The emergence of rhotic vowels in Quebec French: a change from below?

Massimo Lipari

McGill University, Canada
massimo.lipari@mail.mcgill.ca

ABSTRACT

Quebec French is undergoing a rarely observed kind of sound change whereby the front mid rounded vowels become rhotic (produced with a bunched tongue or retroflex gesture resulting in low F3, like English /ɚ/). Previous work has suggested that this is a change from below—i.e., one which speakers are not conscious of and which is thus not socially marked. These previous studies, however, have not had the breadth to robustly demonstrate this finding throughout the entire speech community, and the extent of interspeaker variation (especially according to age, gender, and geographic origin) remains largely unknown. This study, which employs a sizeable corpus of parliamentary speech, adds to the documentation of the change, while directly examining the effects of the aforementioned predictors. The findings are consistent with the hypothesis of change from below (with some caveats), although additional data from new speakers is likely needed to draw any strong conclusions.

Keywords: rhotic vowels, change from below, Quebec French, sound change, corpus phonetics.

1. INTRODUCTION

Rhotic vowels, like English /ɚ/, are a rare class of sounds, occurring in less than 1 percent of the world's languages [1]. Their retroflex or bunched tongue articulation creates their signature low third formant [2]. In Quebec French, rhotic-sounding realizations of the mid front rounded vowels /ø/ and /œ/ were first noted by Dumas [3] more than 50 years ago in the speech of a handful of young working-class male Montrealers—particularly in final open syllables and before rhotic codas. This discovery, however, received little attention until a pair of studies by Mielke [4, 5], which found a change in progress also affecting the /œ̃/ phoneme.

The first of these is a corpus study of 75 speakers from the Gatineau-Ottawa region (~170 km west of Montreal, straddling the Quebec-Ontario border) [4]. A significant fall in F3 was found (in apparent time) for each of the three vowels, beginning with speakers born in the 1960s—the drop occurs slightly earlier and steeper for /ø/ than for the others. The effects of gender and socioeconomic class were not significant.

A follow-up production study with 23 speakers from the same region (save for a few) [4, 5] then showed, via ultrasound imaging, that there exists a high degree of variability in how the mid front rounded vowels are realized articulatorily, both between and within speakers. That is, varying degrees of tongue bunching (the much more commonly used gesture compared to retroflexion) were observed, with tokens lying on a continuum from not bunched (fully non-rhotic) to English /ɚ/-like (fully rhotic).

The non-significance of the effects of the social predictors leads Mielke to reject Dumas' (somewhat speculative) suggestion that rhotacization results from deliberate borrowing from English for stylistic or prestige reasons.[1] Instead, he suggests it is a change from below: a kind of change which speakers are not conscious of during its progression, and which is accordingly not socially marked (cf. [6]). He also cites in favour of his view anecdotal evidence of native speakers generally failing to perceive a salient difference between rhotic and non-rhotic productions of the vowels when asked directly. Mielke proposes a perceptual motivation for the change, noting that rhotacization may serve to enhance the acoustic cues produced by lip rounding.[2] Rhotacization would thus serve to increase perceptual distance between the front mid rounded vowels and their unrounded counterparts /e, ɛ, ɛ̃/ in a crowded vowel space.[3]

While the theory of change from below is certainly consistent with Mielke's results, the studies have some limitations. For one, the model reporting is sparse. In particular, nothing is said about the random effects: those results, however, could be illuminating, given the finding of much inter-speaker variability. For another, the corpus used is geographically circumscribed. Since we know from Dumas that rhotacization is present in other regions of Quebec, data from these other areas would allow us to test how robustly the theory of change from below holds up across the speech community.

The research questions of the present study are thus twofold. The first set is descriptive: does the state of rhotacization look similar throughout Quebec to the situation in the Gatineau-Ottawa region? The second aims to determine whether the evidence from other dialects is consistent with a change from below. Were this the case, we may expect to find no effect of gender, and either no effect of dialect or perhaps some evidence of gradual, wave-like diffusion throughout the territory (since a change from below would not be consciously associated with any particular group).

2. DATA & METHODS

The data used in the present study are from a combination of two corpora of proceedings of the National Assembly of Quebec (ANQ), Quebec's national parliament. The first corpus was developed by Milne [9]: it consists of approximately 9 hours of recordings from 61 parliamentarians,[4] which was previously force-aligned with purpose-built software. The other is a new corpus in the same style, developed for this study, containing almost 7 hours of additional data from 26 speakers (including some of the same speakers as Milne's) and aligned using the *Montreal Forced Aligner* [10]. Neither corpus originally contained biographical data (other than speaker names): gender, year of birth and place of birth were thus obtained manually from the ANQ website [11]. The first two predictors were added as-is, whereas place of birth was used to construct a ternary DIALECT predictor. The traditional assumption that Quebec is roughly divided into two dialect regions, West and East, by a line perpendicular to the Saint Lawrence River somewhere between Montreal and Quebec City (cf. the discussion in [12] on the historic apical vs. dorsal rhotic isogloss) was followed, and speakers in the ambiguous region between those cities were coded as speaking a "Central" dialect.[5] Non-native speakers and speakers with missing biographical data were excluded. The remaining 70 speakers are distributed by gender and dialect per Table 1.

	West	Central	East	**Total**
Female	11	4	8	**23**
Male	22	5	20	**47**
Total	**33**	**9**	**28**	**70**

Table 1: Breakdown of speakers by gender/dialect.

Formants were measured at regular intervals using the "refined" formant analysis function of the Python package *PolyglotDB* [13]. This algorithm, described in [14], iteratively searches for the optimal parameters (number of LPC coefficients and frequency ceiling) to use with the *Praat* [15] formant measurement function on a per-phoneme and per-speaker basis by comparing the results at each step to a set of language-specific input vowel prototypes.[6] The measured F3 values were then vowel-extrinsically normalized using the Neary2 method [16], with the modifications suggested in [17].[6] Phones less than 50 ms in length and/or from stop words were discarded, leaving a total of 5,438 tokens: 1,885 of /ø/, 1,812 of /œ/, and 1,741 of /œ̃/. The first and last 20 percent of duration of each token were then discarded to avoid coarticulation effects from preceding or following segments, and finally only the minimum F3 value was retained for modelling (for the sake of comparability, as this closely mimics Mielke's methodology in [4]).

A mixed-effects linear regression model of minimum F3 was fitted using the *lmerTest* package in *R* [18, 19]. In addition to predictors for GENDER, YEAR OF BIRTH (YOB), and DIALECT, as well as a ternary PHONEME predictor distinguishing between the three vowels, a pair of controls were added. First, SYLLABLE POSITION (σ POS) marks out word-final syllables from others, since these bear accent in the language. Second is a control for following CONTEXT: this is needed because four consonants (/ʁ, v, z, ʒ/) are known to trigger an allophonic lengthening and diphthongization process which may interact with rhotacization (cf. [20]). Since the rhotic may by its very nature have a particular effect on F3, the factor has three levels: ʁ, other lengthening C, and other. Rounding out the fixed effects are a series of interactions terms—in particular, GENDER:YOB, DIALECT:YOB, and GENDER:DIALECT, as well as all binary interactions between PHONE and the social predictors. As for random effects, by-speaker intercepts and slopes for PHONE and CONTEXT, as well as by-word intercepts and slopes for YOB (controlling for lexical effects), were included.

All binary predictors were centred around the mean and scaled by 2σ. Ternary predictors used weighted Helmert contrasts (due to the unbalanced nature of corpus data). Accordingly, the contrasts of PHONE represent /ø/ vs. /œ/ and nasal vs. non-nasal; those of DIALECT represent West vs. Central and non-East vs. East; those of CONTEXT represent /ʁ/ vs. /v, z, ʒ/ and lengthening C vs. other.

3. RESULTS

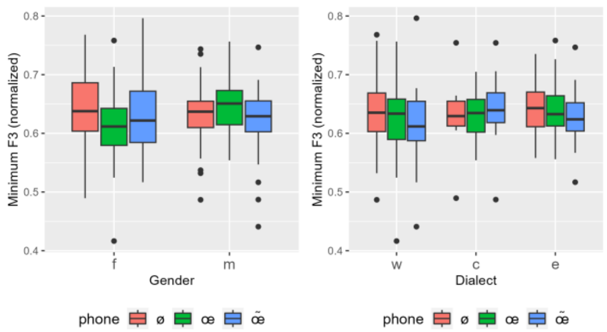

Figure 1: Empirical plots of GENDER (left) and DIALECT (right) by PHONEME, using speaker means.

Before examining the model predictions, it is useful to consider the empirical distribution of the data along the dimensions of interest.[8] In the following discussion and figures, by-speaker averages were first computed in order to prevent statistics from being skewed by a few prolific speakers.

The average minimum F3 values across the dataset for each phoneme are quite similar, at 0.64 for /ø/, 0.63 for /œ/ and 0.62 for /œ̃/, with equally similar standard deviations of 0.056, 0.055, and 0.058, respectively—indicative of a sizeable amount of inter-speaker variation. As seen in Figure 1, however, these figures obscure some differences between the GENDER and DIALECT groups. While the mean of /œ̃/ is not clearly different for women and men ($\mu = 0.63 \pm 0.13$ vs. $\mu = 0.62 \pm 0.10$), there do appear to be differences in the other two vowels: for women, /ø/ appears to be slightly higher than /œ/ (0.64 ± 0.13 vs. 0.61 ± 0.14), while this pattern is reversed for men (0.63 ± 0.10 vs. 0.64 ± 0.08)—although note the large confidence intervals. As for DIALECT, West and East look quite similar, both exhibiting the pattern /ø/ = /œ/ > /œ̃ (0.63 ± 0.12, 0.63 ± 0.12, 0.61 ± 0.12 vs 0.64 ± 0.08, 0.64 ± 0.09, 0.63 ± 0.09), albeit with there being somewhat less variability in the East. The Central dialect, conversely, shows very little difference between the phonemes (0.64 ± 0.14, 0.63 ± 0.09, 0.64 ± 0.09 for /ø/, /œ/ and /œ̃/, respectively).

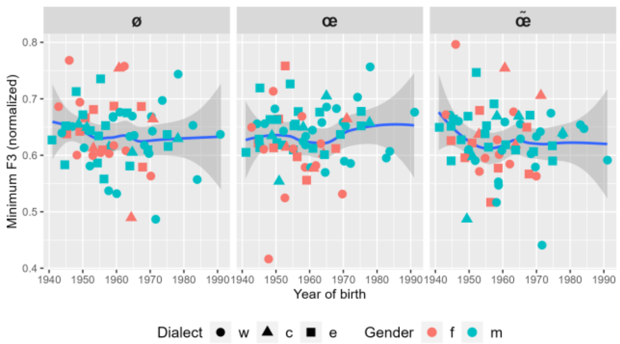

Figure 2: Empirical plot of YOB vs. speaker-averaged minimum F3, faceted by PHONEME.

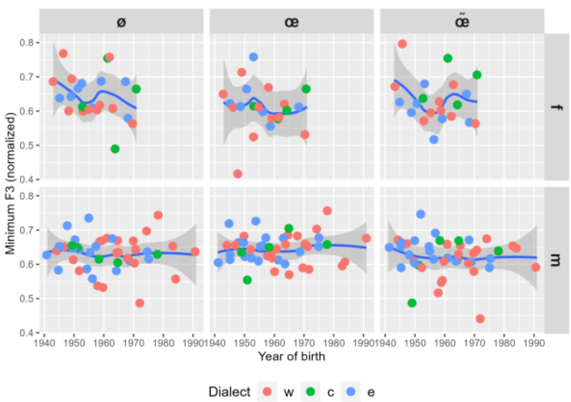

Figure 3: Empirical plot of YOB vs. speaker-averaged minimum F3, faceted by GENDER and PHONEME.

Figures 2 and 3 show individual speaker means for each phoneme (with points labelled by both gender and dialect), as well as their evolution through apparent time. Here, we see that all dialect and gender groups have both very rhotic and very non-rhotic speakers. Across the whole dataset, the YOB smooths by phoneme (Figure 2) do not reveal much in the way of change: there may be a general downward trend (at least for /ø/ and /œ/), although this is very uncertain (largely due to the lack of speakers towards the end of the range, after about 1975). When broken down by GENDER (Figure 3), the pattern becomes a bit clearer: the results are suggestive of a decrease across all three phonemes for women (although the confidence intervals still include 0), whereas no change is manifest for the men. No pattern over time across dialects emerges.

Fixed effects				
Coefficient	Est.	SE	*t*	*p*
(Intercept)	0.633	0.0072	88.08	< 0.001
PHONEME1	0.026	0.0099	2.65	0.009
PHONEME2	-0.013	0.0137	-0.91	0.362
YOB	-0.006	0.0123	-0.48	0.633
GENDER	0.005	0.0126	0.43	0.670
DIALECT1	0.014	0.0172	0.82	0.418
DIALECT2	0.001	0.0129	0.05	0.959
GENDER:DIALECT1	-0.024	0.0341	-0.71	0.478
GENDER:DIALECT2	0.021	0.0258	0.82	0.414
YOB:GENDER	0.039	0.0303	1.28	0.204
YOB:DIALECT1	0.058	0.0366	1.59	0.120
YOB:DIALECT2	-0.033	0.0271	-1.21	0.231
PHONEME1:YOB	0.009	0.0140	0.66	0.514
PHONEME2:YOB	-0.006	0.0159	-0.38	0.710
PHONEME1:GENDER	0.037	0.0141	2.60	0.012
PHONEME2:GENDER	-0.017	0.0141	-1.17	0.245
PHONEME1:DIALECT1	-0.004	0.0179	-0.23	0.817
PHONEME1:DIALECT2	0.004	0.0143	0.31	0.761
PHONEME2:DIALECT1	0.010	0.0183	0.55	0.588
PHONEME2:DIALECT2	-0.009	0.0144	-0.61	0.543
Random effects				
	Variance		σ	
SPKR (INTERCEPT)	0.0015		0.039	
SPKR (PHONEME1)	0.0011		0.033	
SPKR (PHONEME2)	0.0014		0.038	

Table 2: Abridged model table.

Table 2 lists the key model estimates.[9] Only two predictors are found to be significant:[10] the difference between /ø/ and /œ/ ($p = 0.009$), and the interaction between that predictor and GENDER ($p = 0.012$). The combined effects of these terms are visible in Figure 4. Together, they mean that while there is an overall significant /ø/–/œ/ difference, this is due entirely to the men (0.62 ± 0.02 vs. 0.65 ± 0.02), whereas for women there is no clear difference (0.63 ± 0.03 vs 0.63 ± 0.02). The relationship between /œ̃/ and either of its oral counterparts is unclear both overall and for each gender group due to the greater variability.

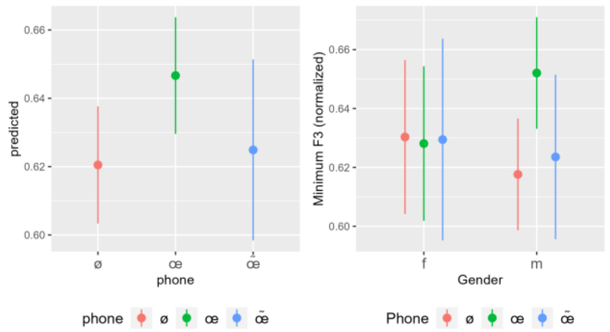

Figure 4: Left: Predicted minimum F3 by phoneme. Right: Interaction between PHONEME and GENDER.

Finally, the overall effect of YOB, while trending downward, is not significant. Figure 5 shows the predicted effects of the interaction between YOB on the one hand and GENDER (left) and DIALECT (right) on the other. A hint of a GENDER difference is visible here: while there is no clear effect for men (the very slight positive slope may be due to a few high-leverage speakers), a decrease in F3 over time can indeed be seen for women. This effect, however, is not significantly different from 0, due to the sparsity of speakers (especially female ones) towards the right end of the YOB range. Similarly, for DIALECT, a similarly sized negative effect is predicted for both the West and East dialects over time, although neither is significant. (Although a positive effect is predicted for the Central dialect, it is very uncertain: it is likely an artefact of the small sample size of this group.) This is likely due to the large amount of inter-speaker variability: the standard deviations of the by-speaker random intercepts (σ = 0.04), as well as those of the by-speaker random slopes for the PHONEME effects (0.03 for /ø/–/œ/ and 0.04 for oral–nasal), are larger than most of the predicted fixed effect coefficients. There is thus a high level of variability that is not conditioned by the social predictors examined here.

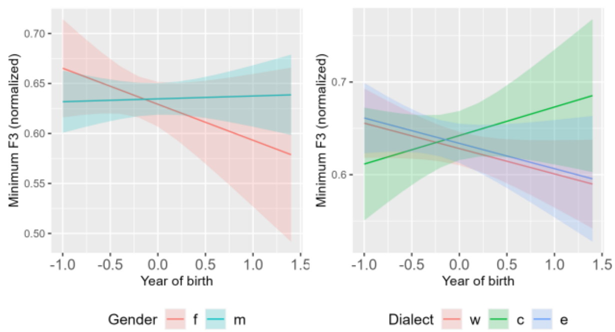

Figure 5: Predicted effects of GENDER (left) and DIALECT (right) by YOB

4. DISCUSSION

The present study confirms one of the major findings of Mielke's studies and extends it to all regions of Quebec: the existence of a large range of inter-speaker variability with respect to rhotacization of the front mid rounded vowels. The evidence for change in progress, however, is not as strong here as it was there, since the effect of YOB is not significant—not even in any of the subgroups. This is perhaps not surprising, as the number of speakers in the corpus used here thins out just as Mielke finds rhotacization begins in earnest in Gatineau-Ottawa. It is possible, then, that rhotacization—despite existing elsewhere, at least in a minority of speakers—only began to progress at an appreciable rate anywhere in Quebec with the cohort of speakers born after 1965. (In fact, it has been argued that a high degree of inter-speaker variability is a prerequisite for sound change [23], which could mean this study has correctly gauged the state of affairs just before the rhotacization took off.) However, it is also possible that the negative effects of YOB observed in 2 of the 3 dialects could become clearer with more data. In any case, the present study unfortunately does not bring clarity to the question of where rhotacization originated, nor to whether its spread is consistent with change from below.

There is also here a hint of a gender effect in two places. One of these is the negative YOB effect found in women but not in men, *contra* Mielke's findings. To be sure, this difference is not significant. Even if it were to be borne out, however, despite contradicting one of the initial hypotheses, it does not necessarily pose a problem for the account of change from below: women in particular have been found to lead changes from below by 1-2 generations in several different studies (cf. [24]). The results here (however tentative) therefore do not contradict Mielke's findings, although more data must be brought to bear on this question.

A further expansion of the corpus with 20-30 additional younger speakers (especially women and speakers from the Central dialect) is currently in progress, and may be sufficient to answer the research questions in much more depth than was possible here. A further goal for future work is to address the issue of the source of rhotacization as a change, which was not taken up here. Finally, there may be additional signal to be found in modelling formant trajectories: this is worthy of investigation in any follow-up study.

Acknowledgments This work was funded by FRQSC B1Z award #321056, as well as Morgan Sonderegger's SSHRC grant #435-2017-0925. Many thanks to MS for his guidance and insight, and to Irene Smith for her comments and support.

5. REFERENCES

[1] Ladefoged, P., Maddieson, I. 1996. *The sounds of the world's languages.* In: *Phonological Theory*. Oxford, UK: Blackwell.
[2] Delattre, P., Freeman, C. 1968. A dialect study of American r's by x-ray motion picture. *Linguistics* 44, 29–68.
[3] Dumas, D. 1972. Le français populaire de Montréal : description phonologique. *Université de Montréal* M.A. thesis.
[4] Mielke, J. 2013. Ultrasound and corpus study of a change from below: Vowel rhoticity in Canadian French. *University of Pennsylvania Working Papers in Linguistics* 19(2), 141–149.
[5] Mielke, J. 2015. An ultrasound study of Canadian French rhotic vowels with polar smoothing spline comparisons. *Journal of the Acoustical Society of America* 137(5), 2858–2869.
[6] Labov, W. 1994. *Principles of Linguistic Change,* vol. 1: Internal Factors. Chichester, UK: Blackwell.
[7] Stevens, K. 1998. *Acoustic phonetics*. Cambridge, MA, USA: MIT Press.
[8] Tousignant, C. 1987. Les variantes du /R/ montréalais: Contextes phonologiques favorisant leur apparition. *Revue québécoise de linguistique théorique et appliquée* 6(3), 73–113.
[9] Milne, P. 2014. The variable pronunciations of word-final consonant clusters in a force aligned corpus of spoken French. *University of Ottawa* Ph. D. Thesis.
[10] McAuliffe, M., Socolof, M., Mihuc, S., Wagner, M., Sonderegger, M. 2017. Montreal Forced Aligner: trainable text-speech alignment using Kaldi. *Proc. 18th Interspeech*, 498–502.
[11] Assemblée nationale du Québec. Dictionnaire des parlementaires du Québec de 1764 à nos jours. *https://www.assnat.qc.ca/fr/membres/notices/index.html*
[12] Saint-Amant Lamy, H. 2016. Kilomètre zéro : /R/ et le statut épicentral de Montréal dans le domaine linguistique laurentien. *Cahiers internationaux de sociolinguistique* 10, 101–134.
[13] McAuliffe, M., Stengel-Eskin, E., Socolof, M., Coles, A., Mihuc, S., Goodale, M., Sonderegger, M. 2019. PolyglotDB [Computer program]. Version 0.1.0, https://github.com/MontrealCorpusTools/PolyglotDB.
[14] Mielke, J., Thomas, E.R., Fruehwald, J. McAuliffe, M., Sonderegger, M., Stuart-Smith, J., Dodsworth, R. 2019. Age vectors vs. axes of intraspeaker variation in vowel formants measured automatically from several English speech corpora. *Proc. 19th ICPhS* Melbourne, 1258–1262.
[15] Boersma, P., Weenink, D. 2023. Praat: doing phonetics by computer [Computer program]. Version 6.3.09, http://www.praat.org/.
[16] Neary, T.M. 1978. *Phonetic feature systems for vowels*. Bloomington: Indiana University Linguistics Club.
[17] Voeten, C.C., Heeringa, W., Van de Velde, H. 2022. Normalization of nonlinearly time-dynamic vowels. *Journal of the Acoustical Society of America* 152(5), 2692–2710.
[18] Kuznetsova, A., Brockhoff, P.B., Christensen, R.H.B. 2017. lmerTest Package: Tests in Linear Mixed Effects Models. *Journal of Statistical Software* 82(13), 1–26.
[19] R Core Team. 2023. R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna. https://www.R-project.org/.
[20] Dumas, D. 1974. Durée vocalique et diphthongaison en français québécois. In: *Cahiers de linguistique de l'Université du Québec* 4. 13–55.
[21] Lüdecke, D., Aust, F., Crawley, S., Ben-Shachar, M.S. 2022. ggeffects: Create tidy data frames of marginal effects for 'ggplot' from model outputs [R package]. Version 1.1.4, https://strengejacke.github.io/ggeffects/.
[22] Wickham, H. 2016. ggplot2: Elegant graphics for data analysis. Springer-Verlag, New York. https://ggplot2.tidyverse.org.
[23] Baker, A., Archangeli, D., Mielke, J. 2011. Variability in American English s-retraction suggests a solution to the actuation problem. *Language Variation and Change* 23, 347–374.
[24] D'Arcy, A. 2015. Variation, transmission, incrementation. In: *The Oxford Handbook of Historical Phonology*, Honeybone, P., Salmons, J. (eds.), 583–602.

[1] Cf. [3] : « […] l'usage de la rétroflexion prend en montréalais une valeur stylistique propre, celle de l'affirmation de soi par référence à l'anglais, qui est sans doute perçu comme 'argument d'autorité' ». (p. 105)
[2] In particular, the lowered F3 which results from creating a constriction at a displacement antinode [7].
[3] Alternatively, at least for /ø/ and /œ/, phonologization of the coarticulatory effects of a following rhotic consonant (which also weakens and deletes), a context where these vowels occur before disproportionately often, may be at cause (cf. [8]).
[4] Milne reports there being 62 speakers, but in the course of this study it was discovered that one was originally double-counted.
[5] More specifically, this Central dialect includes speakers born in the Mauricie, Centre-du-Québec and Estrie administrative regions.
[6] The prototypes are composed of by-phone (but across-speaker) mean formant frequencies, bandwidths and amplitudes, as well as the covariance matrix of these measures. The algorithm then attempts to minimize the Mahalonbis distance between the measured formants and the prototype distributions.
[7] This method was chosen both for its strong performance in [17]'s comparative testing and for comparability purposes with [4], since it is also used there.
[8] All empirical plots are made in R with the *ggplot2* package [21], and all model prediction plots are made with the *ggeffects* package [22].
[9] The full model is available at this project's OSF page, available online at https://osf.io/p3gr7/.
[10] While this might be thought to suggest the model is overparametrized, subsets of this model had essentially the same results with respect to significance.