Two-Sided Cultural Fit: The Differing Behavioral Consequences of Cultural Congruence Based on Values Versus Perceptions

Richard Lu
University of California, Berkeley

Jennifer A. Chatman
University of California, Berkeley

Amir Goldberg
Stanford University

Sameer B. Srivastava
University of California, Berkeley

How do people establish and maintain cultural fit with an organization? Prior research has offered divergent and seemingly incongruous answers. One perspective focuses on the role of values, while another emphasizes perceptions. We propose that value congruence—the match between one's values and those that prevail and are normatively reinforced in an organization—shapes behavior when one periodically steps back from day-to-day interactions, assesses one's identification with an organization, and determines whether to stay or voluntarily depart. In contrast, we argue that perceptual congruence—one's understanding of an organization's cultural norms at a given point in time—influences behavior when one engages in routine interactions with peers. We therefore theorize that these two forms of cultural fit—one based on values and the other on perceptions—relate to distinct behaviors, voluntary exit and real-time linguistic congruence with peers, respectively. Drawing on email and survey data from a mid-sized technology firm, we find support for our theory and discuss implications of our findings for research on person-culture fit, dual-process models of culture and cognition, and the pairing of surveys with digital trace data.

Introduction

Whether assimilating to a country or adapting to a new school, people typically seek to fit in culturally with their social groups. The benefits of conformity, as well as the sanctions and penalties that come with failed cultural integration, are particularly stark in contemporary organizations. Indeed, prior work has consistently demonstrated that high levels of individual cultural fit are associated with increased productivity, stronger commitment, and less turnover (Kristof-Brown et al. 2005, Chatman and O'Reilly 2016). Moreover, employers are increasingly screening and selecting new hires based on their anticipated cultural fit rather than just their skills (Chatman 1991, Meyer et al. 2010, Rivera 2012). At the same time, as the average tenure in organizations has declined (Hall 1996), workers must frequently retool themselves culturally as they move from one organization to the next. Yet people vary considerably in how well they fit into and adapt to a given organization (Chatman 1989, Srivastava et al. 2018). How do people establish and maintain cultural fit in an organization and what are the behavioral consequences?

Existing research offers two different, and seemingly inconsistent, answers to this question. The first focuses on values. This line of work, echoing a long tradition in sociology and psychology, sees the locus of culture in the degree to which people embrace their group's behavioral norms. Fitting in therefore implies having preferences that are consistent with the norms that prevail in an organization.

A second explanation largely rejects the notion that values affect behavior, positing instead that culture shapes action through situational cues. This approach shifts focus from individuals' preferences to their readings of situations, arguing that behaviors are primarily driven by the cultural scripts invoked through interaction with others. An employee's decision to use polite language in a meeting, for example, often reveals little about her underlying preference for civil discourse but instead reflects the norms she observes in the behavior of other meeting participants. Indeed, people pursue action for which their "cultural equipment is well suited" (Swidler 1986, p. 277), suggesting that those who fit in are those whose readings of the cultural code lead them to behave in normatively appropriate ways.

These two perspectives appear to provide contrasting explanations for the sources of cultural fit. Whereas the former suggests that cultural fit is the result of internalizing and embracing prevailing values and norms, the latter views it as the product of correctly deciphering the normative code. In other words, two approaches to understanding cultural fit—one focused on values and the other on perceptions—make very different predictions who will fit into the culture of an organization. We propose that this theoretical ambiguity can be resolved by recognizing that culture operates at the individual level via two distinct forms of cognition—declarative versus non-declarative (Lizardo 2017)—that correspond to different types of behavior. Declarative knowledge, which reflects facets of culture that can be easily expressed and codified, is at play when people make more deliberative and conscious decisions, whereas non-declarative knowledge, which corresponds to aspects of culture that are difficult to articulate, shapes habitual and non-reflective behavior.

Given that values represent a form of declarative knowledge, we first propose that *value congruence*—the match between one's values and those that prevail and are normatively reinforced in an organization (Chatman 1989, Alba and Nee 2009)—shapes behavior when one periodically steps back from day-to-day interactions, assesses one's identification with an organization, and determines whether to stay or voluntarily depart. In contrast, because perceptions of multilayered social situations correspond to non-declarative knowledge, we argue that *perceptual congruence*—one's understanding of an organization's cultural norms at a given point in time—influences behavior when one engages in routine interactions with peers. Thus, we anticipate that these two forms of cultural fit—one based on values and the other on perceptions—will relate to distinct behaviors, voluntary exit and real-time *linguistic congruence* with peers, respectively.

To evaluate these ideas, we employ a multi-method empirical strategy that draws on survey data, eight years of internal email data, and personnel records from a mid-sized technology firm. We use the Organizational Culture Profile (Chatman 1991), a validated culture assessment, to measure value congruence and perceptual congruence. Linguistic congruence is measured by applying the interactional language-use model to a corpus of internal email messages (Srivastava et al. 2018, Goldberg et al. 2016).

We begin by reporting cross-sectional results that are consistent with our hypotheses. Yet, recognizing that cultural fit is likely to play out over time and that prior studies of cultural fit have focused on measures collected only once or a handful of times, we also employ a novel machine learning-based method to impute value congruence and perceptual congruence for individuals over time. Although this method still needs to be validated in other empirical settings where researchers have access to multiple waves of survey data and can therefore assess the degree to which the relationship between survey responses and communication behavior is stable over time, it nevertheless gives us a preliminary window into within-person changes in value congruence and perceptual congruence and thus allows us to estimate longitudinal models that corroborate our cross-sectional results. We conclude by discussing the implications of our findings for research on person-culture fit, dual-process models of culture and cognition, and the pairing of surveys with digital trace data.

Theory and Hypotheses

Cultural Fit Based on Values Versus Perceptions

Values—enduring beliefs about desired or undesired ways of acting—feature prominently in scholarship on culture and its consequences in organizations. Indeed, work on this topic has tended to conceptualize individual cultural fit through the prism of *value congruence*: the match between a person's values and those that predominate and are normatively reinforced in her social group. People whose ideal preferences are compatible with those prevalent in their organizational environment exhibit higher subjective well-being and enjoy greater attainment (Chatman and O'Reilly 2016).

Work that focuses on value congruence as the primary dimension of cultural fit has identified two core mechanisms that link values to individual outcomes in organizations. The first relates to self-perceptions. Individuals whose values are compatible with those prevalent in an organization are more likely to self-identify with that organization (O'Reilly and Chatman 1986, Cable and Judge 1996, Judge and Cable 1997). Such identification, in turn, leads to greater attachment, heightened motivation, stronger commitment, and higher productivity (Chatman 1991, Baron et al. 2001). The second relates to the ease of interpersonal interaction and coordination. Individuals who share similar values find it easier to interact with one another because they have mutually compatible expectations of behavior leading, potentially, to greater coordination within an organization (Morrison 2002, Elfenbein and O'Reilly 2007, Sørensen 2002). For example, employees who value detail-orientation will likely check in with their peers less frequently and expect them to deliver more thoroughly performed tasks than those who value speedy execution. Consequently, employees who differ in these value orientations will find it difficult and frustrating to interact with one another.

The notion that values are fundamental drivers of human behavior has a long history in sociology (Parsons 1968) and psychology (Schwartz 1992, Hofstede 2001). This research demonstrates, for example, that values are associated with cross-national and regional differences in economic growth (Inglehart and Baker 2000) and violence (Nisbett and Cohen 1996), as well as with individual lifestyle (Miles 2015), financial (Keister 2008), and occupational (Alesina et al. 2015) choices. Yet a growing body of research finds that people's stated values are, in many cases, poor predictors of their behavior (Greenwald and Banaji 1995). Economically disadvantaged high school students, for example, tend to express mainstream attitudes on educational achievement and sexual behavior but adopt behaviors that appear to be inconsistent with these ideals (Harding 2007). In organizations, too, people's behaviors are often incongruent with their stated beliefs: self-reported values on cross-functional collaboration, for example, are largely unrelated to individuals' propensity to build network ties that span functional boundaries (Srivastava and Banaji 2011).

Research in cultural sociology has therefore tended to downplay the role of values in shaping behavior. This work often relies on two fundamental and interrelated assumptions. The first is that "people know more culture than they use" (Swidler 1986, p. 277), namely, that they subscribe to multiple, and potentially inconsistent, cultural logics and value systems. Given this multiplicity, the same setting can elicit different interpretations, leading to inconsistent behavioral responses. The second assumption is that people's behavior is situationally driven. Subtle contextual cues in other people's behavior serve as signals about how to interpret a situation and, consequently, what kind of behavior is appropriate. Because these meanings emerge through interaction (Childress and Friedkin 2012, Gibson 2011), value assignment often occurs retroactively (Boltanski and Thévenot 2006).

This constructivist understanding of culture shifts focus from what people value to how they interpret their experiences of the world and produce meaning through interaction. Culture, according to this approach, systematically shapes behavior through what Eliasoph and Lichterman (2003) call "group styles:" idiosyncratic cultural codes that connect symbols, actions, and vocabularies to meaningful categories. Consider, for example, the perennially disgruntled employees in Weeks' (2004) ethnography of a British bank. To an outsider observing people habitually complaining, it may have seemed that these employees were fundamentally rejecting the organization and its culture. As Weeks artfully demonstrates, however, employees were instead partaking in rituals intended at reaffirming their bonds and their commitment to the bank.

Fitting in to an organizational culture depends on possessing the tacit and layered knowledge necessary for accurately deciphering this intricate cultural code. We refer to this ability as *perceptual congruence* and argue that it arises from two underlying processes.[1] The first relates to the person's construal of a situation, by which we mean the mental representation that she conjures when making sense of others' behaviors (DiMaggio and Goldberg 2018). A colleague's cynical joke in a meeting, for example, can be interpreted as a friendly attempt to establish rapport or as a derogatory comment aimed at undercutting others. An observer's capacity to correctly construe the meeting as friendly or adversarial depends on the compatibility between her and others' interpretations of participants' behaviors. Second, the person's reading of the norms that are prevalent in the organization shapes what behaviors she deems appropriate in light of her construal. Her ensuing behavior will be circumscribed by her understanding of the situation and what kinds of action it normatively affords.

Two-Sided Cultural Fit

A challenge raised by juxtaposing these two approaches to understanding cultural fit, one focused on perceptions and the other on values, is that they make very different predictions about what kinds of individuals will fit in culturally. Whereas the former is centered on the accuracy of perceptions, the latter emphasizes the importance of compatibility in values for behaving in culturally conforming ways. What explains differences in individuals' ability to exhibit cultural fit: the degree of alignment between their values and those that prevail in the environment, or the extent to which they can accurately read social situations and respond appropriately in a given context?

Drawing on advances in cognitive science, cultural sociologists have proposed that, at the individual level, culture generally operates via two distinct forms of cognition: declarative and non-declarative (Lizardo 2017). Declarative knowledge, which corresponds to facets of culture that are readily distilled and articulated, is at play when people make more deliberative and conscious decisions, whereas non-declarative knowledge, which refers to cultural understandings that are more nuanced and difficult to express in simple terms, shapes habitual and less reflective behavior. Recognizing that different processes are involved in habitual versus reflective behavior requires rethinking cultural fit as a two-sided construct rather than one that is determined solely by values or perceptions. Accordingly, we argue that values, which are a form of declarative knowledge, are consequential for reflective decisions that are less driven on normative cues given off by others. In contrast, perceptions, which correspond to non-declarative knowledge, matter for routine behavior that is guided in part by observing others and inferring what is appropriate.

Most activities in organizations occur routinely, in settings that provide high situational clarity given people's familiarity with the setting and the availability of habituated behavioral responses within it (Davis-Blake and Pfeffer 1989). We therefore posit that perceptual congruence will be consequential for individuals' ability to exhibit culturally compliant behavior in routine, day-to-day interaction. To productively participate in ritualistic complaining, for example, the employees in Weeks' (2004) ethnography of *BritArm Bank* had to complain at the appropriate level: not too much so as to avoid rocking the boat, but enough to signal membership and belonging with the group.

We refer to the linguistic expressions of such conformity to normative expectations as *linguistic congruence*.[2] We focus on linguistic congruence because it relies on real-time, fine-grained, and routine interaction among peers and thus represents one of the most non-declarative aspects of cultural understanding. Given that people are generally motivated to fit into their social groups, in part to avoid the social sanctions that arise when they deviate from normative expectations, we assume that those with more perceptual congruence will, all else equal, behave in more normatively compliant ways.

We further argue that value congruence will, in contrast, be less consequential for a person's capacity to conform to her group's routine normative expectations. Although people whose values are more congruent with their organization's may be motivated to behave in normatively compliant ways, they may still lack the knowledge needed to do so. It is one thing to prefer, for example, a cooperative work environment and another to understand which behaviors signal cooperativeness in a specific cultural context.

Instead, we expect that value congruence will predict behavior when people periodically assess their place in an organization and contemplate whether they want to stay or instead exit. When people make such decisions, they respond less to what types of appropriate behaviors the situation activates and more to their beliefs about what is desirable. Moreover, such deliberation often occurs in private contexts in which colleagues' behavioral cues and normative expectations are not on display and thus less salient. Together, these arguments lead us to formulate the following two hypotheses:

Hypothesis 1. *Perceptual congruence is positively related to the behavioral outcome of real-time linguistic congruence in routine interactions.*

Hypothesis 2. *Value congruence is positively related to the behavioral outcome of long-term attachment to the organization.*

Method

Overview

Previous work on cultural fit in organizations has, by and large, relied exclusively on self-reports to assess both cultural and behavioral variables. This approach has three major limitations (Gerald and George 2010). First, self-reports predominantly elicit, by design, deliberative cognition (e.g., subjective well-being or retroactive behavioral accounts). Second, habitual decision-making and the day-to-day behaviors it produces are difficult to detect through surveys. Previous work has therefore largely examined the relationship between self-reports and outcomes (such as promotion or departure), assuming that it is mediated by unobserved behaviors. Third, it is usually impractical or too costly to collect self-reports on a frequent basis. Consequently, they are not well-suited to measuring subtle changes on a granular timescale.

To address these limitations, we employ a multi-method approach that draws on both survey and email communication data. We begin by testing our hypotheses using cross-sectional data. We then use a machine learning technique to impute time-varying measures from cross-sectional data and to estimate longitudinal models with individual fixed effects that account for time-invariant unobserved heterogeneity.

Data

Our empirical setting is a mid-sized technology firm, from which we obtained three types of data:

Personnel Records—We received monthly extracts from the firm's human resource information system. These extracts included demographic information such as age and gender, organizational status such as departmental affiliation and start date, and information about individual outcomes such as monthly bonus received.

Email Data—We collected eight years of email data from the organization, including not only metadata (i.e., who sent messages to whom and when) but also raw message content. Given our focus on cultural dynamics within the organization, we excluded emails exchanged between employees and the outside world. We also eliminated automatically generated messages and, per instructions from the company's in-house lawyers, messages sent from or to members of the (small) legal department. The resulting data set included over five million unique emails.

Organizational Culture Profile—All employees were invited to complete an Organizational Culture Profile (OCP) (Chatman et al. 2014) assessment about the organization's current culture. We also asked a randomly selected half of employees to complete the assessment based on their own personally desired cultural characteristics.[3] As described below, our measure of perceptual congruence is based only on the assessment about the current culture, which 440 individuals completed. Our measure of value congruence entails a comparison of others' reports about the current culture with an individual's own preferences. Value congruence is therefore defined for the 238 people who completed the assessment about their personally desired culture.

Archived email data and personnel records were collected in multiple batches starting in 2015 and concluding toward the end of 2016. The OCP was implemented in October of 2016. Once we matched the raw email data to personnel records and removed identifying information, the resulting data set consisted of 29,255 person-month observations, spanning the period from 2008 to 2016.

Dependent Variables

Linguistic Congruence—Hypothesis 1 predicts a positive relationship between perceptual congruence and linguistic congruence. We operationalized linguistic congruence as the similarity between an individual's language and her reference group's, using the Interactional Language Use Model (ILUM) (Goldberg et al. 2016, Srivastava et al. 2018). Although language is not the only means through which culture is enacted—for example, culture also manifests in dress and various forms of nonverbal communication—it is a dominant medium through which cultural information is exchanged (Lazear 1999). Given that linguistic similarity can sometimes reflect alignment for non-cultural reasons—for example, two people coordinating on a shared task might use similar language even when they are culturally incompatible—we focus on the similarity of linguistic style between an individual and her reference group. Drawing on previous sociological work on culture (Bail et al. 2017, Doyle et al. 2017), ILUM uses the well-established and widely used Linguistic Inquiry and Word Count (LIWC) lexicon (Pennebaker et al. 2007) to measure linguistic style. LIWC is a semantic dictionary that maps words into 64 high-level emotional, cognitive, and structural categories. A comprehensive body of work demonstrates that the linguistic units identified by LIWC relate to a wide and universal array of meaningful psychological categories (Tausczik and Pennebaker 2010).

Using LIWC allows us to focus on expressions that are inherently cultural, while downplaying linguistic exchange that is organization- or context-specific or primarily related to functional coordination between organizational members. Imagine, for example, an organization with an aggressive and competitive culture. Such a culture might manifest linguistically in expressions of certainty, negation, and the use of swear words and other forms of non-deferential language. Contrast such a normative environment with one characterized by politeness and the use of tentative and inclusive language, indicating a collaborative and non-confrontational culture. LIWC is specifically designed to capture such culturally meaningful dimensions.

To derive our measure of linguistic congruence, we first translated raw emails into LIWC category counts. We then aggregated each individual's incoming and outgoing emails into monthly time periods and represented each person-month observation as two probability distributions of outgoing and incoming communication over LIWC categories. We used the Jensen-Shannon divergence metric (inverse and log-transformed) between these two probability distributions as the measure of linguistic congruence.

Intuitively, when the outgoing and incoming distributions are nearly identical, the divergence approaches zero, suggesting high linguistic congruence; conversely, greater deviation between the probabilities of usage of LIWC categories translates to greater divergence and thus implies lower linguistic congruence. Thus, the more an employee's use of cognitive, emotional, and structural terms in sent emails matches the use of those terms in received emails, the greater her linguistic congruence in a given month.

We discuss the technical details of this measure in Appendix A, which also reports the results of two validation checks. The first compares LIWC and OCP categories to demonstrate that our language-based measure reflects culturally meaningful content. The second reports the results of a simulation analysis, which reveals that our measure is robust to the exclusion of arbitrary sets of LIWC categories. In other words, even if we assume that given sets of LIWC categories are culturally meaningless, their exclusion would have a negligible effect on the resulting measure.

Voluntary Exit—In Hypothesis 2, we predict that value congruence will be negatively related to a person's chances of departing voluntarily. We identified voluntary exit based an employee's departure date. We used company records to distinguish between voluntary and involuntary exit.

Work Performance—To help validate our measures of value congruence, perceptual congruence, and linguistic congruence, we report below results of models in which we examine their relationship to individual work performance. We used monthly bonus payments as the measure of individual work performance. For people in job roles such as sales or operations in which productivity could be objectively assessed, the company established a formula that linked specific productivity indicators—for example, a sales person's conversion of leads into revenue—to monthly bonus payments. Given that the distribution of bonuses was skewed, we logged this measure in the analyses reported below.

Independent Variables

Perceptual Congruence—We used the OCP to derive our measure of perceptual congruence. The OCP consists of 54 value statements (e.g., fast moving, being precise) that emerged from a review of academic and practitioner-oriented writings on culture (O'Reilly et al. 1991). Using the Q-sort methodology (Block 1961), respondents are asked to array these 54 statements into nine categories, with a specified number of statements in each category. The required distribution of statements across categories is 2-4-6-9-12-9-6-4-2, so that, for example, respondents rating the current culture of their organization would place two value statements each in the "most characteristic" and "most uncharacteristic" categories, respectively, four value statements each in the "quite characteristic" and "quite uncharacteristic" categories respectively, and 6 statements each in the "fairly characteristic" and "fairly uncharacteristic" categories respectively, and so on, until all 54 value statements were categorized. Unlike a Likert-format scoring scheme in which many or all items can be rated as high or low, or a ranking process, which, with 54 value statements to rank, would be unwieldy for human raters, this semi-idiographic approach forces respondents to choose cultural value statements that are most and least characteristic of their organization.

To derive our measure of perceptual congruence, we focused on an OCP question that was asked of all respondents: "To what extent do the value statements characterize the organization as a whole?" We defined perceptual congruence as the match between an individual's current culture profile and those of a reference group of peers. To make this measure comparable to our measure of linguistic congruence, we chose the same reference group—that is, the set of colleagues a person had email contact with in a given month weighted by communication volume.

Value Congruence—For value congruence, we focused on participants' OCP responses to the following: "To what extent do the value statements characterize your personally desired values, that is, the values you desire in an organization?" We defined value congruence as the correspondence between an individual's personal culture profile (what she prefers) and the reference group's current culture profile (the culture that actually exists in the organization). For consistency, we chose the same reference group for value congruence as we did for perceptual congruence and linguistic congruence.

Imputing Perceptual Congruence and Value Congruence Over Time

The procedure above yields cross-sectional measures of perceptual congruence and value congruence. Models based on such measures cannot account for time-invariant, unobserved heterogeneity—for example, stable personality traits and dispositions that might be related to our outcomes of interest.

We therefore undertook a procedure to transform our cross-sectional measures of value congruence and perceptual congruence into longitudinal measures. Taking inspiration from Salganik's (2017) notion of *amplified asking*—that is, combining surveys with digital trace data to infer responses for people who cannot be feasibly surveyed or whose responses are missing—we undertook a procedure based on machine learning techniques to identify from raw email content (rather than the higher-level LIWC categories used to derive our measure of linguistic congruence) the "linguistic signature" of perceptual congruence and value congruence.

We assumed that, if language reflects internal processes of cognition (Pinker 2007), then there should be an identifiable relationship between email communication and the two dimensions of perceptual congruence and value congruence. Specifically, we used a random forest model to help uncover this underlying link between language and cognition (Ho 1995, Friedman et al. 2001). Random forest models have several beneficial characteristics for this task: they can detect arbitrary, nonlinear relationships; they typically require fewer observations than do other machine learning methods to produce comparable results; and they are inherently robust to overfitting, or incorrectly inferring signal from idiosyncratic noise in the data. Figure 1 provides a conceptual overview of this procedure. Further procedural details are provided in the Appendix B. Because we only had access to a single administration of the OCP assessment, we were unable to validate a core assumption underpinning this imputation approach: that the relationship between survey responses and communication behavior remains relatively stable over time. Analyses based on our imputed measures should therefore be interpreted with some caution—pending validation of this assumption in future research.

Control Variables

We estimated both within-person and between-person models for our analyses. In within-person models, which were based on our imputed measures of value congruence and perceptual congruence, we included three time-varying controls that prior research suggests are relevant to the study of cultural conformity. First, we included (lagged) managerial status since employees may be more likely to accommodate the behaviors, and specifically the language use, of interlocutors who possess greater structural power (Mayer et al. 2009). Next, we included an indicator for an employee's first year in the organization given that this is typically a period of intense socialization and cultural learning. Finally, we included departmental dummies since departments vary in relative centrality and power, which may in turn influence the degree to which their members are motivated to conform to behavioral norms (Thompson 1967, Salancik and Pfeffer 1974).[4] For our between-person models, we included additional control variables for age, and gender.

[FIGURE 1 ABOUT HERE]

Analytical Approach

We tested Hypothesis 1, which suggests that perceptual congruence will be positively related to linguistic congrunce, by estimating OLS regressions on cross-sectional data, as well as fixed effect regressions based on longitudinal data (including imputed measures of perceptual congruence and value congruence). Hypothesis 2, which predicts that value congruence will be negatively related to voluntary exit, was estimated using Cox proportional hazard models. We standardized all variables in the regression models reported below. We use lagged predictors in longitudinal models to address (though not fully resolve) reverse causality.

Results

Main Results

Table 1 provides a test of Hypothesis 1. The first three models report results from cross-sectional data in which perceptual congruence and value congruence are derived directly from the Organizational Culture Profile (OCP). Both measures are imputed in the longitudinal models that follow.

Models 1 to 3 report results from cross-sectional data, with linguistic congruence averaged over three months preceding the administration of the OCP. In support of Hypothesis 1, perceptual congruence is significantly related to linguistic congruence, while value congruence is not; moreover, these patterns hold whether the value congruence and perceptual congruence are modeled separately (Models 1 and 2) or jointly (Model 3).

[TABLE 1 ABOUT HERE.]

Table 1, Models 4 to 6, echo the results from the cross-sectional analyses in longitudinal specifications that include individual, department, and year fixed effects. The longitudinal results provide further support for Hypothesis 1 given that perceptual congruence is significantly related to linguistic congruence, while value congruence is not. As individuals' perceptual congruence increases, their linguistic congruence correspondingly increases. Changes in value congruence, in contrast, are unrelated to changes in linguistic congruence.

Of the control variables included in the models, only managerial status and tenure are significant. We conjecture that managers exhibit greater linguistic congruence than do individual contributors either because their general tendency toward cultural congruity was conducive to their past promotion into management or because subordinates are more likely to linguistically accommodate their communication style. Consistent with previous work on enculturation (Srivastava et al. 2018), we also find that individuals exhibit significantly lower linguistic congruence during their first year in the organization.[5]

Table 2 reports tests of Hypothesis 2. Our competing risks Cox hazard models focus on voluntary exit as a function of value congruence and perceptual congruence (with involuntary exit serving as the competing risk).

As Table 2 indicates, value congruence is associated with a decreased risk of voluntary exit, while perceptual congruence is not. The importance of value congruence in affecting voluntary departures, based on the imputed longitudinal measure, is consistent with prior work based on a cross-sectional measure of value congruence that predicted departure from firms up to two years later (Chatman 1991).[6]

[TABLE 2 ABOUT HERE.]

Supplemental Analyses—Evaluating the Variables of Interest

Here we summarize two additional analyses that sought to evaluate the validity of our measures of value congruence and perceptual congruence, including the imputed versions of these measures, as well as linguistic congruence, which was not imputed. First, given that we theorized that value congruence is relatively stable over time while perceptual congruence is more susceptible to change, we traced the two imputed measures over a person's tenure in the organization. We restricted this analysis to the first 36 months of employment given that only about 10% of employees had tenure exceeding 36 months during our observation period. We separately estimated OLS and fixed effect regressions of value congruence and perceptual congruence using indicators for each month (up to month 36 of employment). These results are depicted in Figure 2. According to both models, when employees first enter the organization, they have relatively high value congruence and relatively low perceptual congruence. Through approximately the first year of employment, however, perceptual congruence increases sharply and continues a more gradual ascent thereafter. In contrast, value congruence increases—albeit not as steeply—in the first four months of employment and then remains mostly stable over the remaining months. These results support our contention that value congruence is relatively stable, while perceptual congruence is more malleable.

[FIGURE 2 ABOUT HERE]

Second, in Table 3 we report the results of OLS regressions with individual, department, and year fixed effects, where the dependent variable is bonus (logged) and independent variables—linguistic congruence, perceptual congruence (imputed) and value congruence (imputed)—are lagged. The fixed effects specification with lagged predictors allows us to estimate the effects of within-person change in the three congruence measures on subsequent productivity.

Whether modeled independently or together, all three measures are significantly positively related to productivity. Thus we find, consistent with prior work, that linguistic congruence (Srivastava et al. 2018) and value congruence (Chatman 1991) are positively related to positive job performance—even when we use imputed longitudinal measures of value congruence and perceptual congruence. We also demonstrate the novel finding that perceptual congruence is related to performance independent of its effects on linguistic congruence. Indeed, the coefficients for liguistic congruence and perceptual congruence are of similar magnitude, and the two variables retain their significance even when included together in Model 4. In contrast, the association between value congruence and bonus is more modest. This result is consistent with our expectation that value congruence remains more stable over time. Given that the unwavering component of value congruence is subsumed in the individual fixed effect, it is not surprising that its time-varying component accounts for less of the variance in job performance. Overall, these supplemental analyses help to validate the value congruence and perceptual congruence measures derived from our imputation methodology.

[TABLE 3 ABOUT HERE.]

Discussion and Conclusion

Adjustments to new and changing cultural environments are a fixture of modern life. People's identities in contemporary society typically intersect many social boundaries—including ethnic, religious, political, occupational, and organizational. This crisscrossing of boundaries requires ongoing effort. The contemporary workplace—with its growing emphasis on culture on the one hand and employees' declining average tenure on the other—is a central arena in which these cultural transitions play out. Navigating the cultural heterogeneity across and within organizations involves maintaining multiple and partial commitments to different cultural orders, which in turn requires cultural awareness and adaptability (Friedland and Alford 1991, Morris et al. 2015, DiMaggio and Goldberg 2018).

Prior research has offered competing explanations for why some people fit in better than others. One perspective has highlighted the importance of alignment between individual and group values in shaping behavior, while another has emphasized the role of situational cues and the ability to read the group's cultural code. We develop a theoretical account that reconciles these competing perspectives. Drawing on dual-process theories of culture and cognition, we develop a two-sided theory of cultural fit that encompasses both values and perceptions. We argue and find empirical support for the notion that value congruence matters for the deliberative choices such as whether or not to voluntarily exit an organization, while perceptual congruence instead shapes routine, habitual behavior—specifically, real-time linguistic conformity with peers.

Although we develop a novel theoretical account of cultural fit and bring together disparate forms of data and analytical methods, we also acknowledge that the study has certain limitations. First, it is based on data from a single organization, which raises questions about the extent to which the findings would generalize to other settings. Next, our imputation models rely on the implicit assumption that the relationship between language use and the relevant cultural fit variables is stable over time. As such, future studies that include multiple administrations of the OCP are needed to validate this assumption. Finally, even with the inclusion of individual fixed effects in our longitudinal models, we acknowledge that our estimates are not causal.

One possibility for pinning down a causal relationship between perceptual congruence and linguistic congruence would be to implement a field experiment in which employees take an OCP, with a treatment group receiving feedback about how their perceptions differed from the actual perceptions of their interlocutors and a control group receiving no such feedback. Assuming such an intervention resulted in an increase perceptual congruence in the treatment group but not in the control group, researchers could then examine whether it led to subsequent increases in the treatment group's linguistic congruence relative to the control group's.

These limitations notwithstanding, our theoretical framework and concomitant findings offer three contributions. The first is in advancing person-culture fit theory. Specifically, we demonstrate that the behavioral consequences of cultural fit vary with different modes of cognition. Declarative knowledge is operative when people contemplate how their values match those that prevail in the organization and shape outcomes such as the choice to leave an organization. In contrast, non-declarative knowledge is at play when people determine how to communicate in real-time with their peers. Together, these insights open the door to further investigations of the role that the dual modes of cognition play in shaping how people fit into social groups. Next, we demonstrate that both value congruence and perceptual congruence, as well as the behavioral manifestation of the latter, linguistic congruence, enable people to reap positive career rewards. Indeed, all three of our fit measures are positively linked to individual productivity, as indicated by bonus payments.

The conceptual separation of cultural fit into value congruence and perceptual congruence also paves the way for investigations into: (a) how these two dimensions relate to each other dynamically; and (b) the degree to which they are influenced by individual-level characteristics versus broader structural factors. On the former question, we speculate that value congruence may provide a motivational channel through which a person is more or less vigilant in achieving and maintaining perceptual congruence. We similarly conjecture that people with chronically low value congruence may be able to maintain high perceptual congruence for a finite period of time but that doing so may, over time, adversely affect their identity and sense of self-worth (cf. Hochschild 2012). Conversely, even if those with high perceptual congruence and low value congruence do not experience intrapsychic conflict, they may still experience the deleterious effects of being judged by others as inauthentic. Alternatively, we speculate that such individuals may—through self-perception and attribution processes (Ross 1977)—begin to experience an increase in value congruence. Examining the interrelationships between value congruence and perceptual congruence over time is a fruitful avenue for further developing theories of person-culture fit.

On the latter question, we anticipate that perceptual congruence arises in part through the quality of a person's social network in an organization: those who are connected to peers who are perceptually accurate are more likely to update their potentially flawed interpretations of the culture and converge to a more accurate understanding than are those whose peers are miscalibrated. We see great potential in future research takes advantage of exogenous shifts in network structure—stemming, for example, from an unanticipated organizational restructuring that is undertaken for reasons unrelated to internal collaboration patterns—to causally identify the link between the perceptual congruence of peers and that of the focal individual. It remains to be explored the degree to which one's perceptual congruence is a function of individual-level characteristics such as extraversion (John et al. 1999) or self-monitoring orientation (Snyder 1979) versus the interpersonal transmission of culture between individuals (Herrmann et al. 2013).

Next, we contribute to dual-process theories of culture and cognition (Vaisey 2009, Miles 2015, Lizardo et al. 2016) in two key ways. First, we make a conceptual link between modes of cognition—based on declarative versus explicit and non-declarative knowledge (Lizardo 2017)—and the types of behavior a person engages in within organizational settings. Whereas previous work in this tradition has thought about the link between values and behavior in binary terms—i.e., values either do or do not shape behavior—we develop a more nuanced account of the relationship by fusing dual-process models with a theory of organizational behavior. Our results indicate that values matter for some kinds of behavior (voluntary exit) but not others (linguistic conformity). This insight paves the way for exploring more generally how values matter for a broader range of behavior, especially when there is variation in the social reference groups that people perceive to be relevant (Diehl and McFarland 2010). Second, although dual-process theories of culture in action have proliferated, the empirical evidence in support of their link to concrete behaviors remains scant. We add to this evidence base by establishing a clear link between cultural fit constructs that are tied to declarative versus non-declarative cognition and consequential behaviors such as how people communicate with their colleagues, their choice of voluntary exit, and their level of work productivity (as reflected in bonus payments).

Finally, through this work, we make two methodological contributions. First, we bring together an established culture assessment approach (i.e., the Organizational Culture Profile) with an unobtrusive data source (i.e., internal email communication) that provides a window into complementary facets of culture. In doing so, we demonstrate how the integration of disparate approaches to assessing culture can yield richer and more complete insights into the complex cultural dynamics that occur within organizations. Second, building on Salganik's (2017) notion of "amplified asking," we demonstrate an empirical approach that transforms a one-time self-report into a longitudinal data set. Such an approach is of course, selectively appropriate, with requirements that include having a sufficient number of survey observations, access to rich communication content, protocols and safeguards to protect individual privacy and company confidentiality, and significant computational bandwidth. Nevertheless, given the ubiquity of digital trace data, the increasing difficulty of collecting survey data (particularly over time and from a large number of organizations), the widespread dissemination of off-the-shelf machine learning tools, and the declining cost of processing capacity, we anticipate that the pairing of self-reports and digital trace data will become increasingly common in social science research (Evans and Aceves 2016, McFarland et al. 2016, Lazer and Radford 2017). We see great potential for such work to more fully illuminate how different facets of culture relate to one another and jointly shape organizational life.

Endnotes

[1]We distinguish perceptual congruence from two seemingly related constructs: cultural intelligence and self-monitoring. Cultural intelligence is defined as "an individual's capability to function and manage effectively in culturally diverse settings" (Ang et al. 2007, p. 336). It is a relatively stable individual difference that focuses on an individual's efficacy "in situations arising from differences in race, ethnicity, and nationality" (Ang et al. 2007, p. 336). Although some researchers have applied it conceptually to organizational culture (Earley and Mosakowski 2004), the construct is "uniquely relevant to intercultural contexts rather than monocultural contexts" (Van Dyne et al. 2019, p. 1). Similarly, high self-monitors (Snyder 1979) are consistently responsive to social cues of situational appropriateness (Snyder 1979, Kilduff and Day 1994, Sasovova et al. 2010). They tend to regulate their behavior given their read of what is expected of them, whereas low self-monitors adhere to their sense of self, irrespective of the situation. Self-monitoring is also related to a persistent capacity for deep-acting, the ability to adapt emotions to organizational expectations, leading to more genuine displays of cultural congruence (Grandey 2000, Scott et al. 2012). Like cultural intelligence, self-monitoring is an individual difference that is often described as a personality characteristic (Ang et al. 2006, Snyder 1979). Whereas cultural intelligence and self-monitoring are stable psychological traits, perceptual congruence is a state that can ebb and flow across a person's tenure in an organization. It is quite likely that individuals who are higher in cultural intelligence or self-monitoring will exhibit, all else equal, greater perceptual congruence. Yet perceptual congruence ultimately reflects an individual's exposure to culturally relevant signals given off by others. Even those with high levels of cultural intelligence or self-monitoring will misinterpret the cultural code if the peers they learn the code from are behavioral misfits. Moreover, the signals that give rise to perceptual congruence are independent of the dimensions that matter for cultural intelligence—chiefly race, ethnicity, and nationality.

[2]Although we follow Goldberg et al. (2016) and Srivastava et al. (2018) in how we operationalize linguistic congruence, we depart from them in how we label this construct. They refer to linguistic

conformity with peers as a behavioral measure of "cultural fit." Given that we consider multiple manifestations of cultural fit in this paper, to avoid confusion, we generally refer to the specific constructs of value congruence, perceptual congruence, and linguistic congruence.

[3]The other half completed an assessment of the cultural characteristics needed for the organization to be successful in the future. We shared the results of this latter assessment with organizational leaders as a condition of gaining access to the organization as a research site; however, we do not report these results here because they do not pertain to our theory and hypotheses.

[4]Managerial status and departmental affiliation can be estimated in fixed effect models because some employees get promoted from individual contributor to managerial roles and because some employees move across departments.

[5]Tenure has a curvilinear relationship with linguistic congruence, steadily increasing during the first six to twelve months and gradually stabilizing thereafter. Because individuals vary significantly in their rate of enculturation, we use a binary indicator for early tenure.

[6]Neither perceptual congruence nor value congruence is significant in predicting involuntary exit when we use the same framework with voluntary exit as the competing risk.

References

Alba R, Nee V (2009) *Remaking the American Mainstream: Assimilation and Contemporary Immigration* (Harvard University Press).

Alesina A, Algan Y, Cahuc P, Giuliano P (2015) Family values and the regulation of labor. *Journal of the European Economic Association* 13(4):599–630.

Ang S, Van Dyne L, Koh C (2006) Personality correlates of the four-factor model of cultural intelligence. *Group & Organization Management* 31(1):100–123.

Ang S, Van Dyne L, Koh C, Ng YK, Templer KJ, Tay C, Chandrasekar AN (2007) Cultural intelligence: Its measurement and effects on cultural judgment and decision making, cultural adaptation, and task performance. *Management and Organization Review* 3(3):335–371.

Bail CA, Brown TW, Mann M (2017) Channeling hearts and minds: Advocacy organizations, cognitive-emotional currents, and public conversation. *American Sociological Review* 0003122417733673.

Baron JN, Hannan MT, Burton MD (2001) Labor pains: Change in organizational models and employee turnover in young, high-tech firms. *American Journal of Sociology* 106(4):960–1012.

Block J (1961) *The Q-sort Method in Personality Assessment and Psychiatric Research* (Thomas Springfield, IL).

Boltanski L, Thévenot (2006) *On Justification: Economies of Worth* (Princeton University Press).

Cable DM, Judge TA (1996) Person–organization fit, job choice decisions, and organizational entry. *Organizational Behavior and Human Decision Processes* 67(3):294–311.

Chatman JA (1989) Improving interactional organizational research: A model of person-organization fit. *Academy of Management Review* 14(3):333–349.

Chatman JA (1991) Matching people and organizations: Selection and socialization in public accounting firms. *Administrative Science Quarterly* 36(3):459.

Chatman JA, Caldwell DF, O'Reilly CA, Doerr B (2014) Parsing organizational culture: How the norm for adaptability influences the relationship between culture consensus and financial performance in high-technology firms. *Journal of Organizational Behavior* 35(6):785–808.

Chatman JA, O'Reilly CA (2016) Paradigm lost: Reinvigorating the study of organizational culture. *Research in Organizational Behavior* 36:199–224.

Childress CC, Friedkin NE (2012) Cultural reception and production: The social construction of meaning in book clubs. *American Sociological Review* 77(1):45–68.

Davis-Blake A, Pfeffer J (1989) Just a mirage: The search for dispositional effects in organizational research. *Academy of Management Review* 14(3):385–400.

Diehl D, McFarland D (2010) Toward a historical sociology of social situations. *American Journal of Sociology* 115(6):1713–1752.

DiMaggio P, Goldberg A (2018) Searching for homo economicus: Variation in americans' construals of and attitudes toward markets. *European Journal of Sociology* 1–39.

Doyle G, Goldberg A, Srivastava S, Frank M (2017) Alignment at work: Using language to distinguish the internalization and self-regulation components of cultural fit in organizations. *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, 603–612.

Earley CP, Mosakowski E (2004) Cultural intelligence. *Harvard Business Review* 82(10):139–146.

Elfenbein HA, O'Reilly CA (2007) Fitting in: The effects of relational demography and person-culture fit on group process and performance. *Group & Organization Management* 32(1):109–142.

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

Evans JA, Aceves P (2016) Machine translation: Mining text for social theory. *Annual Review of Sociology* 42:21–50.

Friedland R, Alford RR (1991) *Bringing Society Back in: Symbols, Practices and Institutional Contradictions* (University of Chicago Press).

Friedman J, Hastie T, Tibshirani R (2001) *The Elements of Statistical Learning*, volume 1 (Springer series in statistics New York).

Gerald JH, George SH (2010) Self-report: Psychology's four-letter word. *American Journal of Psychology* 123(2):181–188.

Gibson DR (2011) Avoiding catastrophe: The interactional production of possibility during the cuban missile crisis. *American Journal of Sociology* 117(2):361–419.

Goldberg A, Srivastava SB, Manian GV, Monroe W, Potts C (2016) Fitting in or standing out? the tradeoffs of structural and cultural embeddedness. *American Sociological Review* 81(6):1190–1222.

Grandey AA (2000) Emotional regulation in the workplace: A new way to conceptualize emotional labor. *Journal of Occupational Health Psychology* 5(1):95.

Greenwald AG, Banaji MR (1995) Implicit social cognition: Attitudes, self-esteem, and stereotypes. *Psychological Review* 102(1):4.

Hall DT (1996) *The Career Is Dead–Long Live the Career. A Relational Approach to Careers. The Jossey-Bass Business & Management Series.* (ERIC).

Harding DJ (2007) Cultural context, sexual behavior, and romantic relationships in disadvantaged neighborhoods. *American Sociological Review* 72(3):341–364.

Herrmann PA, Legare CH, Harris PL, Whitehouse H (2013) Stick to the script: The effect of witnessing multiple actors on children's imitation. *Cognition* 129(3):536–543.

Ho TK (1995) Random decision forests. *Document Analysis and Recognition, 1995., Proceedings of the Third International Conference on*, volume 1, 278–282 (IEEE).

Hochschild AR (2012) *The Managed Heart: Commercialization of Human Feeling* (Univ of California Press).

Hofstede G (2001) *Culture's Consequences: Comparing Values, Behaviors, Institutions and Organizations Across Nations* (SAGE Publications).

Inglehart R, Baker WE (2000) Modernization, cultural change, and the persistence of traditional values. *American Sociological Review* 65(1):19–51.

John OP, Srivastava S, et al. (1999) *The Big-Five Trait Taxonomy: History, Measurement, and Theoretical Perspectives*, volume 2 (University of California Berkeley).

Judge TA, Cable DM (1997) Applicant personality, organizational culture, and organization attraction. *Personnel Psychology* 50(2):359–394.

Keister LA (2008) Conservative protestants and wealth: How religion perpetuates asset poverty. *American Journal of Sociology* 113(5):1237–1271.

Kilduff M, Day DV (1994) Do chameleons get ahead? the effects of self-monitoring on managerial careers. *Academy of Management Journal* 37(4):1047–1060.

Kristof-Brown AL, Zimmerman RD, Johnson EC (2005) Consequences of individual's fit at work: A meta-analysis of person-job, person-organization, person-group, and person-supervisor fit. *Personnel Psychology* 58:281–342.

Kullback S, Leibler RA (1951) On information and sufficiency. *The Annals of Mathematical Statistics* 22(1):79–86.

Lazear EP (1999) Culture and language. *Journal of Political Economy* 107(S6):S95–S126.

Lazer D, Radford J (2017) Data ex machina: Introduction to big data. *Annual Review of Sociology* 43:19—39.

Lin J (1991) Divergence measures based on the shannon entropy. *IEEE Transactions on Information Theory* 37(1):145–151.

Lizardo O (2017) Improving cultural analysis: Considering personal culture in its declarative and nondeclarative modes. *American Sociological Review* 82(1):88–115.

Lizardo O, Mowry R, Sepulvado B, Stoltz DS, Taylor MA, Van Ness J, Wood M (2016) What are dual process models? implications for cultural analysis in sociology. *Sociological Theory* 34(4):287–310.

Mayer DM, Kuenzi M, Greenbaum R, Bardes M, Salvador RB (2009) How low does ethical leadership flow? test of a trickle-down model. *Organizational Behavior and Human Decision Processes* 108(1):1–13.

McFarland DA, Lewis K, Goldberg A (2016) Sociology in the era of big data: The ascent of forensic social science. *The American Sociologist* 47(1):12–35.

Meyer JP, Hecht TD, Gill H, Toplonytsky L (2010) Person–organization (culture) fit and employee commitment under conditions of organizational change: A longitudinal study. *Journal of Vocational Behavior* 76(3):458–473.

Miles A (2015) The (re)genesis of values: Examining the importance of values for action. *American Sociological Review* 80(4):680–704.

Morris MW, Chiu Cy, Liu Z (2015) Polycultural psychology. *Annual Review of Psychology* 66:631–659.

Morrison EW (2002) Newcomers' relationships: The role of social network ties during socialization. *Academy of Management Journal* 45(6):1149–1160.

Nisbett RE, Cohen D (1996) *Culture of Honor: The Psychology of Violence in the South* (Westview Press).

O'Reilly CA, Chatman J (1986) Organizational commitment and psychological attachment: The effects of compliance, identification, and internalization on prosocial behavior. *Journal of Applied Psychology* 71(3):492.

O'Reilly CA, Chatman J, Caldwell DF (1991) People and organizational culture: A profile comparison approach to assessing person-organization fit. *Academy of Management Journal* 34(3):487–516.

Parsons T (1968) *The Structure of Social Action* (Free press).

Pennebaker JW (2013) *The Secret Life of Pronouns: What Our Words Say About Us* (Bloomsbury USA), ISBN 9781608194964, URL `https://books.google.com/books?id=p9KmCAAAQBAJ`.

Pennebaker JW, Chung CK, Ireland M, Gonzales A, Booth RJ (2007) *Linguistic Inquiry and Word Count (LIWC): LIWC2007* (Austin, TX: LIWC.net).

Pinker S (2007) *The Stuff of Thought: Language as a Window Into Human Nature* (New York, NY: Viking).

Potts C (2011) Sentiment-aware tokenizer. *Creative Commons Attribution-NonCommercial-ShareAlike 3.0 Unported License: http://creativecommons.org/licenses/by-nc-sa/3.0/* .

Rivera LA (2012) Hiring as cultural matching: The case of elite professional service firms. *American Sociological Review* 77(6):999–1022.

Ross L (1977) The intuitive psychologist and his shortcomings: Distortions in the attribution process. *Advances in Experimental Social Psychology* 10:173–220.

Salancik GR, Pfeffer J (1974) The bases and use of power in organizational decision making: The case of a university. *Administrative Science Quarterly* 453–473.

Salganik MJ (2017) *Bit by Bit: Social Research in the Digital Age* (Princeton University Press).

Sasovova Z, Mehra A, Borgatti SP, Schippers MC (2010) Network churn: The effects of self-monitoring personality on brokerage dynamics. *Administrative Science Quarterly* 55(4):639–670.

Schwartz SH (1992) Universals in the content and structure of values: Theoretical advances and empirical tests in 20 countries. Zanna MP, ed., *Advances in Experimental Social Psychology*, volume 25, 1 – 65 (Academic Press).

Scott BA, Barnes CM, Wagner DT (2012) Chameleonic or consistent? a multilevel investigation of emotional labor variability and self-monitoring. *Academy of Management Journal* 55(4):905–926.

Snyder M (1979) Self-monitoring processes. *Advances in Experimental Social Psychology* 12:85–128.

Sørensen JB (2002) The strength of corporate culture and the reliability of firm performance. *Administrative Science Quarterly* 47(1):70–91.

Srivastava SB, Banaji MR (2011) Culture, cognition, and collaborative networks in organizations. *American Sociological Review* 76(2):207–233.

Srivastava SB, Goldberg A, Manian GV, Potts C (2018) Enculturation trajectories: Language, cultural adaptation, and individual outcomes in organizations. *Management Science* 64(3):1348–1364.

Swidler A (1986) Culture in action: Symbols and strategies. *American Sociological Review* 273–286.

Tausczik YR, Pennebaker JW (2010) The psychological meaning of words: Liwc and computerized text analysis methods. *Journal of language and social psychology* 29(1):24–54.

Thompson JD (1967) *Organizations in Action: Social Science Bases of Administrative Theory* (Transaction publishers).

Vaisey S (2009) Motivation and justification: A dual-process model of culture in action. *American Journal of Sociology* 114(6):1675–1715.

Van Dyne L, Ang S, Tan ML (2019) Cultural intelligence. *Oxford Bibliographies* (Englewood Cliffs, NJ: Prentice-hall).

Weeks J (2004) *Unpopular Culture: The Ritual of Complaint in a British Bank* (University of Chicago Press).

FIGURES

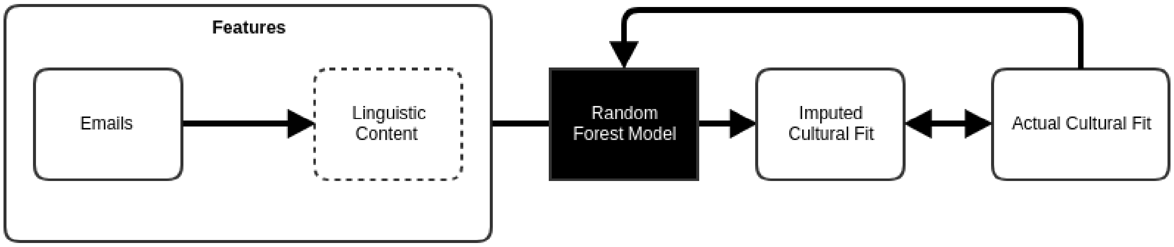

Figure 1 **Conceptual Overview of the Machine Learning Process**

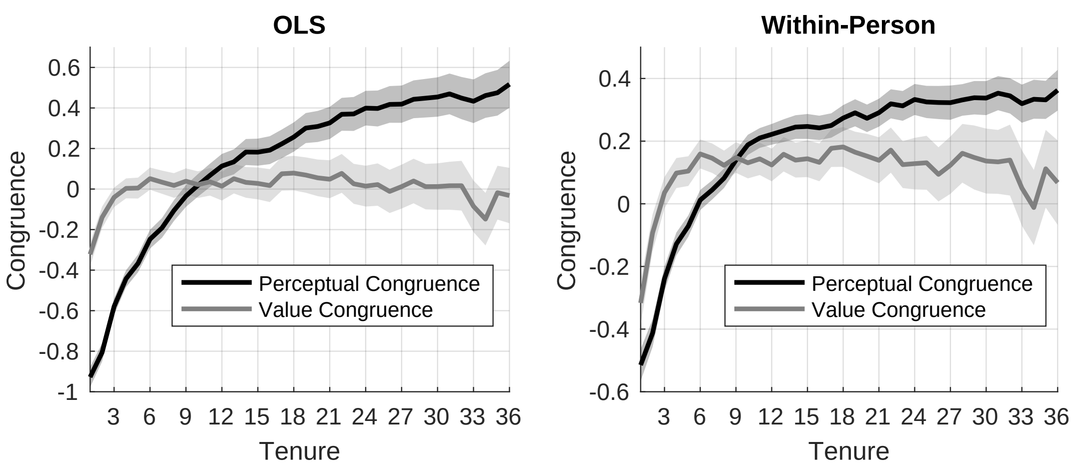

Figure 2 **OLS and fixed effect regressions of perceptual congruence and value congruence, with indicators for each tenure month up to 36 months in the company.**

TABLES

Table 1 Cross-Sectional and Longitudinal Fixed Effects Regressions of Linguistic Congruence

	Cross-Sectional			Longitudinal		
	Model 1†	Model 2†	Model 3†	Model 4	Model 5	Model 6
Perceptual Congruence‡	0.122***		0.149***	0.046**		0.046**
	(3.56)		(3.37)	(2.81)		(2.79)
Value Congruence‡		-0.008	-0.040		0.013	0.012
		(-0.17)	(-0.86)		(1.35)	(1.29)
Manager	0.613***	0.599***	0.555***	0.293***	0.297***	0.292***
	(6.73)	(4.20)	(3.92)	(5.42)	(5.47)	(5.40)
First Year	-0.246**	-0.351***	-0.317**	-0.074*	-0.082**	-0.074*
	(-3.20)	(-3.49)	(-3.13)	(-2.54)	(-2.81)	(-2.53)
Female	0.043	-0.033	-0.065			
	(0.62)	(-0.35)	(-0.68)			
Age	-0.003	-0.002	0.001			
	(-0.84)	(-0.30)	(0.10)			
Constant	0.345*	0.223	0.183	-0.142	-0.145	-0.145
	(2.37)	(1.13)	(0.93)	(-1.14)	(-1.11)	(-1.17)
Individual FE	No	No	No	Yes	Yes	Yes
Department FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	No	No	No	Yes	Yes	Yes
Observations	386	209	202	24215	24215	24215
R^2	0.275	0.235	0.279	0.107	0.075	0.107

t statistics in parentheses; standard errors clustered by individual when individual fixed effects are used

† Linguistic congruence is averaged over 3 months, ‡ Imputed and lagged measures in Models 4-6

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 2 Competing Risks Model of Voluntary Exit

	Model 1	Model 2
Perceptual Congruence	1.005	
	(0.07)	
Value Congruence		0.876*
		(-2.30)
Manager	0.833	0.864
	(-0.77)	(-0.62)
Female	1.386*	1.392*
	(2.53)	(2.56)
Age	0.901**	0.902**
	(-3.23)	(-3.23)
Age^2	1.001**	1.001**
	(3.20)	(3.22)
Num. Employees	1.002***	1.002***
	(9.46)	(9.96)
Department Dummies	Yes	Yes
Observations	27467	27467
χ^2	172.161	177.689
Log-Likelihood	-1320.27	-1318.36

Exponentiated coefficients; t statistics in parentheses
Standard errors clustered by individual; Sample weights by tenure
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 3 Fixed Effect Regressions of Bonus (logged)

	Model 1	Model 2	Model 3	Model 4
Linguistic Congruence†	0.131***			0.122***
	(4.45)			(4.14)
Perceptual Congruence†		0.144***		0.122**
		(3.97)		(3.05)
Value Congruence†			0.056**	0.046*
			(3.18)	(2.37)
Manager	-0.194	0.025	0.063	-0.180
	(-1.12)	(0.13)	(0.31)	(-1.02)
Constant	5.642***	5.394***	5.299***	5.666***
	(28.18)	(26.63)	(25.68)	(28.47)
Individual FE	Yes	Yes	Yes	Yes
Department FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Observations	4785	6379	6379	4780
Num. Inidividuals	1058	1304	1304	1057
R^2	0.059	0.043	0.040	0.065

t statistics in parentheses; standard errors clustered by individual

† lagged variables, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

APPENDIX A: LINGUISTIC CONGRUENCE

The Interactional Language Use Model

We implement the procedure detailed in Goldberg et al. (2016) and Srivastava et al. (2018) to measure behavioral fit. We begin by using LIWC to translate each individual's outgoing and incoming messages in each period t (defined as a calendar month) into probability distributions over the 64 LIWC categories. Specifically, we define $\overrightarrow{m}_{it}$ as each email individual i sends at time t and $\overleftarrow{m}_{it}$ as each email individual i *receives* at time t. We then define the set of LIWC categories as L and the set of all times in any given month as T. Our procedure iterates over all emails sent and received and produces $\overrightarrow{m}^l_{it}$ and $\overleftarrow{m}^l_{it}$ for the count of terms in email $\overrightarrow{m}_{it}$ and $\overleftarrow{m}_{it}$ in LIWC category $l \in L$, respectively. Then, by aggregating all individual email counts $\overrightarrow{m}^l_{it}$ and $\overleftarrow{m}^l_{it}$ for $t \in T$, it produces sent and received LIWC counts in month T, $\overrightarrow{m}^l_{iT}$ and $\overleftarrow{m}^l_{iT}$. We normalize each LIWC count in each month by the total of all LIWC counts in that month to transform the LIWC probability distribution to a standard probability distribution. We use the notation, O^l_{iT} to denote the outgoing normalized probability and I^l_{iT} to denote the incoming normalized probability.

$$O^l_{iT} = \frac{\overrightarrow{m}^l_{iT}}{\sum_{l\in L} \overrightarrow{m}^l_{iT}} \tag{1}$$

$$I^l_{iT} = \frac{\overleftarrow{m}^l_{iT}}{\sum_{l\in L} \overleftarrow{m}^l_{iT}} \tag{2}$$

We define an individual i's linguistic congruence in month T as the negative log of the Jensen-Shannon (JS) divergence (Lin 1991) metric between i's outgoing and incoming normalized distributions:

$$BF_{iT} = -log\,(JS(O_{iT} \parallel I_{iT})) \tag{3}$$

where the JS-divergence between two probability distributions is defined as a symmetric measure built by first taking the mean probability distribution between the normalized outgoing and incoming distributions, $M_{iT} = \frac{1}{2}(O_{iT} + I_{iT})$, and summing the Kullback-Leibler (KL) divergence (Kullback and Leibler 1951) of the outgoing and incoming distributions from that mean probability distribution.

$$JS(O_{iT} \parallel I_{iT}) = \frac{1}{2} KL\,(O_{iT} \parallel M_{iT}) + \frac{1}{2} KL\,(I_{iT} \parallel M_{iT}) \tag{4}$$

$$KL(D_{iT} \parallel M_{iT}) = \sum_{l\in L} D^l_{iT} \log_2 \frac{D^l_{iT}}{M^l_{iT}} \tag{5}$$

Validation of Linguistic Congruence

We have argued above that the LIWC lexicon, on which the linguistic congruence measure is based, is a useful categorization scheme for measuring culturally meaningful behaviors. Indeed, as previous work demonstrates (e.g. Goldberg et al. 2016, Srivastava et al. 2018), this measure of linguistic congruence is effective at predicting individual attainment in an organization. Since this is the first time our measure of linguistic congruence has been related to a validated measure of organizational culture, the OCP, we also sought assurances that the LIWC categories contained face valid connections to the existing OCP dimensions. Therefore, we conducted two types of analyses to further establish the behavioral measure's construct validity.

First, we compared respondents' language use to their responses to the OCP survey. Recall that we asked respondents to describe their desired culture (personal culture survey) and their perception of the organizational culture (current culture survey). We expected there to be a systematic relationship between people's desired and perceived cultures on the one hand and their linguistic behaviors on the other. For example, it would seem plausible that a preference for a people-oriented cultural environment would be reflected in greater use of affective words. Thus, we expected to observe a systematic relationship between people's cultural preferences and perceptions, as reflected in their explicit responses to the OCP and their use of language as captured by LIWC.

To examine this, we compared individuals' rankings of the 54 OCP categories with their LIWC category frequencies in outgoing email communication in a 3-month period close to the OCP survey administration. For the personal culture survey, we found 229 significantly correlated ($p < 0.05$) pairs of OCP and LIWC categories (with sample size of 231 individuals). For the current culture survey, we found 583 significant correlations (for 414 individuals). We found an even greater number of significant OCP/LIWC pair correlations when comparing the current culture survey to respondents' incoming email communication, suggesting that—consistent with our hypotheses—individuals' perceptions of the culture are inherently related to the behaviors they observe. We also compared LIWC frequencies to the eight high-level OCP categories (such as collaborative or detail-oriented, see Chatman et al. (2014) for details). For the personal cultural survey we find that 34% of LIWC categories are correlated with at least one high-level dimension, and that 85% of LIWC categories are correlated with at least one high-level dimension in the current culture survey. Together, these analyses indicate that LIWC use significantly and substantially co-varies with desired and perceived culture.

As illustration, we examine the link between language use and a preference for a people orientated culture. We find that respondents who value people orientation tend to include more affect words (e.g., happy, cry, abandon), perceptual process words (e.g., observe, hear, feel), positive emotion words (e.g., love, nice, sweet), and second-person words (e.g., you, your) in their outgoing communication.[7]. We refrain from substantively interpreting these findings, but we view them as qualitative evidence for the cultural meaningfulness of LIWC use and leave a systematic exploration of the complex relationship between stated beliefs and naturally occurring linguistic behaviors to future work.

In our second test of the construct validity of our linguistic congruence measure, we recognized that LIWC was originally developed as a means to identify the linguistic signatures of psychological, rather than purely cultural categories. Whereas some linguistic categories contained in the LIWC lexicon, such as swearing, are clearly inherently related to culture, others, such as the use of articles, are more ambiguously cultural. Thus, we sought to understand whether our linguistic congruence measure represented a meaningful and relevant set of culturally oriented linguistic categories.

Before discussing these analyses in detail we highlight why we assume that LIWC categories are culturally meaningful. Specifically, while some LIWC categories may initially appear to be unrelated to culture, extensive research by Pennebaker (2013) suggests that the categories are meaningful at both a psychological and sociological level. For example, the use of articles such as *a*, *an* or *the*—each of which seemingly represents a minute technical linguistic decision—actually reflects the speaker's emotional stability, organization, and conservatism (Pennebaker 2013). A group that uses a linguistic style that emphasizes articles might therefore be indicative of a rule-oriented culture that emphasizes attention to detail.

Thus, rather than requiring a typology that distinguishes non-cultural from cultural LIWC categories and that maps the latter to underlying cultural dimensions, we assumed that all LIWC categories are culturally meaningful and that the same category might vary in its cultural meaning across contexts. Our measure of behavioral cultural fit therefore takes all LIWC categories into account and does not privilege certain categories over others.

To test our assumption, we analyzed the measure's robustness to LIWC category inclusion. Let $k < 64$ be the size of a subset of LIWC categories used to generate an alternative measure of linguistic congruence, labeled BF_k. We randomly selected k LIWC categories and constructed the measure as we did above (according to equation 3), using only this subset of categories. We repeated this process 1,000 times for each value of k (because $\binom{64}{k}$ is extremely large for most values of k, we could not realistically explore all possible subsets). For each BF_k that we generated, we identified its correlation with the original BF measure based on all 64 categories.

We report the average correlation between BF_k and BF for all 1,000 random samples in Figure A1. As the plot clearly indicates, the linguistic congruence measure is robust regardless of whether LIWC categories are removed. The measure remains effectively unchanged even if half of the LIWC categories are removed. We interpret these results as an indication of two properties. First, linguistic congruence is not driven by one or a handful of LIWC categories. It is therefore not merely a reflection of a specific linguistic feature or style. Second, the pattern illustrated in Figure A1 indicates that even if certain LIWC categories are culturally irrelevant in this context, their inclusion in the measure construction does not bias its value. In other words, even if we were to conclude that half of the LIWC categories are non-cultural (a conclusion that, for the reasons stated above, we believe is unwarranted) and decide to remove them from the measure, we would still recover near-identical values.

APPENDIX B: MACHINE LEARNING PROCEDURE

Overview

The procedure consisted of five major steps, which are documented at a conceptual level in Figure 1 in the main manuscript and described in greater detail below.

Our first step was to translate the raw email data into a format that is usable by the random forest model. We tokenized and stemmed all words in the body of email messages. Tokenization involves separating the text into distinct terms, for which we used the TwitterTokenizer designed for linguistic analysis Potts (2011). Stemming involves reducing each term to a root form, for which we used the Porter Stemmer from the python nltk package. We removed all characters that could not be encoded into unicode, such as "\x00," and split the text into n-stems, where n is in the set [1,2,3]. Given that language use tends to follow the power law, in which few terms are used frequently and many terms are used infrequently, we then undertook steps to reduce the dimensionality of the data to make it computationally tractable. We retained all n-stems in emails sent from individuals, but only uni-stems in emails sent to individuals. Additionally, we retained only those n-stems that were used by at least 1% of employees in a subsample of emails. Finally, we used principal component analysis (PCA) to further reduce dimensionality, retaining only the top 3,000 PCA components for each type of n-stem. These resulting components served as the feature inputs to our model.

The second step was to transform our measures of cognitive cultural fit into categories that are more conducive to classification given the relatively small number of observations from which we had to fit the model. Recall that perceptual congruence and value congruence were computed as correlations, ranging from 0 to 1. We transformed these continuous measures into three discrete categories–low, medium, and high. Intuitively, this allowed our model to detect distinctive features of belonging to each category, an important characteristic to which we will return when we discuss the testing of our model. For perceptual congruence, we set the cutoffs for low fit at 20% and for high fit at 80%, with everything else considered medium fit. For value congruence, for which we had even fewer observations, we had to set more extreme cutoffs at 10% and 90% to achieve strong model fit.

The third step was to use our feature inputs and their now-discrete mappings to cognitive cultural fit to train a random forest model. The random forest model is an ensemble method, which means it aggregates and blends multiple independent decision trees (Ho 1995, Friedman et al. 2001). After several such decisions according to specific features of the input, all of the inputs are sorted into decision leaves. The random forest model then collects those independent trees and their leaves and predicts results for new observations. New observations get sorted into resultant leaves depending on their own features, and their probabilities of being predicted as a certain class depend on the other data points sorted into that leaf in the trained model. In a simplistic model, imagine that the only decision is that PCA1 $> .5$ and that all observations with PCA1 $> .5$ are high in cultural fit. Then, a new observation whose PCA1 $> .5$ would also get sorted into the same leaf and would then be classified as high cultural fit.

The fourth step was to evaluate the trained model. To do so, we assessed the model's predictions compared to the original continuous values. Random forest models produce, along with the classifications of input, probabilities of observations belonging to each class. Conceptually, this means that if an observation has certain characteristics that correspond to a given class, it will have a higher probability of being in that class. For example, if an individual's email communication has indicators of low, medium, and high cognitive cultural fit, but more indicators of high cultural fit than the others, then his or her output from the random forest model might indicate a 0.2 probability of low fit, a 0.3 probability of medium fit, and a 0.5 probability of high fit. We can then take a weighted sum of these probabilities to generate a measure that is conceptually analogous to the original continuous measure. We used a mix of methods to evaluate the model, including the area under the curve of the receiving operating characteristic curve (ROC AUC), precision-recall, and separation between low and high cognitive cultural fit with respect to the original continuous values. As reported in Appendix C, the final models we used performed well on these evaluations.

The final step was to impute perceptual congruence and value congruence using their corresponding random forest models for all individuals in all time periods for which we had corresponding email data. To do this, we followed the first step above to retrieve the input feature vector for each individual over time and used all the linguistic data for each individual up to a certain month to impute perceptual congruence and value congruence for that individual in that month.

There were a total of over five million unique emails. Each email can be sent from an individual and several other individuals (via the to/cc/bcc lines). We included both messages sent to and received from the focal individual in our final model.

Dimensionality Reduction of Features Considering the size of our potential feature vector, we used dimensionality reduction techniques to make our process computational tractable. In particular, we used a discriminative heuristic to determine which n-stems to keep, since there is a tradeoff between keeping frequent and non-frequent terms: frequent terms allow for discrimination to the extent that they are used differently among a large population of people, while non-frequent terms allow for discrimination to the extent that some people use them and others do not. Given this trade-off, we retained those n-stems that were used by at least 99% of all employees, regardless of their objective frequency. To retain as much information from this pared down set of n-stems, we used principal component analysis (PCA). This allowed us to reduce the hundreds of thousands of features to only a few thousand per n-stem, while still retaining a large part of the variance of the original data. Because of the exponential size of the "to" stems compared to the "from" stems, we ended up using the top 3,000 PCA components from the "from" uni-, bi-, and tri-stems, and from the "to" uni-stems.

Random Forest Model Specification We selected the random forest model because of several favorable characteristics. First, random forest models allow for nonlinear relationships between input and output. Decision trees in general, of which random forest is a collection, thus allow for arbitrarily complex relationships, which we would assume govern the relationship between linguistic data and cognitive cultural fit. Second, random forests are ensembles of decision trees, which inherently reduce overfitting and increase robustness. Since there is the potential for a link between linguistic data and cognitive cultural fit to be extremely idiosyncratic (e.g., use of a certain phrase or way of communicating), it greatly helps that we use a more robust method. Third, random forest models do not require as much training data as neural networks. Deep neural networks have the same, if not better, ability to pick up complex relationships, but require far more training data, depending on the depth of the model. As a result, random forest models are simpler and tend to require fewer training data for comparable results.

We split the data into the usual training, development, and testing sets, with 56% of the original data in the training set, 14% in the development set, and 30% in the testing set. Because of the way the random forest algorithm is implemented, it is strongly vulnerable to the "class imbalance" problem. Specifically, if the input to the model from the training set were 10% class 0, 80% class 1, and 10% class 2, then the model would err towards predicting most new observations as class 1. To overcome this, we used a bootstrapping procedure that randomly samples with replacement the lesser classes until they reach the amount of the most populated class. This procedure ensured that, on average, input classes were balanced and therefore class prediction depended more on the splits than on the original balance of the input classes. In addition to searching the hyperparameter space, we also tested varying N for bootstrapped samples.

Test Set Metrics

Because of the way we constructed our pseudo-continuous imputed cultural fit, we needed to use a set of test metrics that accurately capture what it means to have a "good model." The choice of bounds for the continuous to discrete distribution is forced; it is an educated guess that produces empirically validated results. Therefore, observations that lie just on one side may not differ substantively from observations that lie just on another side. Concretely, observations that are on the high end of the medium cultural fit may be very similar to observations that are on the low end of the high cultural fit, given that we had set the cutoff ourselves. Therefore, our measures should focus less on perfect categorization (i.e., precision, recall), and more on separation of low and high cultural fit and predictive power of imputed results on actual results. As a result, our performance metrics are a mix of the traditional machine learning metrics, as well as novel metrics we developed ourselves.

For the traditional test metrics, we present the pairwise precision and recall measures on the test set. We provide the pairwise precision recall rather than an F score, because we differentially care about the pairwise results. That is, we care the most about the precision recall between the high and the low cultural fits and less about the precision recall between the mid and either high or low cultural fits, as per our previous discussion.

[TABLE B1 ABOUT HERE.]

A better metric might be to directly examine the separation between groups. If we link the original continuous values with the classifications, then we would see a split like in the figure below.

[FIGURE B1 ABOUT HERE.]

We then used the means and standard deviations of each group to see if the classifier successfully split the observations into statistically distinct groups. We find that the models appear to appropriately distinguish between the low and high groups.

[TABLE B2 ABOUT HERE.]

Finally, we used the receiver operating characteristic curve (ROC) that has become popular in machine learning. Since the ROC works with threshold probabilities of classification, mapping the true positive rate versus the false positive rate at different thresholds, it conceptually measures the extent to which the rank-ordering of predicted values is in line with expectations. For a perfect area under the curve (AUC), the rank-ordering would be monotonically increasing such that all actual values of 1 would have higher probabilities of being classified as 1 than all actual values of 0, and vice versa. Since we have three classes versus the regular binary classification, we use the micro-averaged ROC curve, which takes into account this structure. The ROC curves with their AUC's are presented below.

[TABLE B3 ABOUT HERE.]

APPENDIX FIGURES

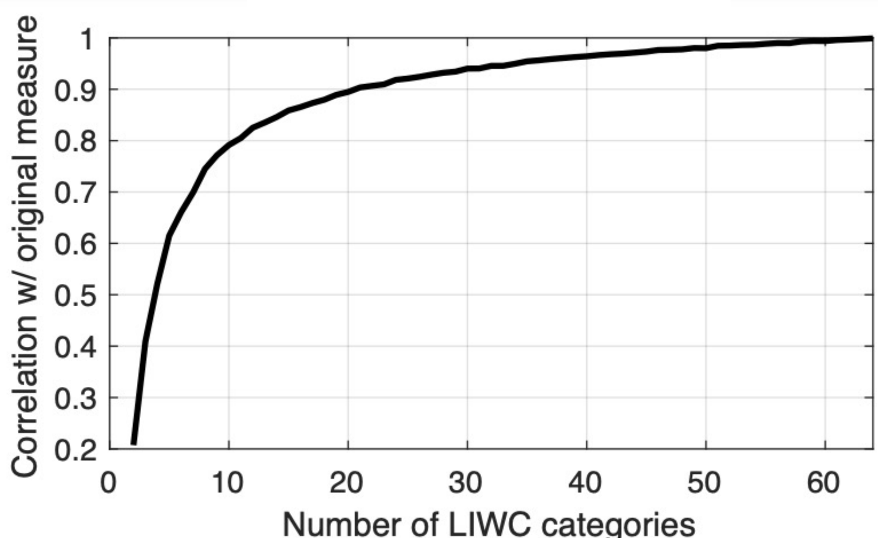

Figure A1 **Robustness of the linguistic fit measure to simulated changes in LIWC category composition**

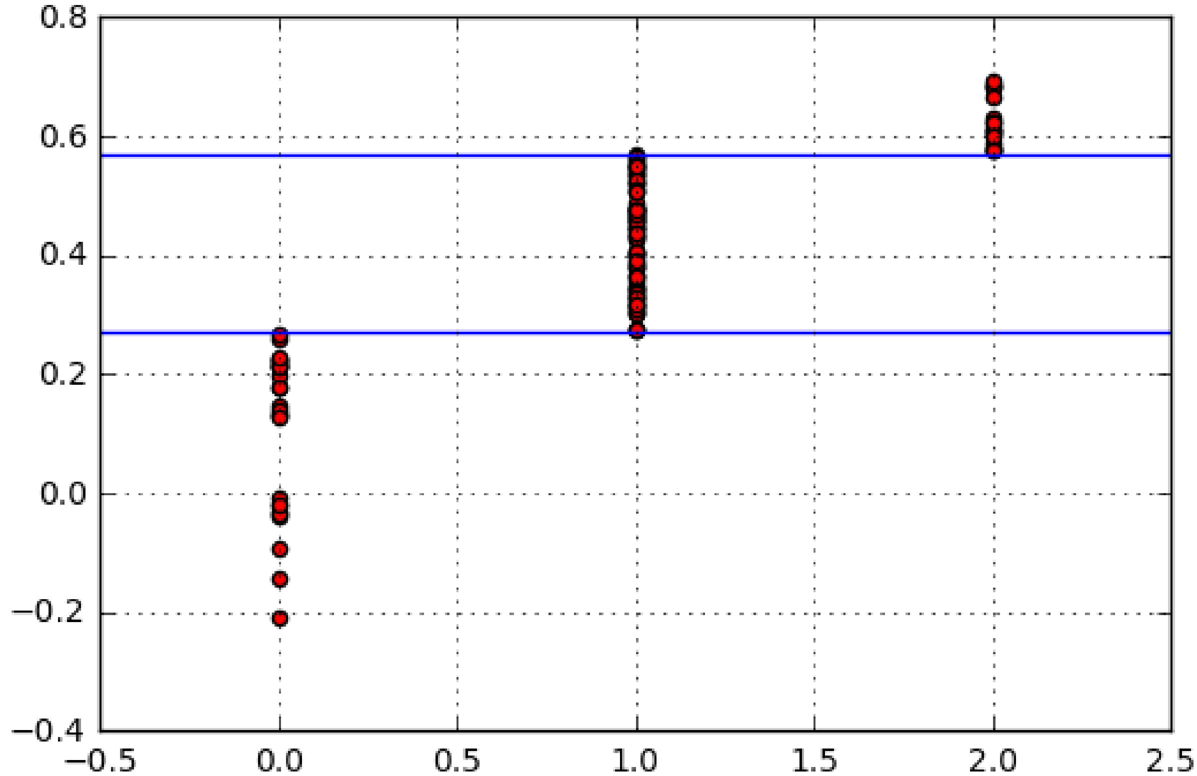

Figure B1 **Division of Continuous Cultural Fit into Classes**

APPENDIX TABLES

Table B1 Test Set Precision-Recall Metrics for Imputations

	Precision Low-High	Precision Low-Mid	Precision Mid-High	Recall Low-High	Recall Low-Mid	Recall Mid-High
Perceptual Congruence	0.857	0.726	0.767	0.267	0.651	0.711
Value Congruence	1	0.952	0.950	0.667	0.952	0.934

Table B2 p-Values for Difference in Means between Low and High

	P-Value
Perceptual Congruence	$2.661e{-3}$
Value Congruence	$8.500e{-6}$

Table B3 Areas under the ROC Curve

	ROC AUC
Perceptual Congruence	0.740
Value Congruence	0.950