CONSCIOUS AND UNCONSCIOUS ACCESS TO GRAMMATICAL GENDER IN HEBREW

by

Tamar Hela Gollan

A Dissertation Submitted to the Faculty of the

DEPARTMENT OF PSYCHOLOGY

In Partial Fulfillment of the Requirements
For the Degree of

DOCTOR OF PHILOSOPHY

In the Graduate College

THE UNIVERSITY OF ARIZONA

1 9 9 8

INFORMATION TO USERS

This manuscript has been reproduced from the microfilm master. UMI films the text directly from the original or copy submitted. Thus, some thesis and dissertation copies are in typewriter face, while others may be from any type of computer printer.

The quality of this reproduction is dependent upon the quality of the copy submitted. Broken or indistinct print, colored or poor quality illustrations and photographs, print bleedthrough, substandard margins, and improper alignment can adversely affect reproduction.

In the unlikely event that the author did not send UMI a complete manuscript and there are missing pages, these will be noted. Also, if unauthorized copyright material had to be removed, a note will indicate the deletion.

Oversize materials (e.g., maps, drawings, charts) are reproduced by sectioning the original, beginning at the upper left-hand corner and continuing from left to right in equal sections with small overlaps. Each original is also photographed in one exposure and is included in reduced form at the back of the book.

Photographs included in the original manuscript have been reproduced xerographically in this copy. Higher quality 6" x 9" black and white photographic prints are available for any photographs or illustrations appearing in this copy for an additional charge. Contact UMI directly to order.

UMI

A Bell & Howell Information Company
300 North Zeeb Road, Ann Arbor MI 48106-1346 USA
313/761-4700 800/521-0600

CONSCIOUS AND UNCONSCIOUS ACCESS TO GRAMMATICAL GENDER IN HEBREW

by

Tamar Hela Gollan

A Dissertation Submitted to the Faculty of the

DEPARTMENT OF PSYCHOLOGY

In Partial Fulfillment of the Requirements
For the Degree of

DOCTOR OF PHILOSOPHY

In the Graduate College

THE UNIVERSITY OF ARIZONA

1 9 9 8

THE UNIVERSITY OF ARIZONA ®
GRADUATE COLLEGE

As members of the Final Examination Committee, we certify that we have read the dissertation prepared by Tamar Hela Gollan

entitled Conscious and Unconscious Access to Grammatical Gender in Hebrew

and recommend that it be accepted as fulfilling the dissertation requirement for the Degree of Doctor of Philosophy

Merrill F. Garrett	11-14-97 Date
Boyd W. Kaas	11/14/97 Date
K Doulas	11/14/97 Date
Susan Brown	11/14/97 Date
James Nicol	11/14/97 Date

Final approval and acceptance of this dissertation is contingent upon the candidate's submission of the final copy of the dissertation to the Graduate College.

I hereby certify that I have read this dissertation prepared under my direction and recommend that it be accepted as fulfilling the dissertation requirement.

Merrill F. Garrett — Dissertation Director

April 8, 1998 — Date

STATEMENT BY AUTHOR

This dissertation has been submitted in partial fulfillment of requirements for an advanced degree at The University of Arizona and is deposited in the University Library to be made available to borrowers under rules of the Library.

Brief quotations from this dissertation are allowable without special permission, provided that accurate acknowledgment of source is made. Requests for permission for extended quotation from or reproduction of this manuscript in whole or in part may be granted by the copyright holder.

SIGNED: 

ACKNOWLEDGMENTS

This dissertation work was an intense learning project which also led to many days of exciting and fulfilling activity. It could not have been completed without the unique contribution of several people whom I would like to thank.

The idea to use grammatical gender as a tool for evaluating models of language processing came from my advisor Merrill F. Garrett. Throughout my time in Tucson I was impressed with Merrill's optimism, and his knack for recognizing when his reassurance and problem-solving actions were needed. Both Merrill and Alfred W. Kaszniak played an essential role in promoting my interest in pursuing both clinical and cognitive neuropsychology, and in firmly but sensitively preventing me from biting off more than I could chew. Kenneth I. Forster provided hours of idea-sharing, and discussion, and also shared his passion for science and method for pursuing it. Ken also served (and continues to do so) as an always-available father in the psycholinguistics lab which became a home-base for the duration of my 5 years in graduate school. Both Ram Frost and Ken seemed to have endless patience for explaining, teaching, and guiding. To Ram I am especially thankful for adopting me completely into his lab at the Hebrew University in Jerusalem, and for showing me how not to miss the forest by obsessing about the trees.

As far back as I can recall, my parents and brother always encouraged me to pursue my dreams while simultaneously keeping my feet planted solidly on a practical ground. They protected me from, and taught me how to face on my own, the challenges that life presents. It is not possible to put into words the various ways in which I am grateful to them. Nina B. Silverberg provided years of fun, silliness, and unconditional emotional support both in solving day-to-day problems, as well as through major life struggles. As a friend and colleague she has been able to challenge my ideas in ways nobody else could. Linda Radzik and Thomas D. Bontly shared with me their strategies for surviving graduate school, and made me believe that I too have what it takes.

I could not have done this work without all of you. Thank you.

DEDICATION

For Ruth, Arye, and Amir.

TABLE OF CONTENTS

LIST OF TABLES

ABSTRACT

Access to grammatical gender in Hebrew was examined using gender decisions, lexical decision, two-word lexical decision, and grammaticality judgments (with noun-adjective pairs, noun-verb phrases, and plural noun-plural adjective pairs). In the gender decision task, the role of word-form was dramatic. Nouns with an explicit gender marker (regular feminine) were classified most easily, and next were nouns in the default gender (unmarked masculine). In contrast, exception nouns (unmarked feminine nouns) produced extremely slow RTs and unusually high error rates. These same exception nouns, however, produced radically different results when syntactic context was provided; in the detection of gender agreement (i.e., "yes" decisions in grammaticality judgments) exception nouns did not produce longer RTs, and also error rates were not higher than those for regular nouns. In contrast, there was a strong effect of regularity in the detection of gender agreement violation (i.e., "no" decisions in grammaticality judgments). This same pattern replicated even when all elements of the phrase were explicitly marked (i.e., plurals grammaticality judgments), and also in the context of a task that placed both grammatical and ungrammatical pairs in the context of the same decision type.

Because comprehension tasks necessarily begin with the analysis of word-forms where gender-marking regularity is most conspicuous, it is suggested that the results provide very strong evidence in favor of modular access mechanisms to lemma and lexeme representations. It is further suggested that access to the mental lexicon does not provide conscious access to grammatical gender. Efficient access is available only when

the syntactic processor is invoked, and this processor does not rely on regularity at the word-form level. Without a syntactic context grammatical gender is retrieved through a variety of strategies and heuristics. Finally, it appears that gender-marking regularity does play a role in post-lexical processing mechanisms that are designed to detect grammaticality violations, and the latter mechanisms, as well as syntactic processing itself, cannot be suppressed.

CHAPTER I - *Introduction*

In order to construct a grammatical sentence many languages require grammatical gender agreement. Grammatical gender must be retrieved for each noun because it must subsequently drive agreement with other sentential elements which refer back to each noun. For example, in the following Hebrew sentence ***ha-ugah hagdolah simxah et ha-yeladim haktanim*** *(The large cake made the small children happy)*, every open class element is gender marked. First, *cake* ***(ugah)*** is marked as a feminine noun with the suffix ***-ah***. Because *cake* is a feminine noun, the adjective *large* ***(gdolah)*** which modifies *cake* is also marked feminine, and so is the verb *made happy* ***(simxah)***. The next noun *children* ***(yeladim)*** is marked with the masculine plural form ***-im*** to indicate that the children are either all male or a mixed group of males and females (masculine gender is always the default in Hebrew), and finally the adjective *small* ***(ktanim)*** which modifies *children* is also marked with the masculine plural suffix. This simple sentence easily demonstrates how gender marking permeates many aspects of language production. Without the proper gender classification for each noun, a speaker of Hebrew is virtually unable to utter grammatical sentences. Although different languages vary with respect to which sentential elements must be gender marked (e.g., articles, pronouns, adjectives, and verbs), and with respect to how exactly the gender is marked, it is clear that rapid access to the correct grammatical gender is necessary in all of these languages for fluent speech production.

Second language learners seem to have significant difficulties with gender processing. Frequently they cannot access the correct gender specification, and they also

produce errors in agreement. In contrast, native speakers seem to access grammatical gender quickly and effortlessly. This raises the question as to how native speakers access the correct grammatical gender, and how it is represented in the mental lexicon. Accessing gender for animate nouns seems relatively easy because in most cases the grammatical gender coincides with conceptual knowledge. For example, the Spanish word for *woman* is the feminine noun *mujer*, and the word for *man* is the masculine noun *hombre*. Although the grammatical gender of these nouns is not obvious on the basis of word-form (in Spanish the suffixes ***-o*** and ***-a*** usually mark the masculine and feminine genders respectively), the conceptual knowledge in this case provides the necessary cue.

On occasion, however, grammatical gender is not correlated with semantic information even for animate nouns. For example, in German the words for *mother*, *aunt*, and *daughter* are feminine and are marked with the feminine definite article *die* (***die*** *Mutter*, ***die*** *Tante*, and ***die*** *Tochter*), and the words for *father*, *uncle*, and *son* are masculine and are marked with the masculine definite article *der* (***der*** *Vater*, ***der*** *Onkel*, and ***der*** *Sohn*). However, the word for *girl* is a conceptually feminine noun which happens to have the neutral grammatical gender and thus takes the neutral definite article *das* (***das*** *Madchen*). Similarly, the word for *boyhood* is neutral (***das*** *Knabenalter*), the word for *fatherhood* is feminine (***die*** *Vaterschaft*), and there are two words which mean both *woman* and *wife* one is feminine (***die*** *Frau*) and one is neutral (***das*** *Weib*). All other sentential elements which refer back to these exception nouns must be marked for grammatical gender (e.g., all gender-marked words which modify *girl* will be marked neutral). Thus, these kind of animate nouns are like the inanimate nouns in the sense that

semantic knowledge does not provide a correct cue to the grammatical gender.

In contrast to animate nouns, inanimate nouns, bear an entirely arbitrary relationship to semantic information. This fact is best exemplified by words which are very close in meaning, but are nevertheless assigned different grammatical genders. For example, the Spanish word for *bed* is assigned the feminine gender (*cam**a***), but the synonym (*lech**o***) also means *bed* and is masculine. A further demonstration of the arbitrary relationship between grammatical gender and semantics is that the gender of translation equivalents varies greatly from language to language. For example, the Spanish word for *rain* is feminine (*lluvi**a***), whereas the Hebrew word for *rain* happens to be masculine (*geshem*).

In many languages the grammatical gender can be identified regardless of semantic opacity because of gender-correlated markings on the word-forms themselves. For example, most Spanish nouns are marked with the above mentioned suffixes *-o* (masculine) and *-a* (feminine), and are therefore *transparent* with respect to grammatical gender. However, there are many *opaque* cases in which the grammatical gender cannot be inferred from markers on the word-form. For example, *leche* the Spanish word for *milk* is feminine and thus correct agreement produces *leche sabros**a*** (*tasty milk*), whereas the similarly ending *coche* (*car*) is masculine, and thus correct agreement produces *coche barat**o*** (*cheap car*). Worse yet, there are several cases in which the marking on the word-form actually conflicts with the grammatical gender which drives agreement. For example, the Spanish word for *hand* ends in the typically masculine *-o (mano)*, however its grammatical gender is feminine, and hence correct agreement calls for a feminine

adjective as in *mano bonita (pretty hand)*.

In some languages, like German and French, only some nouns are marked with one of many gender-predictive endings, and the only perfectly reliable cue to gender is the gender-marked article. For example, in French words that end in ***-ie, -ette, -ue,*** and ***-tion*** tend to be feminine, and words that end in ***-i -u,*** and ***-age*** tend to be masculine. However, most inanimate nouns in French (and German) are opaque. Since gender agreement is not optional in any of these languages, the correct grammatical gender must be accessed in all cases, even when it cannot be recovered from either word-form or on the basis of conceptual knowledge. A major theoretical question therefore is how grammatical gender is stored and accessed in the course of language processing.

Most theories of language production make a separation between a stage of processing in which semantic and syntactic lexical features are accessed, and a second stage in which the corresponding lexical forms are accessed (Bock, 1982; Butterworth, 1989; Dell, 1990; Fromkin, 1971; Garrett, 1975; Levelt, 1989). Although many 2-stage models were developed to account for production processes, similar claims and debates are frequently discussed in models of language comprehension (see Balota, 1992; Frauenfelder & Tyler, 1987, and Levelt, 1992 for reviews). Moreover, comprehension processes may have implications for production processes, and there are likely to be intimate connections between the two (Cutler, 1987; MacKay, 1987). The motivation for two-stage access is to maximize the efficiency of lexical search in both directions - production and comprehension (Fay & Cutler, 1977); for language production there is a lexicon which can only be searched with semantic properties, and for comprehension

there is a lexicon which can only be accessed via a form-based search.

The set of representations accessed in the first stage of language production is often called the lexicon of *lemmas*, and the set accessed in the second stage is called the lexicon of *lexemes* (Kempen & Huijbers, 1983). Because language production proceeds from propositional content to surface form, the connection from conceptual memory to lemma representations must involve a semantic search. Therefore, the lemmas are organized according to semantic principles (see Garrett, 1992a for a discussion about the possible nature of the semantic organization of the lemma lexicon). Once accessed the lemmas provide the lexical syntactic roles that are necessary for grammatical phrase construction, and also an address (or pointer) which eventually links the lemmas to the lexemes. Because the lemmas provide word-specific syntactic features, the lemma is the logical level to consider for the representation of grammatical gender.

In contrast, to production, comprehension processes must connect the words of a language with their conceptual meaning. Thus, the word-forms, or *lexemes*, are stored in a separate lexicon that is organized by principles of surface form whether phonological or orthographic (e.g., Marslen-Wilson, 1987; Forster, 1994; Grainger, 1992; and for a detailed discussion of 2-stage models see Levelt, 1989; Bock & Levelt, 1994). Note that with respect to grammatical gender the precise manner by which information flows between lemmas and lexemes is of particular interest when language comprehension (rather than production) is examined. For if lexemes often provide opaque (or even explicitly misleading) gender information, then how is the correct gender ultimately retrieved?

If gender is represented at the lemma level, then its ultimate retrieval during speech or print recognition necessarily depends on the manner by which lexemes and lemmas exchange information. Jescheniak & Levelt (1994) offered a very specific characterization of gender representation; in their model the lexicon is a 3-layer network (conceptual, syntactic/lemma, and phonological/lexeme) with spreading activation. The representation of gender is modeled with a single node for each gender at the lemma level, and each noun of a given gender class is then linked to the appropriate gender node. This particular model is modular; in language production information flows unidirectionally from lemmas to lexemes, and in comprehension it flows from lexemes to lemmas without interaction or feedback between levels (Garrett, 1976;1980b, Levelt, 1989; Butterworth 1989). In this architecture, then the quasi-regular gender information that can be inferred from the lexeme level plays no role in accessing gender during language production. Moreover, for any given noun in language comprehension the lexeme can activate the appropriate gender node only by first activating the lemma node. Because the gender node does not receive direct activation from the lexeme level, this model therefore predicts that the computation of agreement for transparent and opaque nouns should be equally easy even in language comprehension despite the rather obvious differences in their surface form.

Several studies, however, have suggested an interactive access mechanism, such that feedback from the lexemes can influence processing at the lemma level, and vice versa (Dell, 1986, 1988; Dell & O'Sheaghda, 1991,1992; Stemberger, 1985). In this model explicit markings at the level of form would activate the appropriate gender node

at the lemma level. Thus, activation of the appropriate gender node would be greater for transparent nouns, and access to gender would proceed more quickly (in both language production and comprehension) for transparent nouns. Conversely, the gender of opaque nouns would be accessed more slowly because there is nothing at the lexeme level to facilitate access to the appropriate gender node. Further, nouns for which there is a conflict between lemma level gender information and the gender indicated at the form level (e.g., ***la mano***) should be even more difficult to access than opaque nouns. By the same argument, access to gender should also be easier for animate nouns (in both production and comprehension) because semantic information at the lemma level provides the necessary cue. Finally, it is interesting to note that the interactive model also predicts that the computation of gender agreement should be easier overall in a language with mostly transparent nouns (e.g., Spanish) relative to a language with mostly opaque nouns (e.g., German).

Another crucial difference between the modular and interactive models is in the role of gender node. An interactive model of gender representation could allow for activation to spread from the form level directly to the gender node. Thus, explicit gender marking would not only facilitate access to the gender node rather it could also facilitate access to the lemma itself (i.e., lexical access in general). In the modular model, on the other hand, the gender node plays no role in either lemma or lexeme access; its sole role is to activate the appropriate gender marked closed-class items (e.g., articles) and to specify the appropriate gender-marked affixes within the syntactic frame.

Presently, there is relatively little evidence that provides unequivocal support for

either the interactive or the modular approach. Studies that investigated access to grammatical gender in tip-of-the-tongue (TOT) states in Italian showed that access to gender was preserved even when information about word-form was entirely absent (Badecker, Miozzo, & Zanuttini, 1995; Vigliocco, Antonini, & Garrett, 1996). These findings support the notion of an abstract representation of grammatical gender at the lemma level that can be accessed independently of access to lexemes. More critically, Badecker et al., (1995) and Vigliocco et al. (1996) also found that access to gender for was preserved equally for transparent and opaque nouns suggesting that gender marking regularity at the word form level does not affect the retrieval of gender information. This outcome seems more consistent with the modular view. These results, however, are perhaps not entirely convincing because they come from cases of retrieval failure. It may be possible that feedback from the lexeme level could influence access to the lemma level when intact language production proceeds.

To evaluate the influence of gender on lexeme access, Bates, Devescovi, Pizzamiglio, D'Amico, & Hernandez (1995) compared transparent and opaque nouns in two tasks (i.e., speeded auditory word-repetition and gender decision) that they showed to be sensitive indices of word recognition. They reasoned (as suggested above) that transparency could be used as a cue in the process of lexical access itself (not just in accessing grammatical gender). In this case transparent nouns should be repeated more quickly because the gender of these nouns is more easily detected than the gender of opaque nouns. However, they found no difference between opaque and transparent nouns, and hence they concluded that gender marking does not influence access to the

lexeme. In the same study, however, a large difference between opaque and transparent nouns was found in the gender-decision task (in which subjects were asked to explicitly report the gender of nouns). To accommodate both results, the authors concluded that gender is accessed at the lemma level *after* the lexeme has already been accessed. Further, because of the task-specific nature of the transparency effect, they attributed it to a post-lemma access-checking mechanism, and therefore argued that their results are more consistent with the modular view.

Similar conclusions were reported by Jescheniak & Levelt, 1994 who evaluated word-frequency as a predictor of article-decision times in Dutch (this task is very similar to gender decision; it requires subjects to indicate the correct gender-marked article for each noun presented). They assumed that access to gender should be frequency sensitive if gender is accessed at the lexeme level. In this study, word-frequency effects in gender decision attenuated rapidly (in fact they disappeared entirely) after just one repetition of the items, whereas the same words in tasks which did not require access to gender showed stable frequency effects over several repetitions. Based on these data the authors concluded that the initial effect of frequency on gender-decision could be attributed to a recency effect in the retrieval of gender from the lemma. This effect reaches maximum strength after a single presentation and subsequently decays slowly. Moreover, they argued that since word-frequency is the hallmark of lexical access, the absence of a stable frequency effect in the gender decision task suggests that grammatical gender is not accessed at the lexeme level.

For entirely independent reasons both Bates et al. (1995) and Jescheniak et al.

(1994) concluded that access to gender occurs post-lexically. Recently Taft and Meunier (1995) have produced further data that support this conclusion. In this study, frequency and transparency were examined simultaneously using gender decisions in French. As mentioned above, certain suffixes in French can be used as a cue to grammatical gender. For example, nouns that end in *-ette* tend to be feminine, and nouns that end in *-age* tend to be masculine. However, there are exceptions to these rules and hence like in Italian, orthographic cues to gender in French are quasi-regular. Taft and Meunier (1995) reasoned that if transparency (a lexeme level property) influences access to the lemma level, then an interaction between frequency and regularity should be obtained whenever access to gender is required. Although there were robust effects of frequency and regularity in gender decisions, there was no interaction between these effects.

The absence of an interaction between frequency and regularity suggested that these two factors influence access to gender in distinct processing stages. This conclusion is more consistent with the modular view because the interactive models are committed to predicting a **pre**lexical effect of regularity on access to grammatical gender. Although there is no doubt that robust effects of regularity have been obtained in gender decision, the critical question in terms of what can these effects do to distinguish between models is whether these effects occurred pre- or post-access. Frequency becomes relevant in this context as a measure of pre- vs. post- lexical effects; an interaction between predictors (e.g., frequency and gender marking) suggests they may play a role in the same access stage. Thus, if regularity and frequency interact in gender decision then there is reason to think that, like the effects of frequency, the effects of regularity are

pre-access and this provides support for the interactive view. On the other hand, if there is no interaction then gender-marking must be affecting decisions in a post-lexical stage.

A further way that Taft and Meunier (1995) showed that access to gender occurred post-lexeme access was by showing that the decisions depended on retrieval of the definite article (which obviously has to come after retrieval of the noun itself). That is, gender decisions were severely disabled for nouns that call for gender-neutral articles (e.g., *l'australie*; gender-marked articles in French almost always provide unequivocal gender information except in cases when the article is abbreviated). This finding suggested that gender decisions were computed with an article-retrieval strategy, thereby implying that access to the lexeme does not provide conscious access to the gender specification.

Thus, by demonstrating that gender information is retrieved <u>following</u> access to lexemes latter studies seem to support the notion of an abstract representation of gender at the lemma level. In addition, they also seem to support the modular view by demonstrating that transparency affects access to gender only in post-lexical processing stages. These conclusions, however, are not universally accepted. For example, Bates, Devescovi, Hernandez, & Pizzamiglio, (1996) showed that subjects were faster to repeat a noun when it appeared after a concordant adjective (e.g., ***brutt<u>a</u> cas<u>a</u>***, both are feminine thus having the same suffix), than if it appeared after a neutral adjective (e.g., ***grand<u>e</u> cas<u>a</u>***), or discordant adjective (e.g., ***brutt<u>o</u> cas<u>a</u>***). Thus, the phrase priming tasks did not focus on transparency, rather they showed that gender facilitates lexeme access (for both transparent and opaque nouns alike). This same basic result has been replicated in two-

languages and four different tasks: gender-decision, auditory word-repetition (Bates et al., 1996), a gating paradigm, and also in a simple lexical decision task (Grosjean, Dommergues, Cornu, Guillelmon, & Besson, 1994). These data seem to demonstrate effects of transparency on both lexical access itself, as well as on access to grammatical gender. Although Grosjean et al., (1994) raise the possibility that the effects can be explained as post-lexical syntactic integration, they ultimately dismiss this possibility primarily based on evidence from the gating paradigm. They argued that the gender primes appeared to be altering the actual access path because with a gender prime responses were limited to the target's gender only. Bates et al. (1996) also suggested that the gender primes seem to facilitate lexical access itself. This was a conclusion which was at odds with Bates et al., (1995). However, they ultimately also dismissed the distinction between pre- and post-lexical effects entirely, and proposed to eliminate the "distinction between grammar and the lexicon." In either case, if it is assumed that gender is stored at the lemma level, and as Grosjean et al., (1994) suggested syntactic integration cannot account for the phrase-priming effects, then these results are more consistent with the interactive view.

To summarize, both the interactive and the modular accounts of access to grammatical gender gain some support from the experimental evidence. The critical test of these models concerns the differences (or similarities) in accessing the grammatical gender of transparent and opaque nouns. To this end, data from several languages are useful because languages display a wide variety in the way gender is marked in the level of form, and in the regularity and reliability of the marking system. In the present paper

this issue was addressed by examining access to gender information in the Hebrew language.

This investigation was aimed first to investigate whether access to the lexeme level provides conscious access to grammatical gender, and second to assess access to gender in the context of tasks that vary the extent to which explicit attention to gender is required. More specifically, the goal was to assess the effects of word-form transparency on access to the gender (i.e., at the lemma level). The results are discussed with regard to their relevance for models of grammatical encoding in both language production and comprehension.

Characteristics of gender-marking in Hebrew: All nouns in Hebrew are assigned one of two grammatical genders. There is no neutral gender in Hebrew. Masculine singular nouns are unmarked (e.g., *sefer*, the word for *book)*, whereas feminine singular nouns are usually marked by the suffixes ***-ah***, ***-et***, or ***-it*** (e.g., *sapah*, *rakevet*, *zavit*, the words for *couch*, *train*, and *angle* respectively). Verbs, pronouns, and adjectives are gender inflected usually in the form of the above suffixes which are attached to the base form whenever a feminine noun is modified. In addition, gender marking appears also in the plural forms. For those, however, marking exists in both masculine and feminine nouns. Masculine nouns are usually pluralized with the masculine plural suffix ***-im*** (e.g., ***sfarim*** the word for *books*), and feminine nouns are usually pluralized with the feminine plural suffix ***-ot*** (e.g., *sapot*, *rakavot*, *zaviyot*, the words for *couches*, *trains*, and *angles* respectively).

As in many other languages, the gender-marking system in Hebrew is quasi-

regular; there are many exceptions to the rules described above. First, there is a subset of feminine nouns that are unmarked like the typically masculine nouns (e.g., *nefesh* and *even*, the words for *soul* and *stone*, are feminine yet they do not end in a feminine suffix). There are also very few masculine nouns that end in the typically feminine suffixes. However, common errors in gender inflection with these nouns seem to suggest that native speakers are not aware of their masculine gender (e.g., *tsomet* the word for *junction* is masculine although in ends in a typically feminine suffix, however, many Hebrew speakers are not aware of this fact). There is also a large number of exceptions in plural suffixing; many masculine nouns take the feminine plural suffix *-ot* (e.g., *yerek* --*>yerakot*, the word for *vegetables*), and a number of feminine nouns take the masculine plural suffix *-im* (e.g., *beytsah* --> *beytsim*, the words for *egg* and *eggs*). In addition, some unmarked feminine nouns take the masculine plural suffix (e.g., *even* --> *avanim*, the irregular feminine word *stone* also takes the masculine plural thus making it further irregular), and some unmarked feminine nouns take the feminine plural suffix (e.g., *nefesh* the word for *soul* becomes *nefashot* in the plural form). Table 1 provides examples and a rough approximation of the probabilities of encountering regular and irregular gender marking in Hebrew.

An important feature of gender marking in Hebrew is that gender cannot be inferred from a definite article preceding the noun as in French, Italian, German, Dutch, or Spanish, and thus there is no simple strategy for accessing gender information by considering both article and noun as a single unit. Adjectives are always gender marked appropriately in Hebrew, however, obviously many adjectives could qualify a given noun

and hence it seems unlikely that there could be a stored propositional unit composed of a given noun and a given adjective that would reveal the noun's gender. Thus, in contrast to article attachment (see Taft & Meunier, 1995), gender-marked adjectives cannot easily provide a simple strategic route to gender information. The following experiments were designed to explore the role of form transparency on access to gender in Hebrew, which like many other languages, is characterized by a quasi-regular gender-marking system.

CHAPTER II - *Simple Gender Decision*

In Experiment 1 the gender decision task was used to examine access to grammatical gender for transparent and opaque nouns in Hebrew. Subjects were presented with transparent and opaque nouns and were instructed to make fast but accurate decisions concerning their gender. Gender decision studies in French and Italian have revealed strong effects of transparency, but have produced conflicting results with respect to the effect of frequency on this task (Jescheniak & Levelt, 1994; Bates et al., 1995; Taft & Meunier, 1995). This experiment therefore was designed to reevaluate the effects frequency and transparency in the gender decision. If gender is accessed at the lexeme level, then robust effects of frequency should be obtained in gender decisions.

Method

Participants: Forty undergraduates from the Hebrew University in Jerusalem participated in the experiment for course credit or for payment. All subjects were native speakers of Hebrew.

Stimuli and design: In this first gender-decision experiment 200 Hebrew nouns were presented in the singular form. The nouns were selected to reflect the quasi-regular Hebrew gender-marking system in both the singular and the plural forms. Masculine nouns: There were 50 regular (unmarked) masculine nouns which take the regular masculine plural suffix ***-im*** and 50 regular (unmarked) masculine nouns which take the

irregular (i.e., feminine) plural suffix *-ot*. Feminine nouns: There were 50 regular (marked) feminine nouns which take the regular feminine plural suffix. In addition included 20 irregular (unmarked base form) feminine nouns, and 16 regular (marked) feminine nouns which take the irregular (masculine) plural suffix were included.

Because nouns with grammatical gender that could be inferred on the basis of semantic information were not included, and because extremely low-frequency words were excluded, the imbalance of regular vs. irregular items (both in base-form and in plural suffixation of feminine nouns) could not be avoided. This merely reflects the restricted number of nouns available in these respective categories. However, in order to create a list with half masculine, and half feminine nouns (i.e., to achieve a probability of 0.5 for each response in the gender decision task) 14 additional regular feminine nouns were included as filler items. As mentioned above, irregular masculine nouns are too few in Hebrew and therefore they were not included in the experiment.

Thus, overall six types of nouns were included in the stimulus list roughly according to the extent that they occur in the Hebrew language. The six types were masculine nouns with regular or irregular plural suffixation, marked feminine nouns with regular or irregular plural suffixation, and unmarked feminine nouns with regular or irregular suffixation. Since there is no frequency-count in Hebrew, 50 undergraduates (who did not subsequently participate in the gender-decision experiment) were tested on a questionnaire which assessed noun-frequency ratings on a scale from 1 to 7 with seven being "very frequent" and one being "very infrequent". The results produced an average frequency count of 4.03, with a standard deviation of 1.04, and the range was 1.78-6.57.

Procedure: Each participant was tested individually, in a soundproof room. Verbal instructions indicated that nouns would appear one after another on the computer screen, that she/he was to decide if the nouns was masculine or feminine, and to respond by pressing the appropriate key (masculine on the right, and feminine on the left). Both accuracy and speed were stressed as important. Nouns appeared centered on the screen of a Macintosh SE computer. A millisecond clock which recorded response times (RTs) was time-locked to the presentation of each noun. Each noun remained on the screen either until a response was recorded, or until 2 seconds had elapsed. There was a 2 second pause between each trial. The experiment began with 10 practice items, and subsequently proceeded with the 200 test items.

Results.

Mean RT and standard deviations for correct responses were calculated for items in each noun type. Within each subject/condition combination, RTs that were outside a range of 2 SDs from the respective mean were excluded, and the mean was recalculated. Outliers accounted for less than 5% of all responses. Gender (masculine, feminine), transparency of singular form (regular/irregular marking), regularity of plural form (regular and irregular suffixation) and frequency were evaluated as predictors of RTs and error rates with bivariate correlations, and in the case of significant results, analyses of variance (ANOVAs) were subsequently computed as well.

The correlation of gender with RTs was robust $r=+.56$, and also with errors $r=+.67$. The ANOVA also revealed a main effect of gender in both RTs $F(1,184)=7.85$,

MSE = 4,036 p<.01 and errors *F(1,184)=13.83, MSE = 181 p<.01*, with masculine nouns producing an average RT of 728 ms and 9.7% errors, and the average RT for feminine nouns was 701 ms, and 17.0% errors. However, a strong effect of transparency was also found in a separate analysis of the feminine nouns only, indicating that the main effect of gender was confounded with transparency. Whereas decisions to regular marked feminine nouns were very fast and accurate (671 ms and 9.4% errors), gender decisions to the irregular (unmarked) feminine nouns were much slower (784 ms) and also produced an exceedingly large error rate (37.9% errors). These differences in performance were significant in both the analysis of RTs *F(1,84) = 97.56, MSE = 2,188 p<.01* and in the analysis of errors *F(1,84) = 82.05, MSE = 167 p<.01*. RTs and errors of masculine nouns, marked feminine nouns, and unmarked feminine nouns are presented in Table 2.

The effect of noun-frequency on decision latencies was assessed treating frequency as a continuous variable and plotting its correlation with RTs and errors. Neither of these was found to be correlated with word frequency (*r=+.07*, and *r=+.02* respectively). Further, an ANOVA in which only the 30 highest frequency nouns were compared to the 30 lowest frequency nouns produced no significant effects for both RTs *F(1, 58)=2.97, MSE = 3,415 p = 0.09* and for errors *F(1,58) = 0.13, MSE = 154 p = 0.72*.

Finally, plural-suffix regularity was evaluated as a predictor of RTs and errors. Although nouns were always presented in the singular form, gender decisions for nouns with an irregular plural suffix were slightly slower and less accurate than decisions to

nouns with a regular plural suffix (see Table 1). These differences, however, did not reach significance $F(1,182) = 1.57, MSE = 4,176 p = 0.21$ in the analysis of RTs and $F(1,182) = 1.60, MSE = 194 p = 0.21$ in the analysis of errors.

Discussion

The results of Experiment 1 revealed a dramatic effect of transparency. Marked feminine nouns produced the fastest RTs, masculine nouns (which are always unmarked) also produced relatively fast responses, and irregular (unmarked) feminine nouns produced the slowest RTs. A very similar pattern emerged in the analysis of errors; regular nouns (i.e., marked feminine and unmarked masculine) produced about 10% errors, whereas irregular nouns (i.e., unmarked feminine nouns) produced nearly 40% errors.

This unusually large proportion of errors for irregular nouns calls into question whether the gender decision task allows access to grammatical gender as it occurs in the course of natural language processing. It suggests instead, that at least in Hebrew, the gender decision task generates responses which are mainly based on strategies related to word form analysis. More specifically, the results seem to demonstrate that subjects based their decisions on a search for the feminine suffix ***-ah***. This strategy resulted in a high probability of a "feminine" response to regular feminine nouns, a high probability of a "masculine" response to all unmarked nouns, thereby producing a very high error rate in responses to irregular (unmarked) feminine nouns. A number of other aspects of the results are also consistent with this nonlexical strategy hypothesis. First, RTs to marked

feminine nouns were fastest, suggesting that the default strategy involved searching for feminine suffix marking. Second, there was no effect of frequency on response times and errors. This outcome suggests that a large proportion of the responses were made based on word form analysis rather than on lexical search.

That said, however, if word-form analysis were the sole factor contributing to the computation of a response then unmarked feminine nouns should have produced 100% errors, and clearly this was not the case. Therefore word-form analysis could not have been the sole predictor of response times. In order to produce correct responses to any of the unmarked feminine nouns some lexical level involvement must have been present. One possibility is that the word-form analysis was always checked against lexically generated gender information. Such a mechanism requires the assumption that the gender-decisions actually involved two steps. First, the form based nonlexical response was generated, and subsequently it was checked against lexically generated information. Note, that since only a small proportion of nouns were irregular in Experiment 1 (about 10%), the nonlexical strategy produced correct responses in most cases and it is therefore unlikely that the suffix-checking strategy would have been dropped altogether.

This two-step task analysis of the gender-decision task predicts different outcomes for regular and irregular nouns in the gender decision task. First, as responses to regular nouns via the nonlexical strategy were compatible with the information retrieved from the lexicon, correct decisions could have been given very soon after the lexical representation had been accessed. In contrast, for irregular nouns the nonlexical

strategy (that relies on form analysis) would have biased responses in one direction (e.g., "masculine"), whereas subsequent access to the lexical representation would have produced the opposite response (e.g., "feminine"). By this view, the reaction times and error rates for irregular nouns, were caused by this conflict (which could only be resolved by an additional lexical check).

The difference between regular and irregular nouns within this two-step model of gender decision leads to an interesting prediction. Responses to regular nouns could have been made immediately after the lexical entry was accessed, whereas responses to irregular nouns must have been delayed until well after the entry was accessed to allow time for the conflict to be resolved. For this reason, we might expect to find a frequency effect for regular nouns, but not for irregular nouns in Experiment 1. Consistent with this argumentation, a post-hoc ANOVA of regular nouns only in which the 30 highest frequency nouns (mean RT 695 ms) were compared to the 30 lowest frequency nouns (mean RT 728 ms) revealed a significant effect of frequency $F(1, 58) = 4.13$, MSE = 5185 $p = 0.05$ in the analysis of RTs (but not in the analysis of errors). Also consistent with this 2-step gender decision task analysis were the results of an additional post-hoc analysis of plural suffix regularity. Because of the abundance of exceptions in plural suffixation in Hebrew, it is highly likely that access to the plural suffix would be lexically mediated rather than rule-generated. In this respect, the plural suffix effect can be taken as yet another indicator (in addition to frequency) of lexical involvement in response computation. In order to test this idea, an additional post-hoc analysis of plural suffix regularity was evaluated for regular nouns only. As predicted, the effect of plural suffix

regularity was indeed significant in this reanalysis $F(1,161) = 4.34$, $MSE = 3,484$ $p < .05$ in both the analysis of RTs, and also in the analysis of errors $F(1,161) = 21.37$, $MSE = 54$ $p<.01$. The presence of an effect of the plural suffix regularity in an experiment that did not include a single item in the plural form is rather intriguing. It suggests that it is difficult to gain conscious access to grammatical gender even after the lexical representation has been accessed, and hence in a language like Hebrew which does not have gender-marked definite articles subjects resort to reliance on the considerably less reliable plural suffix. Recall that there are many exceptions in plural suffixation whereas the definite article in French is gender marked and is 100% reliable. Nevertheless, because it was generated post-hoc, this explanation demands further verification. Note that although it follows from the two-step gender-decision mechanism, the presence of lexical effects (frequency and plural suffix regularity) for regular nouns only is somewhat counterintuitive. A priori, it could have been argued instead that only irregular nouns (or perhaps all the unmarked nouns) should be most likely to produce lexical effects, because for these nouns lexically generated information is the only way to compute a correct response.

Finally, the large proportion of errors for unmarked feminine nouns suggests even if the two-step analysis of the gender decision task is correct, of these two the form-based analysis was weighted more heavily than lexical checking. This is not surprising because only a small proportion of nouns in Experiment 1 (about 10%) were irregular, and thus the form-based analysis would have produced a correct answer in the vast majority of cases. This result also suggests that it is more difficult to access to lexical information

than it is to compute a superficial analysis of word-form.

In light of these conclusions, it seems that the predominance of strategically driven responses in Experiment 1 could have prevented a proper evaluation of the role of regularity in accessing grammatical gender. In Experiment 2 therefore, the possibility of nonlexically based was eliminated by excluding marked feminine nouns from the stimulus list.

CHAPTER III - *Gender Decision Without the Easy Route*

In order to examine access to gender information, Experiment 2 was designed to reduce the possibility of adopting a nonlexical heuristic. The manipulation consisted of simply excluding all the marked feminine nouns, thus creating a design that included only unmarked nouns (half irregular feminine and half masculine). With unmarked nouns only, gender decisions cannot be based on the identification of feminine suffixes ***-ah***, ***-et***, and ***-it***. Thus, the aim in Experiment 2 was to examine whether stronger lexical involvement and consequently a more robust frequency effect, would emerge in the gender decision task. It was further hypothesized that under these conditions (i.e., deeper lexical involvement) there would be a lower rate of errors for unmarked feminine nouns (relative to Experiment 1). Note that although this manipulation effectively eliminated transparency as a variable, it could be argued that regularity was still a variable in Experiment 2 because feminine nouns are usually marked whereas unmarked nouns are typically masculine. Thus, an additional aim was to examine whether this level of regularity affects gender decisions as well.

Method

Participants: Forty undergraduate students from the Hebrew University in Jerusalem participated in the experiment for course credit or for payment. All subjects were native speakers of Hebrew, and none had participated in Experiment 1.

Stimuli and design: In this second gender-decision experiment 27 masculine

nouns, and 27 unmarked feminine nouns were selected and presented in the singular form. All words but 3 of the unmarked feminine nouns were taken from Experiment 1. Of the feminine nouns 15 called for the feminine plural suffix, and 12 for the masculine plural suffix. Similarly, of the masculine nouns 15 called for the masculine plural suffix, and 12 for the feminine plural suffix. The masculine items were frequency matched to the feminine items producing an average frequency count of 4.31 (SD = 1.12) for the feminine nouns, and 4.25 (SD = 1.02) for the masculine nouns. The procedure and apparatus were identical to those in Experiment 1.

Results.

Mean RT and standard deviations for correct responses were calculated for items in each noun type. Within each subject/condition combination, RTs that were outside a range of 2 SDs from the respective mean were excluded, and the mean was recalculated. Outliers accounted for less than 5% of all responses. One item was dropped from the analysis because of an unusually high error rate. The results are presented in Table 3. Responses to masculine (regularly unmarked) nouns were faster than responses to the unmarked (therefore irregular) feminine nouns (847 and 889 ms respectively). Moreover, the unmarked feminine nouns yielded twice as many errors (24.8%) as the masculine nouns (12.4%). An ANOVA revealed that these differences were significant $F(1,51) = 3.81, MSE = 6,183\ p = .056$ for RTs, and $F(1,51) = 10.90, MSE = 187\ p < .01$ for errors.

As hypothesized, and in contrast to the results of Experiment 1, there was a significant correlation between word frequency and RTs $r = +.27$. The ANOVA

calculated on the 12 highest and 12 lowest frequency items also revealed a frequency effect; the mean RTs for frequent nouns was 838 ms, and the mean for nonfrequent nouns was 887 ms, $F(1,22) = 13.91$, $MSE = 1,036$ $p<.01$. The effect of frequency on errors was not significant ($F < 1$).

Finally, as in Experiment 1 mean RTs for nouns with a regular plural suffix were longer than for nouns with an irregular plural suffix (857 ms vs. 880 ms), and also error rates were higher (17.9% vs. 19.2%). Once again, however, these differences were not significant $F(1,51) = 1.09$ in the analysis of RTs, and $F(1,51) < 1.0$ in the analysis of errors.

Discussion

The first critical result in Experiment 2 is the presence of a significant frequency effect. This finding suggests that the elimination of all transparent items from the task was successful in discouraging a nonlexical decision strategy in the computation of grammatical gender, and concomitantly in encouraging a stronger lexical involvement. However, note that in spite of the evidence for lexical search (as revealed by the frequency effect), and despite the fact that half of the items presented in this experiment were unmarked feminine nouns, there was a persistently large percentage of errors (about 25%) in responses to these nouns. If the gender-decision task accurately reflects the normal mechanisms of access to grammatical gender, then the elimination of nonlexical heuristics should also have eliminated the unusually high proportion of errors for these nouns. Moreover, RTs to unmarked (and therefore irregular) feminine nouns were also

considerably longer than RTs to masculine nouns. This effect of gender-marking regularity requires explanation.

In Hebrew, transparency and regularity are confounded because the only nouns that are explicitly gender-marked are feminine. In this sense, it could be argued that unmarked feminine nouns are really more comparable to nouns like *la mano* in Spanish (and Italian) in which the form is actually misleading rather than just neutral. Since unmarked nouns in Hebrew are usually masculine this creates a general underlying bias, that may have encouraged yet another nonlexical strategy in Experiment 2. In this case, rather than first focusing on a search for the feminine suffix, the heuristic would first involve the assumption that unmarked nouns are masculine. To classify a noun as "feminine" then, the decision process could function under a deadline mechanism that waits for evidence contrary to the default "masculine" decision; if no evidence is found, then the response would always be "masculine". Such a nonlexical heuristic would again have produced a higher proportion of errors for the unmarked feminine nouns.

The results of both Experiments 1 and 2 lead to similar conclusions i.e., mere access to the mental lexicon does not provide conscious access to grammatical gender. However, if access to the lexicon does not provide gender information, then how can the frequency effect in gender-decisions be explained? One possible account posits that access to the lexeme level consists of only one component of the gender decision task. In a comprehension task, the lexeme is accessed first and only subsequently does access to the lemma level begin. Hence, although gender is ultimately accessed from the lemma

level, a frequency effect is still expected because at least one component of the task is frequency sensitive. However, because lexeme access is frequency sensitive whereas lemma access is not (see Jescheniak & Levelt, 1994 for a series of experiments that isolate the locus of frequency effect to the lexeme level), the frequency effects in gender decision should be less robust than they are in lexical decision.

The most important conclusion to be drawn from the results of Experiments 1 and 2, however, seems to be a persistent incompetence on the part of native Hebrew speakers in computing gender decisions. This suggests that the gender decision task may not reflect the normal mechanisms of access to grammatical gender. Note that similar concerns about the gender decision task were raised by Jescheniak & Levelt (1994); they too considered (albeit briefly) the possibility that the gender-decision task is not "the royal road to the lemma." What characterizes the gender-decision task is that it requires a conscious report of a syntactic feature of a noun, whereas in normal language processing there is seldom any reason for having explicit access to such information. In fact, it could be argued that conscious access to gender information (or other syntactic information) would merely divert attention and processing resources away from the integration and computation of a correct linguistic message. Because gender information is mandatory for syntactic purposes alone one could reasonably assume, therefore, that efficient access to gender information is contingent on the engagement of syntactic processing. By this view, it would not be surprising that the gender-decision task produced so many errors in Hebrew even under conditions in which all transparent nouns were eliminated.

In Experiment 3 this hypothesis was examined by introducing a task that invokes

syntactic processing, and hence may more accurately reflect the way gender is accessed during the course of normal language processing.

CHAPTER IV - *Noun-Adjective Grammaticality Judgments*

In the next experiment access to gender was evaluated in a task that does not require an explicit gender response. Each of the nouns from Experiment 2 was presented in Experiment 3 as part of a noun-adjective pair. Half the pairs were appropriately gender marked, and the other half violated gender agreement. The task was to assess the acceptability of the pairs (e.g., if a masculine noun was modified by a masculine adjective the correct response was "yes", but if it was modified by a feminine adjective then the correct response was "no"). The aim of this manipulation, was to present the subjects with a task that more closely approximates the conditions under which gender information is normally accessed in natural linguistic processing. Correct decisions in this task require access to the grammatical gender of each noun, however, unlike the gender decision task they do not require a specific gender response. Because the detection of agreement normally plays a role in grammatical encoding, it was assumed that the computation needed to make a grammaticality judgment is a relatively better practiced process, and also that it was more likely to reflect normal mechanisms of access to gender.

Thus, Experiment 3 attempted to evaluate the role of regularity in access to grammatical gender during the course of language comprehension, but under different task conditions. It was hypothesized that the inclusion of a syntactic context would invoke the syntactic processor, thus allowing efficient access to the lemma level where gender is represented. This in turn would allow examination of transparency in the computation of gender agreement, with the goal of evaluating the presence or absence of

an interaction between lexemes and lemmas in language comprehension. A manipulation that adequately reflects the normal mechanisms of syntactic processing would be expected to produce a reduction in the percentage of errors in responses to irregular (unmarked feminine) targets. Such an outcome would demonstrate that the grammaticality judgment task reflects the mechanisms of access to gender more accurately than the gender-decision task.

If a reduction of errors in responses to irregular nouns were established, then it would be possible to return to the question of interactive versus modular models. In the presence of comparable error rates to regular and irregular nouns, the interactive model predicts longer reaction times for grammaticality judgments involving irregular nouns. Conversely, the modular model predicts that access to the lemma level is not influenced by lexeme level regularity, and hence reaction times for grammaticality judgments involving irregular nouns should not be longer than those involving regular nouns. Thus, in order to evaluate the modular vs. interactive models, it is first necessary to find a task that reasonably approximates the computation of gender agreement in normal language processing.

Finally, because no marked nouns were included in this experiment it seems unlikely that nonlexical strategic processing could influence the results. However, the results of Experiment 2 suggested that even in the absence of any explicitly marked targets strategic processing may play a role in response computation if subjects operate under the default assumption that unmarked nouns are masculine. To the extent that strategic processing does contribute to response computation this interferes with the

possibility of using the data to evaluate the normal mechanisms of access to gender, and to distinguish between model types. The results, must therefore, be interpreted with this caution in mind.

Method

Participants: Forty-two native speakers of Hebrew who were undergraduates at the Hebrew University in Jerusalem participated in the experiment for course credit or for payment. None had participated in Experiments 1 or 2.

Stimuli and design: In order to exclude the possibility that access to meaning, or episodic memory could influence the results each noun was assigned an adjective in a random fashion. This produced 48 noun-adjective pairs, half of which included a masculine noun, and half an unmarked feminine noun. All but two of the nouns were taken from Experiment 2 (two irregular feminine nouns which end in ***-at*** were replaced because of their similarity to the regular feminine endings ***-et*** and ***-it***). The random assignment of adjectives produced a vast majority of semantically implausible pairs (e.g., *hungry faucet, sad power,* and *curious city*). Two counter-balanced lists were constructed such that once a noun had been assigned its adjective it was always presented with this same adjective; however the adjective was appropriately gender marked in one list and inappropriately marked in the other list. The lists were introduced as a between subjects variable, thus no participant saw the same noun or adjective more than once.

Procedure: Each subject was tested individually, in a soundproof room. Verbal and written instructions indicated that noun-adjective pairs would appear one after

another on the computer screen, that s/he was to decide if the phrase was grammatical or not, and to respond by pressing the appropriate key ("yes" on the right, and "no" on the left). The instructions also indicated that most pairs would be semantically implausible, and stressed that both accuracy and speed were important. Before each word pair a warning stimulus consisting of nine pound signs (i.e., "#########") was presented for 500 ms. The noun-adjective pairs appeared centered on the screen of an IBM 386 computer. A millisecond clock that recorded RTs was time-locked to the presentation of each noun-adjective pair. Each pair remained on the screen either until a response was recorded, or until 3 seconds had elapsed. There was a 1.5 second pause between each trail. The experiment began with 16 practice items that included feedback ("correct", "incorrect", or "too much time" when no response was recorded before 3 seconds), and subsequently proceeded with the 48 test items that did not include any feedback. The pairs were presented in a different pseudo-random order for each subject.

Results.

Means and standard deviations for correct responses (RTs) were calculated for items in each noun type. Within each subject/condition, all RTs more than two standard deviations above or below the mean for each subject were replaced with the appropriate cutoff value, and the mean was recalculated. Outliers accounted for less than 5% of all responses. One subject who made close to 38% errors averaged over conditions (i.e., nearly random responses) was replaced. The data for grammatical and ungrammatical pairs were analyzed separately in ANOVAs that included two factors: Item groups, and

Gender (masculine versus unmarked feminine). The first factor was a nonrepeated factor that was introduced by the counterbalancing procedure, and simply extracts any variance due to this procedure.

Grammatical pairs ("YES" decisions) produced no main effect of gender $F(1,44) < 1.0$ for RTs, and $F(1,44) = 2.77, MSE = 90\ p = .10$ errors. In contrast, the analysis of ungrammatical pairs ("NO" decisions), produced significant differences between masculine and feminine nouns pairs in both the RT analysis $F(1,44) = 4.77, MSE = 19,479\ p < .05$, and the errors analysis $F(1,44) = 26.70, MSE = 113\ p < .001$. The average RT and error rate for each item type are presented in Table 4.

Discussion

The results of Experiment 3 deviate considerably from the results of the previous experiments. First, they demonstrated that decisions about grammatical pairs ("yes" decisions) were not sensitive to gender-marking regularity at the form level. In contrast, decisions about ungrammatical pairs ("no" decisions) resembled the results of previous experiments by showing a strong effect of regularity. The results for well-formed pairs are particularly impressive because of the dramatic reduction in error rates for unmarked (and therefore irregular) feminine nouns. The average error rate for these items in grammatical pairs was only 10.5% in the present experiment, whereas these same nouns produced a 24.8% error rate in Experiment 2, and a 37.9% error rate in Experiment 1.

This pattern of results suggests that access to grammatical gender is different when an explicit gender response is not required, and in the presence of a syntactic

context. Moreover, the data suggest that lexeme level regularity does not affect access to grammatical gender in the analysis of well-formed inputs. The absence of any significant differences between regular and irregular nouns also suggests that strategic mechanisms were eliminated (or at least minimized enough to prevent significant effects) in this experiment. Note that because a comprehension task necessarily begins with the analysis of word-forms the effect of gender marking regularity should be most conspicuous in tasks involving comprehension rather than production. This pattern of results, therefore, is perhaps the strongest form of evidence in favor of a modular characterization of access to the lemma level.

The most obvious difficulty with this conclusion, however, concerns the robust differences found between masculine and unmarked feminine nouns in responses to ungrammatical pairs ("no" decisions). In this condition the results of the present experiment highly resemble the pattern observed in Experiments 1 and 2; unmarked feminine nouns paired with an unmarked masculine adjective produced larger error rates, and took much longer to reject relative to unmarked masculine nouns paired with marked feminine adjectives. Thus, what needs to be explained is why the analysis of grammatical pairs seems to have been allowed access to the lemma level without influence from the lexeme level, whereas the analysis of ungrammatical pairs was influenced by lexeme level characteristics. One possible explanation for the difference between well- and ill-formed pairs involves a post-lexical checking mechanism that reexamines ungrammatical pairs before allowing their rejection. This question will be addressed in the General Discussion.

In the next experiment, an additional manipulation was introduced to reinforce the conclusion that syntactic context allows efficient access to the lemma level where there is a form-free representation of grammatical gender.

CHAPTER V - *Noun-Verb Grammaticality Judgments*

In this experiment the same nouns used in Experiments 2 and 3 were once again paired with correctly marked and incorrectly marked modifiers. This time, however, the randomly assigned adjectives were replaced with randomly assigned verbs. Presumably, the same syntactic processor that detects noun-adjective agreement also detects noun-verb agreement. Thus, if syntactic context allows access to the lemma representation, then the same pattern of results should replicate despite the switch from adjectival to verbal phrases. As in the previous experiment, correct decisions in this task require access to the grammatical gender of each noun, however, they do not necessarily require explicit or conscious focus on gender.

Method

Participants: Sixty native speakers of Hebrew who were undergraduates at the Hebrew University in Jerusalem participated in the experiment for course credit or for payment. None had participated in Experiments 1-3.

Stimuli and design: In order to exclude the possibility that access to meaning, or episodic memory could influence the results each noun was assigned a verb in a random fashion. This produced 48 noun-verb pairs, half of which included a masculine noun, and half an unmarked feminine noun. All but one of the nouns were taken from Experiment 3 (one noun which takes an irregular plural ***-ayim*** instead of ***-im*** was replaced). The random assignment of verbs produced a vast majority of semantically

implausible pairs (e.g., *the faucet scolded, the power threw,* and *the city left*). Two counter-balanced lists were constructed such that once a noun had been assigned its verb it was always presented with this same verb, however the verb was appropriately gender marked in one list and inappropriately marked in the other list. The lists were introduced as a between subjects variable, thus no participant saw the same noun or verb more than once.

Procedure: Each subject was tested individually, in a soundproof room. Verbal and written instructions indicated that noun-verb pairs would appear one after another on the computer screen, that s/he was to decide if the phrase was grammatical or not, and to respond by pressing the appropriate key ("yes" on the right, and "no" on the left). The instructions also indicated that most pairs would be semantically implausible, and stressed that both accuracy and speed were important. Before each word pair a warning stimulus consisting of nine pound signs (i.e., "#########") was presented for 500 ms. The noun-verb pairs appeared centered on the screen of an IBM 386 computer. A millisecond clock that recorded RTs was time-locked to the presentation of each noun-verb pair. Each pair remained on the screen either until a response was recorded, or until 3 seconds had elapsed. There was a 1.5 second pause between each trail. The experiment began with 16 practice items that included feedback ("correct", "incorrect", or "too much time" when no response was recorded before 3 seconds), and subsequently proceeded with the 48 test items that did not include any feedback. The pairs were presented in a different pseudo-random order for each subject.

Results.

Means and standard deviations for correct responses (RTs) were calculated for items in each noun type. Within each subject/condition, all RTs more than two standard deviations above or below the mean for each subject were replaced with the appropriate cutoff value, and the mean was recalculated. Outliers accounted for less than 5% of all responses. As in Experiment 3 the data for grammatical and ungrammatical pairs were analyzed in ANOVAs that included two factors: Item groups, Gender (masculine versus unmarked feminine).

In contrast to the results obtained in Experiment 3, there was a difference between masculine and feminine noun pairs in the decisions on well-formed pairs. However, this effect was not in the expected direction; RTs to the irregular feminine nouns were significantly *faster* than responses to the regular masculine nouns $F(1,44) = 9.25$, $MSE = 12,718$ $p < .01$. This pattern is exactly opposite to that observed in Experiments 1 and 2. The analysis of errors replicated the pattern observed in Experiment 3; there was no difference between regular and irregular nouns $F(1,44) = 2.77$, $MSE = 90$ $p = .10$. Finally, as in Experiment 3, the analysis of ungrammatical pairs alone, produced significant differences in the expected direction between masculine and feminine nouns pairs in both the RT analysis $F(1,44) = 5.26$, $MSE = 14,409$ $p < .05$, and the errors analysis $F(1,44) = 42.28$, $MSE = 162$ $p < .001$. The average RT and error rate for each item type is presented in Table 4.

Discussion

The most important aspect of the results of Experiment 4 is that they demonstrate again that syntactic context produces reduced error rates in responses to irregular feminine nouns, thereby eliminating adverse effect of form regularity on accessing grammatical gender. This result strengthens the previous conclusion that syntactic context produces a difference in how gender is accessed. The pattern of results obtained in Experiment 4 also confirms other conclusions that were based on the results of Experiment 3. For example, access to lemma level gender seems to be easier when an explicit identification of grammatical gender is not required. More specifically, the results demonstrate no disadvantage in gender retrieval for irregular nouns in the analysis of grammatical inputs (in fact, in Experiment 4 irregular feminine nouns were actually processed more quickly). This outcome is incompatible with the interactive model (outlined above) of access to gender.

Nevertheless, the difference that emerged in both Experiments 3 and 4 between grammatical pairs ("yes" decisions) and ungrammatical pairs ("no" decisions) requires explanation. Decisions about ungrammatical pairs were affected by regularity in the expected direction based on the interactive view. That is, irregular (unmarked) feminine targets paired with masculine verbs were rejected more slowly and with more errors than regular masculine targets paired with (marked) feminine verbs. The difference between grammatical and ungrammatical decisions obtained in both Experiments 3 and 4 suggests that "no" decisions involve different computations than "yes" decisions. This difference may reflect the same processing mechanisms that arise when an error occurs in the

context of natural language processing. However, another possibility is that the difference in decision type (i.e., "yes" vs. "no") spuriously produced different results for grammatical and ungrammatical pairs.

In the next experiment another manipulation was introduced to examine responses to grammatical and ungrammatical pairs when both are placed in the context of the same decision type. An additional aim was to further test the hypothesis that syntactic context allows access to the form-independent representation of grammatical gender.

CHAPTER VI - *Two-word Lexical Decision*

This experiment was designed to evaluate both grammatical and ungrammatical pairs in the context of the same decision type. If the difference between these pairs in Experiments 3 and 4 were an artifact of the different task demands ("yes" vs. "no" decision types) then there should no longer be a difference between grammatical and ungrammatical pairs in this experiment. A second goal was to further test the hypothesis that syntactic analysis allows access to the lemma level. The same nouns used in Experiments 2-4, were once again paired with correctly marked and incorrectly marked modifiers. This time, however, a task that specifically discourages syntactic processing was introduced. Half of the words were paired with a nonword stimulus, and subjects were instructed to respond with "yes" if both items presented were real words, and "no" if one of the items was a nonword. In addition, half the word-word pairs were grammatical (e.g., a masculine noun modified by a masculine adjective), and half were not grammatical (e.g., a masculine noun modified by a feminine adjective). Thus, to effectively perform the task subjects would have to suppress the automatic tendency to reject ungrammatical pairs because, although they violate agreement, ungrammatical pairs are still comprised of two real words. Under these conditions then, suppression of syntactic processing would make the task easier, and because it was hypothesized that syntactic processing allows access to the form-independent representation of grammatical gender it was predicted that regularity effects should reemerge (for grammatical pairs) in this task.

Method

Participants: Fifty-four native speakers of Hebrew who were undergraduates at the Hebrew University in Jerusalem participated in the experiment for course credit or for payment. None had participated in Experiments 1-4.

Stimuli and design: The nouns and adjectives selected were those used in Experiments 3-4, the assignment of adjectives was the same as that used in Experiment 3 (noun-adjective agreement). Each noun was either paired with a correctly marked adjective, an incorrectly marked adjective, or a nonword. Half the nonwords appeared in the noun's position (i.e., first), and half appeared in the adjective's position (i.e., second). Because all nouns were unmarked (i.e., regular masculine nouns, and irregular feminine nouns), the nonwords that appeared in the noun's position were also unmarked (i.e., no nonword ended with any of the feminine suffixes ***-ah***, ***-et***, or ***-it***). Half the adjectives were marked feminine, and half were marked masculine in both the word-word, and the word-nonwords pairs. Obviously, item pairs that included a nonword, were neither grammatical nor ungrammatical; however, because all the items were unmarked, and unmarked nouns are masculine by default, nonwords followed by masculine adjectives may have appeared more "grammatical-like", than nonwords followed by a feminine adjective.

There were 32 word-word pairs, and 32 word-nonword pairs. For the word-word pairs half included a masculine noun, and half an unmarked feminine noun. Half of the word-nonword pairs had a nonword in the noun location (nouns appear before the

modifier in Hebrew), and half had a nonword in the adjective location (after the noun). Three counter-balanced lists were constructed such that once a noun had been assigned its adjective it was always presented with this same adjective. However, each noun appeared as part of a grammatical word-word pair, as part of an ungrammatical word-word pair, and also as part of a noun-nonword pair. The lists were introduced as a between subjects variable, thus no participant saw the same noun or adjective more than once.

Procedure: Each subject was tested individually, in a soundproof room. Verbal and written instructions indicated that word pairs would appear one after another on the computer screen, that s/he was to decide if both items were real words, and to respond by pressing the appropriate key ("yes" on the right, and "no" on the left). The instructions also indicated that most pairs would be semantically implausible, and stressed that both accuracy and speed were important. Before each item pair a warning stimulus consisting of nine pound signs (i.e., "##########") was presented for 500 ms. The item pairs appeared centered on the screen of an IBM 386 computer. A millisecond clock that recorded RTs was time-locked to the presentation of each pair. Each pair remained on the screen either until a response was recorded, or until 3 seconds had elapsed. There was a 1.5 second pause between each trial. The experiment began with 16 practice items that included feedback ("correct", "incorrect", or "too much time" when no response was recorded before 3 seconds), and subsequently proceeded with the 64 test items that did not include any feedback. The pairs were presented in a different pseudo-random order for each subject.

Results.

Means and standard deviations for correct responses (RTs) were calculated for items in each response type ("yes" and "no" decisions). Because responses to both grammatical and ungrammatical pairs in this experiment were "yes" responses, these were combined in a single ANOVA. Within each subject/condition, all RTs more than two standard deviations above or below the mean for each subject were replaced with the appropriate cutoff value, and the mean was recalculated. Outliers accounted for less than 5% of all responses. The data were analyzed in ANOVAs that included three factors: Item groups, Grammaticality (well-formed vs. ill-formed pairs), and Gender (masculine versus irregular feminine). For the comparison of nonwords in first position (the noun's place) with nonwords in the second position (the adjective's place) subject means were computed instead of item means (as reported for every other analysis throughout this paper). This is because the position of the items is determined by the language; the noun is always first, and the adjective second. Hence, counterbalancing across items was not possible for this comparison.

There was a no main effect of grammaticality in RTs $F(1,42) = 2.26$, $MSE = 4,989$ $p > .05$, or errors $F(1,42) = 1.38$, $MSE = 93$ $p > .05$, and there was no main effect of gender in RTs $F(1,42) = 1.09$, $MSE = 3,338$ $p > .05$, or errors $F(1,42) = 1.88$, $MSE = 116$ $p > .05$. However, there was a significant interaction between gender and grammaticality. Responses to grammatical word-word pairs showed a trend towards an effect of gender in the expected direction (RTs to regular masculine noun-adjective were

19 ms faster), whereas responses to ungrammatical word-word pairs produced an effect in the opposite direction (RTs to irregular feminine noun-adjective pairs were 44 ms faster). This interaction was significant in the analysis of RTs $F(1,42) = 7.07$, $MSE = 3,338$ $p = .01$, but not in the analysis of errors ($F < 1$). Separate analyses of grammatical and ungrammatical pairs confirmed this pattern. For well-formed pairs there was no effect of gender; $F(1,42) = 1.14$, $MSE = 3,832$ $p > .05$ in the analysis of RTs, and errors ($F < 1$). Whereas for ill-formed pairs there was a significant effect of gender; irregular feminine nouns produced faster responses in the analysis of RTs $F(1,42) = 5.08$, $MSE = 4,496$ $p < .05$. This difference was not significant in the errors analysis ($F < 1$).

In the analysis of word-nonword pairs there was a main effect of first vs. second position; when the nonword was in the noun's position (the first item) subjects were faster to reject the pair than when the nonword was in the adjective's position (the second item). This effect was significant in both the analysis of RTs $F(1,51) = 63.17$, $MSE = 2,868$ $p < .001$, and in the analysis of errors $F(1,51) = 39.11$, $MSE = 149$ $p < .001$. The average RT and error rate for each item type is presented in Table 6.

Discussion

The results of Experiment 5 are similar to those in Experiments 3 and 4 in that a robust difference was obtained between grammatical and ungrammatical pairs; responses to the latter showed strong differences between regular and irregular nouns, whereas the former did not. At the same time, however, the effect of regularity in responses to ungrammatical pairs was in the opposite direction of that observed in Experiments 3 and

4. The two-word lexical decision task requires subjects to suppress the automatic tendency to reject ungrammatical pairs. It was predicted that under these conditions syntactic processing would be suppressed. Thus, the fact that any significant differences between grammatical and ungrammatical pairs were obtained suggested that it was not possible to suppress syntactic analysis entirely. Moreover, because both grammatical and ungrammatical pairs called for a "yes" response, it suggests that the differences observed between these pair types in Experiments 3 and 4 were not merely produced by the difference in decision type.

Further, it is interesting to note that although there was a difference between grammatical and ungrammatical pairs, this time the gender effect was in the unexpected direction. RTs to irregular feminine nouns were significantly faster in the context of an ungrammatical pair (a similar pattern was found in Experiment 4 but in that case it was in responses to grammatical pairs). This indicates that it was easier for subjects to ignore grammaticality violations when an irregular noun was involved i.e., *nefesh(if) zariz(m)* was less obviously ungrammatical than *maslul(m) tmim<u>a</u>(f)*. This result is perhaps not surprising because irregular feminine nouns essentially "look" masculine. Moreover, it seems consistent with the results of Experiments 3 and 4, in which it was shown that it was easier to detect grammaticality violation for regular nouns, and harder to detect grammaticality violation for irregular nouns. In other words, if it is harder to detect a grammaticality violation for an irregular noun, then it should not be surprising that it is also easier to ignore this same violation when the task demands change (i.e., two-word lexical decision instead of grammaticality judgment).

Also notable was the overall rate of errors made in responses to irregular feminine nouns in this experiment relative to Experiments 3 and 4. In the latter, these nouns produced many more errors relative to regular masculine nouns, whereas in this experiment there were virtually no such effects in the error data (for both grammatical and ungrammatical pairs). This result suggests that although it is more difficult to detect a grammaticality violation with irregular feminine nouns, it is not more difficult to identify that an irregular noun is a word.

Two major conclusions can be drawn on the basis of the results of Experiment 5: 1) syntactic processing cannot be suppressed, 2) grammatical and ungrammatical pairs are processed differently (even when they appear in a task that places them in the context of the same decision type). An account of the difference between grammatical and ungrammatical pairs will be described in the General Discussion. Before returning to this question, however, another test of the modularity of the access mechanisms to lemma and lexeme representations was introduced. In Hebrew the only nouns that are explicitly marked in the singular form are the regular feminine nouns. These nouns were purposefully avoided in Experiments 2-5 to prevent the use of nonlexical strategies. It could be argued, therefore, that no effects of regularity were obtained (in responses to grammatical pairs) because the gender markings were not explicit on all phrasal elements. In other words, lexeme level activation may not to rise to the point where it can either facilitate or interfere with access to the lemma level without both extensive and explicit regularity marking. This factor would have minimized the chance of detecting an effect of lexeme-level characteristics on access to the lemma. The next experiment was

an attempt to maximize the possibility of detecting such an effect by looking at grammaticality judgments for noun-adjective pairs in the plural form. These provide an even stronger test of the modularity of lemma and lexeme access mechanisms because in Hebrew both masculine and feminine words are explicitly marked in the plural form.

CHAPTER VII - *Plurals Grammaticality Judgments*

In this experiment access to gender was tested using nouns in the plural form modified by adjectives in the plural form. Because both masculine and feminine nouns are marked explicitly in Hebrew when in the plural form, this manipulation allows an even stronger test of the modularity of lemma and lexeme access mechanisms. In Experiments 3 and 4, the only the items that were explicitly gender marked were the adjectives that marked the irregular feminine nouns. All other items were unmarked; masculine nouns and adjectives in the singular form are not explicitly gender-marked in Hebrew, and all the feminine nouns in those experiments were irregular (i.e., unmarked). If access to the lemma level is not influenced by form, then this should hold even if all elements of the phrase are explicitly gender-marked. Thus, the modular view predicts that there should be no difference between responses to nouns that take a consistent plural suffix (e.g., a masculine noun that takes a masculine plural suffix) and responses to nouns that take an inconsistent plural (e.g., a masculine noun that takes a feminine plural suffix). The interactive view, on the other hand predicts that there should be even stronger effects of regularity in this experiment relative to those obtained in Experiment 3 because the explicit markings should both facilitate access (for consistent pairs) and inhibit access (for inconsistent pairs).

Both these predictions, however, rely on the assumption that strategic processing can be ruled out as a major factor in response computation. It is important to stress that in this experiment a form-based strategy is clearly available for the first time since

Experiment 1; this time the strategy would involve suffix matching between nouns and adjectives (e.g., noun + *im* and adjective + *im* = grammatical, and noun + *im* and adjective + *ot* = ungrammatical). Thus, explicit gender marking on all constituents of the phrase both maximize the possibility of detecting lemma-lexeme interaction, and also allow for the possibility of nonlexical computation of responses. The latter result effectively destroys the possibility of using the data to examine the mechanisms of access to gender. Because half of the nouns were regular and half irregular the nonlexical strategy would produce 50% errors; thus one way to measure the contribution of nonlexical strategic processing is to look at the overall rate of errors. The inability to suppress nonlexical strategies should produce higher error rates especially for inconsistent plural pairs.

Method

Participants: Fifty-two native speakers of Hebrew who were undergraduates at the Hebrew University in Jerusalem participated in the experiment for course credit or for payment. None had participated in Experiments 1-5.

Stimuli and design: For each type of noun (masculine, feminine regular, and feminine irregular) 10 nouns that take a consistent plural and 10 nouns that take an inconsistent plural were selected. Each noun was then assigned an adjective in a random fashion producing 60 noun-adjective pairs. To balance the number of masculine and feminine nouns, an additional 20 masculine nouns (10 with a consistent plural, and 10 with an inconsistent plural) were added as filler items. Most of the nouns were taken from

Experiment 3, however a number of items were discarded or replaced because some items do not have a plural form in Hebrew (e.g., *esh* the word for *fire*), and to avoid ambiguous items. For example, nouns like the irregular feminine noun, *ez* (the word for *goat*) becomes *ezim* in the plural form (the plural is inconsistent), however, because vowels are generally omitted from the Hebrew orthography *ezim* may also be read as *azim* (which means *strong* or *bold*). Two counter-balanced lists were constructed such that once a noun had been assigned its adjective it was always presented with this same adjective, however the adjective was appropriately gender marked in one list and inappropriately marked in the other list. The lists were introduced as a between subjects variable, thus no participant saw the same noun or adjective more than once.

Procedure: Each subject was tested individually, in a soundproof room. Verbal and written instructions indicated that noun-adjective pairs would appear in the plural form, one after another on the computer screen and that s/he was to indicate if the phrase was grammatical or not by pressing the appropriate key ("yes" on the right, and "no" on the left). The instructions also indicated that most pairs would be semantically implausible, and stressed that both accuracy and speed were important. Before each word pair a warning stimulus consisting of nine pound signs (i.e., "#########") was presented for 500 ms. The noun-adjective pairs appeared centered on the screen of an IBM 386 computer. A millisecond clock that recorded RTs was time-locked to the presentation of each noun-adjective pair. Each pair remained on the screen either until a response was recorded, or until 3 seconds had elapsed. There was a 1.5 second pause between each trial. The experiment began with 16 practice items that included feedback

("correct", "incorrect", or "too much time" when no response was recorded before 3 seconds), and subsequently proceeded with the 60 test items, and 20 filler items without feedback. The pairs were presented in a different pseudo-random order for each subject.

Results.

Means and standard deviations for correct responses (RTs) were calculated for items in each noun type (i.e., consistent plural and inconsistent plural). Within each subject/condition, all RTs more than two standard deviations above or below the mean for each subject were replaced with the appropriate cutoff value, and the mean was recalculated. Outliers accounted for less than 5% of all responses. The data for grammatical and ungrammatical pairs were analyzed separately in ANOVAs that included two factors: Item groups, and Plural Regularity (consistent plural versus inconsistent plural). The first factor was a nonrepeated factor that was introduced by the counterbalancing procedure, and simply extracts any variance due to this procedure. The analysis of Consistent and Inconsistent Plural groups collapsed across gender in the singular form thus including 5 masculine, 5 regular feminine, and 5 irregular feminine nouns in each condition (for both grammatical and ungrammatical pairs).

In the analysis of grammatical pairs ("YES" decisions) RTs to consistent plural pairs were 67 ms faster than RTs to inconsistent pairs. However, this effect was not significant in either the analysis of RTs $F(1,56) = 2.28, MSE = 30,109\ p > .05$, or in the analysis of analysis of errors $F(1,56) = 1.32, MSE = 142\ p > .05$. In the analysis of ungrammatical pairs ("NO" decisions) there was also no significant difference between

consistent plural pairs and inconsistent plural pairs in the analysis of RTs ($F<1$), however there was a trend in the expected direction (responses to consistent plurals were 34 ms faster), and there was a significant effect in the same direction in the errors analysis $F(1,56) = 14.36, MSE = 221\ p < .001$.

Discussion

This experiment was designed to be a stronger test of the modularity of access mechanisms to lemma and lexeme representations. To that end, the results are fairly straightforward. Despite the inclusion of an explicit gender marker on every phrasal constituent the only significant difference between consistent and inconsistent plural pairs was in the errors analysis of ungrammatical pairs suggesting that the explicit gender markings have little to no effect on grammaticality judgments. That said, however, a number of aspects of the results suggest that this simple interpretation does not fully capture the nature of the effects. In particular, in this experiment 42% (22/53) of subjects produced error rates above 20%, whereas in Experiment 3 only 16% (8/49) produced error rates above 20%. Moreover, unlike the data in Experiments 3 and 4 in this experiment there were no significant differences between regular and irregular nouns in the RT data for ungrammatical pairs. One possible explanation for this latter difference is that the overall difference in error rates in this experiment was higher (an average of 23% vs. 13% in Experiment 3). This may have reduced the possibility of detecting an effect by increasing the variability of the RT data (which is taken from correct responses only).

To evaluate this question a separate analysis of 28 subjects with low error rates

(below 20% averaged over conditions) was computed. This re-analysis produced the same pattern found in Experiment 3; the data for "yes" decisions produced no differences between consistent and inconsistent pairs in both the analysis of RTs $F(1,56) = 1.08$, $MSE = 34,371$ $p > .05$ (there was, however, a trend the expected direction 50 ms), and also no significant difference was found in the analysis of errors $F(1,56) = 2.57$, $MSE = 100$ $p > .05$. In contrast, the data for "no" decisions produced significant differences between consistent and inconsistent pairs in both the RTs analysis $F(1,56) = 5.81$, $MSE = 24,831$ $p < .05$, and the errors analysis $F(1,56) = 6.76$, $MSE = 171$ $p < .05$. The average RT and error rate for each item type, and for both groups (all subjects, and low-error subjects) is presented in Table 7. The emergence of the same pattern in this reanalysis increases the plausibility of the conclusion that the same mechanisms were involved in plurals and singular grammaticality judgments. Note that, in Experiment 4 (singular noun-verb grammaticality judgments) a similarly high percentage of errors was observed (an average of 20%), and thus, it seems that the introduction of plural markers may not be the only reason why a higher percentage of errors might be obtained. Nevertheless, the plurals agreement experiment is most comparable to Experiment 3 because adjective modifiers were used in both (in fact they were the same adjectives), and thus the differences in overall error rates across the two experiments still requires explanation.

One possibility is that the plurals agreement task was more difficult because of greater levels of interaction between lemma and lexeme representations. By this view, the overall difference between regular and irregular noun pairs should have been larger than it was in Experiment 3 for two reasons. First, because access to consistent pairs

should have been facilitated by markings at the level of form, and second, because access to inconsistent pairs should have been inhibited by markings at the level of form. Another possibility is that nonlexical strategic processing played a larger role in response computation in this experiment. This too would produce greater differences between regular and irregular pairs than had been obtained in Experiment 3. Some support for these hypotheses can be found by comparing the results in Experiment 3 with those obtained in this experiment (Tables 4 and 7). To support either of theses explanations the difference between regular and irregular pairs should be larger in Table 7 than they are in Table 4. For ungrammatical pairs, there was a larger difference in the expected direction in the reaction time data for grammatical pairs (22 ms vs. 67 ms). However, the differences in the error data were in the opposite direction for both grammatical pairs (15.9% vs. 14.5%), and ungrammatical pairs (4.5% vs. 3.5%), and also the reaction time data for grammatical pairs were in the wrong direction (88 ms vs. 34 ms). Thus, these explanations seem less plausible.

An alternative account is that responses in this experiment may actually have required two computations of agreement. This would produce higher error rates overall (for both regular and irregular nouns, as well as both grammatical and ungrammatical pairs). By this view, two agreement computations would be required to analyze each plural noun-adjective pair. First, the plural suffix attached to the noun would be checked for a match with the noun's gender, and then the plural suffix attached to the adjective would also be checked for a match with the noun's gender. This argument must assume that nouns in the singular form, and nouns in the plural form are not stored in separate

representations because in this case no "agreement" processing would be required to identify a noun in the plural form (all that would be required is the location of the appropriate lexical entry). Although noun-suffix agreement was never violated in the context of this experiment (only noun-adjective agreement was violated), and hence there was no need to compute such agreement, this factor might still have affected the results, particularly if syntactic analysis is automatic and cannot be suppressed (evidence in favor of the latter argument was obtained in Experiment 5). The final experiment was designed to test the nature of simple noun+suffix agreement detection, to determine whether this factor might have contributed to the overall difficulty of the plurals grammaticality judgment task.

CHAPTER VIII - *Plurals Lexical Decision*

This experiment was designed to evaluate the nature of noun-plural suffix agreement. If nouns in the plural form are stored in a separate representation from nouns in the singular form then no agreement computation would be required to process a noun marked with the correct or incorrect plural suffix. Hence, nouns that take a consistent plural suffix would be recognized as easily as nouns that take an inconsistent plural suffix. On the other hand, if nouns are stored in the singular form only, and the addition of a plural suffix requires some agreement computation (e.g., via rule), then nouns that take an inconsistent suffix should take longer to recognize. This prediction assumes that generation of a plural form in Hebrew, is somewhat analogous to the generation of regular and irregular past tense forms in English (Pinker, 1991). According to this theory the computation of a plural form would first begin for all nouns (regular and irregular) with a rule-based mechanism (e.g., if the noun ends in the suffix ***-ah*** then the suffix ***-ot*** should be substituted to create the plural form). In contrast, nouns that take an inconsistent plural would have to contain a flag indicating that the noun is irregular and that an inconsistent suffix should be used instead. Thus, the rule-based computation would be interrupted midway, and the inconsistent plural suffix would be attached to the noun instead (of the consistent suffix). Because of the added processing step, nouns with inconsistent plurals would then require additional processing time. The advantage to such a rule based mechanism is that it is less redundant, and perhaps also more efficient given that most nouns take a consistent plural suffix.

In this experiment half the nouns were combined with a correct plural suffix, and

half were combined with an incorrect plural suffix, thus creating a nonword. The task was lexical decision; subjects were asked to respond "yes" when they saw a correct noun + plural suffix combination, and "no" when they saw an incorrect noun + plural suffix combination. If nouns with inconsistent plural suffixes produced longer response times and error rates, then this would suggest that the generation of a plural form in Hebrew requires an additional agreement processing step. On the other hand, if nouns with inconsistent plural suffixes did not produce longer RTs and more errors, this would suggest that nouns in the plural form are stored in a separate representation from nouns in the singular form, and that no syntactic analysis is required to process plural nouns.

In addition to the inanimate nouns, animate nouns were added to the design to serve as a baseline. Animate nouns can take either masculine or feminine form and thus it was hypothesized that the computation of gender for these nouns might be qualitatively different from inanimate nouns. To balance the number of "yes" and "no" responses plural nonwords were also included in the design; these were created by attaching a plural suffix to a nonexistent but phonologically legal Hebrew root. It was hypothesized that these items might also produce qualitatively different responses because in contrast with all other items in this experiment non-root nonwords have no representation in the mental lexicon (and obviously do not take either masculine or feminine plural suffix).

Method

Participants: Forty-four native speakers of Hebrew who were undergraduates at the Hebrew University in Jerusalem participated in the experiment for course credit or for

payment. None had participated in Experiments 1-6.

Stimuli and design: For each type of noun (masculine, feminine regular, and feminine irregular) nouns that take a consistent plural and nouns that take an inconsistent plural were selected. Half the nouns of each type were paired with the correct plural suffix, and half were paired with the incorrect plural suffix thus creating a nonword. Because many nouns paired with an incorrect plural suffix actually create a different word entirely, a number of items that had been used in previous experiments could not be used in this experiment. For example, the plural of *shana* (the word for *year*) is *shanim*. When the incorrect plural is attached to this root it produces the nonword *shanot*, however, this nonword can also be read as the word *shanut* (which means *variance*). Further, because the number of nouns that take an inconsistent plural is severely limited there was an uneven number of items in each group. There were 12 masculine nouns that take a consistent plural suffix; 12 masculine nouns that take an inconsistent plural suffix; 12 regular feminine nouns that take a consistent plural suffix; only 9 regular feminine nouns that take an inconsistent plural suffix; 12 irregular feminine nouns that take a consistent plural suffix; and only 8 irregular feminine nouns that take an inconsistent plural suffix. Missing values for each condition were replaced with the mean values in each condition.

Twelve animate nouns that were frequency matched to the inanimate nouns were also included in the design. Nouns that are obviously associated with one gender semantically (e.g., *pilots* is a word typically associated with the masculine gender) were either avoided entirely, or were balanced by the inclusion of a noun that is typically

associated with the opposite gender (e.g., *nurses* is a word typically associated with the feminine gender). To balance the number of "yes" and "no" responses 12 nonword roots were also included; half were combined with the masculine plural suffix and the other half were combined with the feminine plural suffix. Two counter-balanced lists were constructed such that any noun (or nonword root) that was paired with the feminine plural suffix in one list, was paired with the masculine plural suffix in the other list. The lists were introduced as a between subjects variable, thus no participant saw the same item more than once.

Procedure: Each subject was tested individually, in a soundproof room. Verbal and written instructions indicated that nouns would appear in the plural form, one after another on the computer screen and that s/he was to indicate if the item was a real word or not by pressing the appropriate key ("yes" on the right, and "no" on the left). The instructions also stressed that both accuracy and speed were important. Before each item a warning stimulus consisting of nine pound signs (i.e., "#########") was presented for 500 ms. The items appeared centered on the screen of an IBM 386 computer. A millisecond clock that recorded RTs was time-locked to the presentation of each item. Each item remained on the screen either until a response was recorded, or until 3 seconds had elapsed. There was a 1.5 second pause between each trial. The experiment began with 16 practice items that included feedback ("correct", "incorrect", or "too much time" when no response was recorded before 3 seconds), and subsequently proceeded with the 89 test items without feedback. The items were presented in a different pseudo-random order for each subject. After completing the experiment each subject was handed a list of

all the nouns presented in the experiment combined with the feminine plural suffix in one column and the masculine plural suffix in the other column. For each noun subjects were asked to choose the correct form. This part of the experiment was untimed.

Results.

Means and standard deviations for correct responses (RTs) were calculated for items in each noun type. Within each subject/condition, all RTs more than two standard deviations above or below the mean for each subject were replaced with the appropriate cutoff value, and the mean was recalculated. Outliers accounted for less than 5% of all responses. Two items were eliminated from the analysis (one from the irregular feminine inconsistent plural condition, and one from the masculine inconsistent plural condition) because they generated more than 50% errors in the untimed written questionnaire indicating that subjects did not know the correct form even in an untimed test. The data were analyzed in an ANOVA that included four factors: Item groups, Lexicality (word vs. nonwords), Gender (masculine, feminine, irregular feminine), and Plural Regularity (consistent versus inconsistent). The first factor was a nonrepeated factor that was introduced by the counterbalancing procedure, and simply extracts any variance due to this procedure.

There was a main effect of lexicality; responses to nouns combined with the correct plural suffix were (237 ms) faster and had fewer errors (19.8%) than responses to nouns paired with an incorrect suffix. This effect was significant in both the analysis of RTs $F(1,11) = 336.17\ MSE = 6,591\ p < .001$, and in the analysis of errors $F(1,11) =$

521.33, MSE = 27 p < .001. There was no main effect of gender in the analysis of RTs *F(2,22) = 2.07, MSE = 6,700 p > .05*, however in the analysis of errors there were fewer errors to masculine nouns *F (2,22) = 4.30, MSE = 70 p < .05*. Similarly, there was no main effect of Plural Regularity in the analysis of RTs *F (1,11) = 3.47, MSE = 7,200 p > .05*, however there was a significant effect in the errors analysis; there were fewer errors in responses to nouns that call for a consistent plural suffix relative to nouns that call for an inconsistent plural suffix *F (1,11) = 12.01, MSE = 133 p < .01.*

Most interestingly, in addition to these main effects a number of interactions emerged, and the main effects are perhaps more meaningful when interpreted as part of these interactions. There was an interaction between Lexicality and Gender in both the analysis of RTs *F(2,22) = 4.11, MSE = 5,000 p < .05*, and in the analysis of errors *F(2,22) = 4.03, MSE = 79 p < .05*; differences between masculine, feminine, and irregular feminine nouns were only apparent in the nonword data (i.e., the "no" responses). There was also a significant interaction between Lexicality and Plural Regularity in both the analysis of RTs *F(1,11) = 16.68, MSE = 4,800 p < .01* and the analysis of errors *F(1,11) = 20.31, MSE = 140 p <.001*; differences between consistent and inconsistent plurals were only apparent in the responses to nonwords. Finally, there was an interaction between Gender and Plural Regularity in both the analysis of RTs *F(2,22) = 5.49, MSE = 13,000 p < .05*, and in the analysis of errors *F(2,22) = 14.58, MSE = 70 p < .001*; for regular nouns (masculine nouns and marked feminine nouns) there was a very strong effect of plural regularity in the expected direction, whereas for irregular feminine nouns (unmarked feminine nouns) there was a nonsignificant trend in

the opposite direction in both the RT and error data.

Nouns in the Animate condition, and items in the Non-root Nonword condition (created by attaching a plural suffix to a nonword-base) do not have an intrinsic gender, and also do not have a consistent or inconsistent plural suffix, and therefore these items were analyzed in a separate ANOVA. For Animate nouns there was a significant main effect of gender in the analysis of RTs $F(1,10) = 5.71$, $MSE = 6,321$ $p < .05$, but not in the analysis of errors $F(1,10) = 2.00$, $MSE = 110$ $p > .05$; RTs to masculine nouns were significantly faster than RTs to feminine nouns. There was a trend in the same direction for the non-root nonwords, however, it was not significant in either the analysis of RTs $F(1,10) = 3.20$, $MSE = 7,522$ $p > .05$, or in the analysis of errors $F(1,10) = 3.83$, $MSE = 14$ $p = .08$. The average RT and error rate for each item type in each condition is presented in Table 8.

Discussion

The most striking aspect of the results of Experiment 7 is the resemblance that they bear to the results of other the agreement experiments (noun-adjective, noun-verb, and plurals noun-adjective agreement; i.e., Experiments 3, 4, and 6). There was a robust difference between the effects obtained in "yes" decisions relative to the effects obtained for "no" decisions; the latter showed strong effects of regularity whereas the former did not. To some extent, the hypotheses outlined in the introduction to this experiment are difficult to evaluate on the basis of these data because no distinction between "yes" and "no" decisions were predicted. It was hypothesized that no effects of plural consistency

(for both legal and illegal noun + plural suffix items) should be obtained if nouns are stored as independent lexical representations in both plural and singular form. This view leaves the data for "no" decisions unexplained because for these items significant effects of consistency were obtained. Similarly, under the assumption that plural forms do require some agreement computation, the data for "yes" decisions could not be explained because for this decision type no difference was found between consistent and inconsistent plurals.

One reason why there might be a difference between "yes" and "no" responses in this experiment is that some words in Hebrew are orthographically distinguishable on the basis of the plural suffix alone. For example, the plural form of *beytsa* (the word for *egg*) is *beytsim* (note that *-im* here is an exception plural, thus although *beytsa* is a feminine noun the plural suffix is masculine). Although *beytsa* + *ot* produces the nonword *beytsot*, this nonword can also be read as the word *bitsot* (which means *swamps*). Thus, to determine whether a "no" response can be made on an item that seems to have the incorrect plural suffix, it is not enough to merely locate the lexical entry of the word. Instead, the subject must determine that there is no noun that calls for the given plural suffix. A "yes" decision on the other hand can be made as soon as a correct analysis is detected. If this analysis is correct, then "no" decisions in this experiment required extra processing steps, and the effects of regularity may have emerged as a consequence of one of these steps. Virtually the same argument was used by Forster and Chambers (1973) to explain why responses to nonwords are always longer than responses to words. In this case, it would be argued that nouns with the incorrect plural suffix (i.e., nonwords)

required an exhaustive lexical whereas words with the correct plural suffix did not. However, this argument is not complete because although it explains why "no" responses were longer (in RTs) it cannot explain why regularity effects were obtained with nonwords only.

One reason why regularity effects may have emerged in the "no" decisions is suggested by the data presented in Table 1. Only 1% of regular feminine nouns take an inconsistent plural suffix compared to 11% of masculine nouns and 32% of irregular feminine nouns. Note in Table 8 that it was most difficult (45.6% errors) to reject regular feminine nouns that take an inconsistent plural suffix (e.g., *pnina* the word for *pearl* calls for an inconsistent plural i.e., *pninim* but in this experiment it was paired with a consistent plural 50% of the time plural e.g., *pninot* producing a nonword) suggesting that, at least to some extent, subjects were basing their "no" responses on the overall likelihood that a noun of a given type could take an inconsistent plural. This explanation also implies that the first step in computing a "no" decision would be to strip the suffix off the base noun (as in prefix stripping; see Taft, 1981), then the base noun would be classified, and the likelihood of inconsistent plural suffixation would be computed. If an exception was likely the lexical search was more extensive and a correct response was therefore also more likely, and if an exception was unlikely the search would have terminated sooner with an error. A similar possibility is that when the lexical search failed (e.g., *pninot* was not found) then the "no" response would first be checked against the statistical properties of the language. By this view, the decision maker would be more likely to assume that the entry (for *pninot*) had been missed if this type of exception

is extremely unlikely. A final possibility is that errors such as *pninot* are simply more likely to occur in the environment thus producing episodic traces of this stimulus and this too would make these items harder to reject.

Another puzzling aspect of the data was the pattern obtained with irregular (unmarked) feminine nouns in the "yes" decisions. These nouns showed a different pattern from regular feminine and regular masculine nouns in both the "yes" decisions, and the "no" decisions. First, rather than producing a difference between "yes" and "no" decisions, the irregular feminine nouns produced the same pattern in both the "yes" and the "no" decisions. Moreover, the effect was in the opposite of the expected direction; nouns with inconsistent plural suffixes produced shorter reaction times and fewer errors. Although the opposite pattern was not robust throughout (e.g., not significant in "no" responses, and not in the errors analysis for "yes" responses; see Table 8), it nevertheless suggests, at the very least, that the computation of plurals agreement for irregular feminine nouns is very different relative to regular masculine and regular feminine nouns.

This result would be predicted if plurals agreement computation proceeds for all nouns primarily on a rule-based mechanism that relies on orthographic features alone. Because irregular feminine nouns essentially "look" masculine, this would explain why the direction of the effect for "no" responses resembled that observed for masculine nouns (i.e., the suffix *-ot* produced faster responses and fewer errors that the suffix *-im*). This explanation also can account for the pattern observed with irregular feminine nouns in the "yes" decisions in a similar fashion, all unmarked nouns (whether masculine or irregular feminine) would be paired with the masculine plural suffix more quickly than

with the feminine plural suffix. The main difficulty with this explanation, however, is that if form-level analysis were the primary mechanism for producing nouns in the plural form suffixes then it is unclear why no effect of plural consistency would have emerged for regular masculine and regular feminine nouns in the "yes" decisions as well. It should be noted that irregular feminine nouns were also unusual in the distribution of plural suffixation type. This can be seen in Table 1. These nouns require plural suffixation evenly split between consistent plurals, inconsistent plurals, and the special plural *-ayim* (which is most frequently used for nouns that have semantically plural characteristic e.g., pants, scissors, and bicycle and especially paired body parts e.g., eyes, legs, ears, hands, knees). Hence, it is perhaps not surprising that these nouns produced a different pattern.

A final aspect of the results that needs to be explained is the data for animate nouns, and nonroot nonwords. Earlier it was suggested that the data for these items might turn out to be qualitatively different from the data for inanimate nouns. For animate nouns, a difference was predicted because they can take both feminine and masculine suffixes, whereas inanimates can only take one or the other. For nonroot nonwords this difference was predicted, whether plurals suffix analysis is a syntactic process or not, because for these items there is no lexical representation and hence response computation is likely to involve different processing stages. To some extent, these differences make it difficult to compare responses to these special items with responses to the inanimate nouns. A related complication is that animate nouns can only be tested in the context of a "yes" decision, and nonroot nonwords can only be tested in the context of a "no" decision, whereas the inanimate nouns were tested in both decision

types. The latter difference turned out to be particularly important because, in fact, a robust interaction was observed for inanimate nouns between plural consistency and decision type (i.e., "yes" vs. "no"). It could nevertheless be argued that the predictions about animate nouns and nonroot nonwords were confirmed in part because the pattern observed was essentially the reverse of that for inanimate nouns (with the exception of the irregular feminine nouns which also produced a different pattern from the regular inanimate nouns). Specifically, the regular inanimate nouns produced no differences between the consistent and inconsistent plural suffix in the context of a "yes" decision, whereas a large difference was observed in the context of a "no" decision. In contrast, animate nouns showed a significant differences in the context of a "yes" decision, and the nonroot nonwords (i.e., "no" decisions) showed no significant effects (although there was a trend in the same direction as that observed for "yes" decisions).

The difference between animate nouns with the masculine plural suffix and those with the feminine plural suffix was unexpected. Nouns with the masculine suffix produced significantly faster reaction times, and a trend in the same direction in the errors analysis. It could be argued that animate nouns have separate representations for all forms (i.e., masculine plural, feminine plural and singular forms), and that differences in frequency accounted for the difference (i.e., feminine plural forms are simply less frequent overall). This seems plausible because whenever there is a group of mixed sexes (men and women together) the masculine suffix is always used, and hence the masculine form is likely to be more frequent. This argument, however, requires independent confirmation (i.e., an extensive frequency count in Hebrew).

Alternatively, it could be argued that this difference also suggests that some computation of agreement is involved in adding a plural suffix, and that for animate nouns the masculine suffix is the default. In this case before an animate noun with the feminine plural suffix could be accepted the lexical entry for the singular noun would be located, the masculine suffix would appear first as one legal variant of the noun, and only subsequently could the feminine suffix be retrieved. One point in support of this argument is that of all the "yes" responses in the experiment, the animate nouns + feminine suffix produced the longest reaction times (and also the highest error rate). This seems to suggest that animate nouns are stored in a representation that contains both plural suffixes, and that the masculine plural is always retrieved first, i.e., it is the "default". Thus, animate nouns with a feminine suffix would take longest to process because two plural suffixes would be retrieved, whereas for inanimate nouns and masculine nouns the first plural suffix retrieved would be the correct match. This interpretation assumes that all plural suffixes are retrieved from the base noun, and that nouns in the plural form are not stored in a separate representation.

A final point is about the absence of a difference between the feminine and masculine suffix when attached to nonroot nonwords. This result suggests that the location of lexical representation for the base-noun is the most critical aspect of analyzing a plural form. Without the location of a representation for the base form, the "default" nature of the masculine suffix (which was significant for animate nouns) is not as robust. Also interesting was that RTs to nonroot nonwords were not faster than RTs to all other "no" responses (real word + incorrect plural suffix). This outcome might have been

expected because for the latter a candidate lexical representation would have been located, and thus presumably some extra time should have been required to process these nonword types. In contrast, nonroot nonwords would not activate any lexical entry and hence should never have reached the stage where the plural suffix would be processed, and therefore it should have been possible to reject these nonwords sooner. This may indicate that whenever a plural suffix is present there is an automatic attempt to parse the plural stimulus (i.e., the suffix is stripped) even if no lexical entry is located. This interpretation suggests the syntactic analysis of plural forms is launched automatically, even when a base-noun representation is not located.

To conclude, the results are considerably more complicated than anticipated and are not easily interpretable within the framework of the intended predictions. Nevertheless, the data overall can be characterized simply as being quite similar to those obtained with grammaticality judgments in Experiments 3-6. This, is perhaps the most conservative interpretation of the present results. Thus, it will tentatively be assumed that to process nouns in the plural form some processing takes place which is analogous to the syntactic analysis that is present in singular grammaticality judgments.

CHAPTER IX — *General Discussion*

The present study investigated the mechanisms of access to the grammatical gender of inanimate nouns in Hebrew. In Experiments 1 and 2 there was a strong effect of transparency in the gender decision task. In Experiment 1 this effect was so strong (39% errors to irregular nouns), that it called into question whether the gender-decision task accurately reflects the mechanisms of gender-access as it occurs in the context of normal agreement processing. Also consistent with this interpretation was the absence of a frequency effect in Experiment 1 suggesting that responses were based on strategic heuristics rather than on information retrieved from the mental lexicon. These heuristics likely included suffix identification (i.e., words that end in ***-ah***, ***-et***, or ***-it*** were classified as feminine, and other words were classified as masculine) despite the quasi-regular correlation between form and gender.

In Experiment 2 all words with a suffix that explicitly marks gender (i.e., feminine regular nouns) were excluded from the item list in an attempt to discourage strategically based responses. Under these conditions a significant frequency effect emerged suggesting that the manipulation was successful in encouraging access to the mental lexicon as a basis for generating responses. Nevertheless, an unusually high proportion of errors (25%) for irregular feminine nouns was still present, again suggesting that different mechanisms are involved during access to gender in natural language processing. In Experiments 3 and 4 a task that included a syntactic context was introduced, and under these conditions a very different pattern of results emerged. In both experiments, there was much smaller rate of errors in responses to irregular feminine

nouns modified by a correctly marked adjective (in Experiment 3) or by a correctly marked verb (in Experiment 4). Moreover, RTs to these nouns were no longer slower relative to masculine nouns as they had been in Experiments 1 and 2 (in fact in Experiment 4 they were actually faster). In contrast to responses to grammatical pairs, responses to ungrammatical pairs in Experiments 3 and 4 produced strong effects of regularity, thus resembling the pattern obtained in Experiments 1 and 2; irregular feminine nouns produced both longer reaction times, as well as unusually high error rates.

In Experiment 5 a task which placed both grammatical and ungrammatical pairs in the context of a "yes" decision, and was intended to discourage syntactic analysis, was introduced. Under these conditions differences between grammatical and ungrammatical pairs persisted, thus ruling out the possibility that these differences (which were found in Experiments 3 and 4 as well) could be attributed to differences between decision types (i.e., "yes" vs. "no"). Moreover, the mere fact that a difference was obtained between grammatical and ungrammatical pairs at all in this experiment suggested that it is not possible to suppress syntactic analysis even when it would be advantageous to do so.

Finally, to provide an even stronger test of the modularity of access to lemma and lexeme representations, grammaticality judgments using plural forms (in which all phrasal elements are explicitly gender marked) were examined in Experiment 6. Despite extensive gender marking the same pattern of results was obtained; there were strong effects of regularity in the processing of ungrammatical pairs, but not in the processing of grammatical pairs. The results of this experiment did however, reveal a much higher rate

of errors relative to Experiment 3 (the experiment to which it was most comparable). Three reasons for this finding were considered: 1) that nonlexical strategic processing was available in this experiment (for the first time since Experiment 1), 2) that there was a greater level of interaction between lexeme and lemma level representations, and 3) that two computations of agreement were required in the plurals grammaticality judgment task rather than just one (as in Experiments 3 and 4). The first two possibilities were rejected because although the overall rate of errors was higher in this experiment, the effects of regularity were not more robust (as both reasons predict). These were also rejected because all regularity effects were limited to the rejection of ungrammatical pairs (and were not present in the identification of grammatical pairs as both reasons 1 and 2 predict.

The third possibility (for the higher error rate obtained in Experiment 6) was evaluated in Experiment 7 by assessing the nature of noun-plural suffix processing (without a modifier). In this experiment, a difference between "yes" and "no" decisions was once again obtained; effects of consistency emerged in "no" decisions but not "yes" decisions. Although a number of unexpected findings were obtained in this experiment, the overall similarity of the pattern of results to those obtained with grammaticality judgments (in Experiments 3, 4, and 6) suggested that nouns in the plural form also require some form of syntactic analysis. Thus, it was concluded that plurals grammaticality judgments that require two syntactic analyses, and that this factor may have produced an increased error rate in the plurals grammaticality judgments task.

The persistent difficulty that native Hebrew speakers demonstrated in accessing

the gender of irregular nouns in the gender decision task (in Experiments 1 and 2) strongly suggests that mere location of a lexical entry does not produce conscious access to gender specification. Further, the results of Experiments 3, 4, and 6 suggest that access to grammatical gender is radically different in the presence of syntactic context, and when explicit report of gender information is not required. Because explicit report of gender is seldom (if ever) required in the context of normal language processing this result is perhaps not surprising. Although regularity effects have been obtained with the gender decision task in other languages (Bates et al, 1995; Bates et al, 1996, Taft & Meunier, 1996) the effects observed in Hebrew appear to be particularly pronounced. They persisted even when reliance on word form was strongly discouraged by eliminating all transparent nouns from the stimulus list. This may be because in Hebrew there is no gender-marked article that can unequivocally yield gender information. This interpretation concurs with the results of Taft & Meunier (1995) who demonstrated that when a gender-marked article is not available in French (e.g., *l'australie*) subjects' gender-decision responses were severely impaired. Therefore, it is suggested that explicit access to gender in the gender decision task may be obtained only with the use of strategies (e.g., article retrieval in French, and suffix identification in Hebrew). Because gender-marking regularity at the form level varies from language to language, this conclusion predicts that relative success in the gender decision task should correlate with the quasi-regularity of gender marking in different languages.

In addition to demonstrating that access to gender is different when an explicit gender response is not required, the results of Experiments 3, 4, and 6 also suggest that it

that the heuristics used in the gender decision task played no role in the detection of gender agreement (in the grammaticality judgment task). Instead, it seems that syntactic processing provides access to a representation of gender that is abstract (i.e., not dependent on the regularities of gender marking at the form level). This conclusion is consistent with those of previous gender studies which also suggested that gender is not retrieved from the level of word-form (Jescheniak & Levelt, 1994; Vigliocco et al., 1996). The main problem with this interpretation of the current results, however, is that if syntactic analysis provides privileged access to gender specification, then this should hold for both grammatical as well as ungrammatical pairs; thus, the difference between grammatical and ungrammatical pairs in Experiments 3, 4, and 6 requires explanation.

One possible account is based on the assumption that, relative to "yes" decisions, "no" decisions involve an additional processing component, or more specifically a "post-lexical check". By this view, whenever an ungrammatical input is detected, the analysis must be checked before a negative response is made. A similar argument was suggested by Bates et al., (1995) to explain regularity effects in Italian gender decisions. The justification for assuming a post-lexical check for "no" decisions, but not "yes" decisions is also consistent with the suggestion that the language system expects inputs to be coherent (or in this case grammatical; e.g., de Groot, Thomassen, & Hudson, 1982; Deutsch & Bentin, 1994). When coherence is violated a check is launched to determine whether an error occurred in the stimulus analysis. Thus, ungrammatical pairs cannot be rejected without first passing a post-lexical check, whereas grammatical pairs concur with the expected direction, and hence are accepted immediately.

The coherence assumption was originally proposed to account for automatic inhibition effects in semantic priming. In that context, unrelated word pairs (e.g., *doctor-tree*) produced slower responses relative to a neutral condition (e.g., *blank-tree*), even in an experiment that did not include any related word pairs (e.g., *leaf-tree*). This finding strongly suggested that the search for coherence between prime and target pairs cannot be suppressed (even when no related prime-target pairs are presented). Further, it suggested that inhibition effects in semantic priming could not be explained solely as resulting from predict-and-match strategies which would be very unlikely in an experiment without any related pairs. The results of Experiment 3 in the present study, suggest that the same (or similar) mechanism operates in a syntactic context as well. Despite the presence of 50% ungrammatical pairs the checking mechanism of ill-formed pairs could not be suppressed.

More specifically, this automatic error checking mechanism might operate as follows: First, it would send a message to the syntactic parser to go back and attempt a second analysis. In addition, it would initiate a backup system that makes use of gender-marking regularity. Because most nouns are regular, the back-up system would produce the correct agreement analysis in most cases as well. It may seem that this two-step checking process is inherently less parsimonious because it calls for two independent processes (lemma access and the computation of regularity) that essentially produce the same analysis (i.e., gender agreement). However, an important advantage to this kind of redundancy is an increased probability of recovering from analysis errors. In fact, such a mechanism may develop specifically to compensate for cases when normal parsing mechanisms have temporarily gone awry. Whenever the lemma becomes temporarily

unavailable (e.g., when it is out to lunch), it is useful to have a backup system that does not rely on lemma access to compute gender agreement instead.

One alternative to this automatic checking mechanism could be that the presence of an error actually alerts the conscious system itself to the detection of a problem. The results of the gender decision experiments showed that there is conscious access to regularities in gender-marking at the level of form. Thus, there could be conscious system that makes use of the same metalinguistic skills (e.g., if it does not end in ***-a*** it is probably masculine) used in the gender decision task. This too would produce a difference between grammatical and ungrammatical pairs. By this view, the rejection of ungrammatical pairs showed regularity effects because marked feminine modifiers (e.g., ***zrizah***) with unmarked masculine nouns (e.g., *maslul)* looked more obviously ungrammatical than unmarked feminine nouns (e.g., *nefesh*; which actually look like most masculine nouns) paired with masculine modifiers (e.g., *zariz*). Also consistent with this conclusion are the results of Experiment 5 which showed that it was easier to ignore a grammaticality violation when an irregular feminine noun was involved. This explanation, however, seems less plausible for two reasons. First, because if it is a conscious system that computes the grammaticality checking, then it also seems that it should have been possible to suppress it strategically (to make the task easier). Further, if the checking mechanism is not launched automatically then recovery from errors would be delayed until after engagement of the conscious system, and this seems implausible because of the efficient and primarily error-free nature of the majority of linguistic processing.

A final possibility is that both automatic and conscious level checking mechanisms are present. One reason to suppose that both might be involved is that ill-formed inputs pose a threat to the detection of a correct linguistic message, and hence it would be advantageous to have several ways to avoid this problem. When both conscious and unconscious processing are directed back to the incoming stimulus, the chances that the error might be detected and corrected are increased. The introduction of explicit attention could function to reinforce the computations of the implicit processing mechanisms. For example, if the syntactic parser has produced an incorrect, implausible, or ambiguous analysis, then the conscious system may nevertheless be able to interpret the sentence by focusing attention on environmental cues.

The coherence assumption account suggests that when a grammatical pair is presented the gender for both the noun and adjective is first identified, and then they are checked for a match. When a match is detected a "yes" response quickly follows. In contrast, when an ungrammatical pair is presented a mismatch is detected, and because mismatches are not expected, the "no" response is delayed and a check is launched. This post-lexical check invokes the use of regularity at the level of form, and hence a regularity effect emerges in "no" decisions, but not in "yes" decisions. Because 50% of the pairs in Experiments 3, 4, and 6 were ungrammatical it further appears that the post-lexical check is automatic and cannot be suppressed, otherwise it would have been much easier for subjects to eliminate this processing stage.

Returning to the question of modular and interactive models recall that the interactive view predicted difficulty in accessing the gender of irregular nouns because

there is nothing at the lexeme level to facilitate access to the lemma level. The modular view, on the other hand, posited that the retrieval mechanisms of lemmas and lexemes do not interact. It is suggested that the results overall are more consistent with a modular characterization of lexical retrieval despite the robust effects of regularity observed in both the gender decision task (in Experiments 1 and 2), and in responses to ungrammatical pairs (in Experiments 3, 4 and 6). This conclusion is based primarily on the data obtained with grammatical pairs in Experiments 3, 4, and 6 (noun-adjective, and noun-verb grammaticality judgments), and also on the above-suggested interpretation which essentially discounted the relevance of the regularity effects observed in gender decision.

Presumably, grammaticality judgments approximate more closely the normal circumstances under which access to gender occurs. In this task no difference was found between regular and irregular nouns in Experiments 3 and 6, and in Experiment 4 the reaction time data actually revealed a significant effect in the opposite direction (a result which neither model predicts but is certainly less consistent with the interactive view). Moreover, because the vast majority of linguistic inputs are error-free, it is presumably the data obtained with grammatical pairs that reflect the gender-processing mechanisms of interest. Finally, it must be stressed that response computation in a comprehension task necessarily begins with the analysis of word-forms where regularity is most conspicuous. Thus, it is important to consider that absence of regularity effects in these experiments (especially Experiment 6 in which all phrasal elements were explicitly gender-marked) constitutes particularly strong evidence of the modularity of lemma and

lexeme retrieval.

Regardless of how the difference between grammatical and ungrammatical pairs in Experiments 3, 4, and 6 is explained, the results demonstrate that the detection of agreement errors, but not the computation of agreement itself, is influenced by regularity. Similar distinctions between error-detection and other processes of grammatical encoding have been discussed previously in the language production literature. For example, Levelt (1993) showed that 57% of phrase final errors were detected and corrected, whereas only 15% of phrase initial errors were. This difference in the probability of correcting an error at the beginning vs. at the end of a phrase suggested that it is easier to switch attention from sentence planning to error detection in phrase final locations. The apparent need for a switch in attention, in turn, motivated a proposal that there is an independent error-checking monitor (but see MacKay, 1987 for a connectionist theory of error detection and correction which places the language monitor inside the production system). Levelt also suggested that the error detection mechanism in language production actually relies on the language comprehension system (and hence he called this mechanism the "perceptual loop").

Levelt's (1983) proposal combined with the data from the present experiments make an interesting prediction about agreement errors in language production. If the language production monitor makes use of the language comprehension system as a mechanism for error detection, then speech errors in gender agreement should occur more frequently with irregular nouns. This is because the data in Experiments 3, 4, and 6, suggest that the language comprehension system is less effective in rejecting gender-

agreement errors for irregular nouns relative to regular nouns, and similarly the data in Experiment 5 suggested that it was easier to overlook a gender-agreement error with an irregular noun. Further, the interactive view predicts that the same pattern of results should emerge in the production of correct agreement (i.e., error-free language production), whereas the modular view predicts that the production of correct agreement should proceed equally fast for both regular and irregular nouns. Thus, RT and error data in language production could serve to further test the distinctions between the two model types.

Similar predictions have been tested, but not in the domain of grammatical gender. For example, Baars, Motley, & MacKay (1975) showed that whole-word substitutions are produced more frequently than phoneme substitutions (the lexical bias effect). They suggested that this is because whole-word substitutions more frequently produce meaningful utterances, and therefore they are harder for the language monitor to detect and repair (but see Dell for a rather different account of this effect 1986). A similar interpretation can explain a result obtained in the production of plural agreement errors in English. Eberhard (1997) showed that more agreement errors were found in the production of unmarked singular pairs (e.g., the key to the cabinets are...) relative to marked plural counterparts (e.g., the keys to the cabinet is...). In this study it was argued that processing of unmarked representations (i.e., singular forms) is relatively impaired in the context of marked counterparts (i.e., plural forms), thus implying that agreement errors for marked forms are harder to detect and repair.

To reconcile the current data and the conclusions based on them with those drawn

by Bates et al., (1996) and Grosjean et al., (1994) alternative interpretations of the regularity effects they observed in phrase priming tasks are needed. In those studies it was found, for example, that a feminine noun (e.g., *casa*) is recognized faster when it appears after a feminine-marked adjective (e.g., *brutta*) than when it appears after a neutral or masculine marked adjective. These effects seemed to indicate that gender plays a role in the process of lexical access, and that regularity influences both lexical access as well as access to gender. Grosjean et al., (1994) suggested an alternative interpretation, namely, that these effects (or that part of the effects) may simply reflect post-lexical syntactic integration (an argument similar to that described above to account for the data obtained with "no" decisions). However, they were not convinced that this interpretation could account for the results they obtained in the gating paradigm (i.e., when a gender prime was present words were identified earlier, and also the candidates proposed were always restricted to the gender of the prime suggesting that the prime had modified the actual candidate list in lexical search). Similarly, Bates et al., (1996) argued that the presence of facilitation (i.e., priming relative to a neutral baseline) in addition to inhibition (i.e., priming relative to a discordant baseline) in their data suggested that at least part of the effect of the gender prime must be an actual lexeme-access effect. It should be noted, however, that the on-line nature of results obtained with use of the gating paradigm has been questioned. Moreover, the facilitation effects that Bates et al., (1996) obtained were significant in the phrase primed word-repetition task only, and not in the phrase-primed gender decision, nor in non-primed word repetition (Bates et al., 1995). Thus, it seems that further investigations are necessary before this debate is

settled.

To summarize, the results of the present experiments demonstrate first that mere access to the lexeme level does not produce conscious access to grammatical gender. Second, they suggest that access to gender is radically different in the presence of syntactic context, and finally, that the computation of gender agreement is not affected by regularity at the level of the word-form whereas this factor does play a role in the detection of agreement errors. When asked for an explicit report of grammatical gender, subjects resort to unreliable heuristics that produce extreme effects of regularity in gender-marking at the level of form (i.e., 25-40% errors in gender decision). This same result is obtained even when there is evidence of contact with the mental lexicon (i.e., a frequency effect), and when the possibility of non-lexical strategies (i.e., suffix detection) is eliminated. From these results a number of conclusions may be drawn. First, the presence of so many errors in the gender-decision task strongly implies that this task does not accurately reflect the mechanisms by which gender is accessed in the course of natural language processing. Second, these results are consistent with models that postulate an abstract representation of grammatical gender at the lemma level that does not depend on, and is not linked to gender-marking at the word form level. Third, these results combined with those obtained in previous studies suggest that the robust effects of regularity that are frequently observed in gender processing tasks may reflect post-lexical processing mechanisms.

APPENDIX - *Tables*

Table 1. *Probabilities of encountering regular & irregular gender marking in Hebrew.*

Noun Type	**Examples**	**% of nouns**	**Plural Type**	**Examples**	**% of suffix type**
Masculine	*masmer*	56.5	regular	*masmerim*	89
	vilon		irregular	*vilonot*	11
	katef		*-ayim*	*ktefayim*	1
Regular Feminine	*mita*	42.2	regular	*mitot*	99
	beytsa		irregular	*beytsim*	1
	--		*-ayim*	--	0
Irregular Feminine	*lashon*	1.3	regular	*leshonot*	39
	even		irregular	*avanim*	32
	regel		*-ayim*	*raglayim*	29

Table 2. *Mean gender-decision times in milliseconds, and error rates in Experiment 1.*

Word Type	M	%error
Regular Feminine	671	9.4
Masculine	728	9.7
Irregular Feminine	784	37.9
Regular plural	710	11.9
Irregular plural	722	14.6

Table 3. *Mean gender-decision times in milliseconds and error rates in Experiment 2.*

Word Type	M	%error
Masculine	847	12.4
Irregular	889	24.8
Feminine		
High frequency	838	18.1
Low frequency	887	18.5
Difference	*49***	*0.4*

* significant p<.05 level
** significant p<.01 level

Table 4. *Mean noun-adjective grammaticality-judgment decision times in milliseconds and error rates in Experiment 3.*

Pair Type	Word Type	M	%error
Grammatical	Masculine	1277	6.0
Pairs	Irregular Feminine	1299	10.5
	difference	*22*	*4.5*
Ungrammatical	Masculine	1368	9.3
Pairs	Irregular Feminine	1456	25.2
	difference	*88**	*15.9***

* significant $p<.05$ level
** significant $p<.01$ level

Table 5. *Mean noun-verb grammaticality-judgment decision times in milliseconds and error rates in Experiment 4.*

Pair Type	Word Type	M	%error
Grammatical	Masculine	1270	14.2
Pairs	Irregular Feminine	1171	12.5
	difference	*-99***	*-1.7*
Ungrammatical	Masculine	1296	14.2
Pairs	Irregular Feminine	1375	38.1
	difference	*79**	*23.9***

* significant $p<.05$ level
** significant $p<.01$ level

Table 6. *Mean Two-word lexical decision times in milliseconds and error rates in Experiment 5.*

Pair Type	**WORD-WORD**	M	%error
Grammatical	Masculine	771	4.0
	Irregular Feminine	790	7.9
	Difference	*19*	*3.9*
Ungrammatical	Masculine	824	7.2
	Irregular Feminine	780	9.3
	Difference	*-44**	*2.1*
Pair Type	**1st & 2nd Word Position**		
	NONWORD-WORD	759	5.4
	WORD-NONWORD	841	20.1
	Difference	*82***	*14.7***

* significant p<.05 level
** significant p<.01 level

Table 7. *Mean plurals grammaticality-judgment decision times in milliseconds and error rates in Experiment 6.*

Subject Group	Pair Type	Plural Type	M	%error
ALL Subjects	**Grammatical**	Consistent	1451	16.7
(n = 50)		Inconsistent	1518	20.2
		Difference	67	*3.5*
	Ungrammatical	Consistent	1560	19.5
		Inconsistent	1594	34.0
		Difference	*34*	*14.5***
Low Error Subjects	**Grammatical**	Consistent	1506	9.2
(n = 28)		Inconsistent	1556	13.3
		Difference	*50*	*4.1*
	Ungrammatical	Consistent	1576	13.0
		Inconsistent	1674	21.8
		Difference	*98**	*8.8***

* significant $p<.05$ level
** significant $p<.01$ level

Table 8. *Mean plurals lexical-decision times in milliseconds and error rates in Experiment 7.*

		YES-decision			NO-decision	
	Plural Type	M	% err		M	% err
Masculine	CP *brazim*	735	3.4	CP *brazot*	887	9.8
	IP *mazalot*	732	3.2	IP *mazalim*	1006	25.4
		-3	*-0.2*		*119***	*15.6***
Feminine	CP *legimot*	735	5.7	CP *legimim*	927	9.8
	IP *pninim*	726	0.6	IP *pninot*	1103	45.6
		-9	*-5.1**		*176***	*35.8***
Irreg Fem	CP *aracot*	767	3.0	CP *aracim*	991	25.0
	IP *gfanim*	716	1.7	IP *gfanot*	916	20.3
		*-51**	*-1.3*		*-75*	*-4.7*
	Animate			**Non-root Nonword**		
	Masc Suffix	710	2.6	Masc Suffix	848	4.5
	Femn Suffix	788	8.7	Femn Suffix	911	4.5
		*78**	*6.1*		*63*	*0.0*

* significant p<.05 level
** significant p<.01 level

REFERENCES

Badecker, W., Miozzo, M., & Zanuttini, R. (1995). The two-stage model of lexical retrieval: Evidence from a case of anomia with selective preservation of grammatical gender. *Cognition*, **57**, 193-216.

Balota, D. (1992). Visual word recognition: The journey from feature to meaning. In M. Gernsbacher (ed.), *Handbook of psycholinguistics* (pp. 303-357). New York: Academic Press.

Bates, E., Devescovi, A., Hernandez, A., Pizzamiglio, L., (1996). Gender priming in Italian. *Perception & Psychophysics*, **58**(7), 992-1004.

Bates, E., Devescovi, A., Pizzamiglio, L., & Hernandez, A. (1995). Gender and lexical access in Italian. *Perception & Psychophysics*, **57**, 847-862.

Bock, J. K. (1982). Toward a cognitive psychology of syntax: Information processing contributions to sentence formulation, *Psychological Review*, **89**, 1-47.

Bock, K., & Levelt, W. (1994). Language production. In M. Gernsbacher (ed.), *Handbook of psycholinguistics* (pp. 945-984). New York: Academic Press.

Butterworth, B. (1989). Lexical access in speech production. In B. Butterworth (Ed.), *Language Production: Vol. 1. Speech and talk*. New York: Academic Press.

Cutler, A. (1987). Speaking for listening. In A., Allport, D.G., MacKay, W., Prinz, & E., Scheerer, (Eds.), *Language perception and production: Relationships between listening, speaking, reading, and writing*. London: Academic Press.

de Groot, A.M.B., Thomassen, A.J.W.M., & Hudson, P.T.W. (1982). Associative facilitation of word recognition as measured from a neutral prime. *Memory & Cognition*, **10**, 358-370.

Dell, G.S. (1986). A spreading-activation theory of retrieval in sentence production. *Psychological Review*, **93**, 283-321.

Dell, G.S. (1988). The retrieval of phonological forms in production: Tests of predictions from a connectionist model. *Journal of Memory and Language*, **27**, 124-142.

Dell, G.S. (1990). Effects of frequency and vocabulary type on phonological speech errors. *Language and Cognitive Processes*, **5**, 313-349.

Dell, G.S., & O'Sheaghda, P.G. (1991). Mediated and convergent lexical priming in language production: A comment on Levelt Et al., (1991). *Psychological Review*, **98**,

604-614.

Dell, G.S., & O'Seaghdha, P.G. (1992). Stages of lexical access in language production. *Cognition*, **42**, 287-314.

Deutsch, A. & Bentin, S. (1994). Attention mechanisms mediate the syntactic priming effect in auditory word identification. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **20**, 595-607.

Fay, D., & Cuter, A., (1977). Malapropisms and the structure of the mental lexicon. *Linguistic Inquiry*, **8**(3), 505-520.

Forster, K.I. (1994). Computational modeling and elementary process analysis in visual word recognition. *Journal of Experimental Psychology: Human Perception and Performance*, **20**, 1292-1310.

Forster, K.I., & Chambers, S.M. (1973). Lexical access and naming time. *Journal of Verbal Learning and Verbal Behavior*, **12**, 627-635.

Frauenfelder, U., & Tyler, L. (Eds.; 1987). *Spoken word recognition*. Cambridge, MA: MIT Press/Bradford Books.

Fromkin, V. (1971). The non-anomalous nature of anomalous utterances. *Language*, **47**, 27-52.

Garrett, M.F. (1975). The analysis of sentence production. In G.H. Bower (Ed.), *Slips of the tongue and language production* (pp. 133-177). New York: Academic Press.

Garrett, M.F. (1976). Syntactic processes in language production. In R. Wales, & E. Walker (Eds.), *New approaches to language mechanisms. A collection of psycholinguistic studies.* Amsterdam: North-Holland.

Garrett, M.F. (1980b). The limits of accommodation: Arguments for independent processing levels in sentence production. In V.A. Fromkin (Ed.), *Errors in linguistic performance: Slips of the tongue, ear, pen, and hand.* New York: Academic Press.

Garrett, M.F. (1992a). In E., Kittay, & A., Leherer (Eds.), *Frames, fields, and contrasts*. Hillsdale, NJ: Erlbaum.

Grainger, J. (1992). Orthographic neighborhoods and visual word recognition. In R. Frost, & L. Katz (Eds.), *Orthography, Phonology, Morphology, and Meaning*. The Netherlands: Elsevier Science Publishers.

Grosjean, F., Dommergues, J., Cornu, E., Guillelmon, D., & Besson, C. (1994).

The gender-marking effect in spoken word recognition. *Perception & Psychophysics*, **56**, 590-598.

Jescheniak, J.D., Levelt, W.J.M., (1994). Word frequency effects in speech production: Retrieval of syntactic information and of phonological form. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **20**, 824-843.

Kempen, G., & Huijbers, P. (1983). The lexicalization process in sentence production and naming: Indirect election of words. *Cognition*, **14**, 185-209.

Levelt, W.J.M. (1983). Monitoring and self-repair in speech. *Cognition*, **14**, 41-104.

Levelt, W.J.M. (1989). *Speaking: From intention to articulation*. Cambridge, MA: MIT Press.

Levelt, W.J.M. (1992). Lexical access in speech production [Special issue]. *Cognition*, **42**, (1-3).

Marslen-Wilson, W.D. (1987). Functional parallelism in spoken word-recognition. *Cognition*, **25**, 71-102.

MacKay, D.G. (1987). *The organization of perception and action: A theory for language and other cognitive skills*. New York: Springer-Verlag.

Pinker, S. (1991). Rules of language. *Science*, **253**, 530-535.

Stemberger, J.P. (1985). An interactive activation model of language production. In A. Ellis (Ed.), *Progress in the psychology of language*(pp. 143-186). London: Erlbaum.

Taft, M. (1981) Prefix stripping revisited. *Journal of Verbal Learning and Verbal Behavior*, **20**, 289-297.

Taft, M., & Meunier, F. (1995). A challenge for neural networks: Gender processing in French. *Paper presented at the 36th annual meeting of the Psychonomic Society*, Los Angeles, November.

Vigliocco, G., Antonini, T., Garrett, M.F., (1997). Grammatical gender is on the tip of Italian tongues. *Psychological Science*, p314-318.

IMAGE EVALUATION TEST TARGET (QA-3)

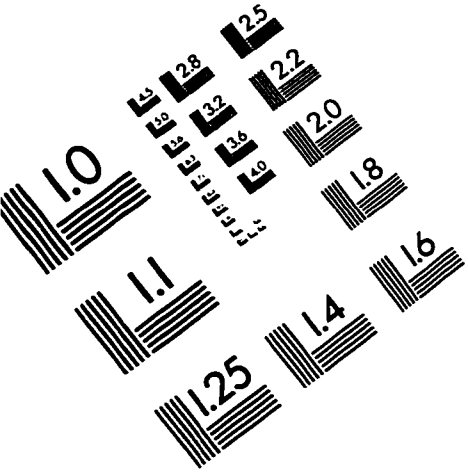

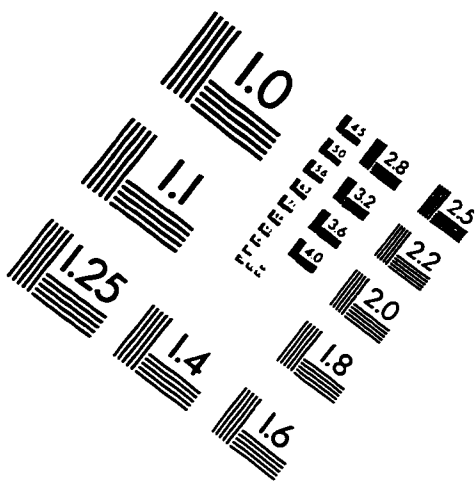

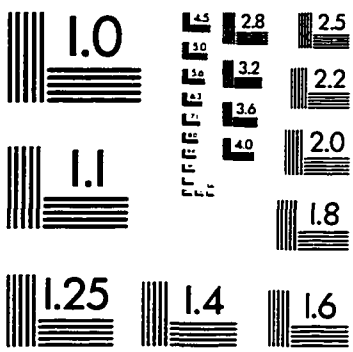

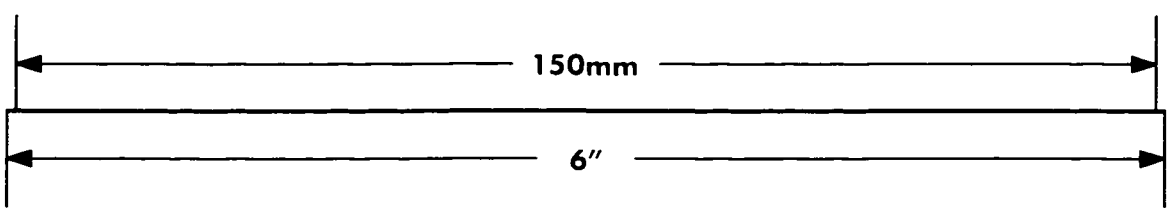

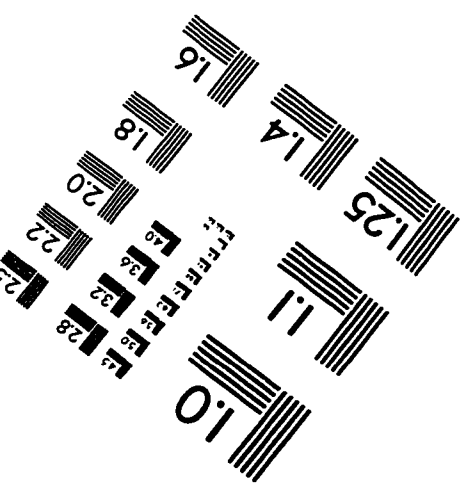

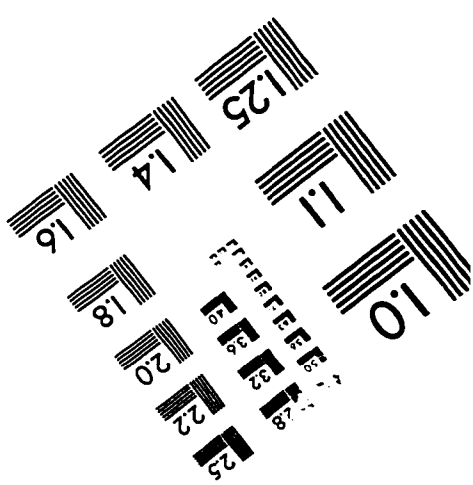

APPLIED IMAGE, Inc
1653 East Main Street
Rochester, NY 14609 USA
Phone: 716/482-0300
Fax: 716/288-5989