

# Situated Cultural Fit: Value Congruence, Perceptual Accuracy, 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

Why are some people more successful than others at fitting in culturally over time? Prior research has offered divergent and seemingly inconsistent answers to this question. One perspective has highlighted the importance of shared values in shaping behavior, while another has emphasized the role of situational cues and the ability to read the group's cultural code. We develop a theoretical account that reconciles these competing perspectives. Drawing on dual-process theories of culture and cognition and the distinction between constrained and unconstrained situations, we develop a situated theory of cultural fit. We argue that values matter for behavior in unconstrained situations—in particular, for the choice to remain at or voluntarily exit from the organization. In contrast, perceptual accuracy matters for behavior in constrained situations—specifically, for the capacity to exhibit real-time linguistic conformity with peers. We further show that a person's behavior and perceptual accuracy are both influenced by observations of others' behavior, whereas value congruence is less susceptible to peer influence. Drawing on email and survey data from a mid-sized technology firm, we use the tools of computational linguistics and machine learning to develop longitudinal measures of cognitive and behavioral cultural fit. We also take advantage of a reorganization that produced quasi-exogenous shifts in employees' interlocutors to identify the causal impact of peer influence. We discuss implications of these findings for research on person-culture fit, cultural change and transmission, dual-process models of culture and cognition, and the pairing of surveys with digital trace data.

---

## Introduction

Whether assimilating to a country or adapting to a new school, people typically seek to fit in culturally with their social groups. The benefits of conformity, as well as the sanctions and penalties that come with failed cultural integration, are particularly stark in contemporary organizations. Indeed, prior work has consistently demonstrated that high levels of individual cultural fit are associated with increased productivity, stronger commitment, and less turnover (Kristof-Brown et al. 2005, Chatman and O'Reilly 2016). Moreover, employers have increasingly emphasized screening, selecting, and socializing new hires on the basis of cultural fit rather than exclusively hiring for skills (Chatman 1991, Meyer et al. 2010a, Rivera 2012). At the same time, as the average tenure in firms has declined (Hall 1996), workers must frequently retool themselves culturally as they move from one organization to the next. Yet people vary considerably in their ability to adapt culturally within a given organization (Srivastava et al. 2018). Why are some individuals more successful than others at cultural adjustment?

Existing research offers two different, and seemingly inconsistent, answers to this question. The first focuses on values. This line of work, echoing a long tradition in psychology and sociology, sees the locus of culture in the degree to which people embrace their group's behavioral norms. Fitting in therefore implies having preferences that are consistent with the norms that prevail in an organization. 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, Edwards 2008)—predicts a variety of individual and organizational outcomes.

A second explanation largely rejects the notion that values affect behavior, positing instead that culture shapes action through situational cues. This approach shifts focus from individuals' preferences to their readings of situations, arguing that behaviors are primarily driven by the cultural scripts invoked by exposure to others. An employee's decision to use polite language in a meeting, for example, often reveals little about her underlying preference for civil discourse but instead reflects the norms she observes in the behavior of other meeting participants. Indeed,

people pursue action for which their “cultural equipment is well suited” (Swidler 1986, p. 277), suggesting that those who fit in are those whose readings of the cultural code lead them to behave in normatively appropriate ways.

These two perspectives appear to provide incompatible explanations for the sources and consequences of cultural fit. Whereas the former suggests that cultural fit is the result of internalizing and embracing prevailing values and norms, the latter sees it as the product of correctly deciphering the normative code. We propose that this theoretical incongruity can be resolved by recognizing that culture operates at the individual level via two distinct forms of cognition—more implicit versus more explicit (Lizardo 2017)—that correspond to different types of behavior. Implicit knowledge shapes habitual and non-reflective behavior, while explicit understanding is at play when people make more deliberative and conscious decisions.

Building on this insight, we propose that whether values or perceptions give rise to cultural conformity will depend on the type of situation a person faces. When situations are *constrained*—that is, others’ behavior provides cues about how to interpret what is going on and accurately signals how one ought to behave—behavior tends to be driven by implicit knowledge, and when situations are *unconstrained*—that is, the setting is unfamiliar, others’ behavior is inconsistent or not salient, or when the choice being considered is highly consequential to the individual and her identity—behavior is more directly tied to explicit knowledge (Leung and Morris 2015). In other words, we reconceptualize cultural fit as a *situated* construct: its antecedents and behavioral consequences vary as a function of the type of situation a person faces.

More specifically, we argue that value congruence, which is based on the accumulation of episodic assessments of the correspondence between an individual’s own normative commitments and the norms that prevail in an organization, predicts behavior in unconstrained situations—in particular, a person’s self-identification with the organization and thus her (voluntary) choice to stay or exit. In contrast, we propose that a heretofore understudied construct—*perceptual accuracy*, or the alignment between an individual’s reading of the cultural code and the readings of other group

members—affects behavior in constrained contexts and thus the capacity to behave in real-time in normatively compliant ways.

We further posit that values are a relatively inert aspect of cognition, whereas perceptions are susceptible to social learning. Specifically, we suggest that exposure to normatively compliant (or non-compliant) behavior among one’s peers can boost (or dampen) one’s own perceptual accuracy and, in turn, one’s capacity for normative compliance—whether or not one subscribes to those norms. In other words, exposure to peers who are normatively compliant: (1) increases perceptual accuracy but not necessarily value congruence; and, as a result, (2) enhances one’s capacity to behave in normatively compliant ways in routine interactions with colleagues.

To evaluate these ideas, we employ a multi-method empirical strategy that draws on survey data, eight years of internal email data, and personnel records from a mid-sized technology firm. First, we use the Organizational Culture Profile (Chatman 1991), a validated culture survey, in two different ways to measure the value congruence and perceptual accuracy dimensions of cultural fit. Following established practice, we assess value congruence by comparing an individual’s self-reported preferences to prevailing values reported by her peers. Departing from prior work, we develop a novel measure of perceptual accuracy by comparing an individual’s reports of widespread values and norms to the ones her peers believe are predominant.

We then use the tools of computational linguistics and machine learning to transform these cross-sectional cognitive measures into longitudinal measures and to develop time-varying measures of (real-time) behavioral cultural fit based on the linguistic style that employees use in email communications with their colleagues. We also take advantage of a reorganization that produced quasi-exogenous shifts in employees’ peer groups to identify the causal impact of social influence—that is, of how a focal actor’s perceptual accuracy and behavioral fit change in response to essentially random changes in the peers to which she is connected. We conclude by discussing how our findings advance theories of cultural fit in organizations.

## Theory and Hypotheses

### Value Congruence and Perceptual Accuracy

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

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

The notion that values are fundamental drivers of human behavior has a long history in sociology (Parsons 1968) and psychology (Schwartz 1992, Hofstede 2001). This research demonstrates, for example, that values are associated with cross-national and regional differences in economic growth (Inglehart and Baker 2000) and violence (Nisbett and Cohen 1996), as well as with individual

lifestyle (Miles 2015), financial (Keister 2008) and occupational (Alesina et al. 2015) choices. Yet a growing body of research finds that people's stated values are, in many cases, poor predictors of their behavior (Greenwald and Banaji 1995). Economically disadvantaged high school students, for example, tend to express mainstream attitudes on educational achievement and sexual behavior but adopt behaviors that appear to be inconsistent with these ideals (Harding 2007). In organizations, too, people's behaviors are often incongruent with their stated beliefs: self-reported values on cross-functional collaboration, for example, are largely unrelated to individuals' propensity to build network ties that span functional boundaries (Srivastava and Banaji 2011).

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

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

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

### **Situated Cultural Fit**

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

Drawing on advances in cognitive science, sociologists of culture have increasingly concluded that, at the individual level, culture generally operates via two distinct forms of cognition: "practical," or implicit knowledge; and "propositional," or explicit knowledge (Lizardo 2017). Practical knowledge refers to schemas, prototypes, and associations that are difficult to articulate. In contrast, propositional knowledge refers to worldviews, ideologies, and orientations that are more readily expressed. These two forms of cognition have differing implications for behavior. Whereas

the former predominantly shapes habitual and less reflective behavior, the latter is at play when people make more deliberative and conscious decisions.

Seen in this light, the question becomes not whether values matter for demonstrations of cultural fit but rather for *what kinds* of behavior values are consequential. We argue that an important distinction missing from previous literature differentiates between situations that trigger habitual versus deliberative forms of action. Following Leung and Morris (2015), we refer to these as *constrained* and *unconstrained* situations, respectively. Constrained situations are ones in which others' behaviors provide consistent cues about the meaning of the situation and, accordingly, which behaviors are desirable. Unconstrained situations, in contrast, occur either when situational cues are absent or ambiguous—whether because the setting is unfamiliar, others' behaviors are inconsistent or behavior is done in private—or when decisions are highly consequential, and perhaps even tinged with moral implications.<sup>1</sup>

Constrained and unconstrained situations induce different forms of action. In constrained contexts behavior is most likely to be driven non-reflectively by situationally activated construals, overriding value orientations. Unconstrained situations, in contrast, activate deliberative decision-making. As Leschziner and Green (2013), for example, demonstrate, chefs are often unable to explain routine food preparation decisions that rely on culinary conventions. These day-to-day decisions are driven by normative expectations about how food should be prepared and presented. But when they are intentionally innovating or deliberately changing in response to economic pressures, chefs provide more explicit rationales. It is during such moments of disjuncture that values matter most, as during such times people reflect on their choices in light of their explicit beliefs on what is worthy and desirable (Miles 2015).

Recognizing that different situations trigger distinct forms of behavior requires rethinking cultural fit as a situated construct rather than one that is determined solely by values or perceptions. Accordingly, we argue that value congruence is consequential for behavior in unconstrained situations, whereas perceptual accuracy shapes behavior in constrained situations.



Most activities in organizations occur routinely, in settings that provide high situational clarity (Davis-Blake and Pfeffer 1989). This situational clarity is commonly a function of the actor's familiarity with the setting and the availability of habituated behavioral responses to it. We therefore posit that perceptual accuracy will be consequential for individuals' ability to exhibit culturally compliant behavior in routine, day-to-day activities, which we define as *behavioral fit*. To productively participate in ritualistic complaining, for example, the employees in Weeks' (2004) ethnography of *BritArm Bank* had to complain at the appropriate level: not too much so as to avoid rocking the boat, but enough to signal membership and belonging with the group. We refer to such conformity to normative expectations as behavioral cultural fit.

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

Instead, we expect that value congruence will predict behavior in unconstrained situations. These types of situations are less common in organizational life. They typically occur when organizational members face unfamiliar settings or in light of events that raise the salience of their self-identification with the organization. Decisions to stay or exit the organization are exemplary of the latter. When people make such decisions, they respond less to what types of appropriate behaviors the situation activates and more to their beliefs about what is desirable. Moreover, such deliberation often occurs in private contexts where colleagues' behavioral cues and normative expectations are not on display and thus less salient.

Together, these arguments lead us to formulate the following two hypotheses:

***Hypothesis 1 (H1).*** *Value congruence will be negatively associated with voluntary exit but unrelated to behavioral cultural fit.*

**Hypothesis 2 (H2).** *Perceptual accuracy will be positively related to behavioral cultural fit but unrelated to voluntary exit.*

### **Situated Cultural Fit and the Interpersonal Transmission of Culture**

Our situated theory of cultural fit also sheds new light on why some people exhibit a greater capacity than others to assimilate into their organizational cultures. Previous work has approached this question from two different and seemingly irreconcilable angles. A prominent line of work has tended to conceptualize cultural fit as a fundamental compatibility between individuals and organizations—a match between the “personalities” of the individual and the group (Schneider 1987, Cable and Judge 1996, Baron et al. 2001). Those who fit in culturally are therefore those who are innately compatible with the organizations’ culture. A parallel stream of work on organizational culture emphasizes cultural fit as a dynamic process. This work demonstrates that individuals are capable of adapting their behavior to the prevailing norms in an organization through (direct and indirect) processes of socialization (Van Maanen and Schein 1979, Chatman 1991, Van Maanen 1975, Ashforth and Saks 1996).

The distinction between value congruence and perceptual accuracy, we contend, explains why cultural fit can be simultaneously fixed and dynamic. Values are, by definition, deeply held and enduring beliefs about what is desirable and appropriate. Given that individuals’ values are encoded in implicit cognition and thus slower to change (Meglino and Ravlin 1998, Vaisey 2009, Srivastava and Banaji 2011), value congruence is likely to remain relatively stable throughout an employee’s tenure in an organization. Indeed, as Vaisey and Lizardo (2016) demonstrate, values and moral attitudes are surprisingly durable throughout people’s adult lifetimes. In other words, cultural change is mostly attributable to generational differences and not to period effects.

Perceptual accuracy is, in contrast, organization-, and even subgroup-, specific and therefore more malleable. Employees joining a new organization are required, by virtue of each group’s idiosyncratic style, to learn the group’s specific cultural code. Recent work demonstrates that individuals exhibit great variability in their ability to adjust their behavioral fit over time. Whereas

some gradually adapt their behaviors to meet their peers' expectations, others fail to do so, thereby lowering their chances of promotion and their likelihood of receiving a favorable performance rating (Chatman and Spataro 2005, Srivastava et al. 2018).

What factors lead some people to increase their behavioral fit over time, while others remain stagnant? One line of work attributes such variance to psychological differences between individuals. For example, a robust literature in social psychology has focused on self-monitoring orientation—a sensitivity and responsiveness to social cues of situational appropriateness (Snyder 1979, Kilduff and Day 1994, Sasovova et al. 2010). High self-monitors tend to regulate their behavior given their read of what is expected of them, whereas low self-monitors hew to their sense of self, irrespective of the situation. Self-monitoring is also related to a 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). High self-monitors, in other words, are more motivated to read the cultural code, conform to it, and be perceived as authentic when they do.

Yet perceptual accuracy is also a matter of context, not just of intrinsic ability. Humans are innately motivated to be attuned to the cultural code prevalent in their immediate social environments (Liebal et al. 2013). Consequently, we argue that perceptual accuracy is dependent not only on inherent differences between people's cultural attentiveness but also on the social context in which they are embedded. Adjusting to the cultural code of a group is, by definition, a process of social learning. The quality of this learning depends not only on the student but also on the peers from whom she learns.

We therefore expect that the composition of a person's network has a significant impact on her ability to correctly decipher the cultural code and to adapt her behaviors accordingly. Experimental work in young children, for example, demonstrates that exposure to multiple and consistent behaviors increases the fidelity and speed of cultural transmission (Herrmann et al. 2013). Similarly, in the workplace, employees' ability to learn and their susceptibility to influence from 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 accuracy (Balkundi and Kilduff 2006).

Importantly, people primarily have access to their peers' behaviors. It is through observing these behaviors that they develop their own perceptions of the cultural environment. We therefore anticipate that peers' behavior—as opposed to their private values or perceptions—will influence the focal individual's own thoughts and behavior. Moreover, because we argue that the ability to behave compliantly is primarily dependent on perceptual accuracy, we also expect that individuals' perceptual accuracy will be influenced through their observations of their colleagues. In contrast, we expect value congruence to remain mostly impermeable to peer influence.

In support of these expectations, an extensive literature has shown that individuals' 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); however, exposure to peers whose deeply held values and beliefs run counter to one's own can also activate biases in information processing such that discordant information is discounted or even rejected (Lord et al. 1979, Dandekar et al. 2013, Liu and Srivastava 2015, Bail et al. 2018). In contrast, expectations of normatively appropriate behavior are strongly shaped by *shared perceptions* that arise through interaction and observation (Friedkin 2001). Taken together, these findings lead to the prediction that a person's perceptions of the cultural order will be more susceptible to social influence than will her deeply rooted values. Overall, we expect:

***Hypothesis 3 (H3).*** *Perceptual accuracy and behavioral cultural fit (but not value congruence) will be susceptible to peer influence. Specifically, as one's peers behave in more (less) normatively compliant ways, one's own perceptual accuracy will increase (decrease) and one's behavioral cultural fit will concomitantly increase (decrease).*

## **Method**

Testing these hypotheses requires access to longitudinal data on value congruence, perceptual accuracy, and behavioral conformity, as well as exogenous variation in the set of peers to which a

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

To address the limitations, we employ a multi-method approach that draws on survey and email communication data from a mid-sized technology firm and that uses machine learning techniques to impute time-varying measures from cross-sectional data. Moreover, we use an instrumental variables methodology, which takes advantage of a reorganization event that produced quasi-exogenous shifts in employees' peer groups, to estimate the causal effect of interpersonal cultural transmission. We detail these methodological choices in this section. First, we explain how we use email and survey data to measure, respectively, behavioral cultural fit and the two cognitive dimensions of cultural fit: value congruence and perceptual accuracy. Second, we provide descriptions of the data and variables, including an explanation of how we use machine learning to transform the one-time survey into imputed, time-varying variables. Finally, we provide an overview of our analytical strategy, with a focus on the instrumental variable approach.

### **Measuring Behavioral and Cognitive Cultural Fit**

Studies of culture often focus on its content, namely, on specific beliefs, interpretations and normative behaviors. In contrast, our approach is distributive (Harrison and Carroll 2006). Rather than asking how specific cultural elements relate to one another and to other variables of interest, we seek to characterize individuals on the basis of their cultural similarity to their groups both behaviorally and cognitively. We define each individual's reference group as her email interlocutors

in a given month, weighted by volume of interaction. Given that subcultures in organizations do not necessarily conform to the contours of formal subunits, this choice of reference group allows us to identify a person's fit in an empirically grounded manner, without having to make assumptions about the boundaries of subcultures in the organization.

*Measuring Behavior*—We operationalize behavioral cultural fit 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—we focus on the similarity of *linguistic style* between an individual and her reference group. Drawing on previous sociological work on culture (Bail et al. 2017, Doyle et al. 2017), ILUM uses the well-established and widely used Linguistic Inquiry and Word Count (LIWC) lexicon (Pennebaker et al. 2007) to measure linguistic style. LIWC is a semantic dictionary that maps words into 64 high-level distinct 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.

*Measuring Cognition*—To assess the two dimensions of cognitive cultural fit, we implemented the widely used Organizational Culture Profile (OCP) (Chatman et al. 2014). Cultural sociologists often rely on self-reports as a means to measure deep-seated values, preferences and beliefs (e.g. Harding 2007, Vaisey 2009, Goldberg 2011, Miles 2015). The advantage of using OCP is that it provides a comprehensive set of cultural elements that have been applied to and validated in a wide variety of organizations. 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, Sarros et al. 2005). Using the Q-sort methodology (Block 1961), respondents are asked to rank these 54 statements into nine categories, with a specified number of statements in each category.<sup>2</sup> This sorting of value statements represents an individual’s cultural profile. Employing our distributive approach, we can use this cultural profile to estimate each individual’s distance from her reference group, as we detail below.

## Data and Variables

Our empirical setting is a mid-sized technology firm. We obtained three types of data:

*Personnel Records*—We obtained 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, date of exit, and reason for exit (voluntary or involuntary).

*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*—We sent two versions of the OCP to the organization, one asking employees to characterize the current culture of the organization and the other asking employees

to characterize their personally desired culture. All employees completed the survey describing the organization's current culture and a randomly selected half of employees completed the survey of their own personally desired cultural characteristics.<sup>3</sup> Overall, we received 440 completed surveys about the current organizational culture and 238 completed surveys about the personally desired culture.

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.

**Behavioral Cultural Fit** We operationalized behavioral fit using ILUM, as applied to internal email communication (Goldberg et al. 2016, Srivastava et al. 2018). To derive this measure, 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. We discuss the technical details of this measure in Appendix A.

Intuitively, when the outgoing and incoming distributions are nearly identical, the divergence approaches zero, suggesting high behavioral fit; conversely, greater deviation between the probabilities of usage of LIWC categories translates to greater divergence and thus implies lower behavioral fit. Stated differently, 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 fit in a given month. For example, an individual using a relatively high proportion of negations in her outgoing communication but who receives a far smaller proportion of negations in her incoming messages would be characterized as having low behavioral cultural fit (at least with respect to this LIWC category). Such an individual would be expressing disagreement, whereas her peers would be refraining from doing so.



Although ILUM has been used in previous work to measure cultural fit, it is still a fairly new methodology. To further validate our measure of behavioral fit we conducted two supplemental analyses. The first demonstrates that LIWC categories reflect culturally meaningful content—for example, that individuals who espouse an innovative culture tend to use more future-tense language. In the second analysis we show that, even if we assume that certain LIWC categories are culturally meaningless, our measure is still robust to the removal of these categories. These additional analyses are reported in Appendix A.

**Perceptual Accuracy and Value Congruence** We operationalized perceptual accuracy and value congruence based on employee responses to the OCP (Chatman et al. 2014). To derive measures of fit, we calculated the correlation between culture profiles. We configured the OCP to yield two separate culture profiles for each respondent: a profile based on her assessment of the current organizational culture and one based on her preferences for each value statement. For the former, we asked: “To what extent do the value statements characterize the organization as a whole?” For the latter, we asked: “To what extent do the value statements characterize your personally desired values, that is, the values you desire in an organization?” Our two measures of cognitive cultural fit are based on the correlation between individual *i*’s cultural profile and a reference group cultural profile.

To make these measures comparable to our measure of behavioral fit, we chose the same reference group—i.e., the set of colleagues a person had email contact with in a given month weighted by communication volume. We defined *perceptual accuracy* as the congruence between an individual’s current culture profile and the reference group’s current culture profile. Similarly, we defined *value congruence* as the correspondence between an individual’s personal culture profile and the reference group’s current culture profile. Note that the reference group profile is identical in both cases. The difference between the two measures stems from the choice of individual culture profile: current culture for perceptual accuracy and personal culture for value congruence.<sup>4</sup>

Our measure of value congruence using OCP is well-established in the literature. Because perceptual accuracy is a new construct, however, it has never been implemented using the OCP.

Perceptual accuracy, as we discuss above, is a function of the compatibility between an individual's and reference group's construals of situations and readings of behavioral norms. The OCP allows us to measure the latter, not the former. Though it is possible that individuals correctly read norms while failing to correctly construe situations, or vice versa, we assume that the two relate to similar cognitive processes and are therefore correlated. As such, our measure is a conservative estimate of perceptual accuracy.

**Imputing Cognitive Cultural Fit Over Time** The procedure above creates cross-sectional measures of perceptual accuracy and value congruence; however, longitudinal cognitive measures are needed to test hypotheses about the dynamic interrelationships among the three fit 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 fit) the “linguistic signature” of perceptual accuracy and value congruence (see also Bail, 2017).

We assumed that, if language reflects internal processes of cognition (Pinker 2007), then there should be an identifiable relationship between email communication and cognitive cultural fit. If this relationship can be discerned through machine learning, then it should be possible to impute perceptual accuracy and value congruence measures for all employees, including those who departed before the OCP was implemented and those who were employed but chose not to participate. Moreover, assuming a relatively stable underlying relationship between language use and cognition, these measures can be imputed for individuals at all points in time for which they exchanged email messages with colleagues. In other words, this procedure allowed us to transform a one-time collection of value preferences and perceptions of the current culture, based on the OCP, into longitudinal measures of cognitive cultural fit.

We used a random forest model to help uncover this underlying link between language and cognition (Ho 1995, Friedman et al. 2001). Random forest models have several beneficial characteristics for this task: they can detect arbitrary, nonlinear relationships; they typically require fewer observations than do other machine learning methods to produce comparable results; and they are inherently robust to overfitting, or incorrectly inferring signal from idiosyncratic noise in the data. Figure 1 provides a conceptual overview of this procedure. Further procedural details are provided in the Appendix B; evaluative analyses regarding model fit are provided in Appendix C.

[FIGURE 1 ABOUT HERE]

**Peer Cultural Fit** After imputing perceptual accuracy and value congruence, we turned next to identifying the distribution of these measures in the network of email contacts surrounding a focal individual as a means to test H3. To do this, we first identified an individual  $i$ 's communication partners  $J$  for each month  $T$ . Then, using our time-varying measures of cognitive cultural fit, as well as our time-varying measure of behavioral fit, we took the mean cultural fit for all communication partners  $J$ , weighted by the volume of incoming communication received from each interlocutor, to generate  $i$ 's peer cultural fit for month  $T$ . We did this for each cultural fit measure, yielding network-based measures that we refer to as peer behavioral fit, peer perceptual accuracy, and peer value congruence.

**Individual Outcomes** To establish the validity of our imputed longitudinal measures, we implemented a supplemental analysis reported below. This was not a direct test of our hypotheses but was designed to assess whether the imputed measures related to career outcomes as would be expected based on theory and prior research. In particular, we focused on monthly bonus payouts, which were received only by those in job roles such as sales or operations, for which productivity could be objectively assessed. For each of these roles, 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.

To test H1 and H2, we focused on voluntary exit, as identified by an employee's departure date. We used company records to distinguish between voluntary and involuntary exit.

**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 tenure since those who have worked in the organization longer are likely to be exposed to more information about the culture. Finally, we included departmental affiliation 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).<sup>5</sup> For our between-person models, we included additional control variables for age and gender.

### **Analytical Approach**

We tested Hypothesis 1, which posits that value congruence will be negatively associated with voluntary exit but unrelated to behavioral cultural fit, using Cox proportional hazard models. We use a competing risks model that extends the Cox model to the case of multiple failures. In our case, involuntary exit is the competing risk.<sup>6</sup>

We tested Hypothesis 2, which posits that perceptual accuracy will be positive related to behavioral cultural fit but unrelated to voluntary exit, using OLS regressions based on cross-sectional data, as well as fixed effect regressions based on longitudinal data (including imputed measures of perceptual accuracy 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.

To test Hypothesis 3, which suggests that perceptual accuracy and behavioral cultural fit (but not value congruence) 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:

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

where  $CF_{idt}$  is the relevant cultural fit measure (behavioral fit, perceptual accuracy or value congruence) for individual  $i$  in department  $d$  at time  $t$ ,  $\langle PeerCF \rangle_{idt-1}$  is the mean peer cultural fit at time  $t - 1$  weighted by number of incoming messages,  $|Peer|_{idt-1}$  is the number of peers at time  $t - 1$ , and  $X$  are time-varying individual attributes. The inclusion of individual fixed effects accounts for stable variation between individuals, such as differences in innate psychological traits, experience, and preferences. Department and year fixed effects account, respectively, for differences between departments (e.g., different demographic compositions) and periods (e.g. variation in turnover rates) that might systematically affect cultural fit.

We lag mean peer cultural fit and number of peers to ensure appropriate temporal ordering. Yet even with individual fixed effects and lagged predictors, this modeling approach does not yield causal estimates. It could be the case, for example, that individuals with high cultural 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 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 an exogenous shock 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 cultural fit and number of peers introduced by the reorganization, and we then used these estimates to predict subsequent changes in cultural fit.

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 cultural 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 cultural fit*,  $\tilde{\Delta}\langle PeerCF \rangle$ , as an instrument.  $\tilde{\Delta}\langle PeerCF \rangle$  is the change induced by the reorganization between periods  $t - 1$  and  $t$ , *assuming peer cultural 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' cultural fit.

In addition to induced change in mean peer cultural 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:

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

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

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

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 mean peer cultural fit, and the interaction between the two—to estimate the model's two endogenous variables, mean peer cultural fit and number of peers. In the first stage we estimated the following regressions:

$$\begin{aligned} \langle PeerCF \rangle'_{idt} = & \beta_0 + \beta_1 NC(I_{it}, I_{it-1}) + \beta_2 \tilde{\Delta}\langle PeerCF \rangle_{idt-1} \\ & + \beta_3 NC(I_{it}, I_{it-1}) \cdot \tilde{\Delta}\langle PeerCF \rangle_{idt-1} + \beta_4 Ind_{\cdot i} + \epsilon_{it} \quad (4) \end{aligned}$$

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

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

$$CF_{idt+1} = \beta_0 + \beta_1 \langle PeerCF \rangle'_{idt} + \beta_2 |Peer|'_{idt} + \beta_3 Year_t + \beta_4 Dept_d + \beta_5 Ind \cdot_i + \eta X_{it} + \epsilon_{idt} \quad (6)$$

where  $X_{it}$  represents time-varying individual controls. We report results from eq. 6 in the tables below.

## Results

### Preliminary Analyses—Evaluating the Variables of Interest

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

[FIGURE 2 ABOUT HERE]

Second, in Table 1 we report the results of OLS regressions with individual, department and year fixed effects, where the dependent variable is bonus (logged) and independent variables—behavioral fit, perceptual accuracy (imputed) and value congruence (imputed)—are lagged. The fixed effects specification with lagged predictors allows us to estimate the effects of within-person change in cultural fit on subsequent productivity.

Whether modeled independently or together, all three cultural fit measures are significantly positively related to productivity. Thus we find, consistent with prior work (Chatman 1991, Srivastava et al. 2018), that behavioral cultural congruity, as well as cognitive alignment, are positively related to positive job performance—even when we use imputed longitudinal measures of cognitive fit. The coefficients for behavioral fit and perceptual accuracy are of similar magnitude. The two variables retain their significance even when included together in Model 4.

In contrast, the effect of value congruence on bonus is more modest. This result is consistent with our expectation that value congruence remains more stable over time. Given that the unwavering component of value congruence is subsumed in the individual fixed effect, it is not surprising that its time-varying component accounts for less of the variance in job performance. Overall, these supplemental analyses help to validate the longitudinal fit measures derived from our imputation methodology.

### Main Results

Table 2 reports tests of Hypothesis 1. Our competing risks Cox hazard models focus on voluntary exit as a function of value congruence and perceptual accuracy (with involuntary exit serving as the competing risk).

As Table 2 indicates, value congruence is associated with a decreased risk of voluntary exit, while perceptual accuracy is not. The importance of value congruence in affecting voluntary departures, based on the imputed longitudinal measure, is consistent with prior work based on a cross-sectional measure of value congruence that predicted departure from firms up to two years later (Chatman 1991).<sup>7</sup>



Table 3 provides a test Hypothesis 2. The first three models report results from cross-sectional data where the cognitive fit measures—perceptual accuracy and value congruence—are derived directly from the Organizational Culture Profile (OCP). Both measures are imputed in the three longitudinal models that follow.

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

[TABLE 3 ABOUT HERE.]

Table 3, Models 4 to 6, echo the results from the cross-sectional analyses in longitudinal specifications that include individual, department, and year fixed effects. The longitudinal results provide further support for Hypotheses 1 and 2 given that perceptual accuracy is significantly related to behavioral fit, while value congruence is not. As individuals' perceptual accuracy increases, their behavioral fit correspondingly increases. Changes in value congruence, in contrast, are unrelated to changes in behavioral fit as measured by language accommodation.

Of the control variables included in the models, only managerial status and tenure are significant. We conjecture that managers exhibit greater behavioral 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. Consistent with previous work on enculturation (Srivastava et al. 2018), we also find that individuals exhibit significantly lower behavioral fit during their first year in the organization.<sup>8</sup>

Table 4 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 accuracy and hence behavioral fit for the focal individual. Model 1 presents estimates

from the baseline fixed effect models with lagged peer behavioral fit, as specified in eq. 1. Individuals exhibit a significant increase in behavioral fit when their peers' mean behavioral fit increases in the preceding month. Importantly, this model includes individual fixed effects and thus accounts for a wide range of time-invariant individual differences—such as self-monitoring or cultural capital—that might also affect a person's capacity for behavioral fit.

[TABLE 4 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 result is reported in Model 2. 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 3. 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, is disruptive to cultural integration. The difference between these coefficients in the OLS (Model 1) and instrumental variable (Model 2) models highlights the importance of causal identification in this context. 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. Models 3

and 4—wherein we estimate the effects of change in peer behavioral fit on the focal individual’s perceptual accuracy and value congruence, respectively—suggest that behavioral adjustment occurs through changes in perceptual accuracy rather than through value congruence. We conjecture 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 accuracy and in peer value congruence on the focal individual’s perceptual accuracy and value congruence, respectively. Both coefficients are insignificant, 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’ cognitive cultural fit. To the extent that they do, for example, when they explicitly discuss their beliefs, it does not appear to be sufficiently potent to translate into changes in their own cognition.

In Table 5, 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 cognitive and behavioral cultural fit 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 and peer behavioral fit measures that were based on the reference group of *all* employees in the organization. Table 5, Model 1, shows that peer behavioral fit, when peers are defined as all other employees in the organization, predicts the focal actor’s behavioral fit relative to this same reference group. This result helps mitigate concerns that our main results are an artifact of our choice to define behavioral 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 5, Model 2, illustrates, our hypothesized effects hold even for this more restricted sample of employees. By removing individuals with supervisory responsibilities, this analysis also offers insight into whether language accommodation, our measure of behavior 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 5 ABOUT HERE.]

## Discussion and Conclusion

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

Prior research has offered competing explanations for why some people fit in better than others. One perspective has highlighted the importance of alignment between individual and group values in shaping behavior, while another has emphasized the role of situational cues and the ability to read the group's cultural code. We develop a theoretical account that reconciles these competing

perspectives. Drawing on dual-process theories of culture and cognition and the distinction between constrained and unconstrained situations, we develop a situated theory of cultural fit. We argue and find empirical support for the notion that values matter for behavior in unconstrained situations—in particular, for the choice to remain at or voluntarily exit from the organization. Perceptual accuracy instead matters for behavior in constrained situations—specifically, for the capacity to exhibit behavioral fit, here assessed in terms of real-time linguistic conformity with peers. We further show that a person’s behavior and perceptual accuracy are both influenced by observations of others’ behavior, whereas value congruence is less susceptible to peer influence.

While we develop a novel, situated theory of cultural fit and bring together disparate forms of data and analytical methods, we also acknowledge certain limitations. First, our study is based on data from a single organization, which raises questions about the extent to which the findings would generalize to other settings. Second, although we theorize about situational constraint, we do not directly observe or measure constraint and instead rely on the assumption that certain behaviors (voluntary exit) are less constrained than others (interactional language use). Future research might identify additional behaviors that are more or less constrained manifestations of fit and misfit. For example, the choice to become a whistleblower when observing organizational misconduct might constitute unconstrained behavior since prevailing norms typically favor displaying loyalty to the organization rather than challenging it. Finally, our imputation models rely on the implicit assumption that the relationship between language use and the relevant cultural fit variables is stable over time. As such, future studies that include multiple administrations of the OCP are needed to validate this assumption.

These limitations notwithstanding, our theoretical framework and concomitant findings offer four contributions. The first is in advancing person-culture fit theory. Specifically, we demonstrate that the antecedents and behavioral consequences of cultural fit vary by the type of situation a person faces and are associated with different modes of cognition. Values matter for cultural conformity in unconstrained situations and, via explicit cognition and propositional knowledge, shape outcomes

such as the choice to leave an organization. In contrast, perceptions are important for cultural alignment in constrained situations and, through implicit cognition and practical knowledge, yield real-time behavioral conformity. Together, these insights open the door to further investigations of the role that situations can play in shaping how people fit into social groups. Next, we demonstrate that both cognitive manifestations of cultural fit—value congruence and perceptual accuracy—as well as its behavioral form—linguistic conformity with peers—enable people to reap positive career rewards. Indeed, all three of our fit measures are positively linked to individual productivity, as indicated by bonus payments. In contrast to prior theoretical formulations of fit, we also demonstrate that different rewards accrue to different forms of cultural alignment. Those who read the code correctly and behave accordingly benefit from being perceived as true and committed group members, while those who identify with and embrace the code enjoy the psychological well-being that comes with a positive self-concept.

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

Our second contribution pertains to cultural change and transmission: 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 perspective-taking 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 also of the peers from whom this code is learned. In this sense, a person’s structural position in an organization is highly consequential regardless of her intrinsic ability to detect the cultural code. If she is connected to peers whose perceptions of the culture are inaccurate and who therefore behave in non-compliant ways, she will find it harder to exhibit normatively compliant behavior herself.

The link we establish between peers’ behaviors and those of the focal actor also contributes to our understanding of cultural diffusion. Previous work has argued that some innate aspects of “cultural intelligence” make individuals sensitive to cultural knowledge in others’ behaviors (Liebal et al. 2013). The literature on social networks, in contrast, has mostly focused on the structural conditions that enable or impede behavioral diffusion. We combine insights from these otherwise disconnected research domains to make two interrelated contributions. First, we theorize and demonstrate empirically that cultural transmission is a function not only of individuals’ attentiveness to cultural knowledge in others’ behaviors but also of the structural conditions that lead and expose them to others. Second, our theory offers a novel perspective on how this process of cultural diffusion operates, first and foremost, by primarily affecting perceptions rather than values.

Third, we contribute to dual-process theories of culture and cognition (Vaisey 2009, Miles 2015, Lizardo et al. 2016) in two key ways. First, we make a conceptual link between modes of cognition—based on implicit and practical knowledge versus explicit and propositional knowledge (Lizardo



2017)—and the types of situations a person faces—constrained versus unconstrained. Whereas previous work in this tradition has thought about the link between values and behavior in binary terms—i.e., values either do or do not shape behavior—we develop a more nuanced account of the relationship by fusing dual-process models with a theory of situations. Our results indicate that values matter for behavior in certain situations (unconstrained) but not others (constrained). This insight paves the way for exploring more generally how values matter when people frame situations in different ways—for example, variation in how people think of others as representatives of person, role, and character categories (Diehl and McFarland 2010). Second, although dual-process theories of culture in action have proliferated, the empirical evidence in support of their link to concrete behaviors remains scant. We add to this evidence base by establishing a clear link between cultural fit constructs that are tied to implicit versus explicit cognition and consequential behaviors such as how people communicate with their colleagues, their choice of voluntary exit, and their level of work productivity (as reflected in bonus payments).

Finally, through this work, we make a methodological contribution that would appear to 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, selectively appropriate, with requirements that include having a sufficient number of survey observations, access to rich communication content, protocols and safeguards to protect individual privacy and company confidentiality, and significant computational bandwidth. 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 cognitive and behavioral arenas of social life relate to one another and jointly shape the life course and the cultures in which it unfolds.

## Endnotes

<sup>1</sup>A robust literature in psychology often refers to situations that restrict participants' behavior as exhibiting high situational strength (Meyer et al. 2010b). We prefer characterizing such situations as “constrained,” rather than “strong,” as the latter implies subjective salience that pressures people into behavioral compliance. Yet a situation—for example, a routine email exchange with colleagues on a team—can be constrained even when a person experiences it as mundane and insignificant. In such cases, the constraint operates through habitual and less reflective, almost instinctual, action. We also note that constrained and unconstrained are not binary categories: situations can vary along a continuum of constraint. For example, some aspects of the decision to exit an organization voluntarily may be constrained. Yet such a consequential decision is, by and large, relatively unconstrained by others' behavioral cues and signals. In contrast, communication with peers is constrained by complex and layered norms of interaction.

<sup>2</sup>The required distribution of statements across categories that range from least to most characteristic of a given value is 2-4-6-9-12-9-6-4-2.

<sup>3</sup>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.

<sup>4</sup>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.

<sup>5</sup>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.

<sup>6</sup>Because including period fixed effects produces unstable estimates in such a model, we instead include the number of employees in the organization as a control. This accounts for time-varying

fluctuations in average value congruence due to firm growth or decline. To account for variation in the number of observations per individual (some individuals remain only a handful of months in the organization, whereas others stay for years), we use overall tenure as a sampling weight.

<sup>7</sup>Neither perceptual accuracy nor value congruence is significant in predicting involuntary exit when we use the same framework with voluntary exit as the competing risk.

<sup>8</sup>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

- Alesina A, Algan Y, Cahuc P, Giuliano P (2015) Family values and the regulation of labor. *Journal of the European Economic Association* 13(4):599–630.
- 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.
- Ashforth BK, Saks AM (1996) Socialization tactics: Longitudinal effects on newcomer adjustment. *Academy of Management Journal* 39(1):149–178.
- Bail CA (2017) Taming big data: Using app technology to study organizational behavior on social media. *Sociological Methods & Research* 46(2):189–217.
- Bail CA, Argyle LP, Brown TW, Bumpus JP, Chen H, Hunzaker MF, Lee J, Mann M, Merhout F, Volfovsky A (2018) Exposure to opposing views on social media can increase political polarization. *Proceedings of the National Academy of Sciences* 115(37):9216–9221.
- 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.

- 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.
- Block J (1961) *The Q-sort Method in Personality Assessment and Psychiatric Research* (Thomas Springfield, IL).
- Boltanski L, Thévenot (2006) *On Justification: Economies of Worth* (Princeton University Press).
- Cable DM, Judge TA (1996) Person–organization fit, job choice decisions, and organizational entry. *Organizational Behavior and Human Decision Processes* 67(3):294–311.
- 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.
- Chatman JA, Spataro SE (2005) Using self-categorization theory to understand relational demography–based variations in people’s responsiveness to organizational culture. *Academy of Management Journal* 48(2):321–331.
- Childress CC, Friedkin NE (2012) Cultural reception and production: The social construction of meaning in book clubs. *American Sociological Review* 77(1):45–68.
- Dandekar P, Goel A, Lee DT (2013) Biased assimilation, homophily, and the dynamics of polarization. *Proceedings of the National Academy of Sciences* 110(15):5791–5796.

- 
- Davis-Blake A, Pfeffer J (1989) Just a mirage: The search for dispositional effects in organizational research. *Academy of Management Review* 14(3):385–400.
- Diehl D, McFarland D (2010) Toward a historical sociology of social situations. *American Journal of Sociology* 115(6):1713–1752.
- DiMaggio P, Goldberg A (2018) Searching for homo economicus: Variation in americans construals of and attitudes toward markets. *European Journal of Sociology* 139.
- 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.
- Edwards JR (2008) Person-environment fit in organizations. *The Academy of Management Annals* 2(1):167–230.
- Elfenbein HA, O'Reilly CA (2007) Fitting in: The effects of relational demography and person-culture fit on group process and performance. *Group & Organization Management* 32(1):109–142.
- Eliasoph N, Lichterman P (2003) Culture in interaction. *American Journal of Sociology* 108(4):735–794.
- Evans JA, Aceves P (2016) Machine translation: Mining text for social theory. *Annual Review of Sociology* 42:21–50.
- 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.
- Friedland R, Alford RR (1991) *Bringing Society Back in: Symbols, Practices and Institutional Contradictions* (University of Chicago Press).
- Friedman J, Hastie T, Tibshirani R (2001) *The Elements of Statistical Learning*, volume 1 (Springer series in statistics New York).
- Gerald JH, George SH (2010) Self-report: Psychology's four-letter word. *American Journal of Psychology* 123(2):181–188.

- Gibson DR (2011) Avoiding catastrophe: The interactional production of possibility during the cuban missile crisis. *American Journal of Sociology* 117(2):361–419.
- Goldberg A (2011) Mapping shared understandings using relational class analysis: The case of the cultural omnivore reexamined. *American Journal of Sociology* 116(5):1397–1436.
- Goldberg A, Srivastava SB, Manian GV, Monroe W, Potts C (2016) Fitting in or standing out? the tradeoffs of structural and cultural embeddedness. *American Sociological Review* 81(6):1190–1222.
- Grandey AA (2000) Emotional regulation in the workplace: A new way to conceptualize emotional labor. *Journal of Occupational Health Psychology* 5(1):95.
- Greenwald AG, Banaji MR (1995) Implicit social cognition: Attitudes, self-esteem, and stereotypes. *Psychological Review* 102(1):4.
- Hall DT (1996) *The Career Is Dead—Long Live the Career. A Relational Approach to Careers. The Jossey-Bass Business & Management Series.* (ERIC).
- Harding DJ (2007) Cultural context, sexual behavior, and romantic relationships in disadvantaged neighborhoods. *American Sociological Review* 72(3):341–364.
- Harrison JR, Carroll G (2006) *Culture and demography in organizations* (Princeton, NJ: Princeton University Press).
- Herrmann PA, Legare CH, Harris PL, Whitehouse H (2013) Stick to the script: The effect of witnessing multiple actors on children’s imitation. *Cognition* 129(3):536–543.
- Ho TK (1995) Random decision forests. *Document Analysis and Recognition, 1995., Proceedings of the Third International Conference on*, volume 1, 278–282 (IEEE).
- Hochschild AR (2012) *The Managed Heart: Commercialization of Human Feeling* (Univ of California Press).
- Hofstede G (2001) *Culture’s Consequences: Comparing Values, Behaviors, Institutions and Organizations Across Nations* (SAGE Publications).
- Inglehart R, Baker WE (2000) Modernization, cultural change, and the persistence of traditional values. *American Sociological Review* 65(1):19–51.
- Judge TA, Cable DM (1997) Applicant personality, organizational culture, and organization attraction. *Personnel Psychology* 50(2):359–394.

- 
- Keister LA (2008) Conservative protestants and wealth: How religion perpetuates asset poverty. *American Journal of Sociology* 113(5):1237–1271.
- Kilduff M, Day DV (1994) Do chameleons get ahead? the effects of self-monitoring on managerial careers. *Academy of Management Journal* 37(4):1047–1060.
- 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.
- Leschziner V, Green AI (2013) Thinking about food and sex: Deliberate cognition in the routine practices of a field. *Sociological Theory* 31(2):116–144.
- Leung K, Morris MW (2015) Values, schemas, and norms in the culture–behavior nexus: A situated dynamics framework. *Journal of International Business Studies* 46(9):1028–1050.
- Liebal K, Carpenter M, Tomasello M (2013) Young children’s understanding of cultural common ground. *British Journal of Developmental Psychology* 31:88–96.
- 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.
- Lizardo O (2017) Improving cultural analysis: Considering personal culture in its declarative and nondeclarative modes. *American Sociological Review* 82(1):88–115.
- Lizardo O, Mowry R, Sepulvado B, Stoltz DS, Taylor MA, Van Ness J, Wood M (2016) What are dual process models? implications for cultural analysis in sociology. *Sociological Theory* 34(4):287–310.

- Lord CG, Ross L, Lepper MR (1979) Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence. *Journal of Personality and Social Psychology* 37(11):2098.
- 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.
- Meglino BM, Ravlin EC (1998) Individual values in organizations: Concepts, controversies, and research. *Journal of Management* 24(3):351 – 389, ISSN 0149-2063, URL [http://dx.doi.org/https://doi.org/10.1016/S0149-2063\(99\)80065-8](http://dx.doi.org/https://doi.org/10.1016/S0149-2063(99)80065-8).
- Meyer JP, Hecht TD, Gill H, Toplonysky L (2010a) Person–organization (culture) fit and employee commitment under conditions of organizational change: A longitudinal study. *Journal of Vocational Behavior* 76(3):458–473.
- Meyer RD, Dalal RS, Hermida R (2010b) A review and synthesis of situational strength in the organizational sciences. *Journal of Management* 36(1):121–140.
- Miles A (2015) The (re)genesis of values: Examining the importance of values for action. *American Sociological Review* 80(4):680–704.
- Morris MW, Chiu Cy, Liu Z (2015) Polycultural psychology. *Annual Review of Psychology* 66:631–659.
- Morrison EW (2002) Newcomers’ relationships: The role of social network ties during socialization. *Academy of Management Journal* 45(6):1149–1160.
- Nisbett RE, Cohen D (1996) *Culture of Honor: The Psychology of Violence in the South* (Westview Press).



- O'Reilly CA, Chatman J (1986) Organizational commitment and psychological attachment: The effects of compliance, identification, and internalization on prosocial behavior. *Journal of Applied Psychology* 71(3):492.
- O'Reilly CA, Chatman J, Caldwell DF (1991) People and organizational culture: A profile comparison approach to assessing person-organization fit. *Academy of Management Journal* 34(3):487–516.
- Parsons T (1968) *The Structure of Social Action* (Free press).
- Pennebaker JW (2013) *The Secret Life of Pronouns: What Our Words Say About Us* (Bloomsbury USA), ISBN 9781608194964, URL <https://books.google.com/books?id=p9KmCAAAQBAJ>.
- Pennebaker JW, Chung CK, Ireland M, Gonzales A, Booth RJ (2007) *Linguistic Inquiry and Word Count (LIWC): LIWC2007* (Austin, TX: LIWC.net).
- Pinker S (2007) *The Stuff of Thought: Language as a Window Into Human Nature* (New York, NY: Viking).
- Potts C (2011) Sentiment-aware tokenizer. *Creative Commons Attribution-NonCommercial-ShareAlike 3.0 Unported License*: <http://creativecommons.org/licenses/by-nc-sa/3.0/> .
- Rivera LA (2012) Hiring as cultural matching: The case of elite professional service firms. *American Sociological Review* 77(6):999–1022.
- Ross L (1977) The intuitive psychologist and his shortcomings: Distortions in the attribution process. *Advances in Experimental Social Psychology* 10:173–220.
- Salancik GR, Pfeffer J (1974) The bases and use of power in organizational decision making: The case of a university. *Administrative Science Quarterly* 453–473.
- Salganik MJ (2017) *Bit by Bit: Social Research in the Digital Age* (Princeton University Press).
- Sarros JC, Gray J, Densten IL, Cooper B (2005) The organizational culture profile revisited and revised: An Australian perspective. *Australian Journal of Management* 30(1):159–182.
- 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.
- Schneider B (1987) The people make the place. *Personnel Psychology* 40(3):437–453.

- Schwartz SH (1992) Universals in the content and structure of values: Theoretical advances and empirical tests in 20 countries. Zanna MP, ed., *Advances in Experimental Social Psychology*, volume 25, 1 – 65 (Academic Press).
- Scott BA, Barnes CM, Wagner DT (2012) Chameleonic or consistent? a multilevel investigation of emotional labor variability and self-monitoring. *Academy of Management Journal* 55(4):905–926.
- Snyder M (1979) Self-monitoring processes. *Advances in Experimental Social Psychology* 12:85–128.
- Sørensen JB (2002) The strength of corporate culture and the reliability of firm performance. *Administrative Science Quarterly* 47(1):70–91.
- Srivastava SB, Banaji MR (2011) Culture, cognition, and collaborative networks in organizations. *American Sociological Review* 76(2):207–233.
- Srivastava SB, Goldberg A, Manian GV, Potts C (2018) Enculturation trajectories: Language, cultural adaptation, and individual outcomes in organizations. *Management Science* 64(3):1348–1364.
- Swidler A (1986) Culture in action: Symbols and strategies. *American Sociological Review* 273–286.
- Tausczik YR, Pennebaker JW (2010) The psychological meaning of words: Liwc and computerized text analysis methods. *Journal of language and social psychology* 29(1):24–54.
- Thompson JD (1967) *Organizations in Action: Social Science Bases of Administrative Theory* (Transaction publishers).
- Vaisey S (2009) Motivation and justification: A dual-process model of culture in action. *American Journal of Sociology* 114(6):1675–1715.
- Vaisey S, Lizardo O (2016) Cultural fragmentation or acquired dispositions? a new approach to accounting for patterns of cultural change. *Socius* 2.
- Van Maanen J (1975) *Breaking In: Socialization to Work* (pp. 67-130 in Robert Dubin (ed.) *Handbook of Work, Organization and Society*).
- Van Maanen J, Schein E (1979) Toward a theory of organizational socialization. *Research in Organizational Behavior* 1:209–264.
- 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.

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

## FIGURES




Figure 1 Conceptual Overview of the Machine Learning Process

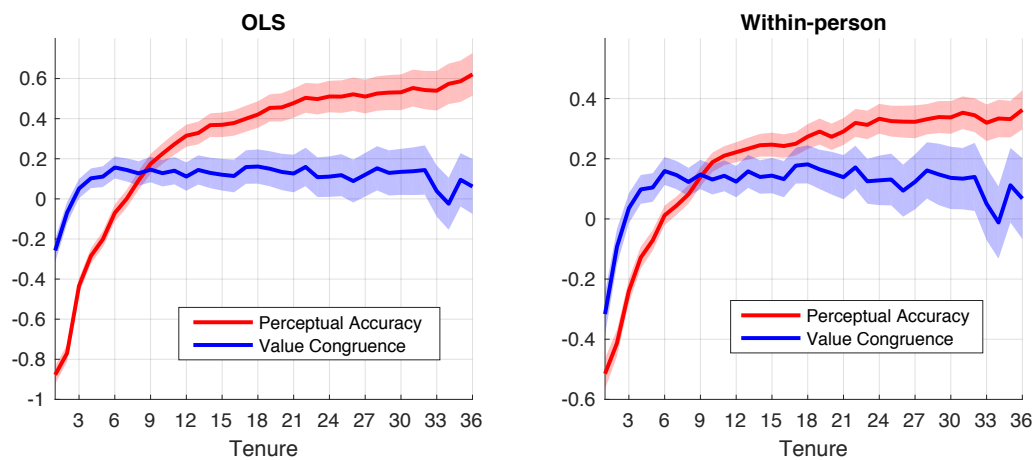



Figure 2 OLS and fixed effect regressions of perceptual accuracy and value congruence, with indicators for each tenure month up to 36 months in the company.



**Figure 3** 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 (blue) or decrease (red) in peer behavioral fit. Shaded areas correspond to 95% confidence intervals.

## TABLES

**Table 1** Fixed Effect Regressions of Bonus (logged)

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

*t* statistics in parentheses; standard errors clustered by individual

<sup>†</sup> lagged variables, \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$

**Table 2** Competing Risks Model of Voluntary Exit

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

Exponentiated coefficients; *t* statistics in parentheses

Standard errors clustered by individual; Sample weights by tenure

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

**Table 3** Cross-Sectional and Longitudinal Fixed Effects Regressions of Behavioral Fit

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

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

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

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



**Table 4** OLS and Instrumental Variables Fixed Effects Regressions of Behavioral Fit

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

*t* statistics in parentheses; standard errors clustered by individual

<sup>†</sup> lagged variables, instrumented endogenous variables in Models 2-6

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

**Table 5 Robustness Checks—Instrumental Variables Fixed Effect Regressions of Behavioral Fit**

	Model 1 Organization	Model 2 Non-Managers
Peer Behavioral Fit <sup>†</sup>		0.235*** (5.78)
Peer Behavioral Fit (Organization) <sup>†</sup>	0.158*** (5.40)	
Num. Peers <sup>†</sup>	-0.003 (-1.85)	-0.013* (-2.10)
Manager	0.133*** (3.57)	
First Year	-0.034* (-2.27)	-0.150** (-3.25)
Constant	2.154*** (26.90)	-0.560 (-1.79)
Individual FE	Yes	Yes
Department FE	Yes	Yes
Year FE	Yes	Yes
N	19938	18097
Num. Individuals	1229	1257
Kleibergen-Paap F	3.03	8.81

*t* statistics in parentheses; standard errors clustered by individual

<sup>†</sup> instrumented and lagged endogenous variables

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

## APPENDIX A: BEHAVIORAL CULTURAL 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  $\vec{m}_{it}$  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  $\vec{m}_{it}^l$  and  $\overleftarrow{m}_{it}^l$  for the count of terms in email  $\vec{m}_{it}$  and  $\overleftarrow{m}_{it}$  in LIWC category  $l \in L$ , respectively. Then, by aggregating all individual email counts  $\vec{m}_{it}^l$  and  $\overleftarrow{m}_{it}^l$  for  $t \in T$ , it produces sent and received LIWC counts in month  $T$ ,  $\vec{m}_{iT}^l$  and  $\overleftarrow{m}_{iT}^l$ . 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,  $O_{iT}^l$  to denote the outgoing normalized probability and  $I_{iT}^l$  to denote the incoming normalized probability.

$$O_{iT}^l = \frac{\vec{m}_{iT}^l}{\sum_{l \in L} \vec{m}_{iT}^l} \quad (7)$$

$$I_{iT}^l = \frac{\overleftarrow{m}_{iT}^l}{\sum_{l \in L} \overleftarrow{m}_{iT}^l} \quad (8)$$

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:

$$BF_{iT} = -\log(JS(O_{iT} \| I_{iT})) \quad (9)$$

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.

$$JS(O_{iT} \| I_{iT}) = \frac{1}{2}KL(O_{iT} \| M_{iT}) + \frac{1}{2}KL(I_{iT} \| M_{iT}) \quad (10)$$

$$KL(D_{iT} \| M_{iT}) = \sum_{l \in L} D_{iT}^l \log_2 \frac{D_{iT}^l}{M_{iT}^l} \quad (11)$$

## 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 oriented 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. We additionally find that those who perceive the organizational culture as results oriented tend to send fewer feel words (e.g., feels, touch) and health words (e.g., clinic, flu, pill) and also tend to receive fewer discrepancy words (e.g., should, would, could) and future tense words (e.g., will, gonna)<sup>9</sup>. 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 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 and his colleagues (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 A1. 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 A1 indicates that even if certain LIWC categories are culturally irrelevant in this context, their inclusion in the measure construction does not bias its value. In other words, even if we were to conclude that half of the LIWC categories are non-cultural (a conclusion that, for the reasons stated above, we believe is unwarranted) and decide to remove them from the measure, we would still recover near-identical values.

## APPENDIX B: MACHINE LEARNING PROCEDURE

### Overview

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

Our first step was to translate the raw email data into a format that is usable by the random forest model. We tokenized and stemmed all words in the body of email messages. Tokenization involves separating the text into distinct terms, for which we used the `TwitterTokenizer` designed for linguistic analysis Potts (2011). Stemming involves reducing each term to a root form, for which we used the Porter Stemmer from the python `nlk` 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 99% 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 accuracy 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 accuracy, 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 accuracy 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 accuracy 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 this.

[FIGURE C1 ABOUT HERE.]

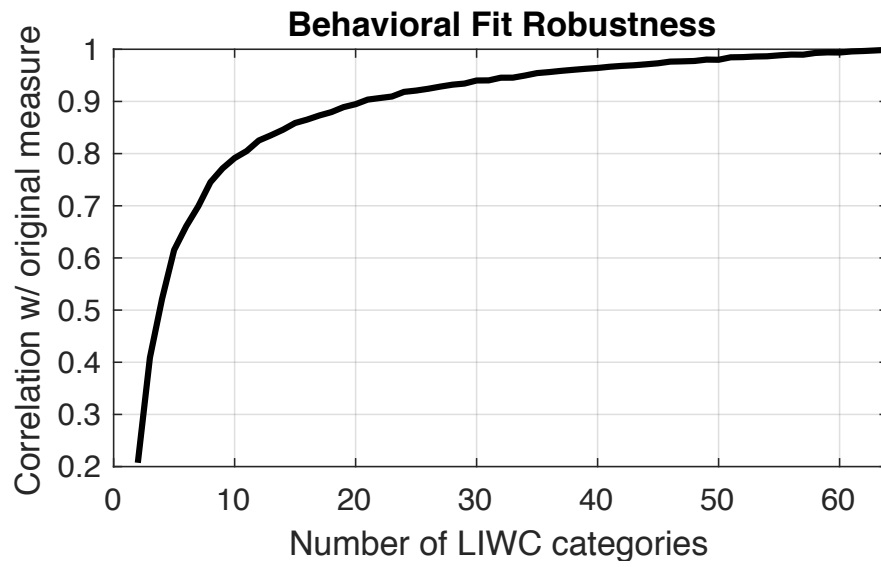
We then used the means and standard deviations of each group to see if the classifier successfully split the observations into statistically distinct groups. We find that the separation between low and high in our models is good.

[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 FIGURES



**Figure A1** Robustness of the behavioral fit measure to LIWC category composition




Figure C1 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
PA-Interloc.	0.857	0.726	0.767	0.267	0.651	0.711
PA-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
PA-Interloc.	2.661e-3
PA-Org.	1.874e-8
VC-Interloc.	8.500e-6
VC-Org.	7.157e-5

**Table C3** Areas under the ROC Curve

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