

Case Attraction in Ancient Greek

Scott Grimm
sgrimm@stanford.edu

Department of Linguistics, Stanford University

Abstract. Case attraction has stood as a puzzling, and elusive, oddity of older Indo-European languages. This paper focuses on attraction in Ancient Greek, establishing both the regularity of the operation and its underlying motivation. A novel method is proposed for grounding case in terms of a feature-based representation of agentivity properties, loosely based on Dowty's proto-role theory, but reformulated in terms of privative opposition and hierarchically organized via a lattice. This structure is then used to model the case system of Ancient Greek and derive a hierarchical ordering on the case system in terms of agentivity. Modelling the interaction between this hierarchy and the other factors involved in case attraction in the Optimality Theory framework yields a full solution, predicting both its distribution and frequencies therein.

The attempt to describe case as a stable, syntactic phenomenon is belied by instances of what is known as *case conflict*. This paper investigates a particular type of case conflict, case attraction, which involves relative pronouns and their antecedents. Case attraction has long been seen as an exception to the general rule of case assignment, but I will argue in what follows that, once the conditions under which case attraction are clarified, and a given case's relation to semantic content is secured, case attraction must no longer be seen as an aberration, but rather as consistent with the general principles of case assignment.

The organization of the paper is as follows. I begin by summarizing the data from Ancient Greek, paying particular attention to the distribution of the phenomenon and the frequencies therein. Case attraction will be seen as crucially linked to the thematic, or agentivity, properties associated with a given case, and a general framework for connecting case and agentivity via a lattice structure will be exposed in section 2. The case system of Ancient Greek is then mapped upon the lattice, determining a hierarchical relation on the case system. Section 3 unifies the foregoing analysis of case attraction with the semantic agentivity properties to account for case attraction and its distribution within the Optimality Theory (OT) framework.

1 Delimiting Case Attraction

Languages which dispose of case systems use case in order to display the syntactic function of lexical items within the clause. Yet, case assignment can also be

subject to agreement constraints, e.g., an adjective modifying a noun must share the noun’s case, if both are marked for case. When the agreement constraints are discordant with the functional role of case assignment, case conflict occurs. This conflict is commonly resolved by selecting to manifest either syntactic function or agreement. In most circumstances, Ancient Greek selects the former, and while a relative pronoun agrees with its antecedent in gender, number, and person, its case is determined by the construction of the clause in which it stands. However, under certain circumstances, Greek prioritizes agreement, as in (1)¹, where the case of the pronoun shifts to agree with the nominal. Although the focus here is on Greek, such fluctuation also occurs in other languages, e.g., Anglo-Saxon, Old High German, and Latin.

- (1) Xenophon, *Anabasis*, 1.7.3

andres	axioi	tēs	eleutheriās	hēs
Men.NOM	worthy.NOM	the.GEN	freedom.GEN	which.GEN
[hēn]	kektēsthe			
[which.ACC]	possess.2nd.PL			

Men worthy of-the freedom which you possess.

I will refer to the type of attraction in (1) as ‘proper attraction’. Additionally, it occurs that it is not the relative pronoun which is provoked into another case, but the antecedent which shifts case to accord with the case of the relative pronoun, as in (2)². Since the roles are reversed, this variety of attraction is known as ‘inverse attraction’.

- (2) Sophocles, *Oedipus Rex*, 449

ton andra	touton	[ho anēr	touto]	hon	palai
the man.ACC	this.ACC	[the man.NOM	this.NOM]	who.ACC	long-ago
zēteis	...houtos	estin	enthade		
search-2nd.	...this one.NOM	is	here		

The man who you long ago searched ... is here.

Inverse attraction “regularly occurs when the antecedent stands at the head of the sentence and precedes the relative clause, which itself precedes the main clause” [4], i.e., the antecedent is in a focus position and distanced from its governing verb.

¹ The examples of case attraction phenomena were gathered from grammars [3], [8] and [13], and cross-checked against a corpus of the relative pronoun in Xenophon’s *Anabasis* gathered from [10].

² It will be noted that this example contains a resumptive pronoun, which would lead one to posit that inverse attraction can be analyzed as left-dislocation, as was done by [11] for Old and Middle High German. However, resumptive pronouns, while possible, do not appear to be the rule in Greek (cf. (4) below), and so left-dislocation will not suffice to explain the data at hand.

The languages which have been the source of discussion for case attraction provide little in the way of variation. If a language exhibits proper attraction, it exhibits inverse attraction as well, as is the case in Latin, Greek, and Old High German.³ This indicates that an account of attraction would optimally posit the same underlying reason for proper and inverse attraction. I now examine the conditions which appear to trigger case attraction based on data from Ancient Greek.

1.1 The Distribution of Cases

The details of the distribution in terms of case provide the main challenges for an account. Proper attraction canonically occurs when the relative pronoun is in the accusative and the antecedent is in the dative or genitive, as in (1).⁴ Less frequently, relative pronouns standing in the nominative or the dative are attracted. Inverse attraction also affects the accusative most frequently, which is then coerced into a genitive or dative. However, it is possible with the nominative as well, which can be realized as an accusative. The possible combinations and attested attractions are summarized in table 1 below.

Table 1. Distribution of Case Attraction

antecedent	relative pronoun	output pair	antecedent	relative pronoun	output pair
nominative	nominative	—	nominative	dative	no change attested
accusative	nominative	no change attested	accusative	dative	no change attested
dative	nominative	(dat, dat)	dative	dative	—
genitive	nominative	(gen, gen)	genitive	dative	(gen, gen)
nominative	accusative	(acc, acc)	nominative	genitive	(gen, gen)
accusative	accusative	—	accusative	genitive	(gen, gen)
dative	accusative	(dat, dat)	dative	genitive	(gen, gen)
genitive	accusative	(gen, gen)	genitive	genitive	—

It is important to note that there are types of attraction not found—one does *not* see a relative pronoun in the nominative attracting an antecedent in the genitive into the nominative.⁵ This can be explained systematically if a case hierarchy which orders the cases is adopted, as was done in [7]:

³ A notable exception is Anglo-Saxon, where one does not find inverse attraction. But this arises for entirely different reasons—the relative particle ‘þe’ is indeclinable, so it does not have a case with which to attract the nominal.

⁴ It must be noted that case attraction, which in certain circumstances is expected, is ultimately optional in Ancient Greek. [3, 51.10.2] notes that a lack of attraction tends to add emphasis, and is primarily found in the works of the Greek orators.

⁵ Note that case attraction does not coincide with the ‘inherent case’ and ‘structural case’ partition [7], since an accusative can attract a nominative (cf. (4) below), yet both are ‘structural’ cases.

nominative < accusative < dative < genitive

This hierarchy can be adduced from the table of the distribution of case attraction. A comparison among the input/output pairs in table 1 makes it evident that attraction only occurs in Ancient Greek when the relative pronoun or antecedent can take a case that is located higher on the case hierarchy. In addition, it is identical to so-called markedness hierarchies of case found in the literature (cf. [14]).

While the case hierarchy has indeed captured a generalization about how case attraction operates, this in itself has not led to a satisfactory account of what underlies case attraction. More generally, earlier attempts to explain case attraction have left room for improvement in two directions. Some accounts, despite their virtues, have not incorporated the case hierarchy, thereby missing an aspect of the phenomenon's regularity and the ability to predict which case attracts which, e.g., [2], who proposes a Principles and Parameters account. [6] models the conflict between a relative pronoun displaying syntactic function versus agreement in terms of Optimal Theoretic parsing, yet without reference to the case hierarchy, their analysis will over-generalize to instances in which attraction has not been attested in Greek. On the other hand, the hierarchy has been used by itself to explain case attraction [7]. Yet, there are two fundamental issues at stake when using this case hierarchy as an explanatory device for case attraction. First, claiming that a certain case is more marked than another leaves open the question of what actually underlies this markedness. In other words, what are the principles upon which the hierarchy is founded? Second, if it were only a matter of blindly applying the hierarchy to clauses conjoined by a relative pronoun, the disparity between the frequency of attraction from the accusative and from the nominative and dative is left unexplained. The first sort of attraction is the most frequent, indeed regular, while the latter two are rare.

1.2 Agentivity and Prominence

Further insight into the factors at play can be gleaned by examining the argument structures of the examples, in particular, regarding what type of thematic content is associated with the attracted items. The pronouns that underwent proper attraction referred to arguments which would have been quite low in agentivity—often referring to the object of verbs such as 'legō' ('to say') or 'echō' ('to have'), i.e., patients, broadly speaking. This observation holds for attracted items of all cases. No pronoun which referred to an accusative argument that would have been high in agentivity was found to undergo proper attraction. A nominative is attracted only when the argument is the subject of a passive or middle verb, where the grammatical subject of the verb is not an agent. Finally, attraction from the dative only seems to occur when the argument represented by the pronoun refers to the theme or beneficiary, as in (3), where attraction affects the direct object of the verb 'entetucheka' ('meet with').

- (3) Plato, Republic 531e

oligoi hōn [toutōn hois] ego entetucheka
 few.NOM who.GEN [those.GEN which.DAT] I.Nom meet.PERF

A few of those whom I have met with.

It can be adduced that proper attraction is contingent on the relative pronoun representing an argument low in agentivity. So it is with inverse attraction which was only observed with subjects that are low in agentivity: subjects of passive constructions or of the “to be” copula, or unaccusatives, as in (4).

(4) Isokrates 6. 48

polīteiān [polīteiā] hoian einai
 Constitution.ACC [Constitution.NOM] of.such.a.sort.ACC to.be.INF
 chrē para monois hēmin estin
 necessary.3rd for alone us.DAT is

We alone have a constitution such as it ought to be.

Intuitively, this is what one would expect. It has often been remarked (cf. [1] and references there) that the subject of a sentence is the least marked and most prominent argument. Conversely, the object, typically associated with arguments low in agentivity, is more marked yet less prominent. So it stands to reason that case attraction, which results in a case-marking which ranks higher on the case hierarchy, would most likely affect the class of arguments that is most apt to be marked, those low in agentivity.

All the instances of attraction share another characterization: all the attracted items are in positions which indicate high discourse prominence—either (topicalized) subjects or heads of relative clauses. Therefore, a generalization arises: attracted items are low in agentivity and high in discourse prominence. This gives cause to suspect that case attraction has a functional explanation—these two competing factors, low agentivity and high prominence, are disharmonious, and set the conditions for attraction to occur. After a detailed examination of the connection between case and agentivity, I will give this generalization a more precise formulation.

2 Case and Agentivity

A full account of case attraction must both incorporate the case hierarchy and at the same time constrain its application, and ideally demonstrate what the case hierarchy is grounded in. Having noted above agentivity constrains the possibility of attraction, it is plausible that a more precise account of the connection between case and agentivity can meet the above requirements. The following section pursues this connection and will ultimately demonstrate that the degree of agentivity associated with a case determines its position on the case hierarchy. This will simultaneously provide an explanation for the observed frequency patterns of case attraction.

The connection between case and agentivity follows from cases' relation to argument structure. An argument structure representation of a predicate states that the predicate requires certain types of participants as its subject, object, etc.—e.g., the verb *hit* in English requires that the subject be an agent, one that performs the action, and that the object be a patient, one who submits to the effects of the action. Marking argument structure is the primary reason for having a case system in the first place—to signal what is the subject, object, etc., of the predicate. Since argument structure is determined by thematic content, i.e., agentivity and affectedness, it follows that case assignment is determined in part by thematic content as well. In order to arrive at a framework capable of modelling fine-grained interaction between the parameters of agentivity and affectedness, I begin by employing a set of event-based properties entailed by the verb, inspired by the approach of [5].

I assume a set of properties which refer to modes of participation in events: *instigation*, *motion*, *sentience*, *volition*, and different degrees of *persistence*. *Instigation* entails any argument effecting the event designated by the predicate. *Motion* is entailed just in case the argument is required to be in motion. *Sentience* designates conscious involvement in the event [12] while *volition* designates deliberate engagement in the event. Agents, then, will typically possess one or more of these properties.

Persistence is a two-tiered notion, for something can persist existentially, that is, its essence remains the same throughout the event/state, or it can persist qualitatively—i.e., it persists in all its particulars. Either of these can obtain at the beginning and/or the end of the event—in terms of features, we have the following set: *existential persistence (beginning)*, *existential persistence (end)*, *qualitative persistence (beginning)*, and *qualitative persistence (end)*.

Establishing agentivity properties in this manner leads to two diametrically opposed classes in privative opposition, one a full agent possessing all the properties, and the other not entailing any, not even independent existence—e.g., arguments of negative existence statements or incorporated/cognate objects (“sing a song”). Affectedness can be reformulated as a lack of persistence during the event; further, this feature configuration is able to capture the different degrees of affectedness with respect to existence. Totally affected patients, e.g., verbs of destruction/consumption (‘destroy’, ‘eat’) entail that their object argument persists existentially at the beginning of the event, but not at the end. Patients which are partially affected (e.g., objects of verbs such as ‘damage’ or ‘move’) persist existentially throughout the event, but do not persist qualitatively, i.e., they are changed in some manner. Unaffected entities, most often agents, persist both existentially and qualitatively throughout the event. The opposition between agents and patients falls out from this feature system in that agents will possess total persistence along with a number of other agentivity properties while patients will generally possess no properties save initial persistence and possibly *qualitative persistence (beginning)*.

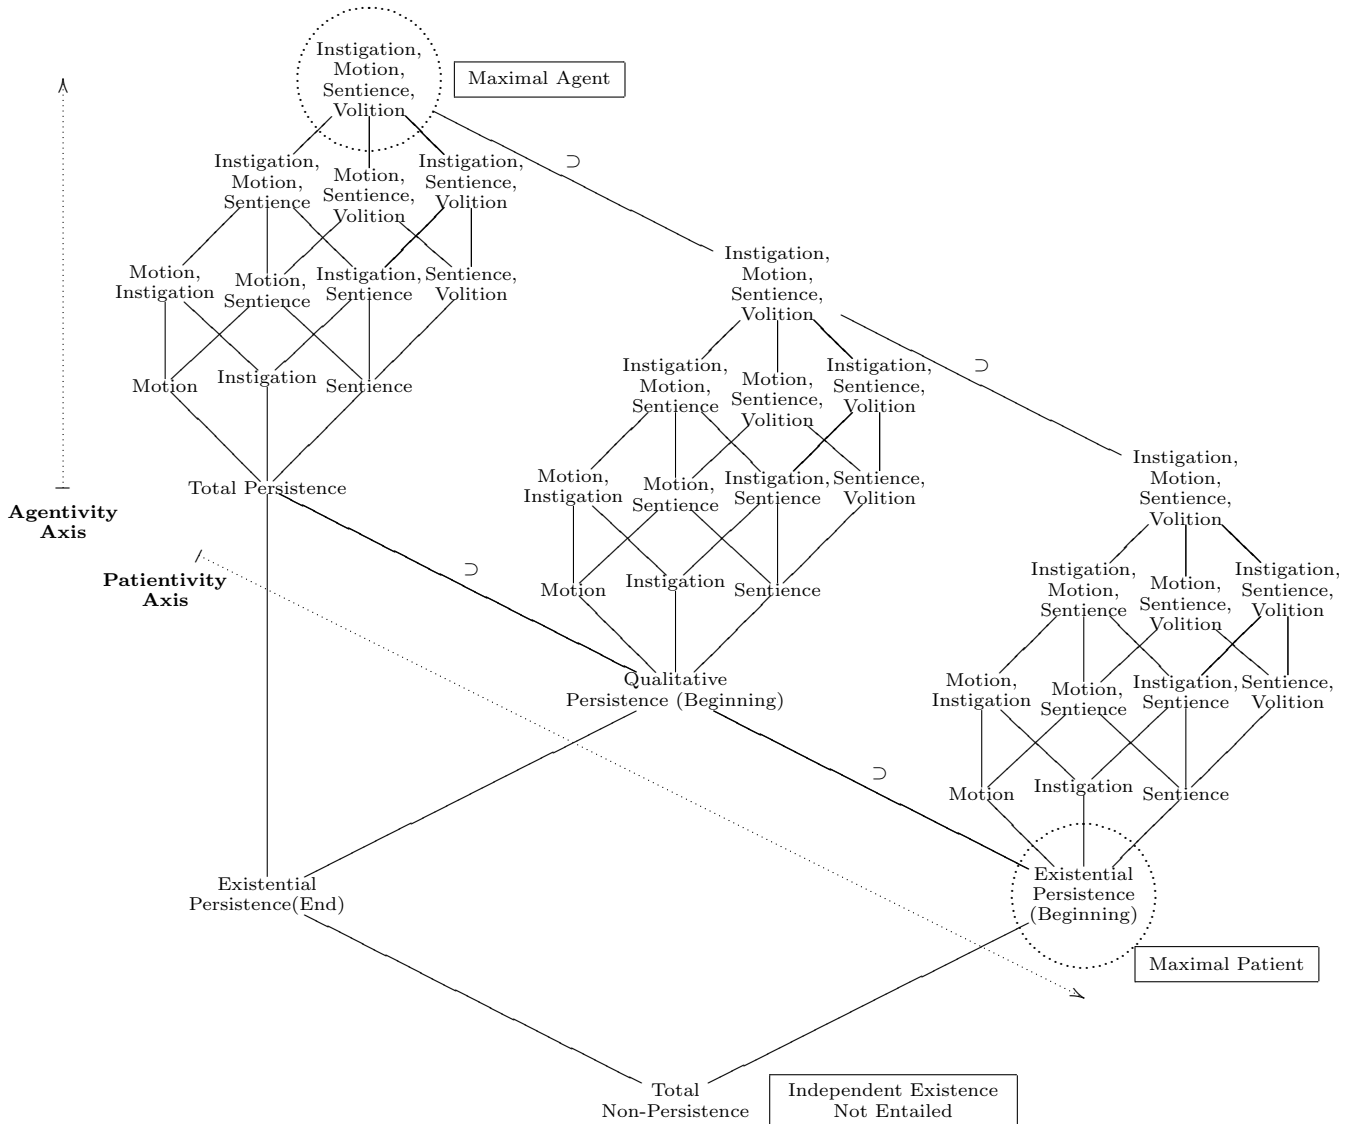


Fig. 1. The Agentivity Lattice

Hierarchization of Agentivity Properties The above has established a set of properties which make up a predicate's argument structure. Logical entailments among the eight features constrain the combinations possible. For instance, *volition* entails *sentience*, since only sentient beings are capable of volition, and *-existential persistence (end)* entails *-qualitative persistence (end)*, since if an entity does not exist at the end of the event, clearly none of its qualities do either. The remaining combinations can then be given greater structure. The sets

of agentivity and persistence properties can be separately ordered by inclusion, giving rise to a lattice structure for each. The Cartesian product of the agentivity and persistence structures results in a larger lattice, shown in figure 1, referred to henceforth as the agentivity lattice. Note that the privative opposition is conspicuous in the structure: the highest node contains all the features (the full agent), the lowest contains none (event internal objects). Further, this lattice exhibits the possible space of argument structure with respect to agentivity.

2.1 Mapping the Cases of Ancient Greek

Turning to mapping the cases of Ancient Greek on the lattice, the methodology is rather straight-forward. First, a case's primary function is located on the agentivity lattice, as will be shown in figure 2. Second, the case-marker is identified with the semantic features of its location. It is then incumbent on those features to provide an explanation for the appropriateness of that case for any secondary uses it has accumulated.

The simplest mapping is the accusative. As the marker of the direct object par excellence, arguments in the accusative are canonically affected by the event and non-agentive, i.e., not possessing any agentive properties. These considerations confine the accusative to the region covering the nodes *Qualitative Persistence (Beginning)* and *Existential Persistence (Beginning)*, which designate that the argument has been affected, but does not possess any agentivity properties.

The dative in its central usage marks the indirect object and “denotes that *to* or *for* which something is done” [13]. The dative appears also as the object of such verbs as ‘benefit’, ‘help’, ‘injure’, ‘meet’, ‘obey’, ‘pardon’, ‘trust’. Notice that in the usual use of these verbs, the object will be sentient and affected by the event. Therefore, the mapping of the core use of the dative is to the *Sentience* node within the *Qualitative Persistence (Beginning)* branch.

In Greek, however, the dative case is syncretic with two other cases, the instrumental and the locative (the latter of which is not relevant here). The former can be located on the lattice in the following fashion. First, prototypical instruments are not sentient, so the possible region is already confined to the lower-third of the agentivity lattice. Second, prototypical instruments are viewed as persisting throughout the event, i.e., if a tree is cut with an axe, the axe persists throughout the cutting event; therefore, prototypical instrumentals would be located on the *Total Persistence* branch of the lattice. Instrumentals would appear equally capable of at least co-instigation along with an understood agent, therefore the instigation node is included.

The genitive in its most central (adnominal) usage marks possession, the sentient possessor being put into the genitive case.⁶ Two further uses of the genitive as a verbal argument reveal a propensity towards high levels of agentivity. First, in most predicates where the object denotes the external cause of the event, the object is put into the genitive, e.g., verbs of emotion and perception (‘to hear’).

⁶ The genitive also expresses partitivity, yet since I am interested in how the genitive relates to agentivity properties, I leave aside the partitive usages.

Second, the passive construction in Ancient Greek puts the demoted agent in the genitive case, preceded by the preposition *hupó*. Thus, when appearing as a verbal argument, excepting partitive uses, the genitive is highly agentive, entailing *sentience* and/or *instigation*.

The adnominal genitive primarily denotes static relations, in which neither the head noun nor the genitive-marked noun undergo any change, e.g., inalienable possession, relations of source, of measure, or of quality. When governed by a verb, the genitive also shows a propensity towards total persistence. For instance, as mentioned, ‘to hear’ takes its object in the genitive. The object heard will persist throughout the hearing event, while the hearer will be affected, and similarly for verbs of emotion. One does not see the genitive marking arguments that are affected or undergo change, unless it falls in with the partitive usage, e.g., ‘to touch’. Therefore, the genitive can be mapped, in its possessive uses to the node of the agentivity lattice containing the combination *Sentient* and *Total Persistence* while its agentive uses are mapped to the node containing the combination *Sentient*, *Instigation* and *Total Persistence*.

The nominative serves to mark subjects and while it ends up most frequently marking agentive arguments, this is only because subjects tend to be agentive. However, the nominative marks subjects of verbs in the passive as well, which are typically patients. Since the nominative can mark any level of agentivity, which is *not* true for the other cases, the nominative is not associated with any particular region of the lattice, i.e., the nominative does not mark agentivity.

Functional Overlap of the Dative and the Genitive The method followed above began with the primary function of a case and derived the semantic properties of the case from that function. In order to ensure that these mappings are valid in general, we must check that secondary uses are in accord with the mappings. An obstacle would appear to arise in that certain functions are designated by more than one case, in particular, the genitive and the dative can both express possessive relations and agents of an event. An examination of the details shows, however, that the properties ascribed to the cases hold, and correspond to the nuances which distinguish the different cases’ uses.

Possession marked by the genitive tends towards inalienable possession, designating a static relation which in some sense characterizes the possessor [3]. In contrast, possession marked by the dative tends towards alienable possession, designating “that something is at the disposal of a person or has fallen to his share temporarily” [13, 1480]. These distinctions correspond to the genitive and dative’s above mappings onto *Total Persistence* and *Qualitative Persistence (Beginning)*, respectively, since inalienable possession is unchanging, while alienable possession is potentially dynamic, e.g., one has acquired something.

There is a similar division between the genitive of agent and the dative of agent. The genitive with the preposition *hupó* is the default case to mark a demoted agent explicitly connected with the event. The dative of agent, when used with persons, is highly restricted, only appearing when the verb is in the perfect or pluperfect tense, and when the subject of the verb is impersonal [13,

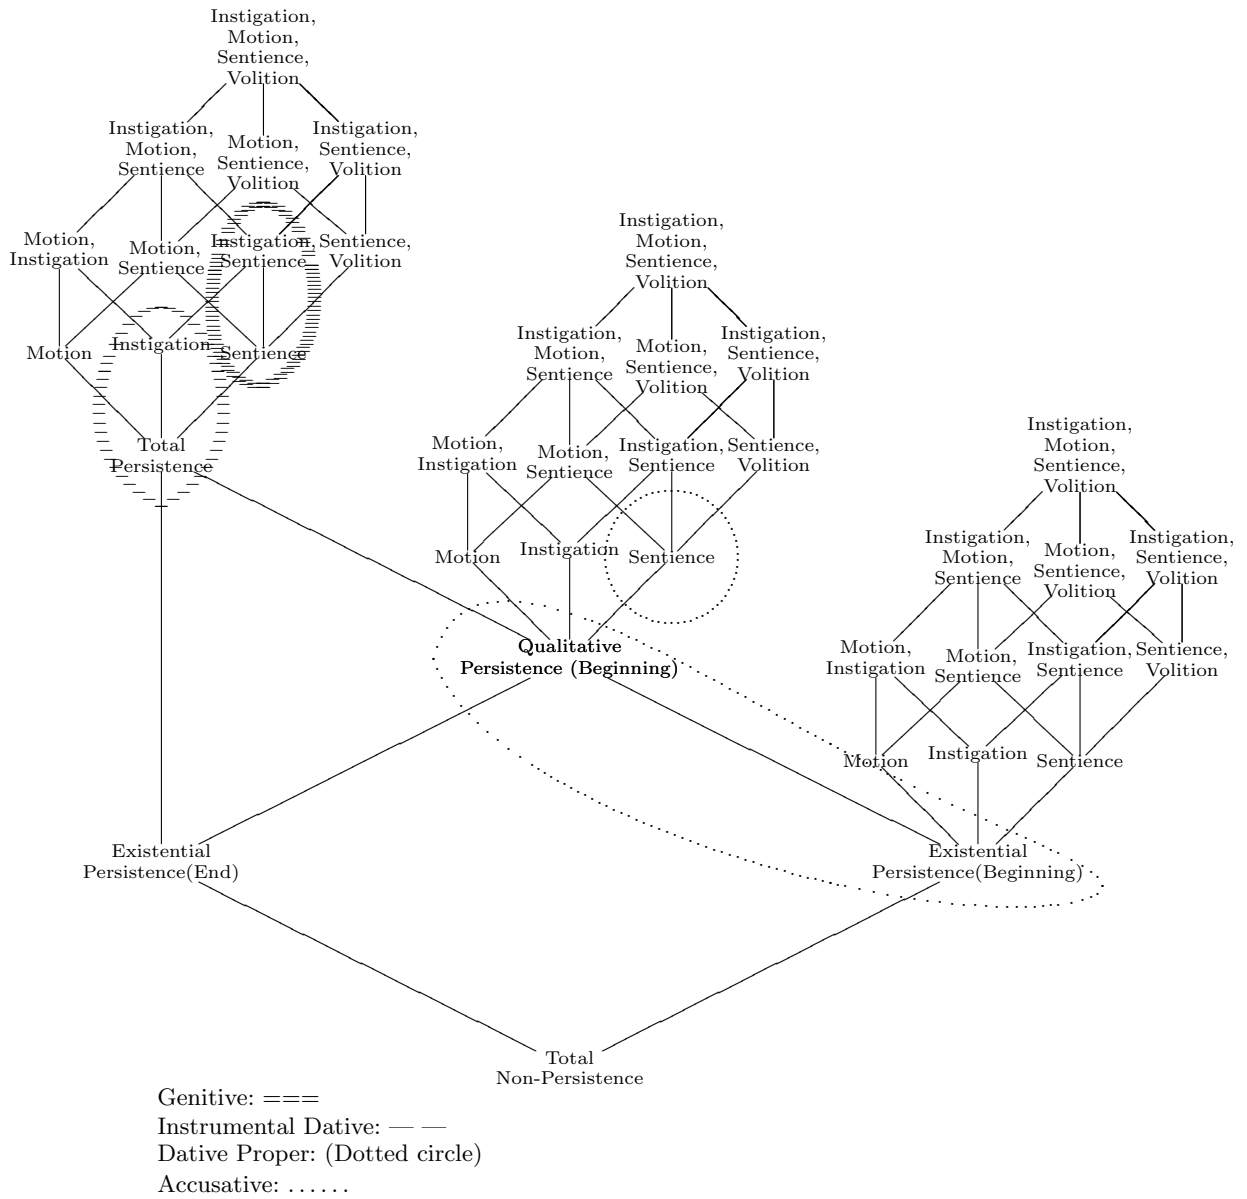


Fig. 2. Cases of Ancient Greek

1492]. This usage, although named the ‘dative of agent’, has much in common with the dative’s more principal use of marking beneficiaries, since “the notion of agency does not belong to the dative, but it is a natural inference that the person interested is the agent” [13, 1488]. Therefore, the dative of agent does not diverge from the region of the lattice ascribed to the core usage of the dative.

When the dative of agent is a thing, the dative is used whether the subject is personal or impersonal, corresponding to its instrumental use.

The above demonstrates that although the dative shares functions with the genitive, the nuances of these usages are consistent with agentivity properties of its primary uses. In summary, the following mappings have been established, corresponding to figure 2:

Genitive:

Possessive uses: *Sentient* and *Total Persistence*

Agentive uses: *Sentient*, *Instigation* and *Total Persistence*

Dative:

Possessive uses: *Sentient* and *Qualitative Persistence (Beginning)*

Agentive uses: *Instigation* and *Total Persistence*

Accusative:

Qualitative Persistence (Beginning) or *Existential Persistence (Beginning)*

2.2 Deriving the Case Hierarchy

The above has established mappings of the cases of Greek onto the agentivity lattice. Since the lattice is ordered by inclusion, then the regions associated with the cases are also ordered by inclusion:

$$\textit{nominative} \subset \textit{accusative} \subset \textit{dative} \subset \textit{genitive}$$

But then, this translates directly into the case hierarchy:

$$\textit{nominative} < \textit{accusative} < \textit{dative} < \textit{genitive}$$

This result independently motivates the case hierarchy in terms of agentivity rather than resting on claims that some cases are inherently more marked or more oblique than others. Founding the case hierarchy in terms of agentivity gives an explanation for the observed frequency of attraction, viz. genitives are never attracted, datives and nominatives are rarely attracted, and accusatives are regularly attracted. Since the genitive case is the most agentive, there is no other more agentive case to which it could be attracted. Datives can only be attracted by the genitive case and free uses of the dative (e.g., instrumental, dative of agent) were not observed to be attracted. The dative as a direct object, i.e., low in agentivity, is only found with a limited number of verbs, thereupon ensuring that attraction of the dative will be similarly limited. Nominatives are most frequently subjects in active clauses, i.e., agents. Less frequently they are subjects low in agentivity (e.g., passive constructions) and, accordingly, less frequently attracted. Accusatives, however, are standardly patients/themes, i.e., low in agentivity, and therefore they are liable to be attracted to more agentive cases.

That agentivity is the underlying force behind the realization of prominent item's case is reasonable given that an analogous situation holds for subjects,

another prominent position. The generalizations behind the various thematic hierarchies have made clear the primacy of agentivity for subject selection, i.e., the most agentive argument is realized as subject. A wider generalization arises: in prominent positions, there is a preference for agentive arguments. Thus, when, say, a relative pronoun is capable of manifesting its relation to its antecedent or its syntactic function, it simply selects whichever is most agentive, highlighting the relation which is most active in the construction.

3 An OT Analysis

In section 1, several characteristics of case attraction were isolated. First, two conflicting responsibilities of relative pronouns and antecedents were noted: to designate their syntactic function and to agree. The presence of such competing factors suggests an analysis within the Optimality Theory (OT) framework, which can model such competition.

Second, it was noted that attracted items are low in agentivity and high in discourse prominence. OT permits a method of modelling the interaction between these two factors via the technique of “harmonic alignment”. The essential idea is that a prominent element in one category combines most harmoniously with a prominent element in another category.⁷

Agentivity and discourse prominence can be put in terms of prominence scales, following [1], where ‘X’ designates an element high in discourse prominence. By harmonic alignment, the harmonic scales in (5) are derived:

	<i>Prominence Scales</i>	<i>Harmonic Alignment</i>
(5)	Agent>Patient X>x	Agent/X>Patient/X Patient/x>Agent/x

Case attraction relates to the first of the two harmonic alignment scales in (5), occurring in contexts where the relative pronoun or subject would be marked as “Patient/X”. Since this situation is disharmonious, it is sought to be avoided. This is the motivation for case attraction to occur in the first place.

Section 2 established the case hierarchy as an agentivity hierarchy, and so this harmonic alignment scale can be expanded into (6).

(6)	Genitive/X>Dative/X>Accusative/X>Nominative/X
-----	---

⁷ The full definition is given in [9, p.21] : Harmonic Alignment : “Given a binary dimension D1 with a scale X>Y and another dimension D2 with a scale a>b> >z, the harmonic alignment of D1 and D2 is the following pair of harmony scales:

$$H_x = X/a > X/b > \dots > X/z$$

$$H_y = Y/z > \dots > Y/b > Y/a$$

The constraint alignment is the following pair of constraint hierarchies:

$$C_x = *X/z >> \dots >> *X/b >> *X/a$$

$$C_y = *Y/a >> *Y/b >> \dots >> *Y/z$$

This alignment scale simply states that if an element is high in discourse prominence, then it is more harmonious to be in the genitive case than in the dative case and so on, which makes the needed bridge between case assignment and the conflict between high discourse prominence and low agentivity.

The competition between designating syntactic function and agreement can be captured by the two following constraints, which compete for assigning case to relative pronouns and their antecedents (see a slightly different take on these constraints in [6]):

FAITH-CASE: lexical items retain the case assigned to them in INPUT (i.e., their structurally assigned case)

AGREE-CASE-Rel.Pronoun-Antecedent: Relative pronouns and their nominal antecedents agree in case

These constraints assume the class of relational structures in INPUT contains information for grammatical and thematic relations. By the principles of OT, these two constraints will be ranked in some order. If FAITH-CASE outranks AGREE-CASE, the case of the relative pronoun will always display the case of its syntactic function within the relative clause. Recalling that attraction is never seen applying to agents, the FAITH-CASE constraint is contextually restricted to agents and non-agents, following the strategy of [14], and FAITH-CASE_{Agent} is ranked highest, ensuring that an agent retains its case-marking. The competition is then seen to be between FAITH-CASE_{Non-Agent} and AGREE-CASE. For Ancient Greek, these two constraints are unordered, so whether the relative pronoun agrees with the antecedent or not is dependent on other factors, and in all cases at least one of the constraints is violated.

These other factors are exactly the case hierarchy in harmonic alignment with the discourse prominence scale, which is ranked below the other constraints. Thus, if an item is disharmonious in that it is low in agentivity and high in discourse prominence, it seeks to become more harmonious by adopting a case with a higher agentivity level, made available by the antecedent in the instance of proper attraction, yet at the cost of violating FAITH-CASE. If the item cannot become more harmonious by agreeing with the case of the antecedent (or relative pronoun for inverse attraction), then it retains its case, since a shift in case would lead to greater disharmony. In this latter scenario, AGREE-CASE is violated, but FAITH-CASE is not. Thus, the winner will be determined by the additional violation marks incurred by way of the case hierarchy. This is shown in tables 2 and 3 which correspond to examples 1 and 4, respectively.

4 Conclusion

The above has proposed a solution to the case attraction puzzle by grounding case in agentivity, postulating a preference for agentive arguments in prominent positions, and viewing case attraction as a resolution of that preference, which can then be represented formally within OT by the technique of harmonic

Table 2. Proper Attraction Tableau

Input	F-CASE _{Agent}	F-CASE _{Non-Agent}	AGREE	*Nom/X	*Acc/X	*Dat/X	*Gen/X
(i) tēs eleutherias hē (NOM)		*	*!	*			
(ii) tēs eleutherias hēs (GEN)		*					*
(iii) tēs eleutherias hēi (DAT)		*	*!			*	
(iv) tēs eleutherias hēv (ACC)			*		*!		

Table 3. Inverse Attraction Tableau

Input	F-CASE _{Agent}	F-CASE _{Non-Agent}	AGREE	*Nom/X	*Acc/X	*Dat/X	*Gen/X
(i) politeia (NOM) hoian			*	*!			
(ii) politeias (GEN) hoian		*	*!				*
(iii) politeiai (DAT) hoian		*	*!			*	
(iv) politeian (ACC) hoian		*			*		

alignment. There are several important advantages of this solution. First, proper attraction and inverse attraction are explained by the same mechanism, and the similarity felt to exist between the two phenomena is justified. The technique of harmonic alignment has led to a functional explanation underlying case attraction which makes this solution more satisfactory than merely stating that cases prefer to agree with more oblique cases if possible, which one is forced to conclude if one relies solely upon the case hierarchy. This functional explanation answers why attraction only occurs in this context and why it only affects items which are low in agentivity, in turn, explaining the frequencies of the cases attracted. Third, a theoretical advance has been made by independently deriving the case markedness hierarchy from agentivity properties. No appeal to “greater obliqueness” or similarly vague concepts need be countenanced to make use of the case hierarchy, for it can now be used with the understanding that it is a notational tool, founded on the same semantic principles which underlie argument selection.

Acknowledgements Many thanks to Henk Zeevat under whose supervision this work was carried and also to Andrej Malchukov for discussion. This work

was presented at the Sixth International Tbilisi Symposium on Language, Logic and Computation (September 2005) and the NWO-DFG Workshop “Modelling Incremental Interpretation” at Radboud University Nijmegen (November 2005). I would like to thank the audiences for their helpful comments on those occasions, and in particular, those of Helen de Hoop, Barbara Partee, Matthias Schlesewsky and Robert Van Valin.

References

- [1] Aissen, Judith: Markedness and Subject Choice in Optimality Theory. *Natural Language and Linguistic Theory* **17** (1999) 673–711
- [2] Bianchi, Valentina: Some Issues in the Syntax of Relative Determiners. In *The Syntax of Relative Clauses*. Artemis Alexiadou, Paul Law, Andre Meinunger and Chris Wilder, eds. John Benjamins (2000) 53–81
- [3] Cooper, Guy: *Attic Greek Prose Syntax*, Vol. 1. The University of Michigan Press (1998)
- [4] Diggel, James: Xenophon, *Anabasis* 3, 1, 6-8 and the Limits of Inverse Attraction. *Studi Italiani di filologia classica* **20** No. 1-2 (2002) 83–86
- [5] Dowty, David: Thematic Proto-roles and Argument Selection. *Language* **67** (1991) 547–619
- [6] Fanselow, Gisbert, Matthias Schlesewsky, Damir Čavar and Reinhold Kliegl: *Optimal Parsing: Syntactic Parsing Preferences and Optimality Theory*. Unpublished Ms. (1999)
- [7] Harbert, Wayne: Case Attraction and the Hierarchization of Case. In *Proceedings of the Eastern States Conference on Linguistics Vol. 6* CLC Publications (1990) 138–149
- [8] Kühner, Raphael and Gerth, Bernhard: *Ausführliche Grammatik der griechischen Sprache*. Darmstadt (1966)
- [9] McCarthy, John: *A Thematic Guide to Optimality Theory*. Cambridge University Press (2002)
- [10] Perseus Digital Library: <http://www.perseus.tufts.edu/>
- [11] Pittner, Karen: The Case of German Relatives. *Linguistic Review* Vol. 12/3 (1995) 197–231
- [12] Rozwadowska, Bożena: Thematic Restrictions on Derived Nominals. In *Syntax and Semantics vol. 21* Wilkins, Wendy, ed. Academic Press (1988) –147-165
- [13] Smyth, Herbert W.: *Greek Grammar*. Harvard University Press (1920)
- [14] Woolford, Ellen: Case Patterns. In *Optimality Theoretic Syntax* MIT Press (2001) eds. Legendre, G., Vikner, S. and Grimshaw, J. 509–543