Understanding Affect in the Workplace via Social Media

Munmun De Choudhury
Microsoft Research, Redmond WA 98052

Scott Counts
Microsoft Research, Redmond WA 98052

{munmund,counts}@microsoft.com

ABSTRACT
We investigate the landscape of affective expression of employees at a large Fortune 500 software corporation via an internal microblogging tool. We present three analyses of emotional expression among employees, based on literature in organizational behavior: its relationship to (1) exogenous/lifestyle and endogenous workplace factors, (2) geography and (3) organizational structure. We find that employees tend to make significant accommodations in affect expression when interacting with others over the organizational hierarchy. We also find that positive affect is expressed through interpersonal communications that connect disparate geographic regions. Our findings have implications for enabling emotional reflection of employees and for management in that they can help uncover emotional patterns associated with episodes of high and low productivity, allowing organizations to improve employee engagement and promote positive attitudes.

Author Keywords
affect, emotion, enterprise, microblog, social media

ACM Classification Keywords
H.5.m [Information Systems]: Information Systems Applications – Miscellaneous.

General Terms
Human Factors; Measurement.

INTRODUCTION
The 'affective climate' of an enterprise constitutes a valuable resource for its operational procedure and its organizational processes and outcomes. For instance, measures of affect of employees may help predict anxiety/ stress levels with respect to company policies, products, or organizational changes. Employment trends including more volatility and insecurity in many industries, more mobility and ability to work remotely, more expectation of or need for active employee participation in decision-making also make it increasingly important for organizations to assess employee affect and reaction to changes. The organization can thereby seek to identify the factors behind employees' joy and satisfaction, as well as their frustration towards difficulty and failure. Ideally such efforts to assess affect culminate in feedback cycles that support action plans for improved emotional counseling, workgroup dynamics, employee collaboration, and organizational structuring in the workplace environment. In short, affect analysis can help enterprises strive for positive change.

Recognizing the value, organizational behavior researchers have conducted numerous ethnographic studies in a variety of workplace types and settings, validating the influence of affect on work outcomes and workplace dynamics, such as task quality, productivity, creativity, group rapport, user focus, and job satisfaction [1,6,15,29,34,39,41]. For instance, a well-established effect is that creativity in problem-solving is heightened by positive affect and is diminished when a user is in a negative emotional state [1].

Beyond self-reports from relatively small samples of workers, employee affect has been difficult to measure and leverage. Managers at organizations are supposed to monitor it in their groups and feed the information forward; human resources personnel try to monitor it and devise awards and morale events to influence it, and perhaps bring in consultants to deal with problems in a specific group. Each of these methods poses unique challenges and many are decentralized (e.g., relying on managers, who are distributed and may vary widely in the amount of attention and sophistication they give to the affect monitoring).

Annual experience surveys, in which explicit anonymous responses are sought from employees who opt-in to respond, are a common method for getting a reading on company-wide affect in a standardized manner. Although these surveys often allow the employees to enter open-ended comments and feedback on their feelings, thoughts, and attitude towards the organization, they face an inherent limitation due to the large temporal gaps in which they are conducted. The cost, staff time and effort associated with the surveys make it difficult to administer them any more frequently. Consequently, the reflections in these surveys tend to be summary in nature, even possibly imprecise, rather than a record of actual data points collected in the flow of daily work. This means that they lose the context that would tie emotional expression to specific aspects of the daily work schedule. When colleagues are separated geographically, additional complications ensue.

CSCW '13, February 23–27, 2013, San Antonio, Texas, USA.

Some organizations today are starting to make sustained use of public, Twitter-like communication channels. Tools are being developed to monitor conversation channels and discern levels of engagement of speakers. In fact, the use of social media (e.g., Facebook, Twitter, Yammer) in enterprises has been shown to impact information sharing [7,10,12,44], collaboration, communication, and interpersonal connectivity [2,8,11, 21, 37] in the workplace. Such data sources can complement smaller sample ethnographic research while simultaneously overcoming the limitations of infrequent employee experience surveys by providing an ongoing record of employee affect.

Although creating individual profiles might alarm employees and alter their use of the channels, monitoring the collective stream of conversation, or conversations involving large groups of workers via these communication modalities, could provide a measure of how different groups interact, or how affect is being impacted by specific events or broad changes in the work processes or environment.

Analysis of employee social media posts, such as Twitter, is tempting given the volume, frequency of use, and easy availability of shared content from employees. However, many of these tools are essentially public, making it unclear how effective they are in capturing organization-specific sentiment of employees. Further given the constant eroding of work-life boundaries [2], employees are likely to use these tools for both personal and professional updates, and are unlikely to discuss company-related information in a public forum.

To this end, we take advantage of the voluntary and casual short updates posted on an *internal* Twitter-like microblogging tool within a large Fortune 500 software corporation based in the United States. Our goal is to characterize affective expression of employees at the workplace in a fine-grained, in-context, and longitudinal manner, and thereby complement traditional mechanisms, such as self-reports and annual experience surveys. The tool, called OfficeTalk was developed and launched in 2008 by a product group inside Microsoft to explore the value of microblogging in an enterprise context. It is hosted on the company's own servers, automatically restricting all content within the company, and thereby allowing sharing of internal information among users. OfficeTalk had been adopted by over 30,000 employees. This promising source of insight into the evolving affective climate in the company can thus be studied and used, while protecting individuals from exposure. The main contributions of this paper are as follows:

- A methodology to mine emotion of employees (Positive Affect or PA and Negative Affect or NA) at-scale, motivated by psycholinguistics literature.
- A series of studies, grounded by theory in organizational behavior research, that describe how affective expression in the enterprise is influenced by the following factors: (1) exogenous and endogenous workplace factors, (2) physical location/geography of organization, and (3) the organizational hierarchy.

Our analyses indicate three major findings. Employees were found to exhibit cyclic rhythms of positive and negative affect over the course of a day; with PA decreasing significantly afterhours, which could be due to employee dissatisfaction with the work-life imbalance of having to work in the evening. Next we observe that interpersonal interactions via the tool generate PA that transcends geographic and cultural homophily, connecting employees across disparate regions known to professionally collaborate. Finally, our results exhibit evidence of "affect accommodation", wherein employees with higher status in the organizational hierarchy were seen expressing increased PA and lowered NA in their conversations with their coworkers and subordinates, likely with the goal of energizing the latter with enthusiasm in order to improve task performance.

Together, our findings bear on the potential of development of systems that can enable assessment of employee affect in an anonymous and aggregated manner, including assessing performance-relevant outcomes, such as job satisfaction, judgments, attitudinal responses, creativity, helping behavior and risk-taking.

BACKGROUND: ENTERPRISE, SOCIAL MEDIA, AFFECT

As newer web-based modalities of communication have emerged, a number of studies began investigating their role in enterprises—including email [8], blogs [11,22], microblogs [12], and instant messaging [21], as well as less direct modes of interaction such as social bookmarking [26], tagging [27], and file sharing [32]. Recent studies have investigated motivations and usage of social networking/media tools in organizational settings. Broadly, these studies have either been about how publicly open tools such as Twitter and Facebook are adopted and used by employees in a workplace [2,10,33,37,43,44], or about user studies of prototypes developed inside an organization for promoting collaboration or networking (notably [7,12]). Prior work also examined and identified power users and their roles in enterprises (e.g., [37]), the traction between social and work benefits, predominantly in light of the impact of these tools on productivity (ref. [33]), as well as how content in emails can indicate the organizational hierarchy [16].

This body of work reveals the value of these social tools in enhancing information sharing among employees in large organizations, helping build common ground, and sustaining a feeling of connectedness among individuals across the organizational hierarchy [43]. Social networking tools especially were observed to help build stronger bonds

via weak ties, enable employees to reach out to unknown colleagues [10] and in finding experts in the company.

However, the examination of affect expression at work, especially in light of the growing use of social tools in the enterprise [2], has received relatively less attention. In [38], a tool called Themail was developed that represented the changing nature of email correspondence between two parties over time; while in [24] an emotional prosthetic, called AffectAura combined automatic labels for valence, arousal and engagement with multimodal contextual information to construct and present affective states to employees in a software company.

THEORETICAL FOUNDATION: WORKPLACE AFFECT

Outside the domain of online social tools and platforms, organizational behavior researchers have showed considerable interest in understanding how "organizations affect their workers, and how workers affect the organizations" [6], with respect to workers' thoughts, feelings and actions. This body of work has examined affective experiences of workers from several perspectives, primarily through ethnographic studies or archived daily records of samples of employees. The purpose has been to study the impact of affect on the quality of work outcomes, productivity, creativity, and job satisfaction [1,4,15,41]. A number of factors, centered around organizational processes have thus been proposed that could lead to the production of affect among workers. In a review work by Brief and Weiss [6], the authors classified these factors broadly as: (1) exogenous factors and conditions at work, (2) physical settings, (3) leaders and organizational roles.

(1) *Exogenous and endogenous workplace factors:* Watson in [39] investigated how cycles of feelings among workers related to lifestyle, biological, and sociocultural factors. Other studies have also found that inherent personal and behavioral traits play important role in stress and job related anxiety/depression [15,20]. A number of organization and work related conditions have also been found to lead to affective expression among workers [45], such as, dealing with administration and supervision [20], time constraints [35], negative job experiences, and workplace socialization.

(2) *Physical settings/geography:* Geographic distancing and work environments [29] have been found to impact collaborative norms that are often manifested through expressions of worker affect [13]. Many enterprises today have global offices, either to access workers with appropriate skillsets, outsource parts of big projects, or for the purpose of sales and marketing. Consequently, it is useful from an organization's point of view to understand the affective expressions of its workers in and between geographic locations, in order to improve communication exchange and foster creative idea sharing and collaboration.

(3) *Leaders, organizational structure:* Emotions among workers are deeply intertwined with the process of leading and leader outcomes [34]. For example, leaders in a negative mood generate a negative group affective tone and more effort from employees, than groups with positive mood leaders. Leaders also exhibit mood convergence with respect to workers, with respect to workplace norms and task and social interdependence [36]. This pattern of convergence is referred to in the socio-linguistics literature as communication accommodation theory [17]. The theory argues that when people of different authority/status levels interact they adjust their speech, their vocal patterns and their gestures, to accommodate to others. A deeper understanding of the manner of affect expression among individuals over the organizational hierarchy can yield insights into promoting healthy collaboration and improved group performance. What is the role leaders are playing via their social media-based interactions with employees? How are employees and leaders accommodating to one another?

Research Questions. The research presented in this paper is motivated by the theory presented above in the organizational behavior literature. As mentioned, the primary method of measuring affect in these works has been through interviews, archived daily records, or other ethnographic tools. We leverage the microblogging activity of employees to investigate, at-scale, how each one of the factors that we discussed impact the affective experiences of the enterprise workers. We frame three studies:

1) *What is the role of exogenous and endogenous workplace factors on affect expression?* [RQ 1]

   We explore how exogenous factors, aggregated into diurnal cycles, impact the emotions of employees. That is, while we have no full accounting of the exogenous factors of employees, we can look for differences in emotional expression over the course of a 24 hour period to characterize broadly the effects of exogenous factors. For instance, employees may show different patterns of affective expression during lunch time or outside of normal working hours. We then examine the endogenous factors within the workplace that trigger positive or negative affect. Finally, we investigate the role of socialization within the workplace context, via microblog conversations, in the expression of affective states.

2) *What is the role of physical location / geographic disparity on affect expression?* [RQ 2]

   Our second study is geared towards understanding the impact of geographic disparity on affect expression among employees. Specifically, we are interested in understanding how employees interacting over disparate geographic regions express emotion towards one another. From an organizational standpoint this can help with understanding interaction norms, employee allocation, and ultimately establish common ground and improve group performance across geographies and cultures.

3) *What are the roles of organizational structure and the diversity of job roles on affect expression?* [RQ 3]

In our final study, we investigate how organizational structure relates to the expression of affect. First, we will explore the affect dynamics between employees having different job roles, based on their conversations on the enterprise microblog. Next, we investigate how employees of various job levels make *accommodations* in their affective expression when communicating with colleagues.

DATA

Source: Enterprise Microblogging Tool

Our source of data is a microblogging tool, called OfficeTalk developed exclusively for internal use by the employees in Microsoft. Potential users of the microblogging tool include all employees of the corporation, who can share short posts and can consume posts from others they follow. All posts by default are public and English is the only supported language. The usernames correspond to employees' corporate email aliases and are associated with their real names. The profile associated with each username reflects the user's job title, as well as his/her position in the organizational chart. Note that the tool also allows user accounts that are non-employees, often an entity related to the company, such as the name of a product, or a team/group. The primary goal of these accounts is to disseminate product/service updates to the larger enterprise community, such as details of new upcoming features, news about future team meetings, releases or trade shows.

Although conceptually the tool is similar to Twitter, there are some key differences. The posts shared via this tool can be up to 420 characters in length. While Twitter promotes the concept of @-replies in tweets to facilitate one-to-one information exchange, this tool allows a threaded structure to the replies.

Data Characteristics

In coordination with the product group who maintained the tool, we collected all posts and their associated information shared through this tool between Dec 1, 2008 and Mar 15, 2012. There were a total of 311,054 posts from 32,813 unique users in our dataset. For the purposes of this work, we consider only the individual user accounts; which leads to 204,284 posts from the 22,968 employees. Thus the mean number of posts per user was 8.9, while the variance was 1.6e+3. Overall, we had 21% users who had more than 10 posts, while 7% who had more than 100 posts. Out of these of posts, 93,256 posts were replies to other users.

From these data we constructed a post data model and a user data model. (1) For each post, the following information defined the post data model: content of post, authoring userid, timestamp of post creation, and id of the recipient user if the post is a reply. (2) At the user level, we obtained the following information to construct the user data model: userid, user's real name, user's corporate email alias (or username), user's job title, and user's workplace location (office/building number/city/country).

Job role	Count	Job role	Count
Program Managers	1441	Sales	1335
Developers	7511	Product/IT Support	994
Testers	7041	Designer	521
Admin	625	Usability	874
Marketing	2439	Researcher	183

Table 1. Number of employees (i.e. microblog users) per job role category.

For the upcoming studies, we then categorize the users of OfficeTalk first into various generic job roles, i.e. the type of responsibilities that entail their work, and next into various levels that define their position in the organizational hierarchy at Microsoft. The job roles we considered, along with the corresponding number of users is shown in Table 1. Note that each role defines a notion of a distinct set of tasks – e.g. a *developer* is primarily entrusted with the job of programming, while a *usability* focused job role would encompass examining the clarity of a human's interaction with a computer system.

Job level	Count	Job level	Count
Executive	24	Individual contributor	13148
Manager	2092	Intern	764
Lead	6940		

Table 2. Number of employees (i.e., microblog users) per job level category.

Table 2 shows the five job levels considered in this paper along with their respective number of users. *Individual contributors* comprise employees who do not have a direct report – for instance, junior *developers* and *testers*. To illustrate the difference between *managers* and *leads*, note that typically in this company, an employee who is a *lead* is someone who supervises one or more *individual contributors*, but does not formally manage them. On the other hand, a *manager* would oversee the entire product group or a team that likely comprises *leads*, *individual contributors* and *interns*. Also note that, the *program manager* role in Table 1 is different from the *manager* in Table 2. *Program managers* are typical in tech roles within product groups, who do not manage people as a rule, they manage feature specifications and have the role of coordinating *developers*, *testers*, *designers*, and so forth, although they do not have authority over these people. Finally, while it is true that *Executives* also serve some managerial role, they are typically higher level officers heading departments and orgs, such as chairman, chief executive officer, chief operating officer, president etc. and with the general responsibility of acting as a communicator, decision maker, leader, manager and executor therein.

MEASURING AFFECT EXPRESSION

Method

We now present our methodology used to measure affect expression of users. We focus on the two broad categories of emotional expression typically used in the sentiment analysis research community [18,28] as well as the organizational behavior community (see [6] for a review) – positive affect (PA) and negative affect (NA).

In order to compute the two types of affect for the users, we utilize the textual content of the microblog posts to capture individuals' affect expressions over time. We use a prominent psycholinguistic lexicon called Linguistic Inquiry and Word Count (LIWC[1]) that identifies over 64 behavioral and psychological dimensions (e.g., "insight", "certainty", "perception") in English text, comprising a total of 4,500 words and word stems. Creators of LIWC [30] had conducted rigorous experimentation to examine how these words related to actual natural language of people – the lexicon has been found to capture around 86 percent of the words people use in written text and speech[2].

High PA post examples
@some_person1 love your shared notes. That's great!!!! Will greatly help me!!
Is ready to celebrate the (product_x) LAUNCH! Awesome times of release. Have a good weekend!
Back to work from the best vacation and happy to see the evolutions of (product_y) build!!! yeah
High NA post examples
Unknown error – seriously. And weird post caching behavior. Hate this…
I sucked at pool so bad at the launch party yesterday that I should lose against myself. Oh brother!
@some_person2 sorry I missed that! Had an emergency and couldn't come to the meeting yesterday. Damn, this sucks

Table 3. Example posts with high PA and high NA.

Turning to the social media domain, considerable work has successfully utilized LIWC to determine PA and NA, as well as a variety of emotional states including 'joy', 'anger' etc. [18,23], validating the use of LIWC for reliably measuring affect expression in short text (Twitter posts).

As in prior literature [18], we focus on the words and word stems available in the emotion categories in LIWC: for PA, we use the dimension "positive emotion", and for NA, we use "negative emotion", "sadness", "anger", "anxiety" and "inhibition". In this manner, we compile positive affect and negative affect lexicons. Based on regular expression matching of the words in each lexicon with the content of a microblog post, we determine measures of PA (or NA) respectively as the ratio of the number of positive (or negative) words in the posts to the total number of words in the post (hence PA and NA values for a given post will lie in the range {0,1}). Some example posts with high PA and high NA respectively are shown in Table 3. Note that usernames and certain terms have been replaced with dummy strings to preserve privacy.

Validation

To demonstrate the effectiveness of LIWC, we perform a validation exercise to determine if our affect measurements could truly capture the nature of affect expressed in a post[3]. For the purpose, we selected randomly a set of 50 posts, and sought feedback from three researchers (English proficient). Each researcher coded PA (and NA) separately in each post using a 1-5 Likert scale, where 1 indicated low PA (or NA) and 5 indicated high PA (or NA). We then combined mean ratings from the three researchers to obtain "ground truth" PA and NA measures (Fleiss' kappa: 0.72). In parallel, to match the Likert scale ratings, we discretized into 5 bins (bin size, 0.2), the LIWC computed affect measurements for the same 50 posts so that for each post, its PA and NA lies in one of the 5 bins. Thereafter, we compared these to the "ground truth" affect measures for all posts – the Pearson's correlation coefficient was found to be 0.83 which provides evidence that our LIWC computed affect measurements are reliable.

Note that we acknowledge a major limitation of lexicon-driven affect measurement is failure to account for negation. However, the availability of hundreds of thousands of posts in our case helps alleviate this issue, as has also been demonstrated in prior work [18]. Also, the aforementioned strong correlation with human coded posts suggests that noise in our data due to negation is minimal.

ROLE OF EXOGENOUS AND ENDOGENOUS WORKPLACE FACTORS [RQ 1]

Exogenous Factors: Rhythmic Cycles of Affect

Recall that we intend to present a summary picture of the impact of exogenous factors by examining emotional expression over the course of a day. We present in Figure 1 the mean daily trends of PA and NA over any typical 24 hour period, computed over all the users and the entire time range of analysis. We notice that NA shows a fairly stable trend throughout the day, with a minor drop in the evening. On the other hand, PA exhibits higher variability, showing a peak during morning hours between 10am and 11am, but gradual decrease through the day, with a sharp drop after 6pm. PA also is considerably greater in magnitude

[1] http://www.liwc.net/

[2] http://www.liwc.net/howliwcworks.php#index6

[3] Note that a similar human coding based validation of LIWC emotions was performed successfully for Twitter posts in [23].

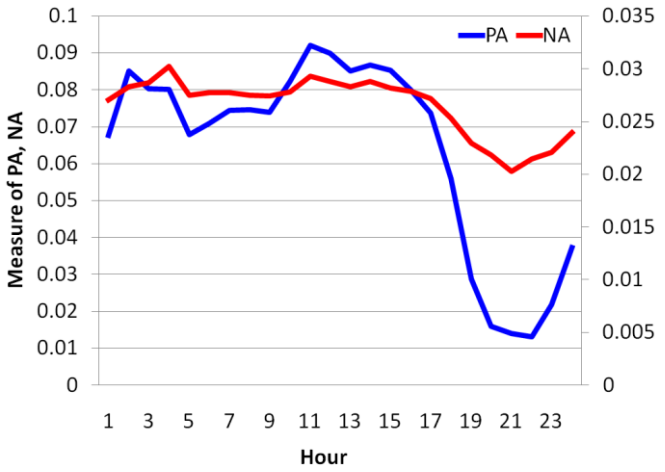

Figure 1. Diurnal trends of PA (left y-axis) and NA (right y-axis) based on enterprise-wide microblog usage. The trends show mean measures across all users and collapsed across all days in our analysis timeframe.

throughout the workday (9am-5pm), compared to NA, at least prior to the sharp drop in the evening. On average, when comparing average values per-hour of the day, employees express greater PA than NA (independent sample *t*-test, $p < .01$; df = 23).

These results reveal several implications about how users express their emotion in a workplace context: people generally express PA throughout the workday around various work-related issues and topics, but then exhibit a sharp drop in PA when working afterhours. It is possible that during afterhours, individuals are dissatisfied, saying fewer positive things, but because they are timid they do not express negativity. The stable expression of NA in our context may also indicate employee caution about expressing extreme NA around work-related topics.

Endogenous Factors: Work Context

We investigate the type of topical content associated with the two affect categories, PA and NA. We define topics here to be non-stopword unigrams present in the text of the posts, although more sophisticated topic models are certainly an option to explore in future work.

In Table 1, we report a list of some of the top keywords that are most frequent for each affect type. We notice distinct topic categories associated with PA and NA – out of the top 100 most frequent keywords for each of the PA and NA distributions, we find only 27.5% overlap. PA tends to co-occur with topics related to the company's products, its corporate partners, business-related updates and customers, as well as broader notions of 'positive feeling' in an enterprise context (e.g., *productivity*, *love*, and the hashtag *#bingo*, used in the sense of "you got it"). NA on the other hand, appears to be associated with work-related problems, updates and topics (e.g., *debug*, *toolbox*, *spec*, *demo*, *error*, *meeting*, *webcast*), vibes of 'negative feeling' possibly around work process or interactions (e.g., *confused*, *blocked*), some mentions of competing products and companies, as well technical issues (e.g., *security*, *servers*, *support*, *drill*, *VPN*). Broadly the observation that workplace issues/contexts increase negative affect expression thus aligns with findings in the literature that suggest that time constraints and the demands of the job responsibilities impact NA production [35].

Together, these observations reveal that the topical patterns around affect expression in the enterprise microblog are almost exclusively work-focused, with positivity shown toward company products and negativity toward IT issues.

	Frequent keywords
PA	windows, xbox, Microsoft, love, #helloworld, weekend, experience, cool, sharepoint, vision, technology, (company_corporate_partner_1), productive, ROI, potential, creative, business, kinect, customers, awesome, kudos, #bingo, media, fix, pretty, design, feedback, idea, thx, #windowsphone, favorite, openness, flawless, amazing, (company_corporate_partner_2), brilliant, money
NA	upgrade, client, discuss, meeting, question, debug, email, servers, support, data, (competing_product_1), network, suck, developers, managers, drill, VPN, security, error, urgent, webcast, interview, frustrating, notes, (competing_company), tired, late, confused, mysterious, toolbox, spec, proxy, users, demo, (competing_product_2), waiting, forbidden, blocked, (competing_product_3)

Table 4. Top keywords in the enterprise microblog posts, associated with PA and NA. Certain keywords are masked to ensure confidentiality (e.g., *competing_product_3*).

Endogenous Factors: Workplace Conversations

A key benefit of enterprise microblogs is that they promote conversations around work related topics. Ideally this fosters a sense of community and helps collaboration.

We study how a measure 'responsivity', defined as the normalized frequency of replies between pairs of users, is associated with expression of PA and NA. Figure 2 shows the measures of PA and NA over the number of replies, averaged over all pairs of conversing users. Overall, affect expression appears to have higher variance and higher extremities for lower responsivity rates (i.e., fewer replies). In fact for very high responsivities, we observe that users share very little emotional content (both PA, NA). The Pearson correlation coefficient between #replies and PA is −.28, and between #replies and NA is −.21. In other words, affect expression appears to decrease with increased frequency of responses. Additionally, it appears that, similar to the daily patterns, PA is comparatively higher than NA, given a certain reply frequency, reflecting the

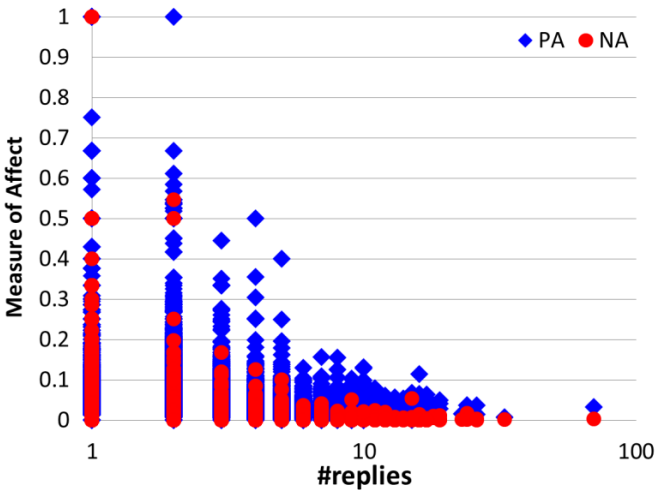

Figure 2. Responsivity (frequency of replies) between pairs of enterprise microblog users, in terms of PA and NA.

general tendency of employees to express less negativity while conversing with their colleagues.

To summarize, we observe that consistent and frequent conversations between employees tend to be focused on objective topics (e.g., discussing a common debugging issue), making conversational relationships mechanisms neutral and more for issue resolution or problem solving than for social engagement or internal networking.

ROLE OF PHYSICAL LOCATION/GEOGRAPHY [RQ 2]

In this section, we investigate the role of geographic proximity between pairs of users in the nature of affect they share through conversations. We study conversational relationships of employees, as manifested in the microblog OfficeTalk, across seven different geographical regions where Microsoft has its offices. To elaborate on our technique used here, given a particular conversing user-pair, we inferred the geographic region of the sender and the recipient, separately based on their respective office locations – we mapped the office location addresses to latitude/longitude pairs, and thereafter obtained the city, country, and region names using Bing maps API. In Table 5, we show the mean PA and NA shared and received between employees at pairs of geographic locations – for brevity, for each region, we show the region which it shares the highest PA (and the highest NA) with – measures of PA (and NA) are respectively shown in boldface font.

Across all geographic region pairs, the PA and NA distributions are significantly different from the mean levels of PA and NA (see Table 6, rows 1 and 2). That is, the distribution of values of PA (or NA) over region pairs differ significantly from the mean level of PA (or NA) expressed between region pairs. We present this to document the variability across region pairs in the expression of affect: expression of PA and NA are not the same across all regions pairs. The PA and NA distributions across region pairs are also different from one another (third row in Table 6). This means that when a region pair expresses higher PA, that pair does not also express higher NA. In fact, the Pearson correlation coefficient between PA and NA in region pairs is only .094.

Sender region	Recipient region	PA	NA
Africa	Asia	0.0211	**0.0016**
Africa	Europe	**0.09**	0.0002
Latin America	U.S./Canada	0.0266	**0.0085**
Latin America	Middle East	**0.1382**	0.0032
Australia/New Zealand	Asia	0.0301	**0.001**
Australia/New Zealand	Europe	**0.0741**	0.0007
U.S./Canada	Latin America	0.0375	**0.005**
U.S./Canada	Asia	**0.094**	0.001
Asia	Latin America	0.0185	**0.005**
Asia	U.S./Canada	**0.0415**	0.001
Europe	Australia/New Zealand	0.023	**0.007**
Europe	U.S./Canada	**0.0679**	0.0021
Middle East	Asia	0.0495	**0.003**
Middle East	Europe	**0.0743**	0.0004

Table 5. Aggregated affect expression between pairs of conversing individuals across geographical locations, based on replies in the enterprise microblog.

Further exploration shows that we do not see higher PA for geographically proximate region conversing user-pairs; neither do we notice lower NA for them (e.g., U.S./Canada → Latin America; Africa → Asia; Middle East → Asia; Table 5). In fact, strong PA is seen in Africa → Europe, U.S./Canada → Asia, and Asia → U.S./Canada relationships. While geographically coarse-grained, these results show that the (geographical/cultural) homophily of living in the same or a relatively more similar country is not a primary driver of positive emotion in workplace interactions.

	p	*t*-stat	df	std dev.
PA^	***	7.9635	46	0.0375
NA^	***	6.7727	46	0.0021
PA × NA^^	***	7.6147	46	0.0374

^ one-sample *t*-test ^^ two-tail paired *t*-test *** p<.001

Table 6. Statistical significant tests for affect expression across pairs of conversing geographical regions (there are *N*=47 region pairs). df in each test is thus based on the number of pairs of geographic regions which have conversing employees.

Rather, higher PA (and lower NA) appear more related to collaboration across regions. Precise amounts of cross-region collaboration are difficult to quantify, but it is known that employees in U.S./Canada interact extensively with those in Asia, and similarly those in Europe with Africa, on a variety of development-related and outsourced products and projects. While this warrants additional exploration, it appears that high PA (or low NA) is not a direct function of geographic or even perhaps cultural homophily, at least at a

country/regional scale. An enterprise context, especially one at Microsoft which is culturally and globally diverse, may however see exchanges of PA more along collaboration ties.

ROLE OF ORGANIZATIONAL STRUCTURE [RQ 3]

Job Roles and Affect Expression

In this section, we examine the role of the organizational structure on employees' affect expression in light of their job roles. We construct network representations of affect shared among the job roles (see Table 1), in which the nodes represent the categories, and a directed weighted edge between two categories represents the mean PA (or NA) over all replies sent by the employees associated with the first job role category to those of the second.

	p	*t*-stat	df	std dev.
PA^	***	11.05	83	0.023
NA^	***	8.21	83	0.0045
PA × NA^^	***	12.84	83	0.0237

^ one-sample *t*-test ^^ two-tail paired *t*-test *** p<.001

Table 7. Statistical significance tests for various job roles pairs' conversational exchange. df for each test is based on the number of directed edges in the job role affect network, connecting pairs of job roles.

We observe differences in affect expression across job role categories. Like the analyses for region pairs, that the amount of expression of PA (or NA) for the job role pairs differ from the mean level of (i.e., that not all job role pairs express the same levels of PA and NA toward one another; rows 1 and 2 in Table 7) and that job role pairs differ in the expression of PA and NA toward one another (row 3 in Table 7). This establishes that job roles of the employees involved in a conversation plays an important role in the nature of affect that is expressed between them.

In Table 8, we present the mean PA (and NA) shared and received by each job role with others. Based on the network representation among the job roles, we also present in Table 8, two network metrics—betweenness centrality (that estimates how the importance of a node/job role in terms of its connectivity in the network) and clustering coefficient (that estimates how well connected a node/job role is with its neighborhood/neighboring roles).

Overall, from Table 8, it appears that shared affect (whether PA or NA) over the ties between job roles is driven strongly by the roles' professional responsibilities i.e., "who works with whom". For example, *admins* and *program managers* are more centrally connected (high betweeness centrality) to a variety of other job roles, sharing higher PA and lower NA. This likely indicates admins' role of coordinating with other employees (e.g. setup a meeting, organize a conference or trade show) and keeping up the morale of people around them. On similar lines, *program managers'* goal is ensuring successful performance of their colleagues. Hence they tend to maintain more positivity in their conversational ties. From the lower values of NA we conjecture that the *program managers*, being relatively higher up in the organizational hierarchy (compared to e.g., *developers, designers, testers* etc.) exercise caution in expression of negative emotion, instead attempting to be friendly and encouraging in their encounters on public communication channels like this enterprise microblog.

PA				
Job role	**Shared**	**Received**	**Betweenness Centrality**	**Clust. Coeff.**
Program Managers	0.0895	0.0824	231.371	0.188
Developers	0.0416	0.0052	35.071	0.386
Testers	0.0327	0.0046	24.038	0.433
Admin	0.0802	0.0782	210.254	0.232
Marketing	0.158	0.103	248.392	0.169
Sales	0.173	0.149	239.458	0.173
Product/IT Support	0.0742	0.0193	234.249	0.192
Designer	0.0374	0.0041	20.094	0.402
Usability	0.0381	0.0042	23.842	0.428
Researcher	0.0524	0.0059	31.096	0.308
NA				
Job role	**Shared**	**Received**	**Betweenness Centrality**	**Clust. Coeff.**
Program Managers	0.0025	0.0016	206.973	0.196
Developers	0.0304	0.0135	28.35	0.391
Testers	0.0273	0.0324	20.942	0.414
Admin	0.0021	0.0019	187.194	0.215
Marketing	0.0024	0.0018	228.486	0.183
Sales	0.0022	0.0013	213.288	0.195
Product/IT Support	0.0082	0.0182	213.349	0.142
Designer	0.0305	0.0245	21.354	0.395
Usability	0.0284	0.0289	20.249	0.293
Researcher	0.0125	0.0117	23.065	0.326

Table 8. Mean PA and NA shared and received by various job roles, based on conversational ties. Also shown are two types of network metrics (betweenness centrality, clustering coefficient) for PA and NA, over the different job roles.

Notable also are the *Sales* and *Marketing* individuals who generally show the highest PA shared and received, and at the same time they also share and receive low NA. They are also highly centrally connected in the network. We note here that *Marketing* and *Sales* people often are selected to be positive individuals. As a result, as a part of their role they maintain positivity, because they are interested in expanding marketing efforts, engage in professional networking and improving interactivity among groups.

On the other hand, we observe that, *designers*, *developers* or *researchers* are not very central in the network (low betweenness centrality) and show high local clustering of their networks in terms of affect expression, perhaps connecting extensively to other individuals of the same job role. Overall, they also show low PA in conversations with other job roles. To elaborate this observation further, with careful inspection we find that these job roles are largely *individual contributors* in the workplace (see Table 2, for job levels) – consequently they are probably less motivated to use the microblogging tool to engage in positive affect-bearing conversations with others. However with the relatively higher NA for these roles, it appears that these employees are perhaps using OfficeTalk to express their negative issues/concerns with others – for either their own problem resolution or for someone's critical task at hand. We also note here that *developers* and especially *testers* are selected to be people who look for faults, which can come across as negative.

To sum up, we note that in an enterprise context, the stratification of organizational structure appears to play an important role in explaining the nature of affect expression. Those that are central in the enterprise network tend to share and receive high positive affect. Conversely, those in individual contributor roles tend to express more negative than positive affect. This finding aligns with observations in prior work: communication patterns and affective content in organizations tend to be structured by its hierarchy [16].

Job Level and Affect Expression

As mentioned in the theoretical foundations section, job level of employees impacts their affect expression--leaders' negative affect generates negativity among their followers [34]. Here we explore this question from the perspective of the five types of job levels in the context of this company (Table 2).

We therefore define a measure called "affect accommodation", that is the degree of adaptation or adjustment a user is likely to make in her expression of affect while conversing with another user, of possibly different status, compared to her baseline level of affect that appears in her non-conversational posts. For instance, we may hypothesize that employees may be more positive than their baseline level of positivity when replying to people of high status in the organizational hierarchy. Generally, let us suppose that an individual u_i sends a reply to another individual u_j. We further suppose that u_i will make a certain degree of adjustment in her affect expression based on u_j's status (in the hierarchy).

Hence, affect accommodation is defined as the mean difference between the levels of affect (PA or NA) expressed by u_i in posts which were replies to u_j, and the level of affect expressed in posts that were not replies to anyone (non-conversational posts assumed as baseline). More generally, for any u_i, u_j, u_k etc., we further define measures of affect accommodation separately for the sender and recipient roles in replies. We investigate (1) what degree of affect accommodation is observed for all conversations sent by u_i and received by individuals, u_j, u_k, ---, of possibly varied statuses and (2) what degree of affect accommodation is observed for all conversations sent by individuals u_i, u_k, ..., of possibly varied statuses, and received by u_j? Finally, by aggregating across the job levels associated with each pair of individuals u_i and u_j involved in conversations, we get affect accommodation measures for the sender and recipient roles respectively for pairs of job levels.

	df	*t*-stat		std dev.	
		PA	NA	PA	NA
Executive × *Manager* **	2,114	6.93	5.73	0.048	5.7e-3
Executive × *Lead* **	6,962	11.15	9.29	0.075	8.3e-3
Executive × *Individual contributor* ***	13,170	14.98	13.73	0.163	1.5e-2
Executive × *Intern* ***	786	4.02	3.14	0.014	3.6e-3
Manager × *Individual contributor* ***	15,238	17.66	14.85	0.224	2.1e-2
Manager × *Intern* ***	2,854	7.35	6.02	0.056	5.2e-3
Manager × *Manager* *	4,182	2.36	1.95	0.019	6.8e-3
Lead × *Individual contributor* ***	20,086	22.35	16.59	0.295	2.8e-2
Lead × *Intern* ***	7,702	12.94	10.26	0.079	7.5e-3

Table 9. Statistical significance tests to compare the degrees of PA and NA accommodation across job levels (* p<0.01; ** *p*<.001; * *p*<.0001). df for each pair is based on the number of employee-employee conversing pairs with the particular two job roles under consideration. Only pairs with significant PA and NA accommodation are shown.**

We begin by examining the differences in PA and NA accommodation shown by various job levels for the above two circumstances: (1) when (employees with) a particular job level is a sender of replies targeted to another job level, and (2) when a particular job title is a recipient of replies sent from (employees with) another job level. Overall, from Table 9, for most pairs we observe statistically significant differences in the degree of accommodation—greater accommodation is observed when the two job levels are distinct.

Further, Figure 3 shows the average measures of affect accommodation for each job level. For affect accommodation with respect to the senders, we find that

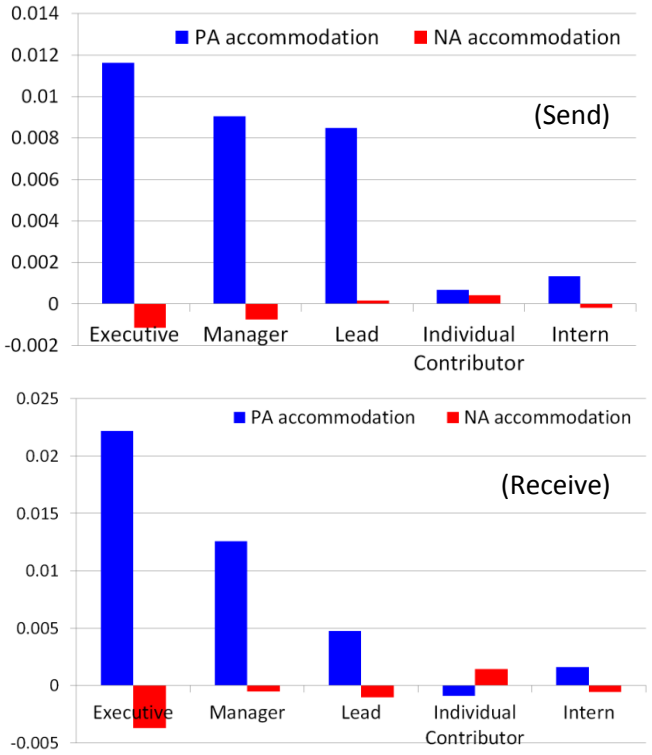

Figure 3. Affect accommodation by employees holding various job levels: (1) Accommodation while sending/sharing affective content (top), (2) Accommodation by others in receiving affective content (bottom).

most job levels tend to express positive accommodation for PA. In particular we observe that higher status job levels such as *Executives*, *Managers*, *Leads* tend to share notably high accommodation in the positive direction for PA and in the negative direction for NA; i.e., they are more positive and less negative while conversing with other employees. These results generally support the theoretical observation that transformational leaders in an organization often bear positivity to help ignite others' aspirations, eliciting enthusiasm, and conveying optimism regarding a desirable future for the workplace [14].

In contrast *Individual Contributors* and *Interns* share less PA and NA accommodation: in fact for the *Individual Contributors* NA accommodation is in the positive direction which implies that they are more negative while engaging with others, than they are on their own.

For affect accommodation with respect to recipients, there is again high PA accommodation in the positive direction for top job levels (e.g., *Executives*, *Managers* etc.)—employees tend to be exceedingly positive in their conversations targeted towards them. At the same time, NA accommodation towards them is negative, meaning individuals tend to be less negative than their baseline while addressing replies to them.

DISCUSSION

Theoretical Implications

The three studies presented in the previous three sections, in a broad sense, provide at-scale empirical evidence in support of the theoretical foundations behind the production of affect at workplace (refer to the three research questions in third section). They also suggest theoretical implications that can be actionable from an organization's perspective.

Exogenous factors. To begin with, the patterns of affect expressed over the course of a day exemplify findings that are actionable on the part of management. That positivity drops off dramatically in the evening is one sign that employees are less satisfied with afterhours work – an aspect that might be critical while employees are meeting deadlines, product pre-release duties, or collaborating with colleagues in distinctively different timezones. Overall, this is not surprising, but the size of the drop in positivity may be surprising or at least informative in the context of work-life balance discussions. Employees do experience *consistent* cycles of positivity throughout the day that are likely to impact their performance levels or how they accomplish targets meant for their particular job role [20]. The consistency is something management could leverage when, for instance, establishing cross-timezone collaborations. While somewhat unexpected, the similar but smaller drop in NA in the evening might indicate that afterhours work does not necessarily result in explicit negative expression. We acknowledge that this could be due to a general reticence to criticize, or simply a feeling that it is not productive to be openly critical or negative.

Endogenous factors. The analysis of topics reveals specific aspects of work that cause both anxiety and negativity, as well as feelings of positivity toward work, broadly helping the organization place employee job satisfaction in the context of workplace conditions and events (social, physical, economic) [29]. From an application standpoint, such ability to align employee emotion with specific contexts is potentially a substantial step toward an actionable, fine-grained feedback loop from employees to management.

For instance, although some of the topics we observed (e.g., the product names) are specific to the company we studied, many of the discomfort points, like the IT related issues (e.g., *upgrading*, *servers*, *support*), are found to be commonplace sources of frustration. Though this observation is not particularly a surprise, but which IT issues get discussed the most or are most associated with negative affect, can guide prioritization of managerial efforts toward amelioration; at the same time the organization can utilize this knowledge to streamline how technical issues are handled internally in order to prevent counter-productivity among the employees. Capturing these issues in the flow of daily work, versus via an annual survey, should thus provide a more accurate reflection of impediments to productivity.

Similarly, words like *urgent, meeting, confused, question* etc.'s frequent co-occurrence with NA might indicate dissatisfaction while attending to the day-to-day demands of a particular job role, which in turn might have implications in how the organization could tackle bureaucratic issues.

Finally, it is well-established that positivity within organizations can be contagious [39] and has been proposed to have diminishing effects on the dysfunctional, counterproductive emotions and behaviors barricading employees from achieving organizational goals. Frequent use of words like *awesome, kudos, #bingo, amazing, brilliant* with PA thus provides helpful evidence to the organization in this direction: that employees are sharing/ propagating their positive perceptions amongst them.

Physical Location/Geography. Our results also have implications regarding the strong role of the workplace's geographic environment [13] on the expression of affect. For instance, we saw significant differences in levels of positivity and negativity in conversational interactions depending on the geographic region of the employees. One explanation is that employees in the same geographic region share affect through other communication mechanisms (e.g. face-to-face interaction, phone conversations), minimizing the need for the use of online social tools for affect sharing. For geographically disparate colleagues, this might be reversed. Since most of their professional interactions tend to be over emails or virtual meetings, they may use the microblog for networking or engagement and thereby express higher positive affect (and lower negative affect), to compensate for not being able to share it via other physical communication means. This is furthermore supported by similar observations in the organization behavior literature – e.g., Kiesler and Sproull [46] described how people corresponding over email in an organization used capital letters to emphasize or stress important concepts, and add emotion back into a more emotionally sterile context through the use of emoticons.

Tentatively we thus conclude that regions that are geographically separated but that collaborate more tend to also share more positive emotion. The variation in affective expression across geography can inform the organization about how to distribute employees across teams who need to collaborate consistently, or how to regulate norms of the remote collaboration/communication environment, which can in turn be valuable in making outsourcing decisions with overseas clients.

Organizational Structure. There is also a considerable difference in affect expression with respect to centrality in the organizational network. Consequently we observe accommodation of affect between employees who are individual contributors versus those who lead or manage groups and teams. As an actionable item, the finding can provide helpful guidelines to managers about how to deal with administration and supervision, and how they can foster and shape the temperament of their subordinates through the use of positive affective cues in their communication, thereby improving group performance [3].

Interestingly, however, while examining the interactive effects of leader affect and group's goals within Cognitive Appraisal Theory, the authors in [14] found that leader negative affect was perceived less favorably when groups were focused on promotion goals, whereas leader negative affect was perceived more favorably in groups with prevention goals. Hence the nature of affect expression we observed across the organizational hierarchy in our studies points to a helpful guideline of how the *Managers*, *Leads* and *Executives* could utilize negative affect expression that could be deemed appropriate in the face of challenging situations.

Together, these influences highlight positive attributes of the enterprise (collaboration generating positivity), offer insight into "soft skill" requirements (*managers* and *sales* people staying positive) and suggest avenues for improvement (perhaps *developers* could be less negative when interacting with others).

Design Implications

We discuss some potential design implications of our studies. One possibility is that an organization's HR/administrative team can be assigned the role to monitor an "affect dashboard", a system that describes the prevalent mood of employees about specific topics in an anonymous manner. The team could look for affect trends over time, both daily and longitudinally, in the context of specific issues and products, and identify the affective influence exerted by different job roles. Further, they could assess emotional reactions among employees in different locations or in different departments without revealing explicit affective scores to individuals. For instance, after a change in health benefits is announced, the organization can use this dashboard to find out what type of reception the news has received among staffers. This would let the organization locate and tune into critical issues and concerns, and propagate them up to suitable authorities for requisite action or change, if necessary.

The other potential design implication could be an individual-focused "application" (e.g. a smartphone app) that the employees could use on their own for self-monitoring and self-evaluation of their stress or anxiety levels (or their coping strategies) in order to promote emotional reflection under different workplace contexts.

Privacy

Finally, we mention the issue of respecting individual privacy in reviewing employee affect via social media. We note that while our monitoring of employee microblog posts may seem a bit "big brother", our analyses as well as the design implications we discussed are aggregated over individual user behaviors where individual identities are excluded from being part of the data. That is, they would

largely prevent one actually deciphering "who expressed what emotion on which topic or about whom". Our hope is that such aggregated findings maintain this minimization of privacy infringement while still providing insight into employee affect that is actionable toward positive organizational directions. However we note that it is possible that knowing this type of affect-monitoring system exists, employees might find it unsettling for themselves. The type of reaction such mining could trigger is, however, hard to gauge without qualitative feedback from them, or a pilot deployment, both being directions for future work.

Limitations

We discuss a few limitations of this work. As a critique of our affect measurement method, we again acknowledge that a lexicon driven approach via LIWC can have limitations. It is worthwhile to examine alternative lexica (e.g., POMS [25], PANAS [40]) that describe emotional states beyond PA/NA. We would also like to mention that measuring arousal or activation of affect beyond pleasure-displeasure can potentially be useful, especially given the Yerkes-Dodson law that there are optimal ranges of arousal for productivity [42]. In the future, we look forward to characterizing workplace affect through a more complex psycho-linguistic instrument such as the circumplex model [31] that can enable a more nuanced investigation of affect in terms of valence and arousal.

CONCLUSION

The emergence of social media as a data source for understanding affect of employees in organizations carries considerable potential—it can open up the "watercooler conversations" and help organizations get a real sense of what is on the mind of employees, stay abreast of their morale or assess reaction to important changes such as a corporate restructuring or a product launch.

Toward this end, we explored the landscape of emotional expression from the employees at the software corporation Microsoft, leveraging the data from an internal microblogging tool called OfficeTalk. Grounded by theoretical foundations in organizational behavior research, we examined three broad factors that have been known to influence affective expression: (1) exogenous and endogenous workplace factors, (2) geography, and (3) organizational structure. We found empirical at-scale evidence for effects of job role and level, geographic region, exogenous factors like daily life patterns, and endogenous factors like IT issues on the expression of affect. Our findings together bear the potential of development of affect assessment systems in enterprises that can reveal aggregated temporal dynamics of how employees feel. In essence, the timescale and context for affect expression in the workplace can provide a more effective tool for assessing key factors and performance-relevant outcomes, such as job satisfaction, judgments, attitudinal responses, creativity, helping behavior and risk-taking, than internal surveys or other traditional methods.

ACKNOWLEDGMENTS

The authors would like to thank Jonathan Grudin for thorough and helpful comments and feedback on early drafts of the paper. They are also thankful to the erstwhile Office Labs colleagues, Chris Slemp and Anne Archambault for help with data gathering and access.

REFERENCES

1. Amabile, T. M., Barsade, S. G., Mueller, J. S. and Staw, B. M. (2000) Affect and Creativity at Work. Administrative Science Quarterly, 50, 3 (2005), 367.
2. Archambault, A. & Grudin, J. (2012). A longitudinal study of Facebook, LinkedIn, & Twitter use. In *Proc. CHI 2012*.
3. Ashforth BE, Humphrey RH. (1995). Emotion in the workplace: a reappraisal. Hum. Relat. 48:97–125.
4. Ashkanasy, N.M., Hartel, C. E. J. and Zerbe, W. J. (2000). Emotions in the Workplace: Research, Theory, and Practice. Westport, CT: Quorum Books, 2000.
5. Bartel C.A., Saavedra R. (2000). The collective construction of work group moods. *Admin. Sci. Q.* 45(2):197–231.
6. Brief, A. P., & Weiss, H. M. (2002). Organizational behavior: affect in the workplace. Annual Review of Psychology, 53(1), 279-307.
7. Brzozowski, M. J. (2009). WaterCooler: exploring an organization through enterprise social media. In *Proc. GROUP 2009*. 219-228.
8. Dabbish, L. A., Kraut, R., Fussell, S. & Kiesler, S. (2005). Understanding email use: predicting action on a message. In *Proc. CHI 2005*. 691-700.
9. De Choudhury, M., Counts, S., & Gamon, M. (2012). Not All Moods are Created Equal! Exploring Human Emotional States in Social Media. In *Proc. ICWSM'12.*
10. DiMicco, J., Millen, D., Geyer, W., Dugan, C., Brownholtz, B. & Muller, M. (2008). Motivations for social networking at work. In *Proc. CSCW 2008*. 711-720.
11. Efimova, L. & Grudin, J. (2007). Crossing Boundaries: A Case Study of Employee Blogging. In *Proc. HICSS 2007*.
12. Ehrlich, K. & Shami, S. (2010). Microblogging inside and outside the workplace. In *Proc. ICWSM 2010.*
13. Eligio, Ulises Xolocotzin, Ainsworth, Shaaron E., Crook, Charles K. (2012). Emotion understanding and performance during computer-supported collaboration. *Computers in Human Behavior*, ISSN 0747-5632, 10.1016/j.chb.2012.06.001.
14. Gaddis, B., Connelly, S., & Mumford, M. D. (2004). Failure feedback as an affective event: Influences of

leader affect on subordinate attitudes and performance. *The Leadership Quarterly*, 15(5), 663−686.

15. George J.M., Brief A.P. (1992). Feeling good doing good: a conceptual analysis of the mood at work-organizational spontaneity relationship. *Psychol. Bull.* 112(2):310–29.
16. Gilbert, Eric. (2012). Phrases that signal workplace hierarchy. In *Proc. CSCW 2012*. 1037-1046.
17. Giles, Howard; Smith, Philip (1979). Accommodation Theory: Optimal Levels of Convergence. In Giles; St. Clair, Robert N. *Language and Social Psychology*. Baltimore: Basil Blackwell.
18. Golder, S. A., & Macy, M. W. (2011). Diurnal and Seasonal Mood Vary with Work, Sleep and Daylength Across Diverse Cultures. *Science*. 30 Sep 2011.
19. Hackman J.R. (1992). Group influences on individuals in organizations. In Handbook of I/O Psychology, ed. MD Dunnette, LM Hough, 3:199–267. Palo Alto, CA: Consult. Psychol. Press. 1095 pp.
20. Hart P.M., Wearing A.J., Headey B. (1995). Police stress and well-being: integrating personality, coping and daily work experiences. J. Occup. Organ. Psychol. 68(2):133–56.
21. Herbsleb, J., Atkins, D., Boyer, D., Handel, M. & Finholt, T. (2002). Introducing instant messaging and chat in the workplace. In *Proc. CHI 2002*. 171-178.
22. Huh, J., Jones, L., Erickson, T., Kellogg, W. A., Bellamy, R., & Thomas, J. C. (2007). BlogCentral: the role of internal blogs at work. In *Proc. CHI EA 2007*.
23. Kivran-Swaine, F., and Naaman, M. (2011). Network properties and social sharing of emotions in social awareness streams. In *Proc. CSCW 2011*.
24. McDuff, D., Karlson, A., Kapoor, A., Roseway, A. & Czerwinski, M. (2012). AffectAura: an intelligent system for emotional memory. In *Proc. CHI 2012*.
25. McNair, D.M., Lorr, M., & Droppleman, L.F. (1971). Profile of mood states. San Diego: Educational and Industrial Testing Service.
26. Millen, D., Feinberg, J., Kerr, B. (2006). Dogear: Social bookmarking in the enterprise. In *Proc. CHI 2006.*
27. Muller, M. (2007). Comparing tagging vocabularies among four enterprise tag-based services. In *Proc. GROUP 2007*.
28. Ortony, A., & Turner, T. J. (1990). What's basic about basic emotions? *Psychological Review*, 97, 315-331.
29. Parkes J. (1990). Coping, negative affectivity, and the work environment: additive and interactive predictors of mental health. *J. Appl. Psychol.* 75:399–409.
30. Pennebaker, J. W., Mehl, M. R., & Niederhoffer, K. (2003). Psychological aspects of natural language use: Our words, our selves. *Annual Review of Psychology*, 54, 547-577.
31. Russell, James A. (1980). A circumplex model of affect. *J. of Personality and Social Psychology*: 39, 1161–1178.
32. Shami, S., Muller, M., & Millen, D. (2011). Browse and discover: Social file sharing in the enterprise. In *Proc. CSCW 2011*.
33. Skeels, M., & Grudin, J. (2009). When social networks cross boundaries: a case study of workplace use of facebook and linkedin. In *Proc.* GROUP *2009*. 95-104.
34. Staw, B. M. and Barsade, S. G. (1993). Affect and managerial performance: A test of the sadder-but-wiser vs. happier-and-smarter hypothesis. *Administrative Science Quarterly*, 38: 304-331.
35. Teuchmann K., Totterdell P., Parker S.K. (1999). Rushed, unhappy, and drained: an experience sampling study of relations between time pressure, perceived control, mood, and emotional exhaustion in a group of accountants. J. Occup. Health Psychol. 4(1):37–54.
36. Totterdell P., Kellett S., Teuchmann K., Briner B.B. (1998). Evidence of mood linkage in work groups. J. Pers. Soc. Psychol. 74(6):1504–15.
37. Turner, T., Qvarfordt, P., Biehl, J. T., Golovchinsky, G. & Back, M. (2010). Exploring the workplace communication ecology. In *Proc. CHI 2010*. 841-850.
38. Viégas, F., Golder, S. and Donath J. (2006). Visualizing email content: portraying relationships from conversational histories. In *Proc. CHI '06*. 979-988.
39. Watson D. (2000). Mood and Temperament. New York: Guilford. 340 pp.
40. Watson, D., Clark, L. A., & Tellegen, A. (1988). Development and validation of brief measures of positive and negative affect: the PANAS scales. *Journal of Personality and Social Psychology*, 54(6), 1063-1070. American Psychological Association.
41. Weiss H.M., Nicholas J.P., Daus C.S. (1999). An examination of the joint effects of affective experiences and job beliefs on job satisfaction and variations in affective experiences over time. Organ. Behav. Hum. Decis. Process. 78(1):1–24.
42. Yerkes RM, Dodson JD (1908). The relation of strength of stimulus to rapidity of habit-formation. *Journal of Comparative Neurology and Psychology* 18: 459–482.
43. Zhang, J., Qu, Y., Cody, J. & Wu, Y. (2010). A case study of micro-blogging in the enterprise: use, value, and related issues. In *Proc. CHI 2010*. 123-132.
44. Zhao, D. & Rosson, M. B. (2009). How and why people Twitter: the role that micro-blogging plays in informal communication at work. In *Proc*. *GROUP 2009*.
45. Zohar D. (1999). When things go wrong: the effect of daily work hassles on effort, exertion and negative mood. J. Occup. Organ. Psychol. 72(3):265–83.
46. Kiesler, S., & Sproull, L. (1992). Group decision making and communication technology. Organizational Behavior and Human Decision Processes, 52, 96–123.