Published in final edited form as:

J Exp Psychol Learn Mem Cogn. 2012 March ; 38(2): 502–512. doi:10.1037/a0025257.

The case for sub-phonemic attenuation in inner speech: Comment on Corley, Brocklehurst, and Moat (2011)

Gary M. Oppenheim

Beckman Institute / Psychology Department University of Illinois at Urbana-Champaign, Urbana, Illinois

Abstract

Corley, Brocklehurst, and Moat (2011) recently demonstrated a phonemic similarity effect for phonological errors in inner speech, claiming that it contradicted Oppenheim and Dell's (2008) characterization of inner speech as lacking sub-phonemic detail (*e.g.* features). However, finding *an effect* in both inner and overt speech is not the same as finding *equal effects* in inner and overt speech. In this response, I demonstrate that Corley *et al.*'s data is entirely consistent with the notion that inner speech lacks sub-phonemic detail, and that each of their experiments exhibits a similarity-by-articulation interaction of about the same size that Oppenheim and Dell (2008; 2010) have reported in their work. I further show that the major discrepancy between the labs' data lies primarily in the magnitude of the main effect of phonemic similarity and the overall efficiency of error elicitation, and demonstrate that greater similarity effects are associated with lower error rates. This leads to the conclusion that successful speech error research requires finding a sweet spot between "too much randomness" and "too little data."

Introduction

Inner speech is a form of imagery that supports many cognitive activities, including reading, planning (*e.g.* Baddeley, Thomson, & Buchanan, 1975), and possibly overt speech production monitoring (Levelt, 1983). Its generation is typically thought to involve a subset of the processing required for speaking aloud, with dispute over precisely how far that parallel extends. According to one recent claim, from Oppenheim and Dell (2008), inner speech corresponds to an abstract phonological processing level (*e.g.* Dell, 1986; Wheeldon & Levelt, 1995) with less robust (*i.e.* weaker or inconsistent) access to sub-phonemic information (*e.g.* featural, phonetic, motoric). It parallels overt production to the point of retrieving and sequencing abstract phonemes, with processing attenuated thereafter. Major empirical support for this *sub-phonemic attenuation hypothesis* (SAH) comes from comparing inner and overt 'slips of the tongue' (*e.g.* REEF → /lif/). Overt slips tend to involve similarly articulated phonemes (the phonemic similarity effect, *e.g.* Nootboom, 1969). For instance, an /r/ to /l/ slip (voiced alveolars, differing in manner of articulation) is more likely than an /r/ to /b/ (both voiced, but differing in place and manner). This reliable overt speech effect is often attributed to the influence of sub-phonemic ('featural') details during speech planning (*e.g.* Dell, 1986), so its size in inner speech, compared to overt, should reflect the relative contribution of sub-phonemic information. If inner speech tends to involve sub-phonemic details to the same extent as overt speech, phonemic similarity should be equally important in determining error patterns, yielding equally strong similarity effects.

But if sub-phonemic information is less important to inner speech (the SAH claim), its similarity effect should be weaker.

A weaker phonemic similarity effect was precisely what Oppenheim and Dell (2008; O&D) found when comparing tongue-twister-elicited errors in inner speech to those in overt, providing initial support for the SAH. Comparable tendencies for both inner and overt slips to create words (lexical bias) suggested robust engagement at the phoneme level for inner speech, in contrast to the differences in the similarity effects. O&D (2010) replicated and extended the work, demonstrating that silently mouthing a tongue-twister elicited an overt-like similarity effect in inner speech, while the tendency in unarticulated inner speech was again significantly diminished. Therefore, the attenuated similarity effect in unarticulated inner speech could not be due to difficulty ‘hearing’ inner slips. Converging support for the SAH came from observations that the influence of articulatory features in phonological working memory hinges on a task’s use of overt articulation (Schweppe, Grice, & Rummer, 2011). Thus, the SAH and associated empirical findings have proven empirically robust and, to judge by recent discussions (*e.g.* Abdel Rahman & Aristei, 2010; Geva, Bennett, Warburton, & Patterson, in press; Harley, 2010; Heuttig & Hartsuiker, 2010; Heuttig, Roomers, & Meyer, in press; Hickok, Houde, & Rong, 2011; Hubbard, 2010; Laganaro & Zimmermann, 2010; Nooteboom & Quene, 2008; Nozari & Dell, 2009; Severens, Janssens, Kühn, Brass, & Hartsuiker, in press; O’Seaghdha, Chen, & Chen, 2010; Stemberger, 2009; Vicente & Martinez-Manrique, in press), theoretically useful.

Corley, Brocklehurst, and Moat (2011; CBM) recently presented evidence that they interpreted as challenging the SAH. In three experiments modeled on O&D’s (2008) task, they replicated the lexical bias findings but additionally observed simple main effects of phonemic similarity in both inner and overt speech, with only “some small signs that there might be numerical trends” (p179) toward a weaker similarity effect in inner speech. Since their similarity-by-overtness interaction was only marginally significant ($p < .09$ in a two-degrees of freedom model comparison), CBM concluded in favor of the null hypothesis that the similarity effects in inner and overt speech were equal:

“Over three experiments, we have shown that overtness does not interact with similarity in predicting the likelihood of an onset substitution” (p169) “Taking data from 18,432 total recitations of four-word tongue twisters by 112 participants, no evidence could be found that any numerical difference in the likelihood of substituting similar phonemes in inner compared to overt speech was reliable.” (p171), “Phonemic similarity consistently influenced the likelihood of reporting errors to a similar extent in inner speech as it did in overt speech.” (p171).

As CBM noted, “Perhaps most surprisingly, when we replicated our experiments using Oppenheim and Dell’s [2010] materials, the results were consistent with our two earlier experiments.” Figure 1 compares results from O&D’s two studies with CBM’s experiment using the same stimuli (their Experiment 3). O&D’s experiments show significant crossover interactions, where the phonemic similarity effect is stronger in overtly articulated speech, but CBM’s crossover interaction was not statistically significant. The figure clearly indicates a comparable interaction, though, so it is important to closely examine it and CBM’s other two studies with different, but comparable stimuli.

In CBM’s Experiment 3, the similarity effect in inner speech consisted of a 13 error difference between the similar and dissimilar conditions, yielding an odds-ratio of 2.1:1. This difference is asserted to be equal to that in overt speech (29 errors difference, an odds-ratio of 3.1:1): “participants were once again much more likely (here, by a factor of 2.7) to substitute similar than dissimilar phonemes, regardless of whether the speech was overt or not,” (p169). More generally, CBM concluded that, “far from being underspecified, our

'inner voice' sounds much like our overt speech, and is produced in much the same way, whether overtly articulated or not," (p172). It is both the specific conclusion that the similarity effects are the same in inner and overt slips, and the general one that inner speech is not underspecified, that I dispute here.

Like the similarity effect, the size of the similarity-by-articulation interaction can be quantified as an odds-ratio. In CBM's Experiment 3, presented in the figure, the interaction effect size is 1.45:1. It was even greater for their other two experiments: 1.70:1 and 1.51:1. This means that CBM's overt-speech similarity effects were 45%, 51%, and 71% larger than those in the comparable inner speech conditions. Do these findings justify asserting the null effect? I will argue that they do not.

Through a statistical and conceptual reconsideration of the error data from five experiments (*i.e.* CBM's three experiments, and O&D 2008 and 2010), this analysis will first demonstrate that the data clearly and consistently support the SAH prediction that phonemic similarity effects are weaker in unarticulated inner speech. It will also show that CBM's central finding—that their data showed no evidence of the expected similarity-by-overtness interaction—is false, and hence their theoretical interpretation—that inner speech is fully specified for phonological features¹—is unsupported, and even a more nuanced version of CBM's claim—that they found a significantly weaker interaction than O&D—lacks statistical support. Second, since CBM could not explain why they found a simple main effect of phonemic similarity in inner speech whereas O&D did not ("We are not able to fully account for Oppenheim and Dell's findings," p171), this response will offer a plausible resolution for the discrepant findings—one that draws on a mathematical consequence of the over-determined nature of speech errors to demonstrate a long-recognized, but underappreciated, aspect of speech error distributions.

Part 1: The phonemic similarity-by-articulation interaction

The first analysis builds on CBM's recognition that power issues for individual speech error experiments may be alleviated by combining data across similar experiments. Laboratory studies of speech errors tend to elicit few errors of the desired type, so even seemingly large studies may lack the power to statistically detect a true effect. For instance, CBM's "18,432 total recitations" only provided 340 target errors, spread over 112 subjects, 128 items and four conditions in three experiments. It is error counts that build the power of a binomial analysis – not the number of trials, subjects, items, or experiments. (Mathematically², increasing the number of 'observations' (trials), while holding constant the number of 'successes' (errors), only *reduces* one's chance of statistically detecting an effect.) With few errors to examine, even reasonably strong, consistent differences can be difficult to detect. Though each of their experiments replicated O&D's similarity-by-articulation interaction (*e.g.* in the sense of Killeen's, 2005, *p_{rep}*), even combining data across experiments only allowed CBM to detect this reasonably large effect with marginal significance.

Of course, the real question is whether the interaction holds overall—not whether it achieves significance in some arbitrary subset. I address this question by extending CBM's combined analysis to include O&D's experiments, recoding O&D's data to match CBM's methods. These five experiments used comparable materials, tasks, and data-coding to evaluate

¹In CBM's words, "inner speech is fully phonologically represented" (p171). CBM's terminology reflects a linguistic tradition where phonemes are considered, if anything, bundles of features (*e.g.* Goldstein *et al.*, 2008), whereas O&D's reflects a psycholinguistic approach where phonemes occupy a level above features (*e.g.* Dell, 1986; Wheeldon & Levelt, 1995).

²For instance, according to formulas such as $SE_p = \sqrt{p(1-p)/n}$. Holding the number of successes constant (where $p = \text{successes}/n$), increasing n only decreases p/SE_p , thus decreasing the power to detect an effect involving p . And since $OR_{p_1p_2} = (p_1/(1-p_1))/(p_2/(1-p_2))$, increasing n decreases $OR_{p_1p_2}$.

similarity effects and their interaction with articulation. The data is first compiled in a single table, providing the first parametric estimates of the interaction and its inter-experiment variability. (Neither O&D's nonparametric statistics nor CBM's stepwise regression allowed such estimates.) Then per-trial data is used to examine whether similarity effects are characteristically weaker in unarticulated inner speech, as O&D proposed, and to evaluate evidence for claims that this interaction varies across labs.

Methods

Analyses considered O&D's two experiments (2008; 2010) and CBM's three. Each used four-word tongue-twisters with ABBA onset phoneme patterns, constructed in matched sets varying featural similarity of the A and B onset consonants (*e.g.* LEAN REED REEF LEECH vs. BEAN REED REEF BEECH; see O&D, 2008, for more details). Each participant memorized and then attempted to recite a single variant of each tongue-twister four times, in a single articulation condition, reporting their errors aloud.

I obtained the dataset from CBM's three-experiment meta-analysis (for which I am grateful) and recoded O&D's (2008; 2010) data with their methods. The main change was that, since CBM counted all errors in a given recitation (up to two target errors per attempt), whereas O&D considered only the first error, I adopted their scheme. Thus the new error counts differ slightly from those in our previous papers. Following CBM's treatment, each trial was binomially coded as the number of target errors ('successes') versus error-free productions ('failures'). Target errors were reported slips from a B onset to an A onset (*e.g.* REEF → /lif/) without other deviations (*e.g.* REEF → /lid/).³

Following CBM, analyses used mixed-effects logistic regression via Bates and Maechler's (2010) lme4 package for R (R Development Core Team, 2010), including crossed random effects for Subject, nested within experiment, and Item, non-nested to reflect some items' reuse across experiments.⁴ Cross-lab analyses treated Lab as a fixed effect (base level is O&D); within-lab analyses used Experiment instead.

Analyses considered three other fixed main effects: Phonemic similarity, Articulation, and Audition. As in the source experiments, Phonemic similarity (two levels) classified trials as Similar or Dissimilar based on the number of feature contrasts between the trial's two onset phonemes (1 vs. >1). Articulation (two levels) coded whether a trial involved overt movements, thus comparing unarticulated inner speech (without overt articulation) to mouthed, noise-masked, and normal overt speech (with overt articulation). Audition (two levels) coded whether a trial involved subject-audible overt speech (the so-called external loop), thus comparing inner speech, silently mouthed speech, and noise-masked overt speech (without external audition) to unmasked overt speech (with external audition). Treating noise-masked and silently mouthed speech as equivalent in this regard draws some justification from Postma and Noordanus (1996), who demonstrated continuity between these conditions across an exhaustive range of self-reported error types.

Fitting began with a full theoretically-justified model, including centered binary predictors for Phonemic similarity, Articulation, Audition, and Lab, plus two-way interactions between Phonemic similarity and the other predictors. To evaluate CBM's claim that their similarity-by-articulation interaction differed from Oppenheim and Dell's, I also included a Phonemic similarity-by-Articulation-by-Lab interaction (and ancillary Articulation-by-Lab interaction). Nonsignificant components were incrementally removed to reduce collinearity,

³Following O&D (2008; 2010), I use all B→A target errors to increase the statistical power, but all claimed significant effects in the full analyses still hold if considering only Word 3.

⁴Coding Item as a nested random effect yields equivalent statistical outcomes.

with the relevant portion of the reduced model then refitted to each lab's individual data. Model coefficients (β) denote log-odds. P-values in the text reflect 1-df likelihood-ratio tests, evaluable under non-directional ($\alpha=.05$ for $\beta=0$) or directional ($\alpha=.10$ for $\beta>0$ XOR $\beta<0$) hypotheses as appropriate; tables provide Wald z-scores.

Results and Discussion

The five experiments provide 912 self-reported target errors (487 in overtly articulated speech, 435 in unarticulated inner speech), summarized in Table 1, to evaluate differences in phonemic similarity effects between overtly articulated and unarticulated (inner) speech. The table suggests three major patterns. First, similarity effects are clearly stronger in articulated speech than unarticulated inner speech, and this holds for each individual experiment, whether considering counts or odds-ratios. Second, although CBM's experiments show stronger similarity effects overall, their similarity-by-articulation interaction is remarkably stable and approximately equal in odds-size to the estimate from O&D's work. In fact the "small signs that there might be numerical trends in this direction" are not small at all: in overtly articulated speech, their subjects reported 130 similar errors to 47 dissimilar, an 83-error difference. In inner speech, they reported 105 similar errors to 58 dissimilar, a 47-error difference. Finally, although the two labs contributed comparable numbers of trials (CBM: 18k vs. O&D: 16k), CBM's elicited far fewer target errors.

Statistical analyses support all of these trends (Table 2). Consistent with CBM's masking analyses, neither Audition⁵ ($\beta=-0.18$, $\chi^2(1)=2.66$, $p=.10$) nor its interaction with Phonemic similarity ($\beta=0.04$, $\chi^2(1)=0.03$, $p=.87$) reached significance. This reinforces O&D's (2010) conclusion that smaller similarity effects in unarticulated inner speech do not merely reflect the absence of an external loop (cf. Levelt, 1989). Neither the Phonemic similarity-by-Articulation-by-Lab interaction nor its ancillary Articulation-by-Lab interaction, approached significance (each $|\beta|<0.1$, $\chi^2(1)<0.01$, $p>.6$), indicating a lack of support for claims that CBM found a smaller similarity-by-articulation interaction. Incrementally removing these nonsignificant predictors reduced the model to the five fixed effects discussed below (Full-Reduced comparison: $\chi^2(4)=2.94 < \text{crit}(\chi^2(1))_{\alpha=.05}$).

The tendency for target errors to involve similarly articulated phonemes ($\beta=0.57$, $\chi^2(1)=52.09$, $p<.0001$) held for both overtly articulated speech ($\beta=0.80$, $\chi^2(1)=56.18$, $p<.0001$) and unarticulated inner speech ($\beta=0.33$, $\chi^2(1)=6.13$, $p<.02$) separately. This pattern reinforces CBM's assertion that inner speech can show similarity effects (contra O&D's 2008 data), implying inner speech can incorporate sub-phonemic information under some conditions (e.g. as O&D, 2010, suggested).

But finding a *similarity effect* in inner speech is not the same as finding *equal similarity effects* in inner and overt speech. In fact, the effect in overtly articulated speech is much greater, about 1.6 times the size of that in unarticulated speech (similarity-by-articulation interaction: $\beta=0.48$, $\chi^2(1)=12.04$, $p<.0006$). As Table 1 suggests, this ratio holds quite well for each lab individually (O&D: $\beta=0.52$, $\chi^2(1)=9.02$, $p<.003$; CBM: $\beta=0.45$, $\chi^2(1)=3.56$, $p<.06$), meaning that their data actually agree on the direction and size of the similarity-by-articulation interaction. Note also that the interaction using CBM's data only would be significant under the directional SAH prediction of a *stronger* similarity effect in articulated speech. Thus, the combined data clearly support the Sub-phonemic Attenuation Hypothesis, and each lab and experiment individually contributes to this support.

⁵The negative Audition coefficient suggests a trend opposite what is typically predicted (Levelt, 1989) and demonstrated (Postma & Noordanus, 1996), so a directional test of the coefficient is not considered. Retaining the predictor, however, would not noticeably change any claimed results.

If the difference in the labs' conclusions does not reflect differences in the size or direction of the similarity-by-articulation interaction, what does it reflect? Here, the analysis statistically confirms two points suggested by visual inspection of Figure 1. First, CBM's experiments elicited target errors at less than half the rate of O&D's (main effect of Lab: $\beta = -0.89$, $\chi^2(1) = 38.47$, $p < .0001$). Second, CBM's data showed stronger similarity effects overall, not just for inner speech (Similarity-by-Lab interaction: $\beta = 0.53$, $\chi^2(1) = 13.02$, $p < .0003$). Together these points appear to resolve the question of why two labs would report contrasting results for the similarity-by-articulation interaction that is the crucial test of the SAH: although both datasets showed a similar-sized interaction, low error rates could make it harder to statistically detect, and the presence of simple main effects of similarity in both inner and overt speech would make the interaction easier to overlook.

To recap, the patterns from each experiment (Table 1) and the statistical results from each lab (Table 2) individually support the prediction of a weaker similarity effect in unarticulated inner speech, indicating that the overt similarity effect is about 60% greater. CBM's failure to statistically confirm the similarity-by-articulation interaction does not reflect a reliable or even substantial discrepancy between the two labs' data on that point, as implied by claims that they failed to detect the interaction despite great effort. And their *simple main effect* of similarity in inner speech—erroneously claimed as evidence against the SAH—is better characterized as a difference in the magnitude of the *main effect*.

Part 2: The simple main effect of phonemic similarity in unarticulated inner speech

I have shown that the discrepancy in the presence of a simple main effect of phonemic similarity in unarticulated inner speech, which CBM could not explain, is merely symptomatic of a difference in the size of the main effect. So this section will propose a theory-derived explanation that both stands on its own and could account for some variation in similarity effects across experiments. The story, in a nutshell, is that the over-determined nature of slips of the tongue has the consequence that their specific causes (*e.g.* dimensions of similarity between interacting representations) should be more evident (as odds-ratios) when their more general causes (*e.g.* stress, time pressure, novelty, priming) contribute less. But with less support from general causes, the resultantly rare events may provide more volatile estimates with less statistical power.

Although speakers may complain that they err too often, researchers complain that they err too rarely. Carefully controlled stimuli often fail to compel unimpaired speakers to produce the kinds of errors that researchers want in the quantities that they need. So researchers rely on seemingly irrelevant aspects of an experiment to help elicit errors, such as time pressure, priming, and otherwise difficult sequences. The classic SLIP procedure (Baars, Motley & MacKay, 1975), for instance, elicits spoonerisms in part by priming a particular onset phoneme sequence and then unexpectedly reversing it (ABABABBA). Researchers assume that such manipulations make phonological encoding less deterministic—shifting productions away from near-ceiling accuracy and hence increasing statistical power—but the resulting errors nonetheless reflect the structure and processes of successful speech production. Their assumption reflects what Freud (1901/1958) described as the “over-determined” nature of speech errors: many factors interact to determine if and how production may miss its mark. Thus a slip from ‘barn door’ to /darn bɔr/ may simultaneously reflect priming of the /d... b.../ onset pattern, featural overlap between /d/ and /b/ onsets, lexicality of the resulting utterance, pressure to respond quickly, and latent feelings of bucolic ennui. An underappreciated property of this over-determination is that, if one factor better supports a slip, the remaining factors become less crucial. For instance, if priming increases the likelihood of substitutions in general, phoneme substitutions will require less support from

shared features (*i.e.* phonemic similarity). Consequently, an error effect like that of phonemic similarity should tend to be larger in log-odds terms when fewer factors promote slips and hence overall error probabilities are lower.

Error researchers have long recognized that when they are rarer, errors are more likely to exhibit more of the properties that promote them. For example, exchange errors (*e.g.* BARN DOOR → /dɔrn bɔr/) are relatively rare compared to other substitutions, but it is in exchanges, that similarity and familiarity effects are the strongest (*e.g.* Dell, 1986; Garrett, 1980). Similarly, phoneme substitution errors usually have obvious sources in the surrounding context, but noncontextual slips show proportionally stronger phonemic similarity effects (Stemberger, 1992), suggesting that similarity is more crucial for slips with less contextual support. At the other extreme, aphasic individuals, who make frequent phonological errors, may show attenuated phonemic similarity effects (*e.g.* Goldrick & Rapp, 2007; Laganaro & Zimmerman, 2010), demonstrating that similarity is less crucial for slips driven by other (*e.g.* lexical) factors. Consider the following analogy: Students at prestigious institutions often distinguish themselves by possessing certain factors (*e.g.* intelligence, ambition, industriousness). A student from a rich family, where attending such an institution is more common, may need only one of these factors. But a student from a poor family, who nonetheless makes it in, would be more likely to possess all of them. Making an error is like making it into college. If making it is rare, then those that do will more consistently exhibit the relevant factors.

For a mathematically transparent illustration of this relationship, consider a common mathematical approximation of a stochastic selection process: the Luce-Shepard choice rule (Equation 1, adapted from Luce, 1963; Shepard, 1957). Models often use this equation to translate continuous activations into expected probabilities for discrete outcomes (*e.g.* Dell, Burger, & Svec, 1997; Gordon & Dell, 2003; Kruschke, 1992; Love, Medin, & Gureckis, 2004; McClelland & Elman, 1986; McClelland & Rumelhart, 1981; Nosofsky, 1986).

$$p(i) = \frac{e^{\mu a_i}}{\sum e^{\mu a}} \quad \text{Equation 1}$$

Here a_i is the activation of outcome i , and $\mu = 0$ makes selection more or less deterministic. To apply the rule to phoneme selection, imagine a target phoneme, /r/, and two competitors, /l/ (similar) and /b/ (dissimilar), with activations, $a_{r,l,b} = \{1, 0.5, 0.25\}$, as might arise from a similarity-sensitive retrieval process. From Equation 1, the probability of erroneously choosing /l/, $p(l)$, is a function of the activation of /l/, $a_{l/l}$, compared to the summed activation of all potential onsets. With the exponential function, $e^{\mu a}$, μ scales selection to a total error rate: larger μ yield more “correct” selections, while smaller μ simulate more random selection. Using the equation to derive probabilities of selecting /l/ or /b/, given the target /r/, we estimate the similarity effect as the odds-ratio for /l/ versus /b/ outcomes. Varying μ produces greater similarity effects when errors are less frequent (Figure 2a).

Dell's (1986) model also instantiates this principle: reducing activation noise (analogous to μ in Equation 1) reduces error rates, yielding stronger similarity effects (Figure 2b). But these simulations also illustrate a down-side to very low error rates: resultant effects are small as counts (Figure 2c) and variable as odds-ratios (Figure 2b). Noise similarly modulates phonological effects in Dell, Schwartz, Martin, Saffran, and Gagnon's (1997) aphasia model, and semantic effects in Oppenheim, Dell, and Schwartz's (2010) lexical retrieval model. In fact, every model of production errors that I know of has the property that lower error rates produce stronger odds-ratio effects. This connection is ubiquitous

because it is a direct consequence of the over-determined nature of speech errors. Factors that make production *more* deterministic (accurate) produce models where errors are infrequent (thus yielding less precise estimates) but error patterns are dominated by ‘good’ errors (and hence greater odds-ratio effects). Factors that make production *less* deterministic lead to greater numerical error effects (at least until randomness dominates the error profiles) and hence more power to statistically detect them. In the current experiments, these influences manifest as larger phonemic similarity effects when slips are rarer, but more power to detect the same-sized effect when slips are more frequent. I claim that both articulated and unarticulated speech show this tendency, but the weaker similarity effect in unarticulated speech is easily overshadowed when production becomes less deterministic. The following statistical analyses demonstrate that the mediating effect of generalized error rate offers a reasonable explanation for the major difference between the labs findings, providing possibly the first empirical demonstration of this long-recognized property of speech errors.

Methods

This analysis used the same dataset and basic methods as the previous. Subjects remain nested within Experiment (and Items remain non-nested) but I replaced the nominal Lab and Experiment predictors with a continuous measure of general error rate—analogue to $-\mu$ in Equation 1. Nominal Lab and Experiment predictors are omitted because, since the analysis offers a theoretical explanation for the discrepancy that these identity-based factors describe, any predictor that successfully explains that discrepancy would necessarily be highly collinear with identity-based descriptions of it. For instance, *General error rate* hypothesizes a particular ordering and spacing of data across experiments; if it perfectly explained the variation across experiments, it would provide exactly the same estimates as the nominal Experiment predictor. Because the previous section addressed the other fixed effects, I restrict discussion to the General error rate predictor and its interaction with phonemic similarity.

General error rate was calculated here as the proportion of Word 2 and 3 attempts, per-experiment, that elicited neither target errors nor correct productions (*i.e.* $1 - p(\text{correct}) - p(\text{target error})$). Given the available data, this method provides an estimate that is consistent across labs, reasonably independent of the target error distributions within an experiment,⁶ and concordant with ordinal rankings of the published total error rates (Kendall's τ for concordance between General error rate and published per-experiment error rates: Overtly articulated speech: $\tau=1$, $p<.03$; Unarticulated inner speech: $\tau=1$, $p<.03$). However, the results do not hinge on this particular estimation method,⁷ and other methods appear to work equally well here (*e.g.* more focused estimates of general onset slip rates, as in the simulations, may prove more robust for cross-paradigm comparisons). Finally, although Figure 2 suggests similarity effects should relate to target error rate via a power function, I treat General error rate as a linear predictor to avoid overfitting the limited dataset and make the resulting coefficients easier to interpret.

⁶The mathematical relation between target error rates (calculated as $[target]/[target] + [correct]$) and non-target errors is inconsequential here—calculating target rates as $[target]/[total]$ minimizes the relation while producing essentially the same coefficients, standard errors, p-values, figures, etc. Note also that, mathematically, greater non-target rates would increase odds-ratio effects (per Footnote 2), contra the predicted interaction.

⁷The estimate should, though, collapse across conditions (*e.g.* similarity). Estimating per-experiment, not per-subject or -item, avoids sparse-matrix artifacts and trivial auto-correlations (*e.g.* with infrequent target errors, anything predicting more target errors may predict stronger effects merely by reducing sampling error and floor effects).

Results and Discussion

With 2,206 non-target errors to estimate *General error rates* (described in Table 1), more error-prone production generally produced more target errors but weaker phonemic similarity effects (Table 3). It is perhaps unsurprising that people reported more target errors when reporting more non-target errors⁸ (main effect of General error rate: $\beta=.13$, $\chi^2(1)=42.43$, $p<.0001$), but this association does support the validity of using non-target errors to index the general error proneness of a study.

Figure 3 shows that larger General error rates in these experiments were also associated with weaker similarity effects (interaction: $\beta=-.07$, $\chi^2(1)=9.86$, $p<.002$). This similarity-by-error rate interaction provides a log-odds slope for the regression lines in the figure, suggesting that a one-percent increase in General error rate (*e.g.* from 5% to 6%) is accompanied by a seven-percent decrease in the log-odds similarity effect. This pattern holds for both articulated ($\beta=-.07$, $\chi^2(1)=4.96$, $p<.03$) and unarticulated speech ($\beta=-.06$, $\chi^2(1)=4.30$, $p<.04$) individually, reinforcing earlier claims that the major quantitative difference between the labs' results lies in the size of the main effect of phonemic similarity, not a simple main effect or interaction with articulation. Visually, the similarity effects in articulated and unarticulated speech form linearly separable sets when paired by experiment, because the theoretically motivated inclusion of a mediating effect of overall error rate gives order to the variation between experiments. This overall pattern is easily simulated by a production model where inner speech involves attenuated sub-phonemic connections (Figure 2b-c). (The model also predicts a stronger similarity-by-error rate interaction for articulated speech—which the data only hint at—but note that in the observed range of error rates it does not necessarily predict a large log-odds difference, particularly when computing General error rate by collapsing across articulation conditions.)

Might this pattern extend to other data? Figure 4 plots similarity effects for the articulated conditions, plus data from four experiments using the SLIP procedure: Nootboom (2005), Nootboom and Quene (2008, two experiments) and O&D (2008, SLIP task).⁹ To compare across procedures with different error profiles, the predictor is now the overall probability of target errors, and following Nootboom and Quene (2008) target errors in the SLIP task are defined as primed-for completed spoonerisms and anticipations.¹⁰ Though methodological differences suggest caution in relating results across paradigms, the SLIP data fits the expected pattern surprisingly well. There are three main points to note. First, the SLIP data spans a wide range of error rates and actually shows the expected similarity-by-error rate pattern on its own. Second, the unprimed conditions in three experiments (plotted separately) elicited few target errors but particularly strong (and variable) similarity effects—precisely as predicted by simulations with Dell's (1986) model (*e.g.* Figure 2b), and more generally consistent with the notion that similarity more strongly constrains errors that lack support from other sources. Finally, where Nootboom and Quene (2008) removed time pressure and warned participants about the task structure—changes intended to maximize editing, that generally eliminated several irrelevant speech error causes—this reduces overall error while eliciting the strongest similarity effect in their test data, following the prediction that particular error causes should tend to be more evident (in odds-ratios) when other error causes contribute less.

⁸Note this does not mean target error rates scale in fixed proportion to non-target error rates.

⁹O&D's (2008) SLIP task used the same stimuli as their tongue-twister task. Data from Nootboom (2005) and Nootboom & Quene (2008) were provided by the authors, who recoded their 2005 similarity distinctions to match their 2008. This is every published experiment I know that tests comparable similarity distinctions, uses CVC or CVCC stimuli, and codes specific primed-for onset substitutions.

¹⁰Errors for SLIP tasks are experimenter-coded. Adding SLIP perseverations would slightly increase target error counts (mostly for Oppenheim & Dell's data) without substantially affecting the plotted effects.

To recap, error rates were used here to estimate generalized tendencies to produce target errors that are unrelated to the manipulations of interest. The influence of general error-proneness does not preclude contributions from more specific error causes, but by indexing the ‘noise’ in the speech production process it provides a powerful, theoretically motivated resolution for a set of seemingly discrepant findings and has predictive value beyond the current data. This point is separate from the power issues discussed earlier, because more accurate production produces fewer errors in the same number of trials, and practically speaking this makes the estimates more vulnerable to sampling error and less able to support statistically detecting the same-sized effect.

The relationship between error rates and error patterns is a consequence of the fact that speech errors are simultaneously determined by multiple causes. It does not compel a particular feedback explanation, and could certainly be described in terms of strategic speech monitoring (e.g. overworked monitors are less effective). But, as a domain-general phenomenon, it seems appropriate to posit a domain-general mechanism, and these patterns naturally arise from the kinds of stochastic selection algorithms that have played a crucial role in models of cognition for over half a century.

Conclusion

The first analysis demonstrated a robust similarity-by-articulation interaction, similar-sized in both labs’ data. Thus, the SAH is supported by both datasets and generally speaking has strong support. It also identified the major discrepancies between the labs’ data: CBM elicited stronger main effects of similarity and fewer errors in general, making it easy to overlook the similarity-by-articulation interaction.

The second analysis explained these discrepancies together as a consequence of the over-determined nature of speech errors. Phonemic similarity effects, like other ‘good’ error patterns, are generally greater when error likelihood (indexing support from other error-causing factors) is low. Both inner and overt speech show this tendency, but the weaker similarity effect in inner speech is more easily obscured as phoneme selection grows less deterministic. These results are clearly predicted by a model where inner speech involves attenuated phoneme-feature connections, and more generally suggest that successful speech error research requires finding a sweet spot between ‘too much randomness’ and ‘too little data’.

Corley *et al.* questioned Oppenheim and Dell's claim that inner speech involves less robust access to sub-phonemic information. The data however, fully support that claim by showing that phonemic similarity effects are consistently weaker in unarticulated inner speech.

Acknowledgments

Supported by NIH DC000191 and HD44458. Thanks to Gary Dell, Sieb Nooteboom, Kay Bock, Simon Fischer-Baum, Scott Fraundorf, Matt Goldrick, Audrey Kittredge, Tuan Lam, Nazbanou Nozari, and Ying Wang for valuable discussions and other contributions. I also thank Paul Brocklehurst and colleagues for inspiring this response and providing data from their 2011 article.

References

- Baars BJ, Motley MT, MacKay DG. Output editing for lexical status in artificially elicited slips of the tongue. *Journal of Verbal Learning and Verbal Behavior*. 1975; 14(4):382–391.
- Baddeley A, Thomson N, Buchanan M. Word length and the structure of short-term memory. *Journal of Verbal Learning and Verbal Behavior*. 1975; 14:575–589.

- Bates, D.; Maechler, M. lme4: Linear mixed-effects models using S4 classes. R package version 0.999375-35. 2010. <http://CRAN.Rproject.org/package=lme4>
- Chen JY, Chen TM. Proximate units in word production: Phonological encoding begins with syllables in Mandarin Chinese but with segments in English. *Cognition*. 2010; 115(2):21.
- Corley M, Brocklehurst PH, Moat HS. Error biases in inner and overt speech: Evidence from tongue twisters. *Journal of experimental psychology: Learning, memory, and cognition*. 2011; 37(1):162–75.
- Dell GS. A spreading-activation theory of retrieval in sentence production. *Psychological review*. 1986; 93(3):283–321. [PubMed: 3749399]
- Dell GS, Burger LK, Svec WR. Language production and serial order: A functional analysis and a model. *Psychological Review*. 1997; 104(1):123–147. [PubMed: 9009882]
- Dell GS, Schwartz MF, Martin N, Saffran EM, Gagnon D. Lexical access in aphasic and nonaphasic speakers. *Psychological Review*. 1997; 104(4):801–838. [PubMed: 9337631]
- Freud, S.; Brill, AA. *The psychopathology of everyday life*. Macmillan; New York: 1958. (Original work published 1901)
- Garrett, MF. Levels of processing in sentence production.. In: Butterworth, B., editor. *Language production*. Academic Press; New York: 1980. p. 177-220.
- Geva S, Bennett S, Warburton E, Patterson K. Discrepancy between inner and overt speech: Implications for post-stroke aphasia and normal language processing. *Aphasiology*. 2011; 25(3): 323–343.
- Goldrick MA, Rapp B. Lexical and post-lexical phonological representations in spoken production. *Cognition*. 2007; 102(2):219–260. [PubMed: 16483561]
- Gordon JK, Dell GS. Learning to divide the labor: An account of deficits in light and heavy verb production. *Cognitive Science*. 2003; 27(1):1–40.
- Harley, TA. *Talking the talk: language, psychology and science*. Psychology Press; New York: 2010.
- Hickok G, Houde J, Rong F. Sensorimotor Integration in Speech Processing: Computational Basis and Neural Organization. *Neuron*. 2011; 69(3):407–422. [PubMed: 21315253]
- Huetting F, Hartsuiker RJ. Listening to yourself is like listening to others: External, but not internal, verbal self-monitoring is based on speech perception. *Language and Cognitive Processes*. 2010; 25(3):347–374.
- Hubbard TL. Auditory imagery: empirical findings. *Psychological bulletin*. 2010; 136(2):302–29. [PubMed: 20192565]
- Huetting F, Rommers J, Meyer AS. Using the visual world paradigm to study language processing: A review and critical evaluation. *Acta psychologica*. in press.
- Killeen PR. An alternative to null-hypothesis significance tests. *Psychological science*. 2005; 16(5): 345–53. [PubMed: 15869691]
- Kruschke JK. ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review*. 1992; 99(1):22–44. [PubMed: 1546117]
- Laganaro M, Zimmermann C. Origin of phoneme substitution and phoneme movement errors in aphasia. *Language and Cognitive Processes*. 2010; 25(1):1–37.
- Levelt WJM. Monitoring and self-repair in speech. *Cognition*. 1983; 14(1):41–104. [PubMed: 6685011]
- Love BC, Medin DL, Gureckis TM. SUSTAIN: a network model of category learning. *Psychological Review*. 2004; 111(2):309–32. [PubMed: 15065912]
- Luce, RD. Detection and recognition.. In: Luce, RD.; Bush, RR.; Galanter, E., editors. *Handbook of Mathematical Psychology*. Vol. I. Wiley; New York: 1963. p. 103-189.
- McClelland JL, Elman JL. The TRACE model of speech perception. *Cognitive Psychology*. 1986; 18:1–86. [PubMed: 3753912]
- McClelland JL, Rumelhart DE. An interactive activation model of context effects in letter perception. Part 1. An account of basic findings. *Psychological review*. 1981; 88:375–407.
- Nooteboom, SG. The tongue slips into patterns.. In: Sciarone, AG.; van Essen, AJ.; van Raad, AA., editors. *Nomen: Leyden studies in linguistics and phonetics*. Mouton; The Hague: 1969. p. 114-132.

- Nooteboom SG. Lexical bias revisited: Detecting, rejecting and repairing speech errors in inner speech. *Speech communication*. 2005; 47(1-2):43–58.
- Nooteboom SG, Quené H. Self-monitoring and feedback: A new attempt to find the main cause of lexical bias in phonological speech errors. *Journal of Memory and Language*. 2008; 58(3):837–861.
- Nosofsky RM. Attention, similarity, and the identification-categorization relationship. *Journal of experimental psychology: General*. 1986; 115(1):39–61. [PubMed: 2937873]
- Nozari N, Dell GS. More on lexical bias: How efficient can a “lexical editor” be? *Journal of Memory and Language*. 2009; 60(2):291–307. [PubMed: 20126302]
- Oppenheim GM, Dell GS. Inner speech slips exhibit lexical bias, but not the phonemic similarity effect. *Cognition*. 2008; 106(1):528–537. [PubMed: 17407776]
- Oppenheim GM, Dell GS. Motor movement matters: the flexible abstractness of inner speech. *Memory & cognition*. 2010; 38(8):1147–60.
- Oppenheim GM, Dell GS, Schwartz MF. The dark side of incremental learning: a model of cumulative semantic interference during lexical access in speech production. *Cognition*. 2010; 114(2):227–252. [PubMed: 19854436]
- Postma A, Noordanus C. Production and detection of speech errors in silent, mouthed, noise-masked, and normal auditory feedback speech. *Language and Speech*. 1996; 39(4):375–392.
- R Development Core Team. R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing; Vienna, Austria: 2010. <http://www.R-project.org>
- Rahman RA, Aristei S. Now you see it ... and now again: Semantic interference reflects lexical competition in speech production with and without articulation. *Psychonomic Bulletin & Review*. 2010; 17(5):657–661. [PubMed: 21037163]
- Schweppe J, Grice M, Rummer R. What models of verbal working memory can learn from phonological theory: Decomposing the phonological similarity effect. *Journal of Memory and Language*. 2011; 64(3):256–269.
- Severens E, Janssens I, Kühn S, Brass M, Hartsuiker RJ. When the brain tames the tongue: Covert editing of inappropriate language. *Psychophysiology*. in press.
- Shepard RN. Stimulus and response generalization: A stochastic model relating generalization to distance in psychological space. *Psychometrika*. 1957; 22(4):325–345.
- Stemberger, JP. The reliability and replicability of naturalistic speech error data: A comparison with experimentally induced errors.. In: Baars, BJ., editor. *Experimental slips and human error: Exploring the architecture of volition*. Plenum; New York: 1992. p. 195-215.
- Stemberger JP. Preventing perseveration in language production. *Language and Cognitive Processes*. 2009; 24(10):1431–1470.
- Vicente A, Martinez Manrique F. Inner Speech: Nature and Functions. *Philosophy Compass*. 2011; 6(3):209–219.
- Wheeldon LR, Levelt WJM. Monitoring the time course of phonological encoding. *Journal of Memory and Language*. 1995; 34(3):311–334.

Figure 1.

Error distributions for the same stimuli in three experiments testing phonemic similarity effects in overtly articulated and unarticulated (inner) speech. O&D's experiments (a, b) show larger similarity effects in articulated speech compared to unarticulated. The same crossover interaction fails to reach significance in CBM's Experiment 3 (c), where, in their words, "participants were once again much more likely (here, by a factor of 2.7) to substitute similar rather than dissimilar phonemes, regardless of whether the speech was overt or not," (p169). As CBM note, the (c) results "were consistent with [their] two earlier experiments."

Figure 2.

Models predict stronger odds-ratio similarity effects when errors are rare. a.) With the Luce-Shepard choice rule, more deterministic selections (larger μ) yield stronger error effects. Following the experiments, Similar and Dissimilar slip probabilities are calculated on separate trials. Parameters: $a_{\text{Target}}=1$, $a_{\text{Similar}}=.5$, $a_{\text{Dissimilar}}=.25$, $\mu=0:10$. b.-c.) Modulating activation noise in Dell's (1986) model produces similar behavior, additionally demonstrating that (b) odds-ratio effect estimates grow less precise when one has less data with which to estimate them (i.e. at lower error rates), and (c) higher error rates produce numerically larger effects (at least until selection approaches chance accuracy). Modulating phonological priming (not shown) has similar effects. Each point represents a 3,000-trial

replication, reflecting the expected variability between similar-sized experiments. Parameters: Lexical activation=1, Activation noise= $\mathcal{U}(0.2, 1.7)$, Competitor phoneme activation=0.1, Decay=.4, Connection weights=0.2. Representing inner speech (triangles) with less robust phoneme-feature connections (weights set stochastically per-'subject': $0 \leq \mathcal{N}(0.1, 0.1) \leq 0.2$), the model shows the same patterns with weaker effects.

Figure 3.

Similarity effects in articulated (circles) and unarticulated (triangles) speech, plotted against general measures of erroneous production. Regression lines depict fixed effects of the fitted model. The main effect of similarity places lines above 1.0 on the y-axis. The similarity-by-articulation interaction specifies a superior line for articulated speech. The similarity-by-error rate interaction specifies the lines' slope, suggesting similarity effects are stronger when production is less error-prone

Figure 4.

The similarity-by-error rate interaction generalizes to new data. The regression line, fit only to the articulated tongue-twister data (circles), uses a power function (based on the simulations in Figure 2). Similarity effects from several SLIP tasks (squares) follow the same function.

Table 1

Aggregated data for the five experiments.

Study	n	Trials per cell	p(Target error)	p(Other error)	Articulated			Unarticulated			Interaction	
					Sim	Dis	OR	Sim	Dis	OR	OR	OR
Oppenheim & Dell	128	4096	.017	.114	196	114	1.72	133	129	1.02	1.69	
2008	48	1536	.016	.098	69	38	1.79	40	52	0.75	2.38	
2010	80	2560	.018	.123	127	76	1.68	93	77	1.20	1.41	
Corley <i>et al.</i> , 2011	112	4608	.009	.049	130	47	2.88	105	58	1.84	1.57	
Experiment 1	32	1536	.007	.050	34	12	3.01	28	16	1.76	1.70	
Experiment 2	32	1536	.012	.057	53	21	2.62	50	29	1.73	1.51	
Experiment 3	48	1536	.008	.041	43	14	3.13	27	13	2.16	1.45	
Total	240	8704	.013	.064	326	161	2.06	238	187	1.27	1.62	

Table 2

Regression summaries for Part 1.

All experiments	Odds-ratio	Coef β	SE(β)	z
(Intercept)	0.01	-4.67	0.09	-49.77
Similarity	1.77	0.57	0.08	7.49
Articulation	1.11	0.10	0.07	1.43
Similarity:Articulation	1.62	0.48	0.14	3.44
Lab	0.41	-0.89	0.13	-6.72
Similarity:Lab	1.69	0.53	0.15	3.52
Oppenheim & Dell (2008; 2010)				
(Intercept, Expt 2008)	0.01	-4.35	0.16	-27.07
Expt 2010	1.28	0.25	0.15	1.72
Similarity	1.33	0.28	0.09	3.24
Articulation	1.15	0.14	0.09	1.62
Similarity:Articulation	1.69	0.52	0.18	2.98
Corley <i>et al.</i> (2011)				
(Intercept, Expt 1)	0.00	-5.36	0.20	-26.74
Expt 2	1.64	0.49	0.27	1.81
Expt 3	0.98	-0.03	0.28	-0.09
Similarity	2.30	0.83	0.12	6.79
Articulation	1.05	0.05	0.12	0.38
Similarity:Articulation	1.57	0.45	0.25	1.84

Table 3*General error rate regression summaries.*

Combined	Odds-ratio	Coef β	SE(β)	z
(Intercept)	0.01	-4.68	0.09	-50.60
Similarity	1.76	0.56	0.08	7.38
Articulation	1.11	0.10	0.07	1.44
Similarity:Articulation	1.62	0.48	0.14	3.44
General error rate*100	1.14	0.13	0.02	7.00
Similarity:General error rate*100	0.93	-0.07	0.02	-3.09
Articulated				
(Intercept)	0.01	-4.72	0.11	-44.05
Similarity	2.23	0.80	0.11	7.25
General error rate*100	1.14	0.13	0.02	5.77
Similarity:General error rate*100	0.93	-0.07	0.03	-2.15
Unarticulated				
(Intercept)	0.01	-4.81	0.11	-44.94
Similarity	1.37	0.32	0.11	2.87
General error rate*100	1.14	0.13	0.02	5.15
Similarity:General error rate*100	0.94	-0.06	0.03	-2.01