naturally fall out from the parameters of the model, while many unattested patterns are excluded. Case therefore joins agreement, movement, and long-distance phonotactics in exemplifying the close link between formal typology and computational complexity, further strengthening the TSL-2 hypothesis.

This does not at all mean that the investigation of the subregular complexity of case is complete. As mentioned at the outset, the typological predictions of a multi-tier system of case assignment are not nearly as clear as those for a single tier. The same can be said of movement, agreement, and the interactions among these phenomena. Working this out will require a close examination of the grammars of individual languages, similar to the study of Japanese case in Hanson (2023). We should also carefully investigate the parallels between case and agreement, particularly complex DAM effects, which have been argued by theorists as a reason for unification under the Agree operation (cf. Clem and Deal 2024). The present perspective also argues for unification, but of a very different sort: given that case and agreement are built upon the same computations and both can involve multiple DPs which are adjacent on a tier, the existence of such parallels is exactly what we expect.

In addition, the present study highlights several directions for further investigation of the formal model. These include treating case in the mapping from the dependency tree to the phrase structure tree, determining how best to reconcile the tree tier model with the c-command constraint, and exploring other formal constraints such as monotonicity which may complement subregular complexity. Now that we have a clearer understanding of the aspects of case which are explained well by the TSL-2 model, this should help with identifying good candidate factors to handle the rest.

Acknowledgments

This work was partly supported by the Institute for Advanced Computational Science at Stony Brook University. My thanks to three anonymous reviewers, whose comments led to many improvements in the presentation of this paper.

References

Baker, Mark C. (2015). *Case*. Cambridge Studies in Linguistics 146. Cambridge University Press. DOI: 10.1145/1463891.1464000.

Baker, Mark C. and Nadya Vinokurova (2010). Two modalities of case assignment: Case in Sakha. *Natural Language & Linguistic Theory* 28.3, pp. 593–642. DOI: 10.1007/s11049-010-9105-1.

[14]This proposal builds on a suggestion by a reviewer that the order of cases must not be uniformly reflexive or irreflexive. The former constraint is too strong (just one case with iteration is reflexive) and the latter is too weak, as described in the text.

Bárány, András (2017). *Person, Case, and Agreement*. Oxford University Press, Nov. DOI: 10.1093/oso/9780198804185.001.0001.

Berro, Ane and Ricardo Etxepare (2017). Ergativity in Basque. *The Oxford Handbook of Ergativity*. Ed. by Jessica Coon, Diane Massam, and Lisa Demena Travis. Oxford University Press, Aug., pp. 782–806. DOI: 10.1093/oxfordhb/9780198739371.013.32.

Bickel, Balthasar and Johanna Nichols (2008). Case Marking and Alignment. *The Oxford Handbook of Case*. Ed. by Andrej L. Malchukov and Andrew Spencer. Oxford University Press, pp. 304–321. DOI: 10.1093/oxfordhb/9780199206476.013.0021.

Bittner, Maria and Ken Hale (1996). The Structural Determination of Case and Agreement. *Linguistic Inquiry* 27.1, pp. 1–68.

Blake, Barry J. (2001). *Case*. 2nd ed. Cambridge University Press.

Bobaljik, Jonathan (2008). Where's Phi? Agreement as a Postsyntactic Operation. *Phi Theory*. Ed. by Daniel Harbour, David Adger, and Susana Béjar. Oxford: Oxford University Press, May 1, pp. 295–328. DOI: 10.1093/oso/9780199213764.003.0010.

Boston, Marisa Ferrara, John T. Hale, and Marco Kuhlmann (2010). Dependency structures derived from Minimalist Grammars. *The Mathematics of Language. 10th and 11th Biennial Conference, MOL 10, Los Angeles, CA, USA, July 28-30, 2007 and MOL 11, Bielefeld, Germany, August 20-21, 2009, Revised Selected Papers*. Ed. by Christian Ebert, Gerhard Jäger, and Jens Michaelis. Lecture Notes in Computer Science 6149. Berlin, Heidelberg: Springer, pp. 1–12. DOI: 10.1007/978-3-642-14322-9_1.

Burness, Phillip Alexander et al. (2021). Long-distance phonological processes as tier-based strictly local functions. *Glossa: a journal of general linguistics* 6.1. DOI: 10.16995/glossa.5780.

Butt, Miriam (2006). *Theories of case*. Cambridge University Press.

Chomsky, Noam (1959). On certain formal properties of grammars. *Information and control* 2.2, pp. 137–167. DOI: 10.1016/S0019-9958(59)90362-6.

Chomsky, Noam (1981). *Lectures on government and binding: The Pisa lectures*. Berlin, New York: Walter de Gruyter.

Clem, Emily and Amy Rose Deal (2024). Dependent Case by Agree: Ergative in Shawi. *Linguistic Inquiry*, pp. 1–47. DOI: 10.1162/ling_a_00529.

Daniel, Penelope (2024). Unifying Differential Argument Marking through Interpretable Features. *Proceedings of the 41st West Coast Conference on Formal Linguistics*. Somerville, MA: Cascadilla Press, pp. 105–112.

De Santo, Aniello and Thomas Graf (2019). Structure Sensitive Tier Projection: Applications and Formal Properties. *Formal Grammar: 24th International Conference*. FG 2019. Ed. by Raffaella Bernardi, Gregory Kobele, and Sylvain Pogodalla. Berlin, Heidelberg: Springer, pp. 35–50. DOI: 10.1007/978-3-662-59648-7_3.

Goldsmith, John (1976). "Autosegmental phonology". PhD thesis. Cambridge, MA: Massachusetts Institute of Technology.

Graf, Thomas (2018). Why movement comes for free once you have adjunction. *Proceedings of CLS 53*. Chicago, IL: The Chicago Linguistic Society, pp. 117–136.

Graf, Thomas (2019). Monotonicity as an effective theory of morphosyntactic variation. *Journal of Language Modelling* 7.2, pp. 3–47. DOI: 10.15398/jlm.v7i2.211.

Graf, Thomas (2020). Monotonicity in Syntax. *Monotonicity in Logic and Language*. Ed. by Dun Deng et al. Berlin, Heidelberg: Springer, pp. 35–53. DOI: 10.1007/978-3-662-62843-0_3.

Graf, Thomas (2022a). Subregular linguistics: bridging theoretical linguistics and formal grammar. *Theoretical Linguistics* 48.3–4, pp. 145–184. DOI: 10.1515/tl-2022-2037.

Graf, Thomas (2022b). Typological implications of tier-based strictly local movement. *Proceedings of the Society for Computation in Linguistics (SCiL) 2022*. Amherst, MA: University of Massachusetts Amherst, pp. 184–193. DOI: 10.7275/gb65-ht31.

Graf, Thomas (2023). Subregular Tree Transductions, Movement, Copies, Traces, and the Ban on Improper Movement. *Proceedings of the Society for Computation in Linguistics (SCiL) 2023*. Ed. by Tim Hunter and Brandon Prickett. Amherst, MA: University of Massachusetts Amherst, pp. 289–299. DOI: 10.7275/TK1N-Q855.

Graf, Thomas and Kalina Kostyszyn (2021). Multiple Wh-Movement is not Special: The Subregular

Complexity of Persistent Features in Minimalist Grammars. *Proceedings of the Society for Computation in Linguistics (SCiL) 2021*. Ed. by Allyson Ettinger, Ellie Pavlick, and Brandon Prickett. Amherst, MA: University of Massachusetts Amherst, pp. 275–285. DOI: 10.7275/2dfv-v219.

Graf, Thomas and Connor Mayer (2018). Sanskrit n-Retroflexion is Input-Output Tier-Based Strictly Local. *Proceedings of SIGMORPHON 2018*. Association for Computational Linguistics, pp. 151–160. DOI: 10.18653/v1/W18-5817.

Graf, Thomas and Nazila Shafiei (2019). C-command dependencies as TSL string constraints. *Proceedings of the Society for Computation in Linguistics (SCiL) 2019*. Ed. by Gaja Jarosz et al. Amherst, MA: University of Massachusetts Amherst, pp. 205–215. DOI: 10.7275/4rrx-x488.

Hanson, Kenneth (2023). A TSL Analysis of Japanese Case. *Proceedings of the Society for Computation in Linguistics (SCiL) 2023*. Ed. by Tim Hunter and Brandon Prickett. Amherst, MA: University of Massachusetts Amherst Libraries, pp. 15–24. DOI: 10.7275/xqhr-r404.

Hanson, Kenneth (2024a). A Tier-Based Model of Syntactic Agreement. *Proceedings of CLS 60*. Chicago, IL: The Chicago Linguistic Society. To appear.

Hanson, Kenneth (2024b). Strict Locality in Syntax. *Proceedings of CLS 59*. Ed. by Kutay Serova and M. K. Snigaroff. Chicago, IL: The Chicago Linguistic Society, pp. 131–145.

Hanson, Kenneth (2025a). "The Subregular Complexity of Case and Agreement". PhD thesis. Stony Brook University.

Hanson, Kenneth (2025b). Tier-based strict locality and the typology of agreement. *Journal of Language Modelling* 13.1, pp. 43–97. DOI: 10.15398/jlm.v13i1.411.

Heinz, Jeffrey (2018). The computational nature of phonological generalizations. *Phonological Typology*. Ed. by Larry M. Hyman and Frans Plank. Phonetics and Phonology 23. De Gruyter Mouton, pp. 126–195. DOI: 10.1515/9783110451931-005.

Heinz, Jeffrey, Chetan Rawal, and Herbert G. Tanner (2011). Tier-based strictly local constraints for phonology. *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*. Ed. by Dekang Lin, Yuji Matsumoto, and Rada Mihalcea, pp. 58–64.

Hiraiwa, Ken (2001). Multiple Agree and the Defective Intervention Constraint in Japanese. *Proceedings of the MIT-Harvard Joint Conference (HUMIT 2000)*. Cambridge, MA: MITWPL, pp. 67–80.

Kishimoto, Hideki (2018). On exceptional case marking phenomena in Japanese. *Kobe Papers in Linguistics* 11, pp. 31–49.

Kobele, Gregory M. (2012). Eliding the derivation: A minimalist formalization of ellipsis. *Proceedings of the International Conference on Head-Driven Phrase Structure Grammar*. CSLI Publications, Oct. DOI: 10.21248/hpsg.2012.23.

Kuno, Susumu (1973). *The Structure of the Japanese Language*. Cambridge, MA: MIT Press.

Mahajan, Anoop Kumar (1990). "The A/A-bar distinction and movement theory". PhD thesis. Cambridge, MA: Massachusetts Institute of Technology.

Malchukov, Andrej and Andrew Spencer, eds. (2008). *The handbook of case*. Oxford: Oxford University Press.

Marantz, Alec (2000). Case and licensing. *Arguments and Case: Explaining Burzio's generalization*, pp. 11–30.

Mayer, Connor and Travis Major (2018). A challenge for tier-based strict locality from Uyghur backness harmony. *Formal Grammar 2018: 23rd International Conference*. FG 2018. Berlin, Heidelberg: Springer, pp. 62–83. DOI: 10.1007/978-3-662-57784-4_4.

McMullin, Kevin and Gunnar Ólafur Hansson (2016). Long-Distance Phonotactics as Tier-Based Strictly 2-Local Languages. *Proceedings of the Annual Meetings on Phonology*. Ed. by Adam Albright and Michelle A. Fullwood. Vol. 2. Washington, DC: Linguistic Society of America. DOI: 10.3765/amp.v2i0.3750.

Poole, Ethan (2022). Improper case. *Natural Language & Linguistic Theory* 41.1, pp. 347–397. DOI: 10.1007/s11049-022-09541-6.

Rezac, Milan, Pablo Albizu, and Ricardo Etxepare (2014). The structural ergative of Basque and the theory of Case. *Natural Language & Linguistic Theory* 32.4, pp. 1273–1330. DOI: 10.1007/s11049-014-9239-7.

Saito, Mamoru (2015). Cartography and Selection: Case Studies in Japanese. *Beyond Functional Sequence*. Oxford University Press, May,

pp. 255–274. DOI: 10.1093/acprof:oso/9780190210588.003.0014.
Saito, Mamoru, T.-H. Jonah Lin, and Keiko Murasugi (2008). N'-Ellipsis and the Structure of Noun Phrases in Chinese and Japanese. *Journal of East Asian Linguistics* 17.3, pp. 247–271. DOI: 10.1007/s10831-008-9026-8.
Stabler, Edward P. (1997). Derivational minimalism. *Logical Aspects of Computational Linguistics. First International Conference, LACL '96, Nancy, France, September 23-25, 1996. Selected Papers*. Ed. by Christian Retoré. Lecture Notes in Computer Science 1328. Berlin, Heidelberg: Springer. DOI: 10.1007/BFb0052152.
Stabler, Edward P. (2011). Computational perspectives on Minimalism. *Oxford handbook of linguistic Minimalism*. Ed. by Cedric Boeckx. Oxford: Oxford University Press, pp. 617–643. DOI: 10.1093/oxfordhb/9780199549368.013.0027.
Vu, Mai Ha, Nazila Shafiei, and Thomas Graf (2019). Case assignment in TSL syntax: A case study. *Proceedings of the Society for Computation in Linguistics (SCiL) 2019*. Ed. by Gaja Jarosz et al. Amherst, MA: University of Massachusetts Amherst, pp. 267–276. DOI: 10.7275/sywz-xw23.
Witzlack-Makarevich, Alena and Ilja A. Seržant (2018). Differential argument marking: Patterns of variation. *Diachrony of differential argument marking*. Ed. by Ilja A. Seržant and Alena Witzlack-Makarevich. Studies in Diversity Linguistics 19. Berlin: Language Science Press. DOI: 10.5281/zenodo.1219168.
Yip, Moira, Joan Maling, and Ray Jackendoff (1987). Case in Tiers. *Language* 63.2, pp. 217–250.
Zaenen, A., J. Maling, and H. Thráinsson (1985). Case and Grammatical Functions: The Icelandic Passive. *Natural Language & Linguistic Theory* 3.4, pp. 441–483.