

Coherence across social and temporal scales

Meredith Tamminga (University of Pennsylvania)

Lacey Wade (University of Pennsylvania)

[To Appear in *The Coherence of Linguistic Communities: Orderly Heterogeneity and Social Meaning*
(*Routledge Studies in Sociolinguistics*) Editors Karen Beaman & Gregory Guy]

Abstract

What do we mean when we say “coherence,” and what are we looking for in our data when we search for it? This chapter takes coherence as a broadly-defined umbrella term for a class of covariation phenomena and puts forward some descriptive scaffolding to organize the study of these phenomena. Specifically, we argue that coherence research should specify the **unit** and **scope** (together, the **scale**) of analysis. We discuss how coherence can dissociate across scales of different social and temporal granularity. Finally, we sketch some connections between questions of how coherence aligns (or not) across scales and the theoretical implications that can be drawn from coherence evidence.

1 Introduction

An active frontier in variationist sociolinguistics is the integration of individual-level behavior and community level patterns. One topic attracting increasing inquiry along these lines is **coherence**. In their influential introduction to a coherence-themed *Lingua* issue, Guy and Hinskens define coherence as follows: “To the extent that linguistic variables systematically co-vary, they can be characterized as displaying coherence” (2016, 1). We believe that the study of coherence holds significant promise for our understanding of the relationship between the sociolinguistic individual and the speech community. However, empirical research in this area has been held back by a lack of consensus around both the conceptual definition of coherence and the question of what kinds of quantitative patterns should be taken as empirical evidence for coherence. What do we mean when we say “coherence,” and what are we looking for in our data when we search for it? In this chapter we take the position that adopting a very broad definition like that just given from Guy and Hinskens (2016) is the right approach, but that doing so naturally leads us to think of coherence as an umbrella term for a class of covariation phenomena, instantiated in different ways at different social and temporal scales. In this chapter we offer some steps toward a more systematic descriptive scaffolding to help make sense of this range of phenomena.

Specifically, we argue that to promote comparability across studies of coherence at various scales, researchers should explicitly define what we will call the **unit** and the **scope** of the analysis. To understand what we mean by these terms, envision a set of points in a scatterplot, with a best-fit line representing the correlation between those points. The points themselves typically represent some kind of summary statistic, such as an average rate of a variant from a single speaker during a single interview. We use the term **unit** to refer to the definition of what goes into each point, as well as any particular point itself. To fully define the unit, we have to say something about *whose* data we are averaging over and a *time span* that data comes from. We therefore refer to the **social dimension** and the **temporal dimension** in specifying the granularity of the unit. We also use these dimensions to describe the granularity of the **scope** of the analysis, our term for the set of points

the correlation line is fit to.¹ For example, if we computed a single correlation between a sample of white Philadelphians in a study, we could change the scope by adding participants of other races or computing separate correlations for participants of different genders. In Section 2, we unpack these distinctions and give some toy examples of what can happen to covariation patterns as we change the social and temporal granularity of the unit and the scope. We will use the term **scale** to refer to different combinations of unit and scope.

While there are many logically possible scales at which coherence could be investigated, some are more familiar to sociolinguists – and in some cases, so familiar as to have assumed an almost-invisible default status. The very idea of picking out a language variety (whether dialect, ethnolect, style, etc.) as an object of study assumes coherence at some scale, with multiple features each tending to occur more in one variety than another. Coherence at this large scale, such as that reflected in isogloss bundles separating regional dialects, might blend into our background assumptions, rather than standing out as an example of coherence.² At the same time, there is a particular smaller scale that many researchers treat as prototypical in the coherence literature: where the unit is the individual speaker’s average use of each variant (typically over a single interview, for methodological reasons), and the scope is a fairly large social grouping such as a speech community. As Oushiro puts it, “The main question is whether multiple variables in a community correlate in individual speakers’ usage: do speakers who tend to employ variant x of variable A also tend to employ variant y of variable B, or are variables independently embedded in language and society?” (2016, 116). The investigation of coherence at this scale typically asks whether interspeaker correlations mirror the broad associations between variables that characterize the varieties themselves; we suspect some would argue that this is the single phenomenon that the term “coherence” is intended to refer to.

Certainly it would be possible to define coherence narrowly to match the prototypical view of coherence just described, and some researchers have come close to doing so. We share the literature’s interest in coherence at this scale, believing that work in this vein brings useful new evidence to questions like how social and linguistic information are connected in our mental representations (e.g., Guy 2013) and which language users innovate and propagate language changes (e.g., Nevalainen et al. 2011; Tamminga 2019). However, even studies that seem to operate mostly under this understanding of coherence continue to pose questions or use the term in ways that involve changing the scale of the analysis. For example, Oushiro (2016) computes correlations within many different social categories (thereby changing the scope), and Guy writes that “when varying their speech styles, speakers might be expected to synchronize their choices on all of these variables at the same time” (2013, 64), suggesting a coherence analysis involving intraspeaker units. These are perfectly reasonable and interesting things to do, but they go beyond a possible narrow definition of coherence. We are also not the first to explicitly use the term coherence to label covariation patterns at different scales. For example, Gregersen and Pharaoh (2016) entertain the possibility that coherence could arise within an individual, in communicative groups, or in geographic regions. Once we allow for the term coherence to apply across different scales, it is not clear why we should arbitrarily stop at one granularity or another. The definition of coherence proposed by Guy and Hinksens, as *systematic co-variation of linguistic variables*, gives us no basis on which to constrain our understanding of coherence as a sociolinguistic construct to

¹ By “set of points” we mean the definition of that set, with the actual points included in the analysis being a sample of that set.

² Work in dialectometry concerned with the consistency of dialectometric data sets might be considered investigations into coherence at this larger scale. For example, computing Chronbach’s α for inter-item correlation bears similarities to the type of analyses conducted in work on coherence.

a certain scale or subset of scales. While it would be *possible* to give a definition of coherence that specified limits on these scale manipulations, to this point nobody has done so.

Our approach – to maintain a broad definition of coherence and then specify the scale at which it is investigated – has a number of consequences that go beyond concerns of definitional precision or methodological clarity. First, we make the point in Section 2 that the identification of any given entity as “coherent” or “incoherent” is only possible in reference to a particular unit and scope. Second, Section 2 makes clear that coherence at one scale may or may not correspond to coherence at another scale. This intuition already lies at the heart of the coherence literature, but without defining coherence analyses in terms of their unit and scope, it is difficult to articulate questions and results about these possible correspondences. The practice of defining unit and scope, then, is critical for drawing accurate connections between quantitative covariation data and larger theoretical frameworks. We can think of this in cautionary terms: without attention to questions of granularity, we risk incorrectly concluding that varieties or speakers are “incoherent” when we’re simply looking at the wrong scale. But the positive flip-side is that this approach opens up new territory for investigation, generating many possible empirical questions about the relationship between coherence *across* multiple scales. We believe those questions hold promise for the integration of different strands of the sociolinguistic literature, with the goal of moving toward a more complete understanding of the relationship between the individual and the group and between synchronic and diachronic perspectives on language.

2 Coherence at different scales

There are a range of empirical patterns at different social and temporal *scales* that would fall under Guy and Hinskens’ definition of coherence, some more common in the literature than others. Before we unpack these possibilities in Sections 2.1 and 2.2, we offer some additional terminological and conceptual clarifications. First, what it means for a given entity to be considered “coherent” requires some clarification. The term “coherent” as an adjective is used in multiple ways in the literature. A language variety is often described as coherent if there is interspeaker covariation of speaker means in some subset of the features that are taken to characterize the language variety (e.g., “If varieties are coherent, the variables associated with them should co-vary in the usage of individuals,” Guy and Hinskens 2016, p. 4). Alternatively, individuals are often described as coherent if they exhibit high (or low, or intermediate) average usage of some or all of these features (e.g., “Fourteen speakers have both High and Low assignments – these are speakers who we might categorize as “incoherent,” or lacking co-occurrence of the three features,” Becker 2016, p. 94). We accept this ambiguity between different senses, but we are now in a position to make some of these notions more precise. On our view, it is only possible to describe a unit as coherent within a specified scope, and only possible to describe a scope as coherent relative to a particular unit. We would say that a *unit* is coherent if it falls close to the line of best fit in our scatter plot correlating the relevant variants. Notice that we can only know whether a unit falls close to the best-fit line if we draw that line, and drawing the line requires us to define the scope. Relatedly, we cannot say whether an individual uses “high” or “low” rates of a variant without knowing the rates of other speakers for comparison. We will describe a *scope* as coherent if the best-fit line has a significant slope, but we can’t know the slope of the line without knowing what datapoints (i.e., units) the line is fit to. Notice that this approach still leaves us with an additional ambiguity to be cautious about: it is fairly intuitive to apply the adjective “coherent” to speakers as units (who fall close to the best-fit line in a given scope), or to speakers as scopes (where there is coherence across intraspeaker units within a single individual). Although we think our

definitions can help, the unresolved terminological ambiguity is a reason to be as specific as possible when using “coherent” as an adjective.

Second, we want to highlight that there are many decisions to be made in any investigation of coherence, including the language variety being targeted, the features being correlated, the population included in the analysis, the other factors being controlled, and the mathematical approach to aggregation over individuals, groups, or other units. While much could be said about these other decisions and their implications, we will focus on why we believe it is necessary to define the *scale* of any analysis in order to connect any empirical findings of (in)coherence within a language variety with theoretical claims.

When we talk about the *unit* of the analysis, the level of granularity at which we examine covariation can broadly be broken into individuals, groups of individuals, or intraspeaker points. Typically, each point is an average over some socially- and temporally-defined chunk. When the social granularity of the unit is the individual, the temporal granularity of that unit is often all of the tokens in an interview with that individual. If the social granularity of the unit is a group, the temporal granularity might be all of the tokens within all interviews for some socially-demarcated group. The unit may also be an intraspeaker chunk such as a short stretch of speech in a larger conversation, perhaps on a single topic or with a particular interlocutor. Merely saying that a correlation is intraspeaker, interspeaker, or intergroup does not fully define the unit: these units may still differ in their exact social granularity (e.g., the groups compared in an intergroup correlation may be broad regional categories or narrower categories like a community of practice) as well as their temporal granularity (e.g., units may comprise the average of each individual’s tokens elicited within an entire interview, or may more narrowly average over, say, conversational turns) and should be specified for both dimensions.

Independent of the social and temporal granularity of the unit of analysis, there are also various levels of social and temporal granularity for the *scope*, or the set of units included in the analysis. The social scope is often a pre-delimited speech community chosen to align with the variety under investigation. For example, in an investigation of coherence within New York City English, the scope of the data might be restricted to native English speakers who are from New York City. Of course, the methodological considerations in choosing a population, any narrower social grouping of that population, and the relevant features for that variety may not be straightforward and should be chosen carefully. The temporal scope of the analysis can also range from broader diachronic trajectories to smaller microtemporal slices. In the following subsections, we give a non-comprehensive discussion of some different social and temporal granularities of scope (Section 2.1) and unit (Section 2.2), taking a special interest in the ways that coherence at one scale may or may not align with coherence at other scales.

2.1 Social and temporal granularity of the *scope*

To illustrate the effect of social and temporal *scope* on coherence patterns, we hold the social and temporal *unit* constant at the level where a point represents a single individual’s average use of each variable throughout the course of a single conversation and vary the granularity of only the scope. We focus here on hypothetical instances where what appears to be non-coherence under one scope turns out to be compatible with coherence under a broader or narrower scope. We start by focusing on the *social granularity* of the scope to which we restrict the analysis. This could be a social grouping for which there is a 1-to-1 mapping between that variety and the individual or group (e.g., social class or dialect background), or some grouping of contexts that a single individual may move across frequently such that there exist many social mappings between a single

individual and a range of social categories (e.g., a persona, stance, or orientation). Take, for instance, the lack of interspeaker coherence shown in the left facet of Figure 1. If we *broaden* the scope of the correlation by including individuals from two other social groups, we see that there is in fact coherence observable at this broader social scale. However, the interspeaker coherence that we see is actually driven solely by group membership, outlined by the differed colored points, rather than the individuals themselves. Group A uses the least amount of both variants, while the Group C uses the highest amount, with Group B intermediate between these two; however, within any given sub-group the individuals exhibit no such correlations. In this example individuals are constrained by their social group patterns but still able to produce high rates of variant X and low rates of variant Y (or vice versa) relative to the other members of their group. We refer to this as group-driven interspeaker coherence because it is individuals' group membership that underlies the interspeaker correlation. Note that in some cases, when researchers are interested in detecting interspeaker correlations that are *not* group-driven, they may control out known social factors prior to the coherence analysis.

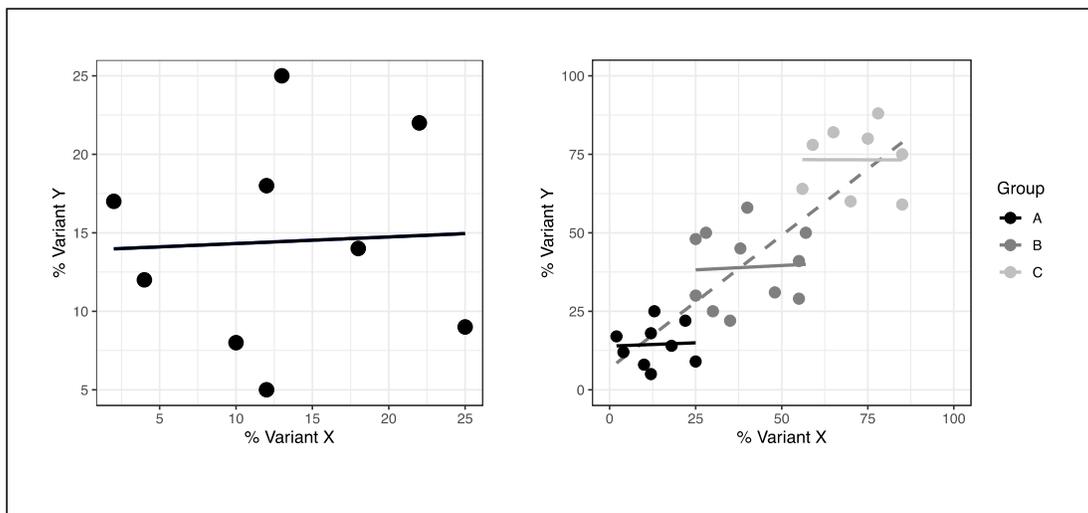


Figure 1: Lack of interspeaker coherence with smaller scope (left) compatible with group-driven interspeaker coherence at larger scope (right)

Similarly, if we *narrow* the social scope to smaller groups that comprise the larger population, we may also find coherence that was otherwise not apparent. As Figure 2 shows, interspeaker incoherence when the scope is one broader group may mask coherence at a narrower scope, perhaps indicating that the smaller groups are more relevant to the chosen features than the larger group. We refer to it as individual-driven interspeaker coherence when we do not have reason to believe that the interspeaker correlation is driven by group membership. Of course, knowing when this is the case is a hard question; in our example in Figure 2 the combination of different groups actually obscures the interspeaker correlations that occur *within* groups, so it seems clear that they are not the source of the correlations, but in principle there could always be another subdivision we have missed. This analysis could just as easily focus in on only one of the three groups and, having sufficiently narrowed the social scope of the data set, would still exemplify individual-driven interspeaker coherence. Of course, it is also possible that narrowing in on smaller social groups

would reveal different interspeaker patterns of (in)coherence within each group. This is essentially what Oushiro (2016) finds in Brazilian Portuguese. Local variables only exhibited coherence across a subgroup of speakers whose parents were born in São Paulo, whereas both supra-regional variables exhibited coherence across speakers in the community, regardless of social grouping. This again shows why it is necessary to choose variants carefully in light of the social scope of the analysis and also illustrates the importance of accompanying investigations of coherence with good ethnographic work.

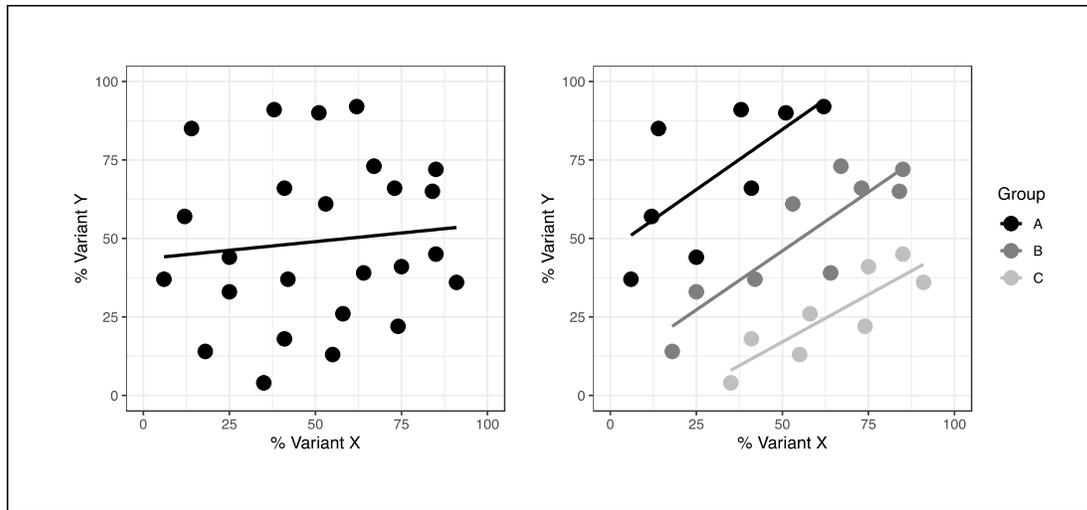


Figure 2: Lack of interspeaker coherence at larger scope (left) compatible with intragroup interspeaker coherence at smaller scope (right)

Finally, as Figure 3 shows, it is possible to have interspeaker coherence that is both individual-driven and group-driven, which we call community coherence. Here we see a positive interspeaker correlation between Variant X and Variant Y, mirrored by the intergroup pattern. If the social scope were not sufficiently broad to encompass all three groups, however, this would look no different from individual-driven interspeaker coherence within any one group. Note that the distinction between individual-driven, group-driven, and community coherence crucially relies on the *relationship* between two different scopes, highlighting the importance of examining the relationship across different scales.

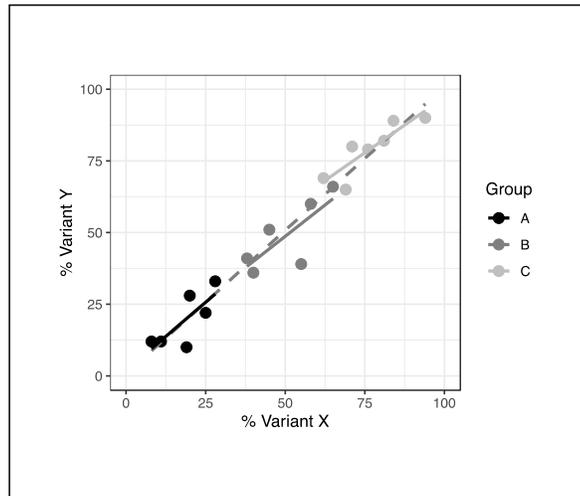


Figure 3: Alignment between interspeaker correlations and group-driven correlations yields community coherence.

So far we have been outlining possibilities for changing the social scope of our coherence analyses. We must also consider the temporal scope of the data, which may span many years, a single generation, a single conversation, or some smaller temporal dimension. The temporal scope for the prototypical scale of coherence analysis we've discussed is often a time period (perhaps a year or two) in which interviews were collected, treated as a single synchronic moment. However, these usually consist of a range of speaker birth years that are representative of the community at a given time, so treating these as a single narrow temporal scope may elide diachronic trajectories that, under the apparent-time assumption, could be present in the data. If we take the apparent-time construct seriously, the temporal scope of most data sets is fairly broad, consisting of multiple generations of speakers representing the state of the language variety at various points in time.

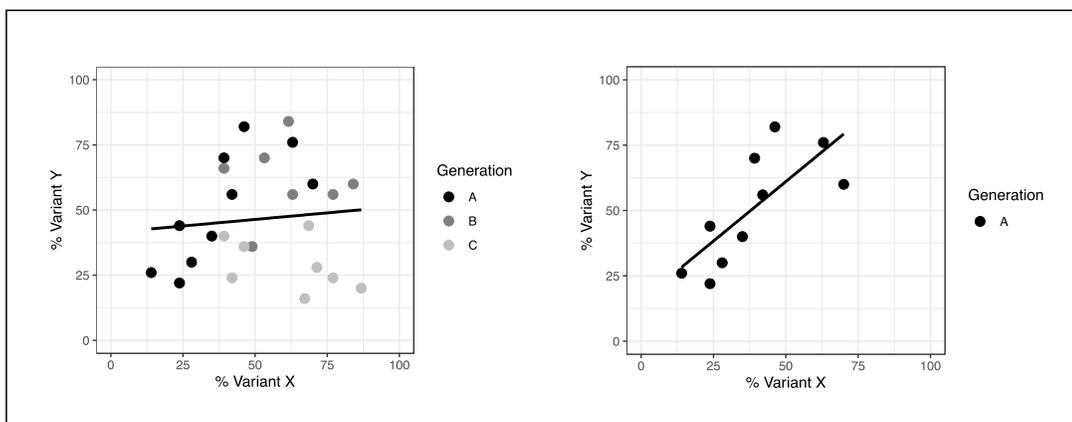


Figure 4: Larger diachronic chunk (left) vs. single generation (right).

As with the social scope, there are cases where such a large temporal scope may mask interspeaker coherence at a smaller granularity. For instance, take the lack of interspeaker coherence in the left facet of Figure 4, which encompasses three generations that might span the entire birth year range of a typical data set from a single community. When we narrow in on only the earliest generation, Generation A, we see that there is in fact interspeaker coherence that was masked by the diverging diachronic trajectories of the other generations. Such a pattern might occur if, for instance, Generation A grows up at the beginning of a linguistic change that eventually goes on to partially reverse. If we take the partial reversal of Philadelphia vowel shifts as a hypothetical example, we can imagine that as generation A shifts toward greater usage of DOWN raising and FACE raising simultaneously, we might see a positive correlation between these variants at the interspeaker level.³ Generation B might then continue on with FACE raising but begin to reverse DOWN raising, and Generation C may continue this reversal further. At the larger temporal scope, then, there are various patterns of covariation because there are different linguistic changes taking place within different time spans. Really, we might consider a *negative* correlation between DOWN and FACE raising in Generation C to be coherence because this would indicate the trajectory of the change *for this subset of speakers*⁴. This is yet another example of how the variants chosen for investigation in a coherence analysis should be closely considered in relation to the scope of the analysis. In cases of ongoing linguistic change in particular, different feature patterns may be representative of different temporal chunks in the community. This is further complicated by the idea that language change is inherently disruptive (Labov, 2001b, 4–6) and therefore might lead to incoherence. For example, Guy et al. (2021) find that in Brazilian Portuguese, the newest variant, diphthongal EN, is the only variant that doesn't correlate with several others, which they suggest is because social meanings are not inherently attached to linguistic features but come to be associated with them through use over time. Note, however, that Tamminga (2021) finds that some novel changes *do* covary at the outset of a change, potentially even foreshadowing later developments. Defining the temporal scope, then, is just as important as defining the social scope and has implications for the appropriate variants to investigate as potentially cohering within the population.

2.2 Social and temporal granularity of the *unit*

Just as with the *scope* of the analysis, broadening or narrowing the *unit* of analysis impacts whether we observe coherence and has important methodological and theoretical implications. So far, we have focused on interspeaker coherence, looking at datapoints representing individual means. Notice, though, that some of our discussion in Section 2.1 did implicitly make reference to units of a broader social granularity. In our toy version of group-driven interspeaker coherence, for example, the relationship we described between the average rates of variants X and Y in Groups A, B, and C properly requires reference to group-level units. We could also narrow the social granularity of the unit so that there are multiple points per individual in the analysis, perhaps representing different topics, interlocutors, or contexts. Here we unpack some examples of how the interspeaker covariation patterns we looked at in Section 2.1 might relate to intergroup patterns

³ Note that we may want to narrow the temporal scope even further to determine if any interspeaker coherence we see within Generation A is driven by the span of birth years (i.e., reflecting the change over time) or if it truly picks up on differences in individuals at a synchronic moment in time (i.e., the leaders of the new changes).

⁴ This is a hypothetical example; see Tamminga (2019, 2021) on how Philadelphia vowel changes actually correlate at different points in time.

when we broaden the social granularity of the unit from an individual to a social group, saving some briefer discussion of narrower units for the end of the subsection.

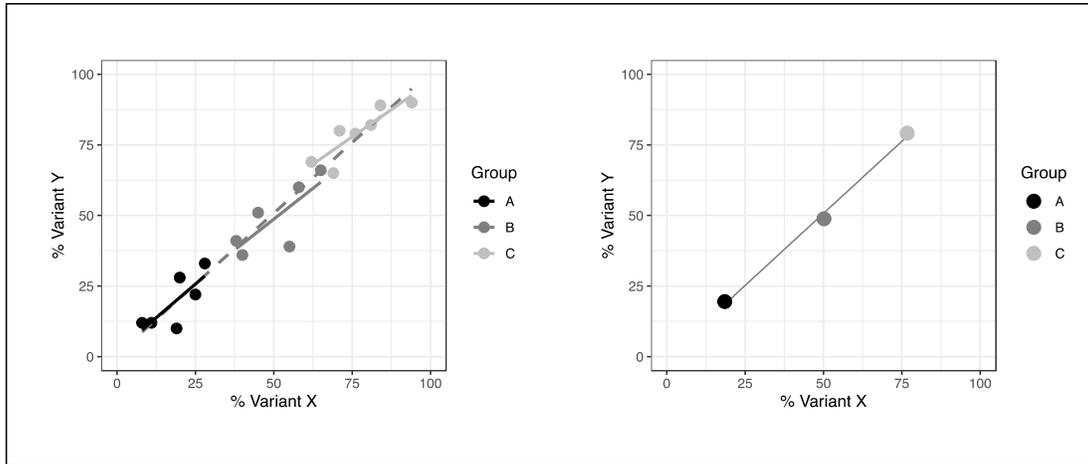


Figure 5: Community coherence (left), intergroup coherence (right).

In the idealized case, interspeaker and intergroup covariation patterns will align. This is what we see in the cases of both community coherence and group-driven interspeaker coherence. Figure 5 shows community coherence, with a positive correlation between variants X and Y at the interspeaker level in the left facet. When groups rather than individuals are the social unit, X and Y again positively correlate. Similarly, individuals and groups may align *without* the type of intragroup coherence we see with community coherence. Figure 6 shows group-driven interspeaker coherence on the left, which is essentially the same empirical finding as intergroup coherence on the right, as the pattern exemplified in both stems from the groups rather than the individuals. In these first two cases, while interspeaker and intergroup correlations align, broadening the social unit obscures the difference, namely whether the interspeaker coherence patterns are group-driven or individual-driven. Only if we consider multiple levels of social granularity can we distinguish between these two types of coherence.

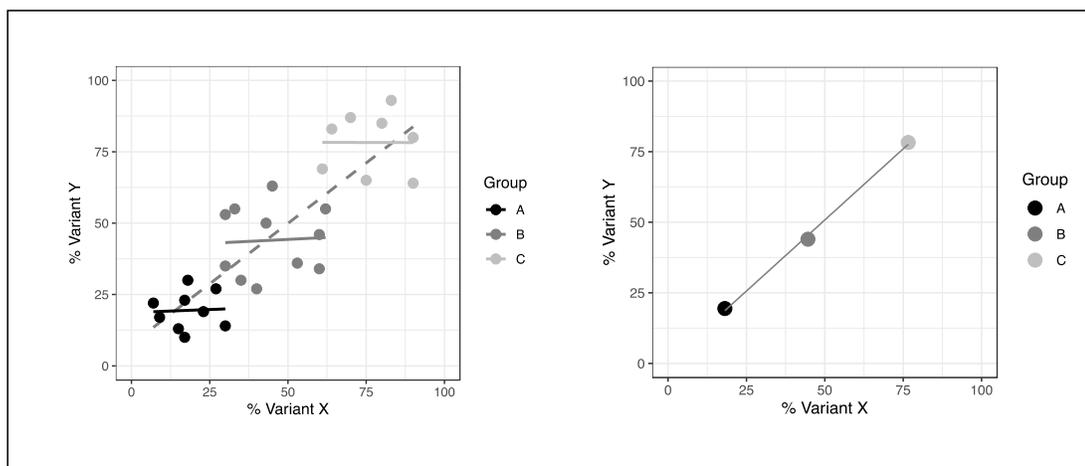


Figure 6: Group-driven interspeaker non-coherence (left) equates to intergroup coherence when the granularity of the unit is broadened.

It is also of course possible that intragroup and intergroup patterns do not align, as shown in Figure 7. The left facet, repeated from Figure 2 above, shows interspeaker coherence at a sufficiently narrow social scope (i.e., within social groups), but what we did not highlight in the earlier discussion of this pattern is that the accompanying intergroup correlation is in the opposite direction. This might occur when individuals who are part of different subgroups adhere to the same community pattern, but their inclusion in a subgroup poses certain constraints on their overall usage of one or both variants. That is, X might increase (or decrease) in tandem with Y, but one subgroup might generally use more X, while a different subgroup uses more Y. This is close to what Gregersen and Phrao (2016) find, where Danish raising of /EN/ and /æ/ positively correlate across individuals within the regional subgroups of Vinderup and Næstved, yet there is a negative correlation across these two subgroups because the region with the greatest /EN/ raising (Næstved) is not the region with the greatest /æ/ raising.

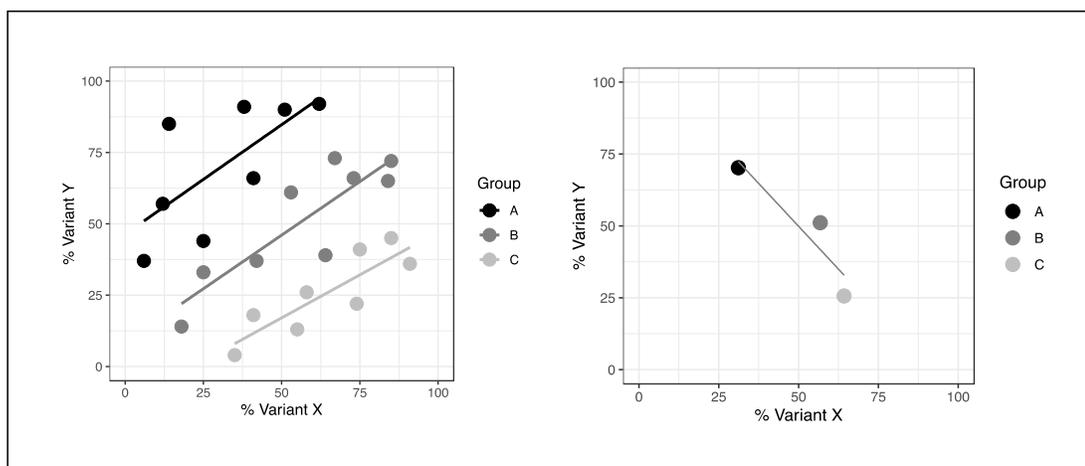


Figure 7: Interspeaker correlations at the narrower social scope within groups are positive, but the intergroup correlation is negative.

Another logical possibility for the relationship between interspeaker and intergroup covariation patterns is that different groups may exhibit different directions of interspeaker coherence within different subgroups. One interesting possibility is that the variant usage patterns that define a social group may be reflected in interspeaker correlations within that group. Figure 8 shows such a case where interspeaker patterns differ within the different social groups. Group A shows a positive correlation between variant X and Y while groups B and C show negative correlations. This may appear to be an example of incoherence, particularly since the overall trend for all speakers is a lack of correlation, evidenced by the dotted line. However, when we compare these interspeaker patterns to the intergroup pattern in the right facet of the graph, we see that individual speakers are actually mirroring the patterns that define the group, which may be indicative of different social groups that actually reflect different language varieties. To make this point more concrete, imagine that variants X and Y are associated with the Polish NYC ethnolect (this is a hypothetical example built on top of a real study from Newlin-Łukowicz 2016). If variant X reflects orientation toward Poland but variant Y reflects orientation toward NYC, we might expect that NYC social groups with differing orientations toward Poland and the U.S. would differently utilize or reject these variants, while X and Y might still define the “Polish NYC ethnolect” as a whole, because it lies at the intersection of NYC features and Polish features. If we take Group A to be firmly oriented toward the Polish NYC community, it makes sense that they would use high rates of both variants overall and that additionally these variants would positively correlate at the interspeaker level (with speakers who have the strongest Polish NYC orientation using the highest rates of both). Group B, then, may be U.S. oriented speakers who reject the Polish-associated X variant (but still likely use it more than non-Polish New Yorkers) but still use high rates of the NYC Y variant, thus exhibiting a moderate negative correlation at the interspeaker level. On the other hand, Group C may be Poland-oriented and exhibit the opposite pattern, using high rates of the Polish-associated X variant and rejecting the NYC Y variant (though likely using higher rates of it than non-New Yorkers), also exhibiting a negative correlation at the interspeaker level.

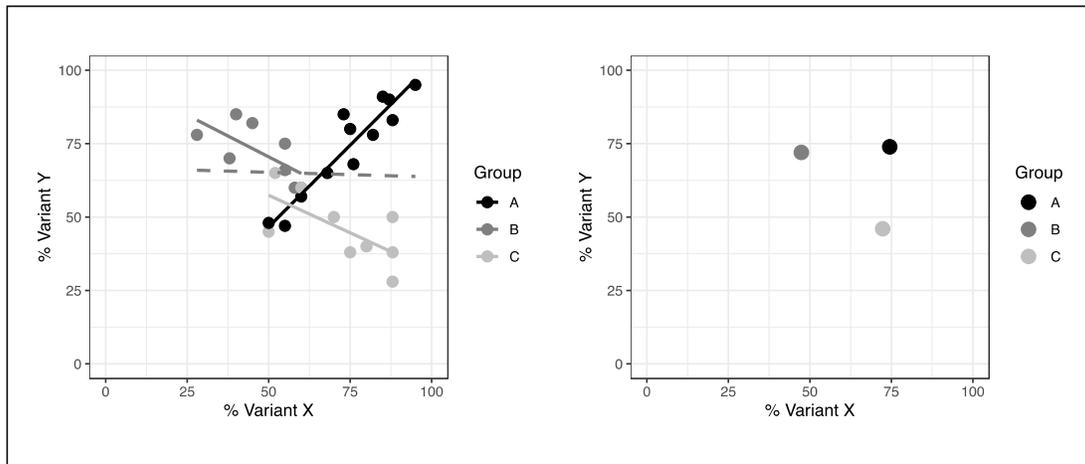


Figure 8: Different patterns of interspeaker coherence that align with the features that define each group within the community.

While on one scale this looks like speakers in the three groups differently draw from the variants available in the broader community, on our view it is still coherence when both the social scope and the corresponding language variety are accordingly narrowed. Newlin-Łukowicz (2016) reports findings that bear some resemblance to the group relationships in our hypothetical example. While interspeaker coherence was not apparent for the “Polish ethnolect” in NYC, smaller social groups within this did exemplify particular linguistic patterns, with “the conceptually intermediate group (i.e., speakers with an orientation toward Polish New York City) overlapping in their choice of variants with each of the remaining groups [(i.e., Speakers oriented toward America and speakers oriented toward Poland)]” (Newlin-Łukowicz, 2016, 110). Newlin-Łukowicz suggests that this finding provides evidence against coherence and for bricolage at the level of the ethnolect, citing as evidence the finding that Polish NYC speakers utilize features from NYC English and Polish English apparently independently. However, narrowing the social scope and language variety driven by a cross-section of regional and ethnic orientation could potentially reveal coherence that we do not see at the level of the larger ethnolect (Newlin-Łukowicz does not compute these intragroup interspeaker correlations, so we can’t know whether this possibility would be realized). This again illustrates the potential value of considering the relationship between the language variety and the features taken as representative of that variety.

Our examples in this subsection have focused on broadening the granularity of the social unit from the individual to the social group. We could also narrow the social grain of the unit to an intraspeaker measure. These intraspeaker units might be matched across individuals (e.g., each individual has a single point for each of X topics) or not (e.g., each individual has multiple points that don’t necessarily correspond to the division of points of another individual). Much of the classic style-shifting literature can be thought of as involving intraspeaker units within the social scope of a single individual. Coupland’s (1980) study of a Cardiff travel agent’s stylistic use of multiple variables is a classic example. Note that it is also possible to use intraspeaker units but a scope that includes multiple individuals. For example, Labov (2001a) divides a collection of individual sociolinguistic interviews into casual and careful speech on an utterance-by-utterance basis, then presents average rates of ING and DH for careful and casual speech, showing that both

ING and DH respond similarly to that stylistic division. In these cases, notice that the temporal granularity of the unit was also narrowed from the length of a conversation, which feels intuitive because the study of intraspeaker variation in sociolinguistics tends to be about microtemporal patterns. Note, however, that it would be possible to investigate intraspeaker coherence with units at a larger temporal grain: for example, we could ask whether multiple variables shift in tandem across decade-sized chunks of an individual's lifespan.

Finally, if we turn to the relationship between scales involving interspeaker and intraspeaker units, we might expect a similar range of possibilities as we saw when we compared interspeaker to intergroup coherence: we could find that what looks like interspeaker incoherence is simply the aggregate of multiple coherent intra-individual patterns, or that strong interspeaker coherence is not paralleled by intraspeaker coherence at some granularity, or that in fact interspeaker coherence and intraspeaker coherence on some scale might align.

2.3 Additional considerations

Unit and *scope* are not the only two dimensions in which the study of coherence would benefit from additional terminological, methodological, and theoretical attention. Here we briefly sketch some additional dimensions around the study of coherence that we believe are under-theorized, to encourage systematic treatments of those areas in parallel to this chapter's attempts to systematize the scale of the analysis.

Multiple variables The literature we have discussed mostly focuses on pairwise linear comparisons between variables. It is likely desirable to model relationships between multiple variables simultaneously instead of in a pairwise fashion (see e.g., van Meel et al., 2016). Some obvious methodological candidates that could be further explored for this purpose are Principal Components Analysis (PCA) and hierarchical clustering analyses. Such approaches could reveal more complex coherence patterns among multiple variables in a multidimensional space. Another relevant model might be Biber's (1992) use of multidimensional analysis for discourse complexity.

Dealing with continuous dimensions Defining the scope of a coherence analysis, in the terms we have given here, requires specifying a clearly-demarcated group of points. But the groupings we have entertained as examples may abstract over social dimensions that are in reality continuous: discrete generations are an abstraction over the continuous dimension of when a speaker was born, discrete socioeconomic classes might represent an abstraction over a more continuous socioeconomic status dimension (or multidimensional space), and so on. We are not sure what it would look like to operationalize coherence such that it can vary along a continuous independent variable dimension without binning that independent variable, but it is an interesting question.⁵

The question of whether coherence can be modeled over continuous dimensions also presents itself when we think about how microtemporal units are sequenced in time. We think that cases where the moving averages of two variants appear to track each other closely as time elapses during a conversation, as illustrated in Figure 9, could be thought of as a type of coherence. Tamminga et al. (2016) call this microcovariation and give a couple of real examples. In Figure 9, the apparent relationship between the rates of variant X and variant Y is offset in time and therefore may not show up as a simple correlation between rates in time slices if those rate observations are dissociated from the order in which they occur. As a result, capturing the full picture of the

⁵ Thanks to Greg Guy for asking it.

relationship between these moving averages may require techniques from, for example, time series analysis. Relatedly, we could look for coherence at a very fine-grained scale by looking at the microtemporal co-occurrence of individual variant tokens, with no unit-based averaging at all, as in the “style clusters” of Podesva (2008).

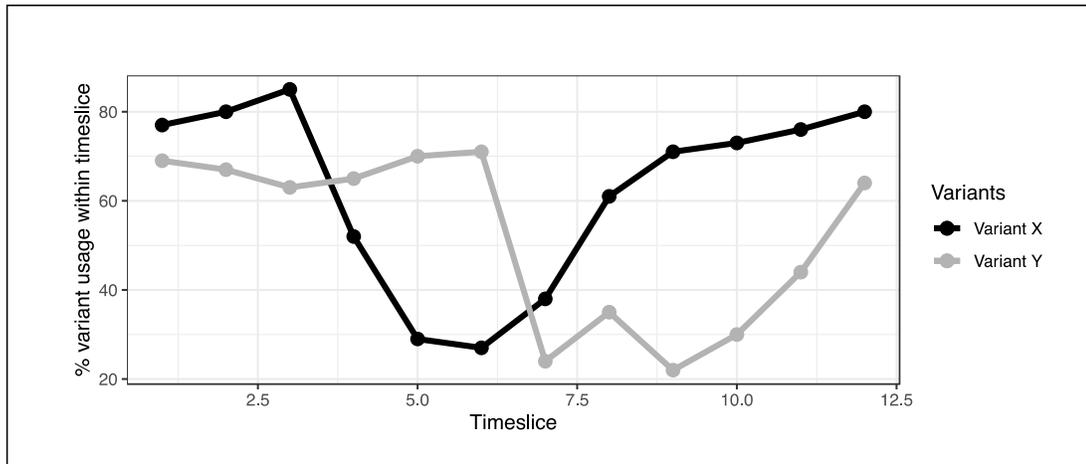


Figure 9: Microcovariation between two variables as time elapses

Beyond single point measures For all of the different possible units we have been discussing, we have only considered cases where each unit is characterized by a single point measure, such as a variant frequency or a central tendency. These measures can be made more accurate by controlling for other factors conditioning the variation before producing the unit-level summary statistics; Tamminga (2021), for example, uses speaker random intercepts taken from linear mixed effects regressions as the unit measures. More interestingly, we could move beyond single point measures into the study of how units might exhibit similar or different linguistic systems in some more complex way.⁶ This would open up a new set of questions: for example, one could ask whether speakers who exhibit a certain pattern of following segment conditioning for Variable A are more likely to exhibit a related pattern for Variable B, or whether speakers who allow Syntactic Construction X across the board are more likely to use Sociolinguistic Variable Y than speakers who use Construction X only in certain environments. Here we think the coherence literature connects up with the individual differences literature (for example, see MacKenzie 2019 for a sociolinguistic perspective and Yu and Zellou 2019 for a recent review in the domain of phonetics and phonology) because it raises the possibility of hidden subpopulations with qualitatively different linguistic systems.

3 Theoretical implications of coherence within and across scales

So far we have focused on making terminological and conceptual distinctions, mostly in fairly abstract terms. The importance of this enterprise comes into sharper focus when we consider the implications of coherence for sociolinguistic theory. Guy and Hinskens (2016) propose that

⁶ Thanks to Benedikt Szmrecsanyi for this suggestion.

different theoretical frameworks give rise to different predictions about coherence. The first perspective they outline is what they call the “orderly heterogeneity” view, in which a central phenomenon of interest is the parallel stratification of multiple sociolinguistic variables across demographic groups and stylistic contexts. They suggest that a sociolinguistic view giving primacy to facts about parallel correlations between class and multiple variables, or between style and multiple variables, should lead us to expect coherence: “variants (or rates of use of variants) that index a given style, status, or a social characteristic should co-occur. Coherent middle class speakers would use all the variants associated with their status, and speakers who are coherently signaling a ‘casual’ style would use all the ‘casual’ variants” (Guy and Hinskens, 2016, 2). The second perspective that Guy and Hinskens outline is what they call the “bricolage” view, which could also be called a Third Wave view. On this view, a central phenomenon of interest is the stylistic practice of individual speakers, who dynamically construct their social identities in part through the use of socially-meaningful linguistic variants. Guy and Hinskens suggest that this view does not generate the same strong coherence predictions that the orderly heterogeneity view does, saying: “Given that indexicalities, identities, and purposes are all profoundly diverse, the particular mix of variants used by a particular speaker on a particular occasion will be drawn from a selection that is essentially infinite. Hence there is no reason to expect, a priori, that such mixes will be coherent in the sense defined above” (Guy and Hinskens, 2016, 2).

Guy and Hinskens themselves acknowledge that this dichotomy is too stark, in part because as Eckert points out, “since a stylistic move is to be put out into a community for the purpose of being interpreted, speakers select resources on the basis of their potential comprehensibility in that community” (Eckert, 2004, 44) — in other words, the choices a speaker makes in stylistic practice are far from “essentially infinite.” However, the basic dichotomy has been widely adopted into the framing of recent research on coherence. In the framework we’ve advocated for here, sociolinguistic theories can be differentiated not in terms of whether they do or do not predict coherence, but in terms of the scale at which they generate coherence expectations. We therefore briefly revisit Guy and Hinskens’ predictions to illustrate how they might be recast within an understanding of different scales of analysis.

Notions of covariation play an integral role in Third Wave conceptions of stylistic practice. Guy and Hinskens (2016) note that the California Style Collective defines style as “a clustering of linguistic resources and an association of that clustering with an identifiable aspect of social practice” (1993). Bundles of features across different variables constitute different persona styles that can be shared across multiple speakers (as in Zhang’s 2005 study of Beijing yuppies or King’s 2018 study of African American personae in Rochester, NY) but can also structure intraspeaker variation within a single individual (as in Podesva’s 2008 case study of distinct personae used by a gay medical student). Moreover, Podesva (2008) points to “style clusters” that appear when features cluster together microtemporally, as utterances unfold in real time. Guy and Hinskens discount these kinds of clustering as “more or less temporary and situation-specific” (Guy and Hinskens, 2016, 3) and therefore not coherent. But in our view, persona styles and microtemporal style clusters reflect coherence at smaller scales than the “default” conception of coherence. The discussion of Figures 8 and 9 gave some ways of thinking about coherence patterns at these smaller scales.

Guy and Hinskens’ suggestion that orderly heterogeneity might be expected to give rise to interspeaker covariation, meanwhile, is derived from the observation that multiple variables each separately covary with demographic categories such as socioeconomic class. These correlations might be realized as intergroup correlations between the variables themselves such as those shown in the right-hand facets of Figures 5 and 6, which in turn suggests the hypothesis that the slope of

the interspeaker correlation in the corresponding left-hand facets should be positive. What is not clear is whether orderly heterogeneity might predict community coherence (as in Figure 5) or intergroup coherence (as in Figure 6) at this particular scale. On a view such as Labov's that "the individual does not exist as a unit of linguistic analysis" (2012, 267), we might predict an intergroup-without-intragroup coherence pattern like in Figure 6, where further individual differences within the group (if it is defined at the right social granularity!) reflect only noise. Becker's (2016) discussion of individual coherence, on the other hand, provides an example of an analysis at this scale that connects individuals to the social stratification of orderly heterogeneity yet predicts community coherence as seen in Figure 3. These questions come down to differences in the mechanism by which factors like socioeconomic class are thought to shape sociolinguistic behavior. However, if you start from the theoretical premises of the orderly heterogeneity view as Guy and Hinskens characterize it, there is at least one plausible interpretation of those premises that says that coherence should *not* arise at some smaller scales, such as smaller social scopes or smaller temporal units.

Our goal in rehashing these influential points from Guy and Hinskens is not to offer firm predictions of our own, but rather to illustrate that connecting theoretical frameworks to predictions about coherence depends crucially on issues of the social and temporal scale of the analysis. If we understand coherence as a phenomenon that can arise at many different scales, we need to take those scales into consideration when we think about when and why we expect coherence. For any given research question about coherence, the social and temporal granularity of the scope and unit of analysis must be chosen deliberately alongside the language variety and variants being investigated, because those choices affect our ability to detect coherence. Without careful attention to these issues as methodological choices and explicit theorizing of why coherence should arise at what scales, we risk drawing conclusions that miss clustering behavior at other, possibly even neighboring, scales.

The task of investigating coherence across so many different possible scales (not to mention varieties and variables) may seem herculean. But we believe the complexity of this space is a reason to engage, rather than pull back. It is already apparent that there is no blanket guarantee of parallels between coherence at one scale and coherence at the next scale up or down. But knowing that coherence patterns can dissociate across scales as we have defined them here puts us in a position to ask a new set of questions about when coherence *does* align across adjacent, or even non-adjacent, scales. To be sure, exploring these questions will require investment in new theoretical reasoning about the mechanisms of linguistic perception and production that might lead to coherence. A consequence of the framework that we've sketched out here is that there is no one-size-fits-all methodological or theoretical approach to the study of coherence. But by pursuing a better understanding of the relationships between coherence patterns at different scales, we stand to make new progress in integrating research aimed at different scales of analysis. In doing so, we may shed new light on the relationship between individuals' sociolinguistic identities and broader varieties such as dialects or ethnolects, as well as the relationship between language use in moments of face-to-face interaction and centuries-long language change.

References

- Becker, Kara. 2016. Linking community coherence, individual coherence, and bricolage: The co-occurrence of (r), raised BOUGHT and raised BAD in New York City English. *Lingua* 172–173:87–99.
- Biber, Douglas. 1992. On the complexity of discourse complexity: A multidimensional analysis. *Discourse Processes* 15(2):133–163.
- California Style Collective. 1993. Personal and group style. Paper presented at the 22nd annual conference on New Ways of Analyzing Variation in English, Ottawa.
- Coupland, Nikolas. 1980. Style-shifting in a Cardiff work-setting. *Language in Society* 9(1):1–12.
- Eckert, Penelope. 2004. The meaning of style. In *Proceedings of the Eleventh Annual Symposium about Language and Society–Austin*. Texas Linguistic Forum 47. 41–53.
- Gregersen, Frans. and Phrao, Nicolai. 2016. Lects are perceptually invariant, productively variable: A coherent claim about Danish lects. *Lingua* 172–173:26–44.
- Guy, Gregory. R. 2013. The cognitive coherence of sociolects: How do speakers handle multiple sociolinguistic variables? *Journal of Pragmatics* 52:63–71.
- Guy, Gregory. R. and Hinskens, Frans. 2016. Linguistic coherence: Systems, repertoires and speech communities. *Lingua* 172–173:1–9.
- Guy, Gregory. R., Oushiro, Livia., and Mendes, Ronald. (2021). Indexicality and coherence. In *The coherence of linguistic communities: Orderly heterogeneity and social meaning (Routledge Studies in Sociolinguistics)*, K. Beaman and G. Guy (eds). London/New York: Routledge.
- King, Sharese. 2018. *Exploring social and linguistic diversity across African Americans from Rochester, New York*. PhD thesis, Stanford University, Palo Alto, CA.
- Labov, William. 2001a. The anatomy of style-shifting. In *Style and Sociolinguistic Variation*, P. Eckert and J.R. Rickford, (eds). Cambridge: Cambridge University Press. 85–108.
- Labov, William. 2001b. *Principles of Linguistic Change, Vol. 2. Social Factors*. Oxford: Blackwell.
- Labov, William. 2012. What is to be learned: the community as the focus of social cognition. *Review of Cognitive Linguistics* 10(2):265–293.
- MacKenzie, Laurel. 2019. Perturbing the community grammar: Individual differences and community-level constraints on sociolinguistic variation. *Glossa: A journal of general linguistics* 4(1):1–23.
- Nevalainen, Terttu, Raumolin-Brunberg, Helena, and Mannila, Heikki. 2011. The diffusion of language change in real time: Progressive and conservative individuals and the time depth of change. *Language Variation and Change* 23:1–43.
- Newlin-Łukowicz, Luiza. 2016. Co-occurrence of sociolinguistic variables and the construction of ethnic identities. *Lingua* 172–173:100–115.
- Oushiro, Livia. 2016. Social and structural constraints in lectal cohesion. *Lingua* 173:116–130.
- Podesva, Robert. J. 2008. Three sources of stylistic meaning. In *Proceedings of the Symposium about Language and Society - Austin 15*. Texas Linguistic Forum 51. 134–143.
- Tamminga, Meredith. 2019. Interspeaker covariation in Philadelphia vowel changes. *Language Variation and Change* 31(2):1–15.
- Tamminga, Meredith. 2021. Leaders of language change: Macro and micro perspectives. In *Language Variation – European Perspectives VIII: Selected Papers from the Tenth International Conference on Language Variation in Europe (ICLaVE 10)*, H. van de Velde, N.H. Hilton, and R. Knooihuizen (eds). Amsterdam: John Benjamins. 270–289.

- Tamminga, Meredith, MacKenzie, Laurel, and Embick, David. 2016. The dynamics of variation in individuals. *Linguistic Variation* 16(2):300–336.
- van Meel, Linda, Hinskens, Frans, and van Hout, Roeland. 2016. Co-variation and varieties in modern Dutch ethnolects. *Lingua* 172–173:72–86.
- Yu, Alan C. and Zellou, Georgia. 2019. Individual differences in language processing: Phonology. *Annual Review of Linguistics* 5:131–150.
- Zhang, Qing. 2005. A Chinese yuppie in Beijing: phonological variation and the construction of a new professional identity. *Language in Society* 34:431–466.