

Overcoming Structural Preference: Effects of Context on the Interpretation of the Chinese reflexive *ziji*

David Cheng-Huan Li and Elsi Kaiser

University of Southern California, Department of Linguistics,
90089 Los Angeles, California
{lfdc, emkaiser} @ usc.edu

Abstract. We report two experiments that investigate the referential properties of the Chinese simple reflexive *ziji*. These experiments have two main aims: (i) to investigate whether and to what extent different kinds of information guide reference resolution in Mandarin Chinese, and (ii) to shed light on the strength and time-course of contextual effects. Our results show that contextual constraints guide people’s final interpretation and on-line processing of the simple reflexive *ziji*. These findings are best captured with an interactive-parallel-constraint model, in which various constraints determine the final activation level of an antecedent candidate. In particular, the present study demonstrated that the interpretation of reflexive anaphors is guided not only by structural factors but also by discourse contextual information outside the present sentence. Our results show that effects of discourse context emerge rapidly during online processing and contribute to our understanding of how structural and discourse-level information interact during reference resolution.

Keywords: Chinese, reflexives, *ziji*, anaphor resolution, context, sentence processing, psycholinguistics

1 Introduction

Resolving referential dependencies is an integral part of language understanding. Earlier formal and experimental research on the interpretation of reflexive pronouns has tended to focus mainly on how grammatical knowledge influences interpretation. A prominent theory based on the structural relationships among syntactic constituents is Chomsky’s Binding Theory [1] (see also [2] for an experimental demonstration). In addition to syntactic constraints, the set of antecedent candidates is also narrowed down by the morphosyntactic features of the referring expression, including number and gender [3,4,5,6]. However, relatively little experimental work has examined how different types of constraints interact with one another to guide the interpretation of reflexives (but see e.g. [7] on the interaction of structural and semantic constraints).

The two experiments presented in this paper investigated the role of structural and contextual constraints in the interpretation of the Chinese reflexive *ziji*, which can be interpreted as coreferential with either a local or a long-distance antecedent. Our results support the general finding that contextual information interacts with structural

constraints in guiding the resolution of referential dependencies. Furthermore, because the contextual manipulation in our experiments was located in the preceding sentence, our experiments show that sentence-external information can guide the interpretation of *ziji*. Our results also have implications for our understanding of how potential antecedents compete with each other and how this competition is modulated by contextual plausibility.

1.1 Previous Studies

In an early study investigating the time-course of structural constraints on reference resolution, Nicol and Swinney [8] used sentences like ex. 1 to investigate the interpretation of reflexives like *himself* and pronouns like *him*.

- (1) a. The boxer told the skier that the doctor for the team would blame *himself* for the recent injury.
- b. The boxer told the skier that the doctor for the team would blame *him* for the recent injury.

The reflexive *himself* in ex. 1a is subject to Principle A of the Binding Theory (BT), which requires a reflexive to be bound by a c-commanding antecedent in its local domain. Thus, in ex. 1a, the local subject *doctor* is an accessible antecedent for *himself*. The pronoun *him* in ex. 1b, on the other hand, is subject to Principle B, which requires a pronoun to be free within its local domain. Therefore, the local subject *doctor* is not an accessible antecedent for *him*.

Nicol and Swinney were interested in testing whether Binding Theory acts as an ‘initial filter’, i.e., whether it is the case that, even at the earliest moments of processing a pronoun or a reflexive, only BT-accessible antecedents are activated. Or could it be the case that, early on, non-BT-licensed referents are also considered as part of the candidate antecedent set? On the basis of data from cross-modal lexical decision, Nicol and Swinney argued for the former option: If an antecedent candidate violates one of the binding principles, it is ‘filtered out’ from the initial set of antecedent candidates.

However, a growing body of recent work argues against this view. For example, in a series of self-paced reading experiments, Badecker & Straub [9] found that in sentences like ex. 2, the gender of structurally inaccessible local subject also played a role in participants’ self-paced reading performance – even though, according to Principle B, the local subject is not a possible antecedent for *him*.

- (2) a. John thought that *Bill* owed *him* another chance to solve the problem.
- b. John thought that *Beth* owed *him* another chance to solve the problem.

Badecker and Straub employed the self-paced reading paradigm, a widely-used method in which participants read the sentences in a word-by-word fashion and the reading times for each word are recorded. Slowdowns in reading time are interpreted as indicative of processing difficulties and increased processing load. They found that the gender of the BT-inaccessible local subject (*Bill/Beth*) influenced processing of

the pronoun: Participants' reading times were slower in the postpronominal region when the local subject (as well as the matrix subject) matched the pronoun in gender (*Bill*) than when it did not (*Beth*). Badecker & Straub attribute this slowdown due to increased processing load resulting from the processing system struggling with two possible antecedents, what they call the *multiple-candidate effect*. Badecker & Straub also replicated this 'multiple-candidate' slowdown for reflexives when they manipulated the gender of the structurally-inaccessible matrix subject. They conclude that Binding Theory does not act as an initial filter: Instead, even BT-inaccessible antecedents are considered during processing, as long as they sufficiently prominent. More broadly, their findings fit with the view that various constraints act in parallel and interact with one another during the process of evaluating antecedent candidates (see also [10,7], but see [11] for evidence that structural constraints are privileged and apply very early during processing).

1.2 Aims of the Present Research

The two experiments presented in this paper investigate the off-line and on-line processing of the simple reflexive *ziji* in Chinese. A growing body of psycholinguistic research suggests that Binding Theory does not act as a strict initial filter and that both syntactic and semantic information guide the interpretation of anaphoric expressions. However, given that existing psycholinguistic studies have tended to focus on English, there is relatively little work on anaphoric forms which are known to take both local and long-distance antecedents, such as *ziji* in Chinese.

To better understand how and whether contextual and structural factors influence the interpretation of such forms, Experiment 1 investigated whether *ziji* has a baseline preference for the local or long-distance antecedent. At the same time, the referential flexibility of *ziji* – the fact that, grammatically speaking, it can refer to either the local or the long-distance antecedent – allows us to address our second aim, which is to study of the effects of discourse context and timing of these effects. The structural properties of the *ziji*-containing sentence were kept constant, and the preceding discourse context was varied. With this set-up, any change in the interpretation of *ziji* can only be attributed to effects of contextual information that are sentence-external. Experiment 1 probed for off-line context effects, whereas Experiment 2 used self-paced reading which allows us to investigate when potential context effects emerge during real-time processing.

2 Experiment 1

Experiment 1 used an offline survey to investigate (i) the baseline referential properties of the simple reflexive *ziji*, as well as (ii) the extent to which the interpretation of *ziji* can be influenced by preceding discourse context.

As shown in ex. 3, when the simple reflexive *ziji* occurs in the object position, it can be interpreted as coreferential with either the local subject (*Zhangsan*) or the long-distance (matrix) subject (*Lisi*) [12,13].

- (3) Lisi shuo Zhangsan yaoqiu *ziji* xiage libai wancheng zhege yue de yewu
 Lisi say Zhangsan ask *self* next week complete this month DE quota
 “Lisi said Zhangsan asked *self* to complete this month’s quota this week.”

However, although both the local and the long-distance antecedents are possible, to the best of our knowledge it has not yet been investigated whether one is preferred over the other. Thus, the *first aim* of Experiment 1 was to test whether *ziji* prefers the matrix subject or the local subject in a neutral discourse context where both antecedents are equally plausible (as shown in ex. 4a). If a participant reads the context sentence in ex. 4a, followed by the test sentence in ex. 3, how will s/he interpret the ambiguous reflexive *ziji*?

The *second aim* was to assess whether the interpretation of *ziji* could be influenced by information in the preceding discourse. Because we had a hunch that, in a neutral context, *ziji* would prefer the local antecedent, we decided to investigate whether contextual factors could act against this preference. Thus, we investigated biasing contexts which made the *long-distance antecedent* more likely than the local antecedent, as shown in ex. 4b.¹ When a participant reads the biasing context in ex. 4b, followed by the test sentence in ex. 3, will this influence the interpretation of *ziji*? In a situation where Lisi is known to be troubled by Zhangsan’s orders, will participants tend to interpret Lisi as being the one who is asked to complete the task?

- (4) a. Zhangsan han Lisi shang libai tandao gongzuo de shiqing [NEUTRAL]
 Zhangsan and Lisi last week talk work DE business
 “Zhangsan and Lisi talked about work-related business last week.”
- b. Zhangsan dui xiashu de guiding ling Lisi gandao kunrao. [BIASING]
 Zhangsan to subordinates DE rules make Lisi feel troublesome
 “Zhangsan’s order to his subordinates has troubled Lisi.”

In this particular example, the local antecedent in the test sentence (*Zhangsan* in ex. 3) is realized as the possessor of the subject in the biasing context (ex. 4b), and the long-distance antecedent in the text sentence (*Lisi* in ex. 3) is the object of the biasing context sentence. Both *Zhangsan* and *Lisi* are realized (as a conjoined constituent, with *Zhangsan* linearly initial) in the subject position of the neutral context. Thus, in this particular example, the noun that turns out to be the local antecedent in the test sentence linearly precedes the long-distance antecedent in both the neutral context and

¹ One could also investigate biasing contexts which make the local antecedent more plausible. However, in this study we opted for contexts which bias towards the non-local antecedent. This was done to maximize the detectability of potential context effects. Given our hunch that *ziji* might have a baseline preference for local antecedents, there is a concern that contexts which bias towards the local antecedent may not end up having a detectable effect, due to potential ceiling effects resulting from a baseline local-antecedent preference. Thus, in this research, we opted for contexts which bias the non-local antecedent. This also allowed us to keep the number of conditions within this particular experiment down to a manageable size. However, the question of whether the preference for the antecedent can be strengthened by a context manipulation is an important question for future work.

the biasing context. However, this was counterbalanced across the target items in Experiment 1: On half of the targets, the local antecedent noun is mentioned before the long-distance antecedent (as in ex. 4), but in the other half of the targets, the neutral and biasing contexts were constructed such that the contexts mention the long-distance antecedent noun before the local antecedent noun. Because the local and long-distance antecedents' locations in the context sentences were controlled in this way, any potential effects of context cannot be attributed simply to a preference for *ziji* to refer to an antecedent in a particular linear or structural position.

In sum, Experiment 1 aims to test whether (i) *ziji* has a preference for either the long-distance or the local antecedent in a neutral context, and whether (ii) this baseline preference can be 'pushed around' or weakened by information in the preceding discourse.

2.1 Methods

Participants. Thirty native Chinese speakers participated. All participants were born in Taiwan.

Materials and Design. Participants read two-sentence sequences consisting of a context sentence (e.g. ex. 4a, 4b) and a test sentence (e.g. ex. 3). The experiment consisted of 24 target items and 36 fillers. After each target item, participants were asked a question as shown in ex. 5. We manipulated whether the question asked about the local subject (5a), the matrix subject (5b) or both (5c).² Participants were also asked to rate their confidence in their answer on a scale of 1 (most confident) to 5 (least confident). Thus, our study had a 2 x 3 design (context x question type). The experiment was conducted over the internet, via a webpage where participants entered their responses.

- (5) a. LOCAL ANTECEDENT QUESTION:
qing wen xiage libai dei wancheng zhege yue de yewu de ren shi *Lisi* ma?
"Is the person who needs to complete this month's quota this week *Lisi*?"
ANSWER CHOICES: YES or NO
- b. LONG-DISTANCE ANTECEDENT QUESTION:
qing wen xiage libai dei wancheng zhege yue de yewu de ren shi *Zhangsan* ma?
"Is the person who needs to complete this month's quota this week *Zhangsan*?"
ANSWER CHOICES: YES or NO
- c. BOTH ANTECEDENTS QUESTION:

² This three-question setup has the advantage of measuring the participants' response to both preferred and potentially dispreferred interpretations. A forced choice task that asks participants to select *either* the long-distance *or* local antecedent might mask the existence/availability of dispreferred but nevertheless grammatical interpretations.

qing wen xiage libai dei wancheng zhege yue de yewu de ren you keneng
 shi *Lisi* huo *Zhangsan* ma?
 “Can the person who needs to complete this month’s quota this week be
 either *Zhangsan* or *Lisi*?”
 ANSWER CHOICES: ONLY ONE IS POSSIBLE or BOTH ARE
 POSSIBLE

2.2 Results

Neutral context. There was a preference to interpret *ziji* as referring to the local antecedent when the context was neutral, as shown in Fig. 1. The long-distance antecedent question elicited only 26% ‘yes’ responses, whereas the local antecedent question resulted in 83% ‘yes’ responses in neutral contexts (significantly higher, $p < .001$). Thus, we see a strong preference for the local antecedent. When the ‘both’ question was used, 53% of participants responded that both antecedents are possible.

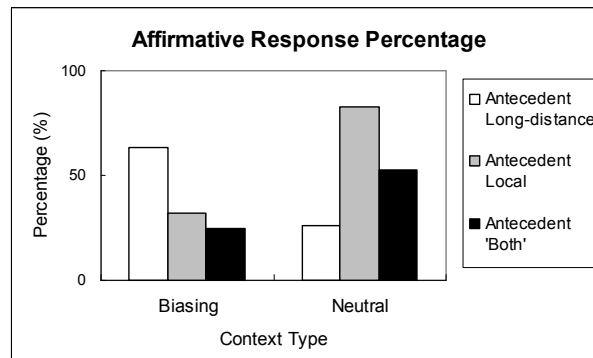


Fig. 1. Percentage of ‘yes’ responses to questions probing the acceptability of the long-distance antecedent, local antecedent, and both antecedents, grouped by discourse context

Biasing context. However, the biasing context conditions show that the referential preferences of *ziji* are subject to the influence of context. When the context favored the long-distance antecedent, we found that the long-distance antecedent question elicited 63% ‘yes’ responses (significantly more than in the neutral context, $p < .001$). In the biasing context, the local antecedent question elicited only 32% ‘yes’ responses (significantly less than in the neutral context, $p < .001$). The difference in the rate of ‘yes’ responses to the long-distance antecedent question and the local antecedent question was also significant ($p < .05$). The ‘both’ question resulted in 25% ‘both’ responses, which is significantly less than in the neutral context ($p < .001$).

In sum, we found that the long-distance-antecedent question elicited a lot of ‘yes’ answers when the context biased towards the long-distance antecedent, but the rate of ‘yes’ answers for the long-distance antecedent was low in the neutral contexts. In contrast, the local-antecedent question elicited a higher number of ‘yes’ responses in the neutral context than in the biasing contexts, suggesting a default preference for the local antecedent.

Confidence ratings. In terms of the participants' confidence level regarding their response, the participants were *more confident when the contexts were biasing* (Fig. 2) Confidence ratings in the biasing contexts did not differ significantly – they were uniformly high: The average rating for the long-distance-antecedent condition was 1.6, for the local antecedent condition it was 1.5, and for the 'both' condition it was 1.6.

However, the participants were not as confident in the neutral contexts. They were the least confident in the long-distance-antecedent condition (average rating: 2.2), more confident in the local antecedent (average rating: 1.9), and the most confident in the 'both' condition (average rating: 1.7). The decrease in confidence from the biasing contexts to the neutral contexts was significant for both the long-distance antecedent condition and the local-antecedent condition. For the 'both' condition, however, the difference was not significant ($p = .2$). The lowest confidence ratings were in the long-distance antecedent condition, which elicited significantly lower confidence ratings than the local antecedent condition or the 'both' condition (p 's $< .05$). However, there was no significant difference between the local antecedent condition and the 'both' condition ($p = .1$).

Thus, the confidence rating data suggest that when the contexts were biasing, the confidence levels in all three antecedent conditions were relatively low. When the contexts were neutral, the confidence levels were higher.

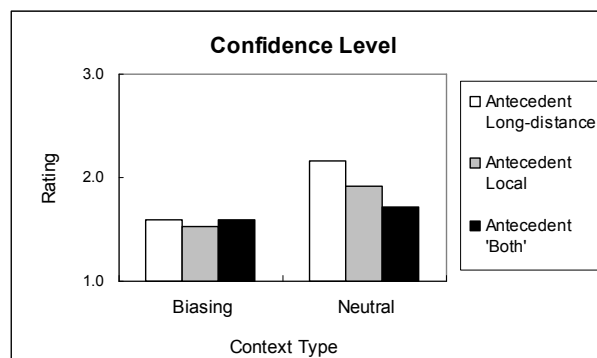


Fig. 2. Confidence levels with respect to prior response in the biasing and neutral contexts when the long-distance antecedent, local antecedent, and both antecedents were used to probe the participants' offline reference resolution.

3 Experiment 2

To investigate whether (and when) the real-time processing of *ziji* is sensitive to context effects, we conducted a self-paced reading experiment. Self-paced reading will allow us to see whether context effects emerge early, or whether the influence of context is something that only influences processing later on.

Recall that Experiment 1 showed that (i) in a neutral context, *ziji* has a default preference for the local antecedent but that (ii) contexts which bias the long-distance antecedent result in a significantly higher number of long-distance choices than local choices. Moreover, participants' confidence level ratings showed that the presence of a biasing context increased confidence levels, suggesting that there is *less competition between the two possible antecedents in biasing conditions than in neutral conditions*.

On the basis of these findings, we formulated the following predictions for Experiment 2: If presence of a biasing context boosts the activation of the long-distance antecedent and decreases competition between the local and the long-distance antecedent, as indicated by the confidence level ratings in Experiment 1, then one might predict that the reading times in Experiment 2 would be shorter in the biasing contexts than in neutral contexts. Importantly, by looking at when potential reading time differences occur, we can also tell when the effect of context 'kicks' in. Thus, we can assess whether the plausibility information in the preceding discourse context guides processing as soon as *ziji* is encountered, or whether contextual information only steps in at a later point in time.

3.1 Methods

Participants. Sixteen native Chinese speakers participated.³ All participants were born in Taiwan.

Materials and Design. In this experiment, we manipulated context type (neutral vs. biasing towards the long-distance antecedent), as in Experiment 1. In contrast to the yes/no questions of Experiment 1, in this experiment participants were asked forced-choice questions after each target which required making a choice between the long-distance antecedent and the local antecedent:

- (6) qing wen xiage libai dei wancheng zhege yue de yewu de ren shi shei
“Who is the one that needs to complete this month's quota this week?”
ANSWER CHOICES: LISI or ZHANGSAN

The context sentences and the test sentences in Experiment 1 and Experiment 2 were the same. More specifically, the test sentences were designed to have a 'spillover' region of at least six words following the reflexive, to ensure that a potential slowdown could be detected. Within an experimental item, the spillover region always consisted of the same words. In addition, to keep the initial spillover region maximally similar across items, the first two words always indicated time by phrases such as *zhe libai* (“this week”), *shangge yue* (“last month”), or *zhege xinqisan* (“this Wednesday.”)

³ One additional participant was excluded from subsequent analysis due to strikingly slow reading times on *ziji*, which were 2957 ms on average, in contrast to the other participants' average of 734 ms. Such an extreme slowdown suggests strategic reasoning/processing rather than normal reading for comprehension, and this was also confirmed by the participant's post-experiment comments.

Procedure. Experiment 2 was run using the LINGER 2.88 software by Doug Rohde (available at <http://tedlab.mit.edu/~dr/Linger/>) for collecting self-paced word-by-word reading times on a PC. Each item consisted of two sentences. After the participants pressed the spacebar for the first time, the entire first sentence (the context sentence) appeared. Then, when they pressed the space bar for the second time, the second sentence (critical sentence) was presented one word at a time in a noncumulative, self-paced, moving-window format. When the participants finished reading each item, a comprehension question appeared with two possible answer choices. The *f* key always referred to the long-distance antecedent and the *j* key referred to the local antecedent.

3.2 Results

Choice of Antecedent. Consistent with the results of Experiment 1, there was a preference for the local antecedent in neutral contexts. In the neutral contexts, participants overwhelmingly chose the local antecedent over the long-distance antecedent as the referent of *ziji*, 74% to 26%. In the biasing contexts, we observed the opposite pattern: participants preferred the long-distance antecedent over the local antecedent, 67% to 23%. Analyses of variance confirm that the effect of the context manipulation is highly significant ($p < .001$).

Reading times. An analysis of participants' reading times shows clear effects of the context manipulation. Overall, participants' word-by-word reading times were significantly slower in the neutral context than in the biasing context, as shown in Fig. 3. This effect emerged at the reflexive (word position 5; $p < .01$) and persisted during the following words (word positions 6, 7, 8; p 's $< .05$), exhibiting a spill-over effect that is common in self-paced reading. Although there was no significant difference at word position 9, the difference was again significant at word position 10 ($p < .05$). There were no significant differences between the neutral and biasing conditions in the regions before the reflexive.

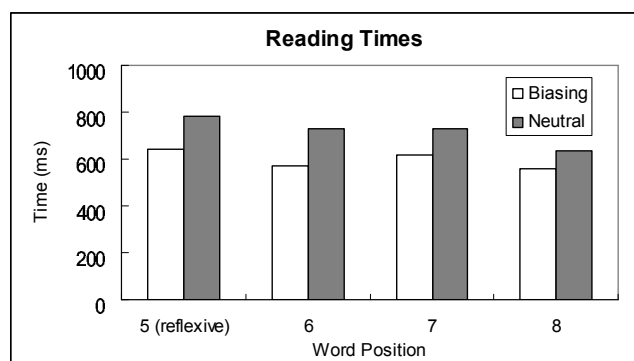


Fig. 3. Reading times in milliseconds for the reflexive *ziji* and the three words after it, in biasing and neutral contexts

4 Discussion

The results of Experiment 1 and Experiment 2 both showed that contextual constraints play a role in the interpretation of the Chinese simple reflexive *ziji*. In the offline survey of Experiment 1 we saw that a biasing context could overcome *ziji*'s default preference for the local antecedent, and in Experiment 2 we saw that contextual effects occur very early, starting at the reflexive itself. As a whole, our results fit well with the view that anaphor resolution is driven by multiple constraints, starting at the earliest moments of processing.

In Experiment 1, we also found that participants were more confident in their choice of antecedent in biasing contexts. Given that the reading times in the biasing contexts were also shorter in Experiment 2, we hypothesize that the biasing context reduced competition between the two antecedent candidates and thus weakened the slowdown that would otherwise be triggered by the presence of multiple candidate antecedents. More specifically, we suggest that the biased context increases the activation level of the long-distance antecedent and thereby decreases the amount of competition between the two possible antecedents – and this results in faster reading times.

Although there has been a lot of psycholinguistic research on how contextual factors influence pronoun interpretation (see e.g. [14] for an overview), the possibility of cross-sentential context effects on the interpretation of reflexives has received little attention in psycholinguistic work. In particular, our research sheds light on the real-time interpretation of a particularly flexible class of reflexives, anaphoric forms like *ziji* which can be bound by either a local or a long-distance antecedent. Our results demonstrated that although *ziji* has a default preference for a local antecedent, preceding discourse-contextual information can influence the interpretation of *ziji* very rapidly during real-time processing. The cross-sentential context effects emerged very early during processing (starting at the reflexive) and thus cannot be analyzed as a delayed (or second-stage) effect.

References

1. Chomsky, N.: Lectures on Government and Binding. Dordrecht, Netherlands: Foris (1981)
2. Gordon, P., Hendrick, R.: Intuitive Knowledge of Linguistic Co-reference. *Cognition*, 62, 325--370 (1997)
3. Arnold, J., Eisenband, J., Brown-Schmidt, S., Trueswell, J.: The Rapid Use of Gender Information: Evidence of the Time Course of Pronoun Resolution from Eyetracking. *Cognition*, 76, B13--B26 (2000)
4. Erlich, K.: Comprehension of Pronouns. *Quarterly Journal of Experimental Psychology*, 32, 247--255 (1980)
5. Garnham, A., Oakhill, J.: On-line Resolution of Anaphoric Pronouns: Effects of Inference Making and Verb Semantics. *British Journal of Psychology*, 76, 385--393 (1985)
6. MacDonald, M., MacWhinney, B.: Measuring inhibition and facilitation in pronouns. *Journal of Memory and Language*, 29, 469--492 (1990)
7. Kaiser, E., Runner, J.T., Sussman, R.S., Tanenhaus, M.K.: Structural and Semantic Constraints on the Resolution of Pronouns and Reflexives. *Cognition*, 112, 55--80 (2009)

8. Nicol, J., Swinney, D.: The Role of Structure in Coreference Assignment During Sentence Comprehension. *Journal of Psycholinguistic Research*, 18, 5--19 (1989)
9. Badecker, W., Straub, K.: The Processing Role of Structural Constraints on the Interpretation of Pronouns and Anaphors. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 28, 748--769 (2002)
10. Runner, J.T., Sussman, R.S., Tanenhaus, M.K. Assigning Reference to Reflexives and Pronouns in Picture Noun Phrases. *Cognitive Science* 30, 1--49 (2006)
11. Sturt, P.: The Time-Course of the Application of Binding Constraints in Reference Resolution. *Journal of Memory and Language*, 48, 542--562 (2003)
12. Chao, Y.-R.: *A Grammar of Spoken Chinese*. University of California Press, Berkeley, (1968)
13. Huang, C.-T. J., Tang, C.-C. J.: The Local Nature of the Long-distance Reflexive in Chinese. In Koster, J., Reuland, E. (eds.), *Long-distance Anaphora*, pp. 263--282. Cambridge University Press, Cambridge (1991)
14. Garnham, A.: *Mental Models and the Interpretation of Anaphora*. Psychology Press, New York (2001)