On the origin of constraints

Sander Lestrade, Geertje van Bergen, and Peter de Swart

8.1 Introduction

This chapter addresses two debates central to Optimality Theory. The first concerns the nature of constraints (cf. Tesar and Smolensky 2000; Boersma and Hayes 2001; Jäger 2003, 2007). According to one position, the set of constraints is innate and only their ranking has to be determined by the learning child; the other position has it that the set of constraints is functionally motivated and developed by generalizing over input data. The second debate pertains to the nature of the optimization process (Blutner 2000; Blutner et al. 2006). Again two hypotheses can broadly be distinguished. Some argue that the process is unidirectional, going from form to meaning or from meaning to form; according to others the process is bidirectional, going from meaning to form and back again. In this chapter, we will link the two discussions (cf. also Blutner 2007; Zeevat and Jäger 2002). We will argue that specific unidirectional constraints can be derived from bidirectional optimization processes and are thus functionally motivated (i.e., following from the use of language). To prevent this from being an abstract exercise only, we will show how this would work for *differential case marking (DCM)*.

We will first give a brief introduction to DCM. Next, we will discuss the differences between unidirectional and bidirectional optimization procedures. Finally, we will show how unidirectional constraints can be derived from generalizing over bidirectional optimization processes.

8.2 Differential case marking

In many languages, case marking is used only selectively; for example, case marking of a verb's arguments may depend on their animacy values. If the differential use of case concerns the marking of a transitive object, this is called *differential object marking (DOM)*.[1] For example, in Malayalam, all and only animate objects are obligatorily marked, as can be seen in (1).

Malayalam (Dravidian, India; Asher and Kumari 1997: 203)

(1)	a.	*Avan*	*oru*	*paʃuvin-e*	*vaɲɲi.*
		he	a	cow-ACC	buy.PAST

[1] We will restrict our analysis to the differential marking of objects, although everything we say should similarly hold for *differential subject marking (DSM)* too. For a more elaborate introduction into differential case marking, cf. Aissen (2003); de Swart (2007); for a more general discussion on case alternations, cf. Lestrade (2013).

‘He bought a cow.’

b. *ɲaan teeɲɲa vaɲɲi.*
I coconut buy.PAST
‘I bought a coconut.’

If case marking is dependent on properties of a single argument only, as in example (1), this is called *local* DCM. Local DCM contrasts with *global* DCM, in which case marking is dependent on properties of both arguments. For example, in Awtuw, if two unmarked NPs are present, then the one higher in animacy is interpreted as the transitive subject (in which Human > Animate > Inanimate). Equal animacy results in an interpretation with a conjoined subject, as the contrast between (2a) and (2b) shows.

Awtuw (Papuan, Papua New Guinea; Feldman 1986: 110)

(2) a. *Tey tale yaw dæl-i.*
3.F.SG woman pig FAC-bite-PAST
‘The woman bit the pig.’ not: ‘The pig bit the woman.’

b. *Piyren yaw di-kæl-iy.*
dog pig FAC-IMP-bite-IMP
‘The dog and the pig bite.’ not: ‘The dog is biting the pig/The pig is biting the dog.’

An object marker can be used to overrule this interpretational preference. Accusative case is obligatory when the object (*O*) equals or is higher than the transitive subject (*A*) in the Animacy Hierarchy.[2] This is illustrated in the following examples (Feldman 1986: 110):

(3) a. *Tey tale-re yaw dæl-i.*
3.F.SG woman-ACC pig FAC.bite-PAST
‘The pig bit the woman.’

b. *Piyren-re yaw di-k-æl-iy.*
dog-ACC pig FAC-IMP-bite-IMP
‘The pig is biting the dog.’

[2] We follow conventions in the functional-typological literature and use A for the agent (subject) of a transitive clause, O for the patient (object) of a transitive clause, and S for the single argument of an intransitive clause.

Finally, consider global DCM in Fore, in which multiple strategies are used to resolve ambiguity. Standardly, higher animates (humans > animates > inanimates) are interpreted as having higher roles (subject > indirect object > direct object; Scott 1978: 114–16). Differently from the previous language, linear order is decisive in case of a draw (subjects preceding direct objects). Only when the interpretative hierarchy or linear order needs to be overruled, ergative case *-(wa)ma* is used (cf. Donohue and Donohue 1997 for an alternative analysis).

Fore (Trans-New Guinea, Papua New Guinea; Scott 1978: 115–16)

(4) a. *Yaga:-wama* *wá* *aegúye.*
pig-ERG man 3.SG.OBJ.hit.3.SG.SU.IND
'The pig attacks the man.'

b. *Yaga:* *wá* *aegúye.*
pig man 3.SG.OBJ.hit.3.SG.SU.IND
'The man kills the pig.'

Below, we will show how the different DCM systems can be seen as diachronic variants of the same underlying system, in which local DOM uses an extra constraint that is derived from a bidirectional optimization process. First, however, we will show how the above variation is generally dealt with in OT.

8.3 The state of the art

Differential case marking has been a prominent topic in Optimality-Theoretical approaches to grammar, initiated by the work of Aissen (2003). Proposed OT analyses have associated the two case-marking systems introduced above with their own type of optimization procedure. For local DCM, a unidirectional analysis has been proposed (e.g., Aissen 2003); for global DCM, a (semi-)bidirectional analysis has been argued to be necessary (e.g., de Swart 2007; cf. also de Hoop and Malchukov 2007, 2008).

8.3.1 Unidirectional OT and local DCM

Recall the examples of local DCM from Malayalam. For such variation, a simple unidirectional analysis seems to suffice. In unidirectional OT, the goal is to determine the best form for a meaning (*OT syntax*, e.g. Grimshaw 1997) or the best meaning for a form (*OT semantics*, Hendriks and de Hoop 2001). In such an approach, the optimization procedure for the production of (1a) 'He bought a cow' would look as in Tableau 8.1. This tableau essentially presents the analysis of Aissen (2003), but with different names for the constraints involved.[3]

PROD: buy(he, cow)	ANIMO→MARKER	ECON
a. *he cow buy*	*!	
→ b. *he cow.ACC buy*		*

Tableau 8.1 Unidirectional productive optimization procedure for *buy(he, cow)* in local DCM

In this optimization procedure, a general economy constraint ECON(OMY) prefers unmarked form candidates above case-marked ones. Because this constraint is outranked by a constraint that requires animate objects to be case-marked, ANIMO→MARKER, the second candidate becomes optimal and accusative case is used in this situation.

Obviously, form candidates for meaning inputs with inanimate objects will vacuously satisfy ANIMO→MARKER. In these cases, ECON can decide, as illustrated in Tableau 8.2. Thus, Malayalam will not use accusative case on coconuts.

PROD: buy(he, cow)	ANIMO→MARKER	ECON
→ a. *I coconut buy*		
b. *I coconut.ACC buy*		*!

Tableau 8.2 Unidirectional productive optimization procedure for *buy(he, coconut)* in local DCM

Given the use of these forms, the optimization procedure from form to meaning also results in the right outcome. To this end we employ two constraints. The general constraint FAITHL which requires hearers to interpret everything the speaker has said (cf. Zeevat 2000; see also Zeevat and Jäger 2002) is violated by any interpretation in conflict with the provided morphosyntactic information, e.g., assignment of the subject function to an argument marked

[3] A single-headed arrow "→" is used to indicate unidirectional optimality. Note further that we include harmonically bounded candidates in our tableaux, since the choice between case-marked and bare arguments is fundamental to our argument.

with accusative case. The other semantic constraint is PROMINENCE (cf. de Hoop and Lamers 2006; see also the BIAS constraint in Zeevat and Jäger 2002), that tells the hearer to interpret as the A that argument that is highest in animacy. This constraint is violated by interpretations in which O is higher in animacy than A.

When applied to the form *he cow-ACC buy* marked with accusative case (which was found optimal in Tableau 8.1), we arrive at the right interpretation, as Tableau 8.3 shows. This is due to the fact that FAITHL requires the accusative case-marked argument to be assigned the O-function. Likewise, absence of accusative case will result in the right interpretation when the O-argument is coconut, because otherwise PROM will be violated; cf. Tableau 8.4. This interpretation would of course have also come out as optimal if accusative had been used on the object. This shows that use of case marking is indeed not necessary to express this meaning.

INT: *he cow.ACC buy*	FAITHL	PROM
→ a. buy(he, cow)		
b. buy(cow, he)	*!	*

Tableau 8.3 Unidirectional interpretive optimization procedure for *he cow.ACC buy* in local DCM

INT: *I coconut buy*	FAITHL	PROM
→ a. buy(I, coconut)		
b. buy(coconut, I)		*!

Tableau 8.4 Unidirectional interpretive optimization procedure for *I coconut buy* in local DCM

8.3.2 Bidirectional OT and global DCM

Global DCM was illustrated in the previous section with examples from Awtuw. Rather than assuming constraints that require the marking of only those objects that outrank their subjects in animacy, these systems seem to involve an additional interpretational check. The use of overt case marking is determined by taking into account the hearer's perspective to see if the optimal candidate indeed leads to the intended result. Consequently, bidirectional OT analyses have been proposed for these kinds of systems (de Swart 2007; de Hoop and Malchukov 2007, 2008). In bidirectional OT, syntax and semantics are dependent on one another, as the outcome of one direction of optimization constrains the outcome of the other direction (Blutner 2000; Blutner et al. 2006). Different models of bidirectional optimization have been proposed in recent years to account for a wide variety of phenomena (see Beaver and Lee 2004 and the contributions to Benz and Mattausch 2011 for discussion; see contributions to this volume by Hoek and de Hoop, de Swart, Legendre et al., and Hendriks for applications of weak bidirectional optimization). We will employ here an asymmetric version of bidirectional OT similar in spirit to earlier work (Smolensky 1998; Donohue 1999; Zeevat 2000; Wilson 2001; Jäger 2003). See de Swart (2011) for a discussion of the relation of this model with respect to other versions of bidirectional OT. We have used this model in earlier work to model (differential) case marking and word order phenomena (see de Swart 2007, 2011; Lestrade 2010; van Bergen 2011).

In this model the outcome of the production component is constrained by the interpretational component, but not (necessarily) vice versa. More specifically, a form *f* is bidirectionally optimal for a given meaning *m* iff the meaning *m* is uniquely recoverable from that form *f* and there is no form *f′* which is less marked than *f*—i.e., a better form from the viewpoint of productive optimization, and from which *m* is uniquely recoverable. A meaning *m* is uniquely recoverable from a form *f* iff it is the unique optimal candidate in the interpretive optimization of *f*. Hence, a form which is optimal from the production perspective can be rejected as the output candidate when it results in the wrong interpretation, i.e., an interpretation different from the one intended. As a result, a candidate which is suboptimal from the production perspective can become bidirectionally optimal (given that it does express the intended meaning).

Consider the optimization procedure for (2a) 'The woman bit the pig' in Tableau 8.5. We will make use of the constraints introduced in the previous section.

PROD: bit(woman,pig)	ECON	FAITHL	PROM
↔ a. *woman pig bit*			
b. *woman pig-OBJ bit*	*!		
INT$_a$: *woman pig bit*	ECON	FAITHL	PROM
→ i. bit(woman, pig)			
ii. bit(pig, woman)			*!
INT$_b$: *woman pig-OBJ bit*	ECON	FAITHL	PROM
→ i. bit(woman, pig)			

ii. bit(pig, woman)		*!	*

Tableau 8.5 Bidirectional optimization procedure for *bit(woman, pig)* in global DCM

First, at the production stage (PROD), the unmarked candidate is found optimal because of ECON. When considering the interpretation of this initially preferred candidate, PROM(INENCE) excludes the interpretation in which the pig acts on the woman (INT_a). Since this leads to the desired result, the preferred form can be used indeed, which is indicated with the two-sided arrow '↔' at the production stage (as INT_b shows, the case-marked candidate would have led to the right interpretation too).

Now contrast this with the bidirectional optimization procedure for (3a) 'The pig bit the woman' in Tableau 8.6, where an object marker is necessary.

PROD: bit(pig, woman)	ECON	FAITHL	PROM
→ a. *woman pig bit*			
↔ b. *woman-OBJ pig bit*	*!		
INT_a: *woman pig bit*	ECON	FAITHL	PROM
→ i. bit(woman, pig)			
ii. bit(pig, woman)			*!
INT_b: *woman-OBJ pig bit*	ECON	FAITHL	PROM
i. bit(woman, pig)		*!	
→ ii. bit(pig, woman)			*

Tableau 8.6 Bidirectional optimization procedure for *bit(pig, woman)* in global DCM

Again, the unmarked candidate is preferred for economy reasons, as is indicated with the simple arrow at the production stage. This time, however, this form does not lead to the correct interpretation, as its evaluation at INT_a shows: PROM would lead to an interpretation in which the woman is the A. Production candidate *b*, which makes use of an object marker, does yield the intended interpretation that it was the pig that bit the woman (INT_b). By FAITHL the correct interpretation is ensured. Thus, in order to be understood, the unmarked candidate cannot be used but the speaker has to resort to a more explicit way of expression (again indicated by the two-sided arrow at PROD).

In sum, we have shown how local DCM systems can be modeled using unidirectional optimization through the stipulation of a constraint such as ANIMO→MARKER that applies to a subset of arguments. Global DCM, by contrast, requires a bidirectional analysis with an additional interpretational check to ensure the optimal candidate leads to the intended meaning.

We see two problems with the above state of affairs. First, a constraint such as ANIMO→MARKER seems rather ad hoc and should preferably be independently motivated.[4] Second, it is undesirable to propose completely different optimization procedures, unidirectional and bidirectional, for a very similar phenomenon in different languages. Preferably, there is one and the same system that only differs in the (ranking of its) constraints. Fortunately, both problems can be solved at once, as we will show in the following section.

8.4 From bidirectional optimization to unidirectional constraints

In this section we will show how domain-specific constraints need not be stipulated but can be functionally motivated as the result of a generalization over bidirectional optimization procedures in which only very general constraints are used initially. Although we believe that this claim holds true more generally (cf. Haspelmath 1999, Zeevat and Jäger 2002, and Jackendoff 2002 for similar ideas), we will restrict ourselves in this chapter to the application to case marking.

We assume that the following general principles are universally present in the (human) mind. First, there is an economy principle, which is also known as the *principle of least effort* (cf. Zipf 1965). People will naturally choose the path of least resistance to arrive at some goal. This principle is generally accepted in the literature, and probably does not need much further introduction.

Next, we have a desire for communicative success. This principle too should hardly need motivation. The goal of human communication is to transfer an idea to (the mind of) a receiver, and the only natural way to do so is via encoding that idea in a signal. If there was no desire for having the decoded signal come close to what was originally meant and encoded, there would not be any communication in the first place (cf. Pagin 2008 for a more sophisticated discussion). Indeed, it probably is not so much the desire for success when communicating that one may question about. Instead, the real question is whether language basically is for communication or to organize one's thoughts for oneself. Even if one

[4] Aissen (2003) derives a similar constraint on the basis of the process of harmonic alignment. This approach, however, stipulates the (universal) availability of prominence hierarchies and a (broad) set of associated constraints.

assumed the latter (something we would disagree with), we could still say that whenever language is used for communication, people make sure that it works (given what we have just said).

Finally, humans have a disposition to generalize. This disposition is humorously illustrated by the following quote from the RationalWiki entry of *pattern recognition*: “If one thinks one sees something coming at one’s head, it is better to duck and be wrong, than not duck and be wrong.”[5] More serious motivations for the primacy of this principle that one may encounter are that “[categorization] is the basis for the construction of our knowledge of the world. It is the most basic phenomenon of cognition” (Cohen and Lefebvre 2005: 2), “survival and success in the world depend on making judgments that are as accurate as possible given the limited amount of information” (Krynski and Tenenbaum 2007: 430), or that generalization allows for the use of a maximum amount of information with minimal cognitive effort (Rosch 1978).

For language, these principles can be illustrated by the fact that frequent words are short (cf. the principle of economy), the only way to develop a lexicon is by negotiating meaning (cf. the desire for communicative success), and without generalizations, there would not be words in the first place (cf. the disposition to generalize).

Below, it will be shown how we can describe and derive the different types of case-marking systems discussed in section 8.3, using only the just-mentioned general principles (an economy principle, a desire for communicative success, and a disposition to generalize). A local system can be shown to follow from a global system in which one and the same solution is used to resolve ambiguity (i.e., the Awtuw type). Although the origins and historical development of DCM systems in individual languages are often still unclear, some diachronic data can be shown to mirror our developmental sketch (see Morimoto and de Swart 2006 for a discussion of DOM in Spanish).[6]

8.4.1 Simulating the development of a local DCM system

Equipped with the three general principles just introduced, a learning hearer will observe in the input from their language environment that most animates by far are subjects (both intransitive (S) and transitive (A)), whereas most objects are inanimate (cf. Table 8.1; cf. Zeevat and Jäger (2002) and Øvrelid (2004) for concrete corpus counts underlying this abstraction).[7]

[5] <http://rationalwiki.org/wiki/Pattern_recognition>, consulted on June 14, 2013.

[6] There is also evidence for a development from a global to a local function of other types of argument marking, including direction marking (Zúñiga 2006: 248–9) and agent-focus marking (Stiebels 2006).

[7] This assumes that a child can already distinguish animate from inanimate referents. Indeed, this distinction has been shown to emerge in early infancy (Gelman and Opfer 2002).

	S/A	O
animate	very many	few
inanimate	few	many

Table 8.1 Abstract crosstab for observed relation between animacy and function

The most prominent generalization that follows from these data is that animate event participants are likely to be the subject (cf. Comrie's 1989 notion of *natural transitivity* and the PROMINENCE constraint, which was used in the previous section too).

When growing up in a DOM language environment in which there is a single strategy to resolve ambiguity, our learning hearer will next observe that if this marker is used, this is always done to mark objects (Table 8.2).[8]

	S/A	O
Ø	very many	many
marker	-	few

Table 8.2 Abstract crosstab for the use of markers and function in a single-strategy DOM environment

Given our DOM environment, this does not hold the other way around: because of the economy principle, Os remain unmarked whenever possible.

Thus, our learning hearer is left with two generalizations, summarized in (5).

(5) a. Generalization I: animate → S/A
b. Generalization II: marker → O

Generally, language learners are not only listening to input, they will (try to) produce utterances themselves too. Doing so, they will wish to verbalize different types of events, predominantly intransitive events (with an S argument only), transitive events in which the A outranks the O in animacy ($A >_{animacy} O$), and transitive events in which this is the other way around, the O being equal or higher in animacy ($A \leq_{animacy} O$). Because of the economy principle, learning speakers will prefer to leave arguments unmarked. When checking the communicative success of the expressions of their preference, they will observe that for situations in which there is only one animate argument which functions as the transitive subject, the correct interpretation will follow by Generalization I (animate→S/A). However, when there is more than one animate argument or when the single animate argument functions as the object, this generalization does not lead to the right interpretation. For

[8] In this and the following tables, the hyphen is used to indicate that the relevant (object) marker is never used to mark subjects or agents. This does not exclude the possibility of using other markers in passive-like constructions.

communication of such situations, something extra needs to be done. Fortunately, the solution is provided by Generalization II (marker→O): If an additional object marker is used, the correct interpretation should follow.

As the reader may have noticed, this is nothing but the prose version of the bidirectional optimization process that describes *global* DCM discussed around Tableau 8.6: in case of ambiguity, the object needs to be marked; otherwise, the arguments can be left unmarked.

The optimization process by a learning speaker just described can easily be simulated by a computer. We can let a computer algorithm randomly produce sentences, check the expected communicative success for each utterance using Generalization I, use an O marker in cases of failure (Generalization II), and, most importantly, have it track the use of object marking throughout a series of optimization procedures. This last step is crucial for our argument: on the basis of the evolving record of optimization procedures that thus emerges, new generalizations are made that can be used as constraints in later optimization processes. It will be shown that the constraint ANIMO→MARKER that was used in the analysis of local DCM system, can be seen as such a generalization.

In the production of sentences, our computer algorithm will first randomly provide situations to be described following the assumed distribution of interest in (6). Although purely hypothetical, these numbers are inspired by work of Comrie (1989) and corpus research of Zeevat and Jäger (2002) and Øvrelid (2004) (cf. also the contributions in de Swart et al. 2008). The underlying assumption here is that people are mostly interested in other people and their interactions, and it is mostly animates that cause the action in the world, inanimates being acted upon.

(6)	a.	p(Subject=Anim)=.2	(e.g., *John walks.*)
	b.	p(Subject=Inan)=.1	(*The statue stands*.)
	c.	p(Subject=Anim&Object=Anim)=.2	(*John sees Pete*.)
	d.	p(Subject=Anim&Object=Inanim)=.4	(*John sees the statue*.)
	e.	p(Subject=Inanim&Object=Anim)=.05	(*The statue pleases John*.)
	f.	p(Subject=Inanim&Object=Inanim)=.05	(*The statue resembles a car*.)

That is, we assume that people talk about simple properties of animates in approximately 20 percent of the utterances, and about simple properties of inanimates in 10 percent of their utterances, animates are said to act on other animates in 20 percent of the utterances, etc. The exact numbers are of minor importance; it is mostly the relative proportions that matter. Specifically, animates are mostly subjects.

By default, the most economical way of expression is tried, as this is preferred by ECONOMY. That is, at this stage of the production, and at this point in the development of the grammar, the algorithm initially generates an unmarked form candidate for each meaning input.

In the next step, the expected communicative success of the utterance is determined.[9] There are two ways of doing this, both of which lead to the same results. The simpler version would be to assume communicative success with an unmarked expression whenever the A argument outranks the O in animacy, and only to use a marker when this is not the case (the O argument outranking the A, or both arguments being equal in animacy). Given the assumed world in (6), this would mean that a marker is necessary in 30 percent of the utterances, in which an animate O is involved in 5/6 of the cases. This simple version, however, ignores lexical specifications that one may argue play a role. For example, the verb *to see* requires an animate A, the verb *to please* an animate O.

To preempt possible concerns about lexical specifications of the verb involved, the figures below are based on a more complex version taking lexical specifications of the verb into account. Generalization I still applies and could be seen as summarizing the difference in type frequency of the different subcategorization preferences of verbs. This time, however, we do not assume that all unmarked sentences with animate As will be understood correctly, nor that all unmarked sentences with inanimate Os will. First, if there is only one argument, this necessarily is the subject and therefore success is guaranteed (this holds both for the simple and the complex version). But now, even if the two arguments are equal in animacy, a marker may not always be necessary to disambiguate. For example, if a man and a cow are involved in a buying event (cf. (1a)), we can predict on the basis of what we know about buying that it is probably the cow that is bought. Also, whereas the role assignment still mostly follows from Generalization I if the A argument outranks the O in animacy, this does not hold when the verb is specified for the opposite animacy-role distribution. That is, lexical specification but not Generalization I requires a marker in cases in which it is for some reason really meant that a cow buys someone or that John pleases the statue. These situations can be expected to be very rare, but we will assume a rather large proportion of 1 percent of the cases nevertheless to show the robustness of the simulation. The more important difference with the simple simulation, however, is that in most unmarked sentences in which an inanimate A and an animate O are intended we will assume that this is interpreted correctly because of the lexical specification of the verb (again, cf. *to please*). Specifically, we will assume that only in 40 percent of the sentences with an animate O and an inanimate A an object marker is deemed necessary to ensure the correct interpretation. Similarly, we assume that in 5 percent of the cases in which two animate arguments figure, verbal lexical specification can still tell them apart.

Table 8.3 summarizes these revised assumptions (the proportions in the second column follow from (6)).

Sentence type	Proportion of	Success?

[9] Note that we thus implement the desire for communicative success as a bidirectional check for correct interpretation; cf. Levelt (1989) for a comparable claim.

	utterances	yes	no
Single argument	.3	1	0
$A >_{animacy} O$	.4	.99	.01
$A <_{animacy} O$	.05	.6	.4
animate A&O	.2	.05	.95
inanimate A&O	.05	0	1

Table 8.3 Communicative success of unmarked expressions combining Generalization I and verbal lexical specification for different types of sentence types

As said above, in those cases in which communicative success is not achieved with an unmarked sentence, the algorithm will try again using an object marker (cf. Generalization II in (5)). As a result, all situations eventually will be successfully "communicated."

TYPE	ARGUMENT(S)	ROLES CLEAR?	USE MARKER?	LOG ODDS OBJECT MARKER		
				overall	AnimO	InanimO
intrans	S=anim	yes	no	0.00	0.00	0.00
intrans	S=anim	yes	no	0.00	0.00	0.00
trans	S=anim; O=inanim	yes	no	-0.69	0.00	−0.69
trans	S=anim; O=anim	no	yes	0.00	0.69	−0.69
trans	S=anim; O=anim	no	yes	0.41	1.10	−0.69
intrans	S=inanim	yes	no	0.41	1.10	−0.69
trans	S=inanim; O=inanim	no	yes	0.69	1.10	0.00
intrans	S=anim	yes	no	0.69	1.10	0.00
trans	S=anim; O=inanim	yes	no	0.29	1.10	−0.41
trans	S=inanim; O=inanim	no	yes	0.51	1.10	0.00

Table 8.4 First ten rows of randomly produced output of computer simulation of DOM system

Table 8.4 should make the simulation process more concrete. The utterances are generated by the algorithm row by row. The input for the production, given in the first two columns, is a random draw from the world specified in (6). Next, the algorithm checks if the roles of the two arguments can be expected to be understood correctly. If they are, an unmarked form is used; if the roles are unclear, they are specified by an object marker. Finally, as the last block of columns shows, the developing log odds that an object marker is used are tracked for transitive clauses. In order to have the log odds start with zero, prior to the simulation a "history" is assumed of one unmarked and one marked utterance. As only transitive utterances are taken into consideration, the log odds change only after the third utterance in which no marker is used, as a result of which the overall score becomes log(1/2) (i.e., one marked utterance vs two unmarked utterances) = −.69. As this utterance involves an inanimate object, the log odds that an object marker is used for inanimate objects change accordingly.

If we now generalize over many more optimization procedures (cf. our disposition to generalize), keeping track of the type of arguments that require a marker, the pattern illustrated in Figure 8.1 emerges.

[INSERT FIGURE 8.1 HERE]

In this plot, the developing log odds of five series of 1000 optimization procedures are plotted, distinguishing between the overall odds that a marker is used (dark-grey circle), the odds that a marker is used for an animate object (light-grey triangle), and the odds that a marker is used on an inanimate object (black cross; corresponding to the last three columns of Table 8.4). As is clear, the five series yield very comparable outcomes, suggesting that our findings are reliable. Also, it can be seen that the overall log odds and the odds that a marker is used for an inanimate object quickly stabilize. Overall, the log odds that a marker is used settle just below zero. Undoing the logarithmic transformation, this means that these odds are slightly lower than one to one. The odds that a marker is used for animate objects increase quickly and stabilize around two on a logarithmic scale, meaning that a marker is necessary for roughly seven out of eight animate objects.

This pattern can easily be understood. Both animate and inanimate objects sometimes are in need of an object marker. Because of verbal lexical specification, this is mostly the case when the arguments are equal in animacy (if they're not, the verb is likely to tell who does what). As we talk much more often about animates than about inanimates interacting, an object marker on inanimates is hardly ever necessary. Instead, animate objects much more often cause ambiguity and therefore animate objects will often need marking.

The crucial point of Figure 8.1 is that a third generalization presents itself very clearly: animate objects need object marking. When we formulate this third generalization into a constraint, we have functionally derived ANIMO→MARKER, which was stipulated in the analysis of local DCM above.

8.4.2 Simulating the maintenance of a global DCM system

For the development of a unidirectional constraint, it is crucial that there is a convention of using one and the same resolution strategy in cases of communicative failure, i.e., the object marker. Only then, the odds for the use of the marker given a certain type of input can increase to such an extent that a generalization about their relation is possible. As we saw in section 8.2, some languages combine different strategies to achieve communicative success. For example, Fore can use both word order and case marking as cues to tell apart A from O; cf. example (4). In this section, we will show that unidirectional constraints are unlikely to develop from such systems.

Much of the discussion in the previous section remains the same. We do not expect situations being talked about to differ between languages, so we will take the relative frequencies of the assumed world in (6) as a starting point for this simulation. What will differ, however, are the proportions of expected success of unmarked expressions. Note that we use *unmarked* here to refer to expressions in which no object marker is used, irrespective of word order. Consider Table 8.5. The crucial difference between this table and Table 8.3 in the previous version is that we now assume word order to disambiguate between arguments of equal animacy half of the time (cf. the fourth and fifth row). As discussed in section 8.2, sometimes markers are still necessary—for example, when word order is used to serve a different function, e.g. for information-structure purposes. (In fact, “Success?” may not be the best column label here, but we maintain it for ease of comparison with Table 8.3.)

Sentence type	Proportion of utterances	Success? yes	no
Single argument	.3	1	0
$A >_{animacy} O$	.4	.99	.01
$A <_{animacy} O$	.05	.6	.4
animate A&O	.2	.5	.5
inanimate A&O	.05	.5	.5

Table 8.5 Communicative success of unmarked expressions combining Generalization I and verbal lexical specification, and word order for different types of sentence types

If we now again run a simulation making use of these numbers, a very different pattern emerges, as shown in Figure 8.2. Differently from the previous simulation, the log odds for a marker for animate objects now linger around zero. From such a state of affairs, there is no generalization presenting itself: a constraint that directly assigns an object marker to animate objects will never emerge. This shows that our algorithm can model both local and global DCM systems.

[INSERT FIGURE 8.2 HERE]

8.5 Implications for the optimization process

Thus far, we have shown how domain-specific constraints need not be postulated but can be derived from a generalization about bidirectional optimization processes. Obviously, this is highly desirable for reasons of parsimony and cognitive plausibility. As we will show now, there is another important gain. By adding this constraint to our grammar, we can speed up the bidirectional optimization process as we no longer first have to consider an inadequate candidate. In the tableaux below, we will show this gain by going step-by-step through the semi-bidirectional optimization procedure before and after the derivation of the constraint ANIMO→MARKER, forcing the use of the object marker on animate Os.

Consider the following example from Malayalam, a language with a local DCM system (cf. section 8.2).

(7) *Avan* *kuṭṭiy-e* *aṭiccu.*
he child-ACC beat.PAST
'He beat the child.'

First, we will assume the state of affairs before a unidirectional constraint is developed. The PROM constraint of the original unidirectional proposal (section 8.3.1) is substituted with GEN(ERALIZE), which can be seen as the constraint version of our assumed generalization principle and here tells the hearer to use the default distribution of roles in the interpretation of specific utterances. Note that we have retained the use of FAITHL, which can now be understood as the constraint implementation of the hearer's version of the need for communicative success: in order to understand each other, one should be faithful to the conventions in the language that say, for example, that *chairs* denote 'chairs' and object markers mark objects.

First, an unmarked form is produced (PROD) and checked for its interpretation (INT_a, Tableau 8.7). Since both arguments are animate, GEN cannot decide between them, and since no marking is used, neither can FAITHL.

PROD: beat(he, child)	ECON	FAITHL	GEN
→ a. *he child beat*			
b. *he child-OBJ beat*	*		

INT$_a$: *he child beat*	ECON	FAITHL	GEN
→ i. beat(he, child)			
→ ii. beat(child, he)			

Tableau 8.7 Check 1 (before generalization)

As the unmarked form will not lead to success, an object marker is tried next. Since the use of an object marker to mark the object is generally accepted, FAITHL will now penalize any interpretation other than one in which the object-marked argument is the object. As a result, the correct interpretation is obtained (Tableau 8.8).

PROD: beat(he, child)	ECON	FAITHL	GEN
→ a. *he child beat*			
↔ b. *he child-OBJ beat*	*		
INT$_a$: *he child beat*	ECON	FAITHL	GEN
→ i. beat(he, child)			
→ ii. beat(child, he)			
INT$_b$: *he child-OBJ beat*	ECON	FAITHL	GEN
→ i. beat(he, child)			
ii. beat(child , he)		*	

Tableau 8.8 Check 2 (before generalization)

As argued above, after many of such bidirectional optimization processes, the generalization will be made that animate objects need marking, leading to a constraint ANIMO→MARKER. As we will show now, the development of this constraint will save one bidirectional check, as the optimal candidate from the production perspective will directly lead to the right interpretation. In Tableau 8.9, the derived constraint has been included in the grammar. As can be seen, it decides for a marked form candidate directly at PROD, a candidate that will indeed lead to the correct interpretation. Differently from the previous state of affairs, we only need one bidirectional check to confirm communicative success.

PROD: beat(he, child)	ANIMO→ MARKER	ECON	FAITHL	GEN
a. *he child beat*	*			
↔ b. *he child-OBJ beat*		*		
INT$_a$: *he child beat*		ECON	FAITHL	GEN
→ i. beat(he, child)				
ii. beat(child, he)			*	

Tableau 8.9 Gain (after generalization)

There is a drawback to this generalization, albeit a minor one. Once a constraint is developed that requires the marking of all animate objects, this object marker will also be used in situations in which it is not needed for disambiguation (cf. Durie 1995). Consider the following example from Malayalam.

(8) *Avan oru paʃuvin-e vaɲɲi.*
he a cow-ACC buy.PAST
'He bought a cow.'

Obviously, the hearer will know who is buying whom if a man and a cow are involved, also in the absence of an object marker. Indeed, GEN penalizes interpretations that go against the default, as illustrated in Tableau 8.10. Nevertheless, the marker is obligatory because of ANIMO→MARKER. This is the price that has to be paid for the use of a specialized constraint. It can be hypothesized that the costs of this overgeneralization (i.e., the incidental unnecessary use of an object marker) are more than compensated for by the savings from not having to find out for most situations that an economical utterance will not do, and to develop a new attempt subsequently.[10] In sum, generalization leads to accidental redundancy, which is inherent to local DCM systems.

PROD: buy(he, cow)	ANIMO→ MARKER	ECON	FAITHL	GEN

[10] Given the SOV nature of the language, use of the object marker may also be helpful in resolving potential online ambiguity during sentence processing.

	a. *he cow buy*	*			
↔	b. *he cow-*OBJ *buy*		*		
INT$_b$: *he cow-*OBJ *buy*			ECON	FAITHL	GEN
→	i. buy(he, cow)				
	ii. buy(cow, he)			*	*
INT$_a$: *he cow buy*			ECON	FAITHL	GEN
→	i. buy(he, cow)				
	ii. buy(cow, he)				*

Tableau 8.10 Overgeneralization of object marker in local DCM

Differently from the traditional approach to local DCM, we have proposed a semi-bidirectional analysis. Apart from the benefit of the overall parsimony, there is a crucial observation that suggests that this local system is indeed still bidirectional. Recall that in Malayalam all and only animate objects are obligatorily marked. Inanimate ambiguity, however, is also solved with an object marker, as illustrated in the following example:

(9) a. *Kappal tiramaalakaḷ-e bheediccu.*
ship waves-ACC split.PAST
'The ship broke the waves.'

b. *Tiramaalakaḷ kappal-ine bheediccu.*
waves ship-ACC split.PAST
'The waves broke the ship.'

Obviously, our constraint ANIMO→MARKER does not apply here and the object marker has to be motivated in a different way. We propose that it follows from the bidirectional reasoning that we have used throughout our analysis. The development of unidirectional constraints does not turn a bidirectional system into a unidirectional one; it only ensures that the first candidate that is checked will be found bidirectionally optimal.

8.6 Discussion

In this chapter we have shown how a unidirectional constraint can be derived from a generalization over multiple bidirectional optimization processes. We have suggested that our analysis, although applied to DOM here, should hold more generally. Indeed, nothing in our algorithm is specific to object marking. The abstract idea simply is that the odds for a certain output given a certain input increase to such to an extent that a generalization about their relation presents itself. In the process of tracking these optimization procedures (cf. Figure 8.1) at some point the learner will generalize over their outcomes and establish a constraint that directly links input and output. As the reader will have noticed, we have not specified at which point exactly this generalization is made or expected. Although the establishment of such a link sounds very plausible to us, we know too little about the brain to pinpoint the number of optimizations or the odds for a marker that are necessary for this. Comparable questions have been studied experimentally, however, and may give at least an abstract answer. Peterson and Beach (1967: 37–8) discuss the *optional stopping* paradigm, in which "[a]fter each datum the subjects has the option of continuing the sample or of stopping and making his inference." The conclusion drawn from multiple studies using this paradigm is that subjects keep sampling data as long as the costs of a new datum are less than the expected increase in payoff from the information it will provide.

In the simulation presented above, we assumed that the language learner starts out with a grammar that is bidirectionally specified. Although this is perhaps not the standard assumption in the OT literature on child language acquisition (see Hendriks, this volume, for a discussion), we believe it can be reconciled with existing proposals. The distinction between unidirectional and bidirectional grammars in language acquisition figures most prominently in discussions of asymmetries between language production and comprehension, be it when production lags behind comprehension (Smolensky 1996) or vice versa (Hendriks and Spenader 2005/6; de Hoop and Krämer 2005/6). Smolensky (1996) proposes that children use constraints bidirectionally. Both in production and comprehension the same constraints are used in the same ranking. Our proposal comes very close to this (cf. Tableau 9) with the difference that we do not assume a given set of domain-specific constraints. (But cf. Hendriks and Spenader 2005/6—see also de Hoop and Krämer 2005/6—who propose that children start out with a unidirectional grammar and only at a later stage develop a bidirectional grammar in which they optimize over pairs of form and meaning.)

8.7 Conclusion

In this chapter, we simulated the development of differential case-marking systems, showing that unidirectional constraints can be derived from a generalization over bidirectional optimization processes. Our findings bear on two discussions in the OT literature. First, they suggest that not all constraints need to be innate. Domain-specific constraints can be derived from optimization procedures in which initially only very general constraints are used. Second, it was shown that different case-marking systems could be explained using the same bidirectional architecture. With the development of specific unidirectional constraints, the bidirectionally optimal candidate is produced right away more often, because of which costly ad hoc repair strategies are less often necessary.

Acknowledgments

We would like to thank the editors of this book, the anonymous reviewers, and the audience of the workshop Advances in Optimality Theory: Syntax and Semantics at the Johns Hopkins University, whose comments and suggestions helped improve this chapter. The first author gratefully acknowledges financial support from the Netherlands Organization for Scientific Research (NWO; grant 275-78-001, *The Exaptation of Argument Marking*).