

Journal Pre-proof

Later bedtimes predict President Trump's performance

Douglas Almond, Xinming Du

PII: S0165-1765(20)30355-4

DOI: <https://doi.org/10.1016/j.econlet.2020.109590>

Reference: ECOLET 109590

To appear in: *Economics Letters*

Received date: 26 July 2020

Revised date: 15 September 2020

Accepted date: 16 September 2020



Please cite this article as: D. Almond and X. Du, Later bedtimes predict President Trump's performance. *Economics Letters* (2020), doi: <https://doi.org/10.1016/j.econlet.2020.109590>.

This is a PDF file of an article that has undergone enhancements after acceptance, such as the addition of a cover page and metadata, and formatting for readability, but it is not yet the definitive version of record. This version will undergo additional copyediting, typesetting and review before it is published in its final form, but we are providing this version to give early visibility of the article. Please note that, during the production process, errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

© 2020 Published by Elsevier B.V.

Later Bedtimes Predict President Trump's Performance*

Douglas Almond[†] and Xinming Du[‡]

September 15, 2020

do not cite or distribute

Abstract

Technology and social media use are increasingly associated with delays in nightly sleep. Here, we consider the timing of President Trump's official Twitter account posts as a proxy for sleep duration and how it relates to his public performance. The President wakes around 6am, a routine which has not changed since early 2017. In contrast, the frequency of Twitter activity **11pm-2am increased 317%** from under one day per week in 2017 to three days a week in 2020. The President's increased late-night activity is not accounted for by increases in the frequency of his use of social media over time, his travel schedule, or seasonality. On the day following one where he posts late at night, his Twitter followers interact less with his posts, described as "official statements by the President of the United States."¹ He receives 7,400 fewer likes per tweet, 1,300 fewer retweets per tweet, and 1,400 fewer replies per tweet after a late night (drops of 6.5%-8%).

Tweets aside, the President's speeches and interview transcripts have previously been coded for their dominant emotion through text analysis. On the day following a late night, the President's inferred emotion is less likely to be "happy" and nearly three times more likely to be "angry" in his interviews and speeches. Finally, the 2020 election odds of the President's chief opponent also increase after a late night, while the President's are unchanged. The pattern we document is consistent with a progressive shortening of the President's sleep over his first term and compromised performance from sleep deprivation.

*We thank Victoria Li and Diane Huang for research assistance. Danny Bressler, Yi Cheng, Olle Folke, Osea Giuntella, Bhashkar Mazumder, Jeff Shrader, and Michael Ting provided helpful comments and Solveig Asplund proof reading. We thank Suresh Naidu for suggesting we analyze betting markets (Table 4). Any errors are our own.

[†]Columbia University and NBER: da2152@columbia.edu

[‡]Columbia University: xd2197@columbia.edu

¹White House Press Secretary Sean Spicer, June 2017. In July 2019, the Second Circuit Court of Appeals decided unanimously that the President's Twitter account is "official".

1 Introduction and Background

*Sleep that knits up the ravell'd sleeve of care,
The death of each day's life, sore labour's bath,
Balm of hurt minds, great nature's second course,
Chief nourisher in life's feast.*

Shakespeare's Macbeth

Sleep impacts neurobehavioral performance. While historically neglected by social scientists, recent studies have shown a diverse range of sleep deprivation consequences, including to labor market outcomes. These observational findings are consistent with more short-term impacts from randomized control trials, the gold standard for empirical evidence. Given sleep's demonstrated importance and furthermore its malleability, the historical neglect by researchers of sleep is surprising. This neglect may stem, at least in part, from the practical obstacle that nightly sleep duration is typically unobserved by researchers.

Our analysis of a single individual's sleep and performance is nevertheless informative to future research. First, there is a large literature in political economy on politician quality, reviewed by Dal Bó and Finan (2018), to which we add new, high-frequency measures (see also Section 5). Second, technology and social media use appear to be delaying the onset of nightly sleep generally (Exelmans and Van den Bulck, 2016; Bhat et al., 2018; Scott and Woods, 2019). In light of previous design-based findings on sleep's impacts, this secular trend has costs for both the economy and population health. To date, the observational literature on sleep has focussed more on daylight savings policies and time zone boundaries. By contrast, we focus on a commonly-experienced and growing source of sleep curtailment. Third, such nighttime technology use is frequently observed *publicly* via social media use and can proxy for sleep duration. This constitutes a "big data" opportunity for researchers to consider the traditionally neglected role of sleep on (multidimensional) performance. While late-night social media use has already been successfully deployed to consider NBA player performance (Jones et al., 2019), it is of more general interest to social scientists to consider various non-athletic outcomes.² Finally, regular use of public social media alongside daily performance data are increasingly common and available for elected leaders, making large-sample analyses in political economy feasible. Harnessing Trump's frequent use of social media to proxy for sleep and diurnal activity can be viewed as providing "proof of concept" for assessing the relationship between sleep and worker performance in an increasingly digital age.

President Trump's frequent use of social media provides an unprecedented window into one of the world's most impactful sleepers. We follow previous research which has used the timing of social media activity as a proxy for inferring sleep schedules, e.g. Golder and Macy (2011); Jones et al. (2019). We make comparisons across the roughly 1,200 nights since the President's inauguration, inferring which nights he

²For example, one could systematically relate sleep proxies to the tenor of public statements by business leaders through the evolving tools of "text as data" and sentiment analysis. Sleep-instrumented changes in public statements could then be related to asset prices for the relevant firms.

slept less. The President is thus compared only to himself, as large variation in the (unobserved) sleep needs across people would suggest.

If we assume that the President was sleeping around his optimal (albeit short) amount early in his Presidency, does this remain the case in 2020? Our sleep proxy indicates this is decidedly not so: he now sleeps substantially less than in 2017. To our knowledge, we are the first to use nightly data to argue that President's sleep has become noticeably shorter. This trend in sleep motivates our primary analysis: how does the President's daily public performance vary with his inferred sleep on the previous night? This research question merits investigation as the President's bedtime and sleep duration are choices, at least in part. The President's sleep duration does not appear fixed even within his first term, so it is possible that it could change again and perhaps beneficially.

Beyond proxying for his nightly sleep, we will argue that President Trump's frequent Twitter posts also furnish meaningful performance measures. In prohibiting the President from blocking unwanted followers, the 2nd Circuit Court of Appeals concluded in July 2019 his Twitter use was "government speech" and that "the President has consistently used the [Twitter] Account as an important tool of governance and executive outreach". The National Archives and Records Administration concluded his tweets are "official records that must be preserved under the Presidential Records Act". The White House has described his posts as "official statements by the President of the United States."³ Speaking for himself, the President tweeted in July of 2017: "My use of social media is not Presidential - it's MODERN DAY PRESIDENTIAL".

We consider four sets of performance measures and their relationship to nightly sleep. First, we consider the quality of the President's tweets. Because quality assessment can be subjective, we use the President's Twitter followers to arbitrate quality by gauging their reactions to his frequent postings. The US Court of Appeals, Second Circuit noted that the President:

...uses the 'like', 'retweet', 'reply,' and other functions of the Account to understand and to evaluate the public's reaction to what he says and does.

The President has likewise revealed that he monitors these follower interaction metrics.⁴ For the roughly 11,000 Twitter posts the President has made since his inauguration, we analyze the count of likes, retweets, and replies. Second, we use the Washington Post's Fact Checker database to consider the veracity of the President's statements. Third, we consider the text of the President's speeches and interview transcripts: some 1,950 interviews and transcripts since inauguration. These presidential statements are coded independently by Factbase for their

³White House Press Secretary Sean Spicer, June 2017.

⁴For example, on monitoring his followers' response to his tweets, *Quartz* quoted the President:

"It used to go up, it would say 7,000, 7,008, 7,017, 7,024, 7,032, 7,044, right? Now it goes 7,000, 7,008, 6,998," he said. "Does anyone know what I'm talking about with this? I never had that before. I used to watch it. It would be like a rocket ship when I put out a beauty."

dominant emotion using text analysis. We will consider whether the emotional content of his non-Twitter statements varies with his sleep duration on the previous night, following the literature on the large and immediate neurobehavioral impacts of sleep. Finally, we analyze daily betting markets for the 2020 presidential election and whether these change systematically after a late night.

2 Related Literature

Sleep deprivation has been definitively shown to impair performance. In a seminal study, Van Dongen et al. (2003) randomized 4, 6, or 8 hours of time in bed per night. Sleeping 6 hours or fewer per night produced declines in cognitive performance equivalent to roughly 2 nights of total sleep deprivation. Van Dongen et al. (2003) concluded that “even relatively moderate sleep restriction can seriously impair waking neurobehavioral functions in healthy adults”. Design-based observational studies yield similar findings. For example, Gibson and Shrader (2018) use variation in sunset time and find that it impacts both sleep duration and labor productivity. The impact of long-term (persistent) sleep deprivation on earnings is about 5 times as large as the effect of short-run sleep deprivation. Sleep loss has also been found to impair self-control and compromise other neurobehavioral outcomes for all age groups (Christian and Ellis, 2011; Barnes et al., 2011; Pilcher et al., 2015; Mai et al., 2019) and is associated with poorer memory, lower attentional capacity, worse cognitive skills and higher risk of incident dementia among elderly adults (Sterniczuk et al., 2013; Richards et al., 2017; Wams et al., 2017; Sabeti et al., 2018).

Despite these effects, decisions governing sleep may not be obvious to the individual. Van Dongen et al. (2003) found that self-reported sleepiness scores did not correspond well with cognitive effects, suggesting that subjects were unaware of the cognitive deficits. This disconnect “may explain why the impact of chronic sleep restriction on waking cognitive functions is often assumed to be benign” (Van Dongen et al., 2003). Additionally, neoclassical models of sleep choice – beginning with Biddle and Hamermesh (1990) – typically assume that increasing the opportunity cost of sleep reduces the optimal amount of sleep for the individual *ceteris paribus*. The opportunity costs of sleep may be especially high for executives and leaders.⁵ Turning to behavioral economics, given widespread use of personal electronics and social media at night, issues with self-control may cause people to sleep less than they would were sleep-commitment devices available (Avery, Giuntella, and Jiao, 2019). Finally, older individuals may experience decreased sleep duration and lower sleep quality for biological reasons (Scullin and Bliwise, 2015; Li et al., 2018).

Outside of lab settings, direct measures of sleep are typically unavailable.⁶ Unfortunately, we do not have a direct measure of sleep duration in this

⁵It is less clear that the President’s opportunity cost has increased during his first term. While the COVID-19 crisis arguably raises the opportunity cost of sleep, most of the President’s increased late-night tweeting occurred before it: 2019 late-night tweeting is higher than 2018 late-night tweeting which in turn is higher than 2017 late-night tweeting.

⁶Avery, Giuntella, and Jiao (2019); Bessone, Rao, Schilbach, Schofield, and Toma (2020) are notable exceptions.

observational study. Instead we adopt a proxy previously used in the sleep literature: late night social media activity. Bedtime social media use directly delays sleep time, shortens sleep duration, and worsens sleep quality (Exelmans and Van den Bulck, 2016; Bhat et al., 2018; Scott and Woods, 2019). Berger et al. (2019) found that smartphone touchscreen time is strongly-correlated with more direct measure of sleep (wrist-worn accelerometers) for sleep-onset time.

Turning to sleep and performance, Jones et al. (2019) evaluate the late-night tweeting of 112 NBA players, finding that shooting accuracy, points scored, and rebounds are lower the day after late-night tweeting. They interpret these results as reflecting the effect of shortened sleep. Leypunskiy et al. (2018) use the nighttime lull in Twitter activity as a proxy for users' sleep time. They find this lull shifts to later times on weekends relative to weekdays and that "social jet lag" is the lowest over school holidays. Golder and Macy (2011) relate the diurnal cycle of tweeting to the sentiment of tweets. Those active late at night ("night owls") have the lowest negative feeling in the morning but their negative sentiment builds up to a nighttime peak. Golder and Macy (2011) conclude that people are emotionally refreshed by sleep.

The President's tweeting has been the subject of research even before he became a presidential candidate. In focussing on sleep during his Presidency, Kryger (2017)'s study is the most relevant to our own. Using the President's tweets over six months in 2017, Kryger (2017) calculated the average number of tweets per hour (regardless of local time zone) and found that he is quiet online for about five hours a night. Kryger (2017) concluded that the President is a short sleeper and is probably sleep deprived. We will argue below that the President's late-night schedule has changed dramatically since 2017.

Sleep aside, the President's tweets are consequential. This is not surprising given that "the official nature of the [Twitter] Account is overwhelming" (2nd Circuit, US Court of Appeals: July 9, 2019). The President's tweets have been shown to influence financial markets (Bianchi et al., 2019; Ge et al., 2019), and inflame anti-minority sentiment (Hobbs and Lajevardi, 2019; Muller and Schwarz, 2019), among other consequences. Other studies emphasize the President uses Twitter strategically for political objectives. For example, he tries to shape general public opinion (Miles and Haider-Markel, 2018) and tends to tweet more when he has recently garnered less attention in news coverage (Wells et al., 2020). The President's tweets often mention other countries' purported violations of international norms and laws compared to other presidents (Carnegie and Carson, 2019). Through automated text analysis, "text as data" studies find his tweets include "unpleasant" contents (Whissell, 2018), false claims (Ross and Rivers, 2018), and negative sentiments about women (Scotto di Carlo, 2020).

3 Data and Estimation

Sleep schedules of national leaders are typically private information. Tweets have the advantage of including a time stamp, being publicly available, and being

linkable to an individual. We describe our sleep proxy and next-day performance measures below.

3.1 Data sources

Using Twitter’s application programming interface (API), we scraped 36,148 tweets posted by @realDonaldTrump from 2009 to April 10, 2020. The data include the date, time and text of each tweet. Additionally, we observe the number of likes, comments and retweets each tweet received. We do not observe the location from which tweets were posted.

We merge the President’s tweets with his location as taken from his public schedule as maintained by Factbase.⁷ This public schedule records the President’s press briefings, pool call time, “executive time” and other public events beginning January 24, 2017 (four days after inauguration). His entourage and family members are inconsistently recorded. We assume no location changes occur unless documented in the schedule. In addition, we obtain President Trump’s false claims from the Fact Checker database maintained by the Washington Post (clickable link).⁸ The Fact Checker also provides the President’s speech and interview transcripts (link) and documents word count and dominant emotion for each transcript. We use the assigned emotion of transcripts as additional performance outcomes.

Finally, we assess daily betting odds for the 2020 presidential election. Odds data are provided by BetData, which tracks odds for 105 potential candidates since November 2016. We use the implied likelihood of winning for Trump and for his strongest competitor – the candidate with the highest likelihood on each day – as dependent variables.

3.2 Sleep Proxy

We assume the President is awake when there is a Tweet from his official account.⁹ We assign the local time zones to his tweets (using the scheduling data) and identify nights on which he stays up “late”.

The President’s Twitter activity is higher from 6am to 11pm and lower from 11pm to 6am. We deem tweets posted 11pm-2am “late-night” tweets (and 6am-11pm as daytime tweets). 11pm is our threshold for “late” because of the Centers for Disease

⁷<https://factba.se/topic/calendar>

⁸Unfortunately, the vast majority of false claims cannot be merged to false tweets and their timing. The database includes a quotation of the claim, its topic, source (news conference, Twitter, interview, speech, etc.), claim date, and a short analysis on the false or misleading contents. We use claims made on Twitter and assign tweets as false in our tweet sample by merging on the quoted text and claim date.

⁹The authorship of @realDonaldTrump tweets has been widely discussed in the press, for example by Robert Draper, April 16, 2018 in *The New York Times*. *Wired* argued that tweets “sent between 6pm and 10am” are particularly likely to be written by the President himself. Insofar as our nighttime activity treatment measure “late” is concerned, we believe we are getting the President. One proxy for authorship is that staffers are more likely to use hashtags (Andrew, 2017). In Supplementary Material Section 3.3, we drop the roughly 10% of tweets with a hashtag and find similar (and if anything, slightly stronger) results.

Control and Prevention (CDC) guidelines for adequate sleep. The CDC recommends that adults over age 65 get 7-8 hours of sleep per night. Taking the minimum of 7 hours and the fact the President routinely wakes at 6am gives an 11pm start time for a “late night”. We perform robustness checks with a wider window of tweeting between 10pm and 5am in Supplementary Material Section 3.4. We proxy insufficient sleep with both the occurrence and count of late-night tweets on each night.¹⁰

On an average night where the President tweets late, his inferred sleep is below 7 hours. The average time of the President’s late-night tweeting to date is 11:55pm. Thus, his average sleep is inferred to be 6 hours or fewer when he tweets “late”.

To depict his bedtime pattern over time, we calculate the weekly proportion of late-tweeting nights. The weekly proportion is defined as the total number of late-nights per week (0-7 in integers) divided by 7. We further calculate a monthly moving average of the weekly proportion to smooth out noise.

3.3 Estimation

We apply a linear regression model to analyze the President’s likelihood of late-night tweeting and multiple factors that might govern it, and also to assess the overall predictability of late-night tweeting:

$$\text{Late}_t = \alpha_0 + \beta_1 \#Tweets_t + \beta_2 \text{MarApr} * Y2020_t + \text{Year}_t + \text{Month}_t + \text{DOW}_t + u_t \quad (1)$$

where t is the unit of analysis and indicates the number of days into the Trump Presidency. Late_t is either the extensive margin or the number of tweets between 11pm on date t to 2am on date $t + 1$. When the dependent variable is a binary indicator for late-night tweeting (extensive margin), (1) is a linear probability model.

Our independent variables include Year_t dummy variables that capture changes in the annual average of late tweeting relative to the omitted year (2017), a binary variable $\text{MarApr} * Y2020_t$ that is 1 for March and April 2020 and 0 otherwise, and $\#Tweets_t$ which equals the total number of daytime tweets posted 6:30am-9:45pm on date t . The coefficient β_1 captures how his daytime social media activity-level predicts his night tweeting.¹¹ If late-night tweeting reflects the continuation of a busy day of tweeting, then we expect $\beta_1 > 0$. We also replace the linear control for $\#Tweets_t$ with dummy variables for each possible number of daily tweets (ranging

¹⁰We drop retweets without any text added by the President, i.e. we use original tweets and retweets with text in our main analysis. We bring back retweets without added text to infer the timing of sleep in Supplementary Material Section 3.2 and find similar results. Retweets without added text suggest a more severe curtailment of the President’s sleep over time.

¹¹6:30am is used as the start of daytime because the President is not particularly twitter-active between his wake-up time and 6:30am, as shown in Figure 1 bottom panel. He tweets about 0.2-0.4 times every half hour *after* 6:30am (except for the upward spike in the morning). We see a similar ascent using 15-minute intervals, i.e. a larger number of tweets 6:30-6:45 than 6:15-6:30am. In the evening, he returns to the White House at 10:04pm on average (from his public calendar). This is consistent with the smaller number of tweets after 10pm in Figure 1. We consider tweets before 9:45pm as better capturing his daytime social media activity, i.e. tweets 9:45-11pm do not contribute very much.

from 0 to a maximum of 31 tweets per day). β_2 measures the change in late-night tweeting in March and April 2020 when COVID-19 deaths were increasing in the US (the first US COVID-19 deaths were reported on February 29). We control for month fixed effects and day of week fixed effects, noting that they are common across the four years 2017-2020.

We estimate the relationship between late-night tweeting and his tweets after 6am the following day with the following specification:

$$\#Interactions_{t+1} = \alpha + \theta_1 Late_t + \theta_2 \#Tweets_t + Year_t + Month_t + DOW_t + u_t \quad (2)$$

$\#Interactions_{t+1}$ denotes the number of interactions with the President's tweets posted between 6am to 11pm on date $t + 1$. $\#Interactions_{t+1}$ include the number of likes per tweet, the number of retweets, and the number of replies per tweet. The independent variables $Late_t$ and $\#Tweets_t$ are the same as those in Equation (1). We also control for year, month and day of week fixed effects. We include a linear trend in days to account for cognitive decline, behavioral changes, etc. over time – regression results are very similar without the linear trend. (Results with a quadratic time trend are in Tables S11-S13 of the appendix.) We address potential serial correlation in greater detail in Supplementary Material Section 4.

We evaluate additional dimensions of the President's tweet behavior after sleeping by replacing $\#Interactions_{t+1}$ with additional outcome variables including the number of tweets posted, the absolute number and the proportion of false tweets, and in Supplementary Material Table S4, the average sentiment of daytime tweets.¹²

Apart from Twitter performance, we analyze the dominant emotion of his transcripts on the following day and use the daily proportion of happy, fear(some), and angry dominant emotions in place of $\#Interactions_{t+1}$. For example, we assign happy dummy for each transcript which equals one if its dominant emotion is coded as happy by the Fact Checker and zero if not. We use the President's word count in each transcript as weights and calculate weighted sum of happy dummies on each day. We define the daily proportion of happy transcripts as the weighted sum divided by the number of transcripts. With the daily proportion on the left-hand side, we use weighted least square to estimate the relationship with the maximum word count on day $t + 1$ as regression weights to address different precision of measurement across days.¹³ We use an analogous method for the other dominant emotions.

¹²We use the "off the shelf" dictionary in the Python Vader Sentiment Library to calculate the sentiment of each tweet. The library is constructed by assigning "positive", "neutral" or "negative" to commonly used keywords. A team of ten people were asked to evaluate each keyword using a rating scale from -4 (most negative) to 4 (most positive). After that, their responses are averaged and each keyword was polarized to these three categories. The sentiment of a tweet is calculated by the proportion of "positive", "neutral" and "negative" words.

¹³The coded dominant emotion could better capture Trump's emotion if he speaks more. That is, we have a better measure of his emotion in transcripts where his word count is higher. After converting transcript-level emotion to the daily proportion, we think the precision across days should be captured by the longest transcript on each day. For example, if he said 1 word and 99 words in two events on day 1, 50 words and 50 words in two events on day 2, we think the daily measure of dominant emotion is more precise on day 1 than that on day 2.

Finally, we study betting odds on the 2020 presidential election. We replace $\#Interactions_{t+1}$ with the implied likelihood of Trump's winning and that of the leading candidate among his competitors. Betting markets have the virtue of providing a summary, contemporaneous, and "skin in the game" metric of perceived quality. Perhaps unsurprisingly, betting markets generally adjust quickly to politicians' behaviors and the perception thereof. For example, betting odds shifted substantially after the first U.S. presidential debate in September 2016 (Wolfers and Zitzewitz, 2016), as well as when FBI Director James B. Comey made his announcement in October 2016 regarding the status of Hilary Clinton's private email investigation (Halcoussis et al., 2020). While presidential betting markets may seem novel, they have been "large and well-organized markets for betting on presidential elections" stretching back to at least 1868 (Rhode and Strumpf, 2004).

4 Results

Figures on tweet timing

Figure 1 plots the average number of tweets posted at each half hour of the day from January 24, 2017 to April 10, 2020.¹⁴ We use local time zones based on his scheduled locations and plot the average separately by year. As noted above, tweeting consistently starts around 6am and reaches a peak level of one tweet per hour at 8am, before the "In-House Pool Call Time" in the morning. He then continues to tweet about one tweet every two to three hours for the rest of the day.¹⁵

The bottom panel of Figure 1 focuses on the number of tweets posted between 10pm to 7am, when the President is relatively quiet on Twitter. The yearly lines are clustered together around his 6am wake-up time, indicating stability over time. In contrast, the yearly lines "feather out" around his bedtime. This bedtime divergence is monotonic in year – late-night tweeting is more common in 2020 than it is in 2019, 2019 is more common than 2018, and 2018 is more common than 2017. As discussed in Sections 1&2, tweeting late at night can proxy for the duration of sleep.

In Figure 2, we show weekly rates of late-night tweeting. Specifically, we calculate the weekly fraction of nights on which he posts at least one tweet from 11pm to 2am and take a monthly moving average to reduce noise. Red horizontal lines show the yearly averages. His biggest annual increase is between 2018 and 2019, and as noted above, late-tweeting increases each year. So far within 2020, we see evidence of further increases.

¹⁴We add day of week and month fixed effects and plot the residuals in Supplementary Material Figure S1. The diurnal cycle looks similar with and without these fixed effects, suggesting that his tweeting pattern within a day is relatively unaffected by day of week and month. We separately plot the diurnal cycle in 2020 before and after March 1st in Supplementary Material Figure S2, and stratify by late-tweeting days versus other days in Figure S3.

¹⁵There is a notable tweeting peak around 11am in 2020. Despite the smaller sample size and introducing some subjective reading, we find 33 out of 121 tweets posted between 10am to 1pm after March 1st, 2020 are related to COVID-19. Among these 33 tweets, Governor Cuomo and the New York State are mentioned just four times. Thus, it is unclear whether Governor Cuomo's daily coronavirus news conference around 11am (since March 3rd) has encouraged the President to tweet more in response.

Figure 1: Diurnal cycle every half hour over 24 hours (top) and over sleep hours (bottom)

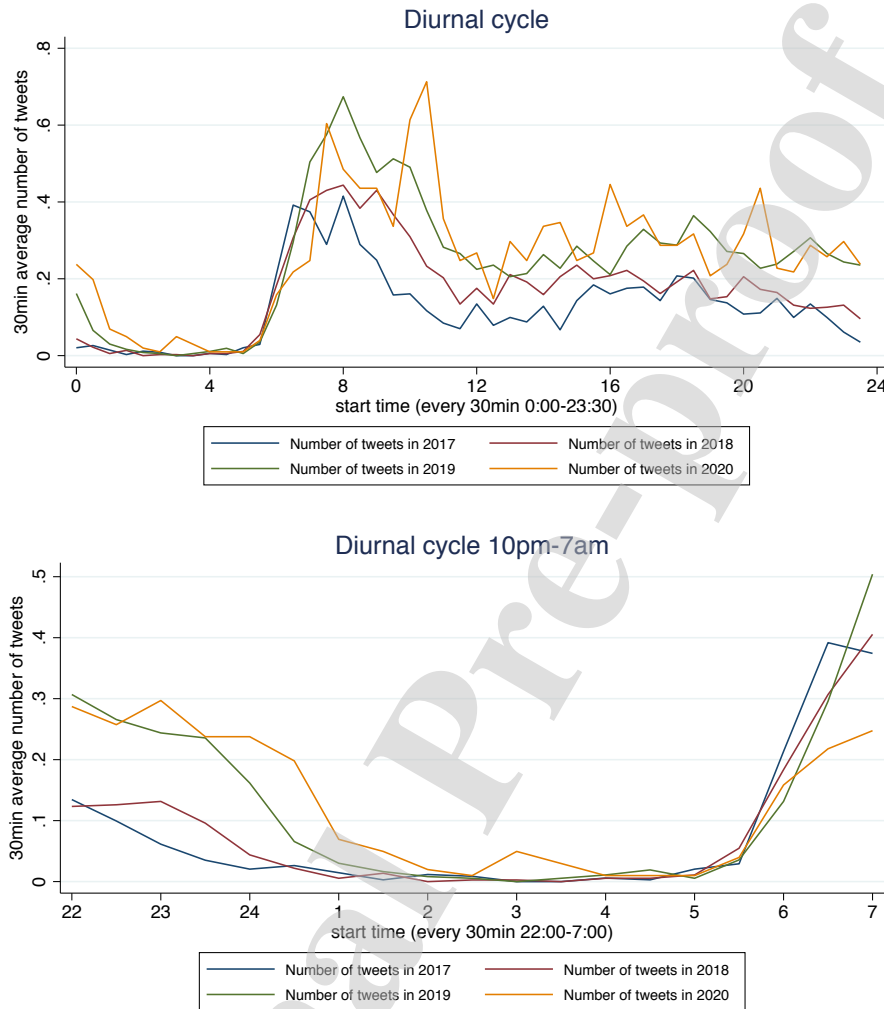



Table 1: Late-night tweeting

Regression results from estimating equation (1) confirm these basic patterns and permit assessment of their statistical significance. The coefficients on the dummy variables in Table 1 indicate the President was more likely to stay up late in 2019 and still more so in 2020 (relative to 2017, see footnote 23 for statistical significance). We use a binary indicator for late tweeting as the dependent variable $Late_t$ in Panel A and the number of late tweets in Panel B. The likelihood of late tweeting increases by 0.22 in 2019 and 0.38 in 2020 relative to the omitted year (2017). This is equivalent to a 183% and 317% increase relative to the 2017 mean, respectively. Additionally, the number of late-night tweets increases over time. He posts roughly one more tweet per night in 2020, a sixfold increase compared with 2017 when he tweeted late about once per week (Panel B).

Figure 2: Fraction of late-tweeting nights from January 24, 2017 to April 10, 2020.

(We use daily binaries of late-tweeting, sum them together every week and divide it by seven to calculate the weekly fraction. Red horizontal lines are the annual mean of the weekly fraction. We sum weekly fraction on that week, four weekly values before and four after, and divide this sum by nine to calculate monthly moving average. The first week of 2018, 2019 and 2020 are marked in vertical dash lines.)



Are these annual increases in late-night tweeting an artifact of increased tweeting activity generally? In Columns 2 and 4, we add $\#Tweets_t$ on the right hand side of the regression and replace the linear control for $\#Tweets_t$ with dummy variables for each possible number of daily tweets (respectively). We make two observations:

1. The intensity of the President's daytime social media use does not help much in predicting his late tweeting. The estimated coefficient on the number of daytime tweets is not distinguishable from 0 in Column 2 Panels A or B. The increase in the (unadjusted) R^2 between Columns 2 and 4 is modest, indicating that the dummy variables for tweet activity do not add a lot of predictive accuracy. Nor do day of the week or calendar month fixed effects contribute much to explanatory power, as these are included in all Table 1 specifications, and R^2 peaks at .122 for Panel A and .146 for Panel B.
2. More importantly, the estimated coefficients on Y_{2020} and Y_{2019} are similar with and without controls for the number of daytime tweets. This suggests that the frequency of Twitter use before sleep does not account for this increase over years (nor do day of week and month of year fixed effects).

We further add a $MarApr * Y_{2020}$ dummy as a control variable in Column 3. While point estimates are positive, the imprecision of the estimates and the almost unchanged R^2 indicate there is no statistically significant change in his late-tweeting pattern after COVID-19 mortality increases in the US, although our test is underpowered.¹⁶

¹⁶We explored robustness to restricting our sample to days in Washington DC. 794 out of 1172 days remain. We report the results in Supplementary Material Section 3.1 Table S5 and they are similar as those in Table 1. We perform analogous robustness checks for other main results in Table S6.

We drop the $MarApr * Y2020_t$ dummy in Column 5 and estimate a Logit model for the extensive margin of late tweeting and a Poisson model for the number of late tweets. In Panel A, logit results show the likelihood of late tweeting increases by 187% in 2019 and 354% in 2020, similar patterns as the linear probability model.¹⁷ In Panel B, the coefficients on $Y2019$ and $Y2020$ indicate the President posts 151% and 244% more late-night tweets relative to the overall mean, and a six- and nine-fold increase relative to the 2017 mean, respectively.

Tables 2-4: Next-day tweets, transcripts, and election odds

Before considering tweet quality, we first assess whether the number of @realDonaldTrump tweets changes the day following a late night. In Supplementary Material Table S1 Panel A, the coefficient on $Late_t$ indicates that President Trump's tweeting frequency is unaffected by the occurrence or a larger number of late tweets. Trump maintains a similar level of activity on Twitter despite staying up late the night before. Thus, we don't believe that the quality difference on the following day results from a quantity change in his tweets.

In Table 2 Panel A-C, tweets after a late-tweeting night receive 7,400 fewer likes, 1,300 fewer retweets and 1,400 fewer replies, or 8%, 6.5% and 7% fewer reactions relative to the mean. We interpret these less-influential postings as lower tweet quality. The coefficients on later years indicate that likes and re-tweets per tweet increase and replies per tweet fall over time. After accounting for the year effects, a late-tweeting night is associated with poorer performances along all three dimensions. The worse Twitter performance is not only a matter of interactions on social media, but also captures longer-term changes in presidential approval ratings. In Supplementary Material Section 2.7, we show more likes means more approval and less disapproval.

Additionally, we find the effects are larger at the lower end of the interactions distribution. In Supplementary Material Section 2.3, the lower deciles of $\#Interactions_{t+1}$ are moved more than the higher ones, suggesting relatively more non-resonant tweets as opposed to fewer "home run" tweets after a late night. Finally, do interactions fall on the subsequent day simply because followers have already interacted more with the President on the previous late night? This might be true if followers had a "liking budget," etc. In Supplementary Material Section 2.6, we show that the President's fewer interactions following a late night are not explained by more interactions the previous (late) night.

In Panel D, we use the proportion of false tweets as the dependent variable and find no significant relationship with late-night tweeting. This suggests sleep deprivation does not herald lower veracity per tweet. That said, the vast majority of false claims cannot be matched to tweets, which limits our power.

¹⁷We get 354% by calculating the likelihood of late tweeting P in 2020 and 2017 holding other controls constant. From the last column of Table 1 Panel A, year dummy $Y2020$ increases $\logit(P) = \ln\left(\frac{P}{1-P}\right)$ by 2.174. With other controls constant, the difference in $\logit(P)$ in 2017 ($=\logit(0.1202)=-1.991$) and 2020 only comes from $Y2020$. $\logit(P)=0.183$ in 2020, and the likelihood of late tweeting is $P=0.5457$. A similar approach is used for year 2019.

The average sentiment of daytime tweets increases by 0.06 in Supplementary Material Table S4. For comparison, the standard deviation of sentiment is .30.¹⁸ The interpretation of sentiment scores for the President's tweets is not straightforward. For example, tweeting about "MAGA" or "Make America Great Again" is coded with positive average sentiment scores by the "off the shelf" Python Vader Sentiment Library dictionary we are using. (The sentiment score for tweet "Make America Great Again!" is 0.66.)

Given the ambiguous implication of sentiment scores, we study the dominant emotion in the President's speech transcripts as classified by the Fact Checker. In Table 3, the proportion of happy transcripts decreases 4.4 percentage points (4.9%) following a late night. Despite his being happy in 88% transcripts, late-tweeting nights and more late tweets appear to make him less happy the following day. In Figure 3, the annual mean of the happy proportion decreased from 90% to 85% after 2019, and this time trend is consistent with the trend of staying up late in Figure 2. Meanwhile, the proportion of angry transcripts increases by 2.9 percentage points after a late night, a nearly three-fold increase compared with the mean 1.1%.¹⁹ We provide text examples of happy, fear(some), and angry transcripts in Supplementary Material Section 5.

In addition to the President's performance, late tweeting does not predict betting on his likelihood of winning, as is shown in Table 4. In contrast, we do find a significant relationship between late tweeting and his competitor's odds. After a late night, more people believe the leading candidate other than Trump is more likely to win and wager on Trump's opponent. The implied chance of his competitor's winning increases by .6 percentage points, or 4.8% relative to the mean.

Interpretation

Empirically and from a casual inference perspective, the relatively low predictive power in Table 1 (R^2 around .1) compared to Table 2 (R^2 around .2) may be viewed as a virtue. To the extent that late-night tweeting behavior is less predictable, it might be more exogenous. And to the extent that variation in late-night tweeting is independent conditional on the other regression controls, then Table 2 will capture the causal effect of late-night tweeting on tweet quality the following day (Angrist and Pischke, 2009). That said, there are many additional potential control variables we do not include in equations (1) and (2) because we do not observe them. In the absence of a natural experiment in late-night tweeting or a deeper understanding of what generates late-night tweeting, we interpret θ_1 in equation (2) conservatively as a partial correlation coefficient.

When we include fixed effects for the # tweets (i.e. 31 separate dummy variables for each possible number of daily tweets, running from 0 to 31), this restricts comparisons to be purely within days t where the number of tweets are identical.

¹⁸Quantiles of the sentiment are 0.063, 0.2860, 0.4713.

¹⁹The opposing patterns captured by tweets' sentiment scores and transcripts' dominant emotion could be due to the differing metrics of text analysis. It may also result from a negative feedback on emotion from lower Twitter interactions, or a substitution of emotion between tweets and speech.

Because our estimated coefficients of interest are unaffected by # tweets fixed effects, we do not think late-night tweeting is simply an artifact of continued busyness, at least as reflected by daytime social media activity. Thus, persistent busyness is not an omitted factor we believe drives our results. Nor do we think the President's travel schedule, seasonal effects, annual time trends, or day of week effects drive our results, given our set of control variables.²⁰

²⁰In the case of travel schedule, we show robustness using a stratification: only considering days in Washington DC.

Table 1: Predicting late-tweeting with daytime tweets before sleep, year, etc.

	Panel A: Late tweeting dummy				
	OLS				Logit
Y2020	0.386*** (0.051)	0.378*** (0.055)	0.337*** (0.070)	0.383*** (0.056)	2.174*** (0.325)
Y2019	0.229*** (0.031)	0.223*** (0.034)	0.219*** (0.034)	0.220*** (0.034)	1.352*** (0.218)
Y2018	0.070** (0.031)	0.067** (0.031)	0.064** (0.032)	0.065** (0.032)	0.509** (0.221)
# daytime tweets before sleep		.00121 (.00275)			
MarApr*Y2020			.105 (.0963)		
Observations	1172	1172	1172	1172	1172
R-square	0.100	0.100	0.122	0.121	0.105
Y-mean	0.2389	0.2389	0.2389	0.2389	0.2389
Y-mean 2017	0.1202	0.1202	0.1202	0.1202	0.1202
	Panel B: Count of late tweets				
	OLS				Poisson
Y2020	1.129*** (0.127)	1.094*** (0.137)	0.987*** (0.173)	1.083*** (0.138)	2.442*** (0.213)
Y2019	0.609*** (0.077)	0.586*** (0.084)	0.574*** (0.085)	0.576*** (0.085)	1.514*** (0.156)
Y2018	0.168** (0.077)	0.158** (0.078)	0.149* (0.079)	0.152* (0.079)	0.659*** (0.167)
# daytime tweets before sleep		.0048 (.00686)			
MarApr*Y2020			.22 (.239)		
Observations	1172	1172	1172	1172	1172
R-square	0.118	0.118	0.146	0.145	0.162
Y-mean	0.4787	0.4787	0.4787	0.4787	0.4787
Y-mean 2017	0.1613	0.1613	0.1613	0.1613	0.1613
DOW FEs	Y	Y	Y	Y	Y
Month FEs	Y	Y	Y	Y	Y
#Tweets FEs			Y	Y	Y

Notes: Pseudo R-square is reported for logit and poisson regression. * significant 10% level; ** significant at 5% level; *** significant at 1% level.

Table 2: Late-tweeting and likes, retweets, replies, proportion of false tweeting after sleep

Panel A: Likes after sleep (in thousands)						
Late dummy	-7.575*** (2.286)	-7.435*** (2.256)	-7.412*** (2.279)			
Late count				-2.524*** (0.917)	-2.429*** (0.905)	-2.374*** (0.917)
# daytime tweets before sleep		-1.181*** (0.211)			-1.178*** (0.211)	
Days	-0.108 (0.105)	-0.102 (0.104)	-0.084 (0.105)	-0.111 (0.105)	-0.105 (0.104)	-0.084 (0.105)
Y2018	48.4 (38.4)	48.6 (37.9)	41.2 (38.4)	49.4 (38.5)	49.4 (38)	41.2 (38.5)
Y2019	93 (76.8)	94.1 (75.8)	80.7 (76.8)	94.9 (77)	95.9 (76)	80.7 (76.9)
Y2020	151 (115)	153 (114)	134 (115)	154 (115)	156 (114)	134 (115)
Observations	1153	1153	1153	1153	1153	1153
R-square	0.200	0.222	0.243	0.198	0.219	0.240
Y-mean	91.67	91.67	91.67	91.67	91.67	91.67
Panel B: Retweets after sleep (in thousands)						
Late dummy	-1.400*** (0.528)	-1.373*** (0.523)	-1.325** (0.530)			
Late count				-0.497** (0.212)	-0.478** (0.210)	-0.452** (0.213)
# daytime tweets before sleep		-0.231*** (0.049)			-0.230*** (0.049)	
Days	-0.025 (0.024)	-0.023 (0.024)	-0.018 (0.024)	-0.025 (0.024)	-0.024 (0.024)	-0.018 (0.025)
Y2018	11.6 (8.88)	11.6 (8.8)	9.58 (8.95)	11.8 (8.89)	11.8 (8.81)	9.59 (8.95)
Y2019	21.6 (17.7)	21.8 (17.6)	18.1 (17.9)	22 (17.8)	22.2 (17.6)	18.1 (17.9)
Y2020	33.7 (26.6)	34 (26.4)	28.7 (26.8)	34.3 (26.6)	34.6 (26.4)	28.8 (26.8)
Observations	1153	1153	1153	1153	1153	1153
R-square	0.159	0.175	0.191	0.158	0.174	0.190
Y-mean	20.46	20.46	20.46	20.46	20.46	20.46
Panel C: Replies after sleep (in thousands)						
Late dummy	-1.484** (0.695)	-1.454** (0.690)	-1.400** (0.699)			
Late count				-0.441 (0.278)	-0.421 (0.277)	-0.381 (0.281)
# daytime tweets before sleep		-0.251*** (0.064)			-0.251*** (0.065)	

September 15, 2020

do not cite or distribute

Days	0.019 (0.032)	0.020 (0.032)	0.024 (0.032)	0.018 (0.032)	0.019 (0.032)	0.024 (0.032)
Y2018	-7.63 (11.7)	-7.6 (11.6)	-9.11 (11.8)	-7.49 (11.7)	-7.47 (11.6)	-9.16 (11.8)
Y2019	-18.1 (23.4)	-17.9 (23.2)	-20.6 (23.5)	-17.9 (23.4)	-17.7 (23.2)	-20.7 (23.6)
Y2020	-28.7 (35)	-28.4 (34.8)	-32.1 (35.3)	-28.3 (35.1)	-28 (34.8)	-32.3 (35.3)
Observations	1153	1153	1153	1153	1153	1153
R-square	0.178	0.189	0.208	0.176	0.187	0.206
Y-mean	19.81	19.81	19.81	19.81	19.81	19.81
Panel D: Proportion of false tweets after sleep (in percentage)						
Late dummy	0.103 (1.502)	0.105 (1.503)	-0.162 (1.526)			
Late count				0.439 (0.601)	0.441 (0.602)	0.415 (0.613)
# daytime tweets before sleep		-0.019 (0.140)			-0.021 (0.140)	
Days	-0.005 (0.069)	-0.004 (0.069)	-0.006 (0.070)	-0.004 (0.069)	-0.004 (0.069)	-0.006 (0.070)
Y2018	14.6 (25.3)	14.6 (25.3)	15.5 (25.7)	14.3 (25.3)	14.3 (25.3)	15.2 (25.7)
Y2019	19.9 (50.5)	19.9 (50.5)	21.8 (51.4)	19.1 (50.5)	19.1 (50.5)	21 (51.4)
Y2020	6.12 (75.7)	6.15 (75.8)	8.47 (77.1)	4.84 (75.7)	4.87 (75.8)	7.17 (77.1)
Observations	1153	1153	1153	1153	1153	1153
R-square	0.160	0.160	0.175	0.160	0.160	0.175
Y-mean	26.90	26.90	26.90	26.90	26.90	26.90
DOW FEs	Y	Y	Y	Y	Y	Y
Month FEs	Y	Y	Y	Y	Y	Y
#Tweets FEs			Y			Y

Notes: The smaller sample size than that in Table 1 is due to days with no daytime tweets. Dependent variable likes, retweets and replies are divided by 1000, proportion is multiplied by 100. * significant 10% level; ** significant at 5% level; *** significant at 1% level.

Table 3: Late-tweeting and dominant emotion of transcripts after sleep

	Proportion of dominant emotion in transcripts after sleep (in percentage)					
	Happy	Fear	Angry	Happy	Fear	Angry
Late dummy	-4.362* (2.349)	-1.444* (0.794)	2.942*** (0.822)			
Late count				-1.411 (0.865)	-0.769*** (0.292)	0.440 (0.305)
Days	.197* (.113)	-.005 (.0384)	-.0341 (.0397)	.203* (.113)	-.00489 (.0383)	-.0406 (.0399)
Y2018	-70.323* (41.481)	2.078 (14.031)	12.525 (14.515)	-72.373* (41.453)	2.045 (13.983)	14.998 (14.596)
Y2019	-146.470* (82.738)	4.943 (27.987)	26.073 (28.952)	-150.520* (82.680)	5.002 (27.889)	31.168 (29.113)
Y2020	-220* (124)	9.87 (42)	39.3 (43.5)	-226* (124)	10 (41.9)	47 (43.7)
Observations	837	837	837	837	837	837
R-square	0.091	0.140	0.126	0.090	0.144	0.114
Y-mean	88.18	0.8727	1.080	88.18	0.8727	1.080
DOW FEs	Y	Y	Y	Y	Y	Y
Month FEs	Y	Y	Y	Y	Y	Y
#Tweets FEs	Y	Y	Y	Y	Y	Y

Notes: The smaller sample size than that in Table 1 is due to days with no transcripts. proportion is multiplied by 100. * significant 10% level; ** significant at 5% level; *** significant at 1% level.

Figure 3: Proportion of happy transcripts from January 24, 2017 to April 10, 2020.

(For each transcript, we assign happy dummy which equals one if its dominant emotion is happy and zero if not. We use Trump's word count in each transcript as weight and calculate weighted sum of happy dummies on each day. We define daily proportion of happy transcripts as the weighted sum divided by the number of transcripts. We take the average of daily proportion to weekly value, and calculate monthly moving average using that week, four weeks before and after. Red horizontal lines are the annual mean of the weekly proportion. The first week of 2018, 2019 and 2020 are marked in vertical dash lines.)




Table 4: Late-tweeting and election odds after sleep

Panel A: Trump's implied probability (in percentage)						
Late dummy	0.073 (0.273)	0.061 (0.271)	0.139 (0.276)			
Late count				0.148 (0.109)	0.140 (0.109)	0.149 (0.111)
# daytime tweets before sleep		0.086*** (0.025)			0.086*** (0.025)	
Days	.0172 (.0125)	.0166 (.0125)	.0213* (.0127)	.0174 (.0125)	.0168 (.0125)	.0214* (.0127)
Y2018	0.412 (4.582)	0.462 (4.561)	-1.234 (4.640)	0.318 (4.579)	0.372 (4.558)	-1.296 (4.637)
Y2019	1.12 (9.15)	1.16 (9.11)	-2.28 (9.27)	.903 (9.15)	.95 (9.11)	-2.43 (9.26)
Y2020	9.78 (13.7)	9.84 (13.7)	4.65 (13.9)	9.42 (13.7)	9.5 (13.7)	4.4 (13.9)
Observations	1172	1172	1172	1172	1172	1172
R-square	0.809	0.811	0.815	0.810	0.812	0.815
Y-mean	36.27	36.27	36.27	36.27	36.27	36.27
Panel B: Competitor's implied probability (in percentage)						
Late dummy	0.617** (0.276)	0.616** (0.276)	0.586** (0.278)			
Late count				0.190* (0.111)	0.189* (0.111)	0.187* (0.112)
# daytime tweets before sleep		0.013 (0.026)			0.013 (0.026)	
Days	.0184 (.0127)	.0183 (.0127)	.0189 (.0128)	.0186 (.0127)	.0186 (.0127)	.0189 (.0128)
Y2018	-6.536 (4.637)	-6.528 (4.639)	-6.674 (4.682)	-6.603 (4.642)	-6.594 (4.643)	-6.676 (4.685)
Y2019	-6.84 (9.26)	-6.83 (9.27)	-7.23 (9.35)	-6.97 (9.27)	-6.96 (9.28)	-7.24 (9.36)
Y2020	-6.22 (13.9)	-6.12 (13.9)	-1.05 (14)	-8.33 (13.9)	-8.2 (13.9)	-1.06 (14)
Observations	1172	1172	1172	1172	1172	1172
R-square	0.693	0.693	0.703	0.692	0.692	0.703
Y-mean	12.24	12.24	12.24	12.24	12.24	12.24
DOW FEs	Y	Y	Y	Y	Y	Y
Month FEs	Y	Y	Y	Y	Y	Y
#Tweets FEs			Y			Y

Notes: * significant 10% level; ** significant at 5% level; *** significant at 1% level.

5 Discussion

There is a large literature in political economy considering the quality of politicians, how they are selected, etc. (Mattozzi and Merlo, 2007; Galasso and Nannicini, 2011; Dal Bó, Finan, Folke, Persson, and Rickne, 2017; Dal Bó and Finan, 2018).²¹ Jones and Olken (2005) consider the impact of changes in national leadership on economic performance and Besley, Montalvo, and Reynal-Querol (2011) link heterogeneity in this growth effect to the educational attainment of leaders. Typically in this literature, quality is assumed to be a fixed characteristic for the individual politician.²²

We expand on this conception of politician quality to include not only time-varying quality measures for the individual, but indeed quality that varies at a very high frequency (daily). A practical virtue of our approach is that comparisons can be restricted to be exclusively within the same politician (a comparison not permitted by the static conception of quality). As candidate quality/valence is multidimensional (Dal Bó, Finan, Folke, Persson, and Rickne, 2017; Dal Bó and Finan, 2018), comparing the politician to only herself/himself accounts for unobserved dimensions of individual quality at the politician level that remain fixed. For example, unobserved integrity is likely correlated with observed measures of quality and may thereby cloud interpretations of static, unidimensional measures observed and analyzed by researchers.

Empirical opportunities to observe time-varying politician quality abound, particularly in the “big data” era. This development allows researchers to harness within-subject designs in their analyses and consider multiple measures of quality/valence, as we do here. A natural extension of our approach would be to consider high-frequency measures of quality/valence for multiple politicians in a panel data design.

A Nobel laureate has recently argued that Economics commits systematic “sins of omission” by ignoring important research topics:

...it is easy for people to agree regarding the hardness/softness of research. In contrast, importance is fuzzy, so that it is relatively easy to disagree regarding its importance.
Akerlof (2020)

Where research is evaluated by committee consensus, including tenure and journal review, this imparts an evaluation bias toward “hardness” over societal importance (Akerlof, 2020).

Although studying sleep may appear “soft”, quaint, or pedestrian, the President’s sleep is important because:

1. Previous research documents large performance impacts from sleep, including modest *randomized* reductions in sleep;

²¹We thank an anonymous referee for suggesting we discuss the political economy literature on politician quality.

²²The work on how term limits change candidate incentives and effort, e.g. Dal Bó and Rossi (2011), is a notable exception.

2. Mr. Trump's performance is immensely important to others given his position as US President;
3. Publicly available data suggest the President is sleep-deprived.

If the President's sleep is sub-optimally short in 2020, it can and should be addressed.

We have leveraged the President's frequent tweeting – roughly 10 times a day on average – to construct the best publicly-available proxy for his sleep, following the existing sleep literature. Looking across 1,200 nights since inauguration, the President appears to be sleeping substantially less as his first term has progressed. In general, gradual developments can be more difficult to notice and thereby react to appropriately, e.g. Moore et al. (2019) on climate change. Figure 2 indicates a fairly gradual increase in late-night Twitter activity. This important trend may have gone unnoticed even at the White House.

The US Centers for Disease Control and Prevention recommends that adults over age 65 sleep 7-8 hours per night. On the 3 nights a week the President tweets late in 2020, his average tweet time is 12:06 am. This suggests that when the President stays up late in 2020, he is asleep fewer than 6 hours on average. If we assume that the President was sleeping his optimal – albeit short – personal amount in 2017 or 2018, this no longer appears to be the case in 2020.²³ In late February-early March 2020, his fraction of nights with a late-night tweet reached a 160 week high of .54.²⁴ This increase is not accounted for by his increased tweeting activity generally. If the President's sleep has fallen below his optimum – and indeed perhaps well below his optimum – this provides context for interpreting the frequent official communications and policy announcements from the Executive Branch. Was the President up late the preceding night? Furthermore, the sleep adequacy of an individual can be addressed at low cost. For example, the measurement of the President's sleep could be improved with a personal activity monitor (e.g. Fitbit) and the White House physician could review these data. The benefit of lengthening sleep may be exceptionally high relative to its cost.

Underscoring the need to evaluate his sleep, we see systematic differences in the President's performance following one of his late nights. This is plausible given both existing research and personal experience. Still, our evidence for this relationship reflects the tradeoff between topic importance and methodological "hardness" described by Akerlof (2020). We argue we are at the frontier for this tradeoff: there is not a "harder" way to explore a question of this importance absent either direct sleep measures or an identification strategy that changes the President's sleep exogenously.

²³We test the equality of coefficients on Y2020, Y2019 and Y2018 dummy variables in Table 1 Column 4. For the binary outcome of late-night tweeting, Y2020 is larger than Y2019 by 0.1632 (p-value=0.001), Y2019 is larger than Y2018 by 0.1554 (p-value=0.000). For Panel B, the differences are 0.5071 (0.000) and 0.4247 (0.000).

²⁴The monthly moving average is described in Figure 3. We "lose" 8 weeks to calculate the moving average at the beginning and end of the sample period.

References

- Akerlof, G. A. (2020, June). Sins of omission and the practice of economics. *Journal of Economic Literature* 58(2), 405–18. 5, 5
- Andrew, M. (2017). A bot that can tell when it's really donald trump who's tweeting. 9
- Angrist, J. D. and J.-S. Pischke (2009). *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton, New Jersey: Princeton University Press. 4
- Avery, M., O. Giuntella, and P. Jiao (2019, November). Why don't we sleep enough? a field experiment among college students. IZA DP No. 12772. 2, 6
- Barnes, C. M., J. Schaubroeck, M. Huth, and S. Ghumman (2011). Lack of sleep and unethical conduct. *Organizational Behavior and Human Decision Processes* 115(2), 169 – 180. 2
- Besley, T., J. G. Montalvo, and M. Reynal-Querol (2011, 07). Do Educated Leaders Matter? *The Economic Journal* 121(554), F205–227. 5
- Bessone, P., G. Rao, F. Schilbach, H. Schofield, and M. Toma (2020, February). The economic consequences of increasing sleep among the urban poor. Working Paper 26746, National Bureau of Economic Research. 6
- Bhat, S., G. Pinto-Zipp, H. Upadhyay, and P. G. Polos (2018). To sleep, perchance to tweet": in-bed electronic social media use and its associations with insomnia, daytime sleepiness, mood, and sleep duration in adults. *Sleep Health* 4(2), 166 – 173. 1, 2
- Bianchi, F., T. Kind, and H. Kung (2019, September). Threats to central bank independence: High-frequency identification with twitter. Working Paper 26308, National Bureau of Economic Research. 2
- Biddle, J. E. and D. S. Hamermesh (1990, October). Sleep and the Allocation of Time. *Journal of Political Economy* 98(5), 922–943. 2
- Borger, J. N., R. Huber, and A. Ghosh (2019). Capturing sleep-wake cycles by using day-to-day smartphone touchscreen interactions. *NPJ Digital Medicine* 2(73), 1–8. 2
- Carnegie, A. and A. Carson (2019). Reckless rhetoric? compliance pessimism and international order in the age of trump. *The Journal of Politics* 81(2), 739–746. 2
- Christian, M. S. and A. P. J. Ellis (2011). Examining the effects of sleep deprivation on workplace deviance: A self-regulatory perspective. Volume 54. 2
- Dal Bó, E. and F. Finan (2018). Progress and perspectives in the study of political selection. *Annual Review of Economics* 10(1), 541–575. 1, 5
- Dal Bó, E., F. Finan, O. Folke, T. Persson, and J. Rickne (2017). Who Becomes A Politician? *The Quarterly Journal of Economics* 132(4), 1877–1914. 5

- Dal Bó, E. and M. A. Rossi (2011, 05). Term Length and the Effort of Politicians. *The Review of Economic Studies* 78(4), 1237–1263. 22
- Exelmans, L. and J. Van den Bulck (2016). Bedtime mobile phone use and sleep in adults. *Social Science & Medicine* 148, 93 – 101. 1, 2
- Galasso, V. and T. Nannicini (2011). Competing on good politicians. *American Political Science Review* 105(1), 79–99. 5
- Ge, Q., A. Kurov, and M. H. Wolfe (2019). Do investors care about presidential company-specific tweets? *Journal of Financial Research* 42(2), 213–242. 2
- Gibson, M. and J. Shrader (2018). Time use and labor productivity: The returns to sleep. *The Review of Economics and Statistics* 100(5), 783–798. 2
- Golder, S. A. and M. W. Macy (2011). Diurnal and seasonal mood vary with work, sleep, and daylength across diverse cultures. *Science* 333(6051), 1878–1881. 1, 2
- Halcoussis, D., A. D. Lowenberg, and G. M. Phillips (2020). An empirical test of the Comey effect on the 2016 presidential election. *Social Science Quarterly* 101(1), 161–171. 3.3
- Hobbs, W. and N. Lajevardi (2019). Effects of divisive political campaigns on the day-to-day segregation of arab and muslim americans. *American Political Science Review* 113(1), 270–276. 2
- Jones, B. F. and B. A. Olken (2005, 08). Do Leaders Matter? National Leadership and Growth Since World War II. *The Quarterly Journal of Economics* 120(3), 835–864. 5
- Jones, J. J., G. W. Kirschen, S. Kancharla, and L. Hale (2019). Association between late-night tweeting and next-day game performance among professional basketball players. *Sleep Health* 5(1), 68 – 71. 1, 2
- Kryger, M. (2017). What can tweets tell us about a person’s sleep? *Journal of Clinical Sleep Medicine* 13(10), 1219–1221. 2
- Leypunskiy, E., E. Kiciman, M. Shah, O. J. Walch, A. Rzhetsky, A. R. Dinner, and M. J. Rust (2018). Geographically resolved rhythms in twitter use reveal social pressures on daily activity patterns. *Current Biology* 28(23), 3763–3775. 2
- Li, J., M. V. Vitiello, and N. Gooneratne (2018). Sleep in normal aging. *Sleep Medicine Clinics* 13(1), 1–11. 2
- Mai, F., Z. Chen, and A. Lindberg (2019). Does sleep deprivation cause online incivility? evidence from a natural experiment. International Conference on Information Systems 2019 Conference. 2
- Mattozzi, A. and A. Merlo (2007, May). The transparency of politics and the quality of politicians. *American Economic Review* 97(2), 311–315. 5
- Miles, M. R. and D. P. Haider-Markel (2018). Trump, twitter, and public dissuasion: A natural experiment in presidential rhetoric. *Presidential Studies Quarterly*. 2

- Moore, F. C., N. Obradovich, F. Lehner, and P. Baylis (2019). Rapidly declining remarkability of temperature anomalies may obscure public perception of climate change. *Proceedings of the National Academy of Sciences* 116(11), 4905–4910. 5
- Muller, K. and C. Schwarz (2019). From hashtag to hate crime: Twitter and anti-minority sentiment. 2
- Pilcher, J. J., D. M. Morris, J. Donnelly, and H. B. Feigl (2015). Interactions between sleep habits and self-control. *Frontiers in Human Neuroscience* 9, 284. 2
- Rhode, P. W. and K. S. Strumpf (2004, June). Historical presidential betting markets. *Journal of Economic Perspectives* 18(2), 127–141. 3:3
- Richards, A., S. S. Inslicht, T. J. Metzler, B. S. Mohlenhoff, M. N. Rao, A. O'Donovan, and T. C. Neylan (2017, 02). Sleep and Cognitive Performance From Teens To Old Age: More Is Not Better. *Sleep* 40(1). zsw029. 2
- Ross, A. S. and D. J. Rivers (2018). Discursive deflection: Accusation of “fake news” and the spread of mis- and disinformation in the tweets of president trump. *Social Media & Society* 4(2), 1–12. 2
- Sabeti, S., Z. Al-Darsani, B. A. Mander, M. M. Corrada, and C. H. Kawas (2018). Sleep, hippocampal volume, and cognition in adults over 90 years old. *Aging clinical and experimental research* 30(11), 1307–1318. 2
- Scott, H. and H. C. W. Woods (2019). Understanding links between social media use, sleep and mental health: Recent progress and current challenges. *Current Sleep Medicine Reports* 5, 141–149. 1, 2
- Scotto di Carlo, G. (2020). Trumping twitter: Sexism in president trump's tweets. *Journal of Language and Politics* 19(1), 48–70. 2
- Scullin, M. K. and D. L. Bliwise (2015). Sleep, cognition, and normal aging: Integrating a half century of multidisciplinary research. *Journal of the Association for Psychological Science* 10(1), 97–137. 2
- Sterniczuk, R., O. Theou, B. Rusak, and K. Rockwood (2013, 07). Sleep disturbance is associated with incident dementia and mortality. *Current Alzheimer research* 10. 2
- Van Dongen, H. P., G. Maislin, J. M. Mullington, and D. F. Dinges (2003, 03). The Cumulative Cost of Additional Wakefulness: Dose-Response Effects on Neurobehavioral Functions and Sleep Physiology From Chronic Sleep Restriction and Total Sleep Deprivation. *Sleep* 26(2), 117–126. 2
- Wams, E., G. Wilcock, R. Foster, and K. Wulff (2017, 05). Sleep-wake patterns and cognition of older adults with amnesic mild cognitive impairment (amci): A comparison with cognitively healthy adults and moderate alzheimer's disease patients. *Current Alzheimer research* 14. 2
- Wells, C., D. Shah, J. Lukito, A. Pelled, J. C. Pevehouse, and J. Yang (2020). Trump, twitter, and news media responsiveness: A media systems approach. *New Media & Society* 22(4), 659–682. 2

Whissell, C. (2018). The medium may be the same but the message is different: Comparing the tweets of u.s. presidents obama and trump. *Advances in Social Sciences Research Journal* 5(6), 33–38. 2

Wolfers, J. and E. Zitzewitz (2016, October). What do financial markets think of the 2016 election? Brookings Institute. 3.3

Highlights for EJ

- President Trump's sleep appears short relative to that early in his Presidency
- "Late nights" are proxied using the timing of the President's social media posts
- After a late night, various performance metrics for the following day change
- In transcript texts, the President is more likely to be "angry" after a late night
- Twitter followers interact less with the President's posts following a late night