Quantifiers in TIME and SPACE : computational complexity of generalized quantifiers in natural language

Szymanik, J.K.

Publication date
2009

Chapter 4

Complexity of Quantified Reciprocals

The reciprocal expressions *each other* and *one another* are common elements of everyday English. Therefore, it is not surprising that they have been extensively studied in the formal semantics of natural language. There are two main approaches to reciprocals in the literature. The long trend of analyzing reciprocals as anaphoric noun phrases with the addition of plural semantics culminates in a paper of Beck (2000). A different tendency — recently represented by Sabato and Winter (2005) — is to analyze reciprocals as polyadic quantifiers.

In this chapter we study the computational complexity of quantified reciprocal sentences. We ally ourselves to the second tradition and treat reciprocal sentences as examples of a natural language semantic construction that can be analyzed in terms of so-called polyadic lifts of simple generalized quantifiers (see Chapter 3 of the thesis).

First, we propose new polyadic lifts expressing various possible meanings of reciprocal sentences with quantified antecedents, i.e., sentences where "each other" refers in a co-reference to the quantified noun phrase (see Dalrymple et al., 1998, Chapter 7). In other words, we will consider quantified reciprocal sentences, like "Five professors discuss with each other", where reciprocal phrase "each other" refers to quantified noun phrase, in this case "five professors". All these lifts are definable in the existential fragment of second-order logic. Therefore, according to the Σ_1^1-thesis formulated in Section 1.8 the model we investigate seems plausible.

Then we study the computational complexity of reciprocal lifts with respect to the quantifiers in the antecedents. Using results from the previous chapter (on the computational complexity of Ramsey quantifiers) we observe a computational dichotomy between different interpretations of reciprocity. Namely, we treat reciprocal expressions as polyadic lifts turning monadic quantifiers into Ramsey-like quantifiers. Differences in computational complexity beetween various interpretations of reciprocal expressions give an additional argument for the robustness of the semantic distinctions established between reciprocal meanings (see Dalrymple et al., 1998).

In particular, we give a sufficient condition for a generalized quantifier to make its strong reciprocal interpretation PTIME computable. Moreover, we present NP-complete natural language quantifier constructions which occur frequently in everyday English. For instance, strong interpretations of reciprocal sentences with counting and proportional quantifiers in the antecedents are intractable. As far as we are aware, all other known NP-complete quantifier constructions are based on ambiguous and artificial branching operations (see Section 3.2 and Chapter 6 for more discussion).

Finally, we investigate the cognitive status of the so-called Strong Meaning Hypothesis proposed by Dalrymple et al. (1998). We argue that if one assumes some kind of algorithmic theory of meaning as we do in Chapter 1, then the shifts between different interpretations of reciprocal sentences, predicted by the Strong Meaning Hypothesis, have to be extended by accommodating the possible influence of differences in computational complexity between various readings of reciprocity.

The considerations of this chapter are based on papers from the Amsterdam Colloquium 2007 (see Szymanik, 2007b), Lecture Notes in Computer Science (Szymanik, 2008), and results proven in Chapter 3 of the thesis.

4.1 Reciprocal Expressions

We start by recalling examples of reciprocal sentences, versions of which can be found in ordinary (spoken and written) English (see footnote 1 in Dalrymple et al., 1998). Let us first consider sentences (1)–(3).

(1) At least 4 members of parliament refer to each other indirectly.

(2) Most Boston pitchers sat alongside each other.

(3) Some Pirates were staring at each other in surprise.

The possible interpretations of reciprocity exhibit a wide range of variation. For example, sentence (1) implies that there is a subset of parliament members of cardinality at least 4 such that each parliament member in that subset refers to each of the other parliament members in that subset. However, the reciprocals in sentences (2) and (3) have different meanings. Sentence (2) states that each pitcher from a set containing most of the pitchers is directly or indirectly in the relation of sitting alongside with each of the other pitchers from that set. Sentence (3) says that there was a group of pirates such that every pirate belonging to the group stared at some other pirate from the group. Typical models satisfying (1)–(3) are illustrated in Figure 4.1. Following Dalrymple et al. (1998) we will call the illustrated reciprocal meanings *strong*, *intermediate*, and *weak*, respectively.

In general, according to Dalrymple et al. (1998) there are 2 parameters characterizing variations of reciprocity. The first one relates to how the scope relation,

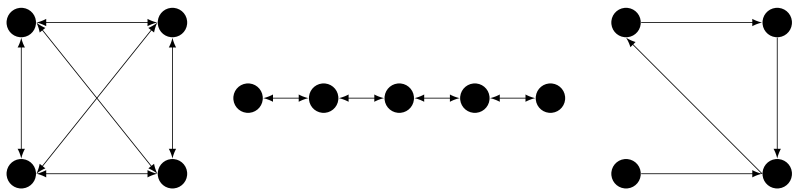

Figure 4.1: On the left is a typical model satisfying sentence (1) under the so-called *strong reciprocal* interpretation. Each element is related to each of the other elements. In the middle is an example of a model satisfying sentence (2) in a context with at most 9 pitchers. This is the *intermediate reciprocal* interpretation. Each element in the witness set of the quantifier Most is related to each other element in that set by a chain of relations. On the right, a model satisfying sentence (3), assuming the so-called *weak reciprocal* interpretation. For each element there exists a different related element.

R, should cover the domain, A, (in our case restricted by a quantifier in the antecedent). We have 3 possibilities:

FUL Each pair of elements from A participates in R directly.

LIN Each pair of elements from A participates in R directly or indirectly.

TOT Each element in A participates directly with at least one element in R.

The second parameter determines whether the relation R between individuals in A is the extension of the reciprocal's scope (R), or is obtained from the extension by ignoring the direction in which the scope relation holds ($R^{\vee} = R \cup R^{-1}$).

By combining these two parameters Dalrymple et al. (1998) gets six possible meanings for reciprocals. We have already encountered three of them: strong reciprocity, FUL(R); intermediate reciprocity, LIN(R); and weak reciprocity, TOT(R). There are three new logical possibilities: strong alternative reciprocity, FUL($R^{\vee}$); intermediate alternative reciprocity, LIN($R^{\vee}$); and weak alternative reciprocity, TOT($R^{\vee}$). Among these, two interpretations are linguistically attested: intermediate alternative reciprocity is exhibited by sentence (4) and weak alternative reciprocity occurs in sentence (5) (see Figure 4.2 for typical models).

(4) Most stones are arranged on top of each other.

(5) All planks were stacked on top of each other.

If we do not put any restrictions on the scope of the relation R, then stronger reciprocal interpretations imply weaker ones — as it is depicted in the left part of Figure 4.3. However, assuming certain properties of the relation some of the possible definitions become equivalent. For example, if the relation in question is symmetric, then obviously alternative versions reduce to their “normal” counterparts

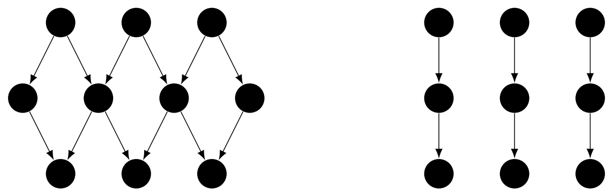

Figure 4.2: On the left is a typical model satisfying sentence (4) under the so-called *intermediate alternative reciprocal* interpretation. Ignoring the direction of arrows, every element in the witness set of the quantifier Most is connected directly or indirectly. On the right is an example of a model satisfying sentence (5) under the so-called *weak alternative reciprocal* reading. Each element participates with some other element in the relation as the first or as the second argument, but not necessarily in both roles.

and we have only three different reciprocal interpretations: $FUL(R) = FUL(R^{\vee})$, $LIN(R) = LIN(R^{\vee})$, and $TOT(R) = TOT(R^{\vee})$. If the relation R is transitive, then $FUL(R) = LIN(R)$ and the classification of different reciprocal meanings collapses to the one depicted in Figure 4.3 on the right.

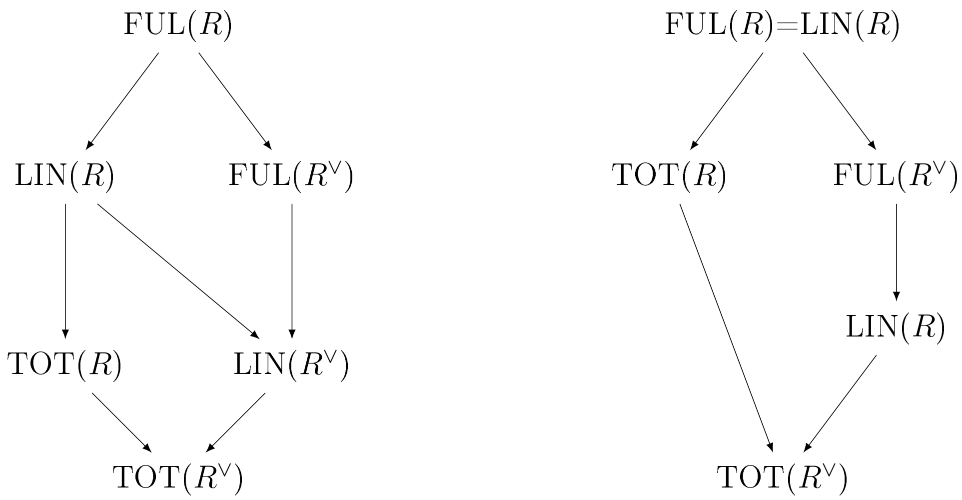

Figure 4.3: On the left, inferential dependencies between the six interpretations of reciprocity. On the right, the situation when the reciprocal relation is transitive. In these diagrams implications are represented by arrows.

4.1.1 Strong Meaning Hypothesis

In an attempt to explain variations in the literal meaning of the reciprocal expressions Dalrymple et al. (1998) proposed the *Strong Meaning Hypothesis* (SMH).

According to this principle, the reading associated with the reciprocal in a given sentence is the strongest available reading which is consistent with the properties of reciprocal relation and with relevant information supplied by the context. Sabato and Winter (2005) proposed a considerably simpler system in which reciprocal meanings are derived directly from semantic restrictions using the SMH.

4.1.1. EXAMPLE. Let us give one of the examples described by Dalrymple et al. (1998) of using SMH to derive proper interpretation of reciprocal statements. Consider the following sentence:

(6) The children followed each other.

This sentence can be interpreted in many ways depending on what is permitted by the context. First, consider:

(7) The children followed each other into the church.

The relation "following into the church" is asymmetric and intransitive disallowing strong (alternative) reciprocal interpretation. Moreover, the intermediate interpretation is impossible since children who go into the church first cannot even indirectly be said to follow children who go into the church later. Additionally, if the group of children is finite then the weak reading is excluded as it is not possible for each child to be a follower; simply put, someone must be the first to go into the church. This analyzes leaves 2 possibilities: the alternative intermediate reading and the weak alternative interpretation. The first suggested that children entered in one group while the later allows more than one group. As the alternative intermediate reading implies the alternative weak reading then SMH predicts that the sentence has the first meaning, assuming that the context does not supply additional information that the children enter the church in multiple groups. However, when you consider the similar sentence:

(8) The children followed each other around the Maypole.

Then unlike in the context described above, the path traversed by the children is circular. Hence, the intermediate reading appears as one of the possible interpretations. This is logically strongest possibility and according to SMH it properly describes the meaning of that sentence.

Our results show that the various meanings assigned to reciprocals with quantified antecedents differ drastically in their computational complexity. This fact can be treated as a suggestion to improve the SMH by taking into an account complexity constraints. We elaborate on this in the last section of this chapter, before we draw some conclusions.

4.2 Reciprocals as Polyadic Quantifiers

Monadic generalized quantifiers provide the most straightforward way to define the semantics of noun phrases in natural language (see Peters and Westerståhl, 2006, for a recent overview; also consult Section 2.2.). Sentences with reciprocal expressions transform such monadic quantifiers into polyadic ones. We will analyze reciprocal expressions in that spirit by defining appropriate lifts on monadic quantifiers. These lifts are definable in existential second-order logic.

For the sake of simplicity we will restrict ourselves to reciprocal sentences with right monotone increasing quantifiers in their antecedents. Recall from Section 2.2.5 that a quantifier Q of type $(1, 1)$ is monotone increasing in its right argument whenever: if $Q_M[A, B]$ and $B \subseteq B' \subseteq M$, then $Q_M[A, B']$. The lifts defined below can be extended to cover also sentences with decreasing and non-monotone quantifiers, for example by following the strategy of bounded composition suggested by Dalrymple et al. (1998) or the determiner fitting operator proposed by Ben-Avi and Winter (2003). The situation is here analogous to problems with collective lifts for non-increasing quantifiers, discussed in Section 5.2.3 of this dissertation.

4.2.1 Strong Reciprocal Lift

In order to define the meaning of strong reciprocity we make use of the well-known operation on quantifiers called *Ramseyfication* (see e.g. Hella et al., 1997, and Section 3.3 of this thesis).

4.2.1. Definition. Let Q be a right monotone increasing quantifier of type $(1, 1)$. We define:

$$\mathsf{Ram}_{\mathsf{S}}(\mathsf{Q})[A, R] \iff \exists X \subseteq A[\mathsf{Q}(A, X) \wedge \forall x, y \in X(x \neq y \implies R(x, y))].$$

■

We will call the result of such lifting a *Ramsey quantifier*.[1] It says that there exists a subset X of the domain A, restricted by a quantifier Q, such that every two elements from X are directly related via the reciprocal relation R.

In the same way we can also easily account for alternative strong reciprocity:

4.2.2. Definition.

$$\mathsf{Ram}_{\mathsf{S}}{}^{\vee}(\mathsf{Q})[A, R] \iff \exists X \subseteq A[\mathsf{Q}(A, X) \wedge \forall x, y \in X(x \neq y \implies (R(x, y) \vee R(y, x)))].$$

■

This expresses an analogous condition to the one before, but this time it is enough for the elements of X to be related either by R or by R^{-1}.

[1] Notice that in the previous chapter we defined Ramsey quantifiers to be of type (2) (see Definition 3.3.1). Hence, to be precise the result of the Ramseyfication gives the relativized (see Section 2.2.5) Ramsey quantifier.

4.2.2 Intermediate Reciprocal Lift

In a similar way we define more lifts to express intermediate reciprocity and its alternative version.

4.2.3. Definition.

$$\begin{aligned}\mathsf{Ram}_{\mathsf{I}}(\mathsf{Q})[A,R] \iff \exists X \subseteq A[\mathsf{Q}(A,X) \wedge \forall x, y \in X\\ (x \neq y \implies \exists \text{ sequence } z_1, \ldots, z_\ell \in X \text{ such that}\\ (z_1 = x \wedge R(z_1, z_2) \wedge \ldots \wedge R(z_{\ell-1}, z_\ell) \wedge z_\ell = y)].\end{aligned}$$

■

This condition guarantees that there exists a subset X of domain A which is connected with respect to R, i.e. any two elements from X are in the relation directly or indirectly.

4.2.4. Definition.

$$\begin{aligned}\mathsf{Ram}_{\mathsf{I}}{}^{\vee}(\mathsf{Q})[A,R] \iff \exists X \subseteq A[\mathsf{Q}(A,X) \wedge \forall x, y \in X\\ (x \neq y \implies \exists \text{ sequence } z_1, \ldots, z_\ell \in X \text{ such that}\\ (z_1 = x \wedge (R(z_1, z_2) \vee R(z_2, z_1)) \wedge \ldots\\ \wedge (R(z_{\ell-1}, z_\ell) \vee R(z_\ell, z_{\ell-1})) \wedge z_\ell = y)].\end{aligned}$$

■

In other words, $\mathsf{Ram}_{\mathsf{I}}{}^{\vee}$ says that any two elements from X are in the relation $R^{\vee}$ directly or indirectly. The property of graph connectedness is not elementary expressible; we need a universal monadic second-order formula. Hence from the definability point of view $\mathsf{Ram}_{\mathsf{I}}$ ($\mathsf{Ram}_{\mathsf{I}}{}^{\vee}$) seems more complicated than $\mathsf{Ram}_{\mathsf{S}}$ ($\mathsf{Ram}_{\mathsf{S}}{}^{\vee}$). However, as we will see, this is not the case from the computational complexity point of view.

4.2.3 Weak Reciprocal Lift

For weak reciprocity we take the following lifts.

4.2.5. Definition.

$$\mathsf{Ram}_{\mathsf{W}}(\mathsf{Q})[A,R] \iff \exists X \subseteq A[\mathsf{Q}(A,X) \wedge \forall x \in X\ \exists y \in X(x \neq y \wedge R(x,y))].$$

■

4.2.6. DEFINITION.

$$\mathsf{Ram_W}^{\vee}(\mathsf{Q})[A,R] \iff \exists X \subseteq A[\mathsf{Q}(A,X) \wedge \forall x \in X \exists y \in X \\ (x \neq y \wedge (R(x,y) \vee R(y,x))].$$

■

The weak lifts say that there exists a subset X of the domain A such that for every element from this subset there exists another element in the subset related by R (or $R^{\vee}$ in the case of the alternative lift).

4.2.4 The Reciprocal Lifts in Action

All reciprocal lifts produce polyadic quantifiers of type (1, 2). We will call the values of these lifts (alternative) strong, (alternative) intermediate and (alternative) weak reciprocity, respectively.

4.2.7. REMARK. Before we continue with an example, notice that all these lifts can be defined analogously for unary quantifiers, just as for type (1, 1). Simply replace condition $\mathsf{Q}(A,X)$ by $\mathsf{Q}(X)$ in the definitions.

4.2.8. EXAMPLE. The linguistic application of reciprocal lifts is straightforward. For example, using them we can account for the meanings of the reciprocal sentences (1)–(5) discussed in Section 4.1. Below we recall these sentences one by one. Each sentence is associated with a meaning representation expressed in terms of reciprocal lifts and quantifiers corresponding to the simple determiners occurring in these sentences.

(1) At least 4 parliament members refer to each other indirectly.

(9) $\mathsf{Ram_S}(\mathsf{At\ least}\ 4)[\mathrm{MP},\ \mathrm{Refer\text{-}indirectly}]$.

(2) Most Boston pitchers sat alongside each other.

(10) $\mathsf{Ram_I}(\mathsf{Most})[\mathrm{Pitcher},\ \mathrm{Sit\text{-}next\text{-}to}]$.

(3) Some pirates were staring at each other in surprise.

(11) $\mathsf{Ram_W}(\mathsf{Some})[\mathrm{Pirate},\ \mathrm{Staring\text{-}at}]$.

(4) Most stones are arranged on top of each other.

(12) $\mathsf{Ram_I}^{\vee}(\mathsf{Most})[\mathrm{Stones},\ \mathrm{Arranged\text{-}on\text{-}top\text{-}of}]$.

(5) All planks were stacked on top of each other.

(13) $\mathsf{Ram_W}^{\vee}(\mathsf{All})[\mathrm{Planks},\ \mathrm{Stack\text{-}on\text{-}top\text{-}of}]$.

It is easy to see that our formulae express the appropriate reciprocal meanings of these sentences, i.e. (alternative) strong, (alternative) intermediate and (alternative) weak reciprocity, respectively. They are true in the corresponding models depicted in Figures 4.1 and 4.2.

4.3 Complexity of Strong Reciprocity

In this section we investigate the computational complexity of quantified strong reciprocal sentences. In other words, we are interested in how difficult it is to evaluate the truth-value of such sentences in finite models. Studying this problem we make direct use of facts proven in Section 3.3 and we refer to the methods of descriptive complexity theory introduced in Section 2.4 of the Mathematical Prerequisites chapter.

Recall that we identify models of the form $\mathbb{M} = (M, A, R)$, where $A \subseteq U$ and $R \subseteq U^2$, with colored graphs and that we consider only monotone increasing quantifiers. Hence, in graph-theoretical terms we can say that $\mathbb{M} \models \mathsf{Ram_S(Q)}[A, R]$ if and only if there is a subgraph in A complete with respect to R, of a size bounded below by the quantifier Q. R is the extension of a reciprocal relation. If R is symmetric then we are dealing with undirected graphs. In such cases $\mathsf{Ram_S}$ and $\mathsf{Ram_S}^{\vee}$ are equivalent. Otherwise, if the reciprocal relation R is not symmetric, our models become directed graphs.

In what follows we will restrict ourselves to undirected graphs. We show that certain strong reciprocal quantified sentences interpreted in such graphs are NP-complete. Notice that undirected graphs are a special case of directed graphs; then our NP-complete sentences are also intractable over directed graphs.

4.3.1 Counting Quantifiers in the Antecedent

To decide whether in some model $\mathbb{M}$ sentence $\mathsf{Ram_S(At\ least\ k)}[A, R]$ is true we have to solve the CLIQUE problem for M and k. Recall from Section 3.3.3 that a brute force algorithm to find a clique in a graph is to examine each subgraph with at least k vertices and check if it forms a clique. This means that for every fixed k the computational complexity of $\mathsf{Ram_S(At\ least\ k)}$ is in PTIME. For instance, $\mathsf{Ram_S(At\ least\ 5)}$ is computable in polynomial time. In general, notice that the strong reciprocal sentence $\mathsf{Ram_S}(\exists^{\geq k})[A, R]$ is equivalent to the following first-order formula:

$$\exists x_1 \ldots \exists x_k \Big[\bigwedge_{1 \leq i < j \leq k} x_i \neq x_j \wedge \bigwedge_{1 \leq i \leq k} A(x_i) \wedge \bigwedge_{\substack{1 \leq i \leq k \\ 1 \leq j \leq k}} R(x_i, x_j) \Big].$$

However, when we consider natural language semantics from a procedural point of view it is natural to assume that people have one quantifier concept

At least k, for every natural number k, rather than the infinite set of concepts At least 1, At least 2, ... It seems reasonable to suppose that we learn one mental algorithm to understand each of the counting quantifiers At least k, At most k, and Exactly k, no matter which natural number k actually is. Mathematically, we can account for this idea by introducing counting quantifiers. Recall from Definition 3.2.4 that the counting quantifier $\mathsf{C}^{\geq \mathsf{A}}$ says that the number of elements satisfying some property is greater than or equal to the cardinality of the set A. In other words, the idea here is that determiners like At least k express a relation between the number of elements satisfying a certain property and the cardinality of some prototypical set A. For instance, the determiner At least k corresponds to the quantifier $\mathsf{C}^{\geq \mathsf{A}}$ such that $\mathrm{card}(A) = k$. Therefore, the determiners At least 1, At least 2, At least 3, ... are interpreted by one counting quantifier $\mathsf{C}^{\geq \mathsf{A}}$ — the set A just has to be chosen differently in every case.

The quantifier $\mathsf{Ram}_\mathsf{S}(\mathsf{C}^{\geq \mathsf{A}})$ expresses the general schema for a reciprocal sentence with a counting quantifier in the antecedent. Such a general pattern defines an NP-complete problem.

4.3.1. Proposition. *The quantifier* $\mathsf{Ram}_\mathsf{S}(\mathsf{C}^{\geq \mathsf{A}})$ *is mighty.*

Proof This fact is equivalent to proposition 3.3.6 from the previous chapter, where we observed that the so-called Ramsey counting quantifier, R_A, is mighty. □

4.3.2. Corollary. *The quantifier* $\mathsf{Ram}_\mathsf{S}{}^{\vee}(\mathsf{C}^{\geq \mathsf{A}})$ *is mighty.*

These results indicate that even though in a given situation checking the truth-value of a sentence with a fixed number, such as (1), is tractable, the general schema characterizing strong reciprocal sentences with counting quantifiers is NP-complete.

4.3.2 Proportional Quantifiers in the Antecedent

We can give another example of a family of strong reciprocal sentences which are intractable. Let us consider the following sentences:

(14) Most members of parliament refer to each other.

(15) At least one third of the members of parliament refer to each other.

(16) At least $q \times 100\%$ of the members of parliament refer to each other.

We will call these sentences *strong reciprocal sentences with proportional quantifiers*. Their general form is given by the sentence schema (16), where q can be interpreted as any rational number between 0 and 1. These sentences say that with respect to the reciprocal relation, R, there is a complete subset $Cl \subseteq A$, where A is the set of all parliament members, such that $\text{card}(Cl) \geq q \times \text{card}(A)$.

Recall that for any rational number $0 < q < 1$ we say that a set $A \subseteq U$ is q-large relative to U if and only if $\frac{card(A)}{card(U)} \geq q$ (see Definition 3.3.7). In this sense q determines a proportional quantifier $\mathsf{Q_q}$ of type (1, 1) as follows.

4.3.3. Definition.

$$\mathbb{M} \models \mathsf{Q_q}[A, B] \text{ iff } \frac{\text{card}(A \cap B)}{\text{card}(A)} \geq q.$$

■

4.3.4. Example. Let us give two examples of proportional quantifiers.

$$\mathbb{M} \models \mathsf{Most}[A, B] \text{ iff } \frac{\text{card}(A \cap B)}{\text{card}(A)} > \frac{1}{2}.$$

$$\mathbb{M} \models \mathsf{At\ least\ one\ third}\ [A, B] \text{ iff } \frac{\text{card}(A \cap B)}{\text{card}(A)} \geq \frac{1}{3}.$$

The strong reciprocal lift of a proportional quantifier, $\mathsf{Ram_S}(\mathsf{Q_q})$, is of type (1, 2) and obviously might be used to express the meaning of sentences like (14)–(16). We will call quantifiers of the form $\mathsf{Ram_S}(\mathsf{Q_q})$ *proportional Ramsey quantifiers*. Notice that a quantifier $\mathsf{Ram_S}(\mathsf{Q_q})$ is simply a relativization (see 3.3.9 for a definition) of the mighty proportional Ramsey quantifier $\mathsf{R_q}$ defined in Chapter 3.3.4. Therefore, it inherits the computational complexity of $\mathsf{R_q}$.

4.3.5. Proposition. *If q is a rational number and $0 < q < 1$, then the quantifier* $\mathsf{Ram_S}(\mathsf{Q_q})$ *is mighty.*

Proof Notice that $\mathsf{Ram_S}(\mathsf{Q_q}) = \mathsf{R_q^{rel}}$ and see the proof of Theorem 3.3.9 in the previous chapter. □

4.3.6. Corollary. *If q is a rational number and $0 < q < 1$, then the quantifier* $\mathsf{Ram_S}^{\vee}(\mathsf{Q_q})$ *is mighty.*

Therefore, strong reciprocal sentences with proportional quantifiers in the antecedent, like (14) or (15), are intractable (NP-complete).

4.3.3 Tractable Strong Reciprocity

Our examples show that the strong interpretation of some reciprocal sentences is intractable. In this section we will describe a class of unary monadic quantifiers for which the strong reciprocal interpretation is tractable (PTIME computable).

Following Väänänen (1997b) we will identify monotone simple unary quantifiers with number-theoretic functions, $f : \omega \longrightarrow \omega$, such that for all $n \in \omega$, $f(n) \leq n + 1$. In that setting the quantifier Q_f (corresponding to f) says of a set A that it has at least $f(n)$ elements, where n is the cardinality of the universe.

4.3.7. Definition. Given $f : \omega \longrightarrow \omega$, we define:

$$(\mathsf{Q}_f)_M[A] \iff card(A) \geq f(\text{card}(M)).$$

■

4.3.8. Example.

- $\exists = (\mathsf{Q}_f)_M$, where $f(\text{card}(M)) \geq 1$.
- $\forall = (\mathsf{Q}_g)_M$, where $g(\text{card}(M)) = \text{card}(M)$.
- $\mathsf{Most} = (\mathsf{Q}_h)_M$, where $h(\text{card}(M)) > \frac{\text{card}(M)}{2}$.

Notice that having a monotone increasing quantifier we can easily find the function corresponding to it.

4.3.9. Definition. Let Q be a monotone increasing quantifier of type (1). Define:

$$f(n) = \begin{cases} \text{least } k \text{ such that:} & \\ \exists U \exists A \subseteq U[\text{card}(U) = n \wedge \text{card}(A) = k \wedge \mathsf{Q}_U(A)] & \text{if such a } k \text{ exists} \\ n+1 & \text{otherwise.} \end{cases}$$

■

4.3.10. Proposition. *If Q is a monotone increasing quantifier of type (1) and function f is defined according to Definition 4.3.9 then*

$$\mathsf{Q} = \mathsf{Q}_f\,.$$

Proof The equality follows directly from the definitions. □

In the previous Chapter we have shown that for every PTIME computable and bounded (see Definition 3.3.15) function, f, the Ramsey quantifier $\mathsf{R_f}$ is also PTIME computable (see Theorem 3.3.16). The strong reciprocal lift — as we mentioned — produces Ramsey quantifiers from simple determiners. Hence, $\mathsf{Ram_S}(\mathsf{Q}_f)$ corresponds to $\mathsf{R_f}$. Therefore, we can claim that polynomial computable bounded quantifiers are closed under the strong reciprocal lift.

4.3.11. PROPOSITION. *If a monotone increasing quantifier Q_f is PTIME computable and bounded, then the reciprocal quantifier $\mathsf{Ram}_{\mathsf{S}}(\mathsf{Q}_f)$ is PTIME computable.*

Proof See the proof of Theorem 3.3.16 from the previous chapter. □

4.3.12. REMARK. Notice that it does not matter whether we consider undirected or directed graphs, as in both cases checking whether a given subgraph is complete can be done in polynomial time. Therefore, the result holds for $\mathsf{Ram}_{\mathsf{S}}{}^{\vee}(\mathsf{Q}_f)$ as well.

4.3.13. COROLLARY. *If a monotone increasing quantifier Q_f is PTIME computable and bounded, then the quantifier $\mathsf{Ram}_{\mathsf{S}}{}^{\vee}(\mathsf{Q}_f)$ is PTIME computable.*

4.3.14. REMARK. Moreover, notice, that the relativization, Q_f^{rel}, of Q_f is the right monotone type (1, 1) quantifier:

$$(Q_f^{rel})_M\ [A, B] \iff \text{card}(A \cap B) \geq f(\text{card}(A)).$$

Thus, the restriction to unary quantifiers is not essential and the result may be easily translated for type (1, 1) determiners.

What are the possible conclusions from Proposition 4.3.11? We have shown that not all strong reciprocal sentences are intractable. As long as a quantifier in the antecedent is bounded the procedure of checking the logical value of the sentence is practically computable. For example, the quantifiers Some and All are relativizations of the PTIME computable bounded quantifiers $\exists$ and $\forall$. Therefore, the following strong reciprocal sentences are tractable:

(17) Some members of parliament refer to each other indirectly.

(18) All members of parliament refer to each other indirectly.

4.4 Intermediate and Weak Lifts

Below we show that intermediate and weak reciprocal sentences — as opposed to strong reciprocal sentences — are tractable, if the determiners occurring in their antecedents are practically computable.

Analogous to the case of strong reciprocity, we can also express the meanings of intermediate and weak reciprocal lifts in graph-theoretical terms. We say that $\mathbb{M} \models \mathsf{Ram}_{\mathsf{I}}(\mathsf{Q})[A, R]$ if and only if there is a connected subgraph in A of a size bounded from below by the quantifier Q. $\mathbb{M} \models \mathsf{Ram}_{\mathsf{W}}(\mathsf{Q})[A, R]$ if and only if there is a subgraph in A of the proper size without isolated vertices. All three are with respect to the reciprocal relation R, either symmetric or asymmetric.

We prove that the class of PTIME quantifiers is closed under the (alternative) intermediate lift and the (alternative) weak lift.

4.4.1. PROPOSITION. *If a monotone increasing quantifier* Q *is PTIME computable, then the quantifier* $\mathsf{Ram_I(Q)}$ *is PTIME computable.*

Proof Let $\mathbb{G} = (V, A, E)$ be a directed colored graph-model. To check whether $\mathbb{G} \in \mathsf{Ram_I(Q)}$ compute all connected components of the subgraph determined by A. For example, you can use a breadth-first search algorithm that begins at some node and explores all the connected neighboring vertices. Then for each of those nearest nodes, it explores their unexplored connected neighbor vertices, and so on, until it finds the full connected subgraph. Next, it chooses a node which does not belong to this subgraph and starts searching for the connected subgraph containing it. Since in the worst case this breadth-first search has to go through all paths to all possible vertices, the time complexity of the breadth-first search on the whole $\mathbb{G}$ is $O(\mathrm{card}(V) + \mathrm{card}(E))$. Moreover, the number of the components in A is bounded by $\mathrm{card}(A)$. Having all connected components it is enough to check whether there is a component C of the proper size, i.e., does $\mathsf{Q}[A, C]$ hold for some connected component C? This can be checked in polynomial time as Q is a PTIME computable quantifier. Hence, $\mathsf{Ram_I(Q)}$ is in PTIME. □

4.4.2. COROLLARY. *If a monotone increasing quantifier* Q *is PTIME computable, then the quantifier* $\mathsf{Ram_I}^{\vee}\mathsf{(Q)}$ *is PTIME computable.*

The next proposition follows immediately.

4.4.3. PROPOSITION. *If a monotone increasing quantifier* Q *is PTIME computable, then the quantifier* $\mathsf{Ram_W(Q)}$ *is PTIME computable.*

Proof To check whether a given graph-model $\mathbb{G} = (V, A, E)$ is in $\mathsf{Ram_W(Q)}$, compute all connected components $C_1, \ldots, C_t$ of the A-subgraph. Take $X = C_1 \cup \ldots \cup C_t$ and check whether $\mathsf{Q}[A, X]$. From the assumption this can be done in polynomial time. Therefore, $\mathsf{Ram_W(Q)}$ is in PTIME. □

4.4.4. COROLLARY. *If a monotone increasing quantifier* Q *is PTIME computable, then the quantifier* $\mathsf{Ram_W}^{\vee}\mathsf{(Q)}$ *is PTIME computable.*

These results show that the intermediate and weak reciprocal lifts do not increase the computational complexity of quantifier sentences in such a drastic way as may happen in the case of strong reciprocal lifts. In other words, in many contexts the intermediate and weak interpretations are relatively easy, as opposed to the strong reciprocal reading. For instance, the sentences (2), (3), (4), and (5) we discussed in the introduction are tractable. Hence from a computational complexity perspective the intermediate and reciprocal lifts behave similar to iteration, cumulation and resumption (discussed in Chapter 3).

In the next section we discuss the potential influence of computational complexity on the shifts in meaning of reciprocal sentences predicted by the Strong Meaning Hypothesis.

4.5 A Complexity Perspective on the SMH

Dalrymple et al. (1998) proposed a pragmatic principle, the Strong Meaning Hypothesis, to predict the proper reading of sentences containing reciprocal expressions. According to the SMH the reciprocal expression is interpreted as having the logically strongest truth conditions that are consistent with the given context. Therefore, if it is only consistent with the specified facts, a statement containing *each other* will be interpreted as a strong reciprocal sentence. Otherwise, the interpretation will shift toward the logically weaker intermediate or weak readings, depending on context (see Section 4.1.1).

The SMH is quite an effective pragmatic principle (see Dalrymple et al., 1998). We will discuss the shifts the SMH predicts from a computational complexity point of view, referring to the results provided in the previous sections.

Let us first think about the meaning of a sentence in the intensional way, identifying the meaning of an expression with an algorithm recognizing its denotation in a finite model.[2] Such algorithms can be described by investigating how language users evaluate the truth-value of sentences in various situations. On the cognitive level this means that subjects have to be equipped with mental devices to deal with the meanings of expressions. Moreover, it is cognitively plausible to assume that we have a single mental device to deal with most instances of the same semantic construction. For example, we believe that there is one mental algorithm to deal with the counting quantifier, $\mathsf{At\ least}\ k$, in most possible contexts, no matter what natural number k is. Thus, in the case of logical expressions like quantifiers, the analogy between meanings and algorithms seems uncontroversial.

However, notice that some sentences, being intractable, are too complex to identify their truth-value directly by investigating a model. The experience of programming suggests that we can claim a sentence to be difficult when it cannot be computed in polynomial time. Despite the fact that some sentences are sometimes[3] too hard for comprehension, we can find their inferential relations

[2] We have argued for this approach in Chapter 1. Now we only recall that it goes back to Frege (1892) and exists in the linguistic literature at different levels of transparency (see e.g. Moschovakis, 2006).

[3] The fact that the general problem is hard does not show that all instances normally encountered are hard. It is the matter for empirical study to provide us with data about the influence of computational complexity on our everyday linguistic experience. However, we believe that it is reasonable to expect that this happens at least in some situations. We refer the reader to Section 1.5.3 for a more substantial discussion.

with relatively easier sentences. See Section 1.8 for more discussion on indirect verification.

According to the SMH any reciprocal sentence, if it is only possible, should be interpreted as a strong reciprocal sentence. We have shown that the strong interpretation of sentences with quantified antecedents is sometimes intractable but the intermediate and weak reading are always easy to comprehend. In other words, it is reasonable to suspect that in some linguistic situations the strong reciprocal interpretation is cognitively much more difficult than the intermediate or the weak interpretation. This prediction makes sense under the assumption that $P \neq NP$ and that the human mind is bounded by computational restrictions. We omit a discussion here but see Chapter 1. We only recall that computational restrictions for cognitive abilities are widely treated in the literature (see e.g. Cherniak, 1981; Chalmers, 1994; Mostowski and Wojtyniak, 2004; Levesque, 1988; Mostowski and Szymanik, 2005). Frixione (2001) explicitly formulates the so-called P-cognition Thesis:

P-cognition Thesis *Human cognitive (linguistic) capacities are constrained by polynomial time computability.*

What happens if a subject is supposed to deal with a sentence too hard for direct comprehension? One possibility — suggested in Section 1.8 — is that the subject will try to establish the truth-value of a sentence indirectly, by shifting to an accessible inferential meaning. That will be, depending on the context, the intermediate or the weak interpretation, both being entailed by the strong interpretation.

Summing up, our descriptive complexity perspective on reciprocity shows that it might not always be possible to interpret a reciprocal sentence in the strong way, as the SMH suggests. If the sentence in question would be intractable under the strong reciprocal interpretation then people will turn to tractable readings, like intermediate and weak reciprocity. Our observations give a cognitively reasonable argument for some shifts to occur, even though they are not predicted by the SMH. For example, the SMH assumes that the following sentence should be interpreted as a strong reciprocal statement.

(19) Most members of parliament refer to each other indirectly.

However, we know that this sentence is intractable. Therefore, if the set of parliament members is large enough then the statement is intractable under the strong interpretation. This gives a perfect reason to switch to weaker interpretations.

4.6 Summary

By investigating reciprocal expressions in a computational paradigm we found differences in computational complexity between various interpretations of reciprocal sentences with quantified antecedents. In particular, we have shown that:

- There exist non-branching natural language constructions whose semantics is intractable. For instance, strong reciprocal sentences with proportional quantifiers in the antecedent, e.g. sentence (14), are NP-complete.
- For PTIME computable quantifiers the intermediate and weak reciprocal interpretations (see e.g. sentences (2) and (3)) are PTIME computable.
- If we additionally assume that a quantifier is bounded, like Some and All, then also the strong reciprocal interpretation stays in PTIME, e.g. sentences (17) and (18).

Therefore, we argue that:

- The semantic distinctions of Dalrymple et al. (1998) seem solid from a computational complexity perspective.
- The Strong Meaning Hypothesis should be improved to account for shifts in meaning triggered by the computational complexity of sentences.

Many questions arise which are to be answered in future work. Here we will mention only a few of them:

4.6.1. Question. Among the reciprocal sentences we found NP-complete constructions. For example, we have shown that the strong reciprocal interpretations of proportional quantifiers are NP-complete. On the other hand, we also proved that the strong reciprocal interpretations of bounded quantifiers are PTIME computable. It is an open problem where the precise border is between those natural language quantifiers for which Ramseyfication is in PTIME and those for which it is NP-complete. Is it the case that for every quantifier, Q, $\mathsf{Ram}_{\mathsf{S}}(\mathsf{Q})$ is either PTIME computable or NP-complete? We stated this question already in the previous chapter.

4.6.2. Question. There is a vast literature on the definability of polyadic lifts of generalized quantifiers (e.g. Väänänen, 1997b; Hella et al., 1997). We introduced some new linguistically relevant lifts, the weak and intermediate reciprocal lifts. The next step is to study their definability. For example, we would like to know how the definability questions for $\mathsf{Ram}_{\mathsf{S}}(\mathsf{Q}_f)$, $\mathsf{Ram}_{\mathsf{I}}(\mathsf{Q}_f)$, and $\mathsf{Ram}_{\mathsf{W}}(\mathsf{Q}_f)$ depend on the properties of f. Another interesting point is to link our operators with other polyadic lifts, like branching.

4.6.3. Question. We could empirically compare the differences in shifts from the strong interpretation of reciprocal sentences with bounded and proportional quantifiers in antecedents. Our approach predicts that subjects will shift to easier interpretations more frequently in the case of sentences with proportional quantifiers. Can we prove it empirically?