1 Introduction

When the scientific discipline of linguistics originated in the eighteenth century, the subdiscipline of historical linguistics originated with it. Indeed, in the prevailing view of the time, the study of language was necessarily historically oriented (Burridge 2013). While the twentieth century saw a methodological shift towards the synchronic study of language and the birth of general linguistics, historical linguistics as it is practiced today still basically operates on the principles developed by a group of scholars that dominated the last quarter of the nineteenth century, the so-called Neogrammarians (German: *Junggrammatiker*). Their main thesis, known as the Neogrammarian Hypothesis, was that all changes in the pronunciation of words that were not due to language contact or analogy (see below) resulted from *regular, phonetically conditioned sound changes*, or *sound laws*. In their own words:

> When we speak of systematic effect of sound laws we can only mean that given the same sound change within the same dialect every individual case in which the same phonetic conditions are present will be handled the same. Therefore either wherever earlier the same sound stood, also in the later stages the same sound is found or, where a split into different sounds has taken place, then a specific cause – a cause of a purely phonetic nature like the effects of surrounding sounds, accent, syllabic position, etc. – should be provided to account for why in the one case this sound, in the other that one has come into being.[1] (Paul 1880: 69, cited in translation in Hale 2003: 343)

Historical linguists in the twenty-first century still adhere to the Neogrammarian Hypothesis because it has proven to work. The assumption that sound change is purely phonetically conditioned has led to the highly detailed reconstruction of

[1] *Wenn wir daher von konsequenter Wirkung der Lautgesetze reden, so kann das nur heissen, dass bei dem Lautwandel innerhalb desselben Dialektes alle einzelnen Fälle, in denen die gleichen lautlichen Bedingungen vorliegen, gleichmässig behandelt werden. Entweder muss also, wo früher einmal der gleiche Laut bestand, auch auf den späteren Entwicklungsstufen immer der gleiche Laut bleiben, oder, wo eine Spaltung in verschiedene Laute eingetreten ist, da muss eine bestimmte Ursache und zwar eine Ursache rein lautlicher Natur wie Einwirkung umgebender Laute, Akzent, Silbenstellung u. dgl. anzugeben sein, warum in dem einen Falle dieser, in dem andern jener Laut entstanden ist.*

unattested proto-languages. In some cases, features of these reconstructions have later been confirmed by their reflexes in newly-discovered languages. On the other hand, the Neogrammarian promise of regularity has encouraged scholars to take a closer look at seemingly random variation in the reflexes of what should be the same sound, uncovering the subtle rules governing this variation in the process. Thus, the hypothesis that sound changes are phonetically regular seems to match the facts. It gives linguists insight into earlier stages of the languages they study and lets them understand how the attested forms came to be.

While the Neogrammarian paradigm is thus largely followed by most modern historical linguists, scholars of the Semitic languages have been, to a certain extent, "a people dwelling alone, and not counting itself among the nations". Perhaps due to the strong philological tradition in the scholarship of such languages as Arabic, Aramaic, and Hebrew, the field of comparative Semitics tends to be more tolerant of loosely formulated sound laws with unexplained exceptions and allows for non-phonetic factors to condition sound change. In the significantly titled 'Non-phonetic conditioning of sound change and Biblical Hebrew', the eminent Hebrew and Arabic scholar Joshua Blau concludes:

> I have tried to show that sound change is not always strictly phonetically conditioned. I have demonstrated on the strength of biblical material that functionally significant sounds may be preserved in positions in which they are, as a rule, omitted (...), and that phonemes are apt to behave differently from phonetically identical allophones (...). On the other hand, I do not consent to the often expressed opinion that sound changes may be limited to certain grammatical classes to the exclusion of others, without any historical, phonetic, or functional reasons, or analogical formation. I am convinced that accurate analysis of such alleged cases is apt to discover special reasons that led to the restriction of a certain sound change to a special grammatical class. (Blau 1979: 14)

As will be clear, this allowance for the functional conditioning (in Blau's case) of sound change contradicts the Neogrammarian Hypothesis, which only allows for sound changes to be phonetically conditioned. Also note his reference to "the often expressed opinion that sound changes may be limited to certain grammatical classes", which violates the Neogrammarian Hypothesis even further. This contradiction suggests that an investigation into the defensibility of the Neogrammarian Hypothesis from a Semitic point of view may be fruitful. If there is sufficient

evidence for non-phonetic conditioning of sound change, as many Semiticists maintain, this is an important finding for the entire field of historical linguistics. If, on the other hand, the Neogrammarian Hypothesis is as reliable as its continued acceptance among non-Semiticists suggests, a thorough application to the Semitic family may yield interesting new results. It is to this controversy about the conditioning factors of sound change that this work seeks to contribute.

This work aims to determine to what extent the development of the vocalic phonemes from Proto-Northwest-Semitic to Biblical Hebrew can be described by exceptionless sound laws. We may motivate this research question by considering its main constituents. The focus on vowels is simply because the development of the consonants – excepting the glides, *w and *y – is uncontroversial and well understood, the only issue being the absolute chronology of some phonetic changes and mergers; the vocalic system of Biblical Hebrew, however, is quite different from that of older and more archaic forms of Semitic, as will be seen in chapter 2, and its development is debated. The other parts of the question require a lengthier introduction, which will be taken up in the following sections. First, why exceptionless sound laws? Second, why Biblical Hebrew? And third, why Proto-Northwest-Semitic?

1.1 Exceptionless sound laws

As the quote from Paul (1880) above indicates, the Neogrammarian Hypothesis states that sound change within a single speech variety is purely phonetically conditioned. In this section, we will consider why this should be the case, taking some more recent scholarship into account, and examine some sources of apparently irregular sound change.

1.1.1 Regular sound change

The purely phonetic conditioning of sound change follows from the normal process of language acquisition (Paul 1920: 69[2], Hale 2003). Mature speakers possess a phonemic representation of the different lexemes and morphemes present in a language. As these phonemic representations are psychological entities, however,

[2]This is the fifth edition of Paul's *Prinzipien der Lautgeschichte* and the last to appear during the author's lifetime. This edition is cited in the rest of this introduction to present the most developed version of the Neogrammarian view of language change.

they cannot be directly communicated to listeners. Rather, the speaker takes the phonetic representations and derives a set of instructions for their speech organs, which then produce a phonetic signal. The way a given phoneme is articulated can be affected by surrounding phonemes, with articulatory motions starting early and persisting for a while. This effect is known as coarticulation. Other factors also affect the phonetic outcome of a certain phoneme. In rapid speech, coarticulation may increase and speech organs may make smaller movements. Additionally, the speaker may realize the same phoneme differently when speaking in different contexts, while different variant realizations may also alternate within the same context (Guy 2003). Together with chaotic variation in articulation and background noise, these factors result in a virtually infinite range of different phonetic realizations of the same underlying forms and phonemes.

A first language learner is then presented with this greatly varied set of phonetic forms produced by speakers in their environment. To be able to produce intelligible utterances (only taking the phonological aspects into account), the learner needs two things, or rather two groups of things: the phonemic representations belonging to individual morphemes and the rules which tell the speech organs how to pronounce a given phoneme in a given context. By recognizing that different phonetic realizations belong to the same word, the learner may infer the underlying phonemic representations. For example, the observation that [kʰæʔ], [kʰɛətˤ], [kʰætʰ], etc. all refer to the same animal allows the learner to posit a single underlying form for the word, e.g. /kæt/. Observing the same variation in the last sound of a word encountered by the learner as [hæʔ], [hɛətˤ], [hætʰ], etc., leads to the same conclusion: both words end in /t/, the latter being stored as /hæt/. Even when different variant forms are found in different words, the learner may identify them as the same phoneme based on patterns of distribution. Although the first sound in *keep* is usually more palatalized than that in *cat*, the difference in this one sound alone never signals a difference in meaning. Moreover, the same predominance of palatalized realizations occurs in *keel*, *keen*, *key*, etc., allowing the learner to identify the following /i/ as the cause of this difference in pronunciation. Accordingly, all of these words can be stored in the learner's mental lexicon with the same initial phoneme, /k/. At the same time, the learner may make a mental note that /k/ before /i/ is to be pronounced with the tongue slightly further forward than when it occurs before /æ/. In this way, the phonemic representations of the lexicon and the phonetic rules producing a given

phoneme's different allophones are transmitted from generation to generation with remarkable faithfulness, considering the complexity of the process.

Thus, the learner tries to acquire both the phonemes present in the language and the rules governing their phonetic realization. If either of these processes occurs with less than 100% accuracy, the learner-become-speaker will produce a significantly different set of phonetic surface forms than the speakers of the previous generation. It may be that the learner is presented with a skewed sample of realizations of a given phoneme in a certain context. Suppose, for example, that /k/ before /i/ is realized on a scale from non-palatalized [k] to highly palatalized [kʲ] in a certain language variety. One learner coincidentally hears an exceptionally large number of [kʲ] realizations. This learner would then internalize this as the normal realization of /k/ in this position. As a result, the learner will produce more palatalized instances of /k/ before /i/ than speakers of the previous generation. Note that this is merely a phonetic change, with no effect on the underlying phonemic representations. As the change concerns the transformation from phonemic representations into phonetic surface forms, it should affect all instances of the same phoneme in the same environment equally. In other words, this kind of phonetic sound change is phonetically conditioned.

Phonemic change, on the other hand, results from a mistaken identification of the underlying representations, as described by John Ohala (1981). As we have seen, learners normally correct for coarticulation, noise, and other factors which give rise to various variant phonetic realizations of a certain phoneme. Ohala notes two ways this can go wrong, which he terms *hypocorrection* and *hypercorrection*.[3] When hypocorrecting, the learner fails to account for coarticulation and other effects. Using the same example of the articulation of /k/ from the last paragraph, the learner may not notice that the words with a large proportion of [kʲ] realizations all have an /i/ following this sound. Instead of identifying these realizations as allophones of /k/, the learner erroneously encodes them as belonging to a separate phoneme, e.g. /kʲ/. This kind of change is frequently accompanied by the loss of a phonemic distinction which originally conditioned the allophonic variation. Thus, if the previous generation had phonemic representations with /ki/, typically realized as [kʲə] when unstressed, /ka/, typically realized as [kə] when unstressed, and /ku/, also typically realized as [kə] when unstressed, the

[3]Not to be confused with the sociolinguistic processes going by the same names, where speakers incorrectly adjust their production to another language variety or register; for some examples of these kinds of hypo- and hypercorrections in Hebrew and related languages, see Blau (1970).

learner may represent these surface forms as /kʲə/ instead of original /ki/ and /kə/ instead of original /ka/ and /ku/. As this example illustrates, hypocorrection can result in both phonemic split (/k/ splitting into /k/ and /kʲ/) and phonemic merger (/i/, /a/, and /u/ all merging into /ə/). Diachronically, the effects of hypocorrection surface as assimilation, lenition, and other forms of reduction.

While hypocorrection involves a lack of necessary correction in the interpretation of a phonetic signal, hypercorrection is the result of the learner correcting for features when this is not necessary. Certain features of the phonetic signal are ascribed to coarticulation, even though they actually belong to the underlying phoneme. For example, if an underlying sequence of /um/ is usually produced as [um], the learner may falsely interpret the roundedness of the [u] to coarticulation, caused by the following [m]. The learner will then hypercorrect this feature and store the phonemic representation as /im/, resulting in less rounded realizations of this sequence in future. Hypercorrection is thus responsible for the diachronic process of dissimilation.

As with phonetic change, described above, the process of phonemic change ensures its purely phonetic conditioning. The phonological learning mechanism is designed to identify one and the same underlying phoneme based on a large, but not random, variation in phonetic realization; recall the example of [kʰæʔ], [kʰɛət̚], [kʰætʰ], [hæʔ], [hɛət̚], [hætʰ], etc. given above, all of which are normally correctly identified as ending in the same phoneme, /t/. Due to the rules governing phonetic production, the speakers providing the learner with input will produce a similar set of variant realizations for the same phoneme in the same environment, regardless of whether the phoneme occurs in a noun or a verb, whether it expresses a functional contrast or a redundant one. Accordingly, if the learner interprets a set of surface forms as deriving from a different underlying representation than the speakers intended due to hypocorrection or hypercorrection, he or she will apply the same mistaken identification to the relevant phoneme wherever it occurs in the same environment. The phonetic context is the only factor that can condition sound change because it is the only factor that affects the phonetic realization of a given phoneme in a consistent manner.

Thus, both phonetic and phonemic sound change are predicted to be phonetically conditioned and regular, as the Neogrammarian Hypothesis states. The latter does allow, however, for irregular change, provided that this is not the result of 'natural' sound change. Two long recognized causes of apparent exceptions to regular sound laws must be mentioned. The first follows from the limitation that sound

change operate regularly within one speech variety. Naturally, a sound change in one language or dialect need not affect any other language or dialect. If a word is borrowed into a speech variety which has undergone different sound changes than that in which it originates, such a word may seem to violate sound laws. The English word *skirt*, for example, has not undergone the regular English change of *sk > *sh*. This is not due to any irregular operation of sound change, however, but because it is a loanword from Old Norse, which had not participated in this change. Its inherited cognate *shirt* shows that the sound law changing *sk to *sh* did regularly apply in this phonetic environment. Where possible, loanwords will be identified and disregarded in the rest of this work.

The second important cause of apparent exceptions to sound laws is analogy and related phenomena. These will be discussed in some depth in the next section.

1.1.2 Analogy and morphological change

Broadly speaking, *analogy* may refer to any linguistic process where the shape of a word is influenced by the shape of another, semantically related word (Fertig 2013: 12–13). For the sake of clarity, however, the Neogrammarian distinction between *analogy* in the strict sense (also 'proportional analogy', 'four-way analogy') and other processes such as *contamination* and *folk etymology* (see below) will be maintained in this work.

Just as regular sound change is a by-product of normal phonological acquisition, analogy is a side effect of the normal operation of morphological rules (Paul 1920: 106–120). In most, if not all languages, different words may be similar to a certain extent in both form and meaning. Often, the difference in form is associated with the same difference in meaning in a number of words. For example, the relationship between *cat* : *cats*, *dog* : *dogs*, *bee* : *bees*, etc. is the same in each case: when compared to the singular, the plural has an additional *-s* (phonemically /-z/) attached to the end. Based on this regularity, learners can formulate a rule: in order to form a plural noun, take the singular and add *-s*. This saves the learner the trouble of having to separately store every single plural form in the language. Additionally, speakers can use such rules to produce forms they have never heard, as in the famous 'wug test' conducted by Jean Berko (1958): when presented with nonsense words like *wug*, both children and adults consistently produce plural forms like *wugs*. Analogy, then, is the application of such a rule to derive a word that did not previously exist in the language. Applying the English pluralization

rule to *man* regularly gives *mans*; should this form become popular and replace the inherited form, *men*, this would be a case of analogical change. Although such analogical changes rely on the presence of perceived rules in a language, it is customary to represent them as proportions, with a pre-existing example of the rule acting as a model for the new creation. For example, the creation of *mans* may be represented as *dog* (singular) : *dogs* (plural) = *man* (singular) : *mans* (plural).

Analogy is relevant to the study of sound change because it can restore phoneme sequences that are regularly lost. A classic example is that of the loss of *w* in the English sequence *swo*: in *sword*, for instance, the *w* is still spelled, but no longer pronounced. The same sound change must have affected *swore*. That *w* is still pronounced in this word today is because it was analogically restored: *bear* (present) : *bore* (past) = *swear* (present) : *swore* (past). Thus, analogy may create words that seemingly violate sound laws. In order to establish the phonetic conditioning of a certain sound change, one must identify exceptions that are due to analogy, as they cannot reveal anything about the sound law itself.

Unlike sound change, analogy does not operate regularly. While some words may be created anew through the application of a certain morphological rule, other irregular forms may well survive, being stored as separate entries in the mental lexicon. It is not the case, however, that anything goes in analogy. To invoke analogy to explain a certain form that seems to contradict a sound law, it must be shown that the supposed analogical change was, in fact, possible, and preferably not too improbable. This raises the question: what limitations are there on the operation of analogy?

First of all, there is the common distinction between *analogical leveling* and *analogical extension*. The difference concerns the alternation of forms of the same morpheme (*allomorphs*), which can either be given up, i.e. leveled, or extended to new cases (Fertig 2013: 48–51). For an example of leveling, compare Old English *frēosan* 'to freeze', *froren* 'frozen' to the Modern English forms, which show the *z* of *freeze* leveled throughout. A common example of extension is that of the verb *to dive*, which originally had the past tense *dived*, as still in most forms of British English; in American English, however, a pattern of stem alternation has been extended: *ride* (present) : *rode* (past) = *dive* (present) : *dove* (past). It is often maintained that leveling and extension are two separate processes, the former being motivated by speakers' preference for non-alternating morphemes (Fertig 2013: 71–76). This does not seem to be the case, however: leveling is

simply the extension of a non-alternating pattern to previously alternating forms. This is convincingly argued by Garrett (2008), who shows that in the reasonably long attested history of English, leveling of present-tense forms to the past tense only ever took place if an analogical model was available. In the case of *to freeze*, for instance, the leveling may be represented as *cleave* (infinitive) : *cloven* (past participle) = *freeze* (infinitive) : *frozen* (past participle). If alternations such as *cleave* : *cloven* had not existed, leveling would necessarily have been based on another model, e.g. *walk* : *walked* = *freeze* : *freezed*; the *z* from the infinitive could not have been leveled to the past participle while leaving the rest of the stem unchanged. No leveling whatsoever could have taken place if there had been no non-alternating forms to serve as a model. Leveling and extension, then, are just subtypes of one and the same process of analogical change.

The next issue is which forms are likely to be targeted by analogical change. Why is an alternation leveled in some cases, but extended in others? And if it is leveled, what determines which form of the morpheme is replaced and which one survives? Frequency is widely acknowledged as the most important factor, but frequency alone is not enough to explain all the observed workings of analogy. Consequently, many other contributing factors have been suggested, such as optimization or simplification of the grammar or certain universal preferences (Fertig 2013: 102–121). The issue is controversial, but the framework proposed by Joan Bybee (1985) will be adopted in this work for its explanatory power and because it lacks the conceptual problems of some other accounts.

Bybee sees *autonomy* as the most important factor in analogy. A word's autonomy determines the chance that it is stored as a separate entry in the mental lexicon, words with higher autonomy being more firmly entrenched. This is determined by the combination of the word's frequency, (semantic) basicness, and morphophonemic irregularity. 'Frequency' here refers to token frequency, how often the individual word is used, regardless of other, related forms. Basicness is a problematic concept in itself, but for present purposes, we may say that a word or morpheme is more basic if it has a broader distribution. *a* is thus more basic than *b* if *a* may occur in more contexts than *b*. In the English verb *to have*, for instance, the form *have* is more basic than *has*, as the latter may only occur in the present tense third person singular, while the former occurs in all other forms of the present, including the infinitive and imperative. The opposite of basicness in this sense is *markedness*: thus, *has* is more marked than *have*. Unlike some other theories of analogy (e.g. Albright 2008), that of Bybee allows for

local markedness (Tiersma 1982): while a certain form of a paradigm may be the unmarked one in most cases, it may be marked in some paradigms, based on semantics. Thus, the singular forms of nouns is normally unmarked, but in words that refer to objects which normally occurs in pairs or groups, like *legs* or *teeth*, the plural may be the unmarked form, making the singular marked.[4] The third constituent factor of autonomy identified by Bybee, morphophonemic irregularity, seems unnecessary. It is unclear how this would contribute to a word's separate storage in the lexicon. Bybee states that "[i]f a word is so irregular that it cannot be derived from any other related words, even if it is semantically marked, it will have to be autonomous" (p. 58), but the obvious alternative is that an insufficiently autonomous irregular word will be replaced by a newly created analogical form.

Should a speaker want to use an insufficiently autonomous form, then, he or she can create it by applying a morphological rule to a more autonomous form, if this is accessible. Bybee argues (pp. 132–134) that the most productive rules are based on the relationships between medium-frequency words. High-frequency words will likely all be individually stored in the lexicon with high lexical strength, which obscures the morphological relationship within the paradigm. Low-frequency words may not have enough lexical strength to cause the speaker to extend a rule that applies to them. In medium-frequency paradigms, however, the speaker probably has separate mental representations of the various parts of the paradigm (e.g. *walk* and *walked*), but he or she is also aware of the morphological relationship between them (*walked* = *walk* + *ed*). Also, as there are more medium-frequency words than high-frequency words, morphological rules that apply to the former will have a higher type frequency than those that apply to the latter. These patterns are thus more likely to be extended.

In summary, analogy is the creation of a new word by applying morphological rules to an existing word. It is most likely to target words with low autonomy, i.e. low frequency and/or high markedness, as these have the lowest chance of being stored in the mental lexicon. If the analogical creation differs from the earlier form of the word, this results in an analogical change. Analogy can only extend patterns which are already present in the language, and the rules that apply to medium-frequency words are most likely to spread to new paradigms.

[4]Note that all words that maintain an umlauting plural in Modern English belong to this category, which explains their resistance to regular plural formation: besides *teeth*, we find *feet*, *geese*, *mice*, *lice*, *men*, and *women* (Fertig 2013: 109).

This narrowly defined concept of analogy is to be distinguished from other categories of morphological change. Like analogy, these processes involve formal changes prompted by semantic similarity with other words, but they work differently and produce different results. The most important types of non-analogical morphological change are folk etymology and contamination (Fertig 2013: 57–70). Both involve the partial adaptation of a morpheme based on confusion with another word.

Folk etymology occurs when a listener mistakenly identifies an existing morpheme as part of a word that did not originally contain it, often because the adapted word is no longer transparent. Thus, Old English *brȳdguma*, literally 'bride man', should yield Modern English ***bridegum*, but folk etymology replaced the second element with a more familiar word, resulting in *bridegroom*. Few cases of folk etymology have been identified in Hebrew, so this concept will not feature prominently in the rest of this work; one likely example is *ṣalmo̊wɛṯ* 'darkness', originally **ṣalmuṯ* but folk-etymologically influenced by *mo̊wɛṯ* 'death'.

Contamination may be distinguished from folk etymology by its effects. While the latter extends entire morphemes to new contexts, contamination can also merely copy one or more phonemes from a related word. Contamination is especially frequent between words with a strong semantic connection, such as synonyms, like Anglo-French *citezein* 'citizen' from earlier *citeien* through contamination with *denizen*; antonyms, like *covert*, originally just an alternate spelling for *covered* but now with a changed pronunciation due to contamination with *overt*; or words that often occur together, like numerals or pronouns. Whereas folk etymology is probably due to misparsing on the listener's behalf, contamination is caused by speech errors. While the speaker prepares to say a certain word, semantically related words also suggest themselves, and in cases of contamination, influence the produced form. Notably, contamination involves the alteration of a previously existing form, while analogy creates completely new forms. This and other crucial differences between contamination and analogy are listed in table 1.1. Bearing these distinctions in mind will allow us to be more rigorous in the identification of morphological change in the following investigation. Most importantly, analogical change should only be posited if the morphological rule in question is already present in the language, and contamination should only be posited if a contaminating form with a close enough semantic association to the affected word can be found.

Table 1.1: Analogy vs. contamination

Analogy	Contamination
by-product of normal language use	result of speech errors
based on morphological similarity	based on semantic association
needs pre-existing model	does not need model
creates new forms	affects existing forms
creates morphologically regular forms	makes forms more similar

We have now surveyed what have traditionally been considered the most important factors affecting a language's phonological development: regular sound change, borrowing, and analogy and other morphological processes. One issue remains, however: taking all of the above into account, there still seem to be some cases of truly irregular sound change. Some of these were already known to the Neogrammarians while others were discovered by twentieth-century scholarship. Both categories will be discussed in the following section.

1.1.3 Irregular sound change

In many cases, sound changes involving assimilation, dissimilation, metathesis and deletion only affect a few words in a certain language, sometimes only one, although they may also occur as regular sound changes (Paul 1920: 63-66). The metathesis seen in the English change from *bridd* to *bird*, for example, was irregular, as shown by its non-operation in *brim*, *bridge*, *brick*, etc. Especially metathesis, haplology (a subtype of deletion), and dissimilation tend to occur in this haphazard fashion. These are then known as *sporadic* sound changes. Interestingly, these sporadic sound changes only have phonemic effects: no cases of sporadic allophony are known, in which a phoneme acquires a new pronunciation in one particular word. Rather, sporadic assimilation and dissimilation always yield previously existing phonemes, while metathesis and deletion necessarily affect phonemes rather than phonetic realizations. Paul (1920: 67) explains these sporadic sound changes as resulting from speech errors. Like sporadic sound changes, attested speech errors only produce phonologically valid forms, and all types of sporadic sound change occur as speech errors (Wells 1951). We may thus safely attribute

sporadic sound change to errors in production, which are then accepted by new generations of speakers.

A more serious challenge to the Neogrammarian Hypothesis is that of *lexical diffusion*. A review of the literature, some convincing examples, and an attempt to reconcile this process with Neogrammarian sound change are given by William Labov (1981). In brief, lexical diffusion refers to observed instances of a sound change apparently spreading from word to word, unpredictably affecting more and more cases. In the English spoken in Philadelphia, for example, /æ/ has shifted to /eːə/ in a number of phonetically defined categories. The change seems to have 'diffused' beyond its original conditioning, however: while the shift does not normally take place before voiced plosives, leaving words like *sad* unaffected, it does occur in *bad*, *mad*, and *glad*. Similarly arbitrary occurrences of this sound change are found in different parts of the lexicon.

This seems like damning evidence against the purely phonetic conditioning of sound change. Labov goes so far as to state that "we have arrived at a situation where no reasonable person can maintain what might be called the Neogrammarian dogma: that sound change is always gradual, always regular, affecting all words at the same time" (p. 271). As Neogrammarian, purely phonetic conditioned sound change clearly has taken place in very many cases, though, we cannot simply throw out the concept of regular sound change altogether. In reaction to the discovery of lexical diffusion, then, scholars have either classified it as a different kind of sound change, to be distinguished from the regular, Neogrammarian kind, or denied that it is a form of sound change in the first place.

Labov himself chooses the first option. Neogrammarian sound change is phonetically gradual and regular, gradually changing all instances of the same phoneme in a certain environment. Lexical diffusion, on the other hand, is phonetically abrupt and unpredictable, changing one phoneme to another in some words, but not in others. To my knowledge, he does not explain how this irregular sound change comes to affect new cases. Lexical diffusion as a process thus remains unexplained.

Recently, though, Bybee (2013) has shed some new light on the matter. She identifies two separate subtypes of lexical diffusion. Some cases of lexical diffusion affect high-frequency words first, then spread to low-frequency words; in other cases, vice versa. Interestingly, this difference in diffusional direction matches a difference in the phonetic change that is diffused. Changes that spread from high-frequency words to low-frequency words tend to be phonetically gradual

and involve reduction, resulting in less articulatory effort, as in the deletion of word-final /t/ in English. Changes that spread from low-frequency words, however, may also be phonetically abrupt, as in the English change of /θ/ to /f/. According to Bybee, lexical diffusion from high-frequency words is simply what phonetic change ('sound change' in her terms) in action looks like. More frequent words have more progressive phonetic realizations, but eventually, these realizations will regularly spread throughout the lexicon. Lexical diffusion from low-frequency words ('phonological change'), on the other hand, is something else. Observing the similarities with analogical change (also noted by Kiparsky 1995), she states that "change affecting low frequency words first indicates that the form of such words presents a challenge to the listener or learner with the result that such words are remade on more familiar patterns" (p. 225). While this is obviously different from analogy proper, the process is reminiscent of another form of morphological change: folk etymology. Hearing an infrequent word with a /θ/, a learner may mistakenly identify the phoneme as /f/, which is itself more frequent than /θ/. Interestingly, this then decreases the relative frequency of /θ/ to /f/, making future changes of this type more likely. Thus, the change may spread from low-frequency words to more frequent ones, eventually affecting all eligible cases and showing the same results as a regular sound change. Alternatively, a few highly frequent words may resist the change indefinitely or perhaps merge with another phoneme, resulting in apparent counterexamples to a sound law.

Thus, it seems that lexical diffusion does not contradict the Neogrammarian Hypothesis: diffusion from high-frequency words is regular sound change in action, while diffusion from low-frequency words is not a form of sound change at all.

1.1.4 Summary

We have encountered a number of different types of change that may affect the phonetic realizations of a word, some phonological in origin, some morphological. All of these changes result from errors, either errors in production or errors in perception.

Phonetic change results from the incorrect acquisition of phonetic rules governing the production of phonemes. It is phonetically regular, as the acquired rules always affect the same phoneme in the same environment in the same way.

Phonemic change results from the incorrect inference of phonemic forms. It is also phonetically regular, as the learner infers the same underlying forms given

the same phonetic surface forms, and the latter are themselves regularly based on the original phonemic representations.

Analogy results from the incorrect (or rather, novel) application of a previously existing morphological rule. It is irregular and tends to create less frequent, more marked words, based on more frequent, less marked parts of the same paradigm.

When speech errors are caused by interference with a semantically related word, they may result in contamination. If there is no semantic motivation for the speech error, the result is sporadic sound change. The former is more frequent, as the existence of a semantically close word will skew the speech errors in a certain word in one direction, making the two words more similar.

Similarly, the incorrect identification of a more frequent morpheme than the one intended results in folk etymology, while the incorrect identification of a more frequent phoneme than the one intended results in lexical diffusion.[5] If the explanation for the latter given above is correct, it should target less frequent words first.

1.2 Biblical Hebrew

Biblical Hebrew is, unsurprisingly, the Hebrew of the Bible. More specifically, it is the language of the Hebrew Bible (minus the Aramaic sections)[6], largely corresponding to the Christian Old Testament. This consists of the Five Books of Moses or Pentateuch (Genesis, Exodus, Leviticus, Numbers, and Deuteronomy), the Prophets (Joshua, Judges, 1 and 2 Samuel, 1 and 2 Kings, Isaiah, Jeremiah, Ezekiel, and the Twelve Minor Prophets), and the Writings (Psalms, Job, Proverbs, Ruth, Song of Songs, Ecclesiastes, Lamentations, Esther, Daniel, Ezra–Nehemiah and 1 and 2 Chronicles), together known as *tanak̲*, an acronym of the Hebrew names of these three sections (*torå*, *nb̲iʾim*, and *k̲tub̲im*).

One of the reasons Biblical Hebrew is such an interesting object for linguistic research is its composite status. The texts of the Hebrew Bible were composed between the late second and late first millennium BCE in the literary Hebrew of the time (Gzella forthcoming a). Due to the nature of the Hebrew writing system, these texts were originally almost purely consonantal, leaving it to the

[5]This last term is a bit of a misnomer to describe a single instance of this type of change, as it refers to the spread of the process, not its initiation. Perhaps *folk phonology* would be more adequate.

[6]Namely, Jer 10:11, Dan 2:4b–7:28, and Ezra 4:8–6:18, 7:12–26.

reader to insert the correct vowels based on his knowledge of Hebrew. After the death of Hebrew as a spoken language, probably in the second century CE,[7] the correct pronunciation of the canonized texts was no longer evident to readers. Rather, it had to be learned, word for word, verse for verse. The reading of the Hebrew Bible was no longer based on a living language, but on a received tradition (Hebrew: *måsorå* or *måsórɛt̲*); the experts in this tradition are known in English as the *Masoretes*. From the sixth or seventh century onward, the Masoretes began to experiment with written vowel signs, as well as accents to indicate the melody of the recitation, known as cantillation signs (Sáenz-Badillos 1993: 77). Different Masoretic systems of vocalization and accentuation were developed, and eventually, that used in Tiberias, on the Sea of Galilee, became widely accepted. After the tenth century, the text of the Hebrew Bible vocalized according to this system, known as the *Masoretic Text*, came to be used as the authoritative version of the Hebrew Bible throughout the Jewish world. The Masoretic Text, particularly the version preserved in the authoritative eleventh-century CE manuscript known as the Leningrad Codex, is also the version of the Hebrew Bible that has featured most prominently in modern scholarship, and its language, which we may more precisely call Tiberian Biblical Hebrew, is the main subject of investigation in this work.

The Masoretic Text is not the only version of the Hebrew Bible to have been preserved. Although the non-Tiberian vocalizations have fallen into disuse, Biblical fragments using these systems have been recovered. The most important alternative vocalization systems are the Palestinian and Babylonian traditions (Sáenz-Badillos 1993: 86–104). Both distinguish a smaller number of vowels than the Tiberian tradition: the latter has seven different vowel qualities, while the Babylonian tradition has six; the Palestinian vocalization has signs corresponding to seven of the Tiberian vowels, but the fluctuation in their usage suggests the Palestinian reading tradition only had five qualitatively differentiated vowels (Sáenz-Badillos 1993: 88–89). Moreover, the non-Tiberian traditions differ from Tiberian Biblical Hebrew on some morphological points.

An older source of information on Biblical Hebrew is found in Greek and Latin transcriptions and translations of the Hebrew Bible (Sáenz-Badillos 1993: 80–86). Mainly personal and place names are attested in the Septuagint, the oldest Greek

[7] But see Gzella (2015: 226–229), who argues for the complete replacement of Hebrew as a spoken language by Aramaic in the Hellenistic period, citing proponents of both the earlier and the later date.

Bible translation. Different books of the Septuagint were translated by different people at different times, but the oldest layer, the Pentateuch, dates to the third century BCE. Hebrew names also occur in the Vulgate, the Latin Bible version made by Jerome in the fourth century CE, who also left us a collection of notes on the Hebrew language. The most important source, however, is the Second Column (Secunda) of the Hexapla ('sixfold'), a six-column edition of the Bible compiled in the third century CE by Origen. Besides the Hebrew text and a number of Greek translations, the Hexapla contains a transcription of the Hebrew text in Greek letters, providing evidence for the vocalism of common nouns, verbs, pronouns and particles. The textually most reliable Hexaplaric transcriptions are found in a number of fragments from Psalms, a linguistic analysis of which can be found in Brønno (1943).

A rather different version of the Pentateuch alone is preserved by the Samaritans. The Samaritan Pentateuch differs from the Masoretic Text in many minor and some major respects, as far as the consonantal text is concerned. Samaritan Pentateuch manuscripts are typically unvocalized, but the Samaritan reading tradition is still passed on from father to son to this day and has been comprehensively described by Ze'ev Ben-Ḥayyim (1955–1977, 2000). The language it reflects, Samaritan Hebrew, is quite distinct from Jewish Biblical Hebrew, and could therefore provide highly valuable insights into their shared ancestor, presumably spoken Hebrew. The authenticity and reliability of the Samaritan reading tradition, however, is far from certain (Macuch 1969: 86–87). As it clearly reflects secondary, artificial forms in a fair number of cases, we cannot simply take any Samaritan Hebrew data at face value. For this reason, the evidence from the Samaritan reading tradition will play a relatively minor role in the following investigation.

Besides a few cuneiform transcriptions of names, the earlier Hebrew corpus consists of epigraphic material. This ranges from brief texts and the occasional inscription from the monarchic period, starting in the eighth century BCE, to the earliest attested Bible texts found together with original compositions at Qumran and in the greater Dead Sea area (second century BCE to second century CE, Sáenz-Badillos 1993: 130–146). These texts can only tell us little about the vocalism of the language at the time of their composition, but they do contain some relevant evidence.

This investigation will thus be concentrated on Tiberian Biblical Hebrew, relying on other forms of Hebrew to supply additional evidence, where relevant. This raises a crucial question. If the Masoretic Text is vocalized according to a system

that was only developed eight centuries after the death of the Hebrew language, how can we be sure of its authenticity? In the wake of the pioneering work on non-Tiberian forms of Hebrew done by Paul Kahle in the early twentieth century, many scholars were highly suspicious of the Masoretic vocalization. In their view, it did not accurately reflect any authentic stage of Hebrew, but was an artificial creation. This view is to be rejected, for two main reasons. First of all, there is the systematic correspondence between the Tiberian vocalization and that reflected in the Second Column of the Hexapla, which bears witness to a reading tradition that was much closer in time to the last phases of spoken Hebrew. This shows that if the vocalization of Biblical Hebrew is an invention, it cannot be a late one. The authentic character of the Masoretic reading tradition is established beyond doubt, however, by the regular correspondence between Tiberian Biblical Hebrew words and their cognates in other Semitic languages. In the words of Aron Dolgopolsky (1999: 22-23):

> The very fact that there are regular phonetic correspondences between Masoretic Hebrew and the cognate Semitic languages (including those not known to the Masoretes as Ugar[itic], Akk[adian], Old South Arabian, Mehri, etc.) and that the [Masoretic Hebrew] phonology can be accounted for by a series of regular sound changes from proto-Semitic (reconstructed on the basis of other Semitic languages) proves that it cannot be artificial. We may stress after Bergsträsser that if the Masoretes did indeed bring about changes attributed to them by Kahle and some other colleagues, they must have been trained as modern Semitic scholars! If we believe Garbini suggesting that the distinction between [*š*] and [*ś*] is an artificial invention of the Masoretes (. . .), how can we explain the exact correspondence between this Masoretic Hebrew distinction and that found in Eth[iopic] between [*s*] and [*ś*], [Old South Arabian] between [s^2] and [s^1] [*sic*], Mehri and Soqotri between *š/h* and *ś*, as well as that between [*s*] and [*š*] in Arabic? In my opinion, the [Masoretic Hebrew] vocalization and pronunciation are based on living tradition of Hebrew (up to the 2nd century C.E.) with subsequent changes according to the laws of phonetic development of the Aramaic dialect(s) spoken by Jews (3rd through 9th century C.E.). This does not exclude the possibility that at different periods (including that of the Masoretes) there were non-systematic individual changes of certain words based on re-etymologization, on tabuistic, euphemistic grounds (just as in any other language), or substitution of high-style variants of forms (. . .) for those of the neutral style (. . .). . . . But taken all this into account, we can nevertheless agree with G. R. Driver: "Although

> the Masoretic vocalization might sometimes be wrong, internal reasons as well as the analogy to the cognate languages testify to the general faithfulness of the tradition" (...).

This is also the position adopted in this work.

1.3 Proto-Northwest-Semitic

The chronological end point of an investigation into the the development of the Biblical Hebrew vowels thus quite naturally falls at the codification of the Masoretic Text, which gives us the most detailed information on Hebrew vocalism available, based on a reliable tradition. But where to start? As no direct ancestor of Biblical Hebrew is attested in vocalized script, any account of its historical phonology must rely on a reconstructed proto-language. In this work, we will limit ourselves to tracing the development of the Hebrew vowel system from that of Proto-Northwest-Semitic, the unattested, last common ancestor of the Northwest Semitic subfamily, which includes Hebrew.

Perhaps unexpectedly, it is easier to get a clear picture of Proto-Northwest-Semitic than of any later ancestor of Hebrew, as the discussion of the Semitic family tree below will show. This is due to the broad attestation of the various Northwest Semitic languages, including some very early cases like Ugaritic and Amarna Canaanite. Together with closely related languages outside the Northwest Semitic subgroup, most importantly Classical Arabic, these allow us to reconstruct Proto-Northwest-Semitic grammar and lexicon with a fairly high level of confidence. Older ancestors of Hebrew, such as Proto-West-Semitic and Proto-Semitic itself, are harder to reconstruct, partially due to the controversy over the genetic subclassification of the Semitic languages; one might say that Proto-Northwest-Semitic is Hebrew's oldest (relatively) uncontroversial ancestor. At the same time, the reconstructed Proto-Northwest-Semitic vowel system is very close to that of Proto-Semitic, so hardly any interesting developments are left out by limiting the scope of investigation to the Proto-Northwest-Semitic–Biblical Hebrew time frame.

The Northwest Semitic languages, spoken in the Levant and Mesopotamia, form a clearly related subgroup of the Semitic family.[8] A number of their shared features will be discussed in chapter 2. The family consists of three main branches, as

[8] See Gzella (2014) for a more detailed description of the social and cultural context of these languages and their neighbours.

well as some minor languages whose classification is debated. Hebrew belongs to the Canaanite subgroup, together with Phoenician, once spoken in modern-day Lebanon, surrounding areas, and colonies, and its later, Carthaginian stage, known as Punic. The Transjordanian languages, i.e. Moabite, Ammonite, and Edomite, are also usually considered to be Canaanite; the latter two are poorly attested, while there is a long inscription in Moabite, as well as some shorter texts. Additionally, there is the invaluable resource of Amarna Canaanite, the collection of Canaanite linguistic features found in the letters sent to the Egyptian court by Canaanite vassal states in the fourteenth century BCE. These letters are written in Akkadian cuneiform, and at first sight, their language looks like Akkadian with some Canaanite elements incorporated; whether this actually reflects a spoken, mixed Canaano-Akkadian language variety (e.g. Izre'el 2012) or an adaptation of the Akkadian writing system to express spoken Canaanite (e.g. Von Dassow 2004) is a moot point.

The historically most important branch of Northwest Semitic is Aramaic, attested from the tenth century BCE onwards and still spoken today. Originating in Syria, Aramaic gradually replaced Akkadian as the dominant language of Mesopotamia, and for about a millennium after the fifth century BCE, it was used as a lingua franca throughout the Near East. Many different varieties of Aramaic are attested. For present purposes, the most important ones are Old Aramaic, mainly preserved in inscriptions; Imperial Aramaic (sometimes grouped under Old Aramaic), the administrative language of the Achaemenid Empire, mainly preserved in letters and other documents written on papyrus; and later, vocalized dialects like Biblical Aramaic, preserved in the Hebrew Bible, and Classical Syriac, the classical language of Middle-Eastern Christianity. Aramaic was the language that exerted the most influence on Hebrew throughout its history, so we will encounter many Aramaic loanwords, as well as some possible cases of contact-induced change.

Scholars have tried to group every other attested Northwest Semitic language variety with either Canaanite or Aramaic, but the evidence shows that Ugaritic, at least, should be considered a separate branch of the Northwest Semitic family tree (Noorlander 2015). Attested in alphabetic and syllabic cuneiform texts from the thirteenth and twelfth centuries BCE found on the Syrian coast, Ugaritic is the oldest West Semitic language in which running texts have been found, providing an important tool in the reconstruction of Proto-Northwest-Semitic. Besides Canaanite, Aramaic, and Ugaritic, several other Northwest Semitic language varieties have been identified, notably Amorite, attested in Mesopotamian names

from the early second millennium BCE; Samʾalian, attested in several inscriptions from the ninth and eighth centuries BCE and seemingly close to Aramaic, if not actually belonging to that family itself (Noorlander 2012); the language of the ninth or eighth century BCE Deir ʿAlla inscription, which shares innovative features with both Aramaic and Canaanite (Gzella 2013a); and possibly Taymanitic, attested in the mid-first-millennium BCE inscriptions from the oasis town of Tayma (Kootstra 2016).

Understandably, the reconstruction of Proto-Northwest-Semitic is mainly based on these Northwest Semitic languages. A number of other Semitic languages will also be referred to in the following chapters, most notably Akkadian, Classical Arabic, and Classical Ethiopic or Gəʿəz. Akkadian was the first Semitic language to reach Mesopotamia. It is subdivided into Old Akkadian, attested in the third millennium BCE, and various periods of the Assyrian (northern) and Babylonian (southern) dialects, attested from the second millennium BCE until the beginning of the Common Era. Akkadian was written in a syllabic cuneiform script, adopted from the Sumerians. Classical Arabic, written with an alphabetic script derived from that of Aramaic, is much later (starting in the seventh century CE), but phonologically one of the most conservative languages of the Semitic family. The history of Arabic is a hotly debated issue in current scholarship, and it is unclear to what degree Classical Arabic reflects any single authentic language variety, which calls for some caution in the use of Arabic comparative evidence. Besides the classical, literary language, a large variety of Arabic dialects are spoken throughout the Middle East and beyond. Gəʿəz, finally, has a similar status as a classical, literary language, used in the cultural sphere of modern Ethiopia from the fourth century CE onwards. As the Ethiopic writing system developed the ability to express vowels early on, there is some certainty that texts from the Aksumite period (fourth to seventh centuries CE) accurately reflect the phonology of the spoken language of the time, not secondary reading traditions. Modern Semitic languages of Ethiopia and Eritrea, such as Amharic and Tigriña, are not directly descended from Gəʿəz, but closely related.

The genetic classification of these and other Semitic languages is debated. It is clear that the Akkadian dialects and another ancient Mesopotamian language, Eblaitic, were the first to split off from Proto-Semitic; the former are thus classified as East Semitic, all other languages as West Semitic. The issue is the internal structure of West Semitic, the controversy being centered around the position of Arabic. Traditionally, Arabic was grouped together with Ethiosemitic (Gəʿəz and

its close relatives) and the Ancient South Arabian and Modern South Arabian languages of Yemen and Oman, forming a subgroup known as South Semitic. Based on morphological criteria, however, an alternative family tree was put forward by Robert Hetzron (1976). In the modified version recently argued for by Huehnergard & Rubin (2011), it connects Arabic and Ancient South Arabian to Northwest Semitic, forming a group known as Central Semitic. Central Semitic, Ethiosemitic, and Modern South Arabian are then three separate branches of West Semitic. This classification hinges on the supposed Central Semitic innovation of a new imperfect, *yaqṭulu, replacing the inherited form of the imperfect, *yaqaṭṭal, which is still reflected in all non-Central-Semitic branches. While the evidence for the Central Semitic hypothesis is not without its problems (Suchard 2015), this is the classification that will be adopted in this work.

1.4 Some previous approaches to the question

Having thus explored the research question, we will now briefly survey the most influential literature to have previously addressed this problem before moving on to the methodological aspects of the current study. We shall discuss each author's aims, his methodology, and assumptions about the Hebrew language, in order to be able to properly assess his proposed solutions later on.

1.4.1 Brockelmann (1908)

Carl Brockelmann's *Grundriss der vergleichenden Grammatik der semitischen Sprachen* can fairly be considered the foundation of modern comparative Semitic linguistics. Even now, more than a century since its publication, it remains an incredible storehouse of data on all the major Semitic languages that had been studied at the time. While the Semitic languages have been subject to comparative investigation longer than any other family (Blau 2010: 13), Brockelmann decided to apply the relatively new historical–comparative linguistic methodology to the field of Semitics. He states that the aim of his *Grundriss* is to gather and organize the available material, in order "to clear the way for the solution of the many questions that are still unanswered in this area" (p. vi).[9] This orientation towards

[9] *Der Lösung der zahlreichen auf diesem Gebiet noch offenen Fragen den Weg zu bahnen ist der vorliegende Grundriß bestimmt.*

the linguistic methods of the time is felt throughout the work, and references to Indo-European linguistics are frequent.

Together with morphology, the phonology of Proto-Semitic and its daughters is treated in the first volume of the *Grundriss* (the second volume being concerned with syntax and therefore less relevant to the present work). After stressing that he considers Proto-Semitic reconstructions to be nothing more than formulas expressing the correspondences between cognates in the various Semitic languages (pp. 4–5), Brockelmann goes on to list a great number of words and forms, providing sound laws to explain them from his reconstruction of their shared ancestor. Unlike some other scholars, he reconstructs a quite consistently triradical version of Proto-Semitic, rather than allowing for biradical roots, corresponding to various classes of weak verbs in later Semitic.

In Brockelmann's view, Hebrew is essentially a form of Canaanite (p. 8), meaning that it does not crucially differ from related dialects like Phoenician and Moabite. From Hellenistic times onward, however, this Canaanite language faced an ever-growing influence from Aramaic (pp. 9–10), resulting in the increase of non-classical features in late Biblical and post-Biblical texts. Nevertheless, he considers the Masoretic tradition to be "relatively faithful" (*relativ treu*, p. 11), and relies on the Masoretic text as his source for Biblical Hebrew data.

1.4.2 Bergsträsser (1918)

While Brockelmann (1908) was concerned with all the Semitic languages known at the time, Gotthelf Bergsträsser's 1918 *Hebräische Grammatik* was nearly equally influential in the more specific field of Hebrew linguistics. It is strongly based on that of Kautzsch (1909), which is itself the reworked, 28th edition of Wilhelm Gesenius' classic grammar. Bergsträsser (1918) thus represents an old tradition, and this grammar with a pedigree is the standard reference work for many later authors. The methodology is similar to Brockelmann's: Bergsträsser, too, appeals to the new concepts developed by the comparative linguists working on Indo-European (pp. 82ff.)

Bergsträsser does not explicitly state the aims of his grammar, but it seems to be meant as a very historically oriented synchronic grammar of Biblical Hebrew. In contrast to Brockelmann (1908)'s immensely wide scope, the *Hebräische Grammatik* gives a more in-depth treatment of its narrower subject matter, describing all the minutiae of Hebrew grammar. The author offers possible explanations of

most features of the language, although he sometimes concedes that a certain phenomenon cannot be explained satisfactorily (e.g. on p. 144).

Unlike Brockelmann, Bergsträsser believes that the Masoretic vocalization is "very *unreliable*" (*sehr wenig zuverlässig*, p. 24; emphasis in original). It is not the reliable result of a pronunciation handed down from generation to generation, but shows a great amount of Aramaic influence and many secondarily constructed forms. The author occasionally uses this unreliability of the vocalisation to explain some otherwise incomprehensible peculiarity of the text. Another feature of this work that is worth mentioning is its complicated account of the prehistory of the Hebrew stress system (pp. 115–116), which is most relevant to chapter 4 of the present work. As this reconstruction of the stress system is fairly ad hoc, as we shall see, it is to be dismissed.

1.4.3 Bauer & Leander (1922)

Hans Bauer and Pontus Leander's *Historische Grammatik der hebräischen Sprache des Alten Testamentes* was the first real historical grammar of the Hebrew language. From a modern point of view, it is a strangely mixed piece of work. On the one hand, it is, like the works of Brockelmann and Bergsträsser before it, a treasure trove of data, and as the authors aim to take a completely comparative linguistic look at Hebrew (p. iii), it offers the reader some invaluable insights into the language's prehistory. On the other hand, it is quite seriously dated and partially rendered obsolete by its incorporation of Bauer's concept of Hebrew as a mixed language.

According to this theory (explained on pp. 16ff.), the notion that Hebrew can be placed in a neat family tree of Semitic, grouped first with Canaanite, then with the other Northwest Semitic languages, and so on, is absolutely false. The authors argue that Hebrew uniquely shares features not only with Aramaic, Arabic and the other West Semitic languages, but also with Akkadian (pp. 7–8). Based on this fact, they reject the East–West split that is generally posited for the Semitic family tree, preferring a distinction between an 'old' group, consisting of Akkadian and the original layer of Hebrew, and a 'new' group, consisting of everything else. The uniquely mixed status of Hebrew, then, was caused by the intrusion of a 'new' speech layer into the 'old' layer that was more like Akkadian. Thus, the authors explain the seemingly irregular operation of the Canaanite Shift (*ā > *ō, see chapter 3), a feature of the 'old' layer in their view, by assigning any forms

that did not undergo it to the 'new' layer; e.g. *måqom* 'place' < *maqāmum (cf. Classical Arabic *maqāmun*) is from the 'old' layer, while *qåm* 'he stood' < *qāma (cf. Classical Arabic *qāma*) is 'young'. This mixed nature of Hebrew need not surprise us, since it is the result of the 'young'-speaking Israelites partially imposing their language on the 'old'-speaking inhabitants of Canaan after conquering them, in accordance with the history recorded in the Hebrew Bible. In fact, though, there is no real reason to suspect that Hebrew is in any way a mixed language (as was already shown by Bergsträsser 1923), and today, Bauer's theory is rejected by the great majority of scholars.

Apart from this assumption about the origins of the Hebrew language, however, Bauer & Leander (1922) is a very useful work. The grammar is not only based on the biblical text, which the authors consider to be written in the highly mixed dialect of Jerusalem (p. 32), but also refers to Amarna Canaanite, the Transjordanian languages, non-Tiberian vocalizations of Biblical Hebrew, etc. Another interesting feature is the authors' reconstruction of Proto-Semitic, which is quite different from that of Brockelmann, allowing for biradical roots and deriving most verbal forms from the imperative, which the authors see as the most basic form of the verb.

1.4.4 Birkeland (1940)

We have seen that the developments in general linguistics influenced the methodology of various works on Semitics and Hebrew, and as the discipline evolved further, so this influence changed as well. The emergence of the subfield of phonology inspired Harris Birkeland to take a new look at the Hebrew language (p. vii) in his *Akzent und Vokalismus im Althebräischen*. As the title suggests, this work is largely concerned with Hebrew stress and accentuation and their interplay with the development of the vowels. This specialized focus leads the author to some interesting and original ideas, which will be examined in the relevant chapters. Methodologically, Birkeland tends to invoke dialect borrowing as an explanation (e.g. on p. 14), and like Bauer & Leander (1922), he traces some developments back to the influence of the supposed original, non-Canaanite language of the invading Israelites (e.g. on p. 17), although the details of his scenario are different from Bauer's.

1.4.5 Blau (2010)

Most recently, there is the critically well-received *Phonology and Morphology of Biblical Hebrew*, written by the prominent scholar cited above, Joshua Blau. Throughout his career, Blau has published many articles on the historical development of the Hebrew language (see Bibliography), and this translation of an earlier work in Modern Hebrew can be seen as a collection and organization of the ideas he has put forward over the years.

As Blau considers Biblical Hebrew to be "profoundly multilayered" (p. 5), he aims for a "diachronic–comparative approach" (ibid.), making this something of a strongly historically-minded synchronic grammar, similar to Bergsträsser (1918). In another parallel to works of the early twentieth century, Blau starts out with a general introduction to historical linguistics, which in Blau's case is more directly related to the Hebrew data. On the whole, he prefers to draw on the Hebrew Bible itself as his main source (p. 5). Like Birkeland, Blau assigns an important role to the workings of the various stress systems in the period leading up to Biblical Hebrew.

1.5 Assumptions and methodology

It will be clear from the previous section that there is no shortage of literature on the historical phonology of Biblical Hebrew. Nevertheless, a new attempt to describe the development of the Hebrew vowel system may well be fruitful. As we have seen above, the present work differs from most previous monographs on the topic in its assumptions. Most importantly, unless the evidence clearly proves otherwise, the following investigation will assume that:

1. the vocalization of the Masoretic Text reflects an authentic form of Hebrew (contra Bergsträsser 1918);

2. Hebrew is not a mixed language, so sound laws should hold across all parts of the lexicon except for loanwords (contra Bauer & Leander 1922, Birkeland 1940);

3. the changes from Proto-Northwest-Semitic to Hebrew should match one of the types of natural phonological and morphological change identified in section 1.1.1 (contra Blau 2010).

1 and 2 have been convincingly argued by Brønno (1950) and Bergsträsser (1923), respectively, while 3 is the issue at stake. These assumptions are also those of Brockelmann (1908); the main differences with his conclusions are expected to stem from the incorporation of new data from Semitic languages and texts discovered in the past century. While many studies of individual problems of Hebrew vocalism have been published in that time, this will be the first work since Brockelmann to approach the historical development of the Hebrew vowels as a whole under the assumptions given above.

This study will largely depend on a database of the Biblical Hebrew lexicon constructed for this purpose. In this database, every attested part of a paradigm is listed as a separate entry. Each entry consists of a transliteration of the Biblical Hebrew form, a rough translation (mainly used to keep homophones apart), and the most likely Proto-Northwest-Semitic reconstructions, with the dummy root *qṭl* replacing the radicals. This made it easy to retrieve, for example, all second person masculine singular *piʿel* perfect forms (all reconstructed as *qaṭṭilta), all words with an unstressed *a* in the first syllable, all reflexes of *qiṭlum nouns with a plosive second radical, etc. The reconstructions were very preliminary and mainly served to identify morphologically similar forms. For practical reasons, the database was based on the words listed in the 1985 reprint of Ludwig Koehler and Walter Baumgartner's *Lexicon in Veteris Testamenti Libros*, with corrections and addenda (Leiden: Brill). The spelling of some suspect words and etymological information was checked against the more accurate revised version (Koehler & Baumgartner 1994–2001). Unless indicated otherwise, the Biblical Aramaic data are also taken from Koehler & Baumgartner (1994–2001); Syriac from Sokoloff (2009); epigraphic Northwest Semitic from Hoftijzer & Jongeling (1995); Classical Arabic from Lane (1863–1893); Gəʿəz from Leslau (1987); and Akkadian from the Chicago Assyrian Dictionary (CAD).

The methodology, then, is fairly straightforward. The sound changes affecting the Hebrew vowels that have not been uncontroversially described will be treated in separate chapters, with small, related changes sometimes sharing a chapter. After a review of the literature since Brockelmann (1908), occasionally taking relevant literature from the nineteenth century into account, we will see which aspects of the sound change have adequately been explained, and what further investigation is needed. This further investigation will then collect the relevant lexical material from the database, consider how these forms should be reconstructed for Proto-Northwest-Semitic, and attempt to establish a plausible phonetic

conditioning for changes that cannot be shown to result from analogy or other types of change we have seen above.

1.6 Outline and conventions

Before the case studies of individual sound changes commence, chapter 2 discusses the reconstruction of Proto-Northwest-Semitic, gives a concise synchronic overview of Biblical Hebrew, and mentions some general changes from the former to the latter. Chapter 3 covers the Canaanite Shift, *ā > *ō. Chapter 4 discusses the development of the stress system and three associated lengthening processes: tonic, pretonic, and pausal lengthening. Chapter 5 looks at the development of postvocalic *w and *y, both in diphthongs and in triphthongs (i.e. intervocalically). Chapters 6 and 7 consider various changes of stressed *í to *á and unstressed *a to *i, respectively. Chapter 8 deals with the preservation or loss of word-final vowels. The interaction of all of these sound changes with the historical morphology of Hebrew is explored in chapter 9. Chapter 10, finally, presents the general conclusions, including a combined relative chronology of all the discussed sound changes.

The Semitic material will be presented in transcription to make this work more accessible to non-Semiticists. The transcriptions of Biblical Hebrew strictly represent the traditional pronunciation only: silent consonants are not represented, nor are so-called *matres lectionis*, consonants used to indicate the presence of a vowel. When necessary, the consonantal spelling of a word will be given in transliteration, indicated by <angled brackets>. Very occasionally, silent consonants are indicated in transcription, in which case they are written in (parentheses). The transcription and transliteration used are based on that of Lettinga (2012), chosen for its transparency and lack of macrons on vowels, which may incorrectly suggest a phonemic length contrast between vowels that are distinguished by quality, not quantity. Tables 1.2 and 1.3 give the Hebrew graphemes with their transliteration, IPA value in the Tiberian reading tradition (when not silent, in the case of the consonants), and name (in the version given in Lettinga). The phonetic realizations are taken from Khan (2013a).

beṯ, *gímɛl*, *dálɛṯ*, *kaf̄*, *pe*, *táw*, and *šin/śin* all represent two separate Biblical Hebrew phonemes. The difference between these is marked in the Tiberian vocalization by a diacritic dot. For the distribution of the allophones of *wáw*, *yoḏ*,

Table 1.2: The Hebrew consonants and their transcription

letter	transcription	IPA value	name
א	ʾ	[ʔ]	*ʾ́ālɛp̄*
ב	*b* or *b̲*	[b] or [v]	*beṯ*
ג	*g* or *ḡ*	[g] or [ʁ]	*gímɛl*
ד	*d* or *d̲*	[d] or [ð]	*dā́lɛṯ*
ה	*h*	[h]	*he*
ו	*w*	[v], [w]	*wåw*
ז	*z*	[z]	*záyin*
ח	*ḥ*	[ħ]	*ḥeṯ*
ט	*ṭ*	[tˤ]	*ṭeṯ*
י	*y*	[j], [ʝ]	*yod̲*
כ, ך	*k* or *k̲*	[k^{h}] or [χ]	*kap̄*
ל	*l*	[l]	*lā́mɛd̲*
מ, ם	*m*	[m]	*mem*
נ, ן	*n*	[n]	*nun*
ס	*s*	[s]	*sā́mɛk̲*
ע	ʿ	[ʕ]	*ʿáyin*
פ, ף	*p* or *p̄*	[p^{h}] or [f]	*pe*
צ, ץ	*ṣ*	[sˤ]	*ṣā́d̲e*
ק	*q*	[q]	*qop̄*
ר	*r*	[ʁ], [rˤ]	*reš*
ש	*š* or *ś*	[ʃ] or [s]	*šin* or *śin*
ת	*t* or *ṯ*	[t^{h}] or [θ]	*tåw*

Table 1.3: The Tiberian vowel signs and their transcription

vowel sign	transcription	IPA value	name
◌ַ	*a*	[aː] or [a]	*páṯaḥ*
◌ֶ	ɛ	[ɛː] or [ɛ]	*sḡol*
◌ֵ	*e*	[eː]	*ṣere*
◌ִ	*i*	[iː] or [i]	*ḥírɛq*
◌ָ	*å* or	[ɔː] or	*qǻmɛṣ (gåḏol)*
	ɔ	[ɔ] (rarely [ɔː])	or *qǻmɛṣ ḥåṭup̄*
◌ֹ	*o*	[oː]	*ḥólɛm*
◌ֻ, ◌ּ	*u*	[uː] or [u]	*qibbuṣ*, *šúrɛq*
◌ְ	–	– or any of the short vowels above	*šwå*
◌ֲ	*ă*	[a]	*ḥåṭep̄ páṭaḥ*
◌ֱ	*ĕ*	[ɛ]	*ḥåṭep̄ sḡol*
◌ֳ	*ŏ*	[ɔ]	*ḥåṭep̄ qǻmɛṣ*

and *reš*, see Khan (2013a). Although it will be argued in chapter 2 that Biblical Hebrew phonology does make a phonemic distinction in vowel length, length will not be indicated in transcription. It is not consistently indicated in the Tiberian vocalization, so any transcription which marks length necessarily imposes a certain interpretation on the data, which is to be avoided. The qualities of *qǻmɛṣ gåḏol* and *qǻmɛṣ ḥåṭup̄* have merged in Tiberian Hebrew, but they are distinguished in many other reading traditions, and also in Modern Hebrew and the academic pronunciation of Hebrew. As ɔ always goes back to *u and *å always goes back to *a or *ā, it may be useful to indicate the distinction in transcription; since it is based on the non-Tiberian reading traditions, rather than the conjectures of linguists, this seems more acceptable than the indication of vowel length in other cases. *šwå* indicates the absence of a phonemic vowel after a consonant. If this would result in a syllable containing two consonants in the onset, this is usually resolved by the insertion of an epenthetic vowel: [i] before /y/, a short vowel with the quality of the next one before a /ḥ ʿ h ɔ/, and [a] elsewhere. Word stress, which is phonemic in Biblical Hebrew, is indicated with an acute accent when it

does not fall on the ultimate syllable, its default position.[10] Geminated consonants are transcribed by the same consonant twice in a row, as in *rabbim*.

The transcription used for Hebrew is also used for Biblical Aramaic and other varieties of Aramaic vocalized according to the Tiberian system. The consonants of Syriac (and other Northwest Semitic languages) are transcribed with the same symbols, but a macron is used to indicate vowel length, as in *sūryāyē*. The transcription used for Classical Arabic is close to the DIN 31635 standard, with the difference that case endings and nunation are also transcribed in isolation. Gəʿəz is transcribed according to the system used in Tropper (2002). The normal conventions have been followed for the transliteration of Akkadian and Ugaritic; note that Ugaritic <a i u> represent the combination of the consonant /ʾ/ and a vowel.

In all languages, vocalized forms are written in *italics*, as are consonantal roots; phonetic realizations, written in the International Phonetic Alphabet (IPA), are written in [square brackets]; phonemic representations are written /between two forward slashes/; and again, consonantal spellings are written in <angle brackets>. Reconstructed forms are preceded by an *asterisk, but written in Roman type; expected forms that are not attested are preceded by an asterisk and written in italics; expected forms that are contradicted by the actually attested forms are preceded by two asterisks and written in italics. Many sound changes have been formulated as $a > b$ (/ c_d), which stands for 'a becomes b (when preceded by c and followed by d)'. In this formulaic notation, capital letters represent multiple phonemes: *C* stands for any consonant (sometimes specified by superscript text in square brackets), *V* stands for any vowel (also subject to further specification), and *W* stands for *w or *y. *$* represents a syllable boundary and # represents a word boundary. To indicate that two sounds are identical, they are marked by the same subscript number, as in C_1C_1; non-identity of two sounds is indicated by different numbers, as in C_1C_2. Deletion of a sound is represented as a change to zero, written as Ø. Analogical change is marked by >> two angle brackets. See table 1.4 for other abbreviations.

[10]The choice to only mark stress on non-ultimate syllables reflects a dictum of my esteemed teacher Dr. Martin Baasten's, known locally as Baasten's Law: "In Biblical Hebrew, the stress always falls on the ultimate syllable, except when this is not the case."

Table 1.4: Abbreviations

1	first person
2	second person
2Ch	2 Chronicles
3	third person
abs.	absolute state
acc.	accusative
Akk.	Akkadian
Arab.	Classical Arabic
Aram.	Aramaic
BA	Biblical Aramaic
BH	Biblical Hebrew
CAD	Chicago Assyrian Dictionary
cs.	construct state
Dan	Daniel
Deut	Deuteronomy
du.	dual
EA	El Amarna letter (see Knudtzon 1915) *or* Amarna Canaanite
Est	Esther
Ex	Exodus
Ezek	Ezekiel
f.	feminine
Gə.	Gəʿəz
Gen	Genesis
gen.	genitive
indep.	independent
inf.	infinitive construct
IPA	International Phonetic Alphabet
ipf.	imperfect
ipf.cs.	consecutive imperfect
Is	Isaiah
JBA	Jewish Babylonian Aramaic
Jer	Jeremiah
Josh	Joshua

Judg	Judges
KTU	Keilalphabetische Texte aus Ugarit (see Manfred et al. 1995)
Lev	Leviticus
m.	masculine
n.	note
nom.	nominative
Num	Numbers
PNWS	Proto-Northwest-Semitic
pf.	perfect
pl.	plural
PS	Proto-Semitic
PWS	Proto-West-Semitic
Prov	Proverbs
Ps	Psalms
ptc.	participle
Ru	Ruth
Sam	Samuel
sg.	singular
SH	Spoken Hebrew
suff.	suffixed
Syll.	Syllabic
Syr.	Syriac
Ug.	Ugaritic
WS	West Semitic