Deciphering the Cultural Code: Cognition, Behavior, and the Interpersonal Transmission of Culture*

Richard Lu[a], Jennifer A. Chatman[a], Amir Goldberg[b], and Sameer B. Srivastava[a]

[a]University of California, Berkeley [b]Stanford University

Abstract

Why are some people more successful than others at cultural adjustment? Research on organizational culture has mostly focused on value congruence as the core dimension of cultural fit. We develop a novel and complementary conceptualization of cognitive fit—perceptual accuracy, or the degree to which a person can decipher the group's cultural code. We demonstrate that the ability to read the cultural code, rather than identification with the code, matters for contemporaneous behavioral conformity. 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 cultural assimilation, the interplay of structure and culture, and the pairing of surveys with digital trace data.

Keywords— Culture, Person-Culture Fit, Cognition, Behavior, Language

July 5, 2018

*Authors after the first are listed in alphabetical order. We thank participants of the International Conference on Computational Computer Science, the Academy of Management Annual Meeting, the Wharton People and Organizations Conference, the Stanford Networks and Culture Conference, and the Berkeley Information-School Research Exchange for helpful comments on prior drafts. We also thank the Ryoichi Sasakawa Young Leaders Fellowship Fund for providing financial support for this work. Any requests can be sent to the corresponding author at richard_lu@haas.berkeley.edu, 404-435-2160, or 2220 Piedmont Ave, Berkeley, CA 94720. The usual disclaimer applies.

Deciphering the Cultural Code: Cognition, Behavior, and the Interpersonal Transmission of Culture

Abstract

Why are some people more successful than others at cultural adjustment? Research on organizational culture has mostly focused on value congruence as the core dimension of cultural fit. We develop a novel and complementary conceptualization of cognitive fit—perceptual accuracy, or the degree to which a person can decipher the group's cultural code. We demonstrate that the ability to read the cultural code, rather than identification with the code, matters for contemporaneous behavioral conformity. 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 cultural assimilation, the interplay of structure and culture, and the pairing of surveys with digital trace data.

Keywords— Culture, Person-Culture Fit, Cognition, Behavior, Language

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., 2010; 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 in adjusting their cultural fit over time?

Existing literature 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 individuals' internalization of their groups' behavioral norms. Those who fit in are therefore those whose preferences are consistent with the norms that are prevalent 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 her 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 others' actions. An employee's decision to use formal language in a meeting, for example, is less a function of her beliefs on the virtues of hierarchy and more a reaction to how others behave in this setting. People pursue action for which their "cultural equipment is well suited" (**?**, 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 approaches appear to provide incompatible explanations for the sources of cultural fit in organizations: whereas the former suggests that cultural fit is the result of internalizing organizational culture, the latter sees it as the product of correctly deciphering the cultural code. But recent work in cognitive psychology and cultural sociology demonstrates that values and interpretations relate to different cognitive mechanisms and therefore affect behavior through different and potentially inconsistent pathways. Expressed values are inherently tied to self-understandings and are most likely to affect behavior when people make reflective non-routine choices. Situational readings, on the other hand, affect routine behavior mostly through habitual action.

Prior work has predominantly explored how the different mechanisms that connect cognition to action relate to fundamental behavioral outcomes, such as teenagers' academic achievements or

lifestyle choices. This theoretical position has been widely influential in the sociology of culture but has hitherto had little impact on the study of culture in organizations, a literature which has mostly explored cultural fit through the lens of value congruence. Extending this debate to organizational contexts, we introduce a new concept—*perceptual accuracy*, which we define as an individual's ability to accurately understand the group's prevailing values and norms.

Drawing on the distinction between value congruence and perceptual accuracy, we make two core arguments. First, we propose that these two dimensions of cultural fit have differing consequences for individual outcomes in the organization: value congruence predicts a person's self-identification with the organization, and therefore her choice to stay or exit, whereas perceptual accuracy affects her ability to behave in normatively compliant ways. Second, we argue that while values are a relatively stable aspect of cognition, perceptions are susceptible to social learning. Consequently, witnessing normatively compliant (or non-compliant) behavior among peers boosts (or diminishes) one's own perceptual accuracy and, in turn, one's capacity for normative compliance regardless of whether or not one subscribes to those norms. Those whose peers' behaviors are culturally appropriate are therefore more likely to exhibit normatively compliant behavior. In contrast, peers are inconsequential for value congruence.

To test our comprehensive theory of culture fit, we employ a multi-method empirical strategy that draws on survey data, eight years of internal email data, and personnel records from a mid-sized technology firm. We use the tools of computational linguistics and machine learning to transform the cross-sectional measures of perceptual accuracy and value congruence, which were assessed through a validated culture survey, into longitudinal measures and to develop measures of 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 identifying how our findings advance theories

of cultural fit in organizations.

WHAT DO VALUES MATTER FOR?

Arguments about culture typically make implicit assumptions about underlying cognitive processes (DiMaggio, 1997). In most everyday settings, one's private cognition is, however, unavailable to others. Rather, one observes others' behavior and then draws inferences—with varying degrees of accuracy—about their beliefs, values, and motivations (Schein, 2010; Sperber, 1996).

What underlying cognitive processes lead some people to behave in culturally appropriate ways more than others? A dominant line of work both in organizational psychology and sociology has highlighted the importance of shared values among organizational members (Ostroff and Judge, 2007; Edwards and Cable, 2009; Baron et al., 2001). By "value," we mean enduring beliefs about desired or undesired ways of working and interacting with others (e.g., "I prefer a friendly work environment"), as distinguished from situation-specific preferences (e.g., "I prefer having lunch before noon") (O'Reilly et al., 1991; Vaisey, 2009; Miles, 2015).

Work that focuses on values 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 (Morrison, 2002; Elfenbein and O'Reilly, 2007).

The notion that shared values lead to behavioral coordination has a long history in sociology and psychology, and ultimately rests on an assumption that people's behvaiors are strongly contrained by their values. But a persistent body of research finds that people's stated values

are poor predictors of their behaviors. Economically disadvantaged high-schoolers, for example, tend to express mainstream attitudes on educational achievement and sexual behavior, but adopt behaviors that appear to be inconsistent with these ideals. Work by sociologists of culture has therefore focused on culture as implicitly acquired practical knowledge that is activated in response to situational cues.

What underlying cognitive processes lead some people to behave in culturally appropriate ways more than others? A dominant line of work both in organizational psychology and sociology has highlighted the importance of shared values among organizational members (Ostroff and Judge, 2007; Edwards and Cable, 2009; Baron et al., 2001). This work has primarily 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 (Morrison, 2002; Elfenbein and O'Reilly, 2007).

But a persistent body of research finds that people's stated values are poor predictors of their behaviors. Economically disadvantaged high-schoolers, for example, tend to express mainstream attitudes on educational achievement and sexual behavior, but adopt behaviors that appear to be inconsistent with these ideals. Work by sociologists of culture has therefore focused on culture as implicitly acquired practical knowledge that is activated in response to situational cues.

Recent work explains this seeming paradox by pointing to the complex ways by which cognition relates to behavior.

Yet, in many cases, people can successfully interact with one another even when they do not share the same values. Work in organizational psychology (Hewlin, 2003; Hewlin et al., 2017)

and sociology (Hochschild, 2012) finds that people often behave in ways that are consistent with their social group's normative expectations even when these norms are incompatible with their own private beliefs. As Willer and his colleagues (2009) demonstrate, this ability to separate beliefs from behaviors can lead to the persistence of unpopular norms. The core distinction is between the beliefs people value personally and those they perceive to be widespread in the social group (cf. Goldstein et al., 2008). When group members believe that a behavior is prevalent—and consequently falsely infer that associated privately held values are also widespread—they accommodate those behaviors themselves and sanction those who fail to conform. The fear of being exposed as inauthentic or deviant motivates them to police the cultural order despite their private disagreement with it.

COGNITIVE AND BEHAVIORAL CULTURAL FIT

Arguments about culture typically make implicit assumptions about underlying cognitive and interpersonal processes (DiMaggio, 1997; O'Reilly and Chatman, 1996). Organizational researchers often define culture as "shared understandings," namely, similarities between individuals' beliefs, value systems, and interpretations.[1] In most everyday settings, one's private cognition is, however, unavailable to others. Rather, one observes others' behavior and then draws inferences—with varying degrees of accuracy—about their beliefs, values, and motivations (Kelley and Michela, 1980; Schein, 2010; Sperber, 1996).

Culture, in other words, resides both in the distribution of inner thoughts and observable behaviors across individuals. Cultural fit, by extension, can be thought of as comprising two related but distinct dimensions: cognitive cultural fit, or the degree of shared understanding between an individual and her peers, and behavioral cultural fit, or the extent to which an individual's behaviors are compliant with the group's normative expectations (Mobasseri et al., 2018).[2]

[1]The definition of culture as an analytical construct has long been a matter of debate by organizational researchers, and we do not attempt to fully resolve this debate here. "Shared understandings," in our view, is a useful shorthand in that it points to two important properties of culture: that it dwells in the similarities between the individuals who constitute a group and that these similarities relate to group members' mental representations of the world. Missing from this useful, albeit simple, definition is the idea that such shared understandings emerge through interpersonal interaction.

[2]We acknowledge that not all individuals seek to fit in behaviorally and that some people are more predisposed than

Previous work has focused on either cognitive or behavioral fit and implicitly assumed that the two correspond highly to one another. An extensive literature in organizational psychology has, for example, examined culture through the lens of person-environment fit, highlighting the importance of shared values among organizational members (Ostroff and Judge, 2007; Edwards and Cable, 2009). This work has primarily identified two core mechanisms that link cognitive cultural fit to individual attainment. 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). The second relates to the ease of interpersonal interaction and coordination. Culturally aligned individuals find it easier to interact with one another because they have mutually compatible expectations of behavior (Morrison, 2002; Elfenbein and O'Reilly, 2007). Findings by organizational sociologists are consistent with this view. Baron et al. (2001), for example, find that organizations that change their models of work and employment experience greater turnover, especially among those most committed to outmoded cultural blueprints.

Yet, in many cases, people can successfully interact with one another even when they do not share the same values. Work in organizational psychology (Hewlin, 2003; Hewlin et al., 2017) and sociology (Hochschild, 2012) finds that people often behave in ways that are consistent with their social group's normative expectations even when these norms are incompatible with their own private beliefs. As Willer and his colleagues (2009) demonstrate, this ability to separate beliefs from behaviors can lead to the persistence of unpopular norms. The core distinction is between the beliefs people value personally and those they perceive to be widespread in the social group (cf. Goldstein

others to engaging in non-compliant behavior. Although the need for uniqueness is most likely hard-wired, it is also balanced by the propensity for compliance and assimilation with important social groups (Leonardelli et al., 2010). Moreover, the tendency to conform is mediated by individual endowments: those with high status or who enjoy structural buffering by virtue of being embedded in a tight-knit community may under some circumstances reap the benefits of culturally non-compliant behavior while limiting its adverse consequences (Goldberg et al., 2016). On balance, however, behavioral conformity is generally beneficial such that people are, by and large, motivated to conform to the normative expectations of their social group (Miller and Prentice, 2016). Thus, we expect individuals to be attuned to their cultural environments and to respond to their peers' behaviors in their attempts to fit in.

et al., 2008). When group members believe that a behavior is prevalent—and consequently falsely infer that associated privately held values are also widespread—they accommodate those behaviors themselves and sanction those who fail to conform. The fear of being exposed as inauthentic or deviant motivates them to police the cultural order despite their private disagreement with it.

To understand how such a situation can arise, it is important to distinguish between two dimensions of cognition: preferences and construals. Whereas preferences define which behaviors are desirable, construals refer to the levels of abstraction and the associated mental representations that a person conjures when making sense of a situation. How an individual construes a social setting affects which of her preferences will be activated and ultimately what action she will pursue (Trope and Liberman, 2010). Shared understandings do not necessarily require that all group members hold the same preferences. Rather, to share understandings is, first and foremost, to construe daily experiences through similar interpretative lenses (Goldberg, 2011; DiMaggio and Goldberg, 2018).

Similar insights derive from symbolic interactionists' studies of interpersonal interaction (Goffman, 1959; Garfinkel, 1967). As long as group members have a shared understanding of a situation—including the social roles it implies, the behaviors appropriate to those roles, and the implicit meanings these behaviors convey—interactions between members can occur relatively seamlessly. Further, even when the group agrees about how a situation is construed, individual members can still craft their self-presentations in a manner that decouples their behavior from their privately held preferences. In the absence of situational agreement, however, interaction breaks down, leading to incompatibilities between one person's expectations and another's behavior. Under such circumstances private cognition is more likely to unintentionally leak into public behavior.

Value Congruence and Perceptual Accuracy

Preferences and construals are aspects of individual cognition; however, they become culturally meaningful when we consider an individual in relation to her social group. Value congruence

represents the cultural manifestation of preferences in that it reflects the match between what individuals prefer and what prevails in the social group. Perceptual accuracy is instead the cultural analogue of construals in that it indicates the degree of alignment between a person's perceptions and those of other group members.

More specifically, *value congruence* is the degree of similarity between an individual's own preferred values and those reported by others as being prevalent in the group. By "value," we mean enduring beliefs about desired or undesired ways of working and interacting with others (e.g., "I prefer a friendly work environment"), as distinguished from situation-specific preferences (e.g., "I prefer having lunch before noon") (O'Reilly et al., 1991; Vaisey, 2009; Miles, 2015). Note that value congruence relates to fit with the normative environment, irrespective of whether other group members privately hold the same preferences. In an "Emperor's New Clothes" dynamic of the kind that Centola et al. (2005) discuss, a person might have low value congruence if she prefers not to blindly defer to hierarchy when the prevailing norm is to defer to more senior colleagues.

People whose ideal preferences are compatible with those prevalent in their social environment find it easier to maintain a positive self-concept (Chatman and Barsade, 1995). Consequently, they identify more strongly with the organization and derive greater satisfaction from their interactions with others. We therefore expect value congruence to be primarily related to motivation and long-term attachment to the organization—as evidenced by a negative association between value congruence and the choice to exit the organization voluntarily.

We anticipate, however, that value congruence will be less consequential for a person's capacity to conform to her group's normative expectations of behavior. 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 information 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 normative context. Moreover, recent work by cultural sociologists suggests that individuals' stated beliefs and motives can be inherently decoupled from their practical and unself-

conscious behavioral decisions. There is often a disconnect between what people ideally desire and what they understand as contextually appropriate behaviors (Vaisey, 2009; Srivastava and Banaji, 2011; Lizardo, 2017).

Organizational researchers have relied heavily on the notion of value congruence in building a theory of person-culture fit (Chatman, 1991). We suggest that value congruence represents a useful yet incomplete conceptualization of person-culture fit. For example, it is possible that members can both interpret the cultural code and comply with it without necessarily holding values that are congruent with those of the organization. We therefore introduce a novel and complementary conceptualization of cognitive cultural fit: *perceptual accuracy*. We define perceptual accuracy as the extent to which an individual's assessment of the behaviors that are or are not normatively compliant with group members' expectations is consistent with the readings of her peers. Note that this accuracy does not relate to peers' private beliefs or preferences. Again, using the "Emperor's New Clothes" (Centola et al., 2005) metaphor, a perceptually accurate individual will correctly decipher that the appropriate behavior is to express admiration for the monarch's clothes, irrespective of whether she correctly perceives that the majority of her peers believe that the emperor is, in fact, naked. As we detail below, we conceptualize and operationalize perceptual accuracy at a high level of construal that relates to group norms and, as such, usefully informs members' behaviors across many relevant group situations.

Perceptual Accuracy and Behavioral Fit

Perceptual accuracy describes an individual's ability to decipher the cultural code implicit in others' behaviors. Although organizations often formalize their idealized values into cultural statements, interpreting the local normative environment is a subtle, complex, and ongoing cognitive task. 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. Correctly construing this behavior requires tacit and layered knowledge that connects behaviors, symbols,

and meanings to abstract cultural categories. Possessing this knowledge is essential to knowing how to behave appropriately. Therefore, we argue that perceptual accuracy is intimately related to the capacity to behave in culturally compliant ways.

Figure 1 illustrates these conceptual arguments and their behavioral implications. Imagine five possible values (labeled *a* to *e*) that people can espouse. The four individuals depicted in the diagram (labeled *A* to *D*) correspond to four hypothetical organizational members. Each individual is characterized by three distributions: her private values (*V*) and perceptions (*P*) and her public behaviors (*B*). The bars in the figure represent the degree to which a person personally espouses a given value (*V*), believes that value is widely shared by other organizational members (*P*), and behaves in accordance with the group's normative expectations related to the value (*B*). As noted above, only the behaviors of others are directly observable; their values and perceptions can only be indirectly inferred.

[FIGURE 1 ABOUT HERE]

Individual *A* in Figure 1 is perceptually accurate but value incongruent: her perceptions of the cultural code (*P*) are consistent with the majority of her peers', but the prevailing values are mostly inconsistent with her own (*V*). Nevertheless, her behavior mirrors her perceptions. Suppose that value *d* is conflict-orientation. Although *A* does not prefer a confrontational environment (her value for *d* is negative), she sees conflict as a common and legitimate behavior in the organization. She is consequently likely to express disagreement and negation in her interaction with others (as reflected in her tendency to exhibit behavior *d*). Individual *D*, in contrast, is also conflict-averse, but unlike *A* she misperceives the prevalence of conflict in the organization. Consequently, her behavior is incongruent with her peers'. She is more likely to be accommodating and apologetic, whereas her peers are confrontational. Although the four hypothetical individuals in the diagram espouse different values, only *D* is a behavioral misfit. Like *A*, individuals *B* and *C* behave in a normatively compliant way because they hold similarly accurate perceptions of the cultural code

despite the latter two being more value congruent than the former.

In sum, we argue that one dimension of cognitive cultural fit—perceptual accuracy—is closely linked to an individual's capacity for behavioral cultural fit, whereas the other dimension—value congruence—does not matter for contemporaneous behavioral fit but is instead related to self-identification and long-term attachment to the organization. Given that the latter expectation has already been established in prior work, our first hypothesis focuses on the novel construct of perceptual accuracy:

HYPOTHESIS 1 (H1): *Perceptual accuracy is positively related to behavioral cultural fit.*

THE INTERPERSONAL TRANSMISSION OF CULTURE

Contending that perceptual accuracy, rather than value congruence, predicts behavioral cultural fit shifts the analytical focus from heterogeneity between individuals' preferences and beliefs to differences in their ability to enculturate—that is, their ability to read and adapt to the cultural code. A prominent line of work has conceptualized 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). This perspective continues to implicitly guide personnel practices in the contemporary workplace. Many organizations emphasize cultural fit in the hiring phase, assuming that only certain individuals possess innate qualities or underlying values that make them a strong cultural match (Rivera, 2012). Yet cultural fit is a dynamic process: individuals are capable of adapting their behavior to the prevailing norms in an organization (Van Maanen and Schein, 1979; Chatman, 1991; Srivastava et al., 2018). People acquire this capability through ongoing socialization (Van Maanen, 1975; Ashforth and Saks, 1996).

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, more inclined to conform to it, and more likely to be perceived as authentic when they do.

Yet perceptual accuracy is also a matter of situational 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 expect perceptual accuracy to be a pliable dimension of cognitive cultural fit that is partially dependent on the social context in which an individual is embedded. Adjusting to the cultural code of a group is, by definition, a process of social learning, and 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 bearing 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 cognition—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 argued above that value congruence is not linked to contemporaneous behavior such as the choice to conform linguistically with discussion partners. It is also likely, we propose, to remain relatively stable given that individuals' deeply held values are encoded in implicit cognition and thus slower to change (Meglino and Ravlin, 1998; Vaisey, 2009; Srivastava and Banaji, 2011; Vaisey and Lizardo, 2016). We therefore expect that value congruence will be less susceptible to peer influence than will perceptual accuracy.

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). 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, beliefs, and preferences.

The causal assumptions informing this model are depicted in the arrows in Figure 1. Individual A observes B's behavior and updates her perceptions accordingly. These perceptions, in turn, affect how she behaves. Her values, in contrast, remain relatively unchanged. Overall, we expect:

HYPOTHESIS 2 (H2): *Perceptual accuracy and behavioral fit are both susceptible to peer influence. Specifically, as one's peers behave in more (less) normatively compliant ways, one's own perceptual accuracy increases (decreases) and one's behavioral fit concomitantly increases (decreases).*

METHOD

Testing these hypotheses requires access to longitudinal data on cognitive and behavioral cultural fit, as well as exogenous variation in the set of peers to which a focal actor is exposed. To meet these criteria, 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 and cognitive cultural fit. 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 on two dimensions: behavioral and cognitive. We therefore need to locate individuals in two cultural spaces—one behavioral and the other cognitive—and measure their distances from the centroids of their respective groups. 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.[3]

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 (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), we use 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.

[3]There are various ways of defining this reference group. Work in organizational culture has traditionally either defined this reference group as the organization as a whole or as the individuals' organizational department. Drawing on Srivastava et al. (2018) we argue that one's group of immediate peers is most consequential for cultural fit. In the robustness tests we conduct below we also use the organization as a whole as the cultural reference group, demonstrating that our findings are not sensitive to this assumption.

Measuring Cognition—To assess 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. 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.[4] 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)

[4]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.

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.[5] 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 the Interactional Language Use Model, 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

[5]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.

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 lower 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 the interactional language use model has been used in previous work to measure cultural fit, it is still a fairly new methodology. To further validate that 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 by translating each value statement into its corresponding category number. For example, if value statement 1 were put in category 7 in one profile and category 2 in another profile, that statement would represent the point (7,2). We similarly computed points for all 54 value statements and calculated the correlation among those points.

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. 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.

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 2 provides a conceptual overview of this procedure. Further procedural details are provided in the Appendix B; evaluative analyses regarding model fit are provided in Appendix C.

[FIGURE 2 ABOUT HERE]

Peer 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. 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 supplemental analyses reported below. These were not direct tests of our hypotheses but designed to assess whether the imputed measures related to career outcomes as would be expected based on theory and prior research. In particular, we derived from the personnel records two individual outcome measures. The first was monthly bonus. Only those in job roles such as sales or operations, for which productivity could be objectively assessed, were bonus eligible. 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. The second outcome was exit, based on 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 (Thomp-

son, 1967; Salancik and Pfeffer, 1974).[6] For our between-person models, we included additional control variables for age and gender.

Analytical Approach

We tested Hypothesis 1 using OLS regressions based on cross-sectional data, as well as fixed effect regressions based on longitudinal data, including the 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 2, we identified the effect of changes in peer composition on the focal individual's cultural fit measures—behavioral fit, perceptual accuracy, and value congruence. 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

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

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:

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

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

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

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

We used these instruments—network change, induced change in 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:

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

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

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:

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

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 three 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 3. 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 3 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.

Finally, in Table 2, we modeled voluntary exit from the organization as a function of value congruence and perceptual accuracy. Although people leave organizations for a variety of reasons, voluntary exit is most likely to be associated with declining attachment. The competing risks model reported in Table 2 is a survival model that extends the Cox Proportional Hazards model to the case of multiple failures. In our case, involuntary exit is the competing risk.[7]

[7]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.

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).[8] Overall, these supplemental analyses help to validate the longitudinal fit measures derived from our imputation methodology.

Main Results

Table 3 provides a test of our first hypothesis: that perceptual accuracy predicts changes in behavioral fit. The dependent variable in all models is behavioral fit. The first three models report results from cross-sectional data where the cognitive fit measures—perceptual accuracy and value congruence (which we analyze because we suggested that it would be less related to behavioral fit than would perceptual accuracy)—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 Hypothesis 1, 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 Hypothesis 1 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.

[8] 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.

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.[9] 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.[10]

Table 4 reports the analyses we used to test Hypothesis 2—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

[9] Although the role of status, whether in the form of managerial status or diffuse status characteristics such as gender, in linguistic conformity is outside the scope of this paper, we see great potential in future research—including both field and experimental studies—that unpacks that mechanisms by which status affects behavioral conformity.

[10] 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.

their own behavioral fit in the following month. The opposite is also true: individuals who, through the reorganization, transitioned into a network of peers with lower behavioral fit experienced a corresponding decline in their own behavioral fit. Interestingly, and likely because reorganizations are disruptive to cultural integration, the majority of employees experienced a decline in peer behavioral fit, and correspondingly, their own behavioral fit during this period.

We illustrate the implications of induced change in peer behavioral fit in Figure 4. The diagram plots the effects of the reorganization on individuals' behavioral fit over time, as estimated by the instrumental variable model. The upper line corresponds to individuals who experienced a half standard deviation positive increase in their peers' behavioral fit, and the lower line corresponds to individuals who experienced a decline of the same magnitude in their peers' behavioral fit. These are substantial changes in peer behavioral fit but not implausible during a period of reorganization. A little over 1% experienced a positive shock at or greater than half a standard deviation, but roughly 35% experienced a decline of that magnitude. Both translate to similarly sized adjustments in the focal individuals' behavioral fit, but in opposite directions. Moreover, both adjustments persisted for roughly two months, after which the effects of the reorganization were no longer apparent and individuals converged toward mean behavioral fit. Because the reorganization was not a true natural experiment, it is worth noting that changes that occurred after its effects were initially felt could have arisen for a variety of reasons that we do not observe in our data. For example, individuals presumably regained more command over whom they interacted with after the reorganization, which would also reintroduce potentially confounding homophily effects. Hence, the period immediately following the reorganization is the appropriate one to consider for this analysis.

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

Changes in the number of peers had a more modest impact: those who experienced an increase in the size of their network due to the reorganization experienced declines in behavioral fit. Forced network growth, in other words, 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).

Organizational research has tended to approach cultural assimilation through the lens of socialization (e.g., Van Maanen and Schein, 1979; Alba and Nee, 2009; Ashforth et al., 2007). Such an approach assumes that cultural adaptation entails a gradual internalization of the group's norms and underlying value system. Prior studies have therefore almost exclusively focused on value congruence as the primary dimension of cultural fit, implicitly equating enculturation with value alignment. We offer a more comprehensive model of fit and enculturation which distinguishes deciphering the cultural code—what we term "perceptual accuracy"—from its internalization, and we demonstrate how these two mechanisms derive from different sources and relate to different aspects of individual attainment.

Our theoretical framework and concomitant findings make two broad contributions to advancing a theory of person-culture fit. First, consistent with previous work, we show that those who learn to fit-in culturally reap positive career rewards (Chatman and O'Reilly, 2016). Indeed our results reinforce the importance of both cognitive and behavioral fit for individual attainment: all three of our fit measures were positively linked to individual productivity, as indicated by bonus

payments. But, in contrast to prior theoretical formulations of fit, we also demonstrate that different rewards accrue to different forms of cultural alignment: whereas perceptual accuracy is related to individuals' capacity to behave in a normatively compliant manner, value congruence is more related to a person's voluntary decision to stay or leave the organization.[11] Those who read the code correctly and behave accordingly benefit from being perceived as true and committed group members, while those who identify with the code enjoy the psychological wellbeing that comes with a positive self-concept. These results offer a theoretical explanation to integrate previously disparate findings that pertain to the cognitive antecedents and subtle behavioral manifestations of cultural fit—such as conformity with norms of how to dress (Rafaeli et al., 1997), the ability to engage in banter about sports at work (Turco, 2010), the enactment of presentational rituals that signal ideological alignment with management (Kunda, 1991), and the use of communication that matches the linguistic style of colleagues (Srivastava et al., 2018).

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

[11] We acknowledge that linguistic fit is not the only way for those with high levels of value congruence to display normative compliance. For example, given the robust link between value congruence and longevity found in previous research, it seems likely that if a member is not involuntarily separated from the organization, she is likely engaging in certain other behaviors that are normatively compliant.

accuracy over time is a fruitful avenue for further developing theories of person-culture fit.

Our second contribution relates to the factors that cause some people to enculturate more successfully than others. 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.

This link we establish between peers' behaviors and those of the focal actor also contributes to a growing body of research on the interrelationships between structure and culture (McLean, 1998; Lizardo, 2006; Goldberg et al., 2016; Askin and Mauskapf, 2017). 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.

At a more general level, our work contributes to organizational research by demonstrating

how cognition and behavior are intertwined in producing and sustaining cultural order. Recent work in cultural sociology has distinguished shared preferences from shared meanings and construals (Goldberg, 2011; DiMaggio and Goldberg, 2018). Our findings show that the latter—that is, agreement in how a situation is interpreted, not necessarily in what is desirable or worthy—is sufficient for an identifiable culture to emerge. Our distinction between perceptual accuracy and value congruence provides an analytical framework for understanding how cognition and behavior can converge or diverge and the consequences for individuals and groups of various combinations of cognitive and behavioral fit. Future work might draw on these foundations to further our understanding of how, despite cognitive fragmentation at the individual level, culture can nevertheless appear to be coherent to the organization as a whole.

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 cognitive and behavioral arenas of social life relate to one another and jointly shape the life course and the cultures in which it unfolds.

References

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

Ashforth, B. E., D. M. Sluss, and A. M. Saks
2007. Socialization tactics, proactive behavior, and newcomer learning: Integrating socialization models. *Journal of Vocational Behavior*, 70(3):447–462.

Ashforth, B. K. and A. M. Saks
1996. Socialization tactics: Longitudinal effects on newcomer adjustment. *Academy of Management Journal*, 39(1):149–178.

Askin, N. and M. Mauskapf
2017. What makes popular culture popular? product features and optimal differentiation in music. *American Sociological Review*, 82(5):910–944.

Bail, C. A.
2017. Taming big data: Using app technology to study organizational behavior on social media. *Sociological Methods & Research*, 46(2):189–217.

Bail, C. A., T. W. Brown, and M. Mann
2017. Channeling hearts and minds: Advocacy organizations, cognitive-emotional currents, and public conversation. *American Sociological Review*, P. 0003122417733673.

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

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

Baron, J. N., M. T. Hannan, and M. D. Burton
2001. Labor pains: Change in organizational models and employee turnover in young, high-tech firms. *American journal of sociology*, 106(4):960–1012.

Baum, C. F., M. E. Schaffer, S. Stillman, 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.

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

Centola, D., R. Willer, and M. Macy
2005. The emperor's dilemma: A computational model of self-enforcing norms. *American Journal of Sociology*, 110(4):1009–1040.

Chan, T. Y., J. Li, and L. Pierce
2014. Compensation and peer effects in competing sales teams. *Management Science*, 60(8):1965–1984.

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

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

Chatman, J. A. and S. G. Barsade
1995. Personality, organizational culture, and cooperation: Evidence from a business simulation. *Administrative Science Quarterly*, 40:423–443.

Chatman, J. A., D. F. Caldwell, C. A. O'Reilly, and B. Doerr
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, J. A. and C. A. O'Reilly
2016. Paradigm lost: Reinvigorating the study of organizational culture. *Research in Organizational Behavior*, 36:199–224.

Dandekar, P., A. Goel, and D. T. Lee
2013. Biased assimilation, homophily, and the dynamics of polarization. *Proceedings of the National Academy of Sciences*, 110(15):5791–5796.

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

DiMaggio, P. and A. Goldberg
2018. Searching for homo economicus: Variation in americans construals of and attitudes toward markets. *European Journal of Sociology*, P. 139.

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

Edwards, J. R.
2008. Person-environment fit in organizations. *The Academy of Management Annals*, 2(1):167–230.

Edwards, J. R. and D. M. Cable
2009. The value of value congruence. *Journal of Applied Psychology*, 94(3):654.

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

Evans, J. A. and P. Aceves
2016. Machine translation: Mining text for social theory. *Annual Review of Sociology*, 42:21–50.

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

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

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

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

Garfinkel, H.
1967. Studies in ethnomethodology.

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

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

Goldstein, N. J., R. B. Cialdini, and V. Griskevicius
2008. A room with a viewpoint: Using social norms to motivate environmental conservation in hotels. *Journal of Consumer Research*, 35(3):472–482.

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

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

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

Harrison, J. R. and G. Carroll
2006. *Culture and demography in organizations*. Princeton, NJ: Princeton University Press.

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

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

Hewlin, P. F., T. L. Dumas, and M. F. Burnett
2017. To thine own self be true? facades of conformity, values incongruence, and the moderating impact of leader integrity. *Academy of Management Journal*, 60(1):178–199.

Ho, T. K.
1995. Random decision forests. In *Document Analysis and Recognition, 1995., Proceedings of the Third International Conference on*, volume 1, Pp. 278–282. IEEE.

Hochschild, A. R.
2012. *The Managed Heart: Commercialization of Human Feeling*. Univ of California Press.

Judge, T. A. and D. M. Cable
1997. Applicant personality, organizational culture, and organization attraction. *Personnel Psychology*, 50(2):359–394.

Kelley, H. H. and J. L. Michela
1980. Attribution theory and research. *Annual Review of Psychology*, 31(1):457–501.

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

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

Kristof-Brown, A. L., R. D. Zimmerman, and E. C. Johnson
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. and R. A. Leibler
1951. On information and sufficiency. *The Annals of Mathematical Statistics*, 22(1):79–86.

Kunda, G.
1991. *Engineering Culture: Control and Commitment in a High-Tech Corporation*. Philadelphia: Temple University Press.

Lazer, D. and J. Radford
2017. Data ex machina: Introduction to big data. *Annual Review of Sociology*, (0):19—39.

Leonardelli, G. J., C. L. Pickett, and M. B. Brewer
2010. Optimal distinctiveness theory: A framework for social identity, social cognition, and intergroup relations. *Advances in Experimental Social Psychology*, 43:63–113.

Liebal, K., M. Carpenter, and M. Tomasello
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, C. C. and S. B. Srivastava
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.
2006. How cultural tastes shape personal networks. *American Sociological Review*, 71(5):778–807.

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

Lord, C. G., L. Ross, and M. R. Lepper
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, W. W., E. Mullen, and A. D. Galinsky
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, P. V. and N. E. Friedkin
1993. Network studies of social influence. *Sociological Methods & Research*, 22(1):127–151.

Mayer, D. M., M. Kuenzi, R. Greenbaum, M. Bardes, and R. B. Salvador
2009. How low does ethical leadership flow? test of a trickle-down model. *Organizational behavior and human decision processes*, 108(1):1–13.

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

McLean, P. D.
1998. A frame analysis of favor seeking in the renaissance: Agency, networks, and political culture. *American Journal of Sociology*, 104(1):AJSv104p51–91.

Meglino, B. M. and E. C. Ravlin
1998. Individual values in organizations: Concepts, controversies, and research. *Journal of Management*, 24(3):351 – 389.

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

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

Miller, D. T. and D. A. Prentice
2016. Changing norms to change behavior. *Annual Review of Psychology*, 67:339–361.

Mobasseri, S., A. Goldberg, and S. B. Srivastava
2018. What is cultural fit? from cognition to behavior (and back). Oxford University Press (forthcoming).

Morris, M. W., C.-y. Chiu, and Z. Liu
2015. Polycultural psychology. *Annual Review of Psychology*, 66:631–659.

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

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

O'Reilly, C. A. and J. A. Chatman
1996. Culture as social control: Corporations, cults, and commitment.

Ostroff, C. L. and T. Judge
2007. *Perspectives on Organizational Fit*. Psychology Press.

Pennebaker, J. W.
2013. *The Secret Life of Pronouns: What Our Words Say About Us*. Bloomsbury USA.

Pennebaker, J. W., C. K. Chung, M. Ireland, A. Gonzales, and R. J. Booth
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/.*

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

Rivera, L. A.
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, G. R. and J. Pfeffer
1974. The bases and use of power in organizational decision making: The case of a university. *Administrative Science Quarterly*, Pp. 453–473.

Salganik, M. J.
2017. *Bit by Bit: Social Research in the Digital Age*. Princeton University Press.

Sarros, J. C., J. Gray, I. L. Densten, and B. Cooper
2005. The organizational culture profile revisited and revised: An australian perspective. *Australian Journal of Management*, 30(1):159–182.

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

Schein, E. H.
2010. *Organizational Culture and Leadership*, volume 2. John Wiley & Sons.

Schneider, B.
1987. The people make the place. *Personnel Psychology*, 40(3):437–453.

Scott, B. A., C. M. Barnes, and D. T. Wagner
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.

Sperber, D.
1996. *Explaining culture: A naturalistic approach*. Oxford, UK: Blackwell Publishers.

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

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

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

Thompson, J. D.
1967. *Organizations in action: Social science bases of administrative theory*. Transaction publishers.

Trope, Y. and N. Liberman
2010. Construal-level theory of psychological distance. *Psychological Review*, 117(2):440.

Turco, C. J.
2010. Cultural foundations of tokenism: Evidence from the leveraged buyout industry. *American sociological review*, 75(6):894–913.

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

Vaisey, S. and O. Lizardo
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. and E. Schein
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.

Willer, R., K. Kuwabara, and M. W. Macy
2009. The false enforcement of unpopular norms. *American Journal of Sociology*, 115(2):451–490.

FIGURES

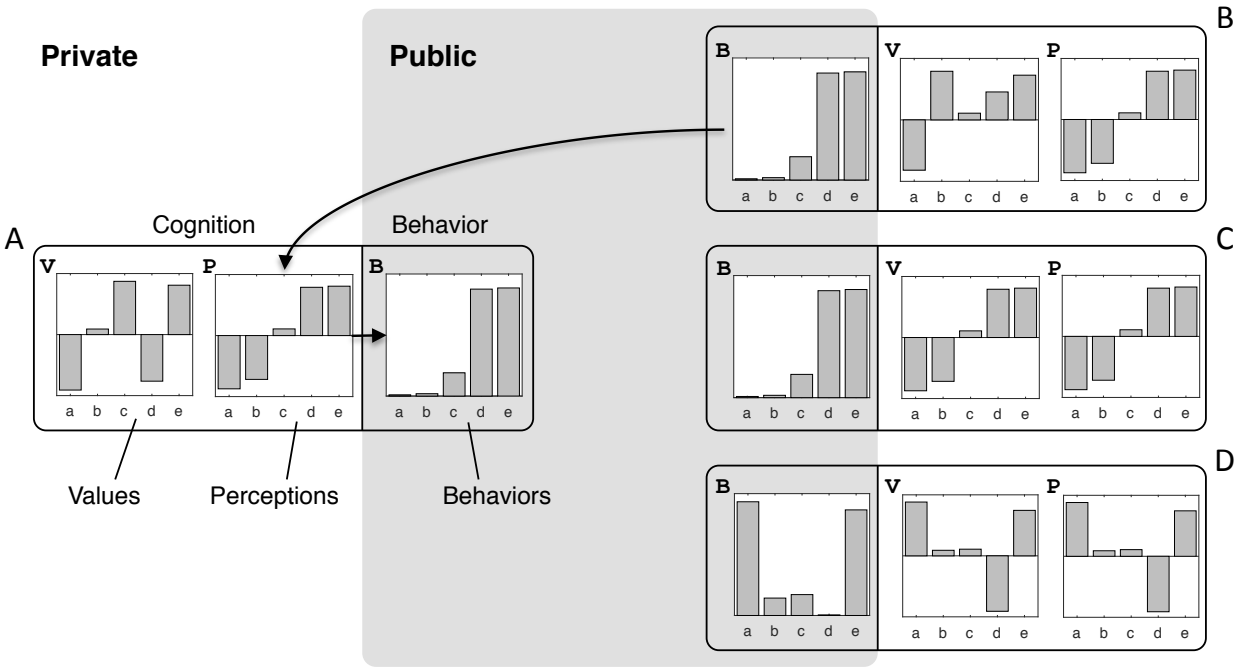

FIGURE 1: A schematic illustration of our theory. Four individuals (A-D) are each characterized by their values (V), perceptions (P) and behavioral probabilities (B). Arrows correspond to causal relationships.

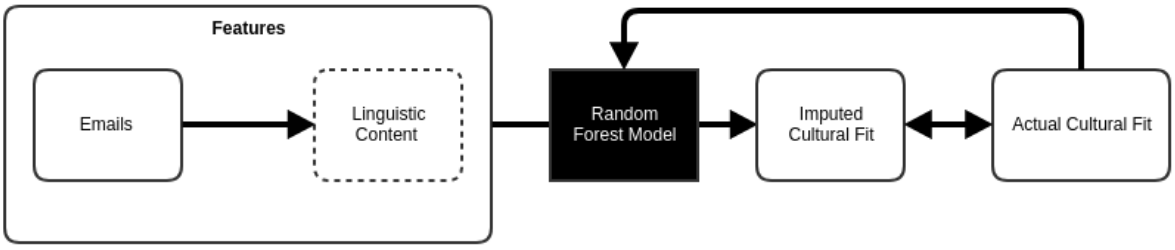

FIGURE 2: Conceptual Overview of the Machine Learning Process

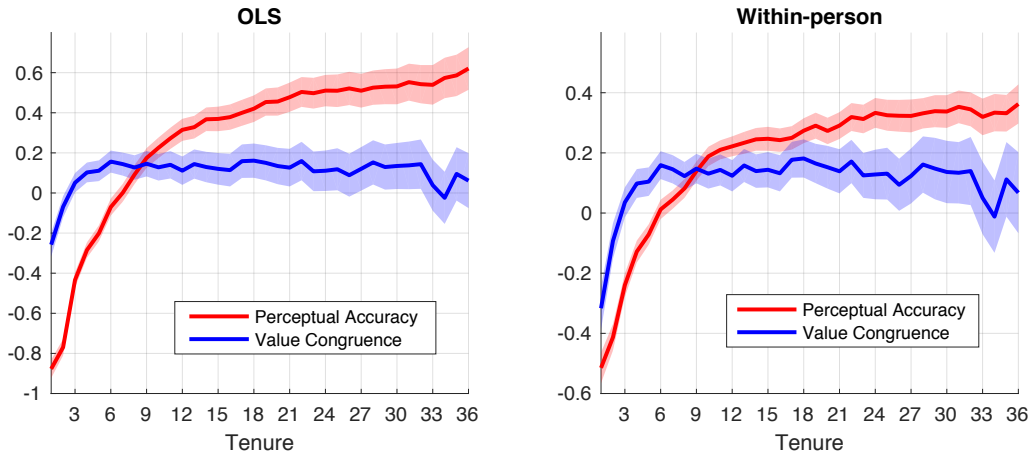

FIGURE 3: OLS and fixed effect regressions of perceptual accuracy and value congruence, with indicators for each tenure month up to 36 months in the company.

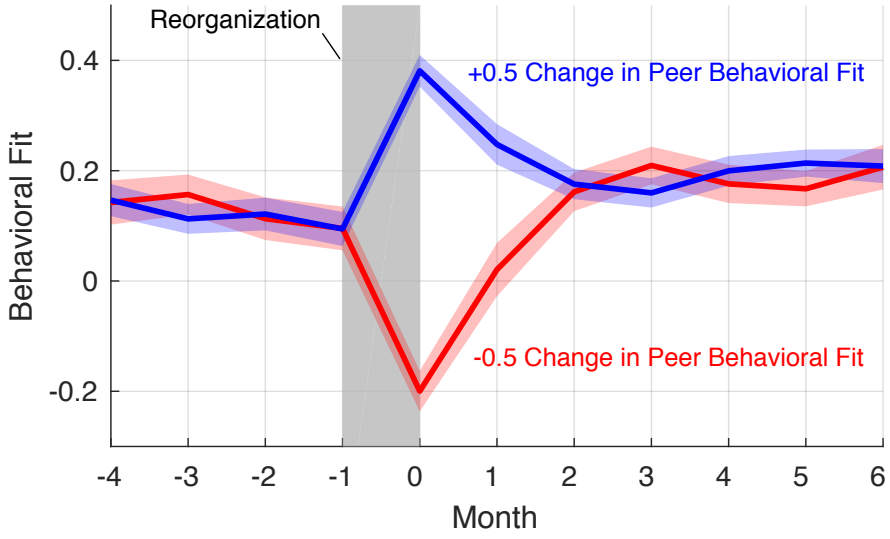

FIGURE 4: Marginal effects, estimated by monthly consecutive instrumental variable models, of change in peer behavioral fit on individual behavioral fit. The two lines correspond to individuals who experienced a 0.5 increase (blue) or decrease (red) in peer behavioral fit. Shaded areas correspond to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

TABLES

TABLE 1

Fixed Effect Regressions of Bonus (logged)

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

t statistics in parentheses; standard errors clustered by individual

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

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	0.864
	(-0.77)	(-0.62)
Female	1.386*	1.392*
	(2.53)	(2.56)
Age	0.901**	0.902**
	(-3.23)	(-3.23)
Age^2	1.001**	1.001**
	(3.20)	(3.22)
Num. Employees	1.002***	1.002***
	(9.46)	(9.96)
Department Dummies	Yes	Yes
Observations	27467	27467
χ^2	172.161	177.689
Log-Likelihood	-1320.27	-1318.36

Exponentiated coefficients; t statistics in parentheses

Standard errors clustered by individual; Sample weights by tenure

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

TABLE 3

Cross-Sectional and Longitudinal Fixed Effects Regressions of Behavioral Fit

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

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

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

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

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

t statistics in parentheses; standard errors clustered by individual

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

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

TABLE 5: Robustness Checks—Instrumental Variables Fixed Effect Regressions of Behavioral Fit

	Model 1 Organization	Model 2 Non-Managers
Peer Behavioral Fit†		0.235***
		(5.78)
Peer Behavioral Fit (Organization)†	0.158***	
	(5.40)	
Num. Peers †	-0.003	-0.013*
	(-1.85)	(-2.10)
Manager	0.133***	
	(3.57)	
First Year	-0.034*	-0.150**
	(-2.27)	(-3.25)
Constant	2.154***	-0.560
	(26.90)	(-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

† 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 $\overrightarrow{m}_{it}$ as each email individual i sends at time t and $\overleftarrow{m}_{it}$ as each email individual i *receives* at time t. We then define the set of LIWC categories as L and the set of all times in any given month as T. Our procedure iterates over all emails sent and received and produces $\overrightarrow{m}^l_{it}$ and $\overleftarrow{m}^l_{it}$ for the count of terms in email $\overrightarrow{m}_{it}$ and $\overleftarrow{m}_{it}$ in LIWC category $l \in L$, respectively. Then, by aggregating all individual email counts $\overrightarrow{m}^l_{it}$ and $\overleftarrow{m}^l_{it}$ for $t \in T$, it produces sent and received LIWC counts in month T, $\overrightarrow{m}^l_{iT}$ and $\overleftarrow{m}^l_{iT}$. We normalize each LIWC count in each month by the total of all LIWC counts in that month to transform the LIWC probability distribution to a standard probability distribution. We use the notation, O^l_{iT} to denote the outgoing normalized probability and I^l_{iT} to denote the incoming normalized probability.

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

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

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

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

where the JS-divergence between two probability distributions is defined as a symmetric measure

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

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

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

Validation of Behavioral Cultural Fit

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

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

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

As illustration, we examine the link between language use and a preference for a people orientated culture. We find that respondents who value people orientation tend to include more affect words (e.g., happy, cry, abandon), perceptual process words (e.g., observe, hear, feel), positive emotion words (e.g., love, nice, sweet), and second-person words (e.g., you, your) in their outgoing communication. 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) [12]. 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

[12]These are the top two correlations among fifteen significant correlations total in the outgoing LIWC categories and the incoming LIWC categories, respectively.

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 $\boldsymbol{BF}_k$. 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 $\boldsymbol{BF}_k$ 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 2 in the main manuscript and described in greater detail below.

Our first step was to translate the raw email data into a format that is usable by the random forest model. We tokenized and stemmed all words in the body of email messages. Tokenization involves separating the text into distinct terms, for which we used the TwitterTokenizer designed for linguistic analysis Potts (2011). Stemming involves reducing each term to a root form, for which we used the Porter Stemmer from the python nltk package. We removed all characters that could not be encoded into unicode, such as "\x00," and split the text into n-stems, where n is in the set [1,2,3]. Given that language use tends to follow the power law, in which few terms are used frequently and many terms are used infrequently, we then undertook steps to reduce the dimensionality of the data to make it computationally tractable. We retained all n-stems in emails

sent from individuals, but only uni-stems in emails sent to individuals. Additionally, we retained only those n-stems that were used by 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

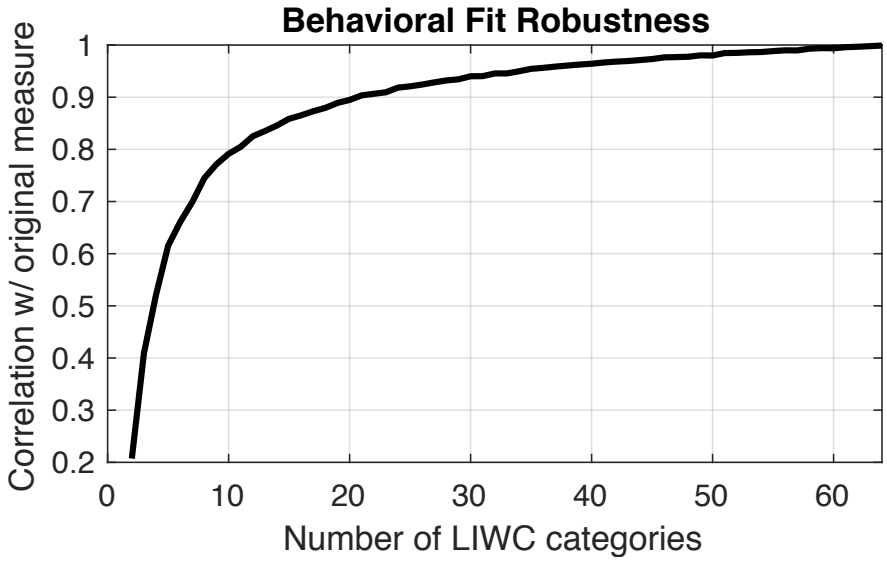

FIGURE A1: Robustness of the behavioral fit measure to LIWC category composition

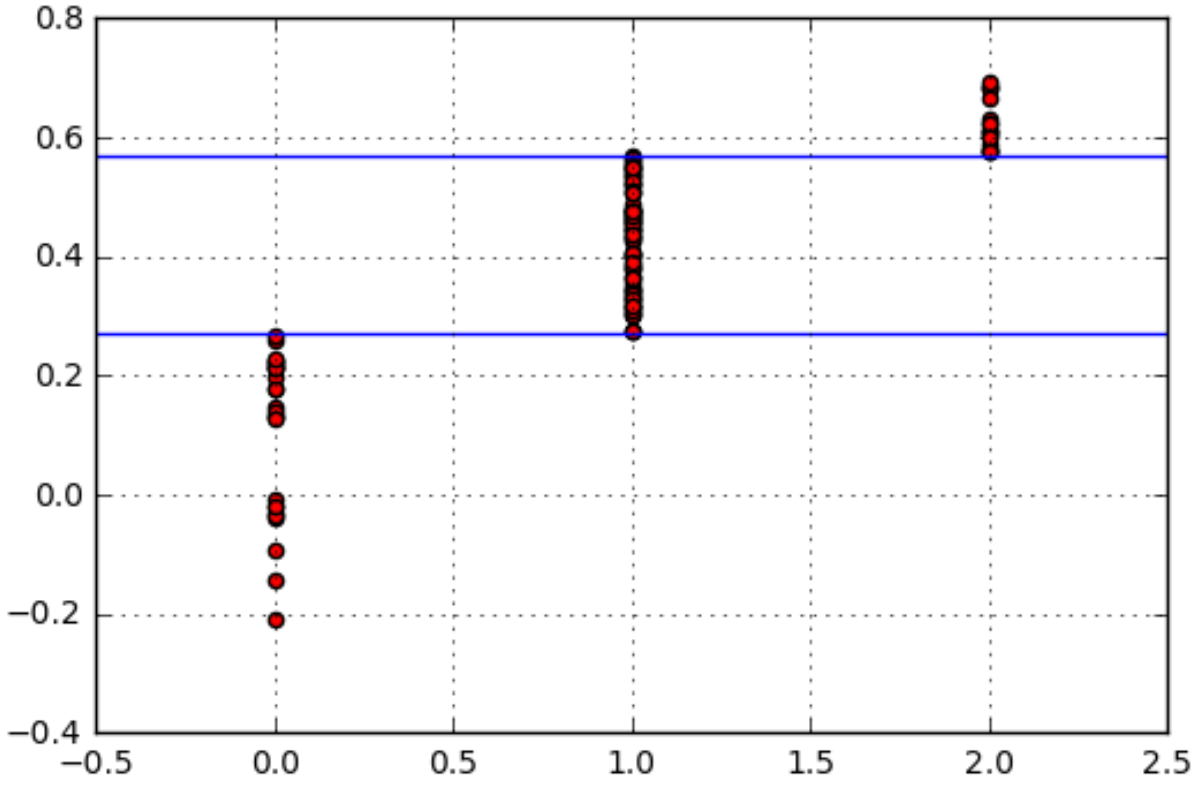

FIGURE C1: Division of Continuous Cultural Fit into Classes

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