Deciphering the Cultural Code: Perceptual Congruence, Behavioral Conformity, and the Interpersonal Transmission of Culture

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? Existing research offers a person-centered and static answer, focusing on fit based on shared values or stable traits. In contrast, we develop a novel and complementary perspective that conceives of fit as a dynamic equilibrium between one's private beliefs and public behaviors. We develop a theory of *perceptual congruence*—defined as a person's understanding of a group's cultural norms at a given point in time—and highlight how culturally congruent behavior arises in response to a person's exposure to and understanding of situational norms. Drawing on survey data, eight years of internal email data, and personnel records from a mid-sized technology firm, we examine the consequences of perceptual congruence for contemporaneous behavioral fit, as indicated by the linguistic similarity between one's email communications and those of one's peers, as well as for individual work performance. Next we consider a core antecedent of perceptual congruence, cultural transmission from peers, which we assess by examining how a quasi-exogeneous reorganization event changed a focal actor's exposure to peers who varied in their own cultural fit. We find support for our hypotheses using cross-sectional data. We also implement a novel machine learning method to impute perceptual congruence for individuals over time. Our longitudinal analyses corroborate our cross-sectional results and provide some insight into the mechanisms that underpin our observed effects. We discuss how these findings advance the theory of person-culture fit in organizations.

Introduction

Whether assimilating to a new organization or adapting to a new team, people typically seek to fit in culturally with their social groups. Above and beyond how technically qualified an employee is for a particular job, the benefits of cultural conformity, as well as the sanctions and penalties that come with failed cultural integration, are particularly stark in contemporary organizations. Indeed, research has consistently demonstrated that high levels of person-culture fit are associated with increased productivity, stronger commitment, and less turnover (Kristof-Brown et al. 2005, Alba and Nee 2009, Chatman and O'Reilly 2016). Moreover, employers have increasingly emphasized screening, selecting, and socializing new hires on the basis of person-culture fit rather than exclusively hiring for skills (Chatman 1989, Meyer et al. 2010, Rivera 2012). At the same time, as the average tenure in firms has declined (United States Bureau of Labor Statistics 2018), and as organizations constantly adapt their cultures to a shifting competitive environment (Denison et al. 2012), workers must frequently retool themselves culturally. In this dynamic context, how do people establish and maintain cultural fit with an organization?

Existing research offers a person-centered and fairly static answer to this question, focusing on how congruent a person's values are with those that prevail in an organization's culture (Chatman 1991, Kristof-Brown et al. 2005), or emphasizing stable personality and motivational orientations to adapting to new settings such as cultural intelligence and self-monitoring (Ang et al. 2006, Harrison et al. 1996). Prior research has tended to think of cultural fit as the degree to which a member has internalized the group's values, beliefs, and norms (Baron et al. 2001, Bourdieu and Passeron 1990, Chatman and O'Reilly 2016). This work has implicitly assumed that individuals who fit into their social environments think and act in ways that correspond to their values. Fitting in therefore implies having stable, often pre-existing, preferences that are consistent with the shared norms that prevail in an organization, or at least a consistent willingness to modify one's values and behavior to fit the culture. Indeed, a robust literature has demonstrated that *value congruence*—the match between a person's values and those that predominate and are normatively reinforced in the social

group (Chatman 1989, Fung et al. 2016)—predicts a variety of outcomes, from stronger attachment to the organization to higher levels of attainment.

Yet cultural fit reflects a dynamic interplay between one's private beliefs and one's public behaviors (Mobasseri et al. 2019). Even when people feel pressure to fit in, how they think and feel about their social group can differ substantially from how they behave when interacting with other members (Sekiguchi and Huber 2011). Thus, theories of person-culture fit based on value congruence are incomplete because they do not account for two common phenomena. First, people can maintain—often for long periods of time—dissonance between conforming to the organization's culture externally and having a different perspective about the culture internally (Salancik and Pfeffer 1978). Second, even when people are motivated to align their private beliefs with their public behaviors, they might easily misinterpret the organization's cultural norms. For example, swearing can, in certain normative contexts, connote camaraderie, while it can represent a form of symbolic violence in others (Willis 2014).

Social psychological and sociological theories have long emphasized how cultural norms shape action through situational cues (Dannals and Miller 2017, Swidler 1986, DiMaggio 1997). Although social constructivist approaches have not been explicitly linked to cultural fit, they have the broad orientation of shifting focus away from individuals' chronic preferences and personality traits to their exposure to relevant cultural information and their understanding of culture at a given point in time. This approach highlights how culturally congruent behavior arises as a response to people's exposure to situational norms and their understanding of those norms at a given point in time (Childress and Friedkin 2012, Salancik and Pfeffer 1978). Building on these insights, we broaden the theory of person-culture fit by highlighting the distinct and largely unexplored dimension of *perceptual congruence*, which we define as a person's understanding of a group's cultural norms at a given point in time. Unlike cultural intelligence and self-monitoring, which are stable personality traits, perceptual congruence can shift over time as a person learns a group's cultural code, as the code evolves, and as membership within the group changes. From this perspective, an employee's

decision to, for example, use deferential rather than aggressive language in a meeting may reveal little about her underlying preference for civil discourse and instead reflect the norms she observes in the behavior of other meeting participants (Goncalo et al. 2015).

We begin by defining perceptual congruence and distinguishing it from other constructs that pertain to deciphering culture. Whereas prevailing research on cultural fit has focused primarily on longer-term outcomes such as retention, we consider the implications of perceptual congruence for short-term behavior. In particular, we propose that perceptual congruence is manifested through behavioral fit—defined as the correspondence between one's observable behaviors in everyday interactions with colleagues and the group's normative expectations. Next we argue that perceptual congruence is positively related to an individual's work performance. Finally, we examine how changes in one's exposure to peers who vary in behavioral fit can influence the accuracy of one's perceptions, thereby shifting one's own perceptual congruence and conformity with the cultural code. This dynamic perspective on person-culture fit contrasts with the predominant view in the literature, which focuses on value congruence as the alignment between the prevailing cultural code and one's stable internalized values and beliefs. In developing these ideas, we make two contributions to advancing the theory of person-culture fit. First, we introduce the construct of perceptual congruence and identify its specific behavioral manifestations and downstream consequences for work performance. Second, we highlight a core mechanism through which perceptual congruence arises: interpersonal cultural transmission that results from an individual's exposure to cultural cues given off by others.

To evaluate the antecedents and consequences of perceptual congruence, we employ a multi-method empirical strategy using survey data, eight years of internal email data, and personnel records from a mid-sized technology firm. We begin by developing a novel measure of perceptual congruence based on the Organizational Culture Profile (OCP) (Chatman et al. 2014), a validated culture assessment. In particular, we compare an individual's report of widespread cultural norms to the ones her peers believe are predominant. We then examine how perceptual congruence is

manifested through behavioral fit, applying the interactional language-use model to a corpus of employee emails (Srivastava et al. 2018, Goldberg et al. 2016). Specifically, we examine the relationship between perceptual congruence and the similarity of linguistic style between employees and the colleagues with whom they communicate.

We report 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 perceptual congruence for individuals over time. This method enables us to trace within-person changes and identify a core antecedent of perceptual congruence: the transmission of culture from peers to a focal actor. We assess interpersonal cultural transmission by examining a restructuring event at our research site that resulted in a quasi-exogeneous shift in employees' peer groups. This shift exposed employees to peers who varied in their behavioral fit. We adapt an established instrumental variables technique to identify the effects of this change on the focal individual's own perceptual congruence and behavioral fit (Waldinger 2012). Our longitudinal analyses corroborate the cross-sectional results and also provide some insight into the mechanisms that underpin our observed effects. We conclude by discussing how our findings advance the theory of person-culture fit in organizations.

Theory and Hypotheses

Defining Perceptual Congruence

Previous research on person-culture fit has focused on value congruence and examined how it relates to behaviors such as voluntary exit (Morrison 2002, Sheridan 1992). This work has implicitly assumed that value congruence is the key pathway to behaving in ways that are similar to others in the organization. This approach does not, however, consider that people can possess the knowledge needed to successfully interact with one another even when they do not share the same values (Hewlin 2003, Hochschild 2012, Mobasseri et al. 2019). In other words, people can be incongruent in values and still find ways to fit, or conversely, they can share values but fail to act appropriately (Goffman 1959). To help explain how this is possible, we introduce the construct of perceptual

congruence, which reflects a person's understanding of the group's cultural code at a given point in time.

Consider the example of a colleague making a cynical joke during a meeting. Such a behavior can be interpreted as a friendly attempt to establish rapport or as a derogatory comment aimed at undercutting others. One's capacity to correctly construe the meeting as friendly or adversarial depends on one's tacit understanding of the organization's cultural code, irrespective of one's ideal preferences. The value congruence tradition would typically elicit one's preferences—for example, for an aggressive or conflict-avoiding organizational culture—and determine that one is congruent insofar as one's values match the normative environment. The perceptual congruence approach, in contrast, focuses on the extent to which one correctly understands the behaviors one observes and conceptualizes congruence as the consistency between one's interpretation of behavioral norms—for example, the tendency to joke in meetings—and those of one's peers.[1]

Before discussing the consequences and a key antecedent of perceptual congruence, we distinguish it 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.

A thought experiment helps distinguish perceptual congruence from cultural intelligence and self-monitoring: Two newcomers, A and B, who have equal levels of cultural intelligence or who are equal in their levels of self-monitoring, enter an organization that is homogeneous with respect to race, ethnicity, and nationality. If A and B are identical except that they are exposed to different peers, the cultural intelligence and self-monitoring perspectives would predict no significant differences in their capacity to decipher this organization's cultural code, irrespective of where in the organization they are situated. In contrast, if A is exposed to peers within the organization whose behavioral fit is higher than the behavioral fit of the peers B is exposed to, the perceptual congruence approach would predict that A would be more accurate than B in deciphering the organization's cultural code. In other words, although perceptual congruence focuses on understanding cultural norms and fit in a similar manner as cultural intelligence and self-monitoring, perceptual congruence is also distinctly and intimately related to one's exposure to peers' perceptions and behaviors. Further, while cultural intelligence and self-monitoring arise from a stable personal interest in learning how to adapt to new cultures and situations, perceptual congruence is based more on the quality of the evidence that one is presented with and how closely the behaviors displayed by peers fit with the

culture. This implies that the quality of the social learning that gives rise to perceptual congruence depends not only on the student but also on exposure to the behaviors of the "teachers," that is, the peers from whom she learns.

We turn next to considering the consequences of perceptual congruence. Most person-culture fit research has focused on distal outcomes such as departure and commitment. In contrast, we begin by examining the link between perceptual congruence and a more proximal outcome: behaving in ways that match peers' normative expectations. Then we turn to perceptual congruence's more distal consequences for work performance.

The Consequences of Perceptual Congruence

Perceptual Congruence and Behavioral Fit. We define behavioral fit as the extent to which an individual's behaviors are compliant with the group's normative expectations (Leonardelli et al. 2010, Goldberg et al. 2016). In general, people seek to behave in ways that are normatively compliant with their social groups because deviations from group norms are typically met with various forms of sanctions (Coleman 1990, Heckathorn 1990, Chatman and O'Reilly 2016). Yet, even when individuals are motivated to fit in, they may simply lack an understanding of how to do so, particularly given that they may belong to myriad social groups within an organization. Attempts to fit in are likely to flounder when employees are culturally miscalibrated—that is, when they exhibit low perceptual congruence.

As a concrete illustration of the connection between perceptual congruence and behavioral fit, consider 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 to reaffirm their interpersonal bonds and their commitment to the bank. To participate in this ritualistic complaining, bank employees 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. Consistent with this depiction, we argue that

perceptual congruence will be intimately related to the extent to which one exhibits behavior—in routine, real-time interaction—that conforms to the organization's prevailing norms.

We use language alignment as evidence of behavioral fit because it focuses on real-time, fine-grained, and routine interaction among peers and thus represents one of the most proximate behavioral manifestations of perceptual congruence. The language one uses when interacting with one's peers provides a window into what one thinks and how one feels (Pinker 2007, Goldberg et al. 2016, Schroeder and Epley 2015). 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 accurate understandings of the cultural code will, all else equal, behave in more normatively compliant ways. Thus, we anticipate a positive relationship between perceptual congruence and behavioral fit:

Hypothesis 1 (H1): *Perceptual congruence will be positively related to behavioral fit.*

Perceptual Congruence and Work Performance. Next we consider the implications of perceptual congruence for individual performance. Prior work has shown that behavioral fit enhances performance by enabling people to better coordinate activity and perform interdependent tasks. For example, Weber and Camerer (2003) used laboratory studies to demonstrate that performance on an interdependent task declines when group members operate through the use of divergent cultural conventions. Similarly, Goldberg et al. (2016) and Srivastava et al. (2018) showed that a behavioral (language-based) measure of cultural fit was positively related to individual career attainment (favorable performance ratings and time-to-promotion). Thus, when people behave in ways that conform to group norms, they are better able to coordinate activity and perform interdependent tasks with their peers (Sørensen 2002, Corritore et al. 2020).

We propose that perceptual congruence will have a positive relationship with work performance independent of a person's behavioral fit. This is because perceptual congruence not only enables people to behave in normatively aligned ways but also to define problems and find solutions using approaches that are valued in the culture. To develop this argument, we draw on Dougherty

(1992), who describes the interpretive barriers that tend to splinter complex organizations into divergent "departmental thought worlds." As a result of this fragmentation, people in one part of the organization often struggle to understand the values, norms, and beliefs in other areas, which in turn impairs their ability to synthesize knowledge, identify cross-cutting problems, and develop novel solutions. Because perceptually congruent individuals have more accurate mental representations of their organization's culture, we propose that they are more likely to overcome such interpretive barriers and produce outputs that will be more highly valued in the overall organizational culture. In other words, perceptual congruence affects performance not only by changing *how* people coordinate with others but also *what* ideas they choose to focus on developing.[2] Thus, we anticipate:

Hypothesis 2 (H2). *Perceptual congruence will be positively related to individual work performance.*

Having laid out a key behavioral manifestation of perceptual congruence and its consequences in terms of work performance, we turn to the question of how perceptual congruence emerges, focusing on how exposure to others influences one's own perceptual congruence. We refer to this process as the interpersonal transmission of culture.

Antecedents to Perceptual Congruence: The Interpersonal Transmission of Culture

Our theory of perceptual congruence is rooted in the notion that culture is learned from others (Levine and Moreland 1991). Building on this insight, we propose that the composition of one's network—that is, the people to whom one is exposed at a given point in time—has a significant impact on the extent to which one can decipher the cultural code and exhibit normatively compliant behavior. Experimental research in young children, for example, demonstrates that exposure to multiple and consistent behaviors increases the fidelity and speed of cultural transmission (Hermann et al. 2013). Similarly, in the workplace, employees' susceptibility to being influenced by others is related to the kinds of colleagues with whom they interact (Chan et al. 2014, Liu and Srivastava 2015). In particular, having colleagues who themselves have a more accurate read of

the cultural environment can help correct one's own misperceptions, thereby improving one's own perceptual congruence (Balkundi and Kilduff 2006).

In contrast to values, which are remarkably stable once one has joined the workforce (Jin and Rounds 2012), perceptions are more susceptible to ongoing social learning (Bandura and Walters 1977, Shepherd 2017). One learns about and forms perceptions of an organization's culture by observing other members' behavior (Levine and Moreland 1991). Even though an individual does not have direct access to others' perceptions, being exposed to colleagues' behavior, particularly over repeated interactions (Beer and Watson 2010), influences one's own perceptions and behavior. Support for this view comes from the extensive literature on diffusion processes, which demonstrates that a person's attitudes can change as a direct consequence of exposure to and interaction with their network contacts (Friedkin and Johnsen 1990, Marsden and Friedkin 1993, Baldassarri and Bearman 2007). Similarly, expectations of normatively appropriate behavior are shaped by perceptions that arise through interaction, observation, and imitation (Friedkin 2001, Chatman et al. 2014, Liu and Srivastava 2015).

Building on these insights, we propose that one draws inferences about the organization's culture based on one's observations of others' behavior. When one is exposed to peers who are well-calibrated and thus behave in normatively compliant ways, one's own understanding of the culture improves. In contrast, when one is exposed to poorly calibrated peers who exhibit low behavioral fit, one's cultural understanding not only does not improve but can even erode if one mistakenly assumes that peers' behavior is normatively compliant. Following the logic of Hypothesis 1 above, these changes in perceptual congruence then translate into corresponding shifts in behavioral fit. We refer to this process as interpersonal cultural transmission.

One behavior that offers ample opportunity for social learning is language-based interaction, such as that which occurs in email exchanges (Srivastava and Goldberg 2017). Language is a behavioral manifestation of culture in that linguistic alignment with peers signals the degree to which an individual is normatively compliant (Pettigrew 1979, Doyle et al. 2017). One observes

others' language and then deduces what they are thinking about the culture, inferring more nuanced information from not only what others say about culture but also how they say it (Schroeder and Epley 2015, Schroeder et al. 2017). These deductions are susceptible to ongoing social learning and behavioral imitation. Thus, we anticipate that one's perceptions of the cultural order, as reflected in perceptual congruence, as well as one's behavioral fit, will be susceptible to social influence through interpersonal cultural transmission. More formally, we expect that:

Hypothesis 3 (H3). *Perceptual congruence and behavioral fit will be susceptible to peer influence. Specifically, as one's peers behave in more (less) normatively compliant ways, one's own perceptual congruence will increase (decrease) and one's behavioral fit will concomitantly increase (decrease).*

Figure 1 provides a visual representation of our theory of perceptual congruence and how the three hypotheses fit together.

[FIGURE 1 ABOUT HERE]

Method

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 completed an Organizational Culture Profile (OCP) (Chatman et al. 2014) describing the organization's current culture. As described in greater

detail below, we used these responses to construct our measure of perceptual congruence. To be able to compare how perceptual congruence relates to the established construct of value congruence, we also asked a randomly selected half of employees to complete the survey based on their own personally desired cultural characteristics. Combining the former and latter sets of responses enabled us to measure value congruence for a subset of employees.[3] Overall, we received 440 completed surveys about the current organizational culture and 238 completed surveys about the 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, and a major reorganization, which we use for our instrumental variables analysis, took place in mid-2015. 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

Behavioral Fit—Hypotheses 1 and 3 focus on behavioral fit, which we operationalized 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. 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 (Brugnoli et al. 2019)—we focus on the similarity of *linguistic style* between an individual and her reference group.[4] 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 behavioral fit, 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 behavioral cultural fit.

Intuitively, when the outgoing and incoming distributions are nearly identical, the divergence approaches zero, suggesting high behavioral cultural fit; conversely, greater deviation between the probabilities of usage of LIWC categories translates to greater divergence and thus implies lower behavioral cultural fit. 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 behavioral cultural fit 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.

Work Performance—To evaluate Hypothesis 2, 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, which uses a comprehensive set of cultural elements that have been applied to and validated in a wide variety of organizations, 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 behavioral fit, 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.[5]

Value Congruence—Although we do not theorize about value congruence, we sought to understand how perceptual congruence relates to this well-established construct that has also been derived from the OCP in prior research. For value congruence, we focused on participants' 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?" Recall that this question was asked of only half the respondents; hence, value congruence is computed for only a subset of our sample. 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 behavioral fit.

Imputing Perceptual Congruence and Value Congruence Over Time

The procedure above creates cross-sectional measures of perceptual congruence and value congruence. Models based on such measures cannot account for time-invariant, unobserved heterogeneity—for example, stable differences in cultural intelligence or other traits and dispositions that might be related to perceptual congruence and our outcomes of interest. In addition, to test Hypothesis 3, which involves assessing how perceptual congruence and ultimately behavioral fit change when an individual is exposed to a different set of peers due to an organizational restructuring event, we need a time-varying measure of perceptual congruence.

For both of these reasons, we undertook a procedure to transform our cross-sectional measures of perceptual congruence, as well as value 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 behavioral cultural fit) 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 2 provides a conceptual overview of this procedure. Further procedural details are provided in the Appendix B; evaluative analyses regarding model fit are provided in Appendix C.

[FIGURE 2 ABOUT HERE]

Peer Perceptual Congruence, Value Congruence, and Behavioral Fit

Given that Hypothesis 3 focuses on peer influence, we first identified a focal actor's i's communication partners J for each month T. We defined peer behavioral fit as the mean of behavioral fit across all of i's communication partners, weighted by the volume of incoming communication received from each interlocutor. We defined peer perceptual congruence and peer value congruence in similar fashion but based on imputed measures given that perceptual congruence and value congruence were measured at just one point in time.

Control Variables

We estimated both within-person and between-person models for our analyses. In within-person models, time-invariant effects (e.g., the role of diffuse status characteristics such as gender and ethnicity) are subsumed by individual fixed effects; however, 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).[6] For our between-person models, we included additional control variables for age, gender, and self-monitoring orientation.

Analytical Approach

We tested Hypothesis 1, which posits that perceptual congruence will be positively related to behavioral cultural fit, using OLS regressions based on cross-sectional data, as well as fixed effect regressions based on longitudinal data (including imputed measures of perceptual congruence and value congruence). 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.

For Hypothesis 2, we estimated ordinary least squares regressions of each focal individual's bonus on the three cultural fit measures. These models included both department and individual fixed effects.

Finally, to test Hypothesis 3, which suggests that perceptual congruence and behavioral fit will be susceptible to peer influence, we identified the effect of changes in peer composition on the focal individual's cultural fit measures. We began by estimating the following basic OLS model, with individual, department and year fixed effects:

$$BF_{idt} = \beta_0 + \beta_1 \langle PeerBF \rangle_{idt-1} + \beta_2 |Peer|_{idt-1} + \eta X_{idt-1} + \beta_3 Year_t + \beta_4 Dept_d + \beta_5 Ind._i + \epsilon_{idt} \quad (1)$$

where BF_{idt} is behavioral fit for individual i in department d at time t, $\langle PeerBF \rangle_{idt-1}$ is peer behavioral fit at time $t-1$ weighted by the number of incoming messages, $|Peer|_{idt-1}$ is the number

of peers at time $t-1$, and X are time-varying individual attributes. As noted above, these models include individual and departmental fixed effects. We lag the peer behavioral fit and number of peers to ensure appropriate temporal ordering.

Instrumental Variables Regression—Even with individual fixed effects and lagged predictors, the modeling approach described above to test Hypothesis 3 does not yield causal estimates. It could be the case, for example, that individuals with high behavioral fit seek to interact with equally culturally integrated individuals. In other words, this modeling approach cannot separate the effects of homophily from those that arise through peer influence.

To help address this problem, we exploited a reorganization event that transpired over a period of two months, roughly seven years after the firm's founding. An ideal test would have included a fully exogenous event that assigned certain individuals to interact with a random set of new peers while others retained their previous network contacts. Such a natural experiment would allow for causal identification of peers' cultural fit on that of the focal individual. In the absence of such an experiment, we relied on this reorganization event, which—although not random—was driven primarily by functional needs arising from rapid growth at the time and which affected all employees to some extent. Moreover, unlike network changes generated by downsizing, the restructuring did not disproportionately affect low-performing or otherwise systematically similar peers.

As such, the reorganization can be thought of as quasi-exogenous in that it introduced significant random variation in employees' network compositions. Recognizing, however, that this event was not a pure natural experiment, we used an extension of an instrumental variable peer effects model first introduced by Waldinger (2012). Using a two-stage least-squares model, we first estimated the random variation in mean peer behavioral fit and number of peers introduced by the reorganization, and we then used these estimates to predict subsequent changes in a focal actor's behavioral fit and perceptual congruence.

In typical instrumental variable designs, the instrument is assumed to only affect the endogenous variable. In the present case, however, the reorganization also affected the focal individuals' peers'

network compositions. Thus, peers also experienced shifts in their behavioral fit, driven by changes in their own peer group after the reorganization and social influence from peers in the month of reorganization. To address this complexity, we follow Waldinger (2012) and use *induced change in peer behavioral fit*, $\tilde{\Delta}\langle PeerBF\rangle$, as an instrument. $\tilde{\Delta}\langle PeerBF\rangle$ is the change induced by the reorganization between periods $t-1$ and t, *assuming peer behavioral fit had remained fixed at its pre-reorganization level*. Defining the measure in this way allowed us to account for the change in peer exposure stemming from the reorganization, while separating out its downstream effects on peers' behavioral fit.

In addition to induced change in mean peer behavioral fit, we also measured the magnitude of change in network composition as an instrument. Let I_{it} be a vector of length N (total number of employees) wherein each cell $I_{it}(j)$ corresponds to the number of messages that i received from interlocutor j during month t. We define i's network change at time t as the cosine distance between i's vectors of incoming messages in two consecutive months:

$$NC(I_{it}, I_{it-1}) = \cos(I_{it}, I_{it-1}) \tag{2}$$

where the cosine distance between two vectors p and q is defined as:

$$\cos(p,q) = 1 - \frac{\sum_{j=1}^{N} p(j)q(j)}{\sqrt{\sum_{j=1}^{N} p(j)^2}\sqrt{\sum_{j=1}^{N} q(j)^2}} \tag{3}$$

Because the number of messages is non-negative, this measure is bounded by 0 and 1.

We used these instruments—network change, induced change in peer behavioral fit, and the interaction between the two—to estimate the model's two endogenous variables, mean peer behavioral fit and number of peers. In the first stage we estimated the following regressions:

$$\begin{aligned}\langle PeerBF\rangle'_{idt} = \beta_0 + \beta_1 NC(I_{it}, I_{it-1}) + \beta_2 \tilde{\Delta}\langle PeerBF\rangle_{idt-1} \\ + \beta_3 NC(I_{it}, I_{it-1}) \cdot \tilde{\Delta}\langle PeerBF\rangle_{idt-1} + \beta_4 Ind._i + \epsilon_{it}\end{aligned} \tag{4}$$

$$\begin{aligned}|Peer|'_{idt} = \beta_0 + \beta_1 NC(I_{it}, I_{it-1}) + \beta_2 \tilde{\Delta}\langle PeerBF\rangle_{idt-1} \\ + \beta_3 NC(I_{it}, I_{it-1}) \cdot \tilde{\Delta}\langle PeerBF\rangle_{idt-1} + \beta_4 Ind._i + \epsilon_{it}\end{aligned} \tag{5}$$

In the second stage we estimated behavioral fit at time $t+1$ (a month after the reorganization) with instrumented mean peer behavioral fit and number of peers as independent variables. These models included individual, department, and year fixed effects. We specified the second stage regression as:

$$BF_{idt+1} = \beta_0 + \beta_1 \langle PeerBF \rangle'_{idt} + \beta_2 |Peer|'_{idt} + \beta_3 Year_t + \beta_4 Dept_d + \beta_5 Ind._i + \eta X_{it} + \epsilon_{idt} \qquad (6)$$

where X_{it} represents time-varying individual controls. We report results from eq. 6 in Model 2 of Table 3 below. Given that Hypothesis 3 focuses on the role of social learning in shifting behavioral fit and perceptual congruence, we followed a similar procedure to estimate how peer behavioral fit relates to a focal actor's perceptual accuracy in the following period. This result appears in Model 3 of Table 3.

Our instrumental variables set-up departs from Waldinger's (2012) in at least two fundamental ways. First, whereas Waldinger focuses on how shifts in peer quality and number of peers that arise from exogenous dismissals (of scientists) affect a focal actor's own quality, in our set-up a focal actor experiences shifts in peer behavioral fit and number of peers through not only the removal of existing peers, but also the addition of new peers in her communication network. For the exclusion restriction to hold, we must assume that the processes by which removals and additions occur are orthogonal to the focal actor's behavioral cultural fit. There are reasons to doubt that this would be the case in a typical reorganization. For example, the architects of the reorganization could have moved people around based on expectations of who would fit better with whom and who would likely increase their fit over time. We partially account for this possibility through a robustness check (described below) that excludes senior employees, whose anticipated fit was more likely to be a factor in architects' specific reorganization choices. Another way in which the exclusion restriction would be violated is if people actively changed their interaction partners during the reorganization, with an eye to anticipated cultural fit in the post-reorganization regime. As we will discuss in greater detail below, our empirical results (see Figure 3) suggest that people do indeed adjust their interaction partners following the reorganization; however, they do so after a lag of over a month.

Thus, in the relevant time period for our instrumental variable regressions, it seems less likely that the exclusion restriction is violated because of self-selection into networks.

Second, unlike Waldinger's case where dismissals are unambiguous events, shifts in peer exposure in our setting are more ambiguous because we only observe the peers to which a focal actor is connected in the email communication network. Thus, we have no way of accounting for changes in peer exposure that occur in other communication media (e.g., phone calls, text messages, face-to-face meetings). In sum, we believe that our instrumental variable regressions improve significantly upon our OLS estimates and, with some caveats, move us closer to estimating the causal influence of peer exposure on perceptual congruence.

Results

Preliminary Analyses

Although our theory does not focus on the content of organizational culture, we used the OCP data to gain insight into the cultural features of our research site. This analysis, the details of which appear in Appendix D, revealed that the culture was highly attuned to customers, results, and integrity, while it focused less on transparency and attention to detail.

Next, we sought to demonstrate the validity of the imputed measures of perceptual congruence and value congruence that were developed using the procedure described in Appendix B. Given that value congruence is relatively stable, while we theorize perceptual congruence to be more malleable, 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 both value congruence and perceptual variables using indicators for each month (up to month 36 of employment). These results are depicted in Figure 3.

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 are consistent with our expectations about temporal variation in these two measures and thus help to validate our imputation procedure.

[FIGURE 3 ABOUT HERE]

Main Results

Table 1 provides a test of Hypothesis 1. The first three models report cross-sectional results, while the next three report longitudinal ones. In support of Hypothesis 1, Model 1 shows a positive and significant relationship between perceptual congruence and behavioral fit. Model 2 demonstrates that value congruence is not significantly related to behavioral fit, and Model 3 reveals that the positive association between perceptual congruence and behavioral fit remains even when controlling for value congruence.

[TABLE 1 ABOUT HERE.]

Models 4 to 6 show the same pattern of results using longitudinal specifications that include individual, department, and year fixed effects. Across both cross-sectional and longitudinal models, we find support for Hypothesis 1: As individuals' perceptual congruence increases, so does their behavioral fit.

Of the control variables included in the models, only managerial status and tenure are significant.[7] Consistent with previous work on enculturation (Srivastava et al. 2018), we also find that individuals exhibit significantly lower behavioral cultural fit during their first year in the organization.[8]

To test Hypothesis 2, which anticipates a positive relationship between perceptual congruence and work performance, we estimated OLS regressions of monthly bonuses on lagged variables of interest with individual, department, and year fixed effects. We report these results in Table 2. In support of Hypothesis 1, Model 1 shows that perceptual congruence is positively related to work performance. Consistent with prior work (Srivastava et al. 2018), Model 2 finds a positive association between behavioral fit and performance. In Model 3, value congruence is also positively

tied—albeit more modestly—to performance. Importantly, Model 4 shows that the positive relationship between perceptual congruence and work performance holds even when controlling for behavioral fit and value congruence.

[TABLE 2 ABOUT HERE.]

Table 3 reports the analyses we used to test Hypothesis 3—that being connected to colleagues with higher (lower) behavioral fit will be associated with corresponding increases (decreases) in perceptual congruence and hence behavioral cultural fit for the focal individual. Model 1 presents estimates from the baseline fixed effect models with lagged peer behavioral fit, as specified in equation 1. Individuals exhibit a significant increase in behavioral fit when their peers' mean behavioral fit increases in the preceding month.

[TABLE 3 ABOUT HERE.]

As noted above, the estimates from Model 1 are not causal given that this empirical approach cannot distinguish the effects of homophily, or seeking out similar others, from those of social influence, or modifying one's own behavior to accommodate others' behavior. We therefore turn to our instrumental variable in the remaining models. The primary results are reported in Models 2 and 3. The coefficient for peer behavioral fit suggests that those who, as a result of the reorganization, transitioned into a network comprising peers with greater behavioral fit experienced an increase in their own behavioral fit in the following month. The opposite is also true: individuals who, through the reorganization, transitioned into a network of peers with lower behavioral fit experienced a corresponding decline in their own behavioral fit. Interestingly, and likely because reorganizations are disruptive to cultural integration, the majority of employees experienced a decline in peer behavioral fit, and correspondingly, their own behavioral fit during this period.

We illustrate the implications of induced change in peer behavioral fit in Figure 4. The diagram plots the effects of the reorganization on individuals' behavioral fit over time, as estimated by the instrumental variable model. The upper line corresponds to individuals who experienced a half standard deviation positive increase in their peers' behavioral fit, and the lower line corresponds to

individuals who experienced a decline of the same magnitude in their peers' behavioral fit. These are substantial changes in peer behavioral fit but not implausible during a period of reorganization. A little over 1% experienced a positive shock at or greater than half a standard deviation, but roughly 35% experienced a decline of that magnitude. Both translate to similarly sized adjustments in the focal individuals' behavioral fit, but in opposite directions. Moreover, both adjustments persisted for roughly two months, after which the effects of the reorganization were no longer apparent and individuals converged toward mean behavioral fit.

Because the reorganization was not a true natural experiment, it is worth noting that changes that occurred after its effects were initially felt could have arisen for a variety of reasons that we do not observe in our data. For example, individuals presumably regained more command over whom they interacted with after the reorganization, which would also reintroduce potentially confounding homophily effects. Hence, the period immediately following the reorganization is the appropriate one to consider for this analysis.

Importantly, the two sets of individuals—positively and negatively "treated"—are indistinguishable in the period preceding the reorganization, suggesting that these adjustments are a result of the imposed change in network composition rather than systematic differences between the two groups. The Kleibergen-Paap F statistic, which is appropriate when using robust standard errors, suggests that the instrument is strong (Kleibergen and Paap 2006, Baum et al. 2007).

Changes in the number of peers had a more modest impact: those who experienced an increase in the size of their network due to the reorganization experienced declines in behavioral fit. Forced network growth, in other words, appears to be disruptive to cultural integration. The difference between these coefficients in the OLS (Model 1) and instrumental variable (Model 2) models is worth noting and suggests that our instrumental variable approach at least partially addresses the endogeneity inherent in our OLS models. During non-turbulent times (Model 1), an increase in number of peers is associated with an increase in behavioral fit. Our results suggest, however, that the increase in network size is driven by improved cultural integration, which facilitates seeking

out more contacts in the organization, and not the other way around. When changes are forced, in contrast, attending to a growing number of peers whom the focal individual does not necessarily choose to interact with appears to undermine cultural adjustment (Model 2).

Our models do not speak directly to how precisely this cultural transmission occurs—for example, whether organizational members explicitly reward and penalize their colleagues for culturally compliant or deviant behavior or whether cultural knowledge is transferred tacitly. Model 3—wherein we estimate the effects of change in peer behavioral fit on the focal individual's perceptual congruence—suggests that behavioral adjustment occurs through changes in perceptual congruence. Overall, Models 1-3 provide support for Hypothesis 3.

Model 4, which estimates the effects of peer behavioral fit on the focal individual's value congruence—suggests that value congruence is not involved in this process of social learning. This suggests that individuals adapt their perceptions, but not their private beliefs, in response to changes in peer composition. Moreover, in Models 5 and 6 we estimate the effects of reorganization-driven changes in peer perceptual congruence and in peer value congruence on the focal individual's perceptual congruence and value congruence, respectively. Neither coefficient is significant, lending further support to our argument that cultural learning occurs through observing peers' behaviors, given that cognition is less directly accessible to others. We suspect that the majority of this cultural transmission happens tacitly. As Models 5 and 6 imply, individuals generally do not have access to their peers' perceptions or underlying values. Given that peer behavioral fit is positively related to one's perceptual congruence (Model 3) but not to one's value congruence (Model 4), the cultural transmission we observe appears to be driven by the diffusion of *perceptions* about the culture rather than the diffusion of *values*.

In Appendix E, we report the results of two supplemental analyses designed to assess the robustness of the results of our instrumental variables analysis. First, given that our measures of interest are all defined with respect to the reference group of an individual's interlocutors in a given month, which people can—to varying degrees—self-select into, we replicated the instrumental variables

analysis using behavioral fit, peer behavioral fit, and peer perceptual congruence measures that were based on the reference group of *all* employees in the organization. Second, to address the possibility that the reorganization did not produce exogenous shifts for all employees, we replicated the analyses using a sub-sample of employees who were not in supervisory roles. We reasoned that, insofar as the reorganization was deliberately intended to change particular peer groups, such social engineering was most likely targeted to the leadership ranks of the company and not to individual contributors. The results of the robustness checks reported in Appendix E are consistent with our main findings.

Discussion and Conclusion

The speed and frequency of cultural change in contemporary organizations, coupled with the increasing importance of person-culture fit, requires that employees regularly decipher and adapt to new cultural orders. Social psychological and sociological research has long shown that people can maintain dissonance between values and actions, enabling them to enact behavioral fit even when their values are incongruent with the culture. This observed reality strains existing, mostly static theories of person-culture fit that are based solely on value congruence and suggests the need for more dynamic and less trait-based orientations to understanding person-culture fit. With this gap in mind, we developed and tested a theory of perceptual congruence, which describes a person's understanding of their group's culture at a particular point in time. Perceptual congruence is (1) manifested through fine-grained, contemporaneous demonstrations of behavioral fit such as interactional language use; (2) enables higher levels of work performance; and (3) is influenced by, and emerges in part from, the behavioral fit of peers to which a person is exposed.

Our theoretical framework and empirical findings offer three key contributions. First, we advance person-culture fit theory by demonstrating that a person's perceptual congruence is influenced by peers' understanding of the culture and their behavioral demonstrations of fit. Our findings from a quasi-exogenous restructuring event in which employees were exposed to different interaction partners who varied in their own cultural fit suggest that a person's perceptual congruence changes

through social learning—that is, by observing peers' behavior and drawing inferences about their underlying perceptions of the culture. Given how quickly people adjusted their behavioral fit right after the reorganization, it seems unlikely that these changes were driven by shifts in value congruence. Indeed, values are known to change slowly, if at all, in adulthood—a point underscored by the relative stability of our imputed value congruence measure. Instead, our results suggest that fitting into an organization's culture involves a more dynamic process in which people decipher the cultural code through real-time and ongoing observations of their peers' behavior.

Within person-culture fit research, our findings also advance theories of organizational socialization. Whereas prior work has highlighted the role of intentional socialization practices such as formal onboarding programs (Cable et al. 2013), we demonstrate how less intentional factors such as the colleagues to which one is assigned or incidentally exposed to can also influence cultural learning. In addition, while prior work on the role of networks in socialization has focused on the initial period when newcomers first enter a new organization (Morrison 2002), our results suggest that structural factors—in particular, shifts in one's network interaction partners—can continue to influence social learning throughout an individual's tenure in an organization.

More specifically, we identify the factors that cause some people to enculturate more successfully than others and illuminate the role of social networks in cultural transmission. Previous work has often assumed that enculturation is a function of individual differences in endowments. Rivera (2012), for example, demonstrates that labor market matching—at least in the elite firms she investigates—is inherently related to the cultural capital that job applicants possess. Separately, research by organizational psychologists has focused on innate differences in psychological traits, demonstrating that stable dispositions such as self-monitoring and cultural intelligence are conducive to cultural adjustment and the benefits it confers (Maddux et al. 2008). In contrast, we use an instrumental variable approach to show that the ability to enculturate is also contextual (cf. Ashforth et al. 2007), accruing to individuals whose peers are themselves successfully enculturated. Cultural adaptation, in other words, is not just a function of the ability to decipher the cultural

code but is also derived from the peers whom one is exposed to and their understanding of the culture.

Our second contribution lies in demonstrating that person-culture fit manifests not only in longer-term outcomes such as voluntary exit but also in more fine-grained behavior. We demonstrate, in particular, that one's perceptual congruence is closely tied to one's linguistic alignment with colleagues over much shorter timescales than prior research has examined. Further, we collected data unobtrusively, from email exchanges, thereby reducing the likelihood that linguistic alignment arose in response to social desirability biases. This avoids a common confound in research on person-culture fit, which is that people seek to appear to fit with their organization's culture (Caldwell et al. 2008). Future research might consider using other unobtrusive indicators of behavioral fit, including the language used in recruiting and performance reviews (Wynn and Correll 2018), adaptations in dress among employees (Rafaeli et al. 1997), patterns of scheduled meetings (Polzer et al. 2018), and adherence to timeliness norms (Dannals and Miller 2017). Examining these granular manifestations of behavioral fit, which had been difficult to capture until the advent of sophisticated computational methods and thus overlooked in much of the past research on person-culture fit, might also reveal conditions under which people not only adapt to existing norms but introduce cultural innovations that others subsequently adopt (O'Reilly and Chatman 2020).

In addition to demonstrating the link between perceptual congruence and the proximate outcome of behavioral fit, we show how perceptual congruence relates to the more distal outcome of work performance. Moreover, we demonstrate a direct link between perceptual congruence and performance, above and beyond that which arises through increased behavioral fit. In short, perceptual congruence enables people to reap positive career rewards: Those who read the code correctly and behave accordingly benefit from being perceived as true and committed group members.

Our third contribution is the development of a novel method that may have wide-ranging application across the social sciences. 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, is only appropriate in certain contexts. Requirements for this approach include: having a sufficient number of survey observations, the availability of rich communication content, the existence of protocols and safeguards to protect individual privacy and company confidentiality, and significant computational bandwidth. In addition, more work—based on multiple administrations of self-reports alongside contemporaneous communications data—is needed to identify the conditions under which the relationships between language use and the self-reports of interest are (as we assume in this paper) stable over time. Yet, 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 the perceptual and behavioral arenas of organizational life relate to one another and jointly shape individual careers and the organizational cultures in which they unfold.

Endnotes

[1]In principle, perceptual congruence could be related to value congruence–for example, if value congruent individuals were also more motivated to learn the organization's cultural code and developed a more accurate understanding because they invested greater effort in doing so. Yet, as we elaborate below, a highly motivated individual may still fail to learn the cultural code if she is mostly exposed to peers who are behavioral misfits. Thus, perceptual congruence is analytically distinct from value congruence.

[2]The link between perceptual congruence and work performance may be stronger for certain job roles—for example, product development managers who must coordinate activities across departments and functions—than for others. Future research identifying the sources of variance in this relationship could be useful for establishing the boundary conditions of our theory. Empirically,

our models account for this variance by including department fixed effects—though this does not address potential variation across roles within a department. Our longitudinal models, however, include individual fixed effects, which account for all role-based differences that are time-invariant.

[3]The other half completed a survey of the cultural characteristics needed for the organization to be successful in the future. We shared the results of this latter survey 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]In the framework developed by Lockwood et al. (2019) to link language to management action, email exchanges between colleagues can be thought of as a form of discourse that relates to action outcomes such as performance and social evaluations.

[5]For robustness checks reported below, we also produced versions of these measures in which the reference group included all employees in the organization rather than just the focal individual's email interaction partners in a given month.

[6]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.

[7]We conjecture that managers exhibit greater behavioral cultural fit 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.

[8]Tenure has a curvilinear relationship with behavioral fit, 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.

References

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

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.

Ashforth BE, Sluss DM, Saks AM (2007) Socialization tactics, proactive behavior, and newcomer learning: Integrating socialization models. *Journal of Vocational Behavior* 70(3):447–462.

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

Baldassarri D, Bearman P (2007) Dynamics of political polarization. *American Sociological Review* 72(5):784–811.

Balkundi P, Kilduff M (2006) The ties that lead: A social network approach to leadership. *The Leadership Quarterly* 17(4):419–439.

Bandura A, Walters RH (1977) *Social learning theory*, volume 1 (Prentice-hall Englewood Cliffs, NJ).

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.

Baum CF, Schaffer ME, Stillman S, et al. (2007) Enhanced routines for instrumental variables/gmm estimation and testing. *Stata Journal* 7(4):465–506.

Beer A, Watson D (2010) The effects of information and exposure on self-other agreement. *Journal of Research in Personality* 44(1):38–45.

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

Bourdieu P, Passeron JC (1990) *Reproduction in Education, Society and Culture*, volume 4 (Sage).

Brugnoli E, Cinelli M, Zollo F, Quattrociocchi W, Scala A (2019) Lexical convergence inside and across echo chambers. *arXiv preprint arXiv:1903.11452* .

Cable DM, Gino F, Staats BR (2013) Breaking them in or eliciting their best? reframing socialization around newcomers' authentic self-expression. *Administrative Science Quarterly* 58(1):1–36.

Caldwell D, Chatman J, O'Reilly C (2008) Profile comparison methods for assessing person-situation fit. *Perspectives on Organizational Fit* 356–361.

Chan TY, Li J, Pierce L (2014) Compensation and peer effects in competing sales teams. *Management Science* 60(8):1965–1984, URL `http://dx.doi.org/10.1287/mnsc.2013.1840`.

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.

Coleman JS (1990) *Foundations of Social Theory* (Cambridge, MA: Harvard University Press).

Corritore M, Goldberg A, Srivastava SB (2020) Duality in diversity: How intrapersonal and interpersonal cultural heterogeneity relate to firm performance. *Administrative Science Quarterly* 65(2).

Dannals JE, Miller DT (2017) Social norm perception in groups with outliers. *Journal of Experimental Psychology: General* 146(9):1342.

Denison D, Hooijberg R, Lane N, Lief C (2012) *Leading culture change in global organizations: Aligning culture and strategy*, volume 394 (John Wiley & Sons).

DiMaggio P (1997) Culture and cognition. *Annual Review of Sociology* 23:263–287.

Dougherty D (1992) Interpretive barriers to successful product innovation in large firms. *Organization Science* 3(2):179–202, URL `http://dx.doi.org/10.1287/orsc.3.2.179`.

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.

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

Friedkin NE (2001) Norm formation in social influence networks. *Social Networks* 23(3):167–189.

Friedkin NE, Johnsen EC (1990) Social influence and opinions. *Journal of Mathematical Sociology* 15(3-4):193–206.

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

Fung HH, Wan Ho Y, Zhang R, Zhang X, Noels KA (2016) Age differences in personal values: Universal or cultural specific? *Psychology and Aging* 31(3):274.

Goffman E (1959) *The Presentation of Self in Everyday Life* (Random House).

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.

Goncalo JA, Chatman JA, Duguid MM, Kennedy JA (2015) Creativity from constraint? how the political correctness norm influences creativity in mixed-sex work groups. *Administrative Science Quarterly* 60(1):1–30.

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

Harrison JK, Chadwick M, Scales M (1996) The relationship between cross-cultural adjustment and the personality variables of self-efficacy and self-monitoring. *International Journal of Intercultural Relations* 20(2):167–188.

Heckathorn DD (1990) Collective sanctions and compliance norms: A formal theory of group-mediated social control. *American Sociological Review* 366–384.

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.

Hewlin PF (2003) And the award for best actor goes to...: Facades of conformity in organizational settings. *Academy of Management Review* 28(4):633–642.

Ho TK (1995) Random decision forests. *Proceedings of 3rd international conference on document analysis and recognition*, volume 1, 278–282 (IEEE).

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

Jin J, Rounds J (2012) Stability and change in work values: A meta-analysis of longitudinal studies. *Journal of Vocational Behavior - J VOCAT BEHAV* 80, URL `http://dx.doi.org/10.1016/j.jvb.2011.10.007`.

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.

Kleibergen F, Paap R (2006) Generalized reduced rank tests using the singular value decomposition. *Journal of Econometrics* 133(1):97–126.

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.

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

Leonardelli GJ, Pickett CL, Brewer MB (2010) Optimal distinctiveness theory: A framework for social identity, social cognition, and intergroup relations. *Advances in Experimental Social Psychology* 43:63–113.

Levine JM, Moreland RL (1991) Culture and socialization in work groups. Resnick LB, Levine JM, Teasley SD, eds., *Perspectives on socially shared cognition*, 257–279.

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

Liu CC, Srivastava SB (2015) Pulling closer and moving apart: Interaction, identity, and influence in the us senate, 1973 to 2009. *American Sociological Review* 80(1):192–217.

Lockwood C, Giorgi S, Glynn MA (2019) "how to do things with words": Mechanisms bridging language and action in management research. *Journal of Management* 45(1):7–34.

Maddux WW, Mullen E, Galinsky AD (2008) Chameleons bake bigger pies and take bigger pieces: Strategic behavioral mimicry facilitates negotiation outcomes. *Journal of Experimental Social Psychology* 44(2):461–468.

Marsden PV, Friedkin NE (1993) Network studies of social influence. *Sociological Methods & Research* 22(1):127–151.

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.

Mobasseri S, Goldberg A, Srivastava SB (2019) What is cultural fit? from cognition to behavior (and back) (Oxford University Press).

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

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.

O'Reilly CA, Chatman JA (2020) Transformational leader or narcissist? how grandiose narcissists can create and destroy organizations and institutions. *California Management Review* 62(3):5–27.

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).

Pettigrew AM (1979) On studying organizational cultures. *Administrative Science Quarterly* 24(4):570–581.

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

Polzer JT, DeFilippis E, Tobio K (2018) Countries, culture, and collaboration. *Academy of Management Proceedings*, volume 2018, 17645 (Academy of Management Briarcliff Manor, NY 10510).

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

Rafaeli A, Dutton J, Harquail CV, Mackie-Lewis S (1997) Navigating by attire: The use of dress by female administrative employees. *Academy of Management Journal* 40(1):9–45.

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

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.

Salancik GR, Pfeffer J (1978) A social information processing approach to job attitudes and task design. *Administrative science quarterly* 224–253.

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.

Schroeder J, Epley N (2015) The sound of intellect: Speech reveals a thoughtful mind, increasing a job candidate's appeal. *Psychological science* 26(6):877–891.

Schroeder J, Kardas M, Epley N (2017) The humanizing voice: Speech reveals, and text conceals, a more thoughtful mind in the midst of disagreement. *Psychological Science* 28(12):1745–1762.

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.

Sekiguchi T, Huber VL (2011) The use of person–organization fit and person–job fit information in making selection decisions. *Psychology and Aging* 116(2):203–216.

Shepherd HR (2017) The structure of perception: How networks shape ideas of norms. *Sociological Forum* 32(1):72–93, URL `http://dx.doi.org/10.1111/socf.12317`.

Sheridan JE (1992) Organizational culture and employee retention. *Academy of Management Journal* 35(5):1036–1056.

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, Goldberg A (2017) Language as a window into culture. *California Management Review* 60(1):56–68.

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).

United States Bureau of Labor Statistics (2018) Employee tenure summary .

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

Waldinger F (2012) Peer effects in science: Evidence from the dismissal of scientists in nazi germany. *The Review of Economic Studies* 79(2):838–861.

Weber RA, Camerer CF (2003) Cultural conflict and merger failure: An experimental approach. *Management Science* 49(4):400–415, URL `http://dx.doi.org/10.1287/mnsc.49.4.400.14430`.

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

Willis PE (2014) *Profane Culture: Updated Edition* (Princeton University Press).

Wynn AT, Correll SJ (2018) Puncturing the pipeline: Do technology companies alienate women in recruiting sessions? *Social Studies of Science* 149–164.

FIGURES

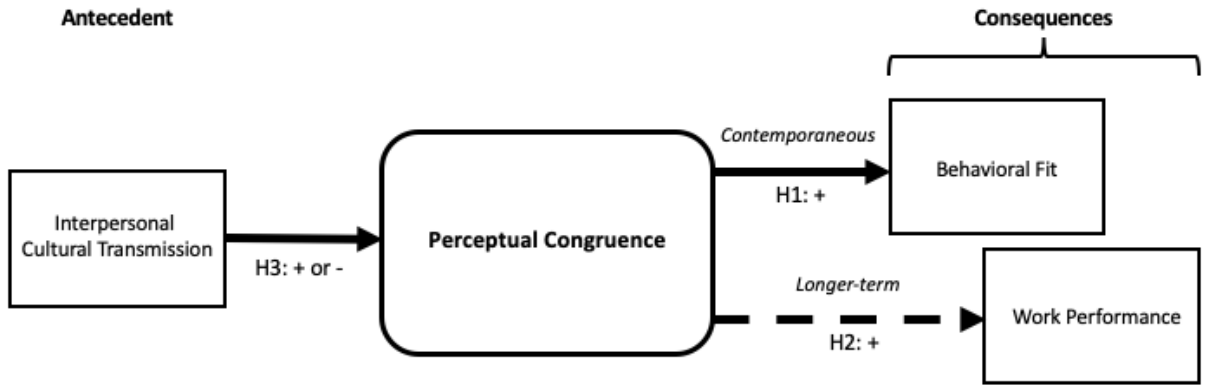

Figure 1 **Schematic representation of the theory. Hypothesis 1 posits a positive relationship between perceptual congruence and contemporaneous behavioral fit. Hypothesis 2 anticipates a positive relationship between perceptual congruence and work performance. Hypothesis 3 suggests that the interpersonal transmission of culture can lead to an increase or a decrease in perceptual congruence depending on the extent to which a focal individual's peers behave in normatively compliant ways and therefore give off accurate cues about the organization's cultural code.**

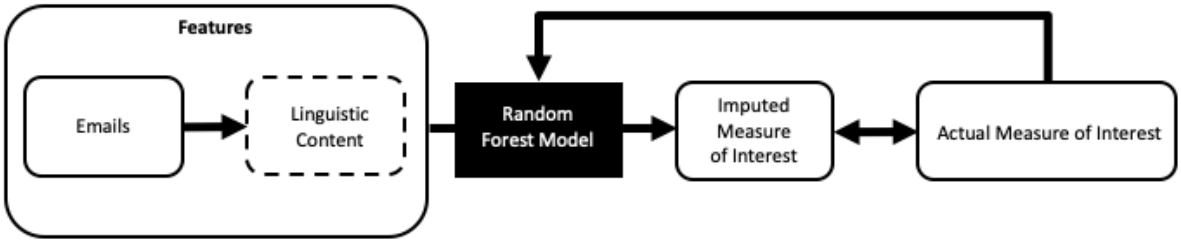

Figure 2 **Conceptual overview of the machine learning technique we use to impute measures of interest. We extract from email content linguistic features, which are inputs to a random forest model that predicts our actual measure of interest—e.g., perceptual congruence. We then use the identified relationship between linguistic features and the measure of interest to impute the measure for individuals in other time periods.**

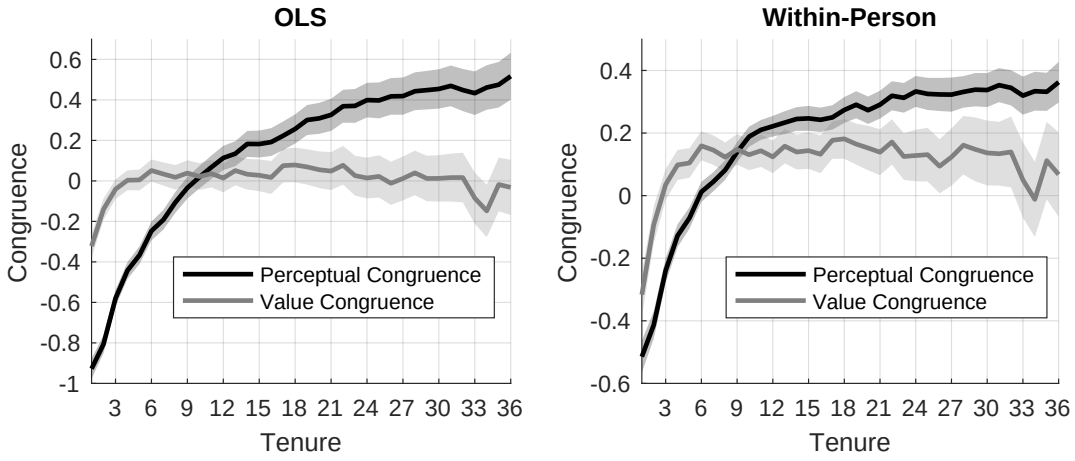

Figure 3 **OLS and Person Fixed Effects regressions of perceptual congruence and value congruence, with indicators for each tenure month up to 36 months in the company. Congruence measures are standardized. Shaded areas correspond to 95% confidence intervals.**

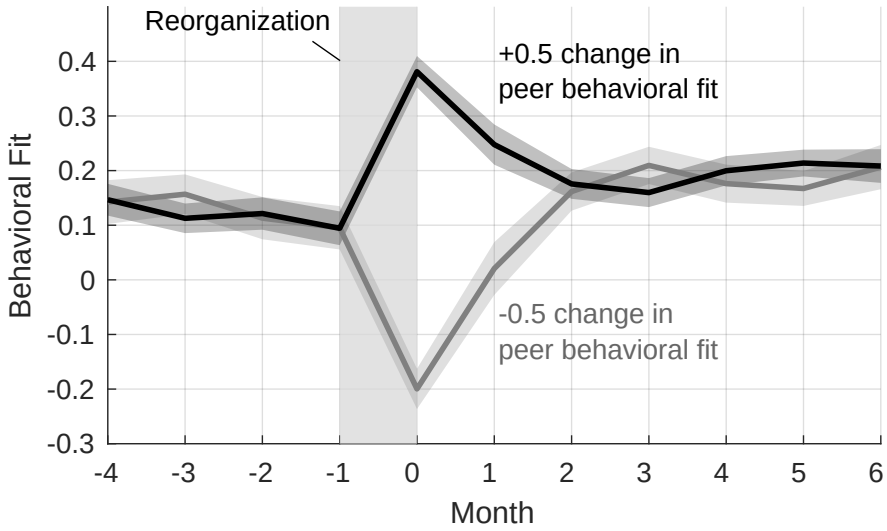

Figure 4 **Marginal effects, estimated by monthly consecutive instrumental variable models, of change in peer behavioral fit on individual behavioral fit. The two lines correspond to individuals who experienced a 0.5 increase (black) or decrease (gray) in peer behavioral fit. Shaded areas correspond to 95% confidence intervals.**

TABLES

Table 1 Cross-Sectional and Longitudinal Fixed Effects Regressions of Behavioral Fit

	Cross-Sectional			Longitudinal		
	Model 1†	Model 2†	Model 3†	Model 4	Model 5	Model 6
Perceptual Congruence‡	0.119***		0.142**	0.043**		0.043**
	(3.48)		(3.18)	(2.81)		(2.79)
Value Congruence‡		-0.009	-0.036		0.013	0.012
		(-0.19)	(-0.76)		(1.35)	(1.29)
Manager	0.630***	0.601***	0.552***	0.299***	0.303***	0.298***
	(6.76)	(4.08)	(3.81)	(5.42)	(5.47)	(5.40)
First Year	-0.239**	-0.358***	-0.337**	-0.076*	-0.084**	-0.075*
	(-3.08)	(-3.48)	(-3.31)	(-2.54)	(-2.81)	(-2.53)
Female	0.047	-0.025	-0.060			
	(0.66)	(-0.25)	(-0.62)			
Age	-0.002	-0.000	0.002			
	(-0.41)	(-0.03)	(0.41)			
Self-Monitoring	0.060	0.087	0.079			
	(1.69)	(1.70)	(1.57)			
Constant	0.291	0.179	0.132	-0.209	-0.218	-0.211
	(1.94)	(0.88)	(0.66)	(-1.64)	(-1.65)	(-1.66)
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	376	203	200	24215	24215	24215
R^2	0.281	0.232	0.281	0.107	0.075	0.107

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

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

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

Table 2 Fixed Effect Regressions of Bonus (logged) on Perceptual Congruence, Behavioral Fit, and Value Congruence

	Model 1	Model 2	Model 3	Model 4
Perceptual Congruence†	0.132***			0.112**
	(3.98)			(3.05)
Behavioral Fit†		0.128***		0.119***
		(4.46)		(4.14)
Value Congruence†			0.057**	0.047*
			(3.17)	(2.36)
Manager	0.025	-0.194	0.063	-0.180
	(0.13)	(-1.13)	(0.31)	(-1.02)
Constant	5.420***	5.652***	5.303***	5.702***
	(26.70)	(28.20)	(25.70)	(28.61)
Individual FE	Yes	Yes	Yes	Yes
Department FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Observations	6377	4782	6377	4778
Num. Individuals	1303	1057	1303	1056
R^2	0.043	0.059	0.040	0.065

t statistics in parentheses; standard errors clustered by individual

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

Table 3 OLS and Instrumental Variables Fixed Effects Regressions of Behavioral Fit and Perceptual Congruence on Peer Behavioral Fit

	OLS	Instrumental Variable				
	Model 1 Behav. Fit	Model 2 Behav. Fit	Model 3 Percep. Congr.	Model 4 Value Congr.	Model 5 Percep. Congr.	Model 6 Value Congr.
Peer Behavioral Fit†	0.221***	0.266***	0.068**	-0.020		
	(12.68)	(6.38)	(3.03)	(-0.47)		
Peer Perceptual Congruence†					0.064	
					(0.63)	
Peer Value Congruence†						0.073
						(0.83)
Num. Peers†	0.001**	-0.013*	0.001	0.008*	0.024	-0.004
	(3.11)	(-2.50)	(0.27)	(2.14)	(1.36)	(-0.38)
Manager	0.365***	0.555***	0.042	-0.096	-0.430	0.136
	(7.67)	(4.34)	(0.77)	(-0.95)	(-1.18)	(0.68)
First Year	-0.154***	-0.204***	-0.163***	0.028	-0.013	-0.043
	(-6.72)	(-4.12)	(-6.28)	(0.65)	(-0.12)	(-0.64)
Constant	-0.065	0.648**	0.259**	-0.257	-0.756	0.257
	(-1.23)	(2.67)	(2.67)	(-1.45)	(-0.99)	(0.63)
Individual FE	Yes	Yes	Yes	Yes	Yes	Yes
Department FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
N	22080	21998	21998	21998	21985	21985
Num. Individuals	1515	1508	1508	1508	1504	1504
R^2	0.28					
Kleibergen-Paap F		8.99	8.99	8.99	0.85	1.79

t statistics in parentheses; standard errors clustered by individual

† lagged variables, instrumented endogenous variables in Models 2-6

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

APPENDIX A: BEHAVIORAL FIT

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.

$$O^l_{iT} = \frac{\overrightarrow{m}^l_{iT}}{\sum_{l \in L} \overrightarrow{m}^l_{iT}} \tag{7}$$

$$I^l_{iT} = \frac{\overleftarrow{m}^l_{iT}}{\sum_{l \in L} \overleftarrow{m}^l_{iT}} \tag{8}$$

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

$$BF_{iT} = -log\left(JS(O_{iT} \parallel I_{iT})\right) \tag{9}$$

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.

$$JS(O_{iT} \parallel I_{iT}) = \frac{1}{2} KL\left(O_{iT} \parallel M_{iT}\right) + \frac{1}{2} KL\left(I_{iT} \parallel M_{iT}\right) \tag{10}$$

$$KL(D_{iT} \parallel M_{iT}) = \sum_{l \in L} D^l_{iT} \log_2 \frac{D^l_{iT}}{M^l_{iT}} \tag{11}$$

Validation of Behavioral Cultural Fit

We have argued above that the LIWC lexicon, on which the behavioral cultural fit 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 behavioral fit is effective at predicting individual attainment in an organization. Since this is the first time our measure of behavioral fit 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.[9]. 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 behavioral fit 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 behavioral fit 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 behavioral fit, labeled BF_k. We randomly selected k LIWC categories and constructed the measure as we did above (according to equation 9), 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 C1. As the plot clearly indicates, the behavioral fit 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, behavioral fit 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 C1 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 2 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.

APPENDIX C: EVALUATING MODEL FIT

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 C1 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 C2 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 C2 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 C3 ABOUT HERE.]

APPENDIX D: CHARACTERIZING THE ORGANIZATION'S CULTURE

To provide context on the content of the organization's culture, we focused on eight OCP dimensions developed in Chatman et al. (2014). This analysis revealed that the organization was most

focused on Customers (mean=6.30), followed by Results (mean=5.82) and Integrity (mean=5.62). In contrast, it was relatively less focused on Transparency (mean=4.66) and Details (mean=4.95). This dimension-level portrait was comparable to the pattern of results for the relative salience of specific value statements, with the six most highly ranked items being: Customer Oriented (mean=6.75, sd=1.8), Making your Numbers (mean=6.32, sd=2.07), Being Results Oriented (mean=6.18, sd=1.78), Listening to Customers (mean=6.13, sd=1.74), Having High Expectations for Performance (mean=6.11, sd=1.85), Being Market Driven (mean=5.99, sd=1.81), and Fast Moving (mean=5.91, sd=2.04). The lowest ranked (Most Uncharacteristic) six items of the 54 item aggregate current culture profile included the items: Predictability (mean=2.97, s.d.=1.65), Stability (mean=3.07, sd=1.78), Security of Employment (mean=3.13, sd=1.77), High Levels of Conflict (mean=3.56, sd=2.11), Being Reflective (mean=3.88, sd=1.53) and Being Calm (mean=4.00, sd=1.67).

APPENDIX E: IV REGRESSION ROBUSTNESS CHECKS

In Table E1, we report the results of two supplemental analyses designed to assess the robustness of the results of our instrumental variables analysis and test the boundary conditions of our theory. First, given that our measures of interest are all defined with respect to the reference group of an individual's interlocutors in a given month, which people can—to varying degrees—self-select into, we replicated the instrumental variables analysis using behavioral fit, peer behavioral fit, and peer perceptual congruence measures that were based on the reference group of *all* employees in the organization. Table E1, Model 1, shows that peer behavioral fit, where peers are defined as all other employees in the organization, predicts the focal actor's behavioral fit relative to this same reference group. Model 3 shows that peer behavioral fit based on all other employees is also positively related to the focal actor's perceptual congruence. These results mitigate concerns that our main results are an artifact of our choice to define behavioral cultural fit relative to a focal actor's interlocutors in a given month.

Second, our instrumental variable approach is predicated on the assumption that the reorganization produced exogenous shifts in focal actors' peer groups. Yet it is possible that the reorganization

was biased toward certain desired shifts in peer groups—for example, distancing leaders and their teams when there was animosity between them or bringing together formal subunits whose heads had compatible management styles. To address such possibilities, we replicated the analyses using a sub-sample of employees who were not in supervisory roles. We reasoned that, insofar as the reorganization was designed in part to change peer groups, such social engineering was targeted to the leadership ranks of the company. For those in individual contributor—rather than supervisory—roles, the reorganization was much more likely to have produced exogenous change in peer networks. As Table E1, Model 2, illustrates, the relationship between peer behavioral fit and the focal actor's behavioral fit holds even for this more restricted sample of employees. Model 4 shows that peer behavioral fit predicts the focal actor's perceptual congruence in the more restricted sample. By removing individuals with supervisory responsibilities, this analysis also offers insight into whether language accommodation, our measure of behavioral fit, is a simple reflection of people aligning to the linguistic style of their most powerful interlocutors. Given the consistency of the findings when supervisors are included or dropped from the analysis, we conclude that this is not likely to be the case.

[TABLE E1 ABOUT HERE.]

APPENDIX FIGURES

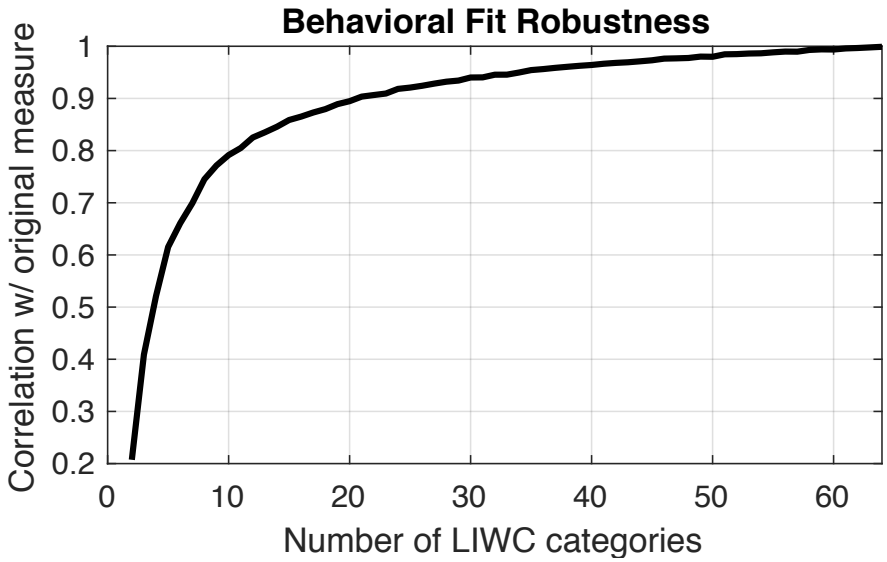

Figure C1 **Robustness of the behavioral fit measure to LIWC category composition**

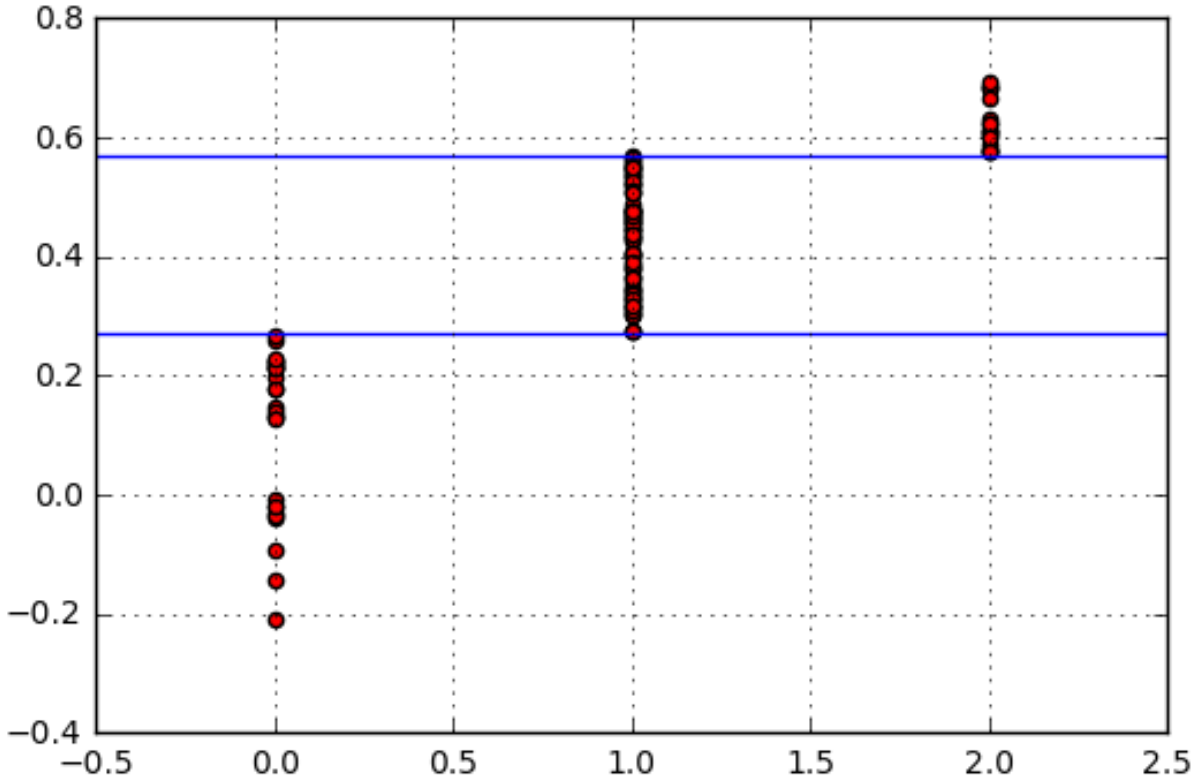

Figure C2 **Division of Continuous Cultural Fit into Classes**

APPENDIX TABLES

Table C1 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
PC-Interloc.	0.857	0.726	0.767	0.267	0.651	0.711
PC-Org.	1	0.875	0.865	0.547	0.867	0.849
VC-Interloc.	1	0.952	0.950	0.667	0.952	0.934
VC-Org.	1	0.923	0.951	0.667	0.923	0.906

Table C2 p-Values for Difference in Means between Low and High

	P-Value
PC-Interloc.	2.661e−3
PC-Org.	1.874e−8
VC-Interloc.	8.500e−6
VC-Org.	7.157e−5

Table C3 Areas under the ROC Curve

	ROC AUC
PC-Interloc.	0.740
PC-Org.	0.910
VC-Interloc.	0.950
VC-Org.	0.930

Table E1 Robustness Checks—Instrumental Variables Fixed Effect Regressions of Behavioral Fit and Percepttual Congruence

	Behavioral Fit		Perceptual Congruence	
	Model 1 Organization	Model 2 Non-Managers	Model 3 Organization	Model 4 Non-Managers
Peer Behavioral Fit†		0.223***		0.080**
		(5.75)		(3.23)
Peer Behavioral Fit (Organization)†	0.147***		0.090*	
	(5.49)		(2.02)	
Num. Peers†	-0.003	-0.013*	0.002	-0.000
	(-1.85)	(-2.09)	(0.50)	(-0.15)
Manager	0.133***		0.059	
	(3.86)		(0.82)	
First Year	-0.031*	-0.163***	-0.182***	-0.213***
	(-2.17)	(-3.45)	(-5.90)	(-6.81)
Individual FE	Yes	Yes	Yes	Yes
Department FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
N	19498	17768	23063	17768
Num. Individuals	1213	1255	1547	1255
Kleibergen-Paap F	2.97	8.80	4.58	8.80

t statistics in parentheses; standard errors clustered by individual

† instrumented and lagged endogenous variables

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