UC Berkeley
UC Berkeley Electronic Theses and Dissertations

Title
Surveying Personality with Behavior: The Situational Influences and Individual Outcomes of Self-Monitoring Behavior

Permalink
https://escholarship.org/uc/item/76x2p2ww

Author
Lu, Richard

Publication Date
2019

Peer reviewed|Thesis/dissertation

eScholarship.org Powered by the California Digital Library
University of California

Surveying Personality with Behavior:
The Situational Influences and Individual Outcomes of Self-Monitoring Behavior

by

Richard W. Lu

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Business Administration

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Sameer B. Srivastava, UC Berkeley, Haas School of Business, Chair
Cameron Anderson, UC Berkeley, Haas School of Business
David Bamman, UC Berkeley, School of Information

Spring 2019

Surveying Personality with Behavior:
The Situational Influences and Individual Outcomes of Self-Monitoring Behavior

Copyright 2019
by
Richard W. Lu

Abstract

Surveying Personality with Behavior:
The Situational Influences and Individual Outcomes of Self-Monitoring Behavior

by

Richard W. Lu

Doctor of Philosophy in Business Administration

University of California, Berkeley

Sameer B. Srivastava, UC Berkeley, Haas School of Business, Chair

Through decades of research, scholars have illustrated the complexity of the relationship between unobservable personalities and observable behaviors. Yet, despite this complexity, the predominant empirical method of assessing personality remains the psychometric scale. In this dissertation, I extend past research on behavioral measures of personality by proposing a novel approach using natural language, and highlight at least three key ways it helps advance our understanding of personality and behavior. First, assessing personality from behavior directly bypasses potential cognitive and perceptual biases involved in the self- and peer-report of scales, allowing for a more objective account. Second, it uncovers temporal and situational variance largely obscured by scales, which capture mean tendencies and more stable, individual differences. Third, it has greater potential to shed light on the processes underlying documented correlations between personality and distal outcomes.

In Chapter 1, I review the personality and behavior literature and draw out two major conclusions, describing how they have led to three different methods of personality assessment. I discuss the limitations of these methods, explain how my approach overcomes these limitations, and introduce the personality construct of self-monitoring through which I illustrate my approach. In Chapter 2, I detail the construction of a behavioral measure of self-monitoring. I explain how characteristics of the technical apparatus map on to the general theory of self-monitoring, outline the design decisions I made in constructing the measure, and provide comparisons with the survey measure as well as robustness checks. In Chapter 3, I conclude with an empirical analysis of the measure, demonstrating that it is responsive to situations as expected. I additionally demonstrate its utility in its association with consequential individual outcomes such as salary bonus and network constraint.

To the Lord my God, the Father of mercies and God of all comfort, who has comforted me in all my affliction and who will perfect the good work He has begun in me until the day of Christ Jesus,

and

To my excellent wife Olivia, whose faithful love to this undeserving sinner has enabled and encouraged me to finish my degree

Contents

Contents **ii**

List of Figures **iii**

List of Tables **iv**

1 The Complex Relationship Between Personality and Behavior **1**
1.1 Personality and Behavior 3
1.2 Assessing Personality 6
1.3 Self-Monitoring as an Illustrative Example 11

2 Measuring Self-Monitoring Behavior **15**
2.1 Traditional Survey-Based Measure of Self-Monitoring 15
2.2 Operationalizing Self-Monitoring Behavior 16
2.3 Research Setting and Collection of Self-Monitoring Data 22
2.4 Comparison Between the Self-Monitoring Survey and Self-Monitoring Behavior 22
2.5 Robustness of LIWC Category Inclusion/Exclusion 25
2.6 Limitations and Future Extensions 33

3 The Role of Context for Self-Monitoring Behavior **35**
3.1 Theory 36
3.2 Methods 41
3.3 Main Results—How Context Affects Self-Monitoring Behavior 43
3.4 Supplementary Results—The Outcomes of Self-Monitoring Behavior 45
3.5 Discussion 47

References **51**

A 85- and 64-LIWC Self-Monitoring Behavior **60**

List of Figures

2.1 Conceptual Diagram of LIWC and WHAM 18
2.2 Correlation of Self-Monitoring Behavior and Cross-Sectional Self-Monitoring by Month 23
2.3 Correlation Between Full and N-Category-Excluded LIWC Self-Monitoring Behavior 26
2.4 Correlation Between Full and One-Category-Excluded LIWC Self-Monitoring Behavior 30
2.5 Correlation Between Single-Category LIWC Measures 31
2.6 Correlation Between Single-Category LIWC Measures (Sub-sample) 32

List of Tables

2.1 Eighteen Item Measure of Self-Monitoring (Snyder and Gangestad, 1986) 16
2.2 Design Decisions with Short Rationales 21
2.3 Correlation of Aggregated Behavioral Self-Monitoring with Cross-Sectional Self-Monitoring (Month 114) 24
2.4 Excluded LIWC Categories by Number of Categories Excluded (Lowest Correlations) 28
2.5 Excluded LIWC Categories by Number of Categories Excluded (Highest Correlations) 29
2.6 Example Words in LIWC Categories 33

3.1 Correlation Table 43
3.2 Individual Fixed-Effects OLS Regression of Contexts on Self-Monitoring Behavior 44
3.3 Individual Fixed-Effects OLS Regression of Self-Monitoring Behavior on Outcomes 46

A1 21 Additional Categories in 85-LIWC 60
A2 Correlation Table (85-LIWC) 61
A3 Individual Fixed-Effects OLS Regression of Contexts on Self-Monitoring Behavior (85-LIWC) 62
A4 Individual Fixed-Effects OLS Regression of Self-Monitoring Behavior on Outcomes (85-LIWC) 63

Acknowledgments

I am grateful to a great number of scholars who have been instrumental in not only the development of my dissertation, but also my own development as a scholar. Though it would be a task in itself to name them all, I wanted to especially acknowledge some.

First, my dissertation chair and advisor, Sameer B. Srivastava, without whom I certainly would not have finished. Your encouragement and support have made the difference, and your guidance and mentorship have made me a better scholar. Thank you for caring enough to tell me and teach me how to improve, and perhaps more importantly, for listening.

Second, my collaborators who have taken the time to teach me as well as work with me, Jennifer A. Chatman, Ming D. Leung, and Amir Goldberg. Thank you for walking me through each step of the research process (or more realistically, perhaps, at times pulling me through it), and thank you for pushing me to think more critically, to write more eloquently, to speak more confidently, and to communicate more gracefully.

Third, my committee members who have provided invaluable input from my oral defense to my finished dissertation, Cameron Anderson, David Bamman, and Toby Stuart. Thank you for your helpful and timely feedback on all stages of the process and your support towards my graduation. It wasn't easy, but you certainly made it easier.

Fourth, my fellow PhD students, without whom it most certainly would have been much harder. Thank you Bernadette Doerr, Brian Reschke, Eliot Sherman, Alex Van Zant, Angus Hildreth, Weiyi Ng, Jennifer Logg, Sanaz Mobasseri, Laura Howland, Nicole Zeng, Daron Sharps, Katherine Ullman, Haochi Zhang, Michael Matheiu, Michael Rosenblum, Derek Schatz, Rachel Habbert, Gauri Subramani, Dan Stein, Derek Brown, Sonya Mishra, Andrew Choi, and Matteo Tranchero.

CHAPTER 1: The Complex Relationship Between Personality and Behavior

It has been said that even mafia bosses are kind to their families. This apparent contradiction stems from the understanding that mafia bosses can not be at once kind and criminal. Kind people are not generally thought of as performing criminal acts; nor are criminals generally thought of as showing kindness. And yet it is evidently possible. However, the more difficult question is to distinguish between the two. Is the mafia boss a generally kind personality who sometimes behaves illicitly? Or is the mafia boss a criminal personality who sometimes behaves kindly? Even more, does his personality matter if he is perceived on the basis of his behavior, which may itself be interpreted in vastly different ways (Padgett and Ansell, 1993)?

These questions regard the definition, assessment, and impact of personality and its relationship with behavior, a relationship which has historically been rather complex. Individuals are assumed to have stable, individual differences which coalesce into personalities (Allport, 1937; Digman, 1990; Goldberg, 1990). These personalities are associated with certain consistent behaviors, although behavior itself is dependent on a multitude of potential influences in addition to one's personality such as the immediate context or certain motivations (Mischel, 1969; Mischel and Shoda, 1995; Fuglestad and Snyder, 2010). Therefore, the abstract notion of personality is ascertained through observable behavior, but many notions and methods exist as to how this should be done (Funder, 1991; Gosling et al., 2002). Finally, the impact of personality is often assessed by its correlation with a variety of outcomes (Roberts et al., 2007), although the mechanistic pathways through which such impact happens is not always as clear.

Despite this complexity in the theory that has developed around personality and behavior, largely where we have ended up empirically is the use of psychometric scales. The international personality item pool (IPIP), a public domain effort to make personality assessment more widely accessible contains over 3,000 items and over 250 scales constructed from those items (Goldberg et al., 2006). Those scales, capturing psychological constructs of individual differences such as the locus of control (Levenson, 1981), the need for cognition (Cacioppo and Petty, 1982), and self-monitoring (Snyder, 1974), among others, have been used in at least 748 studies to date (*Publications that Employ the IPIP*). As a more specific example, there are a number of scales available for measuring the Five Factor Model, or the Big Five, which is perhaps the most widely used taxonomy of personality traits. In addition to scales constructed from the IPIP, there are the Big Five Inventory (John, Donahue, and Kentle, 1991; Benet-Martinez and John, 1998; John, Naumann, and Soto, 2008), the Big Five Aspect Scales (DeYoung, Quilty, and Peterson, 2007), Big Five markers (Goldberg, 1992; Saucier, 2002), the NEO-PI-R and the NEO-FFI (Costa Jr and McCrae, 1992), and the ten item personality index (Gosling, Rentfrow, and Swann Jr, 2003), among others.

My argument is not that psychometric scales are inherently less useful than other po-

tential measures, nor that they should not be used for the purpose of assessing personality. A number of studies have repeatedly demonstrated that personality measured in this way not only significantly correlates with individual outcomes as varied as subjective well being (Emmons and Diener, 1985) to music preferences (Rentfrow and Gosling, 2003), but also social outcomes such as social network position (Fang et al., 2015). Moreover, in a series of longitudinal studies, it has been shown to have effect sizes on important life outcomes as mortality, divorce, and occupational attainment comparable to socioeconomic status and cognitive ability (Roberts et al., 2007).

Rather, my argument is simply that personality as measured by scales only represents a portion of the complexity of personality-behavior theory and as a result is limited in significant ways. To overcome these limitations, I extend past research on behavioral measures of personality (Buss and Craik, 1983; Fleeson, 2001; Furr, 2009) by proposing a novel approach to assessing personality using natural language, and highlight at least three key ways it helps advance our understanding of personality and behavior. First, assessing behavior directly bypasses potential cognitive and perceptual biases involved in the self- and peer-report of scales, allowing for a more objective account. Second, while scales capture mean tendencies and more stable, individual differences, they obscure temporal and situational variance. Assessing behavior unlocks the potential to examine and leverage these sources of variance. Third, assessing behavior allows us to uncover the processes underlying the correlations between personality and distal outcomes (Roberts et al., 2007).

The remainder of this chapter is thus organized accordingly. First, I review the personality and behavior literature and draw out two major conclusions, that personality and behavior are distinct, and that the behavioral consistency assumed of personality is a complex issue. Then, I illustrate how these two conclusions have led to three different methods of assessing personality[1]. The first two, personality *assumed*, where judgments from recollections of aggregate behavior are assumed to equate to one's personality, and personality *predicted*, where behaviors predictive of certain personalities are used to impute one's personality, predominately focus on scales. The third, personality *derived*, where individual behaviors are aggregated in some way to derive one's personality, bypasses scales, yet comes with its own set of limitations. I explain how my method overcomes these limitations, and illustrate it through the personality construct of self-monitoring, or an individual difference in the propensity to monitor self-presentation, expressive behavior, and non-verbal affective displays (Snyder, 1974). Chapter 2 details the construction of a behavioral measure of self-monitoring, and Chapter 3 illustrates the utility of such a measure by considering the effect of context on self-monitoring behavior.

[1]I understand that there are alternate conceptions of personality that incorporate more dimensions of the individual, such as the cognitive affective processing system (Mischel, 2009), than the historical trait perspective of personality would consider. It is beyond the scope of my current work to engage the deeper discussion of what should actually be considered the personality of an individual, and I focus on the traits perspective of personality throughout my dissertation.

1.1 Personality and Behavior

The topic of personality and behavior has been extensively studied in the literature through decades of research. Although there still remain many open questions, two major conclusions seemed to have been reached, whose threads run throughout the literature. The first conclusion is that personality and behavior are distinct. This reality has been posited even from the beginning of personality theory, although the story over time has certainly become more nuanced. The second conclusion is that there is a complexity behind the behavioral consistency that results from personality. While early scholars also theorized that there would never be a perfect behavioral consistency from personality, more recent work has developed this complexity and more sharply characterized how behavioral consistency stems from personality.

The Distinction Between Personality and Behavior

Starting with perhaps the founding fathers of the study of personality, we see that personality and behavior were thought of as distinct. Stern (1921), in his seminal work on personality, described an individual's behavior as comprising a variable portion and a constant purposive portion. This latter portion had been isolated by Stern and developed by Allport (1937) as the trait, and traits, in turn, were thought of as the "fundamental dispositions of personality" (Allport, 1937, pg. 286). These traits, which are characteristic of individuals, then contribute to the constant portion of their behaviors; they are not the behaviors themselves, nor are behaviors directly components of individuals' personalities.

Allport (1937) went on further to theorize that traits are formed from habits. The integration of numerous specific habits, which themselves are more or less routines of behavior, results in the formation of a trait. A young child who learns to brush his or her teeth over times develops brushing into a habit, which, in conjunction with other such habits, such as bathing or washing one's face, develop into a trait of personal cleanliness. In this view we again see the distinction between personality and behavior. While the repeated behaviors in the form of habits may contribute to the formation of certain traits, the traits themselves are distinct from the original behaviors. The young child, once fully grown, may continue to maintain the trait of personal cleanliness without necessarily performing the same exact cleaning behaviors.

Moreover, in his original formulation of traits, Allport (1937) very clearly described his supposed relationship between traits and behavior. "No single trait—nor all traits together—determine behavior all by themselves. The conditions of the moment are also decisive; the special character of the stimulus, the temporary distribution of stresses and tensions within the neuropsychic system, all demand a special form of adaptive response, perhaps never again required in precisely the same way" (Allport, 1937, pg. 313). He supposed that traits surely had an effect on behavior, a complicated one, yet there was more than just the personality so defined that resulted in behavior. So it is clear that during the onset of personality theory,

while there was a close connection between personality and behavior, they were nonetheless thought of as distinct.

Later scholars built upon this foundation of distinctiveness by trying to condense the realm of traits down into a manageable set with which to compare and predict behavior. While Allport himself attempted to condense 17,953 words (primarily adjectives) of the English language describing personal forms of behavior into four columns of trait-names, he did so by rather arbitrary means. Raymond B. Cattell after him brought methodological rigor to the exercise, eventually condensing the realm of personality traits down to 16 by means of factor analysis using self-report questionnaires, observer ratings, and objective behavioral tests (Cattell, 1943; Cattell, 1946; Cattell, 1947; Cattell, 1948). The resulting Sixteen Personality Factor Questionnaire (16PF) has been extensively studied and reanalyzed over decades of research (Cattell and P. Cattell, 1995).

In a similar vein, other scholars have further reduced the number of traits down to five factors. Norman (1963), utilizing the same trait names of Cattell (1947) and scales derived by Tupes and Christal (1961) reduced the personality sphere down to five general factors, extroversion, agreeableness, conscientiousness, emotional stability, and culture (which would later be termed openness to experience). The primary difference between Cattell's sixteen factor structure and the five factor structure was the rotation of the axes; Cattell used an oblique rotation, which allowed for correlation between the individual factors, while the five factor model used orthogonal rotation, causing the factors to be independent. These five factors would become the precursors of the Five Factor Model, or the Big Five (Digman, 1990). While the Five Factor Model has its criticisms (Pervin, 1994) and alternatives (Eysenck, 1991), it is widely used today as a primary framework of personality.

This distillation of traits into five general factors subtly moved the distinction of personality and behavior from an explicit statement to an implicit understanding. As consensus started to build around a framework of personality, traits were solidified as latent variables in their own right, and scholars started moving more towards predicting and correlating behavior to the five general factors. A cursory review of the vast literature reveals that the Big Five have been used to predict behaviors as varied as social status and leadership positions, interpersonal relationships, performance in work groups, health outcomes, etc. (John, Naumann, and Soto, 2008). Even sociologists have used the Big Five in service of predicting social network behavior (Fang et al., 2015). By focusing on how certain behaviors might be related to certain traits, scholars were essentially taking for granted the distinction between personality and behavior; the question then became *what* kinds of behaviors are associated with *which* personality traits.

The Complexity of Behavioral Consistency from Personality

A second major conclusion from personality-behavior theory, that there is a complexity behind the behavioral consistency assumed to arise from personality, has a similarly lengthy history. On the one hand, we have our nagging common intuition that personality must lead to some stable, cross-situational consistency in behavior (Bem and Allen, 1974), but on the

other hand, even early on, there was empirical evidence, namely cross-situational behavioral variance, that this might not be the case (Newcomb, 1929; Dudycha, 1936). This paradox is the starting point for much of the discussion in the field, and after decades of empirical research and theoretical clarification, it still remains as a rather unsolved complexity.

Returning back to when the field of personality psychology was still at its nascent stages, we again find that Allport (1937) argued that the very existence of a trait hinged on "some acceptable method of consistency in behavior" (Allport, 1937, pg. 330). His reasoning for such a proposition was essentially epistemological; he simply claimed that it should be so and provided the thought that some individuals are self-consistent in their decisiveness or fastidiousness. We might consider that the decisive manager would be equally decisive in determining where to dine for dinner, which gift to procure for a wedding anniversary, as well as how to lead his team. Allport, it seemed, like others before him and certainly others after him, simply believed that a singular trait would be highly associated with specific behaviors.

However, this claim would soon find the need to be substantiated. The most vocal critic was perhaps Walter Mischel, who in a series of writings culminating in his 1968 *Personality and assessment*, argued that the empirical evidence suggested otherwise. He demonstrated that, in many of the studies completed at the time of his writing, there was in fact great behavioral variance across situations in areas as varied as attitudes towards authority and peers, moral behavior, rigidity and tolerance for ambiguity, and cognitive avoidance. His findings and arguments led to something of a crisis in the field of personality psychology, as the consistency necessary for traits to "exist" seemed absent.

Note, however, that it was not Mischel's prerogative to argue that personality did not exist, nor was it exactly Allport's arguments that there would be a certain high level of consistency in behavior. Mischel himself stated that it would be "an unfortunate misconstruction" to interpret "a search for a personality psychology that has conceptual room for major variability...as undermining the concept of personality itself" (Mischel, 1969, pg. 1017). And Allport (1937), acknowledging the degree of consistency expected as a somewhat subjective matter, provided at least one major hindrance—albeit not the one picked up by later critics—to observing behavioral consistency, that multiple traits exist in a unitary personality and at times may either dominate one another or act in conjunction to result in seemingly inconsistent behavior. The discussion was not around whether personality existed, as some may have it, but rather around what kind and degree of behavioral consistency could be expected.

From this point forward, many such clarifications have been made. First, scholars have documented that while traits may be less predictive of individual behaviors, they are rather predictive of consistency in the aggregate (Fishbein and Ajzen, 1974). In his seminal article, Epstein (1979) considered this from several points of view, finding across several studies that self-reports of emotions, self-reports of externally observable behavior, and examiner reports of observable behavior were all highly correlated when considering several days' worth of observations versus marginally correlated when considering a single day's worth of observations. Fleeson (2001) similarly finds that when comparing self-reports of behavior by using adjective-based Big Five scales across two randomly selected 3 hour time periods, the

correlations are relatively close to the .30 threshold (Mischel, 1968) and behavior is perhaps even more variable within-person than between-person. However, when comparing the same measures across two different weeks, he finds the correlations near the .90 level.

Second, scholars have considered how different traits may lead to different behavioral consistencies. Moskowitz (1982) considered how two different traits, dominance and dependency, manifested in behaviors across situations. He observed children in preschool for approximately 4 hours each and constructed global measures of the traits based on coding of the behavior observed. The global measure of dominance was relatively generalizable across situations, while the global measure of dependency was not, suggesting that behaviors stemming from some traits may be more cross-situationally consistent than others. Even within the Big Five, Fleeson (2001) demonstrated that there was relative disparity between the behavioral consistency of each component, with within-individual correlations of self-reported behaviors ranging from .28 for conscientiousness to .54 for intellect or openness to experience.

Third, scholars have examined the role of the situation in affecting behavioral consistency. While earlier research did show cases where there was cross-situational consistency (Epstein, 1979), opposing evidence pushed theory towards considering the effect of "strong" situations (Mischel, 1977). The general thrust of the argument is that there are cases where person variables are more influential and conversely cases where situation variables are more influential; the latter is more likely to occur in so-called strong situations where "everyone knows what to do and why and how to do it" (Cooper and Withey, 2009, pg. 63). Although more direct tests (Beaty Jr, Cleveland, and Murphy, 2001; Withey, Gellatly, and Annett, 2005) of how situations moderate the effect of personality on behavior have failed to satisfy the most stringent critics (Cooper and Withey, 2009), they and other studies have nevertheless demonstrated that behavior correlates with characteristics of the situation (Fleeson, 2001).

In short, the discussion surrounding behavioral consistency still remains a complex one, despite decades more of empirical data. Not only do different measures of behavior and different measures of traits affect the consistency uncovered, but also characteristics of the situation constrain and influence behavior. Moreover, there may be additional forces unrelated to the personality-behavior link or the situation that would affect behavioral consistency, such as meaningful changes in one's personality over time (Mund and Neyer, 2014; Roberts, Caspi, and Moffitt, 2001; Lavner et al., 2018). Nevertheless, the echoes of the earliest scholars on the topic remain—there is a certain consistency of behavior due to personality, but the degree of such consistency is highly variable.

1.2 Assessing Personality

With these foundations in place, we can begin to examine how scholars have used the relationship between personality and behavior to assess personality. Returning to Allport, it was explicitly understood that traits were theoretical constructs, that "traits...are themselves never directly observed" (Allport, 1937, pg. 314); the only way personality could be

ascertained was by indirect means through some observable behavior. However, there have been a number of perspectives as to how this was to be done, which have integrated the two prior conclusions of the personality and behavior literature to varying degrees. Each, in turn, has its limitations, which is why I have proposed a fourth method that is both theoretically informed by the literature and methodologically effective in advancing personality research.

Survey-Based Methods

The first perspective by which personality is assessed is that of personality *assumed*. In this vein of literature, there are two primary sources of personality assessment, either the self or an observer. Regardless of the source, the underlying principle is that some aggregate of behavior is assumed to correspond to the levels of certain personality traits (Block, 1989). For example, for self-reports, an individual could assume himself or herself as high in conscientiousness after remembering several instances of planning or following rules. For peer-reports, the observers rate individuals on levels of personality traits after observing behavior ranging from group interaction to the arrangement of office space or bedrooms (Kenny et al., 1994; Gosling et al., 2002). While the general consensus is that both forms of assessment are comparable (Funder, 1991), and for peer-reports particularly, even in "zero-acquaintance" situations and with harder to evaluate traits, there are still moderate levels of correlation between observers (Funder and Colvin, 1988; Ready et al., 2000; Kenny et al., 1994), in both cases the levels of personality traits are still assumed based on interpretations of aggregate behavior.

The second perspective, which in some ways can be seen as a modern methodological extension to the first perspective, is that of personality *predicted*. This approach uses machine learning methods to derive personality, often through complicated relationships, from computationally tractable behavior. The caveat, however, is that the gold standard in many of these studies has been self-reported Big Five questionnaires, and the general approach is to derive features from social media platforms (e.g. Facebook profiles and messages, Twitter profiles and tweets) and train a model that predicts the Big Five levels from the questionnaires (for a review, see Bleidorn and Hopwood, 2018). Although these models have generally demonstrated high accuracy, even outperforming humans in self-other personality agreement and inter-judge agreement (Youyou, Kosinski, and Stillwell, 2015), they are still using digital behavior only to predict traditional self-reported levels of personality traits.

These perspectives focus on the use of surveys or psychometric scales. Clearly, decades of research from psychometricians have produced mature methods and tools that have usefully advanced personality research (Block, 1989). However, there remain at least three major critiques of surveys that limit their potential to, in and of themselves, adequately capture personality and its link with associated outcomes. As a result, any approach that seeks to assess personality solely on the basis of a survey would, for the majority of purposes, remain incomplete.

First, as illustrated by the first perspective, scales are often self- or peer-report, which independently involve a set of assumptions and threats to validity. For self-reports, while

individuals may be able to accurately report on attitudes and beliefs held in short-term memory, they may be less reliable when trying to remember or infer other attributes (Haeffel and Howard, 2010). This is especially pertinent to personality surveys, as they often require individuals to report on general patterns of behavior, of which perhaps only the most recent acts are reliably accessible. Additionally, self-reports are susceptible to cognitive biases, most notably social desirability bias (Donaldson and Grant-Vallone, 2002). Individuals asked to report on a topic as personal as personality may be tempted to respond in self-aggrandizing ways (Bäckström and Björklund, 2013). For peer-report, the basis of personality assessment may be a limited set of observable behaviors which may themselves be interpreted with cognitive biases from the peers.

Second, scales capture individual differences and mean tendencies but obscure situational variance (Mischel and Shoda, 1995; Fleeson, 2004; Fournier, Moskowitz, and Zuroff, 2008). In other words, an individual who ranks high in agreeableness may nevertheless at times be contentious. This is important to consider because this kind of within-person variance has been demonstrated to eclipse between-person variance (Fleeson, 2001; Moskowitz and Zuroff, 2004), highlighting the incomplete picture of the mean tendencies captured by scales, especially when examining how personality relates to consequential outcomes. If differentially agreeable individuals may nevertheless act in similar ways, empirical results may find mixed support.

Third, scales measure more stable characteristics of individuals but not underlying processes or mechanisms. Scales are by nature cross-sectional, and even in longitudinal designs in which individuals fill out the same survey after a period or periods of time, they are still removed from the proximal causal force. For example, one of the more robust findings in the personality and social networks literature has been that high self-monitors occupy brokerage positions (Kilduff and Day, 1994; Mehra, Kilduff, and Brass, 2001; Sasovova et al., 2010; Landis, 2016). However, being a high self-monitor is conflated with acting like one, and the momentary behaviors that high self-monitors engage in to enable their brokerage positions (e.g. how they create and maintain ties) are not uncovered.

In general, these survey-based methods attend to the conclusion that personality and behavior are separate but only one part of the discussion that the behavioral consistency of personality is complex. The consistent portion of behavior picked up by psychometric scales is assumed to be an individual's personality. And even if personality assumed in such a way is predicted by observable behavior, these perspectives still lack a true attention to the more complete picture behavior may reveal about personality.

A Behavioral Method

The third perspective, personality *derived*, has attempted to more directly incorporate behavior into personality assessment (Furr, 2009). In this perspective, personality is understood as some aggregate measure of behavior itself, rather than as a latent variable behavior elucidates. The foundations of this perspective are found in the act frequency approach to personality, which sees dispositional assertions merely as summary statements of individual

behavior. Certain prototypical acts, for example, doing the dishes after dinner or picking up the tab for lunch, are seen as indicative of the dispositional construct of agreeableness. The degree of the personality trait is then derived from the total tally of these acts as well as the variety of acts performed within a certain dispositional category (Buss and Craik, 1983).

Ignoring the more involved limitations of the act frequency approach for a moment (those interested may refer to Block (1989), pages 237-244), there are still three major issues with this perspective. First, it goes against the first major conclusion of the personality and behavior literature that personality and behavior are distinct, instead proposing a view of "dispositions as summary statements," (Buss and Craik, 1983, pg. 108) essentially equating personality constructs with summaries of behavior. Second, the act statements themselves, which are used to construct the act-trend summaries, are largely non-contextual, going against the second conclusion of the complexity of behavioral consistency. For example, it is not very clear if the act statement, "I demanded a backrub," is even related to the supposed personality construct of dominance, especially if the (albeit contrived) context was that someone had demanded the individual to demand a backrub (Block, 1989). Third, and perhaps the most ironic provided that the act frequency approach intended to integrate behavior into the personality discussion, the acts that are studied are not actual behaviors but are self-reported agreements with act statements, meaning that the basis for measurement is still the memory of individuals committing those behaviors.

Despite its limitations, the act frequency approach emphasized the importance of behavior in the study of personality, suggesting that the actual tracking of behavior could help reconcile the disparate findings of personality research. Building on this idea, Fleeson (2001) characterized the entire distribution of trait-relevant behaviors as a density distribution, and demarcated its relevance to the study of personality. Contrary to a tally of certain dispositionally meaningful acts, each behavior of an individual was to be rated on the degree to which it expresses a certain trait, and those ratings could be aggregated to a density distribution. Then, the mean, range, and other summary variables of that distribution become the important units for describing the behavioral patterns of the individual[2].

I deviate from Fleeson's density distribution work in one major way with several implications. Rather than focusing on the state level of traits expressed in a wide range of everyday behaviors, I first describe how a certain trait is manifested in natural language, then measure the degree to which those behavioral expressions take place. For example, for the trait of self-monitoring, I first detail a linguistic measure of self-monitoring, then use this measure to assess the degree to which individuals self-monitor over time. This approach is obviously not without its own limitations, but I believe it does address many of the outstanding limitations of the density distribution approach (Fleeson and Gallagher, 2009) while presenting many

[2]Notably, in keeping with the conclusion that personality and behavior are distinct, Fleeson argued that personality scholars needed "to embrace a new, advanced understanding of traits, realizing that people tend to demonstrate significant flexibility in their behavior and that traits are best used for predicting trends" (Fleeson, 2004, pg. 86). His density distribution approach to personality was an acknowledgement of that significant flexibility in behavior uncovered by decades of empirical work, not a repudiation of relatively stable personality traits.

of its own unique benefits.

First, this approach affords a direct, more objective observation of behavior. In addition to the self-report biases covered in the survey-based methods section, the density distribution work suffers from shared method variance, since both the traits and the state levels are reported by the same individual. Individual differences in reference standards, for example, would affect comparison across individuals. Granted, in later work, the experience sampling approach is used by observers in a lab setting to create density distributions of behavior for focal individuals (Fleeson and Law, 2015), but behavior is still documented and variously interpreted by individuals. Natural language, on the other hand, and in particular, written communication, presents objective, archival documents from which computational methods can derive a consistent measure of behavior.

Second, this approach affords a more precise view of temporal and situational variance. While the density distribution work generally asks respondents to rate their behaviors over a defined period of time (e.g. during the previous hour), when those behaviors occurred and the context surrounding them are lost in consolidation to the singular rating. In contrast, observational records of natural language are often precisely time stamped (e.g. emails), and can be aggregated over multiple periods of time, such as days, weeks, months, or years. Moreover, such a source of behavior can be paired with other similarly precise data, providing a clearer view of the situations surrounding behavior. In fact, I demonstrate this in Chapter 3.

Third, this approach affords more analytical leverage for identifying causal processes and downstream consequences by focusing on a specific manifestation of a trait rather than a rating of multiple behaviors. Using natural language, researchers can identify specific patterns as manifestations of a trait. For example, an agreeable linguistic pattern may be the use of hedging words such as, "I'm not an expert, but...," or, "...isn't it?" These patterns, because of the largely consensual rules of language in a closed cultural system, can be quantified, aggregated, and compared in a way that the behavioral input of the density distribution work can not.[3] Personality traits can be linked to quantifiable personality expressions in behavior which themselves can be linked to downstream outcomes.

In sum, my approach takes inspiration from existing behavioral methods of assessing personality but does not suffer from the same limitations. It represents a more objective and tractable way of identifying how personality impacts meaningful outcomes, especially in ways difficult to examine through traditional survey-based methods. Moreover, it should become decreasingly costly to apply as relevant data becomes more ubiquitous and computational techniques become more powerful.

[3]It bears mentioning that although these patterns can themselves be seen as distinct behaviors which are then "rated" by some aggregation method, their relation relies on the socio-cognitive basis of linguistics rather than the idiosyncratic minds of individuals.

1.3 Self-Monitoring as an Illustrative Example

To illustrate the usefulness of this approach, I focus on the personality construct of self-monitoring. The construct of self-monitoring was introduced by Mark Snyder in his seminal paper (Snyder, 1974) on the self-monitoring of expressive behavior. Building on the perspective of social interaction as a theatrical performance (Goffman, 1959), and the concept of nonverbal expressive behavior (Ekman and Friesen, 1971), Snyder introduced the construct with the idea that there are "striking and important individual differences in the extent to which individuals can and do monitor their self-presentation, expressive behavior, and non-verbal affective display" (Snyder, 1974, pg. 526). He theorized that the goals of self-monitoring have to do with displaying or concealing emotional states and that self-monitoring itself required a certain sensitivity to situational cues. Based on these ideas, he developed a set of self-descriptive statements and validated the resulting self-monitoring scale through a series of four studies.

There are two major reasons why self-monitoring is a good example to illustrate my approach. First, self-monitoring has had a remarkable influence among researchers across disciplines, not only those studying personality, but also those in social psychology, organizational behavior, sociology, and networks. This is likely a result of the consequential outcomes that have been found to be related to self-monitoring, namely that high self-monitors are more likely to have individuals approach them both for expressive resources like friendship and instrumental resources like information and are more likely to broker disconnected friends, leading to further downstream benefits (Fang et al., 2015). By creating a behavioral measure of self-monitoring in the process of illustrating my approach, I hope to provide a tool to help invigorate the well-established study of self-monitoring.

Second, self-monitoring has had a tumultuous history making it especially amenable among the personality constructs to a behavioral approach. In its early genesis as a theoretical construct, it focused on "self-presentation, expressive behavior, and non-verbal affective display," or in other words, the *behavior* of individuals. It seems only fitting that a personality construct designed around individual differences in behavior would be best studied by examining actual instances of behavior. Moreover, the self-monitoring scale and even the concept of self-monitoring itself has changed from its first iteration, largely because of behavioral evidence. If the scale has had such difficulty capturing the concept, perhaps it would be better served by a behavioral approach.

Therefore, I conclude this chapter with the background and history of self-monitoring, consider in the next how we can construct a linguistic measure of self-monitoring behavior, and demonstrate in the third its usefulness in advancing theory, all the while hoping to demonstrate how this general approach can be broadened to consider other personality constructs as well.

A Brief History of the Personality Construct Self-Monitoring

Soon after self-monitoring was theorized as a construct, critiques of the measure came. Perhaps the most well-known of the critiques were that there was a strong disconnect between the theoretical construct and the measurement of it by the scale. Factor analysis of the scale revealed three underlying factors, named "Acting," "Extraversion," and "Other-Directedness," (Briggs, Cheek, and Buss, 1980), which has been replicated even in Snyder's own work (Gangestad and Snyder, 1985). The question was phrased that if the scale tapped three different psychological constructs, how could it capture the unitary idea of self-monitoring? Perhaps self-monitoring as it was currently measured was not actually self-monitoring at all? Adding to the critique, Lennox and Wolfe (1984) devised an alternate self-monitoring scale meant to pick up on a more parsimonious definition of self-monitoring more in line with Snyder's original theorizing. These critiques opened up the possibility that self-monitoring was not as clear a construct as it was intended to be.

Snyder and Gangestad (1986) made two major responses to the building critique of the original self-monitoring measure. The first was that, while it was rather indisputable that there was a non-unitary underlying factor structure of the self-monitoring scale, this did not not mean that the scale did not tap a meaningful general factor. They argued that the first unrotated factor was a factor on which most of the items positively loaded, and therefore it was a general factor that underlied a substantial amount of the variance in the measure. The second was to reformulate the self-monitoring scale to more strongly tap that general factor. Removing items that did not correlate at least +.15 with the single unrotated factor, they ended up with a new 18-item self-monitoring scale. While responding to the critiques, this article marked a transition in the understanding of self-monitoring coming from Snyder himself: "it hardly seems right that the construct should be identified solely with the initial statements, given that the theory concerning the individual differences tapped by the Self-Monitoring Scale has evolved through extension, elaboration, and reemphasis...thus, we would prefer to not directly address whether the self-monitoring construct is 'right,' but instead to ask, 'What *is* true about the individual differences underlying the Self-Monitoring Scale?"' Rather than keep to the original theoretical formulation of self-monitoring, the battle over its definition would now revolve around the empirical data.

This did not end the debate as to how self-monitoring should be measured, and consequently, what self-monitoring actually was. By analyzing six samples versus one sample, Briggs and Cheek (1988) identified that the general factor was not always consistent, and actually there were two general factors that showed up across the samples. In analyzing the new 18-item scale, they again found two general factors, now called public performing and other-directedness, and they argued that in a sense removing seven items from the scale simply shifted the emphasis within these factors but did not resolve the fact that there was no single general factor. John, Cheek, and Klohnen (1996), using an alternate definition of self-monitoring, a Q-sort made by Snyder himself, did not find that either the 18-item scale or the public performing subscale correlated more highly with the alternate definition than additional measures of social potency or extraversion. These findings led them to similarly

suggest that the reformulation of the scale shifted the concept of self-monitoring more towards extraversion and away from other-directedness. It would seem that by readjusting the scale to load more heavily on the first unrotated factor (determined as public performing), Snyder and Gangestad inadvertently shifted the concept of self-monitoring itself.

To this second round of critiques, Gangestad and Snyder (2000) provided a second round of rebuttals. This time around, they relied heavily on the decades of research on self-monitoring that have taken place since the first round of critiques, arguing that the correlates found in various research streams have clarified the identity of the high and low self-monitor. Taking this argument further, they defined a procedure that allowed them to mathematically compare the degree to which external criterion in roughly nine different categories (e.g. expressive control; nonverbal decoding skills; attitude-behavior relations, attitude accessibility, and persuasion; behavioral sensitivity to others' expectations or external cues; behavioral variability; interpersonal orientations; being impressed with physical attractiveness; attention and responsiveness to others; and peer-self trait rating discrepancy) were related to self-monitoring, its two factors, or its three subscales. By transforming factor loadings to correlations (Cronbach, 1951; Gangestad and Snyder, 1991), they were able to determine which external criterion related more to the general self-monitoring axis than any of its factors or subscales. These findings then led them to exclude or include certain aspects of the self-monitoring formulation. Importantly, it was found that the category of attention and responsivity to others was more closely related to the other-directedness axis than the self-monitoring axis, leading them to conclude

> Closely tracking and responding to others' behavior was a core feature within the original self-monitoring formulation. Indeed, the very term self-monitoring was coined to refer to regulating one's presentations of self on the basis of social and interpersonal cues to situational appropriateness, including the behavioral guidelines provided by other people (Snyder, 1974; Snyder, 1979). Ironically, then, the Self-Monitoring axis simply may not measure propensities to "self-monitor" in the precise sense in which that term was defined in the original statements of self-monitoring theory (Gangestad and Snyder, 2000, pg. 545).

It would appear, then, that at this point in the discussion, the construct of self-monitoring as it is used has shifted from the original theoretical impetus for the concept in the first place.

More modern examples of self-monitoring research reflect this ambiguity in the construct of self-monitoring. Not only are clarifications still being made on the basis of empirical data and improved techniques, such as the reanalysis of the scale as better reflecting a dimensional measure versus a categorical measure (Wilmot, 2015), but also when empirical papers do use the self-monitoring concept and accompanying scale, they often cite some of the earliest definitions of the construct (see examples in Day et al. (2002, pg. 390); Wilmot et al. (2016, pg. 335); Tasselli and Kilduff (2018, pg. 806); Fang et al. (2015, pg. 1244). Without discounting the significant amount of scholarly interest in the concept of self-monitoring that continues today (Wilmot et al., 2016), the conclusion from the literature from those

closest to the construct of self-monitoring seems to be that self-monitoring as captured by existing scales remains rather ambiguous, leaving room for my behavioral approach.

CHAPTER 2: Measuring Self-Monitoring Behavior

With the potential benefits of behavioral measures of personality based on natural language outlined, I now turn to considering how one might be constructed. As mentioned in Chapter 1, I focus on constructing a linguistic measure of self-monitoring behavior by using natural language in written communication. I do not attempt to construct a measure encompassing all that self-monitoring may entail—it is impossible, for example, to measure self-monitored physical expressions by examining natural language in written communication alone—but rather seek to construct a linguistic measure corresponding to the personality construct of self-monitoring. In short, I try to construct a measure that captures instances of self-monitoring and allows comparative assessment of the degree of self-monitoring behavior.

In line with these efforts, the rest of this chapter is outlined as follows. First, I review the current survey-based measure of self-monitoring. Second, I detail my operationalization of self-monitoring behavior, including rationale on a unique benefit of natural language, brief descriptions of the technical apparatus, and design decisions in the construction process integrated with formal descriptions of my measures. Third, I demonstrate how the survey-based measure and my behavioral measure compare with one another using data from an organizational setting. Fourth, I include a robustness check for my behavioral measure. Fifth, I describe some limitations and future extensions of the approach.

2.1 Traditional Survey-Based Measure of Self-Monitoring

Self-monitoring has traditionally been measured by a psychometric survey known as the Self-Monitoring Scale (Snyder, 1974). In its original conception, the scale consisted of twenty-five items determined to be most discriminative from forty-one true-false self-descriptive statements. Since then, as a result of critical theoretical and empirical responses to the twenty-five item scale (covered in Chapter 1), Snyder and Gangestad (1986) have reformulated the scale to be more internally consistent, reducing the number of items down to eighteen. This reformed scale is still a widely used measure of self-monitoring for empirical research today (for examples, see Fang et al., 2015, Tasselli and Kilduff, 2018), and is reproduced in Table 2.1 below.

Table 2.1: Eighteen Item Measure of Self-Monitoring (Snyder and Gangestad, 1986)

Item	Key
1. I find it hard to imitate the behavior of other people.	F
2. At parties and social gatherings, I do not attempt to do or say things that others will like.	F
3. I can only argue for ideas which I already believe.	F
4. I can make impromptu speeches even on topics about which I have almost no information.	T
5. I guess I put on a show to impress or entertain others.	T
6. I would probably make a good actor.	T
7. In a group of people I am rarely the center of attention.	F
8. In different situations and with different people, I often act like very different persons.	T
9. I am not particularly good at making other people like me.	F
10. I'm not always the person I appear to be.	T
11. I would not change my opinions (or the way I do things) in order to please someone or win their favor.	F
12. I have considered being an entertainer.	T
13. I have never been good at games like charades or improvisational acting.	F
14. I have trouble changing my behavior to suit different people and different situations.	F
15. At a party I let others keep the jokes and stories going.	F
16. I feel a bit awkward in public and do not show up quite as well as I should.	F
17. I can look anyone in the eye and tell a lie with a straight face (if for a right end).	T
18. I may deceive people by being friendly when I really dislike them.	T

Note: High self-monitoring individuals tend to answer in the keyed direction.

2.2 Operationalizing Self-Monitoring Behavior

A Unique Benefit of Natural Language

In addition to the more general advantages of using natural language to observe personality expression mentioned in Chapter 1, I believe there is one unique advantage considering the nature of the construct of self-monitoring. One primary reason why self-monitoring behavior has not been adequately captured may be the inherent difficulty of observing such behavior. In order to directly observe it, one has to essentially observe two sets of behavior, the non-self-monitored behavior of the individual and the self-monitored behavior. However, in any given instance, such a task is impossible. Natural language, however, in particular, written communication, affords a flexibility in expression that allows pseudo-observation of these two sets. Studies have found that linguistic patterns in written communication are both discriminative enough between individuals to identify authorship (Zheng et al., 2006) and stable enough within individuals to be reliable across time and situations (Pennebaker and King, 1999), suggesting that language can be used to identify the normal, non-self-monitored behavior for individuals. Other research has shown that written communication is sensitive to situational contexts, that individuals will accommodate their language to communication partners (Guo et al., 2015), suggesting that language can also be used to identify self-monitored behavior for individuals.

Technical Apparatus

I leverage two linguistic tools to estimate the components of self-monitoring behavior, the Linguistic Inquiry and Word Count (LIWC; Pennebaker et al., 2015 , and the Word-based Hierarchical Alignment Model (WHAM; Doyle and Frank, 2016). LIWC is a dictionary-based method for categorizing words into various conceptual categories (e.g. positive emotion, certainty, health, present focus, etc.) and was originally developed in 1993 to help identify and group words that tapped basic emotional and cognitive dimensions. Since its original conception, a number of categories have been added and some removed to improve the internal consistency of the dictionary, and the dictionary has been extensively used and validated for research purposes, representing a gold standard of psychometrically validated conceptual categories (Fast, Chen, and Bernstein, 2016).

The general process of creating the 2015 version of LIWC, the version I employ in my measure, consists of seven major steps. First, words for each category were collected from several sources. Second, each of the generated words was rated by a group of judges for their fit. Third, the words were analyzed for occurrence across various sources such as blog posts, spoken language studies, Facebook, novels, etc., and words that did not occur in multiple corpora were discarded. Fourth, high-frequency words that were not previously considered were identified for inclusion and passed through the same judge rating process as before. Fifth, the words were subjected to psychometric evaluation, testing for internal reliability and external validity. Sixth, the dictionary was refined by repeating the first five steps. Seventh, summary variables based on previously published findings were created. The result is a dictionary that translates words to categories representing various linguistic and psychological processes.

WHAM is a generalizable bayesian estimation model that assumes that word use in responses is shaped by the presence or absence of certain markers in the original message. For example, one might assume that the use of formal language in a message elicits formal language in a response which, under normal circumstances, may otherwise not be used. It does this by differentiating between base responses, to original messages that do not contain the marker of interest, and alignment responses, to original messages that do contain the marker of interest. Then, the difference in language use between the two sets of responses represents the level of alignment. The model can then be extended to incorporate the general effects of certain parameters. For example, different usage patterns may exist among different categories of language—positive emotion may generally be used more than negative emotion—or different groups of individuals—managers may use language differently than non-managers.

To be more technically precise, WHAM, being a bayesian probabilistic model, specifies a formal model for the response patterns found in message-reply data. Each response, with corresponding counts for in-category tokens and not-in-category tokens, is modeled as a binomial distribution with probability μ, which depends on the observation's specific parameters—for example, if it is sent by a manager or non-manager. Those parameters influ-

encing μ are modeled as a hierarchical chain of normal distributions[1], which generate linear predictors η, which are inverse logit transformed to the probabilities μ. In this formulation, then, the observed counts of in-category tokens and not-in-category tokens in responses to messages containing and not containing markers in the category of interest are the data informing the estimation (for more technical information, please refer to the original papers Doyle and Frank, 2016, Doyle et al., 2017).

LIWC and WHAM are the linguistic workhorses for developing my measure. Their part in transforming raw message content to a measure of self-monitoring behavior is presented in Figure 2.1 below. I first use LIWC to translate raw messages (in my case, emails) to distributions of LIWC category counts, which are my markers of interest. Then, I use WHAM to estimate base and alignment patterns for individuals in a given month. Exactly how I translate the WHAM output to measures of self-monitoring behavior, I defer to after considering the design decisions in the process.

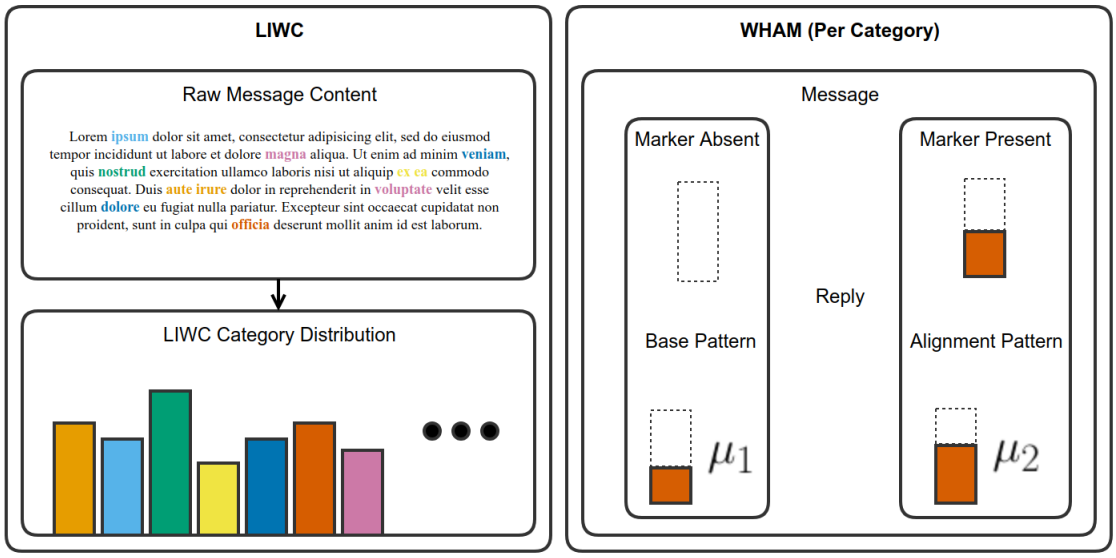

Figure 2.1: Conceptual Diagram of LIWC and WHAM

The mere combination of these linguistic tools is not particularly novel. LIWC and WHAM have been used together in other settings to examine enculturation (Doyle et al., 2017) and emotional alignment (Mobasseri, 2019) in the workplace. The major contribution of my use of this particular technical apparatus here is to leverage its characteristics to capture self-monitoring, particularly that LIWC provides a helpful layer of abstraction, WHAM allows for a clear and thoughtful demarcation of response patterns, and WHAM additionally creates an asymmetric directional measure. These beneficial characteristics help in creating a measure more closely aligned in theory with the construct of self-monitoring.

[1]The heading distribution for the alignment chain is a cauchy distribution.

First, by leveraging the layer of abstraction LIWC provides and by using the majority of LIWC categories versus just a few or one (Doyle and Frank, 2016; Doyle et al., 2017; Mobasseri, 2019), I can create a more general measure of assessment and response that matches the more general theory of self-monitoring. In its abstraction, rather than picking up on the repetition of specific words or phrases, LIWC allows for more flexible patterns of response. For example, if one person smiles at another, a high self-monitor might understand the more appropriate response as a thumbs up rather than a rote mirroring of the smile, paralleling the use of other words in the positive emotion category rather than merely repeating the same positive emotion word or phrase. The ability to aggregate multiple LIWC categories further increases the general patterns that can be identified and may even allow for more complex patterns of self-monitoring.

Second, because WHAM is a *conditional* alignment model, I can create a measure that more closely aligns with the self-monitoring concept of responding to situational cues. WHAM estimates both a base pattern of response and an alignment pattern of response dependent on the presence of a certain marker. That marker is analogous to a certain situational cue to which the self-monitor would respond. Without an understanding of the base pattern of response (when the marker or cue is absent), no reliable inference could be made that the alignment pattern is due to the presence of the marker or cue. There may still be a correlation with behavior (for example, the presence of a positive emotion marker might be correlated with a positive emotion response), but it could not be linked to self-monitoring.

Third, because WHAM is an asymmetrical, directional model, I can create a measure that identifies the locus of action and correctly attributes self-monitoring to an individual at a point in time. Self-monitoring involves a set of actions by individuals to appropriately respond to assessed situational cues. If a symmetrical measure were instead constructed to measure how similar language use is, this identification would not be possible.

Design Decisions and Measures

To create my measure of self-monitoring behavior, first, I extract the set of one-to-one emails with replies from the set of total emails. I believe one-to-one emails provide a cleaner picture of self-monitoring behavior than one-to-many emails, as in the former the self-monitoring behavior is more likely a direct result of conversing with the single individual, while in the latter the cause and effects of the more ambiguous conversational context on self-monitoring behavior is unclear. For example, if multiple people had spoken and an individual was responding to them all in the same message, the cues that influenced the response would be difficult to disambiguate and track. These emails could be a single instance of a message and a reply, or they could be part of a thread, in which case each message-reply would be a separate row in the data set. I additionally remove self-emails, as these would not be relevant for identifying the self-monitoring behavior that I am interested in here. This results in a final data set of 1,178,404 message-reply pairs.

Second, each message-reply is then parsed by the 2015 version of LIWC into corresponding indicators or counts for each LIWC category l, excluding eight aggregate categories due to

the nature of their aggregation, and twenty-one linguistic and grammar categories, as there is evidence these are less susceptible to intentional processes (Ireland and Pennebaker, 2010), leaving sixty-four categories (for a deeper understanding of the implications of this choice, please refer to the robustness check in a later section and Appendix A). In the same row of the data set, messages are parsed in a binary fashion, with a value of 1 for a given LIWC category if the original message contained at least one instance of the category and a value of 0 otherwise, and replies are parsed in a count fashion, with values for the number of words in each LIWC category[2].

Third, this data is then split by month and passed into the WHAM model, which estimates base rates B and alignment rates A for each LIWC category l for ego i based on all messages i replied to in a given month m. I transform the estimates B and A from logit space back to probability space so that the results are more interpretable. The resulting raw probabilities b and a then essentially represent to what degree the LIWC category l will be used in a response[3]. Rather than thinking about the likelihood of using certain categories when markers are present or absent, then, we can interpret the results as some measure of difference between estimated actual category usage patterns. Conceptually, the difference d between these two probabilities represents the extra effort i involves when responding to a message containing LIWC category l. Formally,

$$d_{ilm} = a_{ilm} - b_{ilm} \tag{2.1}$$

Fourth, I aggregate each of these differences over all LIWC categories to obtain a monthly measure of self-monitoring behavior S for i in month m. However, since each category may not be equally represented in the message-reply data set, I weight each difference according to the inverse of the confidence interval of the alignment measure aci for that category, which represents how precisely each rate is estimated. This means that the categories that are less precisely estimated are weighted less in the aggregate measure. Aggregating in this manner allows for a more organic derivation of which categories might be more integral to self-monitoring behavior, which may be dependent on context or time, rather than assuming a priori which may or may not be. For example, it may be that during a tumultuous financial time, the categories of money and tentativeness might be more indicative of self-monitoring behavior and also more prevalent (and more precisely estimated). Therefore, the self-monitoring behavior measure is defined as

[2]Note that each individual word can be counted in multiple LIWC categories. This is not an issue for my current formulation, as the count of terms in any single LIWC category would never exceed the total word count of the email, but in extensions may become an issue if, perhaps, usage of multiple LIWC categories in a response was considered

[3]The more precise understanding is that the raw probability equals the probability of a single independent use of the LIWC category (i.e. the success rate of a single Bernoulli trial). Since the response is assumed to follow a binomial distribution, the number of words in a response multiplied by this raw probability gives the expected number of LIWC category words in the response.

$$S_{im} = \Sigma_{l=1}^{L} d_{ilm} * \frac{\max_L aci_{ilm} + \min_L aci_{ilm} - aci_{ilm}}{\Sigma_{l=1}^{L} aci_{ilm}} \tag{2.2}$$

The various design decisions I made in the process of constructing my measure are briefly summarized in Table 2.2 below.

Table 2.2: Design Decisions with Short Rationales

Design Decision	Short Rationale
1. Use only one-to-one message-reply sets and remove self-emails	One-to-one messaging provides the cleanest way to observe self-monitoring behavior
2. Remove aggregate (e.g. "Clout"), function (e.g. "articles", "prepositions"), and grammar (e.g. "interrogatives", "quantifiers") categories	There is evidence that those categories are not as susceptible to intentional choices in language use (Ireland and Pennebaker, 2010)
3. Use difference in probability estimates as degree of variation in self-monitoring behavior	Provides more easily interpretable results
4. Combine categories using a weighted average by confidence interval size	Some categories, because of usage rates, may not be as precisely estimated

Simply calculating the degree of self-monitoring behavior present can produce misleading results, however. It may be, in this formulation, that the language usage of all ego's communication partners is roughly the same, thereby providing no *opportunity* for ego to self-monitor. In this case, an individual's lack of self-monitoring behavior is different from in the case where there is ample opportunity, and analyses ignoring this distinction would incorrectly estimate not only contextual factors related to self-monitoring behavior but also results of it. Therefore, some measure of the degree of variance in the language use of ego's communication partners must be constructed and controlled for.

In order to do this, I estimate the variance of the base usage probabilities for the LIWC categories among ego's alters *abv*. Taking the set of all email messages E that ego responded to in month m, I first use the WHAM estimates for the base probabilities b for the alters a who wrote those emails, calculate the variance of those base rates across all alters, then sum the variance across all LIWC categories l. In the variance calculation, I include the base rate of the original author for each message, which takes into account communication volume for ego. For example, if individual i sent four emails to which ego responded, then the base probabilities for individual i would be included four times in the variance calculation. Then, for example, if ego responded to one alter once and another alter once in a month, this would

allow for a different level of self-monitoring behavior than if ego responded to one alter fifteen times and another alter fifteen times in a month. Alter base variance, or alter variance for short, is then defined as

$$abv_{im} = \Sigma_{l=1}^{L} Var(b_{alm} \; \forall \; e \in E) \tag{2.3}$$

2.3 Research Setting and Collection of Self-Monitoring Data

To actually construct the measures detailed above, I obtained two sources of data from a mid-sized technology organization. The first type of data was the Snyder and Gangestad (1986) 18-item self-monitoring scale. All employees were asked to complete the survey, answering each question according to a 5-point Likert scale where 5 represented "completely characteristic of me" and 1 represented "not characteristic of me," resulting in a dispersion of responses from 18 to 90. I received in total 436 completed surveys.

The second type of data were email data, which consisted of eight years of not only email metadata (i.e. recipients and timestamps), but also raw message content. Important cleaning steps were removing automatically generated messages and messages sent from or to members of the legal department (per instructions from the company's in-house lawyers). This resulted in an email data set of over five million unique emails. Again, I used the set of one-to-one emails to construct my measures of self-monitoring behavior and alter variance. The limited set of one-to-one emails consisted of over one million unique emails.

2.4 Comparison Between the Self-Monitoring Survey and Self-Monitoring Behavior

The process outlined above allows me to generate monthly values for self-monitoring behavior for individuals, resulting in a granular, longitudinal data set of self-monitoring behavior. Self-monitoring surveys, on the other hand, tend to be cross-sectional. Yet, the time-varying or static nature of the measures does not preclude useful comparison. In the context of the mid-sized technology organization, I consider how collapsed measures of self-monitoring behavior compare to the self-monitoring survey. Because there are significant departures between the two measures in my empirical setting, I then speculate as to why this may be the case, leaving more precise analysis for future work I outline in the limitations and future extensions section.

I begin by correlating monthly values for self-monitoring behavior with the self-report self-monitoring survey values across individuals, which is presented graphically in Figure 2.2 below.

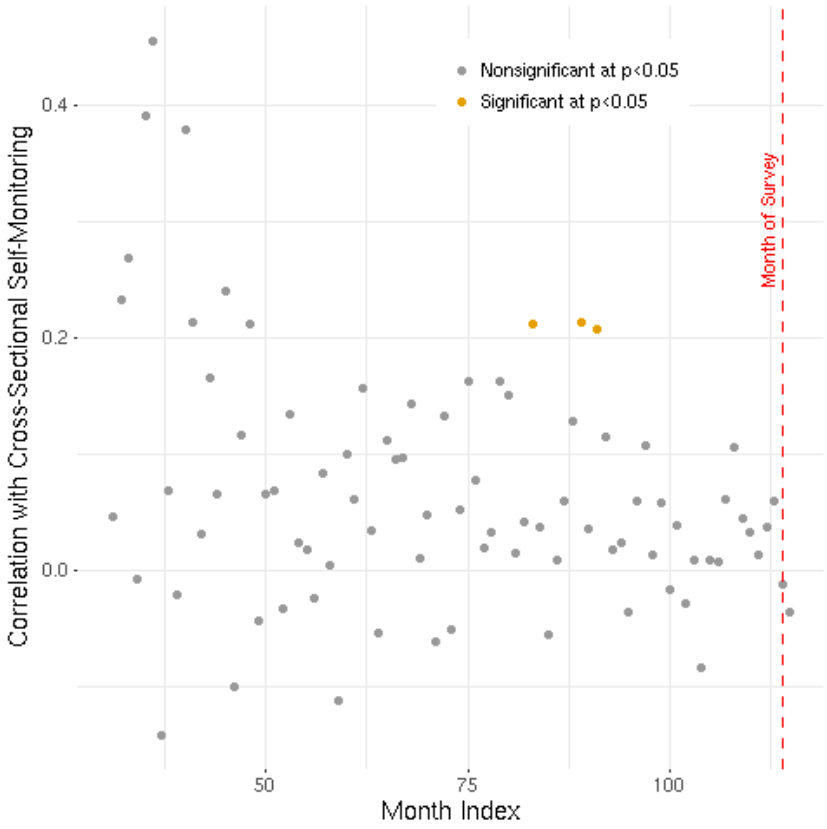

Figure 2.2: Correlation of Self-Monitoring Behavior and Cross-Sectional Self-Monitoring by Month

Each point in the graph represents the correlation between the one-time self-monitoring survey and self-monitoring behavior across all individuals for a single month[4]. In the window of 85 months of observation (months 31 to 115), only three correlations are significant at the $p < 0.05$ level, suggesting that there is a rather stark degree of divergence between the self-report and behavioral measures of self-monitoring. Even considering the months surrounding the month of the survey, in which we might expect the correlation to be the strongest due to tighter temporal connection between behavior and self-report, there are no significant correlations, and the magnitude of the correlations are not above .10.

It may be, however, that a single month's observation of self-monitoring behavior is too volatile to be reliably associated with the self-report self-monitoring. Therefore, in Table 2.3 below, I aggregate self-monitoring behavior across a series of months prior to the start of the survey by taking the mean value for each individual across those months. Aggregation over 114 months represents individuals' mean self-monitoring behavior across the whole observation window.

[4]Note that the number of observations for each correlation does not remain the same due to availability of communication data and entry/turnover of new individuals.

Table 2.3: Correlation of Aggregated Behavioral Self-Monitoring with Cross-Sectional Self-Monitoring (Month 114)

Number of Months Aggregated	Correlation	N	P-Value
3	0.011	269	0.856
6	-0.076	398	0.128
12	-0.046	408	0.352
24	-0.025	410	0.621
36	-0.061	410	0.221
114	-0.085	410	0.085

Again, none of the correlations are significant at the $p < 0.05$ level, and in fact, aside from the first correlation, they are all negative. The most significant correlation, that of self-monitoring behavior aggregated over the whole observation period with the self-monitoring survey, which is significant at the $p < 0.10$ level, is slightly negative. These results, in conjunction with the individual month correlations, suggest that there is some kind of discrepancy between the two methods.

Although I believe these differences are not unreasonable given the whole of personality-behavior theory [5], one reason may be that in an organization and especially in an email network, situational contexts may be strong and would essentially overpower the link between self-monitoring survey levels and self-monitoring behavior. For example, a certain degree of formality and respect may be expected in an email to one's manager, and those considerations may limit the freedom one has to express one's personality—even a low self-monitor may feel the need to use proper punctuation and grammar despite usual habits. If we consider that multiple such layers of context may be operating simultaneously in an organization (e.g. organizational culture, team or department dynamics, macro-economic effects, financial duress, etc.), then it may be reasonable, to use a theoretical cliché, that an individual's self-monitoring behavior is more reflective of the situation than the person. Consequently, because my measure allows for more direct examination of how behavior varies over time and in conjunction with different contexts, I analyze these effects in Chapter 3.

A second reason may simply be that the two measures are measuring different constructs. Considering the multidimensional factor structure in the self-monitoring scale and its developmental history, as well as the fact that I have just introduced my measure without extensive studies to understand and validate it, I believe this is entirely reasonable. I have attempted to design my measure in such a way as to capture true "self-monitoring behavior," but certainly more work can be done to try to reconcile the differences and uncover potential idiosyncrasies in their relationship. To this end, I examine how the inclusion and exclu-

[5] See McClelland, Koestner, and Weinberger (1989) and Leininger and Skeel (2012) for comparable cases of when self-reports differ from behavioral measures.

sion of certain LIWC categories may change my self-monitoring behavior measure to better understand it, and discuss possible future extensions, such as studies of external validation.

2.5 Robustness of LIWC Category Inclusion/Exclusion

There are three primary reasons for considering how the inclusion or exclusion of certain LIWC categories may affect the overall self-monitoring behavior measure. First, I have a made a design choice to remove some categories from the LIWC2015 dictionary in constructing my measure, specifically, eight aggregate categories (e.g. Analytic, Clout, Word Count, Words Per Sentence, etc.) and twenty-one function and grammar categories (e.g. function, pronoun, article, auxiliary verb, etc.). It may therefore be profitable to examine how closely related a full LIWC measure would be with the truncated measure I use in my main analyses. Second, LIWC is a specific framework for understanding underlying psychological processes of word use, and though it has been extensively validated, may be constrained in the categories and words it identifies. Therefore, considering how the self-monitoring behavior measure may be robust to a reduction of categories would provide some indication of how the measure might behave with other "missing" categories or words not included in LIWC. Third, it will provide a window into the internal structure of my measure of self-monitoring behavior. By considering which categories may have disproportionate effects when removed, we can get a better grasp as to what is driving self-monitoring behavior.

To do so, I use a bootstrapping procedure that correlates self-monitoring behavior constructed on more or fewer LIWC categories with my main sixty-four LIWC category self-monitoring behavior[6] (labeled "Full LIWC" in the analysis). To be precise, I cycle through the number of LIWC categories removed, starting from zero all the way to sixty-three, correlating each resulting measure with the original sixty-four LIWC category measure. Because the number of possible combinations increases factorially with the number of categories removed, I randomly sample 100 such combinations when the number of possible combinations exceeds 100. For example, if I were to select three categories to be removed, there would be 41,664 such possible combinations. Then, because my dataset contains 25,817 observations of self-monitoring behavior, instead of using all the observations to calculate the correlation, I use a random sample of 5,000 observations fifty times in order to simulate draws from the population of observations and not just my sample. I use the mean of those fifty draws to estimate the effect of removing those particular categories, then use the mean and standard deviation of all 100 category combinations to estimate and bound the effect of removing that particular number of categories. The resulting graph is presented in Figure 2.3 below.

[6]This approach is conceptually similar to the backward selection procedure in Srivastava et al. (2017) but is distinct in considering the random exclusion of a certain number of LIWC categories and the resulting measure's correlation with the original measure.

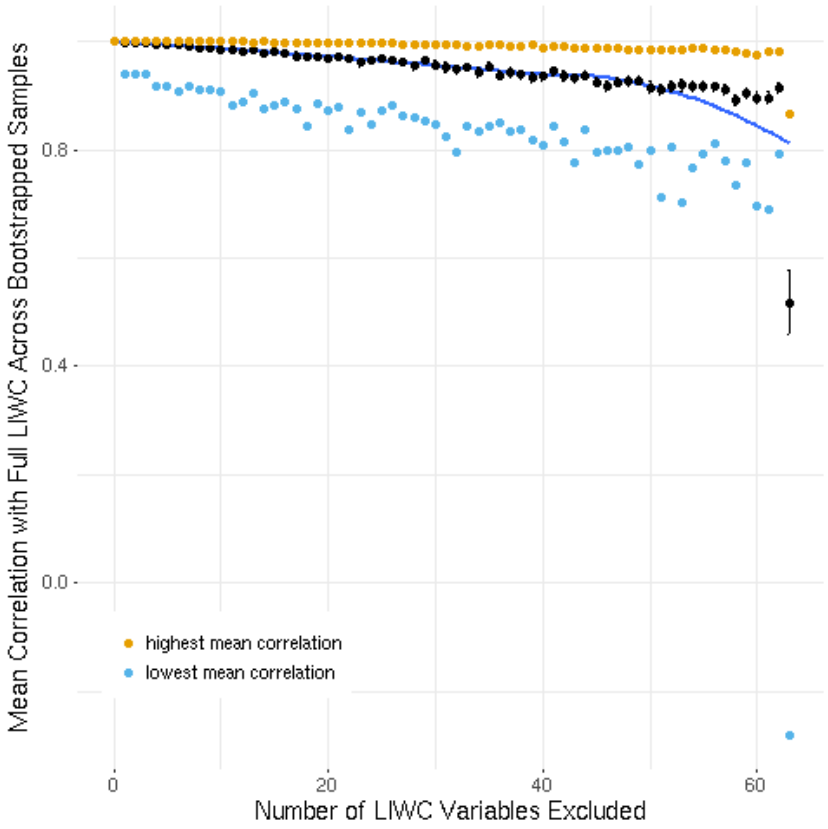

Figure 2.3: Correlation Between Full and N-Category-Excluded LIWC Self-Monitoring Behavior

Although not illustrated in Figure 2.3 above, the correlation between the sixty-four category measure and the eighty-five category measure is approximately 0.9270 (0.9267, s.d. 0.002 using the bootstrapping procedure), indicating that the two measures are rather similar. However, a replication of my results in Chapter 3 with the full eighty-five category measure is presented in Appendix A, and a brief discussion on the differences in results with respect to the twenty-one excluded function and grammar categories is included.

From Figure 2.3, though, it is rather immediately clear that the removal of any random set of categories does impact the measure but not dramatically. Even when only a few categories are included, on average, self-monitoring behavior derived from those categories imitates relatively well the full measure based off of sixty-four categories. The interpretation, then, is that self-monitoring behavior is not purely or even significantly disproportionately driven by one or a few LIWC categories, but is instead capturing a more general pattern of self-monitoring. That being said, however, the blue dots representing the lowest correlations by number of categories excluded demonstrate that some categories *are* more influential than others, as excluding those specific categories results in a greater correlation drop than excluding the same number of randomly selected categories (on average). This could likely

happen for at least two reasons, one being that individuals are aligning more on those categories and therefore their effects are being weighted more heavily in the aggregate measure, and the other being that those categories are generally more used leading to them being more precisely estimated and also therefore being weighted more heavily in the aggregate measure. For either of these reasons, it may be helpful to examine those specific categories.

I approach this in a series of three analyses, which, in conjunction, provide a better picture of what my self-monitoring behavior measure is identifying. First, I take a look at the actual categories excluded for both the lowest and highest correlations by number of categories excluded. Seeing the same categories excluded gives an indication of how relatively important (for lowest correlations) or unimportant (for highest correlations) those categories are. However, it may be that other categories, by random chance, happen to be selected with relatively important or unimportant categories. Therefore, second, I look at the full measure correlated with each single category excluded in turn. And also to that point, third, I look at the correlations among measures constructed by only a single category, as it may not make a significant difference on the overall measure if only one of many highly correlated categories is excluded. For example, if the male and female categories were highly correlated, such that when an individual aligns on male references, he or she also tends to align on female references, then excluding only male or only female from the sixty-four category measure may not make a significant difference in the overall correlation, as the information from the other category would still be included.

Tables 2.4 and 2.5 below indicate the relatively important (Table 2.4) and relatively unimportant (Table 2.5) LIWC categories for my measure of self-monitoring behavior. The categories that tend to be excluded for the lowest correlations are punctuation categories (e.g. OtherP ["other punctuation"], AllPunc ["all punctuation"], Quote, Parenth, QMark, SemiC, Colon, Dash), which is reasonable considering those are visually the most noticeable categories. If an individual is communicating with another individual who tends not to use proper punctuation—it may be easy to leave off apostrophes, periods, or commas in more informal communication—or who does use proper punctuation, then the former may be inclined to reciprocate. We also find that cognitive processes shows up a few times, which may be indicative of the process of understanding and relating, and netspeak shows up early on, perhaps indicative of the strongly informal nature of netspeak and therefore the impetus to align. All this, however, must be interpreted in light of the impact on the correlation, which is not that large. Even though punctuation may be a more central component of my self-monitoring behavior, it is far from the defining feature.

Table 2.4: Excluded LIWC Categories by Number of Categories Excluded (Lowest Correlations)

Excluded	Correlation	Categories Excluded
1	0.9407	OtherP
2	0.9406	OtherP, reward
3	0.9393	cogproc, family, OtherP
4	0.9187	anger, Dash, netspeak, OtherP
5	0.9157	AllPunc, drives, OtherP, Quote
6	0.9080	AllPunc, cogproc, differ, discrep, insight, OtherP
7	0.9165	AllPunc, anger, bio, friend, leisure, OtherP, QMark
8	0.9119	AllPunc, cogproc, focusfuture, friend, hear, OtherP, power, social
9	0.9099	anx, body, Colon, Dash, death, male, OtherP, SemiC
10	0.9079	AllPunc, bio, health, home, motion, OtherP, Parenth, Quote, see

The categories that tend to be excluded for the highest correlations are harder to interpret, as many categories may be roughly equally uninformative. For example, we see that the correlation remains at one, whether the categories of sad, death, or risk are removed (and death and risk jointly). Even when ten categories are removed, the correlation remains almost at one. I would caution, however, against interpreting these categories as wholly uninformative; though in the aggregate their inclusion or exclusion may make little difference on the overall measure, they may make a more significant impact depending on the unit of analysis. For example, during difficult organizational times (i.e. a financial crisis), which may not happen as often, aligning on sad, risk, or negemo ["negative emotion"] may become significant and relatively important indicators of self-monitoring behavior.

Table 2.5: Excluded LIWC Categories by Number of Categories Excluded (Highest Correlations)

Excluded	Correlation	Categories Excluded
1	1.0000	sad
2	1.0000	death, risk
3	0.9999	family, health, negemo
4	0.9999	anx, bio, insight, nonflu
5	0.9999	hear, insight, nonflu, relig, swear
6	0.9999	body, death, feel, friend, insight
7	0.9995	achieve, family, feel, filler, netspeak, Parenth
8	0.9995	cause, death, discrep, feel, hear, percept, QMark
9	0.9995	anger, assent, family, focusfuture, informal, ingest, motion, nonflu, QMark
10	0.9993	anx, body, differ, family, home, posemo, reward, sexual, swear

Figure 2.4, which correlates the full sixty-four LIWC measure with each sixty-three LIWC measure, corroborates many of the inferences drawn thus far. We find that excluding other punctuation makes the most significant difference, with other punctuation categories also high up in the list. We also find cognitive processes in the sixth position and netspeak in the first half. However, it does clarify some points from the first analysis, as we find anger, reward, anx, death, body, and other categories closer to the bottom of the list, even though they were included in Table 2.4. It is likely, then, that they were randomly sampled with more important categories, but are themselves less informative.

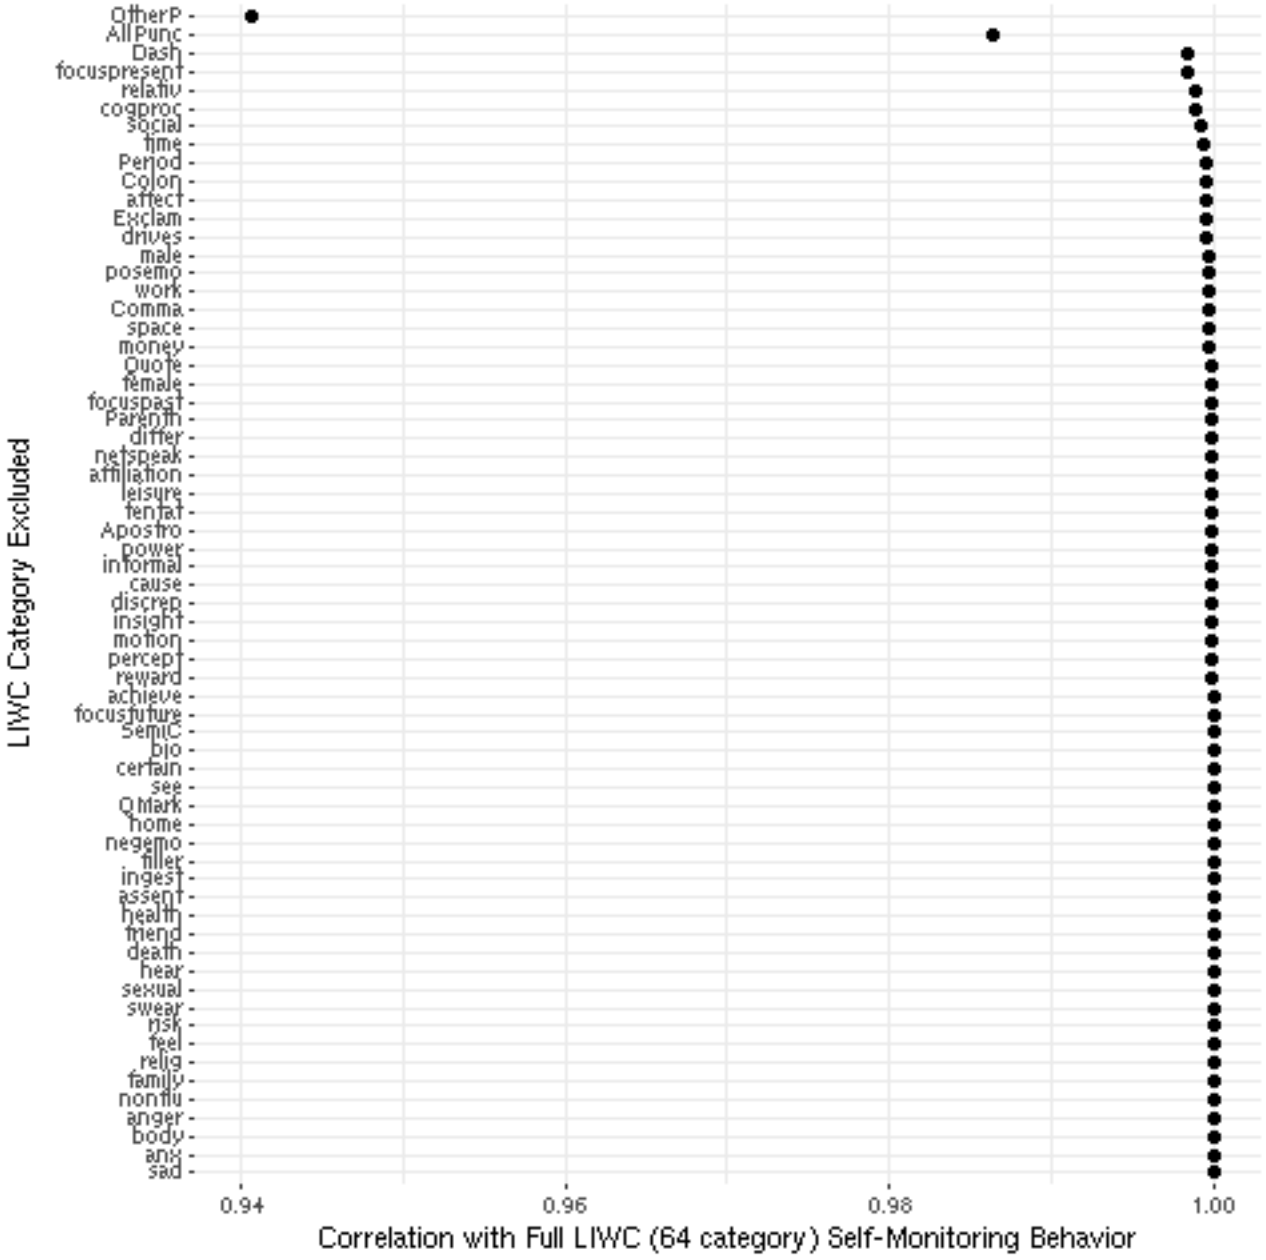

Figure 2.4: Correlation Between Full and One-Category-Excluded LIWC Self-Monitoring Behavior

However, it may be that they are informative in conjunction with other categories, taking into account the underlying correlation structure (though this is far less likely in the context of the first analysis). If we examine Figure 2.5 below, we can see that many of the categories are in fact correlated. For example, aligning on the category of home (when a communication

partner uses more "home" words, you also use more "home" words), is correlated with aligning on the category of friend. Perhaps these categories are often used together, and therefore in the natural unfolding of the communication process, alignment occurs in both categories.

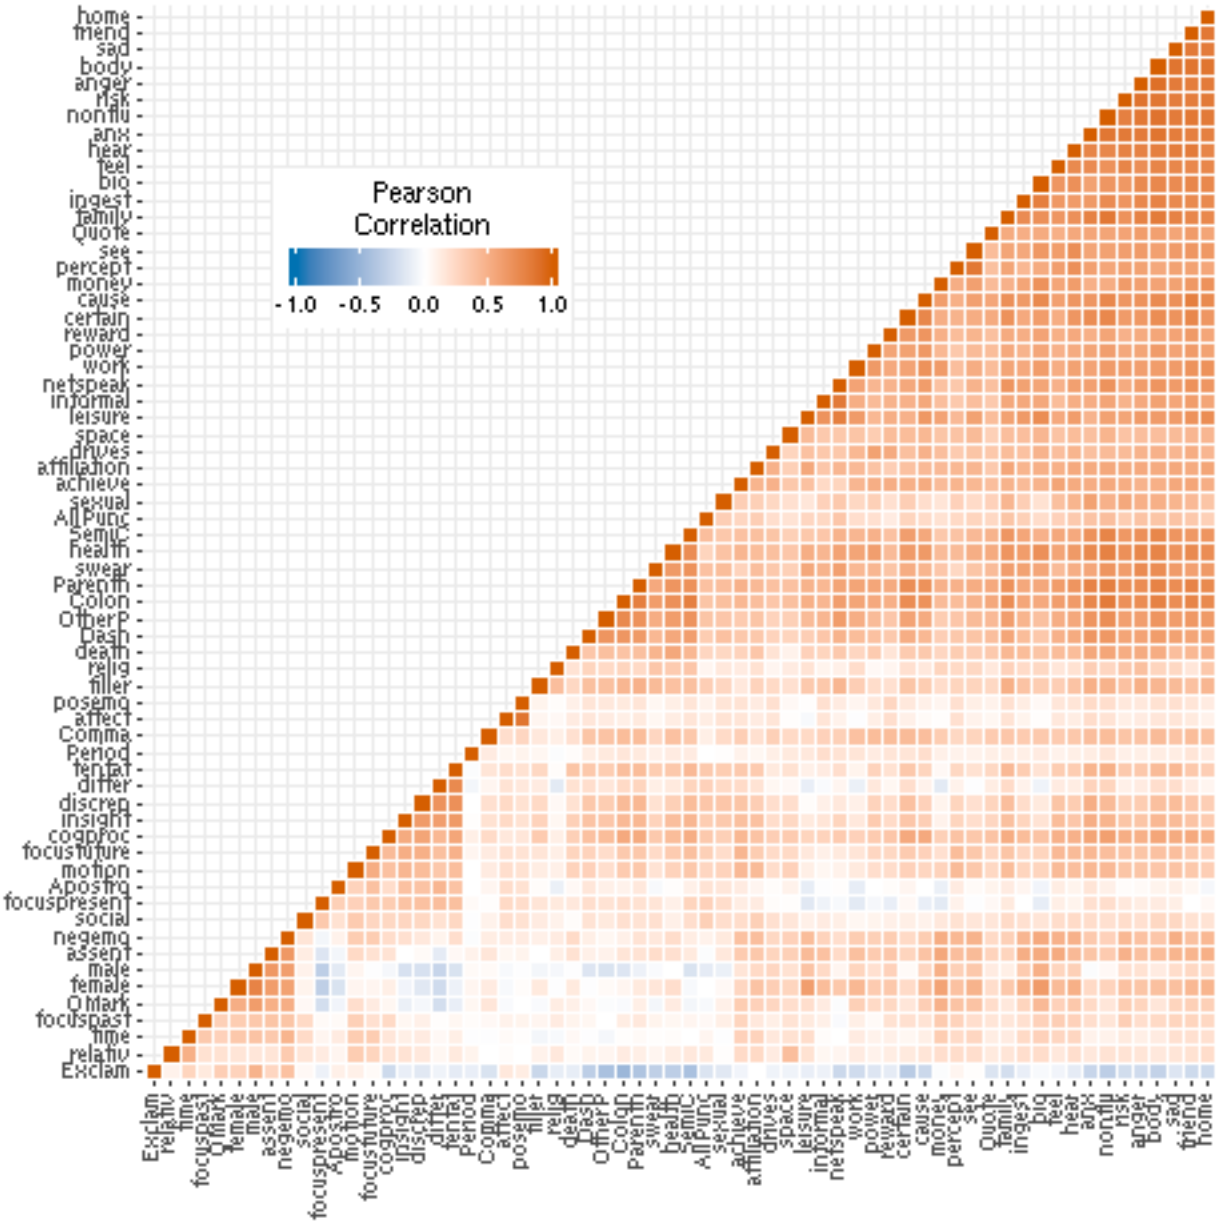

Figure 2.5: Correlation Between Single-Category LIWC Measures

It may be easier to interpret by examining a smaller subsection of the full correlation

graph, presented in Figure 2.6 below (and example words for each category are presented in Table 2.6 below the Figure). If we look at the rightmost column for tentative words, we see that it is highly correlated with differentiation words, discrepancy words, insight words, cognitive process words, and nonfluencies. These in conjunction indicate not just a set of words that are being used, but more than that, a way of communicating or a purpose for communication. If communication contained these categories, it might be an attempt to solve a problem, brainstorm together, express an uncertain opinion, etc. Correlation among alignment for these categories then demonstrates a deeper process of self-monitoring than simply repeating words, one that involves understanding the context and responding appropriately. With the understanding that these correlations exist not only among cognitive process categories, but also many other categories, this interpretation then suggests that my linguistic measure is picking up on complex processes of self-monitoring.

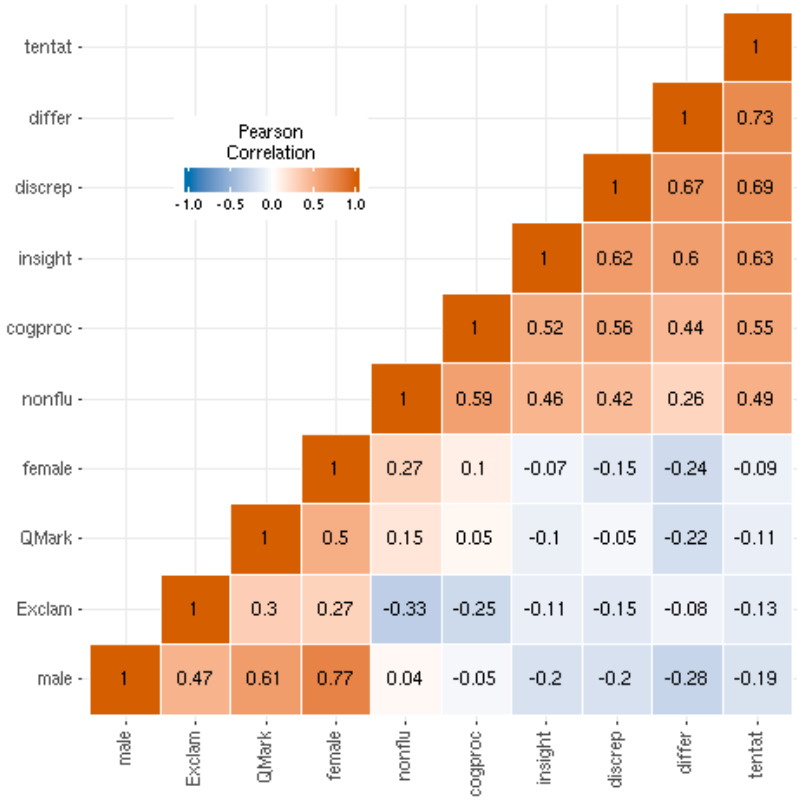

Figure 2.6: Correlation Between Single-Category LIWC Measures (Sub-sample)

Table 2.6: Example Words in LIWC Categories

LIWC Category	Full Name	Examples
male	Male references	boy, his, dad
Exclam	Exclamation marks	!
QMark	Question marks	?
female	Female references	girl, her, mom
nonflu	Nonfluencies	er, hm, umm
cogproc	Cognitive processes	cause, know, ought
insight	Insight	think, know
discrep	Discrepancy	should, would
differ	Differentiation	hasn't, but, else
tentat	Tentative	maybe, perhaps

2.6 Limitations and Future Extensions

I started this chapter by considering how to construct a linguistic measure of self-monitoring behavior, and admittedly, the procedure I have followed is perhaps one of many possible methods. There are at least three notable limitations, which I hope future work will address and move us closer to the idea I presented in Chapter 1 of pairing traditional survey-based methods of personality assessment with behavioral measures to better understand the individual from an empirical standpoint. First, there is an implicit tradeoff that increasing complexity of a model or algorithm often portends decreasing interpretability. Admittedly, even with my attempts in this chapter to diagnose the internal structure of my linguistic measure of self-monitoring behavior, it is still difficult to exactly explain what it is identifying, and more efforts can certainly be pursued in this endeavor. Second, the state-of-the-art in computational analysis is constantly improving, and comparable but alternate specifications, with their own set of assumptions, may improve my measure in getting increasingly closer to the theoretical construct of interest. Third, I have tried to demonstrate how a thoughtful linguistic measure of the personality construct self-monitoring may be developed, in anticipation that similar procedures may be followed for other personality constructs. Yet, I have not made it explicitly clear how one might intend to do so. In this last section of this chapter, I hope to discuss these limitations and potentially profitable future extensions.

To better understand the measure itself, I believe two major steps can be taken. A first step would be to get some sense of external validity. While the emails in my setting are confidential, there are a number of public communication datasets that could be used to validate my measure, the most notable of which is probably the Enron email dataset. A comparable dataset of self-monitoring behavior could be built off of the Enron emails, and the raw emails themselves could even be presented to get peer-reported measures of self-monitoring using the self-monitoring survey. In this way, it could be more objectively identified how my

linguistic measure compares to the survey-based measure of self-monitoring. A second step would be to leverage the possibilities already available in the current technical apparatus to try to get at clearer and cleaner pictures of individual behaviors. For example, as it stands now, the measure considers instances of one LIWC category being used when the preceding message did or did not contain the LIWC category. However, this is easily extendable to considering instances of *other* LIWC categories being used, or a *set* of LIWC categories being used. Perhaps when one receives an email with anger, the response more indicative of self-monitoring is not responding with anger but instead with positive emotion, or positive emotion and perceptual words.

To consider how other computational tools may improve the measure, I will examine two briefly, Analyzing Linguistic Interactions with Generalizable techNiques (ALIGN; Duran, Paxton, and Fusaroli, 2018) and Empath (Fast, Chen, and Bernstein, 2016), and expect a similar thought process would take place for future novel methods. ALIGN is a comparable method of retrieving linguistic alignment values, differentiating between lexical, syntactic, and conceptual alignment. Turn-by-turn exchanges are transformed into vectors of tokens (for lexical alignment), parts-of-speech tags (for syntactic alignment), and high-dimensional semantic representations (for conceptual alignment), and cosine similarity with one's real interlocutor across those vectors is compared to cosine similarity with a surrogate fake interlocutor taken from another conversation to calculate linguistic alignment. While the possibilities of a more interpretable alignment measure and arbitrarily parsing tokens for different kinds of alignment (for example, using LIWC instead of parts-of speech) easily are desirable, the fact that the main method of measuring alignment is, in the end, a similarity measure, versus a conditional response, makes ALIGN less amenable for capturing self-monitoring. It may, however, be a profitable consideration for other personality constructs. Empath, which is a "tool that can generate and validate new lexical categories on demand" (Fast, Chen, and Bernstein, 2016, pg. 1), on the other hand, seems like a parsimonious improvement to using LIWC categories. While the reported correlation between Empath and LIWC on a mixed-corpus dataset is 0.906, Empath offers more flexibility and space for exploration.

Finally, to illustrate how my procedure may be extended to other personality constructs, it may be useful to summarize at a high level, without all the technical details, the process I followed to create a general linguistic measure of personality. First, I would consider what linguistic behaviors are exemplary of the personality construct, turning to the established literature for that personality construct. Second, I would consider what linguistic tools would allow for the most direct capturing of that behavior, also considering interpretative tradeoffs. Third, I would consider all the possible design choices that one could make within that specific technical apparatus, evaluating which might result in a measure closest to the theory. Fourth, I would find ways to validate the measure, comparing it to the survey, and testing for internal and external validity. After these steps, hopefully the use of the measure in different contexts and empirical settings would prove fruitful for advancing research and further validation.

CHAPTER 3: The Role of Context for Self-Monitoring Behavior

The linguistic measure of self-monitoring behavior I have outlined in Chapter 2 opens up a wide range of research possibilities; rather than differentiating individuals on the basis of a static trait of self-monitoring, it allows us to measure the degree to which individuals self-monitor at a given point in time. Returning to Chapter 1 and integrating the setting and data of the current study, the benefits are threefold. First, the evidence in Chapter 2 comparing the survey measure with the linguistic measure of self-monitoring demonstrated that individuals across the spectrum of self-monitoring survey levels can display similar levels of self-monitoring behavior, suggesting that some interpretive processes might have taken place in translating personal behavior to personal assessment. With this new measure, we can avoid the cognitive biases of self-report and more objectively analyze behavioral expressions of personality. Second, by leveraging the natural variation that occurs within an organizational context, we can begin to consider how temporal and situational effects influence self-monitoring behavior. Third, we can expand on existing evidence linking high or low static, self-reported levels of self-monitoring to outcomes by directly observing the self-monitoring behavior associated with those outcomes.

To illustrate these more generalizable benefits of linguistic complements to traditional survey-based measures of personality, I examine in this chapter the effect of situational context on self-monitoring behavior, and the outcomes of self-monitoring behavior for individuals. The impact of context on self-monitoring behavior is important to understand in light of the variety of contexts in which individuals are placed in organizations. An individual may go from a casual conversation with a co-worker to a structured business meeting to a CEO update in the course of a single day, let alone the course of his or her entire tenure in the organization. How that individual conducts him or herself in those situations and how his or her behavior might differ from his or her normal behavior has important implications for the internal stability of the organization and the organization's success at large. Understanding how specific contexts and their characteristics affect self-monitoring behavior, then, can make a significant difference towards better management, as managers may be better able to assess and respond to the true motivations, needs, concerns, etc. of their employees.

As an example, consider a hypothetical, traditional performance review situation. If the performance review only happens once a year, is the main determinant of the employee's annual bonus, and is conducted not by the employee's direct manager but the head of operations, one could imagine that the incentives to cultivate a positive and potentially false image are high. The employee may be unwilling to discuss true areas of need, where the manager could work with and support the employee, leading to decreased job satisfaction, productivity, and perhaps even turnover. It would seem that modern organizations, sensing such contextual constraints, have moved to more continuous performance review programs or dissociated performance review from pay. This is not to say that all self-monitoring is necessarily negative for the organization and the individual, but rather to highlight the importance of understanding the role of context for self-monitoring behavior.

Therefore, to begin to unpack the more systematic effects of context on self-monitoring behavior, I build on Goffman's theory of self-presentation (1959) and work in psychology on situational strength (for a review, see Meyer, Dalal, and Hermida, 2010) to present here a simple framework for considering how situational contexts might affect self-monitoring behavior. Specifically, I focus on two characteristics of the context, uncertainty in the appropriateness of behaviors (to which I will refer to as uncertainty, for short) and severity of consequences for inappropriate behaviors (to which I will refer to as consequences), and theorize how they affect self-monitoring behavior in individuals. While there may be other characteristics of the context or even ways to characterize the context, I believe these two provide a succinct summary of a context meaningful for the analysis of self-monitoring behavior. I examine context from an organizational perspective, specifically as the varying individuals an employee communicates with over email, and tie my theoretical reasoning back to this perspective.

In the remainder of this chapter, I build out the theory of these two contextual characteristics and how I hypothesize they relate to my main construct of interest, self-monitoring behavior. Then, I explain my methodology, describing my research setting and detailing my analytical strategy. I present the main results of the empirical work, finding support for my hypotheses, and supplementary analyses that demonstrate the effect of self-monitoring behavior on individual outcomes. Finally, I conclude with a discussion of the main findings of my work, broader implications of the results, the limitations of this particular study, and future directions for research.

3.1 Theory

A Theoretical Framework of Context

That context constrains behavior is a central theme of social science research. Whether it is found in the sociological literature in the embeddedness of rational action (Granovetter, 1985), where the social context surrounding an individual shapes the economic actions he or she takes, or the psychological literature in the situational or interactionist approach to personality (Snyder and Ickes, 1985), it is apparent that behavior does not happen in a vacuum. Yet, contexts, or as the literature would call them, situations, are not always easy to define, let alone categorize such that their effects on behavior can be analyzed. Every situation is in some way idiosyncratic, which may explain, at least in part, why Goffman's (1959) thesis of self-presentational dynamics has not found broader empirical interest and support (Magnusson, 1981: 13). It may be too insurmountable a task to aggregate innumerable situations into a discrete and succinct set of meaningful ideal types (Frederiksen, 1972).

There has been, however, a rich literature that has tried to do exactly this. As two examples, Block and Block (1981) asked fifteen experienced experimenters to fill out a Q-sort (Block, 1961) for the psychological demand qualities (e.g. "performance in the situation emphasizes speed in behavior," and "the context evokes creativity, imagination") of 54 different situations. Then, by inverse factor analysis, they were able to boil down these psychological

demand qualities into eight distinct factors, and, translating these into summary abstractions, such as structure, convergency, evaluation, etc., defined a set of situational parameters to taxonomize situations. Similarly, Edwards and Templeton (2005) generated a list of adjectives that might describe situations and asked respondents to rate situations using samples from the list. Their factor analysis revealed four general factors, positivity, negativity, productivity, and ease of negotiation, by which they categorized situations. For a more complete list of twenty-one extant situational taxonomies, please refer to Rauthmann et al. (2014).

Nevertheless, while taxonomies may represent some useful conceptual abstraction of situations, they each face important scope conditions that may limit their general usefulness (Meyer, 2015). They often operate within a bounded domain of situations and are especially sensitive to the process by which the taxonomies were generated. The domain of situations in the Block and Block (1981) taxonomy, for example, may be limited to a set of readily testable experimental paradigms, as many of them are simple games or psychological tests (e.g. Peabody Picture Vocabulary Test, Simon Says). Moreover, the situations were categorized on the results from fifteen experimenters, which means that the resulting factor structure may be due to the "perhaps special character of the experimenter population available" (Block and Block, 1981: pg. 92-93). Even with greater capacity to survey more representative samples in large scale today, any taxonomy generated would be susceptible to some degree to this boundedness.

In light of the difficulty of a single situational taxonomy to address general theoretical issues in the situations literature, an alternative approach that has been suggested in the literature has been a *framework* of situations. This framework seeks to define the elements of a situation, namely the cues and characteristics of a situation, which may then group together into classes of situations (Rauthmann, Sherman, and Funder, 2015). Cues represent the composition of situations, or those relatively objectively quantifiable stimuli, such as the persons and relationships of a situation; the objects in a situation; the events and activities taking place in the situation; or the time and location of the situation. Characteristics summarize the psychologically important meanings of perceived cues, such as dutiful, pleasant, negative, etc. Classes represent groupings of situations based upon similarity of characteristics or cues; most situation taxonomies, then, in this framework, organize entire situations into different classes.

Understanding situations in this framework then gives us a vocabulary within which we can start to theorize how situations affect behavior, specifically, self-monitoring behavior. While there may be idiosyncratic effects of the same situation on different individuals provided each individual may perceive different characteristics of the situation, in general, there should be a salient characteristic which can be reliably diagnosed from the available cues in the situation (Gosling et al., 2002). These salient characteristics, in turn, affect behavior, since cues must first be interpreted before they are responded to. And, classes of situations should be expected to affect behavior in similar ways.

In this conception, the same salient characteristics are expected to consistently effect the same types of behavior. However, it may not always be that the same cues lead to the same salient characteristics. For example, the act of someone tripping (cue) may contribute to

an embarrassing (characteristic) situation under normal considerations, but when a camera is present (cue) or a funny sound is played (cue), the same cue may contribute instead to a dutiful (characteristic) situation, where one must trip in order to properly play the part. Nor may it be that only a singular salient characteristic is present in any given situation. A business meeting, for example, can be simultaneously tense, respectful, positive, and collegial, aggregating to a general atmosphere of professionalism. While important to consider, these complexities are outside the scope of the current study. For the purposes of my theorizing, I will consider the independent effects a single salient characteristic may have on self-monitoring behavior, and detail the expected relationship between situational cues and those characteristics within my organizational context.

Finally, given that my main concern in this study is the effect of situations on self-monitoring behavior, Goffman's metaphor and examples will be particularly helpful in bridging self-presentational dynamics and the aforementioned situational framework. It provides a set of easily understandable key terms, including actor, performance, roles, and audience through which cues and characteristics can be hypothesized to affect self-monitoring behavior. Therefore, as I hypothesize which situational characteristics affect self-monitoring behavior, I rely heavily on Goffman's dramaturgy.

Uncertainty in the Appropriateness of Behaviors

The first characteristic I hypothesize affects self-monitoring behavior is the uncertainty in the appropriateness of behaviors. In Goffman's dramaturgical conceptualization of interactional life, there are actors who perform in certain roles, and there are audiences. The roles that the actors play are often a function of their own self-presentational desires in accordance with the interpretation and needs of the audience. For example, in the case of a restaurant waiter, the role of a respectable waiter is played by an individual to an audience of restaurant diners. Because of their payment for services, the diners may expect the waiter to be timely, courteous, and precise. Any deviation from the diners' expectations, such as a mistaken order, results in a breaking down of the waiter's "act," thereby weakening his or her playing the role of a respectable waiter. Immediately, then, we can see that the audience of diners defines a set of appropriate and inappropriate behaviors for the waiter to perform.

Yet, even in this simple example, we can see that there could be a wide range of appropriate behaviors that are expected of a timely, courteous, and precise waiter. It could be that some diners expect the waiter to arrive within 2 minutes of their sitting down to be timely, while others expect the waiter to arrive within 5 minutes, and still others 10 minutes. It could be that some diners consider small talk from the waiter as courteous, while others demand nothing more than listing the specials and taking orders. It could be that some diners expect the waiter to write down their order, even if short-term memory may be more efficient and just as precise. There is a measure of uncertainty for the actor in what constitutes appropriate behavior in the role of a respectable waiter.

Although this measure of uncertainty is fundamentally within the cognition of the actor—different actors may construe different levels of uncertainty given the same situation—there

are factors of the context that should affect the propensity of actors to be uncertain, namely the clarity and consistency of the context (Meyer, Dalal, and Hermida, 2010). The clarity of the context concerns the extent to which situational cues are available and easy to understand (Bowles, Babcock, and McGinn, 2005). The waiter situation may have greater clarity, as the waiter could use the facial expressions of various diners as cues for appropriate behavior; if diners are beginning to get upset, it may have been too long to be considered timely. The consistency of a situation concerns the extent to which situational cues are compatible with one another (Bacharach and Bamberger, 2007). The waiter situation may also have greater consistency, as getting orders correct generally elicits the same positive response in diners. Therefore, these represent the basis by which contexts can be categorized as more or less uncertain with respect to the set of appropriate behaviors.

I hypothesize that a greater uncertainty resulting either from less clarity or less consistency in a context should lead to more self-monitoring behavior. In either case, actors would have more difficulty in ascertaining an appropriate set of behaviors and therefore would be more sensitive and responsive to whatever few, ambiguous, or oppositional cues may be present. In other words, with the original understanding that the construct of self-monitoring involves regulating one's self-presentations to accord with situational cues (see Chapter 1 for a discussion on how the construct may have shifted over time), if the task is made harder *ceteris peribus*, it should involve more work (i.e. more self-monitoring behavior).

In my organizational context, I propose that a key factor influencing this characteristic of the situation is whether an employee is communicating with a new interlocutor, or someone with whom they have never communicated before, or an old interlocutor, or someone with whom they have communicated before. The reasoning is that individuals, especially when they communicate via email, may learn and follow behavior-outcome expectancies (Mischel, 1973), or, they learn what to say and how to say it. For example, knowing that a specific communication partner does not always read long emails may cause a focal individual to always respond in concise statements; this, however, to another communication partner may seem rude and short. These behavior-outcome expectancies, however, are not often known without prolonged exposure to or information about a communication partner. With respect to less clarity then, communicating with new interlocutors results in cues being less easy to understand or even less noticeable (i.e. available). With respect to less consistency, communicating with new interlocutors results in more perceived oppositional cues, as the focal individual may not have the understanding required for synthesizing communication into a more coherent picture. Therefore, it should be that responding to new interlocutors involves a greater degree of uncertainty about appropriate behavior than responding to old interlocutors.

H1: The more a focal individual communicates with new interlocutors versus old interlocutors, the more that individual will display self-monitoring behavior.

Severity of Consequences for Inappropriate Behaviors

Given a certain level of contextual uncertainty, there still may be moderating factors that affect the self-monitoring behavior of individuals. I propose that one key moderating factor is the perceived severity of consequences for inappropriate behaviors. In the dramaturgical metaphor, these inappropriate behaviors or "misperformances" represent the classic faux pas, a breaking of the act of the actor in the role he or she is playing that generally threatens the veracity of the role itself in front of the audience. Again, in the case of the waiter, a more subtle yet important misperformance, at least in American culture, would be failing to refill a diner's water. Because it is commonplace to expect waiters to do so, failing to refill water may challenge the image of the waiter as competent.

The severity of consequences represents the extent to which decisions or actions have important positive or negative implications (Meyer, Dalal, and Hermida, 2010). In an organizational context, although the more prototypical consequences are aligned with the organization's formal bureaucracy, they need not always be so. In addition to bureaucratic consequences such as promotions, layoffs, performance ratings, pay bonuses, etc., there may also be less tangible consequences such as job satisfaction or emotional support. Moreover, while there may be actual consequences for certain behaviors, I propose that, for the purposes of influencing self-monitoring behavior, perceived consequences may be more proximate, although undoubtedly the two are related. Because self-monitoring behavior involves modifying one's behavior, the process underlying the change must relate to the perceived severity of consequences, before the behavior is actually acted out.

To understand how the severity of consequences may vary across situations, I suggest that the consequences of misperformance are at least partially dependent on the audience. The audience can choose to overlook a simple mistake and consider a misperformance nonsignificant or noncharacteristic of the individual, or the audience can consider the misperformance so deleterious to the act that it is no longer believable the individual could play the role. In the former case, the implications are small; often, the actor can continue in his or her role untroubled. In the latter case, however, the implications are severe; the actor may end up losing his or her role permanently.

Again in my organizational context, I propose that two audiences will be particularly distinctive in this respect, interlocutors in a focal individual's department and interlocutors who are more senior (i.e. those with greater tenure in the organization) than the focal individual. For the former, communication misperformances may take the form of not understanding departmental lingo; not responding in a way that demonstrates proper attention, care, or concern; or perhaps being too formal or serious. For these misperformances and others like them, not only may individuals in one's department represent one's in-group and therefore a more psychologically safe environment (Edmondson, 1999), but also there would likely be other situations (e.g. face-to-face conversation, business meetings, collaborations), in which a focal individual could rectify any slights in their emails. At any rate, I would expect that the severity of consequences would be diminished, and therefore the focal actor would display less self-monitoring behavior.

H2: The more a focal individual communicates with interlocutors in his or her department, the less that individual will display self-monitoring behavior.

For the latter, communication misperformance may take the form of not displaying the proper formality or respect, not understanding the potential time constraints on the more senior employee, or not writing according to the more established norms of the organization. More senior interlocutors likely command a certain level of seniority and respect and additionally may have more status or power in the organization (Emerson, 1962), thereby influencing more consequential outcomes in the organization. Therefore, I would expect that the severity of consequences when communicating with this audience would be heightened, and that the focal actor would display more self-monitoring behavior.

H3: The more a focal individual communicates with more senior interlocutors, the more that individual will display self-monitoring behavior.

3.2 Methods

Data

My setting is the same mid-sized technology firm detailed in Chapter 2. In addition to the email data and the self-monitoring survey data described there, I also obtained personnel records, which consisted of monthly extracts from the firm's human resource information system. Key variables include age, departmental affiliation, start date, managerial status, and monthly bonus received.

Operationalization of Context

In order to operationalize my theoretical contexts, I assume that a single one-to-one email exchange constitutes a local context. This local context is distinct from other levels of context, such as an organizational cultural context, or a subcultural context, and is particularly salient at the time the individual crafts their email response. Then, by aggregating these local contexts according to the two more general characteristics in the theory section—uncertainty in the appropriateness of behaviors and severity of consequences for inappropriate behaviors—I derive variation in situational contexts among individuals over months. Formally, I use the percent of an individual's one-to-one emails sent to a certain group for each month as a measure of the overall context.

To test hypothesis 1, I use the percent of communication (by number of emails) to new interlocutors, where again new interlocutors represents individuals with whom a focal individual has not previously communicated at the month of observation. To test hypotheses 2 and 3, I use the percent of communication to interlocutors in the same department as the focal individual and the percent of communication to more senior interlocutors, respectively.

For these, however, I take the additional step of conditioning on only the one-to-one emails sent to new interlocutors, as this provides a cleaner picture of their moderating effects without confounding the degree to which the focal individual may be familiar with old interlocutors.

Control Variables

I include several controls in addition to alter variance as detailed in Chapter 2. First, I control for whether the individual is in their first year in the organization or not, as this generally is a tumultuous time for new employees, which may affect incentives, abilities, and opportunities with regards to self-monitoring behavior. Second, I control for managerial status, as being a manager would affect both who one talks to as well as the self-monitoring behavior in those conversations. Third, I control for log emails as a measure of the volume of email exchange, which would also have an impact on the opportunity for self-monitor behavior. Fourth, I control for department affiliations, as these would affect both communication patterns and self-monitoring behavior through distinct departmental functions and norms.

Outcome Variables

For the supplementary analysis, I use two variables for measures of individual outcomes—log bonus and network constraint. Log bonus represents a measure of reward for individual productivity. Only certain individuals, those for which productivity could be objectively assessed, were bonus eligible, and monthly bonuses were calculated according to a company-established formula linking specific indicators—for example, conversion of leads to revenue—to bonus. I logged this measure since the distribution of monthly bonuses was skewed. Network constraint was calculated using an individual's monthly communication network based on the full set of unique emails according to Burt's (2004) measure of structural holes. Constraint is higher when an individual's communication partners also communicate with one another. I utilized the full set of emails, as this provided the most complete picture of brokerage dynamics within the firm.

Analytical Strategy

I present my findings in two general sections. In the main results, to estimate the effect of context on self-monitoring behavior, I use individual fixed-effect regression models, or within-person models. In these models, time-invariant individual effects that could confound the relationship between context and self-monitoring behavior are partitioned out. If, for example, there are certain stable factors that lead individuals to partake in certain communication contexts as well as display a certain level of self-monitoring behavior, then these would be adequately controlled for by leveraging the within-person variance. In the supplementary results, I again use individual fixed-effect regression models to examine the downstream effects of self-monitoring behavior.

However, though the within-person fixed-effect models take into account stable individual endogenous factors that may affect the relationship between context and self-monitoring

behavior, I can not fully rule out other potential confounding factors. In particular for my study, individuals do not communicate with interlocutors at random, and one's interaction partners are strongly influenced by both self-selection and a number of organizational forces (Biancani, McFarland, and Dahlander, 2014). As a result, there may be additional unobserved factors correlated with one's local communication contexts that affect self-monitoring behavior. Therefore, at this point, I do not make a causal claim to my results, and suggest that future work take advantage of more exogenous sources of variation in one's communication context, such as a pseudo-exogenous restructuring or a pseudo-random shock to the organization.

3.3 Main Results—How Context Affects Self-Monitoring Behavior

In my main results, I test my three hypotheses, that different contexts will systematically affect self-monitoring behavior. The correlation table between the relevant continuous variables is presented in Table 3.1 below.

Table 3.1: Correlation Table

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
(1) Self-Monitoring Behavior	1	0.190	-0.057	-0.218	-0.039	0.372	0.022	-0.104
(2) Alter Variance	0.190	1	-0.007	-0.069	-0.026	0.138	-0.030	-0.083
(3) % Comm. to New Inter.	-0.057	-0.007	1	0.138	0.425	-0.365	0.024	0.150
(4) % Comm. to Same Dept.	-0.218	-0.069	0.138	1	0.010	-0.128	0.039	0.027
(5) % Comm. to More Senior Inter.	-0.039	-0.026	0.425	0.010	1	-0.277	0.047	0.285
(6) Log Emails	0.372	0.138	-0.365	-0.128	-0.277	1	-0.089	-0.368
(7) Log Bonus	0.022	-0.030	0.024	0.039	0.047	-0.089	1	0.221
(8) Network Constraint	-0.104	-0.083	0.150	0.027	0.285	-0.368	0.221	1

Table 3.2 below shows the results of the employee fixed-effects regression models.

Table 3.2: Individual Fixed-Effects OLS Regression of Contexts on Self-Monitoring Behavior

	Dependent variable:				
	Self-Monitoring Behavior				
	(1)	(2)	(3)	(4)	(5)
% Comm. to New Inter.		0.070***			
		(0.007)			
% Comm. to Same Dept. Inter.			−0.052***		−0.048***
			(0.006)		(0.006)
% Comm. to More Senior Inter.				0.064***	0.058***
				(0.007)	(0.007)
Alter Variance	0.105***	0.103***	0.115***	0.119***	0.115***
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
First Year	0.399***	0.382***	0.419***	0.372***	0.373***
	(0.012)	(0.012)	(0.013)	(0.014)	(0.014)
Manager	−0.283***	−0.279***	−0.339***	−0.300***	−0.326***
	(0.020)	(0.020)	(0.022)	(0.022)	(0.022)
Log Emails	0.416***	0.422***	0.420***	0.419***	0.417***
	(0.009)	(0.009)	(0.010)	(0.010)	(0.010)
Department FE	YES	YES	YES	YES	YES
Employee FE	YES	YES	YES	YES	YES
Observations	21,535	21,535	19,224	19,473	19,224
R^2	0.582	0.584	0.593	0.590	0.595
Adjusted R^2	0.545	0.547	0.553	0.550	0.555
Note:					*p<0.05; **p<0.01; ***p<0.001

Model 1 demonstrates the effects of just the control variables on self-monitoring behavior; all appear to be in the expected directions. Greater opportunity for self-monitoring behavior in increased alter variance is positively associated with greater self-monitoring behavior. Employees exhibit greater self-monitoring behavior in their first year, likely as they are learning to adjust to the environment around them. Managers exhibit lower self-monitoring

behavior, likely as a result of their functions and role. And, employees exhibit more self-monitoring behavior the more emails they send, another form of the opportunity explanation.

Model 2 tests hypothesis 1, that the more employees communicate with new interlocutors, the more they will demonstrate self-monitoring behavior. The effect of communicating with more new interlocutors is positive and significant ($p < 0.001$), showing support in favor of hypothesis 1—one standard deviation increase in communication to new interlocutors results in a .07 standard deviation increase in self-monitoring behavior.

Model 3 tests hypothesis 2, that the more employees communicate with interlocutors in their departments, the less they will demonstrate self-monitoring behavior. The effect is negative and significant ($p < 0.001$), showing support in favor of hypothesis 2.

Model 4 tests hypothesis 3, that the more employees communicate with more senior interlocutors, the more they will display self-monitoring behavior. The effect is positive and significant ($p < 0.001$), showing support in favor of hypothesis 3.

Finally, model 5 demonstrates that even when considered together, the independent effects supporting hypothesis 2 and 3 still remain significant ($p < 0.001$) and in the expected directions.

3.4 Supplementary Results—The Outcomes of Self-Monitoring Behavior

In the supplementary analysis, I examine the effect of self-monitoring behavior on two individual outcomes, log bonus and network constraint. Again, the results of employee fixed-effects regression models are in Table 3.3 below.

Table 3.3: Individual Fixed-Effects OLS Regression of Self-Monitoring Behavior on Outcomes

	Dependent variable:	
	Log Bonus	Network Constraint
	(1)	(2)
Self-Monitoring Behavior	0.321^{***}	-0.022^{***}
	(0.085)	(0.003)
Alter Variance	-0.060^{**}	-0.005^{*}
	(0.023)	(0.003)
First Year	-0.664^{***}	0.087^{***}
	(0.069)	(0.006)
Manager	-0.035	-0.143^{***}
	(0.128)	(0.010)
Log Emails	0.136^{**}	
	(0.045)	
Department FE	YES	YES
Employee FE	YES	YES
Observations	14,112	21,535
R^2	0.462	0.675
Adjusted R^2	0.395	0.646
Note:	*p<0.05; **p<0.01; ***p<0.001	

Model 1 shows that self-monitoring behavior is positively associated with log bonus ($p < 0.001$). Results elsewhere have demonstrated that self-monitoring is more highly correlated with subjective measures of job performance (e.g. ratings) than for objective measures (e.g. sales volume, number of promotions), but it could not be ruled out that the subjective and objective measures were unequally reliable (Day et al., 2002). The results from model 1 provide further evidence that self-monitoring behavior is associated with objective measures as well. It may be that for those roles for which bonus is eligible, self-monitoring behavior in communication may lead to advantageous outcomes (e.g. converting leads to sales in the sales department).

Model 2 shows that self-monitoring behavior is negatively associated with network constraint, or in other words, positively associated with brokerage ($p < 0.001$). Although it

is unlikely for self-monitoring behavior to determine the recipients of a certain communication (in fact, earlier results would suggest that the latter influences the former), it may be that greater degrees of self-monitoring behavior help maintain the weaker ties that enable a brokerage network, or that individuals display more self-monitoring behavior when communicating with more disconnected individuals.

3.5 Discussion

In this chapter, I have sought to explain how identifying self-monitoring behavior can provide a lens into the inner machinations of an organization and to demonstrate the utility of the linguistic measure of self-monitoring in picking up that behavior. By understanding the systematic effect various contexts may have on self-monitoring behavior, we can move from a more reactive form of organizational management, based off of the observable behaviors of employees, to a more active form of organizational management that considers and is more sensitive to the subtleties of organizational life. By pursuing this line of research, we can begin to consider how the local contexts we create, whether as co-workers, managers, or subordinates impact those around us in ways previously understudied.

In this particular study, I have theorized that email communication creates meaningful local contexts, and that these local contexts, depending on the recipients of the emails, have certain characteristics that systematically affect the self-monitoring behavior of senders. I have demonstrated through empirical analysis that, when uncertainty with respect to appropriate behaviors is greater, as when communicating with new interlocutors, individuals display more self-monitoring behavior. And, when consequences for misperformance are less or more severe, as when communicating with interlocutors in the same department and more senior interlocutors, respectively, individuals display less and more self-monitoring behavior. As a whole, these results suggest patterns of interactions between individuals that are predictable and consequential in terms of individual earnings (bonus) and network position (brokerage).

These results draw attention to two important considerations that merit further investigation. First, there is the consideration of the importance of context. Not only does each organization have comparable, generalizable contexts—for example, in most organizations employees go through a hiring stage, a socialization stage, engage with an organizational culture, report to a manager—but also those contexts may be more meaningful and influential for an employee than any other aspect of the organization as a whole. Despite a company's ethical code or culture of fair treatment, a prejudiced manager can make his or her work team a miserable place to work.

Second, there is the consideration of the nature of self-monitoring behavior. Whereas previous conceptions of self-monitoring have sought to differentiate individuals on a stable trait, the picture presented here is more nuanced. The basic demonstration is that all individuals can self-monitor to some degree, but their self-monitoring behavior will be influenced by the contexts in which they find themselves. What I have not considered in great depth, however,

are the different motivations and forms of self-monitoring behavior (please see Appendix A for a brief discussion). Self-monitoring behavior, theoretically, can range from deception to sensitivity, which can have varying results on the actors' audience. In light of the need for further study, however, I believe that this notion of self-monitoring behavior can make several contributions to related fields in organizational studies.

Hiring

Undoubtedly, the hiring stage into organizations opens up a great number of possibilities for self-presentational strategies (Rivera, 2012; Kang et al., 2016), as the situational context involves recruiters who are trying to hire the best fit for their organization and potential employees who are trying to present themselves as such. However, while many of the strategies on the supply-side have been considered (e.g. whitening resumes, signaling cultural fit), far fewer studies have considered the demand-side implications of such strategies. The findings from this study on how context may affect self-monitoring behavior, in turn, may provide one perspective on these social processes.

For example, often the hiring process is a highly constrained context, in the sense that most potential employees will realize the necessity of acting professional. Yet, a distribution of professional resumes, photos, interviews, etc. obscures the more natural behavior of potential employees, which is more prone to come out during different organizational contexts, such as the company social. An understanding of how different contexts, or more specifically, the situational cues in those contexts, effect more or less self-monitoring behavior may help recruiters select more appropriately according to their organization's goals. For example, an individual whose resume is relatively unstructured, contains very simple sentences, and lacks professional luster but whose LinkedIn profile picture is pristine may indicate a high degree of self-monitoring behavior. Whether a high degree of self-monitoring behavior in the hiring stage indicates an employee who is willing to put in effort or an employee who presents well but does not have the substance to back it up is a matter of further study.

A better understanding of self-monitoring behavior and how contexts affect self-monitoring behavior then offers potential to improve hiring practices. While measuring self-monitoring behavior linguistically may not be feasible in every situation, recruiters could possibly assess self-monitoring behavior by constructing different contexts in their hiring practice whose effects on self-monitoring behavior is relatively known—recruiters could include one-on-one interviews with variously tenured employees, informal lunch or coffee meetings, or a mock business meeting in the interview schedule, effectively allowing recruiters to assess more precisely what a potential employee's natural behavior might be and the specific manifestations of his or her self-monitoring behavior.

Organizational Culture

There are two primary ways in which this study might contribute to the field of organizational culture. The first is understanding how organizational culture may affect self-monitoring be-

havior in employees, which may in turn affect downstream organizational outcomes. Though I have studied context as the local communication context between two interlocutors in this study and suggested why it may be particularly salient and influential for the behavior I have captured, I do not discount the importance of more overarching contexts, including subcultures and the organizational culture as a whole. Characteristics of these overarching contexts may have similarly systematic effects on self-monitoring behavior. For example, one stream of organizational culture research considers the difference between tight and loose cultures (Gelfand et al., 2011), where tight cultures represent those with stricter constraints such as strong norms and less tolerance for deviant behavior. In this formulation, tight cultures may foster more behavioral similarity, but could also increase the overall self-monitoring behavior of employees. Whether this benefits the organization in the short-term or long-term is yet to be studied.

The second way is understanding how self-monitoring behavior affects organizational cultures. Given that cultures can exist at different levels of analysis (Sackmann, 1992; Schein, 2010), it may be that a culture ostensibly exists because of the self-monitored behaviors of employees, but on a deeper level the culture is far more fragmented. In one conception, for example, employees may be able to assess the culture of the organization and fit in behaviorally, but personally hold differing values (Lu et al., 2019). Would two cultures of similar cultural consensus, or agreement about the prevailing norms, but differing levels of self-monitoring behavior among its employees offer the same benefits to organizations (O'Reilly and Chatman, 1996; Sørensen, 2002; Chatman et al., 2014)? One way to assess this may be to use my linguistic measure of self-monitoring behavior as a moderating factor in relating organizational cultures to broader organizational outcomes.

Limitations

While I believe this study improves our understanding of the construct of self-monitoring, there are still at least two limitations. First, I assume that there are certain cues present in the operationalized situations, and that those cues contribute to the theorized salient characteristics. For example, in communicating with new interlocutors, it is likely that the messages received from those new interlocutors contain cues relating to uncertainty, such as an ellipse which may communicate irritation, confusion, or simply a personal idiosyncrasy. But, I do not actually identify and track these specific cues, instead merely assuming that the more one communicates with new interlocutors, the more those cues would be present. It may be that communicating with some new interlocutors involves a great deal of uncertainty, while communicating with others does not. I believe that this is a reasonable assumption and that departures would be due to random noise, but future studies may be more precise in validating the salient characteristics of relevant situations.

Second, the self-monitoring behavior captured here is only one specific manifestation of self-monitoring behavior, that of one individual in the local context of communicating with another individual over a technology-mediated platform. One can imagine other forms of self-monitoring behavior, for example, face-to-face self-monitoring behavior involving control

of expressive behaviors, self-monitoring behavior in group settings, etc. These other forms, while related, and still in response to situational cues, may be more complicated or be affected in alternate ways and by alternate mechanisms. While this does not negate the individual, email-based self-monitoring behavior captured here, it does open up avenues for further investigation.

Future Directions

These limitations notwithstanding, in addition to the broader implications for other fields of study, two promising avenues for future research are examining the effects of self-monitoring behavior in different situations, and identifying different potential kinds of self-monitoring behavior. The first avenue would involve investigating more of the outcomes of self-monitoring behavior. This measure of self-monitoring behavior provides a framework not only for future extensions of the measure, but also for considering self-monitoring behavior from data other than email corpuses. For example, perhaps by using press releases and other linguistic documents, one can assess the self-monitoring behavior of CEOs of companies. Then, one could relate the variance there to various outcomes of the companies. Or, one could consider the self-monitoring behavior present in transcripts of analyst calls or political meetings, two scenarios in which self-monitoring behavior might be expected to be high, and link that with important outcomes in those settings.

The second avenue would involve investigating more deeply the potential types of self-monitoring. Essentially from the conception of the self-monitoring construct and after empirical evidence had been produced, the unitary structure of self-monitoring had been in question. Perhaps there are clearer demarcations of self-monitoring that could be teased apart by linguistic measures of self-monitoring behavior. For example, self-presentational habits in an effort to be sensitive to others and strengthen relationships may be starkly different from those in an effort to deceive others and superficially maintain relationships. Likely, those two self-monitoring behaviors would have differential outcomes on an individual's network structure, even if perceptually, others may not be able to tell a difference. Perhaps through some of the extensions mentioned in Chapter 2, specifically with regards to using the technical apparatus to identify clearer and cleaner pictures of individual behaviors, we can start to identify and quantify the effects of different self-monitoring behaviors.

[1] Gordon Willard Allport. *Personality: A psychological interpretation.* Holt, 1937.

[2] Samuel B Bacharach and Peter A Bamberger. "9/11 and New York City firefighters' post hoc unit support and control climates: A context theory of the consequences of involvement in traumatic work-related events". In: *Academy of Management Journal* 50.4 (2007), pp. 849–868.

[3] Martin Bäckström and Fredrik Björklund. "Social desirability in personality inventories: Symptoms, diagnosis and prescribed cure". In: *Scandinavian journal of psychology* 54.2 (2013), pp. 152–159.

[4] James C Beaty Jr, Jeanette N Cleveland, and Kevin R Murphy. "The relation between personality and contextual performance in" strong" versus" weak" situations". In: *Human Performance* 14.2 (2001), pp. 125–148.

[5] Daryl J Bem and Andrea Allen. "On predicting some of the people some of the time: The search for cross-situational consistencies in behavior." In: *Psychological review* 81.6 (1974), p. 506.

[6] Veronica Benet-Martinez and Oliver P John. "Los Cinco Grandes across cultures and ethnic groups: Multitrait-multimethod analyses of the Big Five in Spanish and English." In: *Journal of personality and social psychology* 75.3 (1998), p. 729.

[7] Susan Biancani, Daniel A McFarland, and Linus Dahlander. "The semiformal organization". In: *Organization Science* 25.5 (2014), pp. 1306–1324.

[8] Wiebke Bleidorn and Christopher James Hopwood. "Using machine learning to advance personality assessment and theory". In: *Personality and Social Psychology Review* (2018), p. 1088868318772990.

[9] Jack Block. "Critique of the act frequency approach to personality." In: *Journal of Personality and Social Psychology* 56.2 (1989), p. 234.

[10] Jack Block. *The Q-sort method in personality assessment and psychiatric research.* Vol. 457. Citeseer, 1961.

[11] Hannah Riley Bowles, Linda Babcock, and Kathleen L McGinn. "Constraints and triggers: situational mechanics of gender in negotiation." In: *Journal of personality and social psychology* 89.6 (2005), p. 951.

[12] Stephen R Briggs and Jonathan M Cheek. "On the nature of self-monitoring: Problems with assessment, problems with validity." In: *Journal of Personality and Social Psychology* 54.4 (1988), p. 663.

[13] Stephen R Briggs, Jonathan M Cheek, and Arnold H Buss. "An analysis of the self-monitoring scale." In: *Journal of Personality and Social Psychology* 38.4 (1980), p. 679.

[14] Ronald S Burt. "Structural holes and good ideas". In: *American journal of sociology* 110.2 (2004), pp. 349–399.

[15] David M Buss and Kenneth H Craik. "The act frequency approach to personality." In: *Psychological Review* 90.2 (1983), p. 105.

[16] John T Cacioppo and Richard E Petty. "The need for cognition." In: *Journal of personality and social psychology* 42.1 (1982), p. 116.

[17] Raymond B Cattell. "Confirmation and clarification of primary personality factors". In: *Psychometrika* 12.3 (1947), pp. 197–220.

[18] Raymond B Cattell. "Primary personality factors in the realm of objective tests". In: *Journal of Personality* 16.4 (1948), pp. 459–486.

[19] Raymond B Cattell. "The description of personality: Basic traits resolved into clusters." In: *The journal of abnormal and social psychology* 38.4 (1943), p. 476.

[20] Raymond B Cattell and Heather E P. Cattell. "Personality structure and the new fifth edition of the 16PF". In: *Educational and Psychological Measurement* 55.6 (1995), pp. 926–937.

[21] Raymond Bernard Cattell. *Description and measurement of personality.* World Book Company, 1946.

[22] Jennifer A Chatman et al. "Parsing organizational culture: How the norm for adaptability influences the relationship between culture consensus and financial performance in high-technology firms". In: *Journal of Organizational Behavior* 35.6 (2014), pp. 785–808.

[23] William H Cooper and Michael J Withey. "The strong situation hypothesis". In: *Personality and Social Psychology Review* 13.1 (2009), pp. 62–72.

[24] Paul T Costa Jr and Robert R McCrae. "Neo personality inventory–revised (neo-pi-r) and neo five-factor inventory (neo-ffi) professional manual". In: *Odessa, FL: Psychological Assessment Resources* (1992).

[25] Lee J Cronbach. "Coefficient alpha and the internal structure of tests". In: *psychometrika* 16.3 (1951), pp. 297–334.

[26] David V Day et al. "Self-monitoring personality at work: A meta-analytic investigation of construct validity." In: *Journal of Applied Psychology* 87.2 (2002), p. 390.

[27] Colin G DeYoung, Lena C Quilty, and Jordan B Peterson. "Between facets and domains: 10 aspects of the Big Five." In: *Journal of personality and social psychology* 93.5 (2007), p. 880.

[28] John M Digman. "Personality structure: Emergence of the five-factor model". In: *Annual review of psychology* 41.1 (1990), pp. 417–440.

[29] Stewart I Donaldson and Elisa J Grant-Vallone. "Understanding self-report bias in organizational behavior research". In: *Journal of business and Psychology* 17.2 (2002), pp. 245–260.

[30] Gabriel Doyle and Michael C Frank. "Investigating the sources of linguistic alignment in conversation". In: *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Vol. 1. 2016, pp. 526–536.

[31] Gabriel Doyle et al. "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)*. Vol. 1. 2017, pp. 603–612.

[32] George John Dudycha. *An objective study of punctuality in relation to personality and achievement*. 204. Columbia university, 1936.

[33] Nicholas Duran, Alexandra Paxton, and Riccardo Fusaroli. *ALIGN: Analyzing Linguistic Interactions with Generalizable techNiques - a Python Library*. 2018. DOI: 10.31234/osf.io/a5yh9. URL: psyarxiv.com/a5yh9.

[34] Amy Edmondson. "Psychological safety and learning behavior in work teams". In: *Administrative science quarterly* 44.2 (1999), pp. 350–383.

[35] John A Edwards and Angela Templeton. "The structure of perceived qualities of situations". In: *European journal of social psychology* 35.6 (2005), pp. 705–723.

[36] Paul Ekman and Wallace V Friesen. "Constants across cultures in the face and emotion." In: *Journal of personality and social psychology* 17.2 (1971), p. 124.

[37] Richard M Emerson. "Power-dependence relations". In: *American sociological review* (1962), pp. 31–41.

[38] Robert A Emmons and Ed Diener. "Personality correlates of subjective well-being". In: *Personality and Social Psychology Bulletin* 11.1 (1985), pp. 89–97.

[39] Seymour Epstein. "The stability of behavior: I. On predicting most of the people much of the time." In: *Journal of personality and social psychology* 37.7 (1979), p. 1097.

[40] Hans Jurgen Eysenck. "Dimensions of personality: 16, 5 or 3?Criteria for a taxonomic paradigm". In: *Personality and individual differences* 12.8 (1991), pp. 773–790.

[41] Ruolian Fang et al. "Integrating personality and social networks: A meta-analysis of personality, network position, and work outcomes in organizations". In: *Organization Science* 26.4 (2015), pp. 1243–1260.

[42] Ethan Fast, Binbin Chen, and Michael S Bernstein. "Empath: Understanding topic signals in large-scale text". In: *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM. 2016, pp. 4647–4657.

[43] Martin Fishbein and Icek Ajzen. "Attitudes towards objects as predictors of single and multiple behavioral criteria." In: *Psychological review* 81.1 (1974), p. 59.

[44] William Fleeson. "Moving personality beyond the person-situation debate: The challenge and the opportunity of within-person variability". In: *Current Directions in Psychological Science* 13.2 (2004), pp. 83–87.

[45] William Fleeson. "Toward a structure-and process-integrated view of personality: Traits as density distributions of states." In: *Journal of personality and social Psychology* 80.6 (2001), p. 1011.

[46] William Fleeson and Patrick Gallagher. "The implications of Big Five standing for the distribution of trait manifestation in behavior: Fifteen experience-sampling studies and a meta-analysis." In: *Journal of personality and social psychology* 97.6 (2009), p. 1097.

[47] William Fleeson and Mary Kate Law. "Trait enactments as density distributions: The role of actors, situations, and observers in explaining stability and variability." In: *Journal of personality and social psychology* 109.6 (2015), p. 1090.

[48] Marc A Fournier, DS Moskowitz, and David C Zuroff. "Integrating dispositions, signatures, and the interpersonal domain." In: *Journal of personality and social psychology* 94.3 (2008), p. 531.

[49] Norman Frederiksen. "Toward a taxonomy of situations." In: *American Psychologist* 27.2 (1972), p. 114.

[50] Paul T Fuglestad and Mark Snyder. "Status and the Motivational Foundations of Self-Monitoring". In: *Social and Personality Psychology Compass* 4.11 (2010), pp. 1031–1041.

[51] David C Funder. "Global traits: A neo-Allportian approach to personality". In: *Psychological Science* 2.1 (1991), pp. 31–39.

[52] David C Funder and C Randall Colvin. "Friends and strangers: Acquaintanceship, agreement, and the accuracy of personality judgment." In: *Journal of personality and social psychology* 55.1 (1988), p. 149.

[53] R Michael Furr. "Personality psychology as a truly behavioural science". In: *European Journal of Personality: Published for the European Association of Personality Psychology* 23.5 (2009), pp. 369–401.

[54] Steve Gangestad and Mark Snyder. ""To carve nature at its joints": On the existence of discrete classes in personality." In: *Psychological Review* 92.3 (1985), p. 317.

[55] Steven W Gangestad and Mark Snyder. "Self-monitoring: Appraisal and reappraisal." In: *Psychological bulletin* 126.4 (2000), p. 530.

[56] Steven W Gangestad and Mark Snyder. "Taxonomic analysis redux: Some statistical considerations for testing a latent class model." In: (1991).

[57] Michele J Gelfand et al. "Differences between tight and loose cultures: A 33-nation study". In: *science* 332.6033 (2011), pp. 1100–1104.

[58] Erving Goffman. "The Presentation of Self in". In: *Butler, Bodies that Matter* (1959).

[59] Lewis R Goldberg. "An alternative" description of personality": the big-five factor structure." In: *Journal of personality and social psychology* 59.6 (1990), p. 1216.

[60] Lewis R Goldberg. "The development of markers for the Big-Five factor structure." In: *Psychological assessment* 4.1 (1992), p. 26.

[61] Lewis R Goldberg et al. "The international personality item pool and the future of public-domain personality measures". In: *Journal of Research in personality* 40.1 (2006), pp. 84–96.

[62] Samuel D Gosling, Peter J Rentfrow, and William B Swann Jr. "A very brief measure of the Big-Five personality domains". In: *Journal of Research in personality* 37.6 (2003), pp. 504–528.

[63] Samuel D Gosling et al. "A room with a cue: Personality judgments based on offices and bedrooms." In: *Journal of personality and social psychology* 82.3 (2002), p. 379.

[64] Mark Granovetter. "Economic action and social structure: The problem of embeddedness". In: *American journal of sociology* 91.3 (1985), pp. 481–510.

[65] Fangjian Guo et al. "The bayesian echo chamber: Modeling social influence via linguistic accommodation". In: *Artificial Intelligence and Statistics*. 2015, pp. 315–323.

[66] Gerald J Haeffel and George S Howard. "Self-report: Psychologys four-letter word". In: *American Journal of Psychology* 123.2 (2010), pp. 181–188.

[67] Molly E Ireland and James W Pennebaker. "Language style matching in writing: Synchrony in essays, correspondence, and poetry." In: *Journal of personality and social psychology* 99.3 (2010), p. 549.

[68] Oliver P John, Jonathan M Cheek, and Eva C Klohnen. "On the nature of self-monitoring: Construct explication with Q-sort ratings." In: *Journal of Personality and Social Psychology* 71.4 (1996), p. 763.

[69] Oliver P John, Eileen M Donahue, and Robert L Kentle. "The big five inventory–Versions 4a and 54. Berkeley". In: *CA: University of California, Berkeley, Institute of Personality and Social Research* 54 (1991).

[70] Oliver P John, Laura P Naumann, and Christopher J Soto. "Paradigm shift to the integrative big five trait taxonomy". In: *Handbook of personality: Theory and research* 3.2 (2008), pp. 114–158.

[71] Sonia K Kang et al. "Whitened resumes: Race and self-presentation in the labor market". In: *Administrative Science Quarterly* 61.3 (2016), pp. 469–502.

[72] David A Kenny et al. "Consensus in interpersonal perception: Acquaintance and the big five". In: *Psychological bulletin* 116.2 (1994), p. 245.

[73] Martin Kilduff and David V Day. "Do chameleons get ahead? The effects of self-monitoring on managerial careers". In: *Academy of Management Journal* 37.4 (1994), pp. 1047–1060.

[74] Blaine Landis. "Personality and social networks in organizations: A review and future directions". In: *Journal of Organizational Behavior* 37 (2016), S107–S121.

[75] Justin A Lavner et al. "Personality change among newlyweds: Patterns, predictors, and associations with marital satisfaction over time". In: *Developmental Psychology* 54.6 (2018), pp. 1172–1185.

[76] Shelley Leininger and Reid Skeel. "Cortisol and self-report measures of anxiety as predictors of neuropsychological performance". In: *Archives of Clinical Neuropsychology* 27.3 (2012), pp. 318–328.

[77] Richard D Lennox and Raymond N Wolfe. "Revision of the self-monitoring scale." In: (1984).

[78] Hanna Levenson. "Differentiating among internality, powerful others, and chance". In: *Research with the Locus of Control Construct*. Ed. by Herbert M. Lefcourt. New York: Academic Press, 1981, pp. 15–63.

[79] Richard W. Lu et al. "Deciphering the cultural code: Cognition, behavior, and the interpersonal transmission of culture". In: *Working Manuscript* (2019).

[80] David C McClelland, Richard Koestner, and Joel Weinberger. "How do self-attributed and implicit motives differ?" In: *Psychological review* 96.4 (1989), p. 690.

[81] Ajay Mehra, Martin Kilduff, and Daniel J Brass. "The social networks of high and low self-monitors: Implications for workplace performance". In: *Administrative science quarterly* 46.1 (2001), pp. 121–146.

[82] Rustin D Meyer. *Taxonomy of Situations and Their Measurement*. 2015.

[83] Rustin D Meyer, Reeshad S Dalal, and Richard Hermida. "A review and synthesis of situational strength in the organizational sciences". In: *Journal of Management* 36.1 (2010), pp. 121–140.

[84] Walter Mischel. "Continuity and change in personality." In: *American psychologist* 24.11 (1969), p. 1012.

[85] Walter Mischel. "From personality and assessment (1968) to personality science, 2009". In: *Journal of Research in Personality* 43.2 (2009), pp. 282–290.

[86] Walter Mischel. *Personality and assessment*. Psychology Press, 1968.

[87] Walter Mischel. "The interaction of person and situation". In: *Personality at the crossroads: Current issues in interactional psychology* 333 (1977), p. 352.

[88] Walter Mischel. "Toward a cognitive social learning reconceptualization of personality." In: *Psychological review* 80.4 (1973), p. 252.

[89] Walter Mischel and Yuichi Shoda. "A cognitive-affective system theory of personality: reconceptualizing situations, dispositions, dynamics, and invariance in personality structure." In: *Psychological review* 102.2 (1995), p. 246.

[90] Sanaz Mobasseri. "Gender and the Give and Take of Emotions in the Workplace". In: *Working Manuscript* (2019).

[91] Debbie S Moskowitz. "Coherence and cross-situational generality in personality: A new analysis of old problems." In: *Journal of Personality and Social Psychology* 43.4 (1982), p. 754.

[92] DS Moskowitz and David C Zuroff. "Flux, pulse, and spin: Dynamic additions to the personality lexicon." In: *Journal of personality and social psychology* 86.6 (2004), p. 880.

[93] Marcus Mund and Franz J Neyer. "Treating personality-relationship transactions with respect: Narrow facets, advanced models, and extended time frames". In: *Journal of Personality and Social Psychology* 107.2 (2014), p. 352.

[94] Theodore Mead Newcomb. "The consistency of certain extrovert-introvert behavior patterns in 51 problem boys." In: *Teachers College Contributions to Education* (1929).

[95] Warren T Norman. "Toward an adequate taxonomy of personality attributes: Replicated factor structure in peer nomination personality ratings." In: *The Journal of Abnormal and Social Psychology* 66.6 (1963), p. 574.

[96] C.A. O'Reilly and J.A. Chatman. "Culture as social control: Corporations, cults, and commitment". In: *Research in organizational behavior: An annual series of analytical essays and critical reviews*. Ed. by B.M. Staw and L.L. Cummings. Vol. 18. Elsevier Science/JAI Press, 1996, pp. 157–200.

[97] John F Padgett and Christopher K Ansell. "Robust Action and the Rise of the Medici, 1400-1434". In: *American journal of sociology* 98.6 (1993), pp. 1259–1319.

[98] James W Pennebaker and Laura A King. "Linguistic styles: Language use as an individual difference." In: *Journal of personality and social psychology* 77.6 (1999), p. 1296.

[99] James W Pennebaker et al. *The development and psychometric properties of LIWC2015*. Tech. rep. 2015.

[100] Lawrence A Pervin. "A critical analysis of current trait theory". In: *Psychological Inquiry* 5.2 (1994), pp. 103–113.

[101] *Publications that Employ the IPIP*. `https://ipip.ori.org/newPublications.htm`. Accessed: 2018-12-07.

[102] John F Rauthmann, Ryne A Sherman, and David C Funder. "Principles of situation research: Towards a better understanding of psychological situations". In: *European Journal of Personality* 29.3 (2015), pp. 363–381.

[103] John F Rauthmann et al. "The Situational Eight DIAMONDS: A taxonomy of major dimensions of situation characteristics." In: *Journal of Personality and Social Psychology* 107.4 (2014), p. 677.

[104] Rebecca E Ready et al. "Self-and peer-reported personality: Agreement, trait ratability, and the self-based heuristic". In: *Journal of Research in Personality* 34.2 (2000), pp. 208–224.

[105] Peter J Rentfrow and Samuel D Gosling. "The do re mi's of everyday life: the structure and personality correlates of music preferences." In: *Journal of personality and social psychology* 84.6 (2003), p. 1236.

[106] Lauren A Rivera. "Hiring as cultural matching: The case of elite professional service firms". In: *American sociological review* 77.6 (2012), pp. 999–1022.

[107] Brent W Roberts, Avshalom Caspi, and Terrie E Moffitt. "The kids are alright: growth and stability in personality development from adolescence to adulthood". In: *Journal of personality and social psychology* 81.4 (2001), p. 670.

[108] Brent W Roberts et al. "The power of personality: The comparative validity of personality traits, socioeconomic status, and cognitive ability for predicting important life outcomes". In: *Perspectives on Psychological science* 2.4 (2007), pp. 313–345.

[109] Sonja A Sackmann. "Culture and subcultures: An analysis of organizational knowledge." In: *Administrative science quarterly* 37.1 (1992).

[110] Zuzana Sasovova et al. "Network churn: The effects of self-monitoring personality on brokerage dynamics". In: *Administrative Science Quarterly* 55.4 (2010), pp. 639–670.

[111] Gerard Saucier. "Orthogonal markers for orthogonal factors: The case of the Big Five". In: *Journal of Research in Personality* 36.1 (2002), pp. 1–31.

[112] Edgar H Schein. *Organizational culture and leadership*. Vol. 2. John Wiley & Sons, 2010.

[113] Mark Snyder. "Self-monitoring of expressive behavior." In: *Journal of personality and social psychology* 30.4 (1974), p. 526.

[114] Mark Snyder. "Self-monitoring processes". In: *Advances in experimental social psychology* 12 (1979), pp. 85–128.

[115] Mark Snyder and Steve Gangestad. "On the nature of self-monitoring: Matters of assessment, matters of validity." In: *Journal of personality and social psychology* 51.1 (1986), p. 125.

[116] Mark Snyder and William Ickes. "Personality and social behavior". In: *Handbook of social psychology* 2.3 (1985), pp. 883–947.

[117] Jesper B Sørensen. "The strength of corporate culture and the reliability of firm performance". In: *Administrative science quarterly* 47.1 (2002), pp. 70–91.

[118] Sameer B Srivastava et al. "Enculturation trajectories: Language, cultural adaptation, and individual outcomes in organizations". In: *Management Science* 64.3 (2017), pp. 1348–1364.

[119] William Stern. *Differentielle Psychologie*. 1921.

[120] Stefano Tasselli and Martin Kilduff. "When brokerage between friendship cliques endangers trust: a personality–network fit perspective". In: *Academy of Management Journal* 61.3 (2018), pp. 802–825.

[121] Ernest C Tupes and Raymond E Christal. "Recurrent personality factors based on trait ratings". In: *United States Air Force Aeronautical Systems Division, Technical Report* 61-97 (1961).

[122] Michael P Wilmot. "A contemporary taxometric analysis of the latent structure of self-monitoring." In: *Psychological assessment* 27.2 (2015), p. 353.

[123] Michael P Wilmot et al. "Self-monitoring and the metatraits". In: *Journal of personality* 84.3 (2016), pp. 335–347.

[124] Michael J Withey, Ian R Gellatly, and Michael Annett. "The Moderating Effect of Situation Strength on the Relationship Between Personality and Provision of Effort 1". In: *Journal of Applied Social Psychology* 35.8 (2005), pp. 1587–1606.

[125] Wu Youyou, Michal Kosinski, and David Stillwell. "Computer-based personality judgments are more accurate than those made by humans". In: *Proceedings of the National Academy of Sciences* 112.4 (2015), pp. 1036–1040.

[126] Rong Zheng et al. "A framework for authorship identification of online messages: Writing-style features and classification techniques". In: *Journal of the American society for information science and technology* 57.3 (2006), pp. 378–393.

Appendix A: 85- and 64-LIWC Self-Monitoring Behavior

In Chapter 2, I made a design decision to remove twenty-one categories (see Table A1 below for those removed categories) from the 2015 version of LIWC to construct my main measure of self-monitoring behavior, according to findings that those specific categories may be less susceptible to intentional choices (Ireland and Pennebaker, 2010). In this appendix, I revisit that decision, replicating my main and supplementary results with the full 85-LIWC category self-monitoring behavior measure (85-LIWC), and discuss the suggestive implications of the differences. Before comparing the results from the two measures, it will be review the reasoning and clarify what I believe is the major distinction between the 85-LIWC self-monitoring behavior measure and the 64-LIWC measure.

Table A1: 21 Additional Categories in 85-LIWC

Linguistic Processes		*Other Grammar*	
LIWC Category	Full Name	LIWC Category	Full Name
function	Total function words	verb	Common verbs
pronoun	Total pronouns	adj	Common adjectives
ppron	Personal Pronouns	compare	Comparisons
i	1st pers singular	interrog	Interrogatives
we	1st pers plural	number	Number
you	2nd person	quant	Quantifiers
shehe	3rd pers singular		
they	3rd pers plural		
ipron	Impersonal pronouns		
article	Articles		
prep	Prepositions		
auxverb	Auxiliary verbs		
adverb	Adverbs		
conj	Conjunctions		
negate	Negations		

In the formulation of LIWC, different word categories are divided by overarching headers. For example, positive emotion and negative emotion fall under the psychological processes header, and articles and prepositions fall under the linguistic processes header. There are differing levels of such headers in a hierarchy of categories, but ultimately the categories I excluded fell under "linguistic processes" and "other grammar". In their article considering language style matching, Ireland and Pennebaker (2010) were particularly interested in matching on what they considered *style*, which comprised nine nonoverlapping function word categories—personal pronouns, impersonal pronouns, articles, auxiliary verbs, high-

frequency adverbs, prepositions, conjunctions, negations, and quantifiers—in the 2007 formulation of LIWC. Because I had been using the slightly updated 2015 formulation of LIWC, I chose to exclude the categories I thought conceptually similar to theirs.

Their primary study to identify whether these such function word categories could be intentionally manipulated was to ask university students to read short excerpts from best-selling novels and then to complete the story in their own words. The treatment was a differing set of instructions, in which the control group was asked to simply finish the story and be realistic about what likely happened in the scene after the one they had just read, and the treatment group was asked to finish the story the way the author intended, trying to match the voice and style of the author. They found no statistically significant difference in the mean language style matching over the selected categories for the two groups, and concluded that the results suggested that language style matching could not be deliberately controlled.

While there are some reasonable critiques of the study, that perhaps the individuals were already seeking to match the author's style and that the mean levels of language style matching were high in both conditions (.77 out of 1 for treatment and .78 out of 1 for control), the implication is that matching on those categories is more automatic than intentional. Extending that implication with the understanding that my framework for self-monitoring behavior is a little different (for example, due to base rate usages of the function word categories it may be difficult to find cases where messages do not contain articles), it should be that the 85-LIWC measure involves more unintentional self-monitoring than does the 64-LIWC measure. It is from this perspective that I now turn to consider the differences in results.

Recall that the two measures are correlated at approximately .9270, and as presented below, the correlations with other variables (see Table A2 below) and the main findings remain substantively similar. However, there are two major differences worth noting.

Table A2: Correlation Table (85-LIWC)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
(1) Self-Monitoring Behavior	1	0.150	-0.017	-0.176	-0.027	0.301	0.145	-0.052
(2) Alter Variance	0.150	1	-0.002	-0.035	-0.033	0.157	-0.021	-0.121
(3) % Comm. to New Inter.	-0.017	-0.002	1	0.138	0.425	-0.365	0.024	0.150
(4) % Comm. to Same Dept.	-0.176	-0.035	0.138	1	0.010	-0.128	0.039	0.027
(5) % Comm. to More Senior Inter.	-0.027	-0.033	0.425	0.010	1	-0.277	0.047	0.285
(6) Log Emails	0.301	0.157	-0.365	-0.128	-0.277	1	-0.089	-0.368
(7) Log Bonus	0.145	-0.021	0.024	0.039	0.047	-0.089	1	0.221
(8) Network Constraint	-0.052	-0.121	0.150	0.027	0.285	-0.368	0.221	1

First, as shown in Table A3 below, in the test of hypothesis 3, that when communicating more with more senior interlocutors, individuals will display more self-monitoring behavior, the main effect is less significant ($p < 0.05$), and drops out of significance when percent

communication to interlocutors in the same department is included. I interpret this difference as in line with and supportive of my theorizing. Communicating across a status or power difference or with an air of respect may involve more active considerations that may get washed out when more unintentional processes are included.

Table A3: Individual Fixed-Effects OLS Regression of Contexts on Self-Monitoring Behavior (85-LIWC)

	Dependent variable:				
	Self-Monitoring Behavior				
	(1)	(2)	(3)	(4)	(5)
% Comm. to New Inter.		0.064^{***}			
		(0.008)			
% Comm. to Same Dept. Inter.			-0.050^{***}		-0.050^{***}
			(0.007)		(0.007)
% Comm. to More Senior Inter.				0.018^{*}	0.012
				(0.008)	(0.008)
Alter Variance	0.110^{***}	0.108^{***}	0.123^{***}	0.126^{***}	0.123^{***}
	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)
First Year	0.327^{***}	0.311^{***}	0.356^{***}	0.344^{***}	0.346^{***}
	(0.013)	(0.014)	(0.014)	(0.016)	(0.016)
Manager	-0.209^{***}	-0.205^{***}	-0.273^{***}	-0.243^{***}	-0.270^{***}
	(0.023)	(0.023)	(0.025)	(0.024)	(0.025)
Log Emails	0.347^{***}	0.353^{***}	0.355^{***}	0.357^{***}	0.354^{***}
	(0.010)	(0.010)	(0.011)	(0.011)	(0.011)
Observations	21,535	21,535	19,224	19,473	19,224
R^2	0.477	0.479	0.491	0.487	0.491
Adjusted R^2	0.431	0.432	0.441	0.437	0.441
Note:					*p<0.05; **p<0.01; ***p<0.001

Second, while the direction and significance of the supplementary results remain roughly the same, the effect of 85-LIWC self-monitoring behavior on log bonus is slightly heightened,

and the effect on network constraint is slightly lowered (and now significant at $p < 0.01$). More self-monitoring behavior on the function word categories may help facilitate coordination and clarity, which might in turn contribute to higher individual bonuses. And, the process of maintaining ties or speaking to individuals who themselves do not speak to one another may involve more active considerations.

Table A4: Individual Fixed-Effects OLS Regression of Self-Monitoring Behavior on Outcomes (85-LIWC)

	Dependent variable:	
	Log Bonus	Network Constraint
	(1)	(2)
Self-Monitoring Behavior	0.414^{***}	-0.009^{**}
	(0.055)	(0.003)
Alter Variance	-0.031	-0.010^{***}
	(0.023)	(0.003)
First Year	-0.535^{***}	0.079^{***}
	(0.070)	(0.006)
Manager	-0.140	-0.138^{***}
	(0.128)	(0.010)
Log Emails	0.144^{***}	
	(0.041)	
Observations	14,112	21,535
R^2	0.464	0.675
Adjusted R^2	0.397	0.646
Note:	*p<0.05; **p<0.01; ***p<0.001	

Overall, I believe the results highlight the importance of further theorizing on the nature of self-monitoring behavior. Some degree of self-monitoring behavior is necessary to facilitate social interaction, but the questions surrounding the degree to which individuals can and do actively monitor their self-presentation are interesting and I believe understudied.