The adoption of sound change : synchronic and diachronic processing of regional variation in Dutch
Voeten, C.C.

Cover Page

Universiteit Leiden

Author: Voeten, C.C.
Title: The adoption of sound change : synchronic and diachronic processing of regional variation in Dutch
Issue Date: 2020-10-13

Conclusion

7.1 The Polder shift and its adoption

The goal of this dissertation was to find out what factors influence the adoption of sound change, using the Polder shift as a particularly suitable case study of on-going sound change. Chapter 2 investigated the diatopic diffusion of the Polder shift, which was the subject of the first research question in this dissertation: **what is the synchronic diatopic diffusion of the sound changes involved in the Polder shift?** This question was important for practical reasons—to be able to select representative participants for the experiments in the following chapters—but also for theoretical reasons, namely to get a clearer picture of the natures of the four sound changes that are involved in the Polder shift. The corpus study revealed that the phonetic changes diphthongizing /eː,øː,oː/ and lowering /ɛi,œy,ɔu/ are Neogrammarian, whereas the phonological change blocking diphthongization before coda /l/ is based on exemplars. The fourth change, the vocalization and retraction of coda /l/, is of indeterminate status. The synchronic diatopic variation of these changes was found to be mostly homogeneous, with the four changes having all but completed in the Netherlands, but having reached very little of Flanders.

These results paved the way for Chapter 3, the purpose of which was to answer the second research question: (**how**) **do sociolinguistic migrants adopt the Polder shift?** The results from the teacher-corpus study suggested that the ideal participants would be sociolinguistic migrants from Flanders to the

Polder-shift area. Chapter 3 followed ten such sociolinguistic migrants for a total duration of nine months after they had just moved to the Netherlands, and used psycholinguistic experiments to investigate whether and how much they adopted the Polder-shift changes and a control sound change (the realization of coda /r/ as [ɹ]). The results reproduced the differences found in the teacher-corpus data, but were not found to diminish over time; in other words, adoption of the Polder shift by the sociolinguistic migrants was not found.

Chapter 4 followed up on Chapter 3 by investigating the hypothesis that the negative findings in Chapter 3 were not due to failure of the experiment, but because nine months' time may simply not have been long enough to adopt the Polder-shift changes. Previous research has shown that similar changes can be adopted in approximately the same time frame (Evans & Iverson 2007), and hence a larger-scale followup was strongly warranted. The research question answered in Chapter 4 is: **which individuals, after how much time, are more likely to adopt the Polder shift?** This research qustion was answered by means of a large-scale cross-sectional comparison. For this chapter, 18 sociolinguistic migrants were found, who had lived in the Netherlands for various amounts of time ranging from three years to more than two decades. They were compared to suitable control groups of 45 individuals who had lived in the Randstad area of the Netherlands their whole lives and 43 individuals who had lived in Belgium their whole lives. Mixed results were found between production and perception. In production, the migrant group as a whole had moved to be positioned precisely in between the Netherlandic and Flemish control groups. An analysis at the level of the individual showed that this effect was driven the most strongly by ten of the eighteen migrants, who had adopted the Polder shift to such an extent that a cluster analysis grouped them with the Netherlandic control group, rather than the Flemish one. These findings did not directly carry over to the perception data: in perception, expected group differences were found similar to those in production, but an analysis at the individual level did not yield clear results, although its results were partially correlated with the individual-level results in production. This agreed with findings by Evans & Iverson (2007), and suggested that it is not so much the *type* of sound change (phonological change vs. phonetic change) that determines its adoption and subsequent propagation by individuals, but rather its mode of transmission (system-internal vs. contact-driven).

Chapters 5 and 6 probed more specific aspects of the (non)adoption of the Polder shift by the same ten sociolinguistic migrants from Chapter 3 by means of two ERP experiments. Together, they answer the fourth and final research question: (**how**) **is the adoption of the Polder shift reflected in ERPs?** Chapter 5 resolved an open problem regarding the relative roles of perception and representation in the adoption of sound change: is sound change based in misperception (e.g. Ohala 1981) or in misrepresentation (Beddor

2009, Hamann 2009)? The results of this chapter show that the sociolinguistic migrants start out having weaker knowledge about the phonological distribution of the [eː~ei] allophone difference, although they do catch up to the Netherlandic level between the two sessions of the experiment, respectively corresponding to their fourth and their ninth months in the Netherlands. In addition, the Flemish participants did not find the [ɹ] realization of the rhotic as attention-grabbing as the Netherlandic participants did. It was argued that these results reflect phonological and sociolinguistic knowledge, rather than differences in phonetic processing. Thus, they do not lend support to the misperception account of sound change (which is characterized by differences in phonetic processing) but can be explained more readily in terms of the accounts by Beddor (2009) and Hamann (2009) based on differences in cue weighting (in this case, differences in the weighting of phonological and sociolinguistic information).

Chapter 6 used a more exploratory ERP experiment. As with Chapter 5, this experiment compared the perception of Netherlandic and Flemish realizations of the vowels involved in the Polder shift, plus two control conditions consisting of the rhotic and a phonologically (as opposed to just sociolinguistically) illicit vowel realization (/ɛ/ realized [ɛː]). Again, the Polder-shift changes produced no significant ERP differences, nor did the rhotic, but the /ɛ/ control condition resulted in a P600, in the control group only. The P600 was argued to be a logical extension of previous work, particularly by Witteman et al. (2015), who found a behavioral slowdown in similar conditions involving regional accents where a vowel's accented realization crossed a phoneme boundary. If the P600 indeed represents the electrophysiological precursor to this behavioral result, this automatically explains why the same P600 was not found in the Polder-shift conditions or in the rhotic: these all involved changes in *allophones*, not phonemes.

7.2 From compensation to adoption

The results from Chapter 3 and Chapter 4 portray an important contrast. While Chapter 3 failed to show any credible adoption of the Polder shift by the Flemish group in nine months' time, Chapter 4 showed that after multiple decades, some of the Flemish participants in that study had adapted. Taken together, Chapters 3, 5, and 6 provide a chain of evidence that explains what steps are to be taken in the process of adopting an on-going sound change. These chapters discuss the same participants in different experiments. The longitudinal experiments in Chapters 3 and 6 were performed after roughly one, five, and nine months after arrival of the Flemish participants in the Netherlands. The experiment in Chapter 5 was performed during the latter two of these three

occasions: around month five and around month nine. Chapter 5 provided evidence that suggests that there are three independent sources of knowledge to be acquired: listening competence, phonological knowledge, and knowledge about sociolinguistic evaluation. The findings suggest that these three types of knowledge are acquired in a particular order: listening competence first, phonological knowledge secondly, and sociolinguistic evaluation thirdly. This corresponds to these three skills being located at increasing levels of grammatical abstraction. The presence of MMNs indicates that listening competence is already in place by four months after arrival; since these are sound changes and not second languages, it is not unreasonable to assume that this was never an issue to start with for the Flemish students. Specifically concerning the [e:∼ei] allophone pair, passive (receptive) phonological knowledge is acquired sometime between months four and nine. For the rhotic, the required allophonic knowledge appeared to be already in place (as evidenced by the Flemish group reaching the same mismatch negativities as the Netherlandic group in Chapter 5, but also by their high proportion of [ɹ] responses in the perception task in Chapter 3). The Flemish group's sociolinguistic knowledge of the rhotic, however, was not yet completely Netherlandic-like: the rhotic captured less attention (evidenced by a shift of the MMN away from the frontal pole compared to the Netherlandic controls), suggesting that this sound may not have imparted the same sociolinguistic salience in the Flemish group as it did in the controls.

The EEG task in Chapter 5 found that the grammatical knowledge of the Polder shift was largely in place for the Flemish participants, and improved even further by their ninth month of living in the Netherlands. It is telling, then, that this adoption of the Polder shift found electrophysiologically was in no way reflected in the behavioral experiments in Chapter 3, where the differences between the two participant groups were found to be robust and persistent. It was only after much more time than nine months, viz. the timespan covered by Chapter 4, that the same adoption of the Polder shift was also demonstrated behaviorally. The results from Chapters 4 and 5 have shown that these participants are certainly able to acquire these changes eventually, so then there must be competence-extrinsic factors at play which cause the behavioral adoption of the Polder-shift changes to proceed more slowly than the changes in, for example, Evans & Iverson (2007), which departed from a very similar situation and did find adoption of relevant accent differences. Chapter 3 argued that the reason can only be that nine months is simply not enough time to adopt these language changes in a behaviorally-detectable way.

It is possible that the reason for this lies in the phonological status of the changes of the Polder shift. The Polder-shift changes do not result in phonemic mergers or splits, which could well have reduced the pressure on the Flemish participants to adopt them behaviorally. Chapter 6 provides concrete evidence of this type. This exploratory study did not reveal long-term changes, but this

Table 7.1: Log$_{10}$ Bayes factors for the hypothesis that the P600 difference between the groups equals or exceeds that found in the [ε~εː] condition from Chapter 6. Negative values indicate evidence against this hypothesis.

Condition	Bayes factor (log$_{10}$)
[ei~eː]	− 1.93
[øy~øː]	− 1.00
[ou~oː]	− 2.18
[εi~εː]	− 2.03
[œy~œː]	− 1.57
[ɑu~ɑː]	− 3.10
[aːɹ~aːʀ]	− 2.62
[ε~εː]	0.00

was because it turned out that the task used in the experiment (which was simply the production task from Chapter 3) was not sensitive to allophonic violations. This was concluded because a control condition with a phonemic violation, viz. the realization of /ε/ as *[εː], elicited a clear P600, but the experimental conditions involving the Polder-shift changes did not.

Note that this null result does not yet prove that the manipulation used in Chapter 6 does not elicit a P600 for non-phonemic changes; it only shows that this was not found in the experiment reported in that chapter. However, a reanalysis of that chapter's results in a Bayesian framework confirms that the other conditions indeed *did not* elicit a P600, rather than simply having failed to do so. For this analysis, each condition was averaged over the 560–660 ms window, just as the [ε~εː] condition was. R package brms (Bürkner 2017, 2018) was used to fit a model containing the same terms as in Chapter 6's model for the [ε~εː] condition, with fixed-effect priors set to the same values as those obtained from that model. Table 7.1 shows log$_{10}$ Bayes factors for the hypothesis of a P600 difference between the Flemish-Dutch and Netherlandic-Dutch students of at least −2.51 μV. For the [ε~εː] condition, the evidence for this hypothesis is exactly as strong as the prior (which is obvious, given that the priors were set to precisely this condition's Chapter-6 results), but all of the other conditions provide "strong" (Jeffreys 1961) to "decisive" (Jeffreys 1961) evidence that there is *no* P600 of at least this magnitude to be found.

Thus, the results from Chapter 6 show that participants compensate for allophonic variation in on-line auditory speech processing, and fail to do so when this variation crosses a phoneme boundary (this was the case for the [ε~εː] condition, which elicited the P600). For the allophonic changes that constitute the Polder shift, the P600 is not informative: it was absent in both the

NDS and the FDS, despite participants' ability to perceive these changes (Chapter 5) and them not (yet) having had adapted to them behaviorally (Chapter 3). However, the P600 may prove to be informative for sound changes involving phonemic mergers or splits. The significant NDS–FDS difference in P600 amplitude in the [ɛ∼ɛ:] condition shows that the FDS were not yet familiar enough with NDS speech in general to perceive the phonemic violation. The FDS' diminished N400 amplitudes compared to the NDS's point in the same direction. This identifies an avenue for future research: replicate the experiment using ongoing sound changes involving phonemic mergers or splits, instead of the allophonic changes that are central to the Polder shift. For those kinds of changes, the N400 and P600 could be informative of the degree to which the phonemic change has been adopted.

7.3   Salience

Summarizing the previous section, the results of this dissertation imply an important role for *salience*. This is a well-known term within sociolinguistics—in fact, it is the foundation of Trudgill's (1986) change-by-accommodation model (cf. the results in Chapter 3, which are somewhat problematic for this model)—but it has not been well-defined: see Rácz (2013) for a discussion. Auer, Barden, & Grosskopf (1998) observe why salience is so hard to define: salience is a highly subjective attribute, driven by personal-evaluative factors rather than objective phonetic or phonological parameters. Accepting this, this dissertation has oftentimes stated that the rhotic changes are more salient than the Polder-shift changes, referencing Sebregts (2015) when doing so. The results from this dissertation support this distinction, and also support an explanation in terms of different types of salience.

Chapter 5 found that the Flemish participants' perception of the [e:∼ei] allophone distinction changed over time, but their perception of the rhotic distinction did not. In Chapter 6, the non-NDS realization of /ɛ/ as [ɛ:] triggered a P600, but the equally un-NDS-like realization of the diphthongal Polder-shift vowels as monophthongs before non-/l/ did not. These results show that there are two types of salience to be taken into account. The first is *sociolinguistic* salience, the well-known type which is the focus of authors like Auer, Barden, & Grosskopf (1998). This was observed in Chapter 5, where it was shown that the [əɹ] deviants impart less attentional importance for the Flemish group than they do for the Netherlandic group. This is in line with the well-known observation that, for Netherlandic-Dutch speakers, this realization of /r/ is an extremely salient sociolinguistic marker (Sebregts 2015). The results from this dissertation also identify a second type of salience: *phonological* salience. Chapter 6 showed that the P600 is sensitive to phonemic status: violations that

cross phoneme boundaries elicit the P600, violations at the allophonic level do not. This may explain why Chapter 3 failed to find behavioral adoption of the Polder-shift changes: if these changes are not salient enough to the brain, why would the Flemish participants need to adopt them? It must be noted that Chapter 5 did find changes in brain responses to the [e:∼ei] allophone distinction, but this used a mismatch-negativity paradigm, which is pre-attentive, whereas the P600 is post-attentive. Thus, the P600 observed for phonemic violations in Chapter 6 may be indicative of a type of salience that is not primarily sociolinguistic, but rather phonological.

Is it correct to consider phonological category status a type of *salience*, as posited in the previous paragraph? The answer is most likely: yes, in the context of sound change a categorical phonological change may be salient. The argument is provided, indirectly, by Janson (1983). In his study of the sound change from [r] to [ʀ] in Norwegian, Janson argues that this change must have been a change in the underlying form. Note that such a change is not necessarily phonemic: in Janson's (1983) case, it is not a change in the meaning-distinctive-category system of the language, but rather a rule inversion of the type discussed by Hyman (1976). Such changes need to be sufficiently disruptive to the phonological system in order to be adopted, i.e. a change is more likely to be adopted if it is phonologically salient. The realization of /ɛ/ as [ɛ:] in Chapter 6, which elicited the aforementioned P600, is exactly the type of change discussed by Janson (1983). Of the Polder shift, however, we know that it does not meet this criterion of phonological salience: Chapter 2 showed that these changes are either Neogrammarian or exemplar-based, not underlying-form changes. It is possible that this is why these changes did not elicit a P600 in Chapter 6, and why they could not be shown to be adopted within nine months' time in Chapter 3. Their eventual adoption after a longer amount of time has passed (Chapter 4) may then be due to some of the sociolinguistic migrants eventually noticing the differences between their productions and the input they receive (for which the cognitive machinery is already in place; Chapter 5). In this case, these individuals adopt the sound change not because it is intrinsically salient (as was the case for the realization of /ɛ/ as [ɛ:] in Chapter 6), but because it has become salient *for them*.

7.4 Methodology

Section 1.3.2 discussed how psycholinguists have profited from methodological innovations. This dissertation has demonstrated how the study of on-going sound change can reap the same benefits. Specifically, this dissertation has made three types of advances: new methods were developed, old methods were given new uses, and—where necessary—this dissertation used

methods beyond the traditional linear-regression model (which underlies, e.g., ANOVA). It is worth reflecting on how these innovations have aided the interpretation of the data collected in this dissertation, and how they could be applied in other settings than the adoption of the Polder shift.

The only *new* method in this dissertation was Chapter 2's use of the generalized additive mixed model (henceforth "GAMM") to avoid having to segment the gradient boundary between a vowel and a following coda /l/. The backbone of segmental acoustic phonetics has always been the ability to isolate the segment under investigation from the surrounding context, a requirement which is impossible to meet when the segment transitions are fully gradient. This is the case when the coda /l/ is strongly vocalized, which was found to be the case for the Netherlandic varieties of Dutch, but not the Flemish ones (see Section 2.3.4). Chapter 2 used GAMMs to model these VC trajectories as they are, without requiring an *a priori* manual segmentation.

The second type of methodological innovation in this dissertation was the use of well-established methods, but in novel ways. The most significant example is given in Chapter 4. Chapter 4 made use of the mixed-effects model—a completely uncontroversial statistical technique—to capture individual variation. This is an unconventional use of the mixed-effects model: traditionally, random-effect parameters are considered nuisance variables, to be incorporated into a model to absorb variation between participants and items that may interfere with the group-level patterns which are of primary interest. Chapter 4 demonstrates that, for these data, that approach would have been naïve: when the groups are not perfectly homogeneous, as was the case for the sociolinguistic-migrant group in Chapter 4, an analysis that collapses each group into a single β value will misrepresent the data. The data in Chapter 4 were analyzed more appropriately by excluding the (misleading) fixed factor for "Group" from the model, and reconstructing the groups *a posteriori* on the basis of the empirical BLUPs from a full random-slope model. A cluster analysis on the resulting by-participants b coefficients revealed that the sociolinguistic-migrant group was split between more and less innovative participants, which *together* caused the group as a whole to move in between the Flemish and Netherlandic control groups. The degree of adoption of the Polder shift was quantitative, i.e. the sociolinguistic migrants did not differ in whether or not they had adopted the Polder shift categorically, but in the *degree* to which they had done so. If this degree exceeded 0.07 Lobanov units, participants had adopted the Polder shift to a sufficient extent that they were classified as Netherlandic rather than Flemish. Within sociolinguistics, this BLUP-based approach to studying individual variation has slowly begun to take off (Drager & Hay 2012, Tamminga to appear); this dissertation's use of cluster techniques to shed more light on the meaning of these individual differences is only an additional step in these on-going methodological developments.

The interpretation of the data discussed in this dissertation has been made significantly more feasible by the adoption of statistical methods beyond the simple linear-regression model. This was the case in Chapters 3, 5, and 6. Chapters 3 and 5 both dealt with situations where an extremely large number of all-categorical predictors would be required if these chapters' perception data were to be analyzed using traditional regression models. In Chapter 3, the maximal model for a logistic-regression analysis would have contained 48 regression coefficients for each condition, including a four-way interaction "Group (2 levels) × Session (3 levels) × Following consonant (2 levels) × Step (4 levels)". In Chapter 5, the maximal model for an ANOVA like the one in Grosvald & Corina (2012) would have contained 72 coefficients for each of the six conditions, among which is a five-way interaction "Deviant (2 levels) × Group (2 levels) × Session (2 levels) × Hemisphere (3 levels) × Anteriority (3 levels)". Such interactions are impossible to interpret. Section 3 resolved this by making use of mixed-effects regression trees to model exactly this type of data in a much more interpretable way. Similarly, Section 5 resorted to GAMMs, which not only removed the need for "Hemisphere" and "Anteriority" factors, but also provided a much more fine-grained overview of the data in the first place. To the author's knowledge, this has not been done in linguistics before.

Methodological innovations can take place not just in the actual modeling process itself, but also in the reasoning about models that are possible for a given dataset. In Chapter 5, six separate GAMMs were run corresponding to the six different conditions in the experiment, but the MMN ERP was only expected in three of them, given the normally asymmetric nature of this ERP component (Cornell, Lahiri, & Eulitz 2011, Lahiri & Reetz 2010). Therefore, a way was needed to be able to argue not just *against* the null hypothesis, but also in its favor. For this reason, models were evaluated in a Bayesian framework, though still using maximum-likelihood estimates for reasons of computational feasibility, using the approach by Wagenmakers (2007). Bayesian reasoning brings with it a completely different way of thinking—evaluating the likelihood, $p(\beta|y)$, rather than the p-value, $p(y|\beta)$—which is not yet commonplace within linguistics. As another example, Chapter 6 was an exploratory ERP investigation, and hence it was not known *a priori* what ERP component, if any, would be obtained. This is a solved problem in the field of cognitive neuroscience, where permutation testing (Maris & Oostenveld 2007) is the canonical answer. This nonparametric statistical test made it possible to identify both the temporal window and the spatial ROI in which robust differences due to the various manipulated factors arose. This led to the identification of a new putative marker of phonological status, viz. the P600, which can be triggered by phonological violations that cross phoneme boundaries, but could not be detected for within-category violations.

7.5   Conclusions

This dissertation has investigated the adoption of sound change, and the role played by synchronic and diachronic processing of variation within that process. The investigation focused on the Polder shift, an on-going vowel shift in Dutch that has all but completed in the Netherlands, but has not taken place in Flanders (Chapter 2). The adoption of the Polder shift by Flemish sociolinguistic migrants proved to be difficult to detect in the medium term (nine months) using behavioral methods (Chapter 3), but was detected behaviorally in the long term (multiple decades; Chapter 4). Despite the lack of reliable medium-term behavioral evidence, Chapter 5 found robust evidence for the start of sound-change adoption using an EEG experiment. Chapter 6 attempted to extend these findings using a different, novel, paradigm, and found that this was successful, but only for changes that were large enough to cross a phoneme boundary, i.e. not the Polder shift.

The future of sociolinguistics must be sought in the continuing integration of the five fields of historical phonology, sociolinguistics, psycholinguistics, neurolinguistics, and statistics. The different findings from this dissertation demonstrate how the fundamentally sociolinguistic phenomenon of historical sound change can be studied empirically using psycho- and neurolinguistics. The dissertation has additionally demonstrated on several occasions how much the study of on-going sound change can profit from the continuous innovations in the field of methodology and statistics. These have made it possible to analyze data that would previously have been considered unanalyzable (Chapter 2), and have additionally made it possible to draw new conclusions (Chapters 3, 4, 5, and 6). From this dissertation alone, a few possible avenues for future research present themselves. Chapter 2 concluded with the remark that synchronic evidence cannot distinguish between a sound change that is phonetically abrupt and a sound change that has been phonetically gradual but has already completed. The clear Netherlandic–Flemish split on the effect of coda /l/ on the preceding vowel is one such case: future diachronic research is needed to chart exactly how the Netherlandic F2 retraction before coda /l/ developed. An additional remark that was made on the data in Chapter 2 was the low number of words available in the corpus, which might have caused the lexical diffusion of the Polder-shift changes to have been underestimated. Ample future options for synchronic research like Chapter 2's corpus study present themselves here.

As mentioned before, Chapters 3 and 4 supplement each other, in that the former chapter did not find adoption of the Polder shift after nine months, but the latter chapter did find it after multiple years. The obvious research question following from this discrepancy is: when *does* adoption take place? This is not an easy question, especially as Chapter 4 showed that adoption, at the group

level, is gradient, rather than categorical. In addition, this dissertation demonstrated that whether or not one finds adoption of a sound change strongly determines on how one defines "adoption": behaviorally, Chapter 3 did not find any clear adoption of the Polder shift after nine months, but electrophysiologically, Chapter 5 did reveal significant changes. For these two reasons, I would discourage future researchers from devoting their time to the pursuit of broad questions such as "when does adoption take place?". Instead, a more focused and more-thoroughly-operationalized question such as "what is the earliest point in time at which a single sociolinguistic migrant adopts the Polder shift in single-word production?" is more likely to result in positive research outcomes; in the case of this example question, it would identify the empirical amount of time that is *minimally* necessary to adopt the Polder shift. Another point which would be interesting for future study is the question "what is the earliest point in time at which 21.7% of sociolinguistic migrants have adopted the Polder shift in single-word production?", where 21.7% is the critical mass calculated by Yang (2009) for an individual-level sound change to secure its evolution into a group-level sound change.

The neurolinguistics of sociolinguistic variation also deserve further exploration. Chapter 6 revealed a P600 for variation that was particularly salient, viz. crossing the boundary of a phonemic category. This could prove to be a new method for detecting the status of a sound change in progress, with phonemic mergers or splits putatively eliciting a P600, but this needs to be established by research specifically looking into this ERP component as an indicator of phonological status. In addition, this ERP may be the starting point for an objective definition (cf. Auer, Barden, & Grosskopf 1998) of the vague notion of "salience". Future research could explore this further. The same is true for the MMN, of which the topographical distribution was argued in Chapter 5 to index a type of sociolinguistic salience. This, too, needs to be investigated in a more specific manner.

Future research could also proceed from this dissertation in the direction of new methods for investigating linguistic variation. Chapter 2's application of the generalized additive model could be combined with dynamic time warping (Shi et al. 2015) to develop new tools that could aid phoneticians in determining empirical boundaries for speech sounds that are difficult to segment. The method used to investigate individual differences in Chapter 4 could also be extended to GAMMs to directly study individual variation in more complex signals than Chapter 4's point measures. Steps in this direction have already been taken by, for instance, Tamminga, Ahern, & Ecay (2016). It should go without saying that GAMMs in general offer new ways of analyzing data that would have been challenging to analyze in a more traditional way; Chapter 5's use of GAMMs to avoid fitting 6-way interaction models, by smoothing over the entire topographical area present in the data (rather than including many-leveled factors for "Hemisphere" and "Anteriority"), is a prime example.