# Cultural Dimensions in Twitter: Time, Individualism and Power

**Ruth Garcia-Gavilanes**
Yahoo! Research
Universitat Pompeu Fabra
Barcelona, Spain
ruthgavi@yahoo-inc.com

**Daniele Quercia**
Yahoo! Research
Barcelona, Spain
dquercia@yahoo-inc.com

**Alejandro Jaimes**
Yahoo! Research
Barcelona, Spain
ajaimes@yahoo-inc.com

**Abstract**

Previous studies have established the link between one's actions (e.g., engaging with others vs. minding one's own business) and one's national culture (e.g., collectivist vs. individualistic), and such actions have been shown to be important as they are collectively affiliated with a country's economic outcomes (e.g., Gross Domestic Product). Hitherto there has not been any systematic study of whether one's action on Twitter (e.g., deciding when to post messages) is linked to one's culture (e.g., country's Pace of Life). To fix that, we build different network snapshots starting from 55,000 seed users on Twitter, and we do so for 10 weeks across 30 countries (after filtering those with low penetration rates) for a total of 2.34 M profiles. Based on Hofstede's theory of cultural dimensions and Levine's Pace of Life theory, we consider three behavioral patterns on Twitter (i.e., temporal predictability of tweets, engaging with others, and supporting others who are less popular) and associate them with three different dimensions derived from the two theories: Pace of Life, Individualism and Power Distance. We find the following strong correlations: activity predictability negatively correlates with Pace of Life ($r = -0.62$), tweets with mentions negatively correlates with Individualism ($r = -0.55$), and power (e.g, Twitter popularity) imbalance in relationships (between, for example, two users mentioning each other) is correlated with Power Distance ($r = 0.62$). These three cultural dimensions matter because they are associated with a country's socio-economic aspects - with GDP per capita, income inequality, and education expenditure.

## Introduction

Researchers have found that the ways people perceive and accept power differences, interact with each other, and perceive time, drastically differ across countries. For example, in certain countries (e.g., Japan), direct disagreement is synonym of confrontation, while speaking one's mind is a virtue in others (e.g., USA). Also, cultures with a fast Pace of Life (e.g., Germany, Switzerland) tend to give more importance to punctuality and have less flexible schedules; by contrast, cultures with a slower Pace of Life (e.g., Brazil) are more flexible and give importance more to human interactions than to keeping the schedule (Levine 2006).

Cultural variations across countries have been empirically studied using small-scale experiments and surveys in the real world. Geert Hofstede administered opinion surveys to IBM employees in over 70 countries (Hofstede and Minkov 2010). This data, with over 100,000 questionnaires, were one of the largest cross-national databases that existed in 1971. By analyzing it, Hofstede discovered that there were significant differences between cultures: he found that five main factors explained most of the variance in the data and called those factors cultural dimensions, and two of those have been widely studied. The first is *Power Distance* and reflects the extent to which people (especially those less powerful) expect and accept that the power is distributed unequally (e.g., employees would rarely contradict their managers). The second dimension is called *Individualism vs. Collectivism* and reflects the extent to which social relationships are loose (e.g., people look after themselves and are likely to have friends outside their immediate families) as opposed to relationships integrated in strong and cohesive groups (e.g., friends are likely to be within families).

An additional aspect that varies across countries and has been widely studied is Pace of Life. Robert Levine run different experiments to capture Pace of Life in a variety of countries. In 31 countries, he and his students measured the time it takes for people to walk 100 meters in coffee shops, for post clerks to send a parcel, and they also kept track of the accuracy of clocks in public spaces (e.g., in post offices). That resulted into ranking those countries by what they then called Pace of Life (Levine 2006).

Individualism, Power Distance, and Pace of Life have been found to determine how people behave differently in the same situations in the real world. The main goal of this work is to assess the extent to which such differences can also be captured from online interactions. We will see that these differences matter because they are associated with the economic aspects of GDP per capita, income inequality and education expenditure.

To go beyond small-scale experiments and surveys, we consider Twitter, a microblog massively used worldwide, and set out to answer the following research question: Does national culture determine the temporal randomness with which Twitter users post, or the extent to which they men-

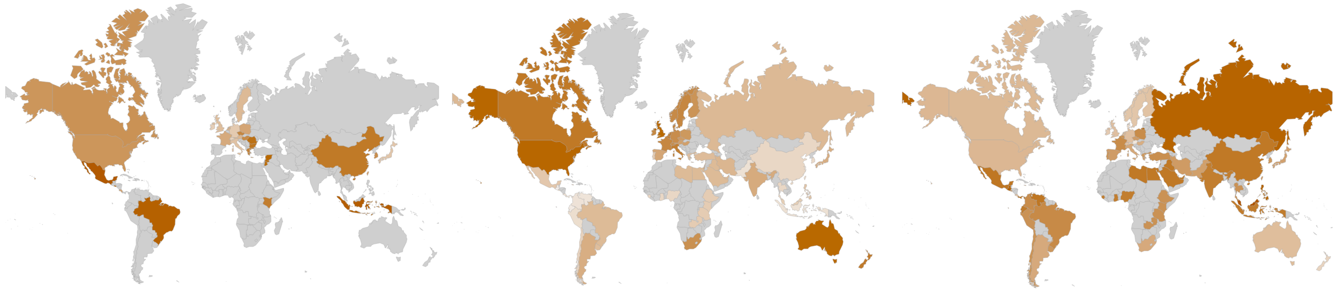

(a) Pace of life ranking (b) Collectivism vs. Individualism (c) High vs. Low Power Distance

Figure 1: Maps showing (a) Levine's Pace of Life ranking, (b) Hofstede's Individualism and (c) Power Distance. Darker colors reflect lower Pace of Life, higher Individualism score and higher Power Distance. Gray areas mark countries that have not been included in Levine's study or Hofstede's.

tion, follow, recommend and befriend others? We crawl more than 2.34 million user profiles (starting from 55K seed users), their tweets during 10 weeks from March to May 2011, their geographic locations, and corresponding time stamps (Section "Data Description"). Upon this data covering the 30 most represented countries in our sample, we test three main hypotheses associated with the three cultural aspects and, in so doing, we make four main contributions:

- We test whether the higher a country's Pace of Life, the more predictable its citizens' temporal patterns (Section "Pace of Life"). The link between Pace of Life and temporal predictability comes from the finding that countries with higher Pace of Life tend to schedule their time in more predictable ways (Levine 2006), (Woods 2003). To test this on Twitter, in our period of ten weeks, we divide each working day into 5 segments and compute the extent to which each user tweets or mentions others in the same daily segments. We aggregate all users in each of the 30 countries, produce a country-level temporal predictability, and correlate it with the country's Pace of Life. The correlations are $r = -0.62$ for tweets' temporal unpredictability, $r = -0.68$ for user mentions', and $r = -0.58$ for tweeting activity within working hours. These consistent results confirm that countries with higher Pace of Life tend to be more predictable not only offline but also online.
- We also test whether people in collectivist countries interact more with each other than those in individualistic countries (Section "Individualism vs. Collectivism"). We do so by computing the percentage of users who mention each other. We find that the correlation between Individualism index (one of Hofstede's cultural dimensions) and the extent to which users mention each other is as high as $r = -0.55$.
- We test whether users in countries comfortable with unequal distribution of power (high power-distance countries) will follow, recommend, and accept recommendations preferentially from users who are more popular (Section "Power Distance"). To this end, we consider three types of relationships: a) who follows whom; b) who recommends whom; and c) who starts to follow whom upon a recommendation. For each relationship, we compute the difference of followers between the pair of users in the relationship, and call that power imbalance. We then correlate country-level imbalance with corresponding Power Distance (another one of Hofstede's cultural dimensions). We find that the correlations are $r = 0.62$ for "who follows whom" relationships; $r = 0.33$ for "who recommends whom" relationships, and $r = 0.42$ for "who starts to follow whom" relationships.
- We finally show that those three cultural dimensions are associated with the three economic indicators of GDP per capita, income inequality and education expenditure. We find correlations as strong as $r = 0.60$.

These strong correlations suggest that cultural differences are not only visible in the real world but also emerge in the way people use social media. To show why these cultural dimensions matter, we will study their relationships with socio-economic indicators, including GDP per capita. We conclude by discussing the theoretical and practical implications of this work (Section "Discussion").

## Related Work

Our goal is to study variations of Twitter use across countries. In a similar way, researchers have already analyzed how a variety of aspects of the online world change across countries. They have studied country-variability of the following aspects:

*Twitter network metrics and use of emotion words.* Poblete *et al.* studied more than 4.5 million Twitter users with corresponding reciprocal links and tweets (Poblete et al. 2011). Network features such as the modularity, density, diameter, assortiativeness and graph degree differed significantly across countries. For example, Indonesia's network modularity is far higher than Australia, and the country with the highest use of positive emotion words is Brazil.

*Flickr pictures about the same thing but in different regions.* Yanai *et al.* used state-of-the-art object recognition techniques to find representative geotagged photos related to a given concept and then studied how photos related to a

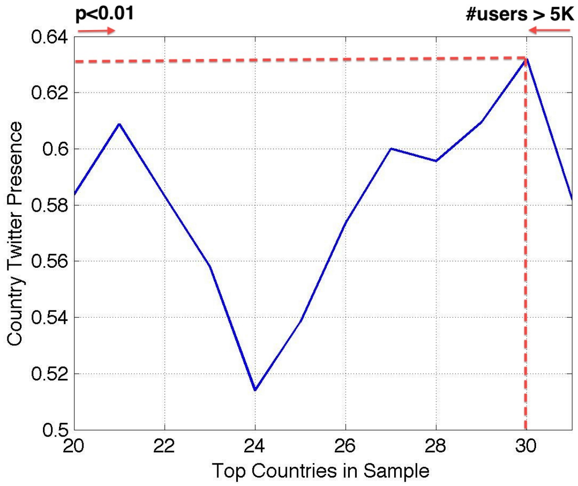


(a) Country-level presence on Twitter.

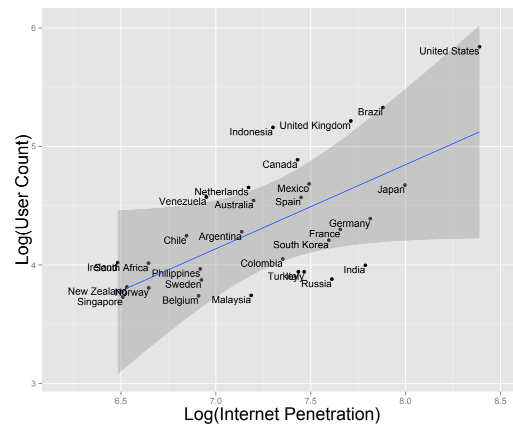


(b) Users in sample vs. Internet users.

Figure 2: *(a)* Plot of country-level presence on Twitter vs. the number of countries in our sample. The highest number of countries for which Twitter presence is significant is around 30. That is, by considering the top 30 countries (by top, we mean countries ordered by number of users our sample has), we strike the right balance between representative presence on Twitter and number of countries under study. *(b)* Number of users in our sample versus number of Internet users in a country. Both quantities are log-transformed.

concept change across countries (Yanai, Yaegashi, and Qiu 2009). They found, for example, that pictures of wedding cakes in US are much taller than those in Europe.

*Travel destinations derived from Flickr posting.* Based on where pictures are taken, Kling *et al.* derived travel patterns of a large number of Flickr users across countries (Kling and Gottron 2011). They then used clustering methods to determine the extent to which any given pair of countries is related. They found that residents in Brazil and Chile have common travel destinations, for instance.

*Color preferences in Instagram pictures.* Hochman *et al.* extracted colors from pictures and found notable differences across pictures of different countries (Hochman and Schwarts 2012). For instance, hues of pictures in New York are mostly blue-gray, while those in Tokyo are characterized by dominant red-yellow tones.

*Download times of research publications.* Wang *et al.* collected real-time data on which publication was downloaded at which time from the Springer Verlag website for 5 weekdays and 4 weekends (Wang et al. 2012). Upon the resulting dataset of 1,800,000 records, they found that downloads during weekends were the most common in Asian countries, and the least common in Germany.

*Use of the Doodle scheduling web tool.* Reinecke *et al.* studied how the use of the web scheduling tool varies across 211 countries (Reinecke et al. 2013). They did so by relating activity features (e.g., consensus, availability) to Hofstede's Colllectivism vs. Individualism dimension, and with Inglehart Survival and self expression values. They found that users of the tool in Germany tended to schedule far ahead of time (around 28 days in advance, while those in Colombia schedule up to 12 days in advance).

*Expression of emotions in Twitter status updates.* Golder and Macy studied the 500 million English tweets that 2.4 million users produced during almost 2 years. Based on their hour-by-hour analysis, they found that offline patterns of mood variations also hold on Twitter: mood variations were associated with seasonal changes in day length, people changed their mood as the working day progressed, and they were happier during weekends (Golder and Macy 2011).

*Query logs.* Baeza-Yates *et al.* studied the geographic locations of users who clicked on 759,153 Internet hosts. They found that users tended to click on hosts in other countries in which people speak the same language, and clicks coming from countries with similar human development index tend to end up into the same countries (Baeza-Yates, Middleton, and Castillo 2009). More recently, Borra and Weber analyzed the queries that resulted into clicks on the top-155 U.S. political blogs to infer users' political leanings (Borra and Weber 2012).

There has not been any work on how cross-country variations of language independent features (e.g., predictability, mentions and subscription activity) in a general-purpose platform (e.g., in Twitter) are associated with indicators well-established in anthropological studies (e.g, cultural dimensions, Pace of Life). That is why we run such a study next.

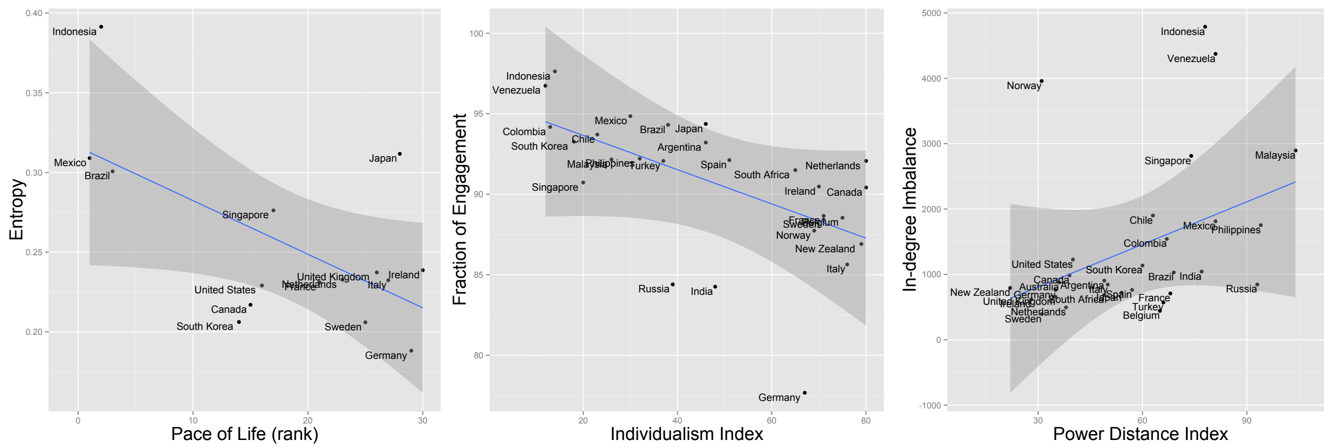


(a) Entropy vs. Pace of Life. (b) Engagement vs. Individualism. (c) Imbalance vs. Power Distance.

Figure 3: *(a)* Entropy of posting and mentioning activities versus Pace of Life. Countries with high pace of life tend to be temporally predictable. *(b)* Fraction of users engaged with others versus Individualism. In countries with low Individualism, users tend to engage with each other more. *(c)* Indegree imbalance between user-followee versus Power Distance users have stronger in-degree imbalance.

## Data Description

From the Twitter stream API, we randomly selected 55K users who tweeted at least once in March 2011 and that had an outdegree and indegree in the range $[100, 1K]$. This choice is imposed by API restrictions but has the side benefit of filtering away less legitimate (e.g., spam) users: the majority of spam users tend to have outdegree and indegree outside the range $[100, 1K]$ (Lee, Eoff, and Caverlee 2011). From these users, we selected those who specified their location in their profile; often these locations are either strings specified by the users themselves or GPS coordinates coming from their mobile devices. We then map these locations into (*long,lat*) points (using Yahoo! PlaceMaker for user-specified strings), resulting into 12.6K seed users that are geo-located. For these seed users, we collected their outbound links (followees) and found that 1.96M of them had location information. Since one of our hypotheses require the study of recommendations (which are made using the follow friday hashtag *#ff* or *#followfriday* in Twitter), we also collected the recommendations (i.e., users to follow) that were made by the followees (outbound links) during the subsequent 10 weeks. This resulted in 362K recommended users with valid geolocations. Overall, we will study 2.34M users (12.6K+1.96M+362K), their timestamped tweets (considering all different timezones), and their locations.

The way we sample our users is convenient and easy to interpret but might be biased by our particular choice of seeds. To partly address this concern, we only considered the top-30 countries in our sample. We choose 30 because it is the highest number of countries in which presence on Twitter highly correlates with presence on the Internet (Figure (2a)), and in which the number of per country users is always more than 5K, ensuring statistically significance of our results. Figure (2b) plots the number of users in our sample as a function of the number of Internet users. Most of the countries follow a straight line. USA deviates considerably from it simply because of its high Twitter penetration rate.

Next, we consider those users and their countries and study their specific cultural aspects in the sections "Pace of Life," "Individualism vs. Collectivism," and "Power Distance".

## Pace of Life

Having this data at hand, we can now start with the first dimension of our analysis: Pace of Life. This differs across countries: for example, Levine found that USA's Pace of Life is higher than Brazil's (Levine 2006). One could order countries by the value residents give to time and would see that Sioux Indians do not have a notion of time (they even do not have a word for it); Brazilians have a 'relaxed' notion of it (e.g., Levine fount that students defined 'being late' as being 33 minutes late on average); and people in USA give high importance to time up to the point of associating it with money (e.g., people experience considerable levels of stress if deadlines are not met).

Another time system was proposed by Hall (Hall and Hall 1990). He divided time perception into two categories: monochronic and polychronic time. Monochronic time refers to paying attention to only one thing at a time. In monochronic cultures, people tend to schedule their activities in a linear way, tend to be less flexible, and perceive time as a measurable, quantifiable entity, something with real weight and value. For these reasons, monochronic countries are also considered more predictable, not least because a fixed time is often allocated to any task or meeting. Such countries include United States, Germany, Switzerland, and Japan. By contrast, polychronic time refers to being involved with many things at once. In polychronic countries, people are more flexible with time, adapt their schedules to others' needs, and see time as a general guideline, something without substance or structure. Consequently, polychromic countries are less (temporarily) predictable. Such countries include France, Italy, Greece, Mexico and some Eastern and African countries.

The problem is that Hall did not provide any country scores we could use for this study but "Levine's Pace of Life research has been indirectly linked to the observations of Hall (1983) to suggest that polychronicity and Pace of Life are negatively related" (Conte and Rizzuto 1999), and that insight was recently used in the study of the scheduling tool of Doodle (Reinecke et al. 2013).

To paraphrase these ideas in the context of Twitter, we hypothesize the following relationship:

$[H1.1]$*The activities (e.g., mentions, status updates) of users in countries with higher Pace of Life are more temporarily predictable.*

To test this hypothesis, since there are several factors that influence people's routine during weekends, we leave them out and analyze activities during working days, during which the differences between monochronic and polychronic cultures are more salient (Centre for Good Governance 2011). After adjusting for the different time zones, we divide each day in five time intervals: sleeping time (00:00 - 05:59), rising time (6:00 - 8:59), working hours (9:00-17:59), dinner (18:00-20:59) and late night (21:00-23:59). This division allows us to separate working hours from the rest of the day and effectively mark changes of activities (Golder and Macy 2011). Then, to capture each user's predictability, we compute the user's entropy in those five intervals, and we choose entropy because it is often used to characterize unpredictability in time series (Song et al. 2010). Specifically, we consider a measure of entropy proposed by (Krumme 2010; Song et al. 2010) called *temporal-uncorrelated entropy* and adapt it to our context. The temporal-uncorrelated entropy calculates the tweeting randomness across time intervals for a given user and is defined by:

$$-\sum_{j=1}^{N_i} p_i(j) log_2 p_i(j) \tag{1}$$

where $p_i(j)$ is the historical probability that user $i$ posted in time interval $j$ and $N_i$ is the number of distinct time intervals in which user $i$ posted his/her tweets. The whole sum reflects the (un)predictability of user $i$ posting across all $j's$ intervals.

This metric is computed for the two main activities of posting updates tweets and mentioning others. After obtaining these entropies for all users for these two activities, we compute the Pearson product-moment correlation between the geometric average of the country-level entropy and its corresponding Pace of Life rank. Pearson's correlation $r \in$ [-1, 1] is a measure of the linear relationship between two random variables, whereby 0 indicates no correlation and +1(-1) perfect positive (negative) correlation. Table 1 and Figure (3a) summarize the results. The higher the Pace of Life (monocronic countries), the lower the tweets' temporal unpredictability ($r_{(15)} = -0.62$) and user mentions' ($r_{(15)} = -0.68$).

| | Entropy (Tweets) | Entropy (Mentions) |
|---|---|---|
| Pace of Life (overall) | -0.62** | -0.68** |
| Pace of Life (walking speed) | -0.56** | -0.61** |
| Pace of Life (post office) | -0.44 | -0.50* |
| Pace of Life (clock accuracy) | -0.45* | -0.51* |

Table 1: Pearson correlation coefficients between the entropy of the activity in twitter and three measures of the pace of life, p-values are expressed with *'s: $p < 0.05$ $(***)$, $p < 0.05$ $(**)$, and $p < 0.1$ $(*)$

Figure (3a) shows the negative relationship between unpredictability and Pace of Life. Thirteen countries follow this relationship but two do not: Japan and Indonesia. Japan's Pace of Life is "one of the most demanding on earth" (Levine 2006) (after Switzerland, Ireland and Germany), and one would thus expect to find predictable (monochromic) temporal patterns for it. Instead, we find high unpredictability, and that matches what Hall found more than 20 years ago (Hall and Hall 1987): Japan is an outlier, in that, it mixes high Pace of Life with strong polychronic characteristics, not least because of, Hall suggested, the importance attributed to social relationships. Also Indonesia (our second outlier) shows considerably higher unpredictability than the remaining countries, and that matches what Levine found when he went to one of Jakarta's post office to buy stamps: "It took us considerably longer than in many other countries to find this out." The postal clerk was more interested in conversing about Levine's life rather than fulfilling his request.

Next, we focus on working hours only. Since people in countries with high Pace of Life schedule their time in a linear way, we expect hat they would tweet less during working hours, in proportion, to avoid any interruption:
$[H1.2]$ *The percentage of a country's users who have tweeted during working hours negatively correlates with the country's Pace of Life.*

The daily fraction of users in a country who tweet during working hours does indeed negatively correlate with Pace of Life ($r_{(15)} = -0.58$).

## Individualism vs. Collectivism

In addition to pace of life, also human relationships differ across cultures. In high collectivist cultures, users tend to focus more on the community to which they belong: for example, peers tend to unconditionally support superiors' opinions. Such countries (e.g., Indonesia) are characterized by "in-groups", and their members are expected to look after each other. By contrast, people from high individualist countries like the U.S. are in a more loosely knit social network, and are generally expected to look after themselves or only after immediate family members (Hofstede and Minkov 2010).

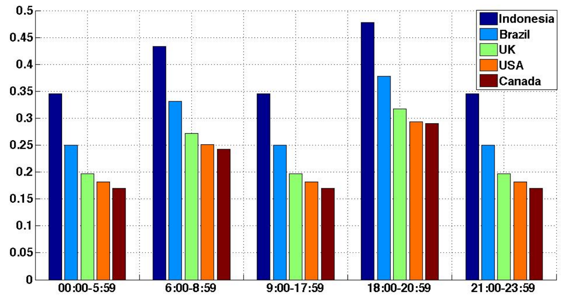


Figure 4: Fraction of users engaging with others at different times of the day. Users in Indonesia and Brazil (collectivist countries) engage with others more than those in UK, USA, and Canada (individualistic countries), and they consistently do so throughout the day.

Another characteristic that differentiate collectivist countries from individualist ones is that the former tend to adopt high-context communication as opposed to low-context. In high-context cultures, people tend to emphasize interpersonal relationships. According to Hall, these cultures prefer group harmony and consensus to individual achievement: "flowery language, humility, and elaborate apologies are typical" (Hall and Hall 1990). Also, the way people acquire information also varies between cultures. According to Hofstede *et al.*, the primary source of information is one's social network in collectivist countries, while it is (news) media in individualistic countries (Hofstede and Minkov 2010). Finally, the right to privacy is relevant in many individualist societies, while letting one's in-group invade one's private life is acceptable in collective societies.

Based on these studies, one should thus expect that people in collectivist countries will engage into public conversations *more* than what people in individualist countries do. We thus hypothesize that:

$[H2]$ *The fraction of users who mention (engage in a conversation with) others negatively correlates with Individualism.*

Using Pearson coefficients, we correlate each country's fraction of users engaging in conversations with the country's Individualism index reported by (Hofstede and Minkov 2010). We find that users in individualistic countries mention others far less than those in collectivist countries (the correlation coefficient is as high as $r_{s(30)}$ =-0.55 ($p <$ 0.005). This consistently holds at different times of the day, and Figure 4 exemplifies that by contrasting two high collectivist countries (Indonesia and Brazil) with three high individualistic countries(USA, UK, Canada). Figure (3b) then shows the correlation between lack of engagement and Individualism, which is high for all countries except for Germany. This result matches that of a previous study on microblogs: German tweets received the least number of mentions out of the 10 most common languages in Twitter (Hong, Convertino, and Chi 2011). Also, in Germany, few comments are left in blogs, and users react to comments lower than what users in less individualist countries such as Russia and China do (Mandl 2009).

## Power Distance

Hofstede defines Power Distance as the "extent to which the less powerful members of institutions and organizations within a country expect and accept that power is distributed unequally." (Hofstede and Minkov 2010). In countries comfortable with Power Distance, subordinates expect to be told what to do: employees tend to prefer to have a boss who decides autocratically (Biatas 2009). As such, hierarchy in organizations and inequalities are expected and desired, and that applies not only to work environments but also to schools and families.

The number of a user's followers (indegree) does not necessarily reflect influence but does reflect popularity (Bakshy et al. 2011; Cha et al. 2010; Kwak et al. 2010)). Therefore, the power relationship between a pair of users is leveled, if their numbers of followers are comparable; while it is imbalanced, if the numbers of followers greatly differ. Based on this observation, we posit that:

$[H3]$ *In countries comfortable with Power Distance, a pair of users who engage in any type of relationship is likely to show indegree imbalance.*

We correlate country-level indegree imbalance with corresponding Power Distance (another one of Hofstede's cultural dimensions). We find imbalance online and Power Distance offline go together for all three types of relationships (Table 2): the correlations are $r = 0.65$ for "who follows whom" relationships; $r = 0.33$ for "who recommends whom" relationships, and $r = 0.42$ for "who starts to follow whom" relationships. Figure (3c) shows the correlation for the "who follows whom" relationship. Norway, Venezuela and Indonesia are outliers. It is difficult to see why this is the case for Norway and Venezuela as there is no previous study for them in this matter. For Indonesia, instead, we found that our results match those on blogs: 27% of all blog trends in this country are about pop and celebrities, which may result in Indonesian users following more celebrities than users in other countries (Silang 2011).

Critics might rightly argue that the number of followers does not necessarily reflect one's popularity, not least because there are "spammers" who accumulate followers by subscribing to random users profiles. This was the case especially in the early years of Twitter (Mowbray 2010). To counter that criticism, in addition to indegree as proxy for popularity, we consider the ratio in-degree/out-degree, correlate the corresponding country-level popularity imbalance with Power Distance, and obtain the same correlation as when considering indegree ($r = 0.67$).

## Why It Matters

We have found strong correlations between country-level behavioral patterns on Twitter and the three cultural aspects of pace of life, individualism and power distance. Those correlations translate into being able to track these three aspects at fine-grained temporal levels - one does not need to wait for the next 10-year effort that replicates Levine's study or Hofstede's; on the contrary, by simply tracking behavioral patterns on Twitter, one could predict the three cultural aspects to a considerable extent for countries that are well represented on Twitter. However, before doing so, one may well wonder why these three aspects matter at all. To see why it would be important to use Twitter to track them, one should consider that the three cultural dimensions have been found to be associated with three main economic indicators: GDP per capita, income inequality and education expenditure (Hofstede and Minkov 2010; Wilkinson and Pickett 2010). We now test whether these economic indicators do also correlated with our three Twitter features: temporal predictability, activity levels during working hours, engagement with others, and popularity imbalance. To ease explanation, we collate the results in Table 3 and comment them next.

**GDP per capita.** High collectivism was found to be related to countries with low national wealth (Hofstede and Minkov 2010). To test whether this holds also for our Twitter features, we get hold of the Gross Domestic Product values for our 30 countries (these values reflect purchasing power normalized by population) and correlate them with our four Twitter features. We find that GDP is associated with three features in a statistically significant way (first row in Table 3): low-GDP countries tend to be temporally unpredictable ($r = 0.55$), be active during working hours ($r = -0.57$), and feel comfortable with popularity imbalance ($r = -0.48$). Figure (5a) shows this relationship and associated outliers - Japan and Singapore. The result for Japan is explained by what we found in the Section "Pace of Life." The result for Singapore is explained by considering that the country has faced rapid economic growth rate and is thus "highly developed and enjoys remarkably open and corruption-free environment, stable prices, and a per capita GDP higher than that of most developed countries" [1]. At the same time, however, it preserves high collective characteristics typical of most Asian countries, and that explains the association of higher entropy with GDP.

**Education Expenditure.** We correlate education expenditure (as percentage of GDP) with our four Twitter features and find that countries with low education expenditure are characterized by the same features as countries with low GPD, even if expenditure is normalized by it. They are (second row in Table 3): temporally unpredictable ($r = 0.58$), be active during working hours ($r = -0.51$), and feel comfortable with popularity imbalance ($r = -0.60$). Figure (5b) depicts high correlation for most countries. Again, Indonesia is an outlier, as one would expect from the previous results.

**Income Inequality.** Power distance was found to be related to the use of violence in domestic politics and to income inequality(Hofstede and Minkov 2010). One widely-used way to measure income inequality is the *Gini* coefficient. This measures the degree of inequality in the distribution of family income in a country (Wilkinson and Pickett 2010).The lower its value, the more equal a society is. We find that unequal countries tend to be (third row in Table 3): temporally unpredictable ($r = -0.53$), be active during working hours ($r = 0.49$), and feel comfortable with popularity imbalance ($r = 0.58$). It should come as no surprise this last result: that the strongest predictor of income inequality is popularity imbalance (popularity inequality) in Twitter. Figure (5c) shows that India, Philippines, Singapore and Brazil are outliers. That is because these countries are characterized by disproportionately high levels of inequality (Wilkinson and Pickett 2010).

## Discussion

Social media sites often assume that people from different countries use their services in very similar ways. By contrast, we find that the use of Twitter considerably changes across them. Fortunately, these changes are not random but are *predictable* so much so that simple country-level behavioral features derived from Twitter strongly correlate with cultural dimensions. Users in monochronic countries tend to be temporarily predictable, those in collectivist countries considerably talk with each other, and those in countries uncomfortable with power distance will not preferentially engage only with popular users. These findings might not only have theoretical implications for future cross-cultural studies but might also have practical implications, including the prediction of country-level economic indicators at fine-grained temporal level, and the design of culture aware recommender system.

**Theoretical Implications.** Twitter is a distal communication modality (distal in the sense that users are separated in space and time), and it has been argued that it is not a social-networking tool but a broadcasting platform of, for example, news and opinions (Kwak et al. 2010). Yet, our cultural analysis suggests that Twitter enjoys social-networking features, and that engagement is predominant among users in collectivist countries. This study not only has suggested the extent to which Twitter use is associated with specific culture dimensions, but also points to the possibility that social media sites could be

[1]Information taken from the Central Intelligence Agency of USA

| Pace of Life | Correlation |
|---|---|
| $[H1.1]$*The activities (e.g, mentions, status updates) of users in countries with higher Pace of Life are more temporarily predictable* | $r_{(15)} = -0.62$** <br> $r_{(15)} = -0.68$** |
| $[H1.2]$*The percentage of a country's users who have tweeted during working hours negatively correlates with the country's Pace of Life* | $r_{(15)} = -0.58$** |
| **Individualism** | **Correlation** |
| $[H2]$ *The fraction of users who mention (engage in a conversation with) others negatively correlates with Individualism index* | $r_{s(30)} = -0.55$*** |
| **Power Distance** | **Correlation** |
| $[H3]$ *In countries comfortable with Power Distance, a pair of users who engage in any type of relationship is likely to show indegree imbalance* | $r_{(30)} = 0.62$*** <br> $r_{(30)} = 0.33$* <br> $r_{(30)} = 0.42$** |

Table 2: Pearson correlation coefficients: (*H1.1*) between Pace of Life and the temporal predictability of users' activity (mentions and tweets); (*H1.2*) between Pace of Life and the percentage's of a country's users tweeting during working hours; (*H2*) between Individualism and the fraction of users engaged with others; and (*H3*) between Power Distance and in-degree imbalance shown in three types of relationships ("who follows whom", "who recommends whom" and "who starts to follow whom"). $p$-values are expressed with *'s: $p < 0.005$ $(***)$, $p < 0.05$ $(**)$, and $p < 0.1$ $(*)$.

| Indicator | $[HP1.1]$ **Predictability** | $[HP1.2]$ **Users (%) in w. hours** | $[HP2]$ **Mentions** | $[HP3]$ **Imbalance** |
|---|---|---|---|---|
| GDP per capita | $\mathbf{r_{(30)} = 0.55}$*** | $\mathbf{r_{(30)} = -0.57}$** | $r_{(30)} = -0.41$* | $\mathbf{r_{(30)} = -0.48}$** |
| Education | $\mathbf{r_{(30)} = 0.58}$*** | $\mathbf{r_{(30)} = -0.51}$*** | $r_{(30)} = -0.24$ | $\mathbf{r_{(30)} = -0.60}$*** |
| Inequality | $\mathbf{r_{(30)} = -0.53}$*** | $\mathbf{r_{(30)} = 0.49}$** | $r_{(30)} = 0.39$* | $\mathbf{r_{(30)} = 0.58}$*** |

Table 3: Pearson correlation coefficients of three socio-economic indicators (first column) with: predictability (second column), activity in working hours (third column), mentions (fourth column) and in-degree imbalance (fifth column). $p$-values are expressed with *'s: $p < 0.005$ $(***)$, $p < 0.05$ $(**)$, and $p < 0.1$ $(*)$.

used to run large-scale cross-cultural studies and could ultimately become tools that promote computational social science. This is is a new discipline that aims at using large archives of naturalistically-created behavioral data (of, for example, emails, tweets, Facebook contacts) to answer social science questions (Lazer, D. et al. 2009; NSF 2011).

**Practical Implications.** Our findings could also be used to design:

*Culture-aware engagement tools.* In collectivist countries, users do engage with each other by exchanging messages and recommending others. One could design country-tailored tools that: promote interactions with strangers in individualistic countries, and with strong ties in collectivist countries; rank status updates based on interestingness in small-power-distance countries, and on popularity in large-power-distance; targets ads in specific time of the day for monochronic countries, and in user-tailored for polychronic countries.

*Culture-aware people recommender.* One increasingly important feature in Twitter is its people recommender system, which suggests people one might know. This tool makes suggestions based on structural features (e.g., common followers) and on content features (e.g., matching one's topics of interest). However, the tool might well benefit from cultural dimensions as well: recommending strangers is fine in individualist countries but not in collectivist ones; or users in large-power-distance are likely to preferentially follow highly-popular users.

**Limitations and Future Work.** Despite the strong correlations, this study suffers from five limitations. First, our sample was collected in a specific time frame. Critics might rightly say that our findings may be co-founded by the days data was crawled. However, the sample spans 10 weeks and, as such, it might be large enough to capture the normal routine of users. Second, we might run the risk to promote stereotyping of individuals based on their countries of origin. This study is about 'mean behavior', and one should consider that there is high variability across individuals in the same country. Third, we naively equated use of mentions with "engagement with others", but that might not be necessarily the case. That is why, in the future, it might be beneficial to propose a taxonomy that will distinguish one's purposes when mentioning others (i.e., conversational, informative, attribution). Fourth, this has been an exploratory study in which causal inference has not been established (and it was not the aim of the study). However, there are two remarks to be made: a) many of the observed relations on Twitter confirm those that are already known in the real world; and b) some of the correlations are weak, but oth-

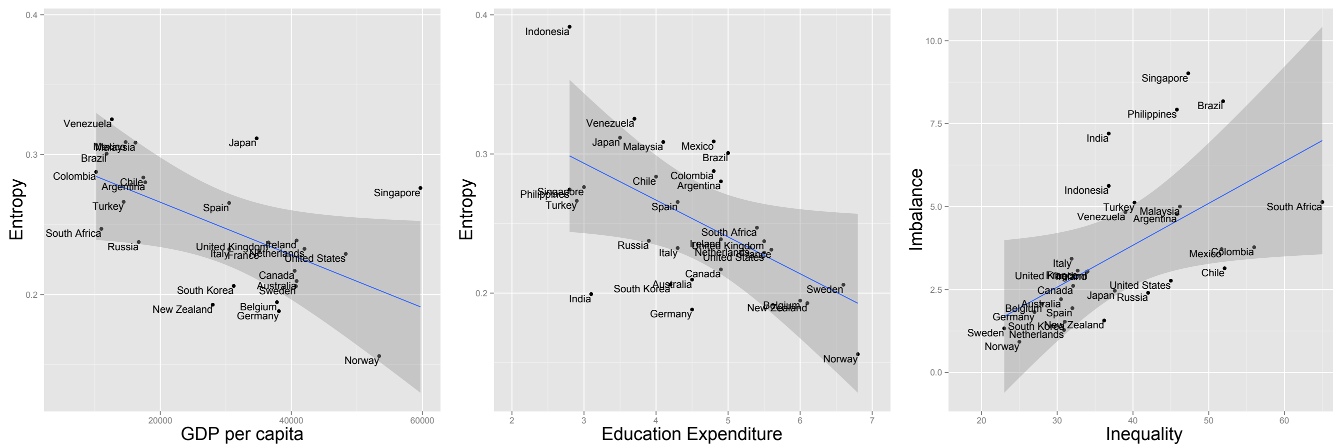


(a) Entropy vs. GDP

(b) Entropy vs. Education Expenditure

(c) Indegree vs. Inequality

Figure 5: The relationships between Twitter features and socio-economic indicators

ers are very strong, suggesting a dose-response form *from* country characteristic *to* behavior on Twitter. Fifth, we have focused on language-independent features. In the future, we will explore how the use of language changes depending on cultural dimensions (Henrich, Heine, and Norenzayan 2003) - for example, do individualistic countries use more singular first-person pronouns (e.g., I, my, mine, yo, eu, moi)?

## Acknowledgments

We thank Eduardo Graells for his beautiful maps, and Amin Mantrach and Neil Ohare for their feedbacks. This research was partially supported by European Community's Seventh Framework Programme FP7/2007- 2013 under the ARCOMEM project; by the Spanish Centre for the Development of Industrial Technology under the CENIT program, project CEN-20101037 (www.cenitsocialmedia.es) "Social Media"; and by the EU SocialSensor FP7 project (contract no. 287975).

## References

Baeza-Yates, R. A.; Middleton, C.; and Castillo, C. 2009. The geographical life of search. In *Web Intelligence*.

Bakshy, E.; Hofman, J. M.; Mason, W. A.; and Watts, D. J. 2011. Everyone's an influencer: quantifying influence on twitter. In *Proceedings of the fourth ACM International Conference on Web Search and Data Mining (WSDM)*.

Biatas, S. 2009. Power distance as a determinant of relations between managers and employees in the enterprises with foreign capital. *Journal of Intercultural Management*.

Borra, E., and Weber, I. 2012. Political insights: Exploring partisanship in web search queries. *First Monday*.

Centre for Good Governance. 2011. Soft skills for public managers: Handbook of time management skills.

Cha, M.; Haddadi, H.; Benevenuto, F.; and Gummadi, K. P. 2010. Measuring User Influence in Twitter: The Million Follower Fallacy. In *Association for the Advancement of Artificial Intelligence*.

Conte, J. M., and Rizzuto, T. E. 1999. A construct-oriented analysis of individual-level polychronicity. *Journal of Managerial Psychology*.

Golder, S. A., and Macy, M. W. 2011. Diurnal and Seasonal Mood Vary with Work, Sleep, and Daylength Across Diverse Cultures. *Science*.

Hall, E., and Hall, M. 1987. *Hidden differences: doing business with the Japanese*. Anchor Press/Doubleday.

Hall, E., and Hall, M. 1990. *Understanding Cultural Differences*. Intercultural Press.

Henrich, J.; Heine, S.; and Norenzayan, A. 2003. Psychological aspects of natural language. use: our words, our selves. *Annual Review Psychology*.

Hochman, N., and Schwarts, R. 2012. Visualizing instagram: Tracing cultural visual rhythms. In *Proc. 6th International AAAI Conference on Weblogs and Social Media (ICWSM)*.

Hofstede, G., and Minkov, M. 2010. *Cultures and Organizations: Software of the Mind, Third Edition*. McGraw-Hill.

Hong, L.; Convertino, G.; and Chi, E. H. 2011. Language matters in twitter: A large scale study. In *The 5th International AAAI Conference on Weblogs and Social Media (ICWSM)*.

Kling, C. C., and Gottron, T. 2011. Detecting culture in coordinates: cultural areas in social media. In *International Workshop on DETecting and Exploiting Cultural diversiTy on the Social Web (DETECT)*. ACM.

Krumme, K. 2010. *How Predictable: Patterns of Human Economic Behavior in the Wild*. Massachusetts Institute of Technology, School of Architecture and Planning, Program in Media Arts and Sciences.

Kwak, H.; Lee, C.; Park, H.; and Moon, S. 2010. What is twitter, a social network or a news media? In *Proc. 19th ACM International Conference on World Wide Web (WWW)*.

Lazer, D. et al. 2009. Computational Social Science. *Science*.

Lee, K.; Eoff, B.; and Caverlee, J. 2011. Seven months with the devils: A long-term study of content polluters on twitter. In *Proc. The 5th International AAAI Conference on Weblogs and Social Media (ICWSM)*.

Levine, R. 2006. *A Geography of Time: The Temporal Misadventures of a Social Psychologist or How Every Culture Keeps Time Just a Little Bit Differently.* University Press.

Mandl, T. 2009. Comparing Chinese and German blogs. In *Proceedings of the 20th ACM Conference on Hypertext and Hypermedia*.

Mowbray, M. 2010. The twittering machine. In *Proc. 6th International Conference on Web Information Systems and Technologies (WEBIST)*.

2011. *Rebuilding the Mosaic*. National Science Foundation.

Poblete, B.; Garcia, R.; Mendoza, M.; and Jaimes, A. 2011. Do All Birds Tweet the Same? Characterizing Twitter Around the World. In *Proc. 20th ACM Conference on Information and Knowledge Management (CIKM)*.

Reinecke, K.; Nguyen, M. K.; Bernstein, A.; Naf, M.; and Gajos, K. Z. 2013. Doodle around the world: Online scheduling behavior reflects cultural differences in time perception and group decision-making. In *Proc. 16th ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW)*.

Silang, S. 2011. Snapshot of Indonesia Social Media Users (2011).

Song, C.; Qu, Z.; Blumm, N.; and Barabási, A.-L. 2010. Limits of predictability in human mobility. *Science*.

Wang, X.; Xu, S.; Peng, L.; Wang, Z.; Wang, C.; Zhang, C.; and Wang, X. 2012. Exploring scientists' working timetable: Do scientists often work overtime? *Journal of Informetrics*.

Wilkinson, R., and Pickett, K. 2010. *The Spirit Level: Why Equality is Better for Everyone*. Penguin Books.

Woods, W. 2003. *Tick Tock!...Who Broke the Clock?* Innovations International.

Yanai, K.; Yaegashi, K.; and Qiu, B. 2009. Detecting cultural differences using consumer-generated geotagged photos. In *2nd International Workshop on Location and the Web (LOCWEB) in CHI*.