# A Language-Based Method for Assessing Symbolic Boundary Maintenance between Social Groups

Anjali M. Bhatt*[1,2], Amir Goldberg[3], and Sameer B. Srivastava[4]

[1]Harvard Business School
[2]Santa Fe Institute
[3]Graduate School of Business, Stanford University
[4]Haas School of Business, University of California, Berkeley

**Abstract**

Symbolic boundaries—the conceptual markers people use to differentiate themselves from others—are at the root of intergroup social boundaries. When the social boundaries between groups are breached, the tendency for people to erect and maintain symbolic boundaries tends to intensify. Drawing on extant perspectives on boundary maintenance, we distinguish between two strategies that people pursue as they seek to maintain symbolic boundaries: *boundary retention*—that is, entrenching themselves in pre-existing symbolic distinctions—and *boundary reformation*—that is, innovating new forms of symbolic distinction. Traditional approaches to measuring symbolic boundaries—interviews, participant-observation, and self-reports—are ill-suited to detecting fine-grained variation in these two forms of boundary maintenance. To overcome this limitation, we use the tools of computational linguistics and machine learning to develop a novel approach to measuring symbolic boundaries based on interactional language use between group members before and after they first come into contact with one another. Specifically, we construct measures of boundary retention and reformation from a set of random forest classifiers that quantify group differences based on pre- and post-contact linguistic styles (as measured by the well-established LIWC lexicon). We demonstrate the utility of this method by applying it to a corpus of email communications from a mid-sized financial services firm that acquired and integrated two smaller firms. Our findings indicate that: (a) evidence of persistent symbolic boundaries can be detected for up to 18 months after a merger; (b) acquired employees exhibit more boundary reformation and less boundary retention than their counterparts from the acquiring firm; and (c) individuals engage in more boundary retention, but not reformation, when their local work environment is more densely populated by ingroup members. We discuss how our conceptualization and measurement of symbolic boundaries can be extended to the study of culture in a wide range of intergroup contexts and highlight implications for computational approaches to measuring culture.

*Corresponding Author: Anjali M. Bhatt, ambhatt@hbs.edu

## Introduction

The most consequential social interactions, whether cooperative or competitive in nature, often occur at the boundaries between social groups. For this reason, a prominent literature in cultural sociology seeks to understand how symbolic boundaries—the conceptual distinctions, interpretive strategies, and cultural traditions that people use to differentiate themselves from others—influence and are shaped by the formation, maintenance, and dissolution of institutionalized social boundaries between groups (Lamont & Molnár, 2002). Even when the social boundaries between groups are relatively stable, group members are constantly engaging in various forms of "boundary work" to preserve or shift symbolic distinctions (Berry et al., 1987; Gieryn, 1983; Lamont et al., 2015; Pachucki et al., 2007). Yet, in many contexts, the social boundaries between groups are breached—for example, when a new racial or ethnic group moves into a previously homogeneous neighborhood, when first-generation students first emerge on a college campus, or when employees from different organizations engage with one another following a merger or acquisition—thereby amplifying people's tendencies to erect and maintain symbolic boundaries. We consider three interrelated questions: (a) What strategies do people use to maintain symbolic boundaries when the social boundaries between groups are in flux? (b) How do these strategies vary across individuals and contexts? and (c) How can we systematically measure ongoing symbolic boundary maintenance given that boundaries are ever-shifting?

To date, research on the strength and persistence of symbolic boundaries has tended to rely on qualitative methods (e.g., semi-structured interviews and ethnography) or quantitative self-reports that provide rich portraits of the boundary itself but are less well-equipped to assess the fine-grained behavior through which individuals maintain symbolic boundaries and how these behaviors vary across individuals and social contexts, as well as over time. Drawing on the tools of computational linguistics and machine learning, in conjunction with increasingly widespread digital trace data, we develop a novel method for assessing how people use symbolic markers of distinction—specifically, styles of discourse that are strongly associated with a given group—to erect, reinforce, and shift symbolic boundaries.

Prior work on symbolic boundaries offers disconnected observations about how people maintain boundaries. One perspective highlights how individuals tend to retain existing symbolic markers of distinction. We label this perspective *boundary retention*. In this view, the utility of symbolic boundaries as tools of social reproduction of group membership and stratification derives from the cultural reproduction of the markers themselves (Bourdieu, 1984; Bourdieu & Passeron, 1970). An alternative tradition instead emphasizes how people constantly negotiate and reform the substance of symbolic boundaries. In this approach, which we refer to as *boundary reformation*, the potency of symbolic boundaries is grounded in the ongoing invention of new symbols (Accominotti et al., 2018; Gieryn, 1983).

We propose that these two forms of symbolic boundary maintenance—one focused on continuity and the other on change—are not mutually exclusive. Indeed, prior work has argued that individuals construct symbolic boundaries not from a singular cultural object but rather from a tapestry of objects that are in turn drawn from a broader cultural repertoire (Bail, 2008; Goldberg, 2011; Swidler, 1986). Building on this insight, we develop the novel theoretical idea that boundary maintenance can simultaneously encompass *boundary retention*—entrenchment in pre-existing symbolic differences—and *boundary reformation*—innovation of new symbolic distinctions.

To investigate this claim, we introduce a method for measuring the two facets of boundary maintenance and then explore how social positions and contexts relate to these processes. Our empirical approach circumvents key limitations in prior studies of boundary work, which have relied on self-reports or ethnographic accounts that are ill-suited to detecting subtle symbolic distinctions that can shift rapidly over time. Self-reports explore symbolic boundaries indirectly, by measuring intergroup affect or a limited set of cultural attitudes (Bail, 2008; Terry et al., 2001; Vila-Henninger, 2015). They are also subject to various forms of social desirability bias (Greenwald & Banaji, 1995; Srivastava & Banaji, 2011). Ethnographic approaches are superior for detecting observable cultural artifacts and behaviors but cannot reveal fine-grained variation in boundary work across individuals and over time (Geertz, 1973; Lamont, 1992).

In contrast, our approach to assessing how people maintain symbolic boundaries has three distinctive features. First, we focus on interactional language use within and between groups as a window into their distinctive "group styles," or the patterns of interaction that arise from shared assumptions and norms (Eliasoph & Lichterman, 2003; Gieryn, 1983). Interactional language—as reflected in our setting of email messages exchanged between employees—represents a source of rich cultural data as it constitutes a frequent, observable behavior that varies across individuals and groups, as well as over time. Building on recent work that has developed language-based measures of individuals' cultural fit in organizations (Goldberg, Srivastava, et al., 2016; Srivastava et al., 2018), we measure group styles using the Linguistic Inquiry and Word Count (LIWC) lexicon (Pennebaker et al., 2007), which reflects such communication tendencies as the use of abstract versus concrete language, the expression of positive versus negative sentiment, and orientations toward the past, present, or future. Whereas other natural language processing techniques such as topic modeling (Blei et al., 2003; DiMaggio et al., 2013) spotlight the specific issues being discussed in a group, LIWC allows us to examine linguistic styles independent of which issues are being discussed.

Second, to determine the characteristic linguistic styles of each social group, we use a binary classification-based approach (rather than continuous regression or clustering) so that we can explicitly model group styles based on the features that distinguish them. We then use this classifier to assess the social group membership of a given individual's emails based on the LIWC-based linguistic styles they contain, effectively projecting a person's emails onto a model of cultural distinction. We choose a random forest classifier for this task because it is the simplest, most intuitive model that does not assume a prototypical group style and because it flexibly allows for interactions between LIWC features. This choice was informed by prior work demonstrating significant variation within group styles (Bail, 2008; Bonikowski, 2016) as well as interviews in our empirical context confirming substantial within-group cultural heterogeneity.

Third, we segment our corpus of interactional language use into different time slices—before two groups first come into contact with one another and after they come together. This allows us to analytically distinguish between the two forms of boundary maintenance—one focused on the retention of historical group markers and the other on the reformation and innovation of these markers. In particular, we measure of boundary retention using a classifier trained on historical linguistic exchanges that occurred prior to intergroup contact, and we assess boundary reformation by combining the results of two classifiers trained on contemporaneous and historical linguistic exchanges, respectively.

The method we develop makes possible a more systematic investigation of the antecedents and consequences of symbolic boundary maintenance. In particular, we explore how individuals' tendency to engage in the two forms of boundary maintenance might be influenced by their social position and the structural context of their interactions. We first consider the role of power differentials. Differences in power can produce asymmetries in the strategies and resources available to groups for boundary maintenance (Lamont et al., 2015). Contrary to the prevailing view that only individuals from groups that wield structural power over other groups will erect and maintain boundaries, we argue that both high-power and low-power groups will engage in boundary work, with low-power groups more prone to exhibiting boundary reformation and high-power groups more apt to displaying boundary retention. Next we examine how variation in a person's local context—specifically, the proportion of ingroup versus outgroup members to which an individual is exposed—can amplify or dampen the tendency to engage in different forms of boundary maintenance through intergroup contact (Brannon & Walton, 2013; Coleman, 1988; Suzuki, 1997).

We apply our methodology and evaluate these ideas using a corpus of 1.5 million internal email communications from a mid-sized financial services firm in the U.S. that acquired and integrated two other financial service firms over the course of two years. As a supplement to archival data collection, structured interviews with employees and senior management informed our specific modeling choices—such as the use of a random forest classifier—and highlighted that the merger context constituted a potent breaching of group boundaries, making legacy organizational identities and symbolic boundaries especially salient to employees. We combine the email communication data with personnel records obtained from the three organizations to investigate variation in boundary maintenance behavior across individuals and contexts.

Several insights emerge from our analyses. First, despite intentions of organizational leaders to culturally integrate employees following the two mergers, we find evidence of both boundary retention and reformation persisting for up to 18 months following a merger. We also find that individuals from acquired firms exhibit greater boundary reformation and lower boundary retention than incumbent individuals, suggesting that power dynamics shape the resources with which people respond to the boundary threat of a merger and that cultural innovation can emerge from low-power groups that are constantly adapting cultural markers of distinction. Consistent with network theories of social closure (Coleman, 1988), our analyses also show that individuals engage in more boundary retention, but not reformation, when their work environments are more densely populated by ingroup members. Finally, since cultural behaviors are known to have important consequences for individual and organizational performance (Srivastava et al., 2018; Weber & Camerer, 2003), we conduct exploratory analyses about the link between boundary work and individual performance and find that retention (but not reformation) is negatively related to performance for individuals from acquired firms.

The contributions of this study are multifold: Methodologically, we develop a computational approach to measuring boundary retention and reformation that allows us to separately model cultural entrenchment and innovation in symbolic boundary maintenance. Theoretically, we are the first to propose and provide systematic evidence that individuals can simultaneously engage in these two forms of boundary maintenance depending on their social position and context. In addition, our methodological choices offer three general insights on measuring culture computationally using digital trace data. First, focusing on relevant and interpretable features, such as interactional language use as manifested through the LIWC lexicon, in reducing the dimensionality of "big data" can enable greater analytical precision in cultural measurement. Second, a subset of black box machine learning models like random forest classifiers provide a robust, non-parametric approach to empirically distinguish group styles while allowing for cultural heterogeneity within groups. Third, thoughtful and creative data segmentation in constructing training and test datasets can be a powerful tool for highlighting different measures of interest in computational research design.

We proceed as follows: first, we provide an overview of existing research on symbolic boundary work, highlighting the conceptual distinction between boundary retention and reformation. Next, we describe our computational, language-based approach to measuring these two forms of boundary work. We then introduce our empirical setting—a pair of organizational mergers—and present our empirical results. Finally, we discuss theoretical implications, as well as how this approach can be readily extended to the study of culture and boundaries across a wide range of social groups.

## Theory

### Symbolic Boundaries and Boundary Maintenance

A large body of sociological research on boundaries has studied how individuals and groups maintain social boundaries—the objectified ways through which people of different groups gain differential access to resources—through the construction and perpetuation of symbolic boundaries—conceptual distinctions, interpretive strategies, and cultural traditions that people use to differentiate themselves from others. Early work in this vein assumed that symbolic boundaries are, for the most part, consistent across social group members and stable over time (Durkheim et al., 2001 [1912]; Grillo, 2003). In this tradition, symbolic boundaries serve to maintain social boundaries across long timescales. Thus, the maintenance of symbolic boundaries results from enduring cultural markers. For example, Bourdieu's influential studies of cultural reproduction argued that it is precisely because cultural distinctions of taste are widely accepted and legitimized that they can serve as vehicles for the intergenerational reproduction of social class (Bourdieu, 1984; Bourdieu & Passeron, 1970). The reproduction of existing cultural symbols is especially important in the face of inevitable turnover in group membership. For example, religious groups, military branches, and college fraternities and sororities all have elaborate and enduring markers—initiation rituals, shared symbols, and specialized jargon—to demarcate and preserve group boundaries. The legitimacy and authenticity of these cultural traditions derives in large part from their durability (Berger & Luckmann, 1966; Carroll, 2015). Moreover, symbols endure because people tend to exhibit cultural inertia (Kiley & Vaisey, 2020; Vaisey & Lizardo, 2016). Bourdieu's notion of *habitus*, for example, is rooted in an understanding of cultural attitudes and behaviors as deeply ingrained and difficult to shift.

In contrast to this view of symbolic boundaries as mostly uniform and static, a separate stream of work has shown that symbolic boundaries vary considerably across individuals and that new cultural markers emerge over time. For instance, Lamont (1992) points out the diverse understandings of high status culture across national contexts, Bail (2008) illuminates multiple configurations of symbolic boundaries that define European nationalism and attitudes towards immigration, and Goldberg (2011) identifies three competing logics of cultural distinction in American musical tastes. In even earlier work, anthropologist Barth (1969) recognized that attempts to classify individuals' culture and ethnicity according to the same unchanging categories was futile because the relationship of cultural behaviors to ethnic groups and individuals' ethnic group membership were in constant flux.

Further support for the view that boundary maintenance involves symbolic innovation—that is, the ongoing creation of new cultural markers—comes from studies that show how boundaries evolve over time. For example, elite consumers of the New York Philharmonic responded to the opening of subscriptions to middle class audiences by erecting new symbolic distinctions of exclusivity through their patterns of attendance (Accominotti et al., 2018). The analysis of baby name trends by Lieberson and Bell (1992) points out that fashion cycles are a form of symbolic innovation that reinforce symbolic boundaries of social class.

In sum, the extant literature makes two implicit and disconnected assumptions about how people maintain symbolic boundaries—through reinforcing existing symbols and innovating new ones. Moving beyond conceptions of symbolic boundary maintenance as involving cultural markers that are either static or constantly evolving, we propose a theoretical integration of these two perspectives. Recent research lends support for this dual approach to boundary maintenance. For instance, Khan (2012) finds that the cultural capital inculcated in modern elite boarding schools blends longstanding highbrow education (e.g., "Beowulf") with more recent elite symbols of cosmopolitan and worldly consumption (e.g., "Jaws"). In a parallel vein, Goldberg, Hannan, et al. (2016) demonstrate how high status tastes in cuisine and film are characterized by both retention of classical genres and more recent boundary refinement according to genre purity. Finally, recent methodological advances in natural language processing have illuminated how the semantic structure underlying gender and class boundaries have simultaneously evolved while key symbolic features remain stable (Garg et al., 2018; Kozlowski et al., 2019).

We label the two forms symbolic boundary maintenance as *boundary retention*—the extent to which individuals sustain and entrench in pre-existing symbolic differences between social groups—and *boundary reformation*—the degree to which individuals alter the substance of symbolic boundaries by creating new cultural markers of distinction between social groups. Conceptualizing these two forms of boundary maintenance allows us to theorize about how different types of individuals will engage in each activity and how their behavior might vary across social contexts. Although we focus on the maintenance of symbolic boundaries between particular social groups that exist in organizations—the legacy organizations people belonged to before their organizations merged—our theory can be readily extended to other social groups both within and outside formal organizations.

### Boundary Retention and Reformation: The Roles of Social Position and Context

We first consider how an individual's social position—specifically the relative power of her group—might influence how she differentially employs the two boundary maintenance strategies. We expect that individuals from high-power groups—in our setting, employees from the acquiring firm—will be more invested in preserving the status quo and less inclined to adapt to a new or changing context. This could be because their group dominance is based on existing cultural markers or have previously invested considerable effort to enculturate to existing styles. We therefore anticipate that members of high-power groups will exhibit more boundary retention of existing symbolic markers. On the other hand, individuals from lower power groups—that is, employees from acquired firms—will tend to engage more in boundary reformation given that they face greater pressure (and control by dominant group members) to adapt to the pre-existing culture of the dominant group. Innovating new symbolic markers of distinction might provide lower-power group members with an alternative source of agency to respond to the identity threat that arises through exposure to dominant individuals.

Next we consider how an individual's social context influences the form of boundary maintenance she engages in. Individuals who interact more frequently with outgroup members may develop more towards them—in large part because intergroup contact promotes cultural openness and exchange (Brannon & Walton, 2013; Suzuki, 1997). In other words, individuals are more inclined to blur symbolic boundaries when they interact more frequently with members of their outgroup. The converse is that infrequent exposure to the outgroup might result in less cultural exchange and greater perpetuation of existing symbolic differences (Coleman, 1988). Comparing boundary retention and reformation, we therefore anticipate that individuals surrounded by a greater proportion ingroup, rather than outgroup, members have more exposure to existing symbolic markers of distinction. As a consequence, they will be more likely to reenact these pre-existing markers—that is, engage in boundary retention—than to culturally adapt and innovate new ones—that is, to engage in boundary reformation.

Finally, we conduct exploratory analyses related to the question of how the two forms of boundary maintenance might relate to an individual's performance in a post-merger organizational setting. Prior work has shown that individuals who do not adapt to the prevailing organizational culture are likely to suffer performance penalties (Srivastava et al., 2018). Building on this insight, we anticipate that boundary retention will result in more negative evaluations for individuals from lower status groups—employees from acquired firms—because retaining pre-existing cultural differences signals the failure to successfully enculturate to the dominant culture. In contrast, when higher status groups—employees from the acquiring firm—exhibit boundary retention, they are still conforming to the dominant culture and will therefore avoid receiving harsh evaluations. As for boundary reformation, it is theoretically unclear whether it will be a boost or detriment to an individual's performance. In some social contexts, symbolic innovation may be perceived positively, while in others, it may be construed as disruptive. Thus, we expect that the performance implications of boundary reformation are likely to be context specific.

We turn next to describing our language-based method for assessing how people maintain symbolic group boundaries and to demonstrating the utility of the approach in the context of organizational mergers.

## Method

### Language as a Measure of Boundary Maintenance

Studies of symbolic boundaries have drawn on various types of data, from in-depth interviews and ethnographic accounts to large-sample surveys (Bail, 2008; Lamont, 1992). Although these methods have yielded valuable insights about how people maintain and manage cultural distinctions, they also have important limitations in uncovering fine-grained variation across individuals and over time. Survey-based approaches, though designed to assess individual-level variation, capture symbolic boundaries at a single point in time and are subject to various forms of self-report bias (Greenwald & Banaji, 1995). Moreover, survey instruments are based on coarse-grained cultural categories that are defined by the researcher and that may not be relevant for some individuals and in particular contexts. Ethnographic approaches are well-suited to inductively detecting subtle symbolic distinctions; however, they are difficult to scale to large groups and therefore require extrapolating from a subset of observations over a focused period of time.

As an alternative, we propose that the natural language people use in communicating with members of their own group and other groups can provide a granular window into how they maintain and manage cultural distinctions over time. Indeed, scholars have long recognized the primary role of language in demarcating boundaries (Gieryn, 1983; Gumperz, 1977). Because language is laden with cultural meaning (Fishman, 2012), it is especially well-suited to inductively identifying symbolic distinctions. Moreover, language-based analyses do not require the researcher to impose a pre-defined set of cultural attributes and are thus a good match for our focal task of disentangling the use of novel versus pre-existing symbolic distinctions. Finally, the tools of computational linguistics and machine learning make it easier than ever to extract individuals' cultural signals from increasingly widespread digital trace data—for example, emails exchanged among employees (Goldberg, Srivastava, et al., 2016), messages posted to online platforms such as Slack (Lix et al., 2022), and worker ratings of their employers on platforms such as Glassdoor (Corritore et al., 2020).

To tease apart symbolic distinctions from the functional aspects of language, we specifically measure the linguistic style of communication as reflected in emotional, cognitive, and cultural categories, rather than context-dependent topics such as particular people or events that different social groups might reference. Our approach is related to and builds on the interactional language use model (Goldberg, Srivastava, et al., 2016; Srivastava et al., 2018), which uses the well-established Linguistic Inquiry and Word Count (LIWC) lexicon to measure linguistic style (Pennebaker et al., 2007). The LIWC lexicon consists of a dictionary of words corresponding to a different dimension of linguistic style, such as Achievement, Anger, or Assent (see Appendix A for a table of these categories). Prior studies have used the interactional language use model to derive a time-varying measure of individual-level cultural fit in an organization. In these studies, one's cultural fit increases with the degree of convergence between one's own linguistic style and that expressed by one's interaction partners at a given point in time. By abstracting away from specific words to higher-order categories, this technique allows one to assess cultural alignment independent of specific issues that are being discussed in a group and without making assumptions about the nature and contours of group culture. In contrast to most other language-based analyses such as topic models that measure *what* content is communicated, a LIWC-based approach probes *how* this content is communicated.

Building on this insight, here we use the LIWC lexicon to assess the extent to which a given communication matches the "group style" of a given social group at a particular point in time (Eliasoph & Lichterman, 2003). We can then aggregate these artifact-based measures at the person-level to derive time-varying measures of how individuals use cultural markers to maintain symbolic group boundaries. Before describing our language-based measures in detail, we begin with a brief description of the empirical setting in which we sought to demonstrate the utility of this approach.

## Empirical Setting

We focus on the context of organizational mergers and acquisitions (M&A). Most mergers and acquisitions involve some level of integration between the two organizations (Bodner & Capron, 2018).[1] Thus, most mergers result in a clear erosion of existing social boundaries between two groups. Moreover, many mergers are implemented with an eye to creating a unified combined culture (Nahavandi & Malekzadeh, 1988; Van den Steen, 2010; Weber &

[1]There are, of course, instances in which one organization may acquire another for purely financial reasons and not seek to assimilate its employees—these are less relevant for the subject at hand.

Camerer, 2003). In some cases, this uniform culture is mostly reflective of the one that prevailed in the acquiring organization, while in other instances the emergent culture is a blend of the pre-existing cultures or a novel recombination of them. Regardless, the erosion of social boundaries between organizations in a post-merger context requires people to engage in various forms of boundary work as they make sense of the new social order and seek to find their place within it (Amiot et al., 2012; Drori et al., 2013).

Our research setting is a mid-sized regional bank in the U.S. ($n_1 = 306$) that acquired and integrated two other regional banks ($n_2 = 247$, $n_3 = 51$) over the course of two years. The second acquisition took place exactly 12 months after the first. Conversations with both employees and senior management, conducted nine months after the second merger and prior to archival data analysis, revealed several key features of the setting that informed our analytical approach.

First, both mergers constituted a sudden blurring of the institutionalized boundaries between organizations. Neither merger event was known to or anticipated by rank-and-file employees until announced. In fact, the vast majority of employees did not experience significant change or threat to their routines or communication patterns until the week the merger went into effect, with all architectural integration occurring on what one of our interviewees referred to as a single "flipping weekend... when all the signage and systems are flipped." This quasi-exogenous discontinuity to individuals' experience inspired key decisions around data segmentation and variable construction discussed below in our measurement approach.

Second, employees became acutely aware of cultural differences following the mergers, making this empirical setting especially appropriate for the study of symbolic boundary work. Given that the rationale for the mergers focused on geographic complementarities between the branch networks of the different banks rather than on their anticipated cultural alignment, interview accounts suggested that employees experienced varying degrees of "culture clash." For example, while one employee perceived an "easy transition as far as the culture... because it was the same," another felt that "the merger was so hard because you thought you were getting a bank with the same culture [but] it was the opposite."

Third, because there was both employee hiring and attrition after each merger, our sample selection for analyses includes those individuals who were present before and after a given merger. Interviews mitigated concerns of any systematic selection of these individuals based on their cultural behavior. Inspecting the archival data confirms that employees who departed at the point of merger or in the following month did so mostly because they were in job roles that were deemed redundant rather than for reasons of cultural fit (i.e., they did not exhibit meaningfully different interactional language).

Finally, significant cultural variation was apparent both before and after the mergers. Because the bank branches were geographically dispersed and staff relied on local knowledge and customer relationships, management acknowledged and allowed for some level of cultural variation within the firm. Variation also existed at the department level, since different occupations and specializations lent themselves to different norms. For example, it was widely acknowledged among employees that aggression and competitiveness were the norm in commercial banking, while these tendencies were frowned upon in retail banking. This cultural variation has two important implications for our analysis: first, we exploit precisely this variation to examine differences between individuals in their symbolic boundary maintenance. Second, it necessitates that we model between-group differences in a manner that accommodates this within-organization heterogeneity. We elaborate on this point in discussing our analytical strategy.

Beyond access to a subset of employees whom we interviewed, the firm provided us with the complete corpus of internal email communication data (totaling roughly 1.5 million messages), including both metadata and hashed content, spanning 23 months—from three months before the first merger through eight months after the second merger. In addition to communication data, we had access to pre- and post-merger personnel records for both the acquiring firm and the two target firms. For ethical and security purposes, all data were deidentified by replacing names, addresses, and other identifying information with codes. To safeguard the content of email messages, we extracted information about the linguistic styles used in the text of emails, which were represented as numerical vectors, and then hashed the message content such that the original text was unrecoverable. Using the LIWC lexicon, we transformed the text of each email into a vector of word counts across 58 stylistic categories (see Appendix A).[2] These LIWC counts form the basis for our cultural measures of interest.

### Analytical Strategy

Our approach to measuring symbolic boundary maintenance is based on the extent to which a given employee's linguistic style is aligned with that of her legacy organization rather than that of the new organization to which she is exposed following a merger. One way to identify communication as stylistically representative of one organization versus the other is to infer what constitutes the general linguistic style of each one—for instance, based on how the "typical" employee uses specific LIWC categories. Yet this approach has a critical drawback: it imposes a monolithic or prototypical group style that prevails in the organization. As discussed earlier, there is ample evidence that cultural boundaries are instead characterized by significant heterogeneity (Bail, 2008). Indeed, as suggested in our interviews and illustrated in Appendix B, there is more within-bank heterogeneity in linguistic style than there is between-bank variation. Traditional regression and clustering models suffer from this drawback because they require a mathematical continuous mapping between the multidimensional space of LIWC features and group style. A second challenge involves accounting for the complexity of language—specifically, the many interactions between LIWC features that characterize group style. For instance, the symbolic implications of language employing future tense might vary based on the presence of positive versus negative emotion. Moreover, these interactions are largely unknown *ex ante*, handicapping any parametric modeling approach. These two issues point to the need for a model of group style that is neither

[2]We omitted superordinate categories in which these lower-level categories are nested in order to avoid collinearity between linguistic features.

parametric nor a continuous mapping.

To accommodate these requirements, we draw on the idea that the quantified output of a flexible machine learning classifier can itself be a measure of divergence between group styles (cf. Gentzkow et al., 2019). We distinguish between group styles by training a classifier to predict the legacy organizational affiliation of the sender of a focal email message based on the distribution of LIWC counts in the message. To obtain a precise, fine-grained measure of boundary work, we choose to model linguistic style at the level of analysis of a single email communication, rather than at the level of a multi-email thread (which involves multiple senders) or bundling a single sender's emails (which might distort our LIWC counts by averaging across distinct styles).

Among classifiers, random forest models are appropriately suited to the complexity of this task: being completely nonparametric and allowing for discontinuities, they are built to deal with multiple heterogeneous patterns within each classification, and they permit nonlinear and interdependent relationships between linguistic features.[3] An ancillary benefit of a random forest model is that it eliminates the need for normalizing LIWC counts of emails by message length. While the relationship between the inputted features (in this case, LIWC counts) and the classification of bank origin is not easily interpretable, the output is a model that inductively maps patterns and interactions of linguistic features to each organization. This mapping is sufficient for our purposes as we are not concerned with the content of each organization's linguistic signature.

### Measurement of Symbolic Boundary Maintenance

To derive our main measures of symbolic boundary maintenance—boundary retention and reformation—we took advantage of organizational merger events as quasi-exogenous breachings of social boundaries and made use of our longitudinal email corpus by segmenting our

[3]Several other black box machine learning models, such as artificial neural networks, also satisfy our criteria. We choose a random forest approach because it is especially well-established and simpler than alternatives; moreover, it aggregates over decision trees which are both interpretable and intuitive.

classifier training data into pre- and post-merger periods. We trained two sets of classification models: one of historical differentiation and another of contemporaneous differentiation (illustrated in Figure 1). Our model of historical differentiation was constructed from a single random forest classifier trained on messages sent in the three months prior to each merger. Intuitively, this model corresponds to the symbolic boundaries that existed when the focal merger was impending but prior to actual integration and significant intergroup contact. In contrast, our model of contemporaneous differentiation was based on multiple monthly classifiers trained on messages sent in the same post-merger month as the test data. Each of the trained models represents the linguistic markers that characterized symbolic boundaries in a given month. We applied these models to paired test datasets from contemporaneous months to measure the extent to which messages could be classified according to contemporaneous symbolic distinctions.

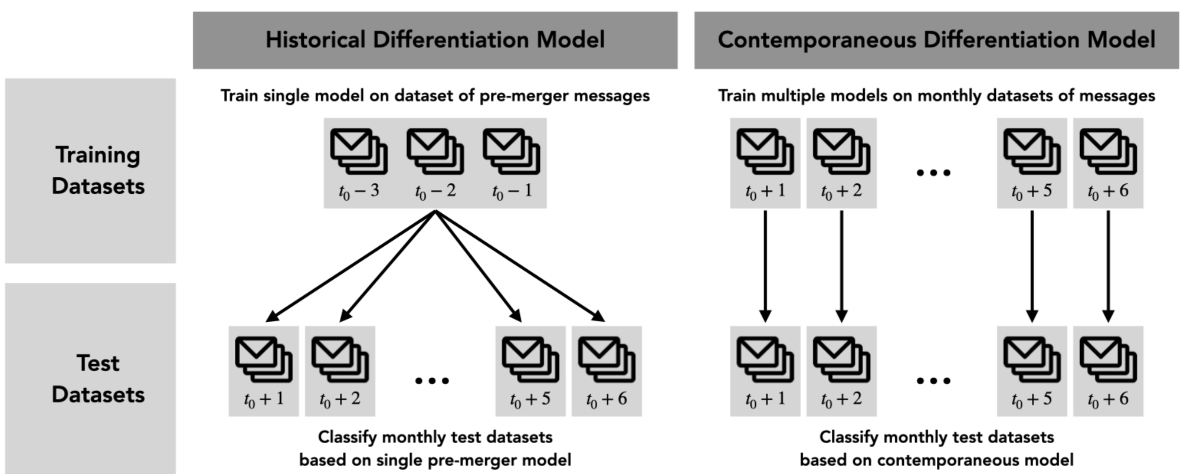


**Figure 1:** Illustration of the Modeling Approach for Historical and Contemporaneous Differentiation

Our measure of boundary retention follows straightforwardly from the historical differentiation model. Theoretically, boundary retention reflects the extent to which pre-existing symbols of distinction are perpetuated. We obtained this empirically by applying the historical differentiation model to test datasets from subsequent months to measure the extent to which post-merger messages reflected the maintenance of pre-merger distinctions. Although the performance of the historical classifier trained on pre-merger data is inevitably worse when applied to post-merger time periods, the objective is to examine the extent to which the historical model retains explanatory power for a given individual in future time periods. We calculated boundary retention as the proportion of an individual employee's ($i$) outgoing messages ($\vec{m}_{i,n}$) in the six months post-merger (in the test dataset) that were classified correctly according to the historical differentiation (HD) model—that is, as more aligned with the pre-merger linguistic style of the focal employee's legacy organization ($org_i$) than of the new organization to which the person was exposed (see Equation 1).

$$BoundaryRetention_i = \frac{1}{N_i} \sum_{n=1}^{N_i} \mathbb{1}(\hat{org}_{HD}(\vec{m}_{i,n}) = org_i)) \tag{1}$$

To gain empirical traction on the concept of boundary reformation, we combined the historical and contemporaneous differentiation models.[4] Boundary reformation involves the innovation of new symbolic distinctions along the same social boundary. One way to identify new symbolic distinctions is to consider post-merger differentiation that does not reflect historical pre-merger distinctions. Correspondingly, we categorized post-merger messages as consistent with symbolic innovation along the boundary if they were classified correctly according to the contemporaneous differentiation (CD) model but not according to the historical differentiation (HD) model. Our measure of boundary reformation is thus the proportion of an individual employee's ($i$) outgoing messages ($\vec{m}_{i,n}$) in the six months post-merger (in the test dataset) that were classified correctly according to the CD model *and* incorrectly

[4]A seemingly simple measure for boundary reformation would parallel that of boundary retention by drawing on the contemporaneous differentiation (CD) model alone. However, this model captures *both* novel and historical symbolic distinctions. Therefore, in order to analytically isolate symbolic innovation, we combined the two models as described. Ideally, we might construct a single multi-class classifier for both forms of boundary maintenance that classifies messages according to a 2x2 of enacting novel and historical symbolic distinctions; however, because we do not have labeled data for training such a model, we developed an alternative approach using data segmentation and multiple classifiers.

according to the HD model (see Equation 2).

$$BoundaryReformation_i = \frac{1}{N_i}\sum_{n=1}^{N_i} \mathbb{1}(\hat{org}_{CD}(\vec{m}_{i,n}) = org_i)) * \mathbb{1}(\hat{org}_{HD}(\vec{m}_{i,n}) \neq org_i)) \quad (2)$$

### Data Sampling and Model Training

We undertook several critical steps to construct the random forest models of historical and contemporaneous differentiation, on which the measures of boundary retention and reformation are based. First, for each of the two mergers, we constructed separate panel datasets of messages sent both before and after the merger. With these two panel datasets, we filtered the communication data along a number of dimensions. We excluded messages from individuals who did not experience the focal merger (i.e., who departed prior or joined following) or who sent fewer than 20 total messages in the six months following the focal merger. We also filtered out idiosyncratic messages with over 500 terms. For the second merger, we filtered out messages sent by individuals who had been acquired in the first merger, since their social group identity was already in flux and potentially ambiguous.

Second, we carefully constructed multiple training and test datasets according to the needs of each theoretical construct. We divided the two panel datasets of messages (one for each merger) into separate monthly datasets and split each dataset of monthly messages into equal-sized training and test sets.[5] The monthly training datasets were downsampled to a maximum of 50 messages per individual so that no single individual's linguistic style dominated the model. We made this choice to avoid overfitting to a handful of prolific individuals. Using training data from the appropriate months (as outlined above and illustrated in Figure 1), we trained multiple random forest classifiers. Models were trained in R using the *caret* and *ranger* packages, with 5-repeated 10-fold cross-validation to tune parameters and conventional downsampling to balance classes. Class balancing helped ensure that, even

[5] The conventional approach in machine learning studies is to use 80 percent of the data for training purposes so as to optimize for model accuracy; however, we choose to include a larger fraction of the data in our test datasets to ensure the reliability of our final measures, which were constructed from the test sets.

in the absence of linguistic information, messages were equally likely to be classified as belonging to either of the two organizations—even if employees in one organization tended to send more messages than their counterparts in the other. See Figure 2 for an illustration of the procedure for sampling training data.

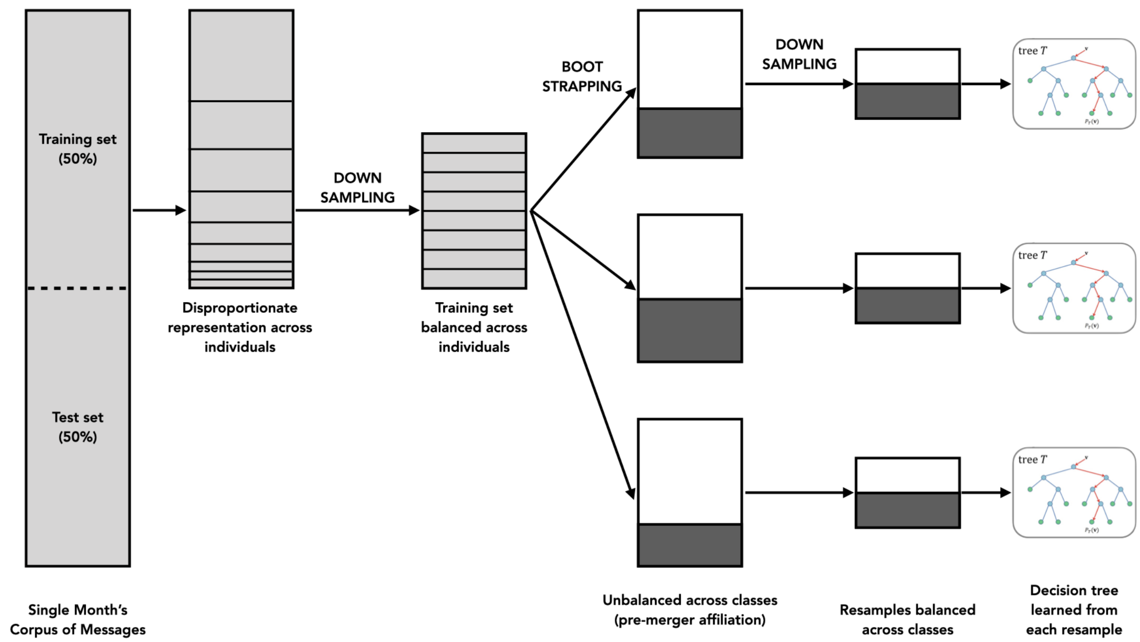


**Figure 2:** Procedure for Sampling Training Data for Each Month

Finally, to validate our models of historical and contemporaneous differentiation, we calculated an aggregate month-level indicator of model fit based on a well-established machine learning metric. Specifically, for both the historical and contemporaneous models, we measured the area under the receiver operating characteristic curve (AUC) for each monthly test dataset (including messages from all individuals). The monthly AUC is an explanatory measure (akin to an $R^2$ metric), representing how well a model can distinguish organizational affiliation of messages in a given month according to their linguistic style. We interpret this ability to distinguish as an aggregate measure of the strength of the historical or contemporaneous symbolic boundary between the merged organizations. Note that this aggregate measure of boundary strength is not the main dependent variable in our analyses; it is separate from and in addition to the individual-level measures of boundary retention and reformation described earlier.

Figure 3 summarizes the entire data processing pipeline.

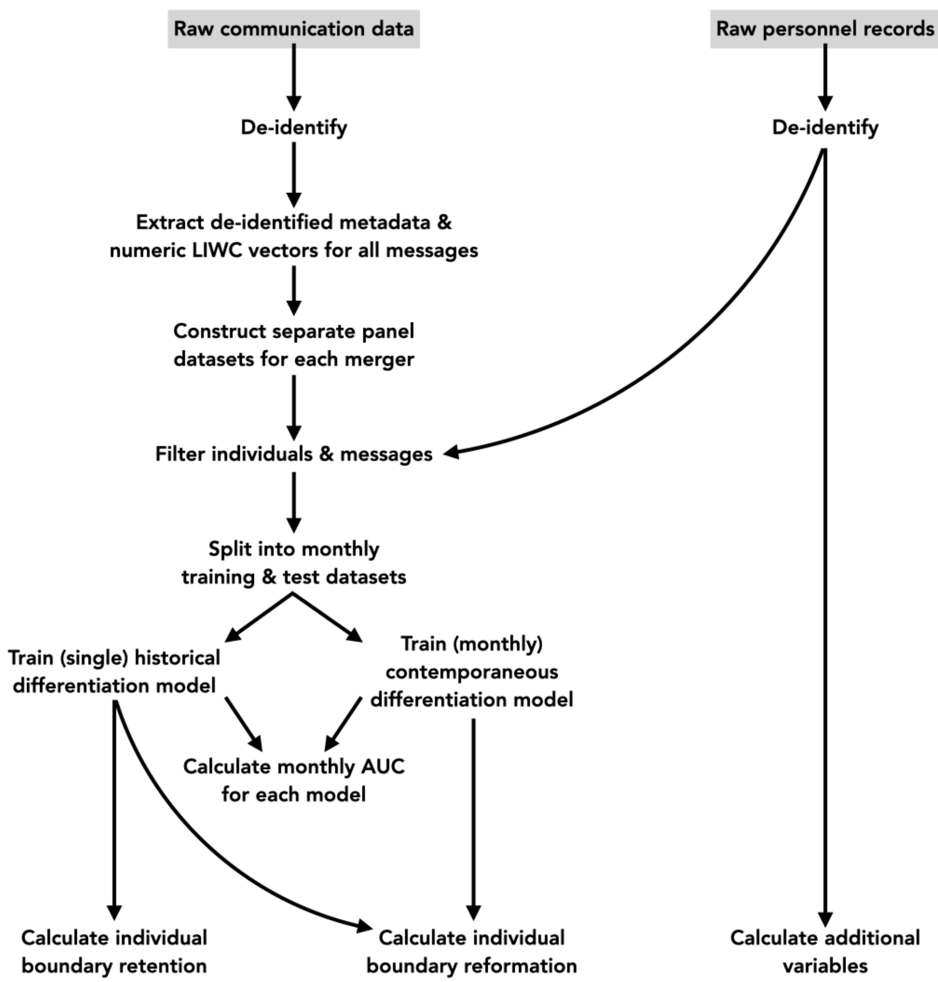


**Figure 3:** Flowchart of the Entire Data Processing Pipeline

## Additional Variables

In addition to these cultural measures, we assembled a number of additional variables for the regression analyses described below. First, because the nature of individuals' boundary work typically differs based on the relative power of their social group, we defined an indicator, *acquired*, which was set to 1 for employees from the acquired (i.e., less dominant) firm and to 0 for employees from the acquiring (i.e., dominant) firm. Given that power also derives from the structural positions one occupies within an organization, we controlled for individuals' hierarchical rank based on their pre-merger job titles. Specifically, we identified whether individuals were senior leaders, middle managers, or individual contributors. *Hierarchical position* is coded ordinally, with 2 indicating senior leaders, 1 middle managers, and 0 individual contributors. Because employees' cultural behavior typically evolves over their tenure in an organization, we also included employee *tenure* as measured in years at the date of the merger (Srivastava et al., 2018). In our models of individual performance, we also included gender (to account for gender bias in evaluations) using an indicator set to 1 for *male* employees.

To understand how an individual's local context shaped their tendency to engage in boundary maintenance through retention versus reformation, we developed a measure of *local ingroup density*, which is defined as the proportion of employees in the same bank branch as the focal individual who share the same pre-merger affiliation (averaged over the six months post-merger). Finally, for exploratory analyses of how boundary maintenance might relate to perceived performance, we used individuals' *post-merger performance ratings*, which were assessed on a 1 to 5 scale in the six months following each merger.[6]

Summary statistics and pairwise correlations of measures are reported in Appendix C.

### Estimation

We estimated ordinary least squares regression models to explain variation in the extent to which individuals engage in boundary reformation and retention. Our dependent variables were boundary reformation and retention, and we included as regressors individuals' local ingroup density, tenure at date of merger, and an indicator of acquired status. We also included controls for individuals' hierarchical position and the interaction between acquired

[6]In supplemental analyses to the ones reported below, we also estimated models with additional control variables, including outgoing message volume, average message length, supervisor status, and department and merger fixed effects. Because the results below were substantively unchanged with the inclusion of these controls, we do not report specifications that include these variables in our main results.

status and local ingroup density.[7]

In our exploratory analyses of the association between boundary maintenance and individuals' post-merger performance ratings, we also estimated ordinary least squares regression models that included boundary reformation and retention as regressors, as well as controls for acquired status, hierarchical position, gender, and tenure. We also included interaction terms between acquired status and the two measures of boundary maintenance.

## Results

Before turning to our main analyses, we first validated our methodology by investigating the extent to which traces of symbolic boundary maintenance can be detected several months after each merger was consummated. To do so, we plotted in Figure 4 the monthly AUC for the historical and contemporaneous differentiation models as applied to our test data for both mergers. As Figure 4 reveals, despite efforts by leaders of the focal firm to culturally integrate employees from the two firms they acquired, both historical and contemporaneous differentiation can be detected in email communication for several months following organizational integration. Unsurprisingly, the contemporaneous differentiation model performs better than the historical differentiation one, since the latter is applied to datasets from a different time period than the training data. Yet the historical differentiation model can still distinguish organizational affiliation several months in the future. In other words, the symbolic boundaries between group members appear to be remarkably resilient and enduring. For up to 18 months following a merger that breaches the social boundaries between organizations, our random forest classifier is still able to detect the subtle cultural distinctions that employees make based just on linguistic styles of communication. Since Figure 4 depicts similar post-merger cultural behavior in both mergers, we combine the datasets for each merger in subsequent analyses.

[7]Because reformation and retention were measured as fractional variables (i.e., the proportion of an individual's messages consistent with specific types of boundary work), we also estimated fractional logistic regressions that yielded comparable results.

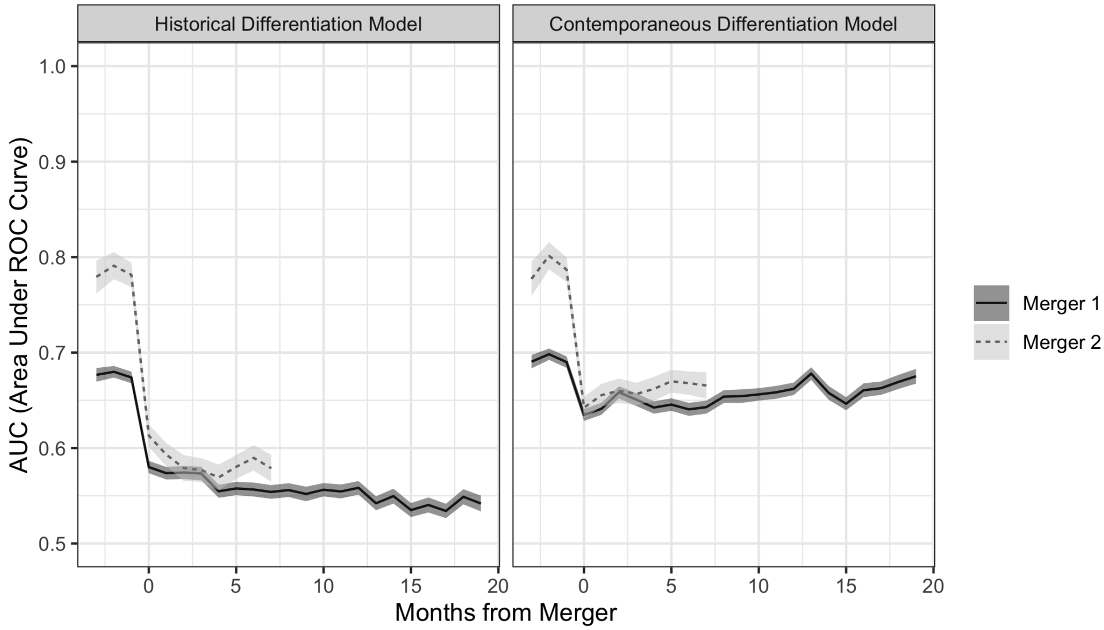


**Figure 4:** Plots of area under the ROC curve (AUC) over time for random forest classification of pre-merger organizational affiliation based on historical (pre-merger training data) and contemporaneous (post-merger training data) differentiation models, for both mergers. Confidence intervals are obtained by bootstrapping test datasets. An AUC above 0.5 corresponds to better-than-random distinguishability of classes; maximal AUC of 1 corresponds to perfect distinguishability. The AUC value is an explanatory model measure and can be interpreted as the strength of the symbolic boundary.

We turn next to considering how a person's social position and their local context influence the extent to which they maintain symbolic boundaries through boundary retention versus reformation. Table 1 reports our main models of interest. Models 1 through 3 characterize boundary retention as the dependent variable of interest, while models 4 through 6 similarly characterize boundary reformation. All models include coefficients for whether individuals are acquired. Models 1 and 4 also include individuals' tenure at the date of merger (continuous). Models 2 and 5 include individuals' local ingroup density, and models 3 and 6 include both of these regressors, as well as hierarchical rank and an interaction between acquired status and local ingroup density.

**Table 1:** OLS Regressions of Post-merger Boundary Maintenance on Employee Characteristics

| | *Dependent variable:* | | | | | |
|---|---|---|---|---|---|---|
| | Boundary Retention | | | Boundary Reformation | | |
| | (1) | (2) | (3) | (4) | (5) | (6) |
| Acquired | $-0.182^{***}$ | $-0.164^{***}$ | $-0.172^{***}$ | $0.144^{***}$ | $0.139^{***}$ | $0.121^{***}$ |
| | (0.009) | (0.009) | (0.029) | (0.006) | (0.006) | (0.020) |
| Tenure | $0.001^{**}$ | | $0.001^{**}$ | $-0.001^{***}$ | | $-0.001^{**}$ |
| | (0.001) | | (0.001) | (0.0004) | | (0.0004) |
| Local Ingroup Density | | $0.056^{***}$ | $0.053^{*}$ | | 0.009 | $-0.008$ |
| | | (0.017) | (0.026) | | (0.011) | (0.018) |
| Acquired x Local Ingroup Density | | | $-0.0001$ | | | 0.030 |
| | | | (0.034) | | | (0.023) |
| Hierarchical Position | | | $-0.006$ | | | $-0.003$ |
| | | | (0.005) | | | (0.003) |
| Constant | $0.632^{***}$ | $0.588^{***}$ | $0.587^{***}$ | $0.175^{***}$ | $0.161^{***}$ | $0.183^{***}$ |
| | (0.005) | (0.015) | (0.024) | (0.003) | (0.011) | (0.016) |
| Observations | 899 | 900 | 898 | 899 | 900 | 898 |
| $R^2$ | 0.337 | 0.340 | 0.345 | 0.411 | 0.403 | 0.413 |
| Adjusted $R^2$ | 0.335 | 0.338 | 0.342 | 0.410 | 0.402 | 0.410 |
| Residual Std. Error | 0.111 (df = 896) | 0.111 (df = 897) | 0.110 (df = 892) | 0.075 (df = 896) | 0.076 (df = 897) | 0.075 (df = 892) |
| F Statistic | $227.407^{***}$ (df = 2; 896) | $230.817^{***}$ (df = 2; 897) | $94.065^{***}$ (df = 5; 892) | $312.639^{***}$ (df = 2; 896) | $302.712^{***}$ (df = 2; 897) | $125.465^{***}$ (df = 5; 892) |

*Note:* + p<0.1; * p<0.05; ** p<0.01; *** p<0.001

Across all specifications, we find that acquired individuals exhibit less boundary retention and greater boundary reformation and than their counterparts in the acquiring organization. To give a sense of the magnitude of these associations, for classifications based solely on pre-merger linguistic distinctions, messages from acquired individuals are 17 percent less likely to be correctly classified relative to those from incumbents. Conversely, acquired individuals' messages are 12 percent more likely to be *both* correctly classified according to contemporaneous (post-merger) and incorrectly classified according to historical (pre-merger) differentiation models relative to messages from incumbents.

These results are consistent with the notion that acquired individuals are forced to adapt to a new organizational context, thereby limiting their ability to consecrate pre-existing cultural symbols. They instead respond to the identity threat of being newcomers to an established and dominant organizational culture by demarcating symbolic boundaries through the innovation of new cultural markers. Conversely, incumbent employees retain and emphasize their pre-existing cultural symbols as a means of preserving the status quo and their dominance within it.

Table 1 also shows depicts the relationship between the two forms of boundary maintenance and tenure. Consistent with explanations of age-based inertia, we find that individuals with greater tenure exhibit greater boundary retention and less boundary reformation (Le Mens et al., 2015). In other words, those who were more immersed in or previously invested considerable effort to learn the cultural symbols of their organization were also more likely to reinforce these existing symbols (and less likely to innovate new symbols) relative to those who more recently joined the organization. This pattern holds even when controlling for hierarchical position.

Regarding social context, we find evidence consistent with intergroup contact theory: individuals demonstrate greater levels of perpetuating existing symbolic distinctions—boundary retention—when surrounded by a greater proportion of ingroup members. Note that while the boundary work is displayed in email communication, contact with ingroup and outgroup members is measured based on physical co-location. Consequently, individuals with high local ingroup density may still be exposed to outgroup linguistic styles via email communication, but they are more likely to retain or innovate linguistic styles associated with their ingroup due to the more proximate social influence of in-person interaction.

However, greater ingroup density is not uniformly associated with an increase in all boundary maintenance behaviors. In line with our expectations that primary consequence of ingroup exposure is exposure to existing cultural practices, rather than motivation to maintain intergroup prejudice, we find that this association is statistically significant for boundary retention but not for reformation—a 10 percent increase in ingroup density corresponds to a 0.5 percent increase in boundary retention, but no meaningful association with boundary reformation. Figure 5 depicts marginal effects plots of the association between local ingroup density and boundary maintenance behaviors (based on Models 3 and 6 in Table 1).

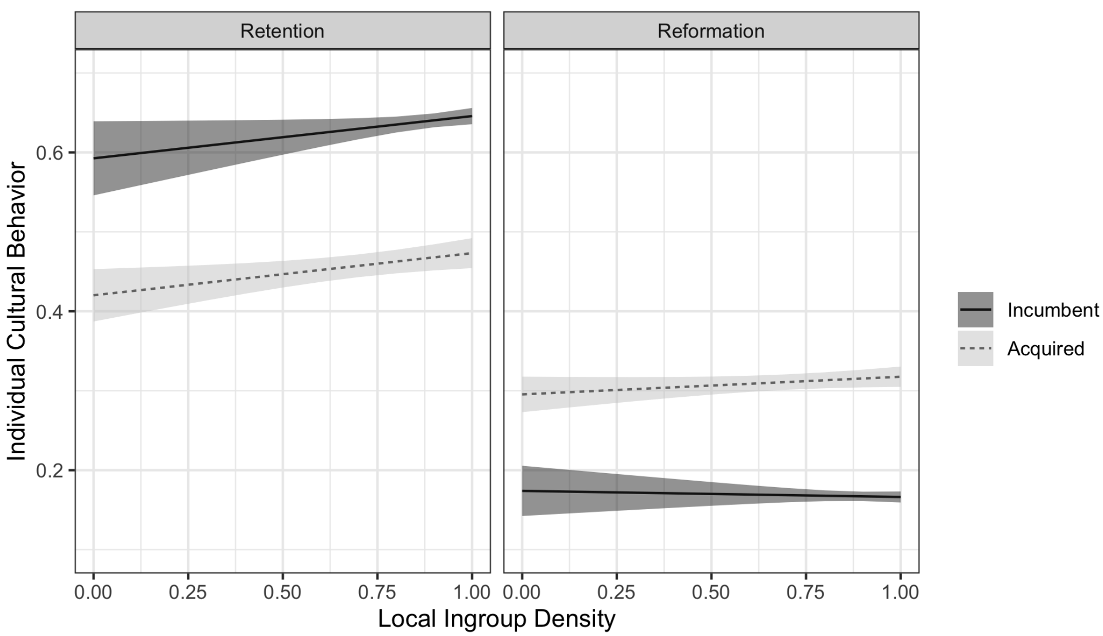


**Figure 5:** Marginal Effect Plots of Boundary Maintenance Behaviors based on Table 1, Models 3 & 6

Finally, Table 2 reports exploratory analyses regarding the relationship between boundary maintenance and individual performance. Model 1 specifies a baseline model with acquired status, hierarchical position, male, and tenure as regressors. Models 2 through 6 progressively add in boundary retention and reformation, as well as interaction terms between acquired status and the cultural measures. Figure 6 depicts marginal effects plots of the association between boundary maintenance behaviors and performance ratings (based on Model 6 in Table 2).

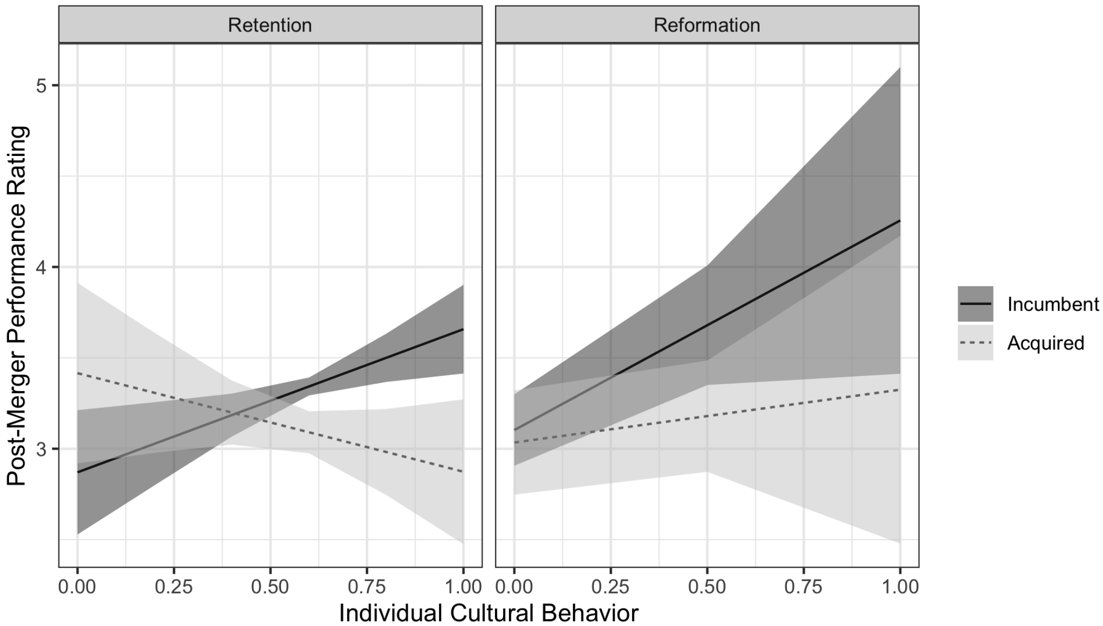


**Figure 6:** Marginal Effect Plots of Performance Ratings based on Model 6 in Table 2

## **Table 2:** OLS Regressions of Performance Rating on Post-merger Acculturation

| | *Dependent variable:* | | | | | |
|---|---|---|---|---|---|---|
| | Post-merger Performance Rating | | | | | |
| | (1) | (2) | (3) | (4) | (5) | (6) |
| Acquired | −0.120** | −0.126* | 0.377+ | −0.185** | −0.374* | 0.722 |
| | (0.043) | (0.051) | (0.194) | (0.057) | (0.146) | (0.446) |
| Retention | | 0.035 | 0.286 | | | 0.788** |
| | | (0.161) | (0.186) | | | (0.295) |
| Acquired x Retention | | | −0.992** | | | −1.330* |
| | | | (0.369) | | | (0.536) |
| Reformation | | | | 0.354 | 0.052 | 1.154* |
| | | | | (0.256) | (0.333) | (0.527) |
| Acquired x Reformation | | | | | 0.750 | −0.863 |
| | | | | | (0.532) | (0.788) |
| Hierarchical Position | 0.131*** | 0.114*** | 0.116*** | 0.114*** | 0.119*** | 0.111*** |
| | (0.024) | (0.024) | (0.024) | (0.024) | (0.025) | (0.025) |
| Male | 0.214*** | 0.220*** | 0.214*** | 0.221*** | 0.223*** | 0.206*** |
| | (0.039) | (0.041) | (0.041) | (0.041) | (0.041) | (0.041) |
| Tenure | 0.007** | 0.007* | 0.007** | 0.007** | 0.007** | 0.007** |
| | (0.003) | (0.003) | (0.003) | (0.003) | (0.003) | (0.003) |
| Constant | 3.104*** | 3.113*** | 2.948*** | 3.075*** | 3.118*** | 2.444*** |
| | (0.029) | (0.107) | (0.123) | (0.053) | (0.061) | (0.260) |
| Observations | 602 | 529 | 529 | 529 | 529 | 529 |
| $R^2$ | 0.153 | 0.157 | 0.169 | 0.160 | 0.163 | 0.177 |
| Adjusted $R^2$ | 0.147 | 0.149 | 0.159 | 0.152 | 0.154 | 0.164 |
| Residual Std. Error | 0.420 (df = 597) | 0.407 (df = 523) | 0.405 (df = 522) | 0.407 (df = 523) | 0.406 (df = 522) | 0.404 (df = 520) |
| F Statistic | 26.903*** (df = 4; 597) | 19.504*** (df = 5; 523) | 17.652*** (df = 6; 522) | 19.948*** (df = 5; 523) | 16.986*** (df = 6; 522) | 13.963*** (df = 8; 520) |
| *Note:* | + p<0.1; * p<0.05; ** p<0.01; *** p<0.001 | | | | | |

In line with expectations, we find that, for acquired employees, boundary maintenance through retention—that is, retaining legacy cultural distinctions that are not aligned with the new, dominant organization—is negatively related to performance: on average, a 50 percent increase in boundary retention corresponds to a 0.27 point decline in performance on a 5 point scale—similar in magnitude to the performance penalty of being female. In contrast, boundary maintenance through retention does not appear to matter for individuals from the acquiring organization—likely because their attempts at cultural retention are already aligned with the dominant culture. Boundary reformation does not appear to have significant or consistent performance implications, perhaps because cultural innovations can have idiosyncratic outcomes for individual performance based on the specific substance of the innovation.

## Discussion

The goal of this paper has been to understand how individuals maintain symbolic boundaries by distinguishing between two forms of the behavior—boundary retention, or entrenchment in existing symbolic markers, and boundary reformation, or innovation of new symbolic markers. In contrast to prior work that implicitly assumes symbolic boundaries are either static or evolving, we instead began with the premise from toolkit theory that the content of symbolic boundaries can both exhibit stability and reflect innovation and change (Swidler, 1986). We drew on the tools of machine learning and natural language processing to provide a systematic method for assessing the extent to which individuals maintain symbolic boundaries through retention versus reformation. We then demonstrated the utility of this approach by applying it to a granular dataset of internal email communications that encompasses two organizational mergers.

Our analyses suggest that individuals continue exhibiting both boundary retention and reformation many months after the social boundaries between two organizations have been breached through a merger. Such a pattern would be difficult to detect using traditional methods such as interviews and surveys given that post-merger contexts are politically sensitive, which increases the vulnerability of self-reports to social desirability bias and Hawthorne effects. Moreover, it is often impractical to collect survey or interview data repeatedly over such a long period of time.

Underlying the aggregate pattern of persistent boundary strength between the acquired and acquiring organizations is considerable variation in how individuals maintain symbolic boundaries based on their group dominance and local context (Bourdieu, 1984; Wimmer, 2008). Consistent with certain mechanisms of intergroup contact theory, boundary retention is positively associated with ingroup interaction—that is, the more local exposure a person has to ingroup members and their pre-existing cultural practices, the more she exhibits retention (but not necessarily reformation). Retention is also more frequently employed by organizational incumbents and those with greater tenure. Conversely, acquired individuals, as well as less tenured people, are more likely to maintain symbolic boundaries through reformation. While retention appears to have negative performance implications for acquired individuals, incumbents experience no such penalty for either retention or reformation.

Notwithstanding questions of generalizability, we believe this study makes two substantive contributions to research on symbolic boundaries. First, it provides a theoretical integration of two perspectives on symbolic boundaries—one that emphasizes cultural entrenchment and the other that focuses on cultural innovation. We demonstrate that people simultaneously engage in both forms of boundary maintenance and identify how their tendency to engage in each form varies as a function of their social position and their social context. Second, it makes a key methodological contribution by developing a novel approach to measuring boundary maintenance. This method can be readily extended to measure boundaries in a number of other contexts that include longitudinal digital trace data of the kind we analyze. Examples of other settings in which the social boundaries between groups are breached include the assimilation of newly arrived immigrants to a new country, the intermingling of social groups on various digital communication platforms, and organizational restructuring that shifts the contours of organizational subunits. In many of these settings, digital trace data of within- and between-group communication are readily available and could be analyzed using an approach very similar to ours.

Focusing on the context of organizational mergers allows us to address a fundamental challenge in studies of boundary work: social boundaries ("who" is demarcated) and symbolic boundaries ( "what" symbols enact the demarcation) tend to co-evolve. The set of individuals classified as native or foreign or as high class or low class frequently changes alongside the set of symbols that demarcate these social categories (Barth, 1969). These simultaneous shifts can make it difficult to interpret individuals' symbolic boundary work given the ambiguity about which side of the social boundary they occupy. By focusing on organizational mergers, we circumvent this problem given that social groups are well-defined based on individuals' pre-merger organizational affiliations. As a result, our analyses can hold social boundaries constant and focus on variation in the symbolic boundary's content.

More broadly, this paper offers a number of general insights for sociologists measuring culture using computational methods and digital trace data (Mohr et al., 2020). With unstructured data such as language, dimensionality reduction is a key step for a computational sociologist. Analytical precision, robustness, and replicability rely on the relevance and conceptual interpretability of these reduced dimensions for theoretical questions of interest. Here we focused on interactional language use as manifested in the LIWC lexicon as a means to isolating group style from the substantives issues discussed among and between group members.

The appeal of combining newly available unstructured data with a wide array of ever-evolving machine learning methods can sometimes obscure the potential for important methodological innovation in data pre-processing. A key step in constructing our measures of boundary maintenance was segmenting the corpus into multiple training datasets corresponding to conceptually distinct mechanisms. Indeed, this study highlights how choices about how to pre-process digital trace data in ways that match one's theoretical aims are critical to the endeavor of the modern computational sociology.

Attempts to create a quantifiable, replicable, and generalizable approach to measuring culture have frequently been criticized for their inability to accommodate and do justice to the tremendous heterogeneity inherent in any culture (Mohr et al., 2020). Culture is, in many ways, defined precisely by its multidimensional and interdependent nature: it cannot be reduced to a single variable such as ethnicity or religious affiliation, or even to a set of independent components. For instance, the same hand gesture may signify approval in one social context and a sharp rebuke in another.

Modern machine learning techniques may play a useful role in addressing this problem of cultural measurement. They offer a variety of tools—both supervised and supervised—that take rich data, like language and images, that are often unstructured and induce ordered patterns from the complexity (Nelson, 2020, 2021). These techniques allow for high dimensionality and interdependencies among cultural elements. In this paper, we employ a relatively straightforward application of such techniques. By imposing two social categories on our data—that is, pre-merger organizational affiliations—we simplify a complex, multidimensional cultural measurement problem into a basic, one-dimensional classification task, for which machine learning is especially well-suited. We chose a random forest algorithm because it is a well-established nonparametric method that accommodates interdependencies yet is robust to overfitting. Although a random forest model is generally uninterpretable, our goal was not to unearth the specific content of the cultural patterns but rather to understand the distribution of these patterns over individuals and time (Rodseth, 1998).

The approach taken in this paper joins a host of other recent efforts to incorporate machine learning methods into studies of culture (Corritore et al., 2020; Kozlowski et al., 2019; Lix et al., 2022). While some of these studies focus on illuminating the substance of culture, others like ours pay primary attention to cultural variation. Overall, studies such as these offer a flexible, inductive approach to pattern recognition that reconciles quantitative efforts to measure culture with the rich tradition of qualitative research.

Finally, we acknowledge that the study has certain limitations, which also point to directions for future research. First, although we believe our analytical approach can be readily extended to a wide range of social groups, it is unclear how our specific findings might generalize to other organizational mergers or more broadly to other contexts in which the social boundaries between groups are breached. Second, our research design reveals the associations between social positions, context, and the two forms of boundary maintenance but does not yield causal estimates. We leave to future research the task of identifying contexts in which people experience exogenous shifts in their social position or local context to establish a causal link between these variables and boundary maintenance. Next, our exploratory analyses of the implications of boundary maintenance for individual career success used a subjective measure of job performance, making it difficult to disentangle potential effects on actual job performance versus perceptions of performance. Future studies in this vein would benefit from having access to objective measures of individual productivity, as well as subjective evaluations of job performance. Lastly, our analytical approach focused on individuals who were present before and after a given merger and ignored the potential role that newcomers who joined after a merger or who were previously integrated into the firm through a prior merger might have played in the dynamics of boundary maintenance.

## Conclusion

This paper provides a conceptual and empirical framework for understanding how symbolic boundaries are maintained: through entrenching in existing symbolic distinctions and innovating new distinctions. It harnesses the tools of computational linguistics and machine learning to develop a method for measuring both forms of boundary maintenance and demonstrates the utility of this approach in the context of post-merger integration. This methodology can be readily extended to the study of symbolic boundaries across the wide range of social groups whose members draw distinctions with one another.

## References

Accominotti, F., Khan, S. R., & Storer, A. (2018). How Cultural Capital Emerged in Gilded Age America: Musical Purification and Cross-Class Inclusion at the New York Philharmonic. *American Journal of Sociology*, *123*(6), 1743–1783.

Amiot, C. E., Terry, D. J., & McKimmie, B. M. (2012). Social Identity Change During an Intergroup Merger: The Role of Status, Similarity, and Identity Threat. *Basic and Applied Social Psychology*, *34*(5), 443–455.

Bail, C. A. (2008). The Configuration of Symbolic Boundaries against Immigrants in Europe. *American Sociological Review*, *73*(1), 37–59.

Barth, F. (1969). *Ethnic Groups and Boundaries: The Social Organization of Culture Difference*. Little Brown & Company.

Berger, P. L., & Luckmann, T. (1966). *The Social Construction of Reality: A Treatise in the Sociology of Knowledge*. Open Road Media.

Berry, J., Kim, U., Minde, T., & Mok, D. (1987). Comparative Studies of Acculturative Stress. *International Migration Review*, *21*(3), 491–511.

Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent dirichlet allocation. *The Journal of Machine Learning Research*, *3*, 993–1022.

Bodner, J., & Capron, L. (2018). Post-merger integration. *Journal of Organization Design*, *7*(1), 1–20.

Bonikowski, B. (2016). Nationalism in Settled Times. *Annual Review of Sociology*, *42*(1), 427–449.

Bourdieu, P. (1984). *Distinction: A Social Critique of the Judgement of Taste*. Harvard University Press.

Bourdieu, P., & Passeron, J.-C. (1970). *La Reproduction: léments pour une théorie du système d'enseignement*. Minuit.

Brannon, T. N., & Walton, G. M. (2013). Enacting cultural interests: How intergroup contact reduces prejudice by sparking interest in an out-group's culture. *Psychological Science*, *24*(10), 1947–1957.

Carroll, G. R. (2015). Authenticity: Attribution, Value, and Meaning. *Emerging Trends in the Social and Behavioral Sciences*, 1–13.

Coleman, J. S. (1988). Social capital in the creation of human capital. *American Journal of Sociology*, *94*, S95–S120.

Corritore, M., Goldberg, A., & Srivastava, S. B. (2020). Duality in Diversity: How Intrapersonal and Interpersonal Cultural Heterogeneity Relate to Firm Performance. *Administrative Science Quarterly*, *65*(2), 359–394.

DiMaggio, P., Nag, M., & Blei, D. (2013). Exploiting affinities between topic modeling and the sociological perspective on culture: Application to newspaper coverage of us government arts funding. *Poetics*, *41*(6), 570–606.

Drori, I., Wrzesniewski, A., & Ellis, S. (2013). One Out of Many? Boundary Negotiation and Identity Formation in Postmerger Integration. *Organization Science*, *24*(6), 1717–1741.

Durkheim, É., Émile Durkheim, M., Cosman, C., Cladis, M., & Cladis, A. (2001 [1912]). *The elementary forms of religious life*. Oxford University Press.

Eliasoph, N., & Lichterman, P. (2003). Culture in interaction. *American Journal of Sociology*, *108*(4), 735–794.

Fishman, J. A. (2012). *The sociology of language*. De Gruyter Mouton.

Garg, N., Schiebinger, L., Jurafsky, D., & Zou, J. (2018). Word embeddings quantify 100 years of gender and ethnic stereotypes. *Proceedings of the National Academy of Sciences*, *115*(16), E3635–E3644.

Geertz, C. (1973). *The Interpretation Of Cultures*. Basic Books.

Gentzkow, M., Shapiro, J. M., & Taddy, M. (2019). Measuring Group Differences in High-Dimensional Choices: Method and Application to Congressional Speech. *Econometrica*, *87*(4), 1307–1340.

Gieryn, T. F. (1983). Boundary-Work and the Demarcation of Science from Non-Science: Strains and Interests in Professional Ideologies of Scientists. *American Sociological Review*, *48*(6), 781–795.

Goldberg, A. (2011). Mapping Shared Understandings Using Relational Class Analysis: The Case of the Cultural Omnivore Reexamined. *American Journal of Sociology*, *116*(5), 1397–1436.

Goldberg, A., Hannan, M. T., & Kovács, B. (2016). What Does It Mean to Span Cultural Boundaries? Variety and Atypicality in Cultural Consumption. *American Sociological Review*, *81*(2), 215–241.

Goldberg, A., Srivastava, S. B., Manian, V. G., Monroe, W., & Potts, C. (2016). Fitting In or Standing Out? The Tradeoffs of Structural and Cultural Embeddedness. *American Sociological Review*, *81*(6), 1190–1222.

Greenwald, A. G., & Banaji, M. R. (1995). Implicit social cognition: Attitudes, self-esteem, and stereotypes. *Psychological review*, *102*(1), 4.

Grillo, R. D. (2003). Cultural essentialism and cultural anxiety. *Anthropological theory*, *3*(2), 157–173.

Gumperz, J. J. (1977). The Sociolinguistic Significance of Conversational Code-Switching. *RELC Journal*, *8*(2), 1–34.

Khan, S. R. (2012). *Privilege: The Making of an Adolescent Elite at St. Paul's School*. Princeton University Press.

Kiley, K., & Vaisey, S. (2020). Measuring Stability and Change in Personal Culture Using Panel Data. *American Sociological Review*, *85*(3), 477–506.

Kozlowski, A. C., Taddy, M., & Evans, J. A. (2019). The Geometry of Culture: Analyzing the Meanings of Class through Word Embeddings. *American Sociological Review*, *84*(5), 905–949.

Lamont, M. (1992). *Money, Morals, and Manners: The Culture of the French and the American Upper-Middle Class*. University of Chicago Press.

Lamont, M., & Molnár, V. (2002). The Study of Boundaries in the Social Sciences. *Annual Review of Sociology*, *28*(1), 167–195.

Lamont, M., Pendergrass, S., & Pachucki, M. C. (2015). Symbolic Boundaries. In J. Wright (Ed.), *International Encyclopedia of Social and Behavioral Sciences* (2nd ed., pp. 850–855). Elsevier.

Le Mens, G., Hannan, M. T., & Pólos, L. (2015). Age-related structural inertia: A distance-based approach. *Organization Science*, *26*(3), 756–773.

Lieberson, S., & Bell, E. O. (1992). Children's first names: An empirical study of social taste. *American Journal of sociology*, *98*(3), 511–554.

Lix, K., Goldberg, A., Srivastava, S., & Valentine, M. A. (2022). Aligning differences: Discursive diversity and team performance. *Management Science*.

Mohr, J. W., Bail, C. A., Frye, M., Lena, J. C., Lizardo, O., McDonnell, T. E., Mische, A., Tavory, I., & Wherry, F. F. (2020). *Measuring Culture*. Columbia University Press.

Nahavandi, A., & Malekzadeh, A. R. (1988). Acculturation in Mergers and Acquisitions. *Academy of Management Review*, *13*(1), 79–90.

Nelson, L. K. (2020). Computational grounded theory: A methodological framework. *Sociological Methods & Research*, *49*(1), 3–42.

Nelson, L. K. (2021). Leveraging the alignment between machine learning and intersectionality: Using word embeddings to measure intersectional experiences of the nineteenth century u.s. south. *Poetics*, 101539.

Pachucki, M. A., Pendergrass, S., & Lamont, M. (2007). Boundary processes: Recent theoretical developments and new contributions. *Poetics*, *35*(6), 331–351.

Pennebaker, J. W., Chung, C. K., Ireland, M., Gonzales, A., & Booth, R. J. (2007). The Development and Psychometric Properties of LIWC2007.

Rodseth, L. (1998). Distributive Models of Culture: A Sapirian Alternative to Essentialism. *American Anthropologist*, *100*(1), 55–69.

Srivastava, S. B., & Banaji, M. R. (2011). Culture, cognition, and collaborative networks in organizations. *American Sociological Review*, *76*(2), 207–233.

Srivastava, S. B., Goldberg, A., Manian, V. G., & Potts, C. (2018). Enculturation Trajectories: Language, Cultural Adaptation, and Individual Outcomes in Organizations. *Management Science*, *64*(3), 1348–1364.

Suzuki, S. (1997). Cultural transmission in international organizations: Impact of interpersonal communication patterns in intergroup contexts. *Human Communication Research*, *24*(1), 147–180.

Swidler, A. (1986). Culture in Action: Symbols and Strategies. *American Sociological Review*, *51*(2), 273–286.

Terry, D. J., Carey, C. J., & Callan, V. J. (2001). Employee Adjustment to an Organizational Merger: An Intergroup Perspective. *Personality and Social Psychology Bulletin*, *27*(3), 267–280.

Vaisey, S., & Lizardo, O. (2016). Cultural Fragmentation or Acquired Dispositions? A New Approach to Accounting for Patterns of Cultural Change. *Socius*, *2*, 2378023116669726.

Van den Steen, E. (2010). Culture clash: The costs and benefits of homogeneity. *Management Science*, *56*(10), 1718–1738.

Vila-Henninger, L. (2015). Understanding Symbolic Boundaries and Improving Quantitative Analysis of Social Exclusion by Improving the Operationalization of Boundary Work. *Sociology Compass*, *9*(12), 1025–1035.

Weber, R. A., & Camerer, C. F. (2003). Cultural Conflict and Merger Failure: An Experimental Approach. *Management Science*, *49*(4), 400–415.

Wimmer, A. (2008). Elementary strategies of ethnic boundary making. *Ethnic and Racial Studies*, *31*(6), 1025–1055.

# Appendix A LIWC Categories

**Table A1:** Linguistic Inquiry and Word Count (LIWC) 2007 Categories Used in Classification Models (58 categories in total)

| Category | Examples | Words in Category |
|---|---|---|
| Word count | - | - |
| 1st pers singular | I, me, mine | 12 |
| 1st pers plural | We, us, our | 12 |
| 2nd person | You, your, thou | 20 |
| 3rd pers singular | She, her, him | 17 |
| 3rd pers plural | They, their, they'd | 10 |
| Impersonal pronouns | It, it's, those | 46 |
| Articles | A, an, the | 3 |
| Common verbs | Walk, went, see | 383 |
| Auxiliary verbs | Am, will, have | 144 |
| Past tense | Went, ran, had | 145 |
| Present tense | Is, does, hear | 169 |
| Future tense | Will, gonna | 48 |
| Adverbs | Very, really, quickly | 69 |
| Prepositions | To, with, above | 60 |
| Conjunctions | And, but, whereas | 28 |
| Negations | No, not, never | 57 |
| Quantifiers | Few, many, much | 89 |
| Numbers | Second, thousand | 34 |
| Swear words | Damn, piss, fuck | 53 |
| Social processes | Mate, talk, they, child | 455 |
| Family | Daughter, husband, aunt | 64 |
| Friends | Buddy, friend, neighbor | 37 |
| Humans | Adult, baby, boy | 61 |
| Positive emotion | Love, nice, sweet | 406 |
| Negative emotion | Hurt, ugly, nasty | 499 |
| Anxiety | Worried, fearful, nervous | 91 |
| Anger | Hate, kill, annoyed | 184 |
| Sadness | Crying, grief, sad | 101 |
| Insight | Think, know, consider | 195 |
| Causation | Because, effect, hence | 108 |
| Discrepancy | Should, would, could | 76 |
| Tentative | Maybe, perhaps, guess | 155 |
| Certainty | Always, never | 83 |
| Inhibition | Block, constrain, stop | 111 |
| Inclusive | And, with, include | 18 |

Continued on next page

**Table A1** – continued from previous page

| Category | Examples | Words in Category |
|---|---|---|
| Exclusive | But, without, exclude | 17 |
| See | View, saw, seen | 72 |
| Hear | Listen, hearing | 51 |
| Feel | Feels, touch | 75 |
| Body | Cheek, hands, spit | 180 |
| Health | Clinic, flu, pill | 236 |
| Sexual | Horny, love, incest | 96 |
| Ingestion | Dish, eat, pizza | 111 |
| Relativity | Area, bend, exit, stop | 638 |
| Motion | Arrive, car, go | 168 |
| Space | Down, in, thin | 220 |
| Time | End, until, season | 239 |
| Work | Job, majors, xerox | 327 |
| Achievement | Earn, hero, win | 186 |
| Leisure | Cook, chat, movie | 229 |
| Home | Apartment, kitchen, family | 93 |
| Money | Audit, cash, owe | 173 |
| Religion | Altar, church, mosque | 159 |
| Death | Bury, coffin, kill | 62 |
| Assent | Agree, OK, yes | 30 |
| Nonfluencies | Er, hm, umm | 8 |
| Fillers | Blah, Imean, youknow | 9 |

## Appendix B Heterogeneity in Linguistic Style

As mentioned before, several studies have demonstrated that there exists significant variation among individuals in a group with respect to cultural behaviors and enactment of symbolic boundaries (Bail, 2008; Bonikowski, 2016). Our data confirms this notion; we find that there is more within-bank heterogeneity in linguistic style than there is between-bank variation. Figure B1 plots the average ratio (across all LIWC categories) of the difference between banks and the standard deviation across individuals within banks. It remains below 1 for both mergers and all LIWC categories.

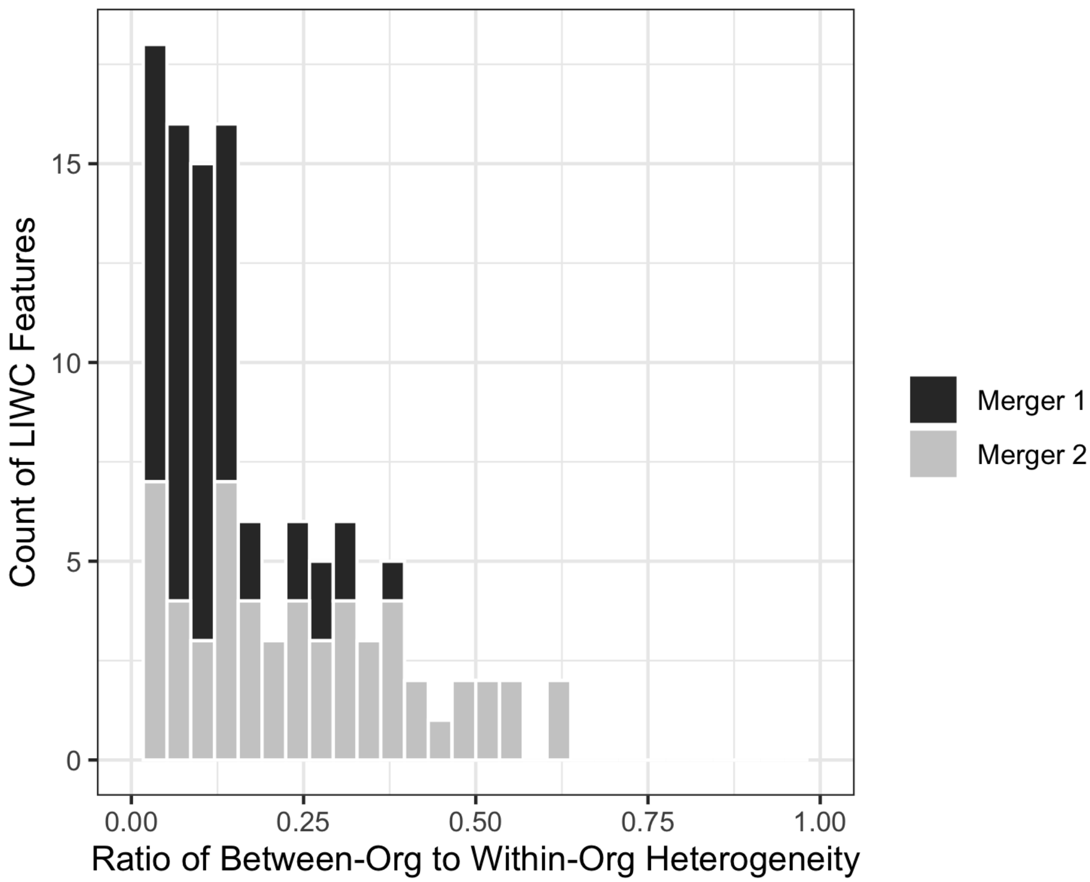


**Figure B1:** Histogram across all LIWC categories of the ratio of between-organization to within-organization heterogeneity, based on email communications in the 3 months prior to merger. The x-axis is defined as $\Delta_l/\overline{\sigma_l}$. $\Delta_l$ refers to the difference across banks between the average employee's use of LIWC category $l$. $\overline{\sigma_l}$ refers to the within-bank standard deviation of employees' use of LIWC category $l$, averaged across the two merging banks.

# Appendix C Summary Statistics

**Table C1:** Summary Statistics and Pairwise Correlations of Key Measures used in Regression Analyses

| | N | Mean | SD | 1. Retention | 2. Reformation | 3. Diff | 4. Acquired | 5. Position | 6. Male | 7. Tenure | 8. Ingroup | 9. Rating |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 1. Retention | 902 | 0.59 | 0.14 | | | | | | | | | |
| 2. Reformation | 902 | 0.21 | 0.1 | -0.8 | | | | | | | | |
| 3. Contemporaneous Diff | 902 | 0.58 | 0.13 | 0.2 | 0.32 | | | | | | | |
| 4. Acquired | 902 | 0.28 | 0.45 | -0.58 | 0.63 | 0.24 | | | | | | |
| 5. Hierarchical Position | 899 | 0.63 | 0.76 | -0.02 | -0.04 | 0 | 0 | | | | | |
| 6. Male | 899 | 0.29 | 0.45 | 0.05 | -0.06 | 0.07 | -0.13 | 0.32 | | | | |
| 7. Tenure | 899 | 6.07 | 7.26 | -0.13 | 0.14 | 0.06 | 0.35 | 0.21 | -0.01 | | | |
| 8. Local Ingroup Density | 900 | 0.84 | 0.24 | 0.28 | -0.2 | 0.04 | -0.34 | -0.02 | -0.04 | -0.06 | | |
| 9. Performance Rating | 595 | 3.29 | 0.44 | 0.07 | -0.05 | 0.16 | -0.12 | 0.28 | 0.32 | 0.08 | -0.12 | |
| 10. Total Outgoing Messages | 902 | 386.07 | 326.42 | 0.08 | -0.06 | 0.06 | -0.2 | 0.21 | 0.02 | 0.07 | -0.08 | 0.28 |