11 The phonological brain

Previous chapters have suggested that the human mind is equipped with core phonological knowledge – a system specialized for the computation of phonological structure. This chapter examines what brain mechanisms mediate phonological computation and evaluate their presumed genetic underpinnings. While the findings suggest that a neural phonological network certainly exists, they cannot determine whether this network is specialized for phonology. An answer to this question hinges on how specialization is defined, and, more generally, how cognitive explanations are linked to neuroanatomical models. Existing neuroanatomical models presently lack an explicit account of that link. I thus conclude that specialization, in general, and the hypothesis of core phonology, specifically, can be presently evaluated primarily at the functional, cognitive level. Neural data can be profitably correlated with functional findings, but they can rarely falsify functional hypotheses concerning specialization.

11.1 Individuating cognitive functions: functional specialization vs. hardware segregation

At the center of this book is the question of specialization: Are human minds equipped with a system specialized for phonological patterning? The previous chapters present several observations that are consistent with this possibility. We have seen that distinct phonological systems share design principles that distinguish them from nonlinguistic systems, that knowledge of grammatical universals is evident even when they concern structures unattested in one's language, and that the capacity for phonological patterning emerges spontaneously, in the absence of a model. Not only are phonological constraints universal and possibly innate, but they are also demonstrably distinct from nonlinguistic pressures, most notably, the phonetic pressures governing the processing of aural stimuli and their

https://doi.org/10.1017/CBO9781139049610.016 Published online by Cambridge University Press

production. Functional specialization, however, should be further mirrored at the neural level. If the mind has a specialized computational system dedicated to phonological patterning, then one would expect this special "software" to require a specialized brain "hardware" that mediates phonological computation. The brain networks that support phonological computation could potentially present another test for the specialization of the phonological mind.

While the expectation that functional specialization should have some correspondence in the organization of the brain is uncontroversial, the precise nature of this correspondence is far less clear. Following Gary Marcus and Hugh Rabagliati (2006), we will distinguish between two views of specialization. A strong position requires a one-to-one isomorphism between cognitive function and biological hardware (see Figure 11.1a). In this view, functionally specialized systems should run hardware circuits that are entirely distinct and non-overlapping. Thus, if some specialized system S_1 exists at the cognitive level, then it should be possible to individuate this system at the level of the brain. In the case of the phonological system, there should be a brain network whose components are exclusively dedicated to the computation of phonology. Moreover, the assembly of this system (in development) and its online operation (in the final, adult state) should be controlled by genes whose entire *raison d'être* is the regulation of language functions in the brain. Specialization at the functional cognitive level should thus be transparently discernible from the organization of the brain and its genetic regulation. More generally, any two systems, S1 and S2, are said to be specialized at the functional level only if they can be segregated from each other at the level of the hardware.

In a second, weaker, hypothesis, functional systems individuated at the functional level are not necessarily segregated at the level of "hardware" (see Figure 11.1b). While this view still requires that functional systems can be each

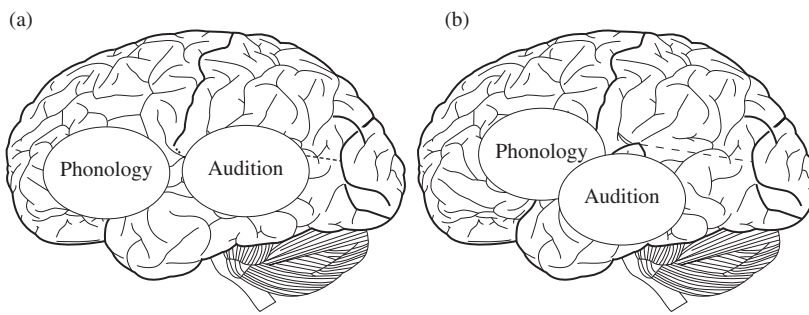

Figure 11.1 Cartoon illustrations of the relationship between two cognitive functions – phonology and audition – and their hardware implementation: either full segregation of the relevant brain substrates (a) or partial overlap (b). Any similarity between localizations in this cartoon and actual brains is purely accidental

linked to distinct brain networks, the relevant networks need not comprise non-overlapping pieces of hardware. A phonological system, for example, might share some (or all) of its components with related nonlinguistic substrates (e.g., audition, motor control), and its operation might be regulated by genes that are expressed at multiple sites, and whose impact is not confined to linguistic brain functions. But despite having no piece of hardware (brain substrates or genes) that is exclusively dedicated to linguistic computations, human brains may be nonetheless innately predisposed to the computation of phonological structure. All healthy human beings would manifest phonological competence that is narrowly constrained by universal principles, this system might preferentially engage an invariant brain network across different individuals, and the functioning of this network could be linked to specific genetic changes unique to humans, such that newly generated random mutations that disrupt these changes will systematically disrupt the language function.

Although this weaker account of neural organization is sometimes considered inconsistent with functional specialization, I believe this view is mistaken. In fact, it's the strong view's insistence on complete neural segregation and its extreme characterization of innateness that are incompatible with modern genetics (see also Marcus, 2006). But for now, I will defer discussion of these claims until the relevant evidence is laid out. I will consider the evidence in two steps (see 1). In the first step, we will review the literature concerning the phonological network and its genetic control. We will begin the discussion by identifying the phonological network of spoken language. We will examine what brain areas mediate phonological computation in healthy individuals and review the consequences of their disruptions in language disorders. After the principal "biological actors" are introduced, we will next move to gauge their specialization for phonology. One test for specialization is the invariance of the system across modalities. To the extent that the system mediating the processing of spoken phonology is dedicated to phonology, rather than to audition or speech per se, then one would expect the key "hubs" of phonological computation to be active across modalities, for both signed and spoken languages. Such similarities, however, could also emerge for reasons unrelated to phonology. Indeed, sign and spoken languages might share computational routines that are domain-general, such as categorization, chunking, hierarchical organization, and sequencing. To illuminate the nature of the overlap, one would like to further evaluate the role of those "phonological" regions in processes that are clearly non-phonological. Music presents a handy baseline. If the neural phonological network is truly segregated, then the mechanisms dedicated to phonological computation should be segregated from those mediating computation in the musical domain. Another test for the specialization of the phonological system concerns its genetic regulation. We will review some hereditary phonological

disorders and examine whether phonological deficits segregate from non-phonological impairments.

(1) Are human brains specialized for phonological computations?
   a. Do human brains include a phonological network?
      (i) Is there a brain network that mediates phonological processing in normal individuals?
      (ii) Are disruptions to that network associated with phonological disorders (congenital or acquired)?
   b. Is the phonological network strictly specialized?
      (i) *Robustness across modalities*: Does the phonological network of spoken language mediate phonological processing in sign language?
      (ii) *Specificity*:
         - Are the components of the phonological network implicated in the processing of music?
         - Do congenital disorders selectively compromise phonological processing?

Foreshadowing the conclusions, there is strong evidence that a phonological network does, in fact, exist. Several brain sites are systematically engaged in the computation of various grammatical phonological structures, these sites mediate phonological computations across individuals and languages, they are partly invariant across modalities – spoken and signed language – and they are further linked to several candidate genes. Nonetheless, none of these brain sites or genes is exclusively implicated in phonology. I will conclude the discussion by considering whether these facts are consistent with the specialization of core phonological knowledge at the functional level.

11.2 The phonological network of spoken language

In view of the very large literature examining the brain mechanisms mediating speech processing (for reviews, see Hickok & Poeppel, 2007; Poeppel et al., 2008), it is striking to see how few have examined grammatical phonological computation. Lacking a concrete model of grammatical phonological computation, specifically, we will thus use speech perception models as a point of departure. Obviously, speech perception and phonology are quite distinct – while speech perception may well be constrained by the phonological grammar, it is also shaped by several other processes, ranging from low-level spectrotemporal analysis to lexical access. Nonetheless, models of speech perception might offer a reasonable first estimate of the grammatical phonological network in spoken language, so we will use them to guide the present discussion.

An influential model by Gregory Hickok and David Poeppel (2007; see Figure 11.2) suggests that the speech stream first undergoes spectrotemporal

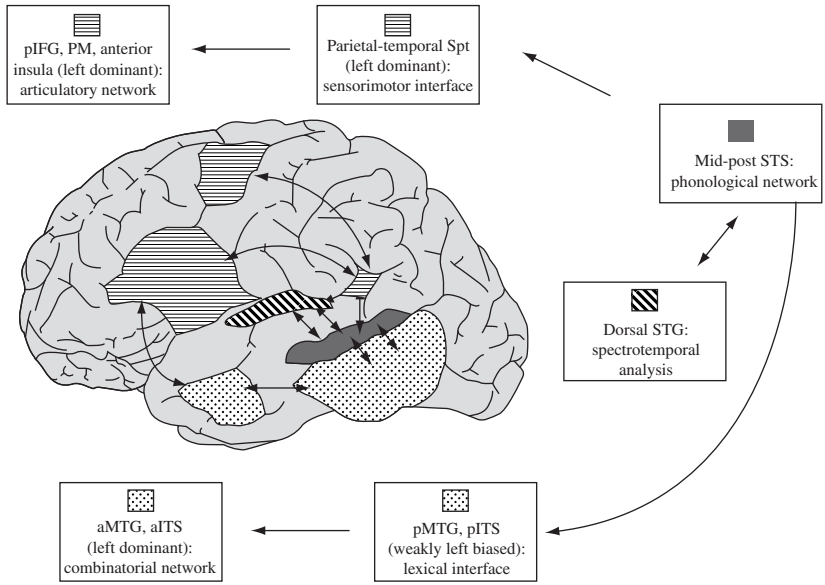

Figure 11.2 Functional anatomy of left hemisphere areas engaged in the phonological processing in spoken language and their interconnectivity (from Hickok & Poeppel, 2007)

analysis through a bilateral activation of the dorsal regions of the superior temporal gyrus (STG) as well as the middle and posterior regions of the superior temporal sulcus. Subsequent processing proceeds along two major neural pathways: ventral and dorsal. The ventral pathway maps auditory inputs onto words stored in the mental lexicon. This stream engages the posterior regions of the middle and inferior temporal gyrus (pMTG/pIFG) as well as anterior and posterior regions of the inferior temporal sulcus and anterior regions of the left middle and inferior temporal gyrus. A second, dorsal stream achieves sensorimotor integration along a left-lateralized pathway involving the left Sylvian parietal-temporal junction (SPT), as well as posterior regions of the inferior frontal gyrus (IFG), including Broca's area, as well as premotor sites and the insula.

Given the behavioral evidence (reviewed shortly), it is clear that some of these sites must be involved in processing grammatical phonological structure, but which sites, specifically, play a role and how they combine to compute phonological structure is rarely addressed. Existing research has typically adopted a rather narrow definition of phonological knowledge. Too often, phonological knowledge is defined by language-particular principles (e.g., knowledge that contrasts English and Russian), and consequently, the

phonological network is equated with mechanisms that mediate language variation (e.g., lexical and phonotactic differences between languages). The behavioral evidence, however, suggests that phonological knowledge might be far broader in scope. Not only does phonological knowledge include grammatical principles that are irreducible to the statistical properties of the lexicon, but there is growing evidence that it might be universal. What brain networks might be involved in grammatical (as opposed to lexical) computation remains largely unknown, and the role of grammatical universals in the brain remains unexplored. To begin addressing these questions, we will thus review the literature in a targeted manner. Rather than asking what is known about phonological computation, generally, we will identify those aspects of phonological computation that are likely to form part of core phonology – phonological primitives and the putatively universal markedness restrictions that might govern their combinations. Our goal for now is to simply identify the regions that mediate the processing of this information – whether those mechanisms are specialized is a question that we defer to subsequent sections.

11.2.1 Phonological primitives

The behavioral evidence reviewed in previous sections suggests that all phonological systems might include several types of representational primitives. All systems apparently represent phonemes, they contrast consonants and vowels, and they encode syllables. Our question here is whether the encoding of those primitives is likewise evident in the brain – in both typical and disordered systems. The next section specifically focuses on segments and the consonant–vowel distinction; the role of syllables is explored along with markedness restrictions in the following section.

11.2.1.1 Phonemes

All phonological systems encode phonemes as discrete elements that are contrasted with each other (e.g., *b* vs. *p*). The processing of such contrasts engages several regions in the STG, and these contrasts have been observed using both functional magnetic resonance imaging (fMRI, e.g., Wilson & Iacoboni, 2006), electrophysiology (e.g., Naatanen et al., 1997; Sharma et al., 1993), and magnetoencephalography (MEG, e.g., Tervaniemi et al., 1999). But while the distinct brain responses to /b/ and /p/, for instance, could reflect phonological processing, these patterns are also amenable to alternative explanations. Stimuli such as /b/ vs. /p/ not only contrast at the phonological level but they also differ on their acoustic and phonetic properties. Accordingly, the observed brain responses could reflect sensitivity to the auditory or phonetic contrast, a possibility that is further supported by the fact that the relevant regions do, in fact, form part of the auditory cortex (broadly defined).

Unlike auditory and phonetic contrasts, however, phonemic categories promote not only differences between speech sounds but also similarities. While members of a single category – say the category of /b/ sounds – can vary in systematic ways across talkers (e.g., in their voice onset time, VOT; e.g., Theodore et al., 2009; Theodore & Miller, 2010), once these distinct members are represented at the phonological levels, the differences between them are erased. A /b/ is a /b/ irrespective of whether its VOT is short (5 ms) or long (e.g., 20 ms), and, within any given language, every phonological generalization true of a short [b] will also apply to a long one. Because the capacity to encode such classes is uniquely phonological, it offers the means to adjudicate between phonological and phonetic/acoustic accounts of discrimination. If the brain encodes speech sounds only at the auditory/phonetic level, then people should respond not only to contrasts between categories (e.g., /b/ vs. /p/) but also to differences among members of each such class (e.g., differentiate [p] with short vs. long VOT). A phonological contrast, however, should register differences between categories, but ignore within-category differences.

The predictions of the phonological account are supported by an MEG study conducted by Colin Phillips and colleagues (2000). In this study, participants were presented with multiple tokens of a single category (e.g., numerous instances of /d/) – the standard – followed by instances of a different category (/t/) – called the deviant. If the brain distinguishes between these two categories, then one would expect the deviant to elicit a change in the brain's magnetic field – a mismatch response. The change associated with the mismatch thus signals the detection of a contrast. By the same logic, however, the mismatch can also gauge the similarity among class members (e.g., of various tokens of a /d/). Indeed, an event can only be considered "deviant" when compared to a "standard," that is, if all other sounds are encoded as members of a single class. In this experiment, however, standard stimuli (e.g., various [d] sounds) were not physically identical, but rather, they slightly differed from each other on their voice onset time. If people encode only the acoustic or phonetic properties of these stimuli, then those various /d/ stimuli should not form a single class (see the top left panel of Figure 11.3), and consequently, no standard should be in place. In the absence of a standard, /t/ should not be recognized as a deviant, so no mismatch response is expected. If, however, people encode those various stimuli phonologically, as members of the /d/ class, then these stimuli should all be considered alike, and the perceived "standard" should give rise to the detection of /t/ as a deviant (see the top right panel of Figure 11.3). This is precisely what was observed: A mismatch response occurred approximately 200 ms after the onset of the auditory stimulus, and it was localized at the superior temporal plane of the left hemisphere (see Figure 11.4).

A second control condition ruled out the possibility that participants relied on a non-phonological acoustic category (see the bottom of Figure 11.3). It is

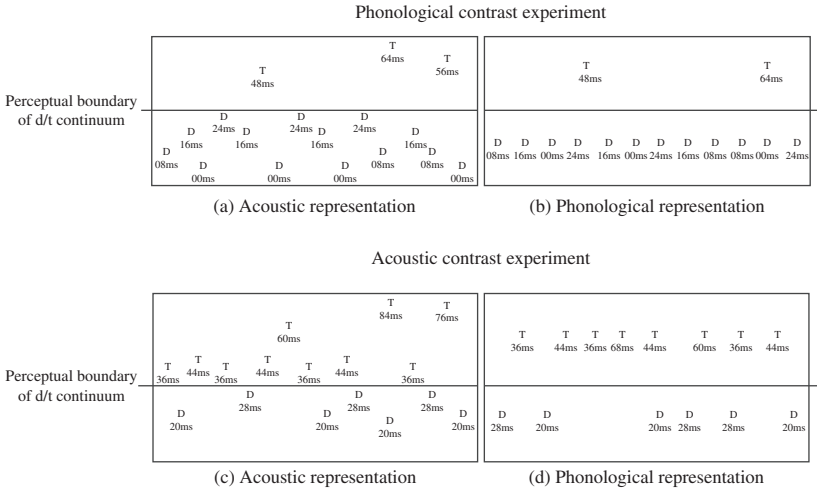

Figure 11.3 The design of Phillips et al.'s experiments (2000)

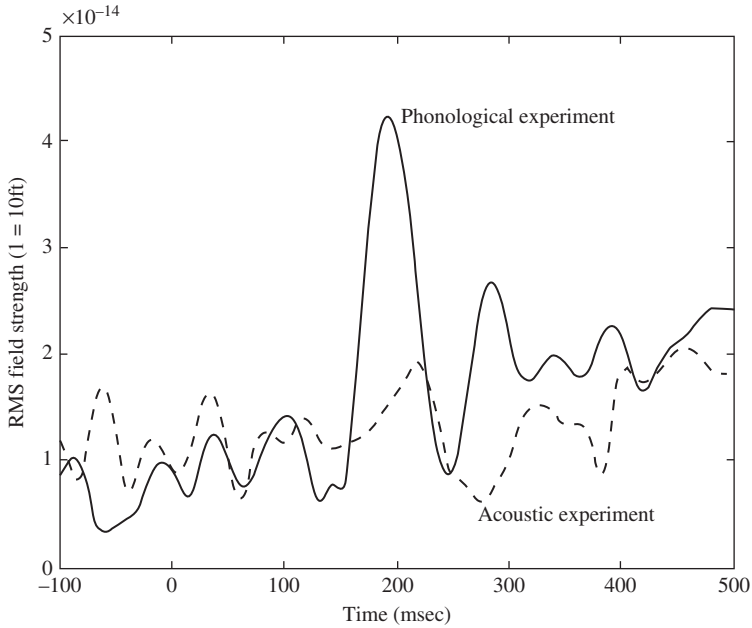

Figure 11.4 Brain responses to the phonological and acoustic control conditions in Phillips et al.'s (2000) experiments

https://doi.org/10.1017/CBO9781139049610.016 Published online by Cambridge University Press

indeed possible that people identified all "long" and "short" VOT stimuli as distinct classes of equivalent members, but the distinction between the classes was based on their mean VOT, rather than by the phonemic contrast between /d/ and /t/, specifically. To address this possibility, Phillips and colleagues increased the VOT of all items by 20 ms. This manipulation still maintained the same VOT difference between "short" and "long" stimuli, but altered their phonological structure. While in the original experiment (at the top of Figure 11.3), "standard" stimuli were identified as the voiced /d/ and the "deviant" was a /t/, most edited items (see the bottom panels) were identified as the voiceless /t/ – either shorter /t/ (for the "standard") or a long one (for the "deviant"). An auditory account would predict a similar outcome in both experiments (see the bottom left panel of Figure 11.3), but if people rely on a phonological representation, then they should now be unable to distinguish the standard from the deviant, and the mismatch effect should be eliminated (see the bottom right panel of Figure 11.3). The outcomes, once again, supported the phonological explanation (see Figure 11.4). Together, these two experiments suggest that regions of the left superior temporal plane mediate phonological processing.

Converging evidence for this conclusion is presented by transcranial magnetic stimulation. While imaging studies reflect typical brain functioning, transcranial stimulation temporarily disrupts activity in select brain areas by transmitting low-level electrical current to electrodes placed on the surface of the cortex. To the extent that that region mediates the function of interest (e.g., the identification of phonemes), then its stimulation should disrupt that function. Results from this procedure converge with the MEG findings to suggest that the middle-posterior region of the STG is involved in both phonetic and phonological processing of consonants (for review, see Boatman, 2004). Other phonological tasks that require phonemic awareness and explicit segmentation (e.g., do *pat* and *bat* share the initial consonant?) implicate additional adjacent regions in the left STG (anterior middle STG, ventral and dorsal portions of posterior STG), and the left inferior frontal lobe (Boatman, 2004) – regions that overlap with Hickok and Poeppel's proposal. And while the right hemisphere does support some limited discrimination between minimal pairs (in both healthy individuals and patients with left hemisphere damage), its effect is gradient and confined to word pairs (e.g., *beach–peach*, but not nonwords, e.g., *beesh–peesh*), whereas discrimination in the left hemisphere is categorical and extends to both words and nonwords (Wolmetz et al., 2011). These observations suggest that the right hemisphere can engage in acoustic-phonetic analysis that supports phoneme discrimination, but it is the left hemisphere that mediates the categorical perception of phonemes (Liebenthal et al., 2005; Wolmetz et al., 2011). The capacity of the left hemisphere to engage in categorical perception might be due to its preferential tuning to the sampling of acoustic information at a temporal window that matches the duration of

individual phonemes (20–50 ms; Morillon et al. 2010; Poeppel, 2003), and this, in turn, might explain the advantage of the left hemisphere in the perception of phonemes.

11.2.1.2 Consonant–vowel distinctions

Consonants and vowels are distinct phonological animals. Many phonological processes specifically target one category while ignoring the other – consonant co-occurrence restrictions apply to consonants across intermediate vowels, whereas vowel harmony is opaque to intervening consonants. Furthermore, consonants and vowels carry distinct roles in the grammar and in language processing (see Chapter 4). These observations suggest that consonants and vowels form distinct functional categories in the phonological mind. In what follows, we show that the processing of consonants and vowels likewise dissociates at the neural level.

One line of evidence comes from an fMRI study of typical individuals (Obleser et al., 2010). This study compared brain activation for consonants (e.g., *da* vs. *ga*) and vowels (e.g., *da* vs. *di*). The findings revealed dissociation between consonants and vowels, but the distinction between them was rather subtle. When the results were examined using a typical subtraction methodology (by comparing brain activation to speech with noise), consonants and vowels did not differ in their pattern of activation, and they both elicited stronger activation in the left anterior-lateral superior temporal cortex. But a more sensitive analysis of the data using a classifier algorithm revealed different patches dedicated for the processing of consonants and vowels that were largely non-overlapping.

If consonants and vowels recruit different brain substrates, then it should be further possible to selectively disrupt the processing of one of these substrates (e.g., consonants) without disrupting the other. This prediction is borne out by the findings from transcranial stimulation. Earlier, we noted the role of the middle posterior region of the STG in processing the distinction between consonants. As it turns out, the same region does not respond to distinctions among vowels (Boatman, 2004). Moreover, the engagement of the STG by consonants is categorical, occurring irrespective of their specification for various phonological features, such as voicing (e.g., *b* vs. *p*) and place of articulation (e.g., *b* vs. *g*). Although the precise location (the middle posterior STG) is more posterior than the fMRI findings (Obleser et al., 2010), they nonetheless converge with other fMRI results (Wolmetz et al., 2011). Together, these two methods reflect a distinction between the neural substrates representing consonants and vowels.

Consonants and vowels likewise dissociate in naturally occurring phonological disorders. Recall (from Chapter 4) that consonants and vowels can be selectively impaired in aphasia. Alfonso Caramazza and colleagues (2000)

describe two conduction aphasic patients who exhibited fluent speech and no motor or articulatory deficits, but their performance on oral repetition and naming was poor. In one patient (IFA), errors concerned mostly consonants, whereas for another (AS) the errors concerned mostly vowels. In accordance with the evidence from transcranial stimulations, the susceptibility of consonants and vowels to errors was independent of their feature constitution, as gauged by their sonority level. Moreover, one of IFA's damaged loci – the left STG – broadly overlaps with the sites implicated by transcranial stimulation and fMRI findings. In contrast, AS showed (unspecified) regions to the left parietal and temporal lobes, as well as a small region to the right parietal lobe.

Consonants and vowels likewise dissociate in reading and spelling (e.g., Cotelli et al., 2003; Miceli et al., 2004). In one extreme case, an Italian patient with a deficit to the left parietal lobe produced spelling responses that omitted all vowels while leaving consonants intact (Cubelli, 1991). The patient, Fondacaro Ciro, spelled his name as FNDCR CR; the city Bologna was spelled as BLG, and the word 'table,' *tavolino*, was spelled TVLN. The most curious case of such dissociations is presented by a healthy woman who manifests strong synesthesia that systematically associates different vowel letters with specific colors (e.g., A is associated with red; E with green, etc.; Rapp et al., 2009). These synesthetic associations are so specific and automatic that they even promote Stroop interference. In Stroop experiments, people are typically presented with words printed in color, and they are asked to name the color while ignoring the word's meaning. People typically experience difficulty when the word spells an incongruent color name (e.g., given the word GREEN printed in red). Remarkably, this synesthetic individual exhibited Stroop interference for vowels: She experienced difficulty when a vowel letter (e.g., A, which is synesthetically associated with red) was presented in a color that was incongruent with the synesthetic percept (A presented in green). Moreover, fMRI scans demonstrated that vowel letters activated brain areas that typically mediate color processing (V4 and V8). Crucially, however, these effects were only present for vowels, not consonants, and as such, they present dissociation between the brain mechanisms dedicated to these two categories.

11.2.2 Markedness restrictions

The findings reviewed in the previous section have identified various brain substrates that are engaged in the representation of phonological primitives. Core phonology, however, includes not only shared representational primitives but also broad, perhaps universal, well-formedness constraints governing their combinations. Accordingly, structures that violate those constraints (i.e., marked structures) are systematically underrepresented across languages and disfavored by individual speakers – both adults and young children. In what follows, we

examine what brain substrates might mediate markedness preferences. Our case study concerns the markedness restrictions on syllable structure.

All languages constrain syllable structure: Onsets are preferred to codas (e.g., *ba*≻*ab*), simple onsets are preferred to complex ones (e.g., *ba*≻*bla*), and onsets with large sonority distances are preferred to those with smaller ones (e.g., *bla*≻*lba*). Our interest here is in the hallmarks of markedness in the brains of individual speakers – both typical participants, and those afflicted with aphasia.

11.2.2.1 Syllable markedness in typical individuals

While syllable-structure preferences have been widely documented in typology and behavioral experiments, the brain mechanisms mediating such computations remain largely unexplored in the imaging literature. The most relevant findings come from an fMRI study by Charlotte Jacquemot and colleagues (2003) that compares syllable-structure preferences of Japanese and French speakers. These two languages differ on their syllable structure: Japanese strictly disallows syllables like *eb*, so inputs like *ebzo* are systematically repaired as *ebuzo*. While the contrast between *eb.zo* and *e.bu.zo* is absent in Japanese, both structures are allowed in French (see 2). French, however, does not contrast vowels in terms of their length (e.g., it does not distinguish *eb**u**za* and *eb**uu**za*), whereas this contrast is present in Japanese. Prior research has further shown that these phonotactic preferences modulate the responses of Japanese and French speakers in behavioral experiments (Dupoux et al., 1999). The goal of the present study was to identify the correlates of those phonotactic preferences in the brain. To this end, this study compared the brain responses of Japanese and French speakers to these two contrasts – the presence/absence of a coda (e.g., *e**b**.za* vs. *e.b**u**.za*) and vowel length (e.g., *eb**u**za* and *eb**uu**za*).

(2) Phonological contrasts in Japanese and French

		Is the contrast attested?	
		Japanese	French
Contrast	*eb.za-e.bu.za*	No	Yes
	ebuza-ebuuza	Yes	No

To distinguish the brain network that mediates phonological computation from purely phonetic substrates, these authors relied on phonological attestation as a guide. Specifically, they reasoned that attested contrasts elicit phonological processing whereas unattested contrasts can only elicit phonetic processing. Following this rationale, they identified phonological areas as those that selectively mediate the processing of attested contrasts, but not unattested ones. These areas included several left perisylvian regions (the IFG, the STG,

supramarginal and angular gyri, and the intraparietal sulcus), bilaterally activated regions including the cingulate cortex, insula, and precentral gyrus, and right hemisphere regions, in the frontal, superior, and middle temporal gyri. While these sites agree with Hickok and Poeppel's model of speech perception, it is unclear that they demarcate a phonological network. Because these sites were defined by contrasting attested and unattested structures, the outcome only narrowly indicates the sites that mediate knowledge of one's own language-particular phonotactics. Behavioral findings, however, suggest that speakers' phonological knowledge may well extend to unattested structures. By equating the phonological network with language-particular phonotactics, one runs the risk of over-emphasizing the contribution of lexical phonological computations to the exclusion of sites supporting core grammatical knowledge. Nonetheless, the conclusions of this pioneering work can guide future research into markedness manifestations in the brain.

11.2.2.2 The role of syllable markedness in aphasia

While imaging studies with typical individuals have not explored markedness restrictions directly, aphasia research has long been interested in the possibility that markedness might play a role in aphasia. Roman Jakobson (1968) has famously asserted that unmarked structures are more likely to be preserved in aphasia. Accordingly, other things being equal, aphasic patients should be more likely to preserve unmarked segments in their speech, and when they produce an error, they should be more likely to produce unmarked structures. This prediction is borne out by the results of several studies.

One recent demonstration of syllable-structure restrictions is presented by the case of BON, reported by Matt Goldrick and Brenda Rapp (2007). BON is an English-speaking female with left hemisphere damage affecting the left superior posterior frontal regions and the parietal lobe. While BON produced many errors in her speech, her production showed a strong effect of syllable markedness. BON was reliably more likely to produce a consonant (e.g., *p*) correctly when it occurred in the onset (e.g., ***p****at*) than in the coda (e.g., *u****p***). Goldrick and Rapp further demonstrate that the advantage of onsets is not due to various extraneous factors, including the status of the consonant as a singleton vs. cluster (e.g., ***p****lay vs*. ***p****ay*), the frequency of the consonant in these two positions, or even the frequency of those syllable structures themselves (e.g., the frequency of CV vs. VC syllables). In fact, words including an onset are *less* frequent than those including a coda.

Another indication of the preference for unmarked syllable structures concerns complex onsets (e.g., ***cl****ip*). It has long been noticed that aphasic patients avert complex onsets – such onsets are frequently simplified by either segment deletion (e.g., *sky*➔*ky*) or the epenthesis of a schwa (e.g., clip➔[kəlɪp]; for review, see Blumstein, 1973; 1995; Rapp & Goldrick,

2006). One concern, however, is that these simplifications reflect not markedness pressures but rather non-grammatical sources – either the patient's inability to encode the auditory input presented to him or her or to plan and execute the relevant articulatory motor commands. A recent study by Adam Buchwald and colleagues, however, addresses this possibility (Buchwald et al., 2007).

Buchwald and colleagues discuss the case of VBR – an English-speaking female who suffered from a large fronto-parietal infarct to the left hemisphere due to a cerebral vascular accident. VBR's word production was characterized by frequent distortions of word-initial onset clusters. Given an input such as ***bl****eed*, VBR would typically produce [bəlid] – separating the onset consonants by appending an epenthetic schwa. A detailed investigation ruled out receptive factors as the source of these errors – VBR was not simply unable to encode the acoustic input or incapable of accessing her mental lexicon. Likewise, detailed acoustic and ultrasound analyses of her production demonstrated that the errors did not result from articulatory failures. By elimination, then, these analyses suggest that epenthetic errors have a specific phonological origin, which potentially concerns the markedness of complex onsets.

Further evidence that the simplification errors in aphasia have a phonological origin is presented by their sensitivity to sonority-sequencing restrictions. Analyzing the case of DB, an Italian patient with a left fronto-parietal lesion, Cristina Romani and Andrea Calabrese (1998a) demonstrated that production errors were constrained by sonority sequencing. Moreover, sonority sequencing constrained both the intended target as well as its erroneous rendition. DB was reliably more likely to simplify onset clusters (i.e., targets) with small sonority rises (e.g., liquid-glide combinations, *rya*) compared to less marked onsets with larger distances (e.g., obstruent-liquid combinations, e.g., *tra*). Moreover, when a complex onset was simplified, DB was reliably more likely to opt for outputs that maximize sonority distance (e.g., *tra*➔*ta* rather than *ra*). These errors were inexplicable by auditory failures (DB's ability to discriminate such auditory syllables fell within the normal range), nor were they due to an inability to produce certain segments (e.g., difficulty with the liquid *r*). Indeed, the errors associated with any given segment depended on its position in the syllable. For example, while DB tended to delete liquids in onset clusters (e.g., *tra*➔*ta*), he did maintain them in codas (e.g., *art*➔*ar*), a pattern consistent with the cross-linguistic asymmetry between onsets and codas (onsets favor a steep rise in sonority, whereas codas manifest moderate falls). Similar cases have been observed in German- (Stenneken et al., 2005) and English-speaking (e.g., Buchwald, 2009) patients. Together, those results suggest that marked structures are more likely to be impaired in aphasia.

11.2.3 Conclusions

Our discussion so far has identified several brain regions mediating phonological processing in normal individuals. These areas (including Broca's area, the posterior STG, superior temporal sulcus, and planum temporale) are implicated in the representation of both phonological primitives (segments, and consonant/vowel categories) and some markedness restrictions on syllable phonotactics. Disruptions of these areas – either transitory disruptions due to transcranial stimulation, or more permanent ones, in the case of aphasia – can selectively impair specific phonological elements (e.g., consonants). Moreover, brain injuries are more likely to spare phonological structures that are universally unmarked.

While these findings do not specifically address the connectivity between these regions, they are certainly in line with the possibility that these loci form part of a network that mediates the computation of core phonological knowledge. Nonetheless, several limitations of these conclusions are noteworthy. First, the available evidence concerning the phonological grammar is scarce. Only a handful of studies with normal individuals have explicitly examined substrates mediating the encoding of grammatical phonological primitives and constraints, so the reliability of these findings and their generality across languages requires further research. Some aphasia studies have concerned themselves with grammatical representations and constraints, and the results are generally consistent with the imaging findings with typical individuals, but in many cases, the localization of the lesions is rather coarse. Finally, the existing evidence is limited inasmuch as it focuses exclusively on the structures attested in one's language. Accordingly, the available data do not address the central question of how the brain encodes core phonological knowledge, including grammatical principles that are potentially universal.

11.3 Is the phonological network dedicated to phonological computation?

Finding that the brain is involved in grammatical phonological computation is hardly surprising given the very large corpus of behavioral data demonstrating that people constrain the phonological structure of their language. Our main question here is not *whether* the brain computes grammatical phonological structure, but rather *how* – is the relevant network specialized for phonological computation, or does it consist of domain-general mechanisms that subserve the processing of auditory sequences and vocal motor control, generally?

An answer to this question would critically depend on one's account of specialization. Earlier in this chapter, we defined two views on how functional specialization should be mirrored in the organization of the brain. A weaker

version merely requires that the cognitive function of interest be associated with a cohesive brain network that is relatively fixed across individuals. A stronger view further mandates that this network and each of its components should all be *exclusively* dedicated for that function alone. The evidence reviewed so far is certainly consistent with the weaker view, but it is moot with respect to the stronger alternative.

In what follows, I evaluate this strong hypothesis from two perspectives. One way to gauge the specialization of phonology is to compare the neural networks that mediate phonological computation across modalities. If the substrates involved in the processing of spoken language phonology are general-purpose engines of auditory sequencing, for instance, then one should not expect it to mediate the computation of phonological structure in sign languages. Specialization, however, should be evident not only by the range of functions that the network mediates but also by the ones it doesn't. A brain network dedicated to phonology should be demonstrably distinct from the systems mediating the processing of non-phonological auditory sequences, such as the ones found in music. The role of the phonological network in non-phonological computations presents a second test for its specialization. Finally, I will examine whether disorders affecting the phonological system are dissociable from non-phonological deficits.

11.3.1 An amodal phonological network?

The discussion in previous chapters has shown that phonological patterning is not confined to speech. Like spoken languages, sign languages have a phonological structure that imposes constraints on the sequencing of meaningless sign elements. We have identified several representational primitives that are shared across modalities, including features and syllables, and reviewed markedness constraints on syllable structure that are possibly amodal. In view of those structural similarities, one wonders whether some of the brain substrates involved in phonological computation might be shared across modalities. Although numerous studies have compared the brain networks in signed and spoken language (for reviews, see Emmorey, 2002), few comparisons specifically concern phonology, and none targets phonological primitives and markedness constraints, in particular. While these crucial questions remain unanswered, there are several indications of common brain regions mediating phonological processing across modalities. Here, I review evidence from two sources: imaging findings from sign language phonology and evidence for phonological transfer across modalities.

11.3.1.1 Brain mechanisms mediating phonological processing in sign language

One region implicated in processing the phonological structure of spoken language is the left planum temporale – a region of the superior temporal

gyrus that forms part of the classical Wernicke's receptive language area. Our question here is why is the planum temporale engaged – does it mediate sensory auditory processing, generally, or phonological patterning, specifically? To address this question, Laura Ann Petitto and colleagues (2000) examined whether this region might support phonological processing in sign language. Using a PET methodology, these researchers compared the brain activity of signers fluent in two distinct sign languages (American Sign Language and Quebec Sign Language) while viewing non-signs (phonotactically legal combinations that are not attested in participants' sign language). To control for the nonlinguistic demands associated with processing such complex visual displays, these signers were compared to non-signer controls. Results showed that signers activated the planum temporale bilaterally in processing sign language. In contrast, non-signers who viewed the same signs showed no engagement of the planum temporale. These results suggest that the planum temporale mediates phonological processing, specifically, irrespective of modality – for either spoken or signed language.

Further evidence for the role of the superior temporal gyrus in phonological processing is presented by another PET study of a deaf individual who was about to undergo surgery for the insertion of cochlear implants (Nishimura et al., 1999). Results showed that, prior to implant, sign language words elicited a bilateral activation of the supratemporal gyri relative to still-frame controls, but once this individual had undergone the implant, auditory inputs triggered activation in the primary auditory cortex (Heschl's gyrus), contra-lateral to the auditory input. These results demonstrate that phonological processing is clearly dissociable from auditory processing.

The amodal phonological network, however, also includes frontal and parietal regions. In a study comparing phonological processing in English and British Sign Language (MacSweeney et al., 2008), bilingual deaf participants were presented with pictures of two objects, and they were asked to perform two judgments regarding their phonological forms. One task required participants to determine whether the signs of those objects in British Sign Language share the same location (e.g., whether both signs are located at the forehead). A second task required participants to judge whether the English names for two objects rhyme, and the performance of these deaf signers was compared to English-speaking monolinguals. Results showed that deaf participants activated common regions in the phonological processing of the two languages (British Sign Language and English), and those regions overlapped with those used by hearing participants. Those regions included medial portions of the superior frontal gyrus, the left superior parietal lobule, the left superior portions of the supramarginal gyrus, and the left posterior inferior frontal gyrus. British Sign Language, however, resulted in greater activation of the left parietal lobule. The greater role of the parietal lobe in sign language is consistent with cortical

stimulation mapping results, showing that stimulation of the supramarginal gyrus results in phonological errors (Corina & Knapp, 2006; Corina et al., 1999). The same study also suggests that sign language phonology recruits posterior aspects of Broca's area (BA 44) as well, as the stimulation of this area resulted in global phonetic distortions. Surprisingly, however, stimulation of the posterior and anterior temporal lobe, including the superior temporal gyrus, resulted in no disruptions of sign repetition (Corina et al., 1999). These results, however, were obtained from an individual with a history of complex parietal seizures, so this factor might explain the discrepancy with the imaging results of phonological processing in neurologically typical signers, where the superior temporal gyrus is often implicated.

11.3.1.2 Cross-modal phonological transfer

Further evidence for an amodal phonological network is presented by the transfer of phonological processing across modalities. It is well known that language acquisition is optimal early in development; later language learners manifest various deficits which are particularly noticeable in the area of phonology – far more so than in lexical access (Newport, 2002). Eric Lenneberg (1967) famously attributed this early window of opportunity to biological constraints on the plasticity of the language system. If such constraints, however, are amodal in nature, then exposure to phonological structure might transfer across modalities. The findings reported by Rachel Mayberry (Mayberry & Witcher, 2005; 2007) are in line with this prediction. The study examined phonological processing using the priming methodology: Participants were asked to determine whether a given "target" display is a real sign in American Sign Language (ASL). Each such target was preceded by a "prime" – a display consisting of another sign that was either unrelated to the target, or phonologically similar, such that the prime and target differed by a single phonological feature. To use an English analogy (see 3), one would compare the processing of *bee* preceded by either the phonologically similar prime *pea* (*bee* and *pea* differ by a single feature – voicing) or the unrelated control *too*. Two questions are of interest. First, are participants sensitive to the phonological properties of the target? If they are, then phonological primes should yield phonological priming – they should facilitate the processing of the target relative to controls. To the extent that people are sensitive to phonological structure, one might further inquire whether these phonological effects depend on the age of acquiring ASL.

(3) Phonological priming in English

Target:	*bee*
Phonological prime:	*pea*
Unrelated control:	*too*

Results with participants who acquired ASL early in life indeed showed phonological priming even when the interval between the prime and target was relatively short (330 ms), whereas participants who acquired ASL later in life did not benefit from the phonological prime; in fact, they showed phonological inhibition – their responses were slower in the presence of the phonologically related prime relative to the control. These results are indeed expected in light of the large body of research showing that the phonological system of late language learners differs from native speakers. The interesting twist comes from a second condition in which the interval between the prime and target was increased (to 1 second), allowing the prime additional time to exert its effect. Remarkably, the effect of this manipulation on late ASL learners critically depended on their prior experience with English. Late ASL learners who had no prior experience with English still showed no benefit from the phonologically related prime; indeed, they continued to show phonological inhibition. In contrast, late ASL learners who were previously exposed to English showed quite a different pattern. Once the long prime-target interval allowed for sufficient processing time, their results now showed a benefit from the phonologically related prime, approximating the performance of the native ASL speakers in the shorter duration. These findings suggest that early exposure to the phonological system of spoken language leads to long-term benefits that transfer to sign language phonology.

Further evidence for cross-modal transfer is presented by imaging results, showing that participants' age of acquiring their first language – British Sign Language – determined not only the regions mediating phonological processing of their native language but also affected the regions engaged in processing English – their second language (MacSweeney et al., 2008). Compared to late learners, native learners of British Sign Language exhibited stronger activation of the left posterior inferior frontal cortex in making rhyme judgments regarding English words (e.g., *chair–bear* vs. *hat–bear*). These results suggest that phonological experience in one modality may transfer to determine both the functional and the neuroanatomical characteristics of phonological processing in the other modality.

11.3.1.3 Conclusion

To summarize, the comparison of signed and spoken languages suggests that several brain substrates might mediate phonological processing across modalities, and that the early engagement of those areas in one modality transfers to the other. These results, however, do not necessarily demonstrate a shared, amodal network dedicated to phonological computation. First, there are several differences between the phonological networks in the two modalities. Such differences, however, are only expected given the functional differences between the

two phonological systems. A second, more significant limitation of the present results is that they do not necessarily show that these shared substrates are, in fact, dedicated to the computation of phonological structure. Indeed, these common regions might mediate processing demands that are shared across modalities – functions such as the formation of feature-categories, the segmentation of the input stream into smaller chunks, the ability to map those chunks into meaning, etc. To assess this possibility, it is necessary to examine not only what the phonological system does, but also what it doesn't do. A strong view of specialization requires that the substrates involved in phonological computation be exclusively dedicated to this purpose. The following sections evaluate this strong hypothesis. We first examine whether any phonological region is specialized for the purpose of phonological processing; next, we examine whether the network as a whole is dedicated for phonological computation.

11.3.2 Do phonological regions mediate musical computations?

A more stringent test for the specialization of the phonological network seeks to dissociate it from non-phonological brain substrates. Strong neuroanatomical specialization would be demonstrated if at least some of the components of this network are uniquely dedicated for this purpose. The comparison of phonological processing and musical processing presents an interesting case study.

Like phonological systems in natural language, music is universally present in all cultures, and the distinct musical systems across the world share some common organizational principles that are quite distinct from the ones governing phonological structure (see Chapter 2; Jackendoff & Lerdahl, 2006; Lerdahl & Jackendoff, 1983; Patel, 2008; Peretz, 2006). The universality of music, its common design, and the emergence of musical abilities early in development render music a good candidate for a specialized knowledge domain, distinct from core phonological knowledge. Given the strong evidence for functional specialization of the two domains – phonology and music – we can now turn to examine whether any of the brain substrates mediating phonological computation are distinct from the ones implicated in music processing. Surprisingly, however, no known component of the phonological network is uniquely phonological.

The strongest candidates for specialized phonological areas are the left hemisphere sites that mediate phonological processing across modalities, spoken and signed. These sites include Broca's area (Corina et al., 1999; Gough et al., 2005; Jacquemot et al., 2003; MacSweeney et al., 2008; Petitto et al., 2000; Sahin et al., 2009); the posterior superior temporal gyrus and superior temporal sulcus (Boatman, 2004; Desai et al., 2008; Gow & Segawa, 2009; Graves et al., 2008; Liebenthal et al., 2003; Liebenthal et al., 2005; Okada & Hickok, 2006; Vouloumanos et al., 2001); and the planum temporale (Jacquemot et al., 2003;

Petitto et al., 2000). Each of these regions, however, is implicated in tasks involving nonlinguistic tonal material. Broca's area is engaged in processing unexpected harmonic progressions (Koelsch et al., 2002; Maess et al., 2001); the left STG and left superior temporal sulcus have been linked to various aspects of pitch and harmonic processing (Koelsch, 2006; Mandell et al., 2007; Tillmann et al., 2006; Wilson et al., 2009), and the left planum temporale has been likewise associated with the perception of absolute pitch (e.g., Schlaug et al., 1995) and singing (Jeffries et al., 2003; Suarez et al., 2010).

Further evidence for the sharing of neural resources across domains is presented by the strong functional links between musical and phonological abilities. Absolute pitch, for example, is more frequent among musicians who speak tonal languages (Deutsch et al., 2006), whereas musicians are more sensitive to linguistic pitch than non-musicians (Bidelman et al., 2009; Wong et al., 2007). Similarly, about a third of the people who suffer from amusia (a disorder affecting the processing of musical pitch) also manifest difficulties in the processing of linguistic pitch information (e.g., in discriminating linguistic declarative statements from questions; Patel et al., 2008). Summarizing, then, no known brain region can be linked to the computation of any specific phonological structure (e.g., the computation of syllable structure; Blumstein, 1995), and, as shown above, each of the key regions mediating segmental phonological computation is shared with musical processing.

11.3.3 The regulation of the phonological network: evidence from hereditary phonological disorders

Although we have so far failed to identify any specialized phonological regions, further evidence for specialization could conceivably come from their genetic regulation. If humans are genetically predisposed to engage in phonological patterning, then phonological ability could be linked to specific genes that have undergone changes in the human lineage, and the disruption of those genes should result in phonological disorders. But as in the case of neuroanatomical specialization, however, the interpretation of those findings depends on one's definition of specialization. A strong view would infer genetic predisposition for phonological computation only if some genetic mutation could be shown to selectively affect phonological competence; a weaker version might require that phonological competence exhibit a one-to-one correspondence with specific genes, but it would not insist on those genes being exclusively implicated in phonological functions.

Existing genetic research on the phonological competence of healthy individuals is extremely limited. The only available findings associate the prevalence of linguistic tones with the frequency of two genes related to brain growth and development in the population (*ASPM* and *Microcephalin*; Dediu & Ladd,

2007; Dediu, 2011). But the specific function of these genes and their role in phonological competence remain unknown. A large literature, however, has examined the genetic mechanisms of phonological disorders. Here, we consider the findings from individuals with various forms of Specific Language Impairment (SLI).

One of the most exciting advancements in the genetics of language has been the discovery of the *FOXP2* gene, and its involvement in speech and language disorders (Lai et al., 2001). The genetic basis of the disorder was strongly suggested by its inheritance pattern across three generations of a single British family – the KE family. About half of the family members (8/15) are affected, and affected individuals manifest a variety of linguistic disorders ranging from morphosyntax to phonology and phonetics. Imaging studies have revealed structural (Vargha-Khadem et al., 1998) and functional (Liegeois et al., 2003) abnormalities in the brain of affected members, including abnormalities to Broca's and Wernicke's areas. Subsequent analyses have linked the disorder to a mutation to a single gene – *FOXP2*, a transcription factor gene located on the long arm of chromosome 7. The mutation was transmitted in a dominant fashion, such that every family member with the mutated gene exhibited the disorder, and this mutation was only present in affected members (for reviews, see Fisher & Marcus, 2006; Fisher & Scharff, 2009; Marcus & Fisher, 2003; Vargha-Khadem et al., 2005). Subsequent studies have linked the *FOXP2* gene to the learning of the motor skills necessary for vocal communication in various other species, including birds (Haesler et al., 2004) and mice (Fischer & Hammerschmidt 2011; French et al., 2007; Shu et al., 2005). The human FOXP2 protein, however, differs from the versions found in mice by three amino acids, two of these changes occurred after the human evolutionary branch split from the chimpanzee, and a comparison of the rate of this change in nonfunctional changes (changes that do not alter amino acids) suggested that the mutation of the human *FOXP2* allele was due to evolutionary selective pressure on the human lineage, occurring within the past 200,000 years – a time that is broadly consistent with the estimated emergence of language in humans (Enard et al., 2002). The link between the *FOXP2* gene and language evolution, on the one hand, and its role in speech and language impairments, on the other, suggest that the gene regulates the assembly of brain networks that mediate linguistic computation.

Behavioral analyses of affected family members further revealed deficits to tasks that require phonemic awareness, including rhyme production (e.g., say what word rhymes with *name*?), phoneme addition (e.g., add *v* to *arg*), and phoneme deletion (e.g., say *varg* without the first sound; Alcock et al.,, 2000; Vargha-Khadem et al., 1995). Affected members also exhibit various types of phonological errors, including epenthesis (e.g., *statistics*➔*sastistics*), metathesis (e.g., *cinnamon*➔*cimenim*), assimilation (*parallel*➔*pararrel*; examples

from Shriberg et al., 2006), and simplification of complex onsets (e.g., *blue➔bu;* Hurst et al., 1990). Although the simplification of complex onsets could reflect markedness pressures, this explanation is countered by the documentation of (phonologically unmotivated) omission errors even with simple onsets (*table➔able*). Indeed, affected members also exhibit severe verbal dyspraxia – an impairment in the performance of the movements necessary for the production of speech – as well as marked deficits to oral facial movements of all kinds, including the production of meaningless noises (e.g., "click your tongue"), singing (e.g., "hum a tune") and various complex movements (e.g., "open your mouth wide, stick out your tongue, and say *ah*"). The production difficulties of individuals with *FOXP2* mutations could thus stem from their apraxia, rather than from a grammatical deficit.

While the grammatical competence of individuals with *FOXP2* mutations has not been fully evaluated, their phonological difficulties mirror some of the impairments seen in the broader group of individuals with linguistic difficulties that are unexpected by their overall intelligence, neurological development and environment, a disorder broadly categorized as SLI. Individuals with this milder and more common disorder typically do not exhibit mutations to the *FOXP2* gene (Balaban, 2006), but SLI is highly heritable (Bishop, 2009; Bishop & Snowling, 2004), and it has been associated with several other candidate genes (Fisher & Marcus, 2006; Newbury & Monaco, 2010). The phonological competence of such individuals has been evaluated far more extensively.

It has been well established that individuals with SLI exhibit a host of phonological disorders, ranging from the categorical perception of phonemes (e.g., van der Lely et al., 2004; Ziegler et al., 2005) to phonemic awareness (i.e., the comparison, discrimination, and segmentation of phonological forms), the use of prosodic information (Marshall et al., 2009), and word production (for reviews, see Bishop & Snowling, 2004). Production errors affect both words (e.g., Bortolini & Leonard, 2000) and nonwords (Bishop, 2006), and they are exacerbated as the number of syllables increases (e.g., Kavitskaya et al., 2011). An interesting study by Nichola Gallon and colleagues specifically links the phonological difficulties of individuals with SLI to phonological markedness (Gallon et al., 2007). While the SLI group did not differ from typical individuals (matched for their linguistic development) on the production of unmarked CVC syllables, individuals with SLI were significantly impaired on the production of more marked syllables (e.g., CCVC and CCVCC). Similar effects of markedness obtained concerning prosodic structure: SLI individuals and controls did not differ on the production of the unmarked trochee (i.e., strong–weak metrical pattern, such as drɛ-pə), but SLI individuals were impaired on the production of marked iambic (i.e., a weak–strong pattern, as in bə-drɛp). The similarity between this pattern and the one observed with younger, typically developing

individuals (see Chapter 9) could suggest that the difficulty of SLI individuals might reflect developmental delays.

Other studies, however, observed no selective impairment with marked syllable structures (Kavitskaya & Babyonyshev, 2011; Kavitskaya et al., 2011), and several aspects of the error patterns point to a non-grammatical explanation for the impairment of SLI children. A markedness account would predict distinct asymmetry in the pattern of errors: Marked, complex onsets should be simplified (e.g., *block*➔*bock*), but unmarked, simple onsets should be produced correctly (e.g., *book*➔*book*). But as it turns out, children with SLI frequently generate complex onsets from simple ones *(book*➔*blook*), and they are significantly more likely to do so than typically developing children (Marshall et al., 2009). Moreover, the simplification of complex onsets and codas by SLI children is unaffected by their sonority profile (Kavitskaya & Babyonyshev, 2011). In view of these findings, one must either assume that the phonological grammar of certain SLI children is impaired (rather than merely delayed), or postulate a secondary extra-grammatical source for the errors.

These putative extra-grammatical deficits of individuals with SLI are not confined to the production system. SLI has long been linked to a series of auditory deficits (Tallal, 2004; Tallal & Piercy, 1973). To be sure, auditory deficits are not present in all affected individuals (Bishop, 2007), nor can they account for the host of grammatical problems seen in individuals at the time of testing (e.g., Marcus & Fisher, 2003; Marshall & van der Lely, 2005). But it is conceivable that individuals with developmental language disorders could manifest those deficits early in development and overcome them later in life (Galaburda et al., 2006). This possibility is indeed consistent with the observation that infants at risk of language disorders manifest deficits in the processing of auditory stimuli presented in fast succession (Benasich & Tallal, 2002; Benasich et al., 2006; Choudhury et al., 2007). Further converging evidence is presented from animal models of dyslexia (discussed in detail in the next chapter). Although dyslexia is defined as an unexplained difficulty in reading, rather than language specifically, this disorder shares important characteristics with SLI, including a strong deficit to phonological processing and a high rate of association (40–80 percent of the individuals with SLI manifest developmental dyslexia; Scerri & Schulte-Körne, 2010). Like SLI, dyslexia has been linked to auditory processing difficulties (Tallal, 2004), and a rodent model has shown that mutations to genes implicated in dyslexia produce deficits to rapid auditory processing only in juvenile animals, but not in adults (Galaburda et al., 2006; Peiffer et al., 2004).

Whether such early auditory deficits are indeed robust in individuals who are later diagnosed with SLI and whether they can specifically predict grammatical phonological development remains to be seen. But in view of the paucity of evidence regarding the phonological system in SLI, the controversy

surrounding its association with auditory processing problems, and the absence of a specific genetic model, it is presently impossible to determine whether the phonological problems in SLI can be linked to any specific genes, let alone genes that are exclusively "phonological."

11.4 Minds, and brains, and core phonology

In this chapter, we have examined the brain mechanisms mediating grammatical phonological computations and their genetic control. Although very few studies have addressed grammatical phonological knowledge specifically, the available evidence implicates several left-perisylvian sites in various phonological computations, ranging from the identification of phonemes to the distinction between consonants and vowels and the sensitivity to markedness restrictions. Not only are these sites implicated in the processing of phonological structure in spoken language, but several of them are also linked to the phonology of sign language, suggesting the possibility of an amodal phonological network.

Whether this network is dedicated to phonological computation, however, is far less clear. None of the implicated areas is uniquely linked to phonology, and several sites have been shown to mediate musical processing. Similarly, phonological disorders are highly co-morbid with nonlinguistic deficits, including articulatory and auditory impairments, and no known gene, including *FOXP2*, is exclusively dedicated to language or phonology (Marcus & Fisher, 2003).

On the face of it, the failure to individuate a phonological system at the hardware level might appear to preclude any functional specialization for phonology. Many authors indeed adopt a strong hypothesis regarding specialization (see 11.1a). In this view, a cognitive system is specialized only if it can be segregated from other systems at the hardware level – if it cannot be linked to a separate piece of hardware that is exclusively dedicated to that function, then this system does not exist at the functional level. And since no known hardware is exclusively dedicated to phonological computation, the view of phonology as a system of core knowledge must be wrong.

Rather than accepting this conclusion, however, I believe one should question its premise – the requirement of "hardware segregation." Hardware segregation is indeed implausible on both biological and cognitive grounds. Considering, first, the neural level, it is well known that human brains manifest a fair degree of plasticity that allows for the recruitment of existing neural circuits at the service of novel functions. Occipital regions, for example, can be temporarily reallocated for the processing of auditory and tactile information in healthy-sighted individuals who were blindfolded for a period of five days (Pascual-Leone & Hamilton, 2001). These observations, however, do not undermine the fact that visual computations are specialized and distinct in

kind from auditory and tactile ones. Clearly, the segregation of neural hardware is not a *sine qua non* for specialization even at the sensory level.

Similar problems afflict the requirement of hardware segregation at the level of genes. Most complex cognitive functions are regulated by multiple genes (Fisher & Marcus, 2006; Fisher & Scharff, 2009). Few, if any, genes are expressed in a single site, linked to a single phenotypic trait, and many significant evolutionary changes to phenotypic traits can be traced to modifications to the regulation of existing genes (Lowe et al., 2011). Consequently, hardware segregation is not merely *generally* unlikely – it is unlikely even for functions that are demonstrably specialized at the functional level. Perhaps the strongest evidence for the innateness of a functional trait is its natural selection – a proof that the trait has a distinct functional advantage that improves replicability. But traits that undergo natural selection are rarely discrete physically (Anderson, 2010; Marcus, 2006). The well-known propensity of natural selection to tinker with existing genes would suggest just the contrary: Novel neurocognitive systems should share most of their components with older systems. So hardware overlap (in both brain substrates and genes) is not merely consistent with functional specialization – it might be, in fact, its defining feature (Marcus, 2004; Marcus, 2006). While we do not know how core phonology has evolved, it is certainly conceivable that it modified gene networks regulating sensorimotor integration and motor-skill learning (Fisher & Scharff, 2009). The adjacency of phonological sites to the oral/aural brain regions and the comorbidity of phonological and orafacial disorders might well be the relics of the evolutionary history of the phonological system. Such overlap, however, is fully expected by the possibility that phonological hardware is innate and specialized, contrary to the segregation requirement.

But the hardware segregation hypothesis runs into a yet deeper problem at the cognitive level. One should not lose track of the fact that hardware segregation and domain-specificity are apples and oranges of sorts – they concern distinct concepts drawn from distinct levels of analysis. Hardware segregation concerns the topology of brain regions that are grossly implicated in phonological processing. Domain-specific systems, however, are functional systems, defined at the cognitive computational level. While we should certainly expect distinct cognitive systems to be implemented by distinct *computational* brain networks, hardware segregation does not effectively test this hypothesis. Computation is the manipulation of information using procedures that operate on symbols – "physical entities that carry information forward in time" (Gallistel & King, 2009, p. 309). Neuroimaging experiments and brain disorders identify brain sites that are grossly related to phonological processing.

Brain regions, however, are not symbols. Accordingly, the activation of these tells us virtually nothing about how these tissues represent phonological information. Finding that the superior temporal sulcus mediates phonological

processing tells us nothing about how this tissue encodes "syllable," how the structure of the physical entities that encodes syllables indicates the fact that "syllable" is a complex symbol (e.g., it includes an onset and a rhyme), and what procedures are used by the brain to operate on such symbols such that the structure of input symbols will determine the structure of the output.

The problem, of course, is not specific to phonology. As Randy Gallistel and Adam Phillip detail (2009), we do not currently know how the brain computes, primarily because we know of no neural mechanisms that allow for long-term representation of symbols in a manner that would support the decoding of information later in time. Since the hypothesis that the mind has distinct systems (distinct systems of core knowledge, or "mental organs," in Noam Chomsky's words) concerns networks that effect computation, these hypotheses can be currently best evaluated at the functional level, not by the localization of its hardware (see also Barrett & Kurzban, 2006; Poeppel, 2011). Accordingly, the strong hypothesis that infers distinct functions (e.g., mental organs) only if these functions are linked to a discrete anatomical site (11.1b) is not only biologically implausible – it is also cognitively untenable. As Randy Gallistel notes (2007: 2):

> Whether this organ resides in a highly localized part of the brain or arises from a language-specific interconnection of diverse data-processing modules in the brain is irrelevant to whether it constitutes a distinct organ or not. Some organs are localized (for example, the kidney) while others ramify everywhere (for example, the circulatory system). The essential feature of an organ is that it has a function distinct from the function of other organs and a structure suited to that function, a structure that makes it possible for it to do the job.

Some readers might find these conclusions disturbing. One might worry that the approach taken here licenses cognitive theory to ignore findings from neuroscience and genetics. Moreover, by releasing neuroscience from the status of an ultimate arbiter on mental architecture, one loses any hope of finding out whether cognitive systems are "real." And if one worries about individuating cognitive systems, such worries would only multiply for systems of core knowledge – systems that come with the additional conceptual baggage of being innately specified: If such systems cannot be linked to any genes that exclusively regulate these functions, then how can we determine whether the functions are innate?

I do not believe such worries are justified. The inability to reduce cognitive explanations, generally, and cognitive specialization, specifically, to the level of neuroscience and genetics illustrates a well-known difficulty in reducing scientific explanations couched at one level of explanation to lower levels of analysis. This problem is neither new nor unique to cognitive science (Chomsky, 2002; Fodor & Pylyshyn, 1988). But when it comes to the study of the mind, the

problem somehow becomes more pressing. Indeed, mentalistic explanations, in general, and nativist mentalistic accounts, specifically, are subject to great mistrust that is deeply rooted in broader philosophical convictions and sociological factors (Pinker, 2002). Many people, laymen and scientists alike, consider claims about mental architecture as quite tentative. Although we rationally understand that mental capacities are intimately linked to brain functioning, we are nonetheless "surprised" to learn that plain mental characteristics can be identified in the activity of our brains. We can easily hear the difference between a concert violinist and an amateur, but we marvel at the finding that the brains of musicians and amateurs function differently – the high visibility of such findings on the pages of both the scientific and the popular press attests to this fact. Our insistence on the hardware segregation of cognitive functions (e.g., music) is the flip side of the same attitude.

Since our marvel at the operations of our brain is just as irrational as our suspicion of mental explanations, we might all benefit from attending to the origins of these emotional reactions and their role in directing scientific inquiry. While some people might explain their quest for "brain confirmation" by the unobservable nature of cognitive constructs, I suspect that this reaction is much deeper, rooted in our strongly dualist view of the world, and most notably, ourselves (Bloom, 2004). It is precisely because we are at pains to reconcile minds and brains that we find such obvious convergences reassuring. But unfortunately, when it comes to specific functional architecture, neuroscience and genetics cannot currently provide us with decisive answers. So if we are to pursue our quest to unveil the design of the phonological mind, then brain and genes cannot serve as the ultimate arbiter.

Relinquishing the decisive status of neuroanatomical evidence, however, does not mean that questions of mental architecture are unfalsifiable or impervious to external evidence. The large body of research reviewed in previous chapters demonstrates how one can evaluate the status of phonological primitives and constraints against a host of behavioral evidence, ranging from the distribution of such structures across languages, their role in linguistic processes, and their status in psychological experiments with humans – both infants and adults. Moreover, findings from neuroscience and genetics can certainly inform our understanding of how cognitive systems are implemented. Consider, for example, the role of markedness scales, such as sonority, or place of articulation. Although such proposals make no specific claims on how markedness is represented in the brain, they do predict that changes in markedness along any given markedness scale should result in a monotonic change in the activation of a single network (as opposed to non-monotonic changes across multiple circuits). Similarly, while core phonological knowledge may be linked to genes that are involved in multiple functions, one would expect to find

hereditary conditions that would affect phonological competence. While there are numerous ways in which predictions associated with domain-specificity can be evaluated against neuroanatomical and genetic data, the final verdict on whether the system is specialized must rely primarily on internal evidence. Ultimately, domain-specificity in cognition is a functional question, and functional questions can only have functional answers.