Productivity losses in the transition to Daylight Saving Time: Evidence from hourly GitHub activity

Andrew Dickinson [1] Glen R. Waddell [1,2]

[1] Department of Economics, University of Oregon
[2] Research Fellow IZA Bonn

July 2023

Abstract

Using data on GitHub users around the world, we estimate the effects of transitions to Daylight Saving Time on worker activity. In daily activity, transitions appear short lived—there is evidence of two days of declines before activity returns to baseline levels. However, hourly analysis reveals a transition to Daylight Saving Time that is much longer—losses appear in the early working hours of work days for up-to-two weeks following the initiation of Daylight Saving Time.

Keywords: daylight saving time; productivity; sleep; time use
JEL: J24; J22

1 Introduction

Motivated by potential energy savings, Daylight Saving Time (DST) was introduced in 1916 and widely adopted by many Western countries soon after.[1] Today, the consensus opinion with respect to energy consumption is that any savings associated with DST are likely small (Kotchen and Grant, 2011; Belzer et al., 2008)—they are on the order of plus-or-minus one percent. The merits of continuing these twice-yearly adjustments to the clocks of 1.6 billion people across 75 countries are therefore important contributors to the continuing debate, even as countries abandon the practice.[2] Most recently, the United States (US) Senate unanimously passed the Sunshine Protection Act in 2021, which would have the United States likewise abandoning the practice.[3]

While the Sunshine Protection Act remains in the US House of Representatives, of note is the House Energy and Commerce Committee's interest in learning of the productivity effects of time changes prior to holding a vote. Our analysis contributes directly to the debate, then, and with immediate implication for policy. To that end, we document significant declines in worker activity following DST, with patterns that suggest a weeks-long period of transition through which workers are adjusting to the shifting of time relative to the cycle of daylight and darkness.

To measure the potential declines in activity induced by DST we exploit publicly available activity data from GitHub, a popular cloud based version-control platform used by collaborative programmers. Observations include detailed information regarding changes to a project's

[1] The shifting forward of the clock increased the availability of natural light in the evenings—at the time, this lessened the need for the carbon filaments and kerosene associated with indoor lighting (Nordhaus, 1997).

[2] China first abandoned the practice in 1991, and several other countries following that path in the mid-to-late 2000s. Other countries discontinuing DST include Pakistan (2009), Russia (2010), Azerbaijan (2015), and Brazil (2019).

[3] This national-level action follows several state-level initiates that similarly moved toward the permanent adoption of DST. For example, Oregon (2019) and Washington (2019) have passed such legislation, and California's Proposition 7 (2018) was approved by over 60 percent of voters.

codebase—including what was changed and who made the change—providing a direct measurement of labor activity. Of particular value, though, is that we observe GitHub user activity in second-level precision. This level of granularity makes GitHub a natural laboratory to consider the effects of DST. In our analysis, we consider activity aggregated to both the daily and hourly levels across cities that experience the Spring transitions to DST differently—those that experience it earlier than others, experience it later than others, or do not experience it at all.

In daily activity we find significant declines on the order of one or two days. While such declines are presumably associated with significant economic costs, this alone suggests that the productivity losses are relatively short-lived. However, in hourly activity we find much longer periods of transition. In particular, in the two weeks following transitions to DST we find significant declines in worker activity in early work hours, between 8am and 10am. That daily measures of productivity evidence significant declines for such a short period of time suggests that the slow morning hours in the data are made up for with increases in activity elsewhere in the work day. However, increases in activity tend not to cluster in particular hours of the day. Rather, the data are consistent with workers recovering from morning productivity declines without a common strategy for dealing with the transition. In the end, it is clear from an hourly analysis that the belief in "one or two days of decline" likely fails to capture the full economic, personal, or potential social costs induced by DST.

Circadian rhythms—that is, the processes that synchronize and regulate the body's sleep-wake cycle around the 24-hour day—are resilient to adjustments to sleep-wake cycles, and the transition to a modified cycle of daylight and darkness is not immediate (Kantermann et al., 2007). Thus, that we find long-lasting disruptions to activity are consistent with this resilience. This is also consistent with laboratory results suggest that transitions to DST are measured in days (Monk and Aplin, 1980; Czeisler et al., 1999), and associated with immediate losses of sleep that do not typically return to normal for roughly seven days (Kantermann et al., 2007;

Lahti et al., 2006). Workplace safety (Barnes and Wagner, 2009; Lahti et al., 2011), road safety (Smith, 2016; Bünnings and Schiele, 2021), and student performance (**?**) have all been shown to decline following DST. Google searches for entertainment-related keywords and phrases in US cities that experience DST (i.e., those other than in Arizona and Hawaii) are differentially higher on the Monday that immediately follows DST transitions than on the Monday prior to DST (Wagner et al., 2012). This has been interpreted as DST-induced increases in "cyberloafing" and suggests that declines in workplace productivity can result from DST. However, to our knowledge, documenting worker activity around DST transitions has thus far been absent in the literature.[4]

While high-frequency observations of activity available on GitHub make for a rare opportunity to learn about such behavior, GitHub users are not representative of the labor force of any country, and certainly not of the world. However, with roughly 83 million users currently, GitHub does account for a large and growing sector of the global labor market. It is also arguable that our sample is representative of a much-larger group of similarly skilled workers who work across industry, government, academic, and policy environments. In many work environments, for example, similarly skilled individuals are working in proprietary tasks that preclude their formal participation in a public code-sharing environment, and in others we would find still more consumers of the public offerings of active GitHub users. In this way, it is reasonable to infer that the effects we identify extend naturally to a larger group of workers.

In Section 2 we describe the data we rely on in our analysis and provide the necessary context for the interpretation of GitHub user data as we consider whether there is evidence of a causal relationship between DST and worker activity. In Section 3 we discuss our methodology

[4] Somewhat more removed from the research question, DST has also been associated with lower well-being and general life satisfaction (Kountouris and Remoundou, 2014; Kuehnle and Wunder, 2016), increases in myocardial infarction (Toro, Tigre, and Sampaio, 2015), reductions in crime (Doleac and Sanders, 2015), and increases in suicide (Berk et al., 2008). (DST has been associated with movements in financial markets, though the evidence is mixed. For example, Kamstra, Kramer, and Levi (2000) suggests that there are negative financial returns following DST weekends, while Gregory-Allen et al. (2010) finds no evidence of a Daylight Saving Time anomaly.)

and report event-study analyses at both the daily and hourly level. We summarize and offer concluding remarks in Section 4.

2 Data

2.1 Sample selection

GitHub is a popular cloud based version-control and code-hosting platform. Widely adopted by software developers, engineers, and scientific coders for hosting and maintaining collaborative projects, GitHub is a web based extension to the most popular distributed version control system, Git.[5] With a current userbase of over 83 million, GitHub is the largest platform of its kind.

Consistent with its origin as a tool for collaboration with excellent version control, the primary feature of GitHub is the facilitation of code updates across multiple collaborative users. In essence, it provides broad access and version control to a remote copy of a project and all of its files, and as progress is made to a project hosted on the platform, contributions are tracked across time and contributor, which generates a record of labor activity metadata.

For the purpose of this study, we collect GitHub activity around the DST episodes from 2013 through 2019.[6] To reflect "active" users we restrict each year's sample to those we observe activity for in each month of that year. To limit the potential presence of "bot" activity, we discard all observations originating from accounts that ever have more than 30 events in one hour or 150 in one day (McDermott and Hansen, 2021). Following these restrictions, the total number of unique users in the sample is 141,899. Having matched each event in the raw data to

[5] Of the 80,000 respondents in the 2021 Stack Overflow Annual Developer Survey, over 90 percent reported using Git.

[6] We draw our sample from two third-party sources (GH Archive and GHTorrent) that have organized the event data from GitHub's public timeline for the purpose of being more accessible to researchers. Our analysis therefore does not include any activity on private repositories. For example, a business producing proprietary software, or a team of data scientists performing analysis on valuable or sensitive data, would host their repositories privately and would not be observed on the public timeline.

the location identified in the user's profile at the time of the event, we aggregate event counts to the city-by-day level (or city-by-hour level) and restrict the sample to the 50 most-active cities. These are reported in Table S2. As expected, the geographic distribution of these workers reflects major cities known for their robust presence in the tech industry.

2.2 Outcomes

All activity on GitHub is recorded as an "event," and categorized as one of 13 types (see Table S1). In some way, they are each potential measures of incremental productivity on shared projects. However, given the inability to distinguish the relative productivities of event types, broadly, we will consider all events, and then separately consider the 42.4 percent of events that are classified as "pushes." Push events are uploaded changes of a local file to the remote copy, which are likely stronger signals of productivity. For example, this analysis will exclude "pull requests," which represent contributors' requests to merge the code changes of others into their remote copy. In the end, we find comparable patterns in activity around transitions to DST in both measures.

2.3 Treatment

While many countries practice DST, there is variation across cities in both the practice of DST and in the timing of DST. Across the 50 cities in our sample, two important sources of treatment variation exist. In North America, clocks transition to DST on the second Sunday of March.[7] In the European Union, clocks transition on the last Sunday of March. Both of these will contribute to identifying the effect of DST on GitHub activity—given the staggered timing of treatment, they each act as a control for the other, and can each be compared to the "never-treated" cities.

[7] The exceptions to this in the 2013–2019 period are Hawaii, parts of Arizona, Puerto Rico, US territories in the Pacific Islands, Saskatchewan, and the Yukon. None of these are among the 50-largest cities, and are therefore not in our sample.

Given the staggered timing of treatment, we follow Sun and Abraham (2021) in the estimation of treatment effects.[8]

In Figure 1 we plot patterns of GitHub events around DST transitions. In Panel A we show daily events in the three weeks before and after DST and aggregate across sample years, separating users by whether they are in North America, in the EU, or in locations that do not experience DST. A large portion of GitHub activity originates in recognizable technology centers, many of which are located in the United States. Thus, activity in North America is higher across the sample period. This is also evident in Panel B, where we plot hourly events. As a general rule, activity is also higher during traditional working hours of 9am-to–5pm, Monday through Friday—a dip in activity during traditional lunch hours is also apparent.

3 Results

To identify the effect of Daylight Saving Time on activity we employ an event-study design. We first consider specifications at the daily level, followed by similarly constructed models at the hourly level. Daily estimates provide insight into the size and duration of any level-decreases in labor activity associated with DST. Hourly estimates, on the other hand, can speak to within-day variation in user activity—this will prove important to understanding the full extent to which transitions to DST are disruptive to productivity, inducing a substitution of activity across hours of the day. In both environments, however, we are interested in identifying the potential changes in activity around the initiation of DST in treatment locations—their dynamics, in particular. In related literatures that assess impacts of DST transitions on outcomes, much of the identifying variation arrises across time. For example, control groups often consist of only parts of Arizona and Hawaii. In such cases, treatment is often captured by estimating a regression discontinuity

[8] In Figure S1 we produce estimates separately for users in North American and European Union cities with users in “never-treated” cities as controls—results are robust to this distinction.

in time (RDiT) (Bünnings and Schiele, 2021; Doleac and Sanders, 2015; Smith, 2016; Toro et al., 2015). Given the worldwide nature of GitHub, however, there is considerable cross-sectional variation available to us in estimating the change in activity induced by treatment—our modeling will include both variation in the timing of treatment among those who experience DST and and a "never treated" group where GitHub activity should be unperturbed by DST.

3.1 Day-level analysis

In estimating the effect of DST on daily activity we estimate

$$\begin{aligned}\ln(\text{Events}_{cyd}) \;=\; \alpha \;+\; \sum_{d=-21}^{-2} \beta_d \;+\; \sum_{d=0}^{21} \beta_d\, \text{DST}_{cyd} \;+\; \psi_c \;+ \\ \lambda_y \;+\; \delta_d \;+\; X_{cyd} \;+\; u_{cyd}\;,\end{aligned} \tag{1}$$

where $\ln(\text{Events}_{cyd})$ is the log number of events in city c in year y on day d. With the day immediately before DST as the excluded category, this allows for estimated $\hat{\beta}_d$ to be interpreted as percentage differences relative to the day before DST. In estimating (1) we also include city fixed effects (ψ_c) to absorb any level differences in average events across cities, day-of-year (δ_d) fixed effects to absorb differences over time and work-week seasonality, and year fixed effects (λ_y) to absorb any average difference in activity across years in our sample. In X_{cyd} we also include controls for day-of-the-week and for national holidays, capturing variation in activity that should not be attributed to DST when likely due to holidays that coincide with Sunday transitions.[9] Errors are captured in u_{cyd}. Given the staggered timing of DST, we follow Sun and Abraham (2021), though the qualitative results do not vary with this consideration.

In Figure 2 we report estimates of (1) for a pooled model (i.e., across all days). However, patterns in Figure 1 suggest that day-of-week effects may vary by city—some cities exhibit more or less weekend activity, for example. We are inclined, then, to allow day-of-week effects

[9] As Easter coincides with treatment in the European Union in 2013 and in 2016, we also differentiate Easter from other holidays. Likewise, Presidents' Day falls within the three-week sample in six out of the seven years of the US cities. (In 2016, President's Day is 27 days prior to DST.)

to vary by city. This is empirically challenging, however, as DST falls on Sundays, which precludes the estimation of separate day-of-week effects by city. To allow for the same flexibility, then, we estimate separate models for each day of the week, which we report in Panel B. (Stratifying the sample in this way allows the city fixed effect to absorb any city-specific day-of-week effects.) For ease of interpretation, we present each of the seven models together in Panel B. With each model using the same day in the week prior as the omitted category, the results are interpreted as percentage differences relative to the same day (e.g., Monday, Tuesday) the week before the initiation of DST.

In the end, in daily measures of activity we find a short-lived decline associated with DST. It is only the first Sunday of DST itself and the Monday following DST for which we can distinguish activity from baseline levels. That said, these effects are economically meaningful, suggesting reductions in the number of events on the public GitHub timeline on the order of 10 percent ($p < 0.0001$, $\sigma = 0.1$) on the Sunday and 4 percent ($p < 0.005$, $\sigma = 0.04$) on Monday following treatment.[10] Results are robust when estimating with a Poisson model. A similar regression in levels produces comparable results—statistically significant estimates of a decline of 72.27 and 51.07 daily events per city for Sunday and Monday respectively.

While longer-lasting effects are statistically indistinguishable from zero, we note a distinct "Sunday effect" in the data following DST, which we interpret as a response to the extra hour of daylight that is coincident with DST inducing more of a weekend out of users in treated cities relative to those in control cities.

In Figure 3 we restrict our sample in two ways, each having the potential to inform how we interpret the systematic relationship between DST and GitHub activity. In Panel A we consider

[10] To ensure this result is not merely reflecting the mechanical loss of one hour on treated Sundays, we impute the activity of this missing hour by adding the number of events that occurred during the same hour in the same city in the same year on the Sunday immediately prior to DST. Without this imputation, which induces a mechanical decline in treated Sundays, the measurable impact for the first treated Sunday is a reduction in activity by 11% percent (instead of 10).

events that occur between 9am and 7pm—these are the most active hours of the day, accounting for 69 percent of activity. Under this restriction we find a similar pattern—a 14 percent reduction in the number of events ($p < 0.0001$, $\sigma = 0.13$) on Sunday and 5 ($p < 0.05$, $\sigma = 0.04$) on the Monday following treatment. In Panel B we restrict observations to only "pushes."[11] Again, we find similar patterns in the number of events—declines of 12 percent ($p < 0.0001$, $\sigma = 0.12$) on Sunday and 3 percent ($p < 0.05$, $\sigma = 0.04$) on the Monday following treatment.

3.2 Hour-level analysis

In Figure 4 we demonstrate the patterns in activity visually, focussing on one week before and after DST transitions. In Panel A we plot weekly mean event counts separately for each group (i.e., cities in North America, in the EU, and those never treated) across each hour of the day. In Panel B we plot the difference in means, subtracting hourly averages by treatment status separately for pre- and post-treatment period. In Panel C we plot the "second difference" in means, subtracting the post-treatment differences by pre-treatment differences (again, separately by group). This suggests declines in GitHub activity in the traditional morning work hours.

To formally model *hourly* GitHub events around transitions to DST we estimate

$$\begin{aligned} \ln(\text{Events}_{cydh}) \ = \ \alpha \ + \ \sum_{d=-21}^{-2} \beta_d \, \text{DST}_{cydh} \ + \ \sum_{d=0}^{21} \beta_{dh} \, \text{DST}_{cydh} \ + \ \psi_c \ + \\ \lambda_y \ + \ \delta_d \ + \ \sigma_h \ + \ X_{cyd} \ + \ u_{cydh} \ , \end{aligned} \tag{2}$$

with $\ln(\text{Events}_{cydh})$ now defined at the city by year by day by hour (h). In daily activity, DST always falling on Sundays prohibited the estimating of separate day-of-week effects by city. To flexibly fit hour effects in 2 we model each hour of the day separately for each day of the week.

In Figure 5 we present 24 plots, corresponding to each hour of the day—each plot can then be interpreted as the estimated differences in hourly productivity relative to the same hour and

[11] Recall that "pushes" are among all events in the earlier analysis, but arguably represent the type of event type that we would think is most representative of productivity as they encompass all changes uploaded to a remote copy.

day in the week prior to DST. Unlike daily estimates, which identify one to two days of significant decreases in activity, the analysis of hourly data reveals longer-lived declines in activity following DST—this is particularly evident in the traditional morning work hours. Relative to baseline activity, the initiation of DST induces significant declines in the 8am, 9am, and 10am hours for roughly two weeks. (Declines in the 8am hour are distinguishable from zero for 12 days, in the 9am hour for 16 days, and in the 10am hour for four days.) In magnitude, activity decreases 28 to 48 percent, and remains evident in the data for roughly two weeks. Without evidence of any systematic increases in activity elsewhere in the day, the overall patterns of daily activity are consistent with users adopting a variety of strategies to cope with the transition.

To synthesize the variation in hourly estimates and better understand the nature of users' transitions to DST, in Figure 6 we consider whether there are discernible patterns in the point estimates from the hourly analysis. Separately for each week, we ask whether the point estimates in Figure 5 exhibit any seasonality across hours of the day. On the left of Figure 6, then, we re-align those same point estimates to instead represent the 168 hours in each week leading into and out of DST. On the right, we then present the hourly variation that is identifiable in that week (i.e., one estimated parameter for each hour of the day).[12] This context makes clear that the "shape" of what is normal variation in hourly activity changes with the arrival of DST. While no systematic patterns exists across hour of the day in the weeks prior to DST, with DST's arrival the systematic declines in activity in the morning work hours are again evident. Immediately following DST, magnitudes of the decline are greatest in the 9am hour, with a 34 percent reduction. Further, this pattern attenuates in the second and third weeks following DST, as workers arguably recover from the disruption to time. Reductions remain largest at 9am but reduce in magnitude to 26 percent and 17 percent in the second and third week following DST.

[12] In each, we've omitted 4am as the reference category.

4 Conclusion

Even though Daylight Saving Time changes the social clocks of over one billion people each year, concerns over the associated economic and social welfare implications continue to be debated. By considering active users of a popular cloud based version-control platform, we exploit variation in the application of DST across technology centers around the world to estimate the productivity implications of shifting clock time this way.

At the daily level, we find significant level declines in the recorded activity of GitHub users who experience DST—on the Sunday of DST itself, and the Monday immediately following the transition to DST. While this suggest that the productivity declines are over quickly, there are two ways that the data support the belief that the transition is more costly to productivity. First, we find persistent "Sunday effects" in treated cities—while daily activity returns to baseline following two-day declines, GitHub activity is persistently lower on Sundays. Second, hourly analysis reveals much-longer-lasting disruptions—morning declines are evident for upwards of two weeks following the initiation of DST.

References

Barnes, C. M. and D. T. Wagner (2009). Changing to daylight saving time cuts into sleep and increases workplace injuries. *Journal of Applied Psychology 94*(5), 1305.

Belzer, D. B., S. W. Hadley, and S.-M. Chin (2008). Impact of extended daylight saving time on national energy consumption. *US Dept. of Energy*.

Berk, M., S. Dodd, K. Hallam, L. Berk, J. Gleeson, and M. Henry (2008). The effect of daylight saving. *Sleep and Biological Rhythms 6*(1), 22–25.

Bünnings, C. and V. Schiele (2021). Spring forward, don't fall back: The effect of daylight saving time on road safety. *Review of Economics and Statistics 103*(1), 165–176.

Czeisler, C. A., J. F. Duffy, T. L. Shanahan, E. N. Brown, J. F. Mitchell, D. W. Rimmer, J. M. Ronda, E. J. Silva, J. S. Allan, J. S. Emens, et al. (1999). Stability, precision, and near-24-hour period of the human circadian pacemaker. *Science 284*(5423), 2177–2181.

Doleac, J. L. and N. J. Sanders (2015). Under the cover of darkness: How ambient light influences criminal activity. *Review of Economics and Statistics 97*(5), 1093–1103.

Gregory-Allen, R. B., B. Jacobsen, and W. Marquering (2010). The daylight saving time anomaly in stock returns: Fact or fiction? *Journal of Financial Research 33*(4), 403–427.

Kamstra, M. J., L. A. Kramer, and M. D. Levi (2000). Losing sleep at the market: The daylight saving anomaly. *American Economic Review 90*(4), 1005–1011.

Kantermann, T., M. Juda, M. Merrow, and T. Roenneberg (2007). The human circadian clock's seasonal adjustment is disrupted by daylight saving time. *Current Biology 17*(22), 1996–2000.

Kotchen, M. J. and L. E. Grant (2011). Does daylight saving time save energy? evidence from a natural experiment in Indiana. *Review of Economics and Statistics 93*(4), 1172–1185.

Kountouris, Y. and K. Remoundou (2014). About time: daylight saving time transition and individual well-being. *Economics Letters 122*(1), 100–103.

Kuehnle, D. and C. Wunder (2016). Using the life satisfaction approach to value daylight savings time transitions: evidence from britain and germany. *Journal of Happiness Studies 17*(6), 2293–2323.

Lahti, T., J. Sysi-Aho, J. Haukka, and T. Partonen (2011). Work-related accidents and daylight saving time in finland. *Occupational Medicine 61*(1), 26–28.

Lahti, T. A., S. Leppämäki, J. Lönnqvist, and T. Partonen (2006). Transition to daylight saving time reduces sleep duration plus sleep efficiency of the deprived sleep. *Neuroscience Letters 406*(3), 174–177.

McDermott, G. R. and B. Hansen (2021). Labor reallocation and remote work during COVID-19: Real-time evidence from GitHub.

Monk, T. H. and L. C. Aplin (1980). Spring and autumn daylight saving time changes: Studies of adjustment in sleep timings, mood, and efficiency. *Ergonomics 23*(2), 167–178.

Nordhaus, W. D. (1997). Do real-output and real-wage measures capture reality? The history of lighting suggests not. *The Economics of New Goods 58*, 29–66.

Smith, A. C. (2016). Spring forward at your own risk: Daylight saving time and fatal vehicle crashes. *American Economic Journal: Applied Economics 8*(2), 65–91.

Sun, L. and S. Abraham (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics 225*(2), 175–199.

Toro, W., R. Tigre, and B. Sampaio (2015). Daylight saving time and incidence of myocardial infarction: Evidence from a regression discontinuity design. *Economics Letters 136*, 1–4.

Wagner, D. T., C. M. Barnes, V. K. G. Lim, and D. L. Ferris (2012). Lost sleep and cyberloafing: Evidence from the laboratory and a daylight saving time quasi-experiment. *Journal of Applied psychology 97*(5), 1068–76.

Figure 1: Mean number of events, by treatment status

A: Total events per day, ± 3 weeks to DST

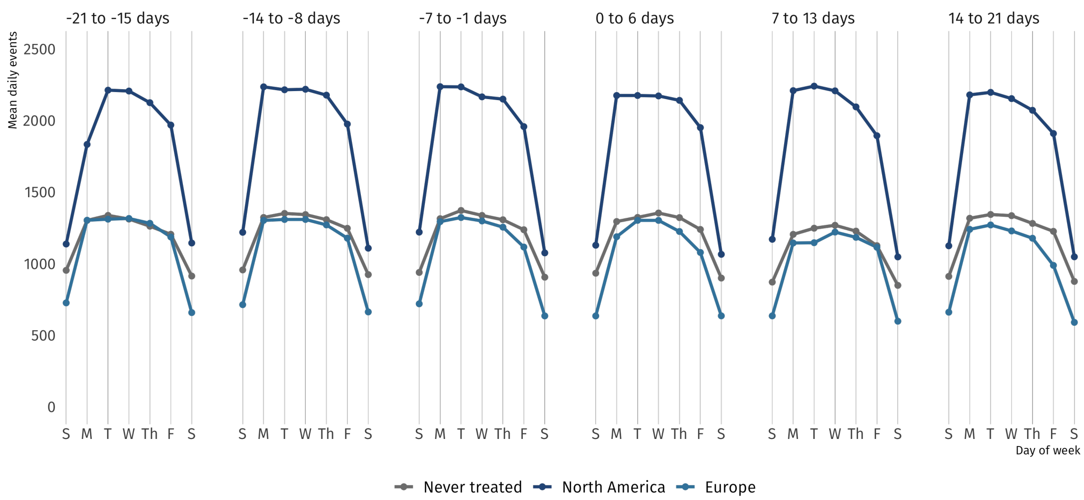

B: Total events per hour, ± 1 weeks to DST

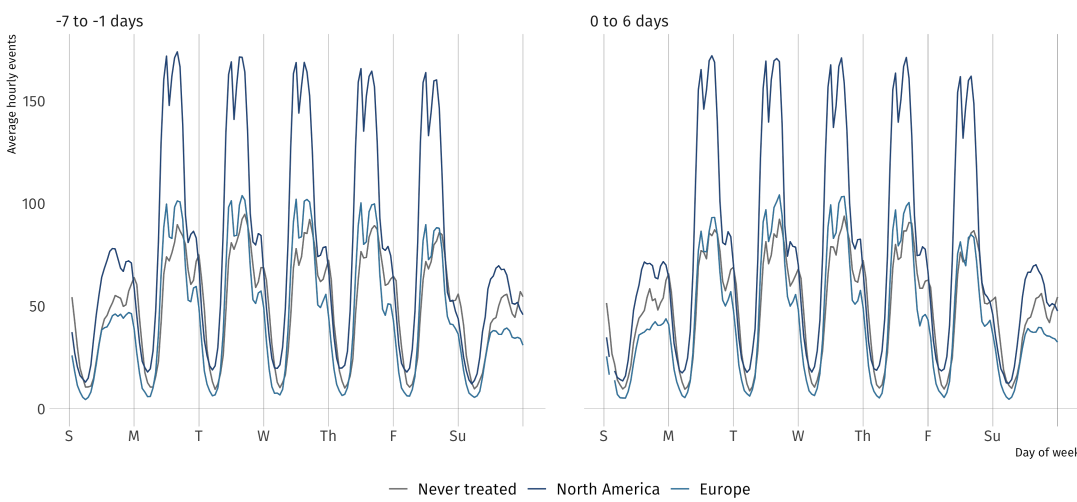

Notes: North America includes the treated cities in the United States and Canada. Europe includes the treated cities in the European Union. See Table S3 for specific dates associated with DST. (In the "-21 to -15" cell of Panel A, the lower level of activity on Monday in North America is attributable to President's day.)

Figure 2: Does DST induce a change in the total number of daily GitHub events?

A: Pooled model (across day of week) normalized to the Saturday prior to treatment

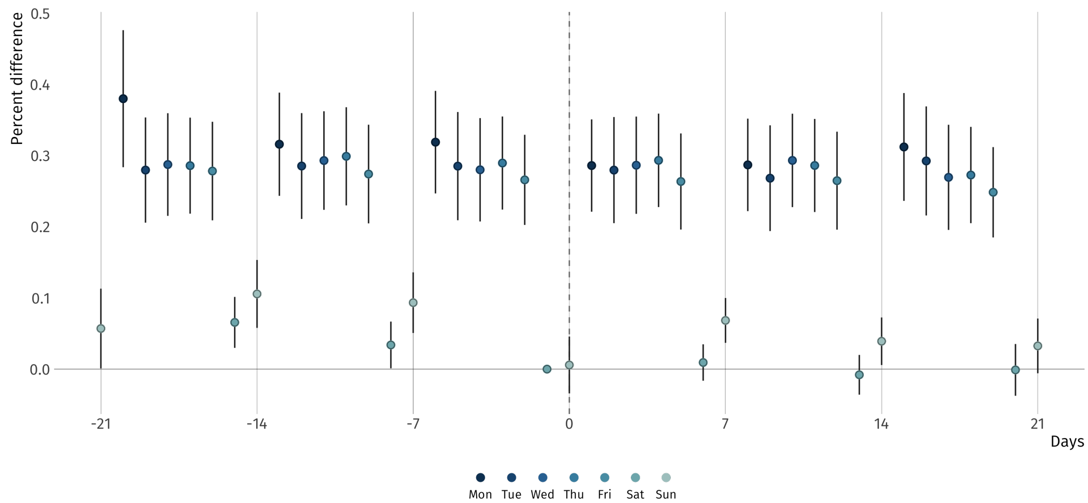

B: Separate models (by day of week), with events normalized to the week prior to treatment

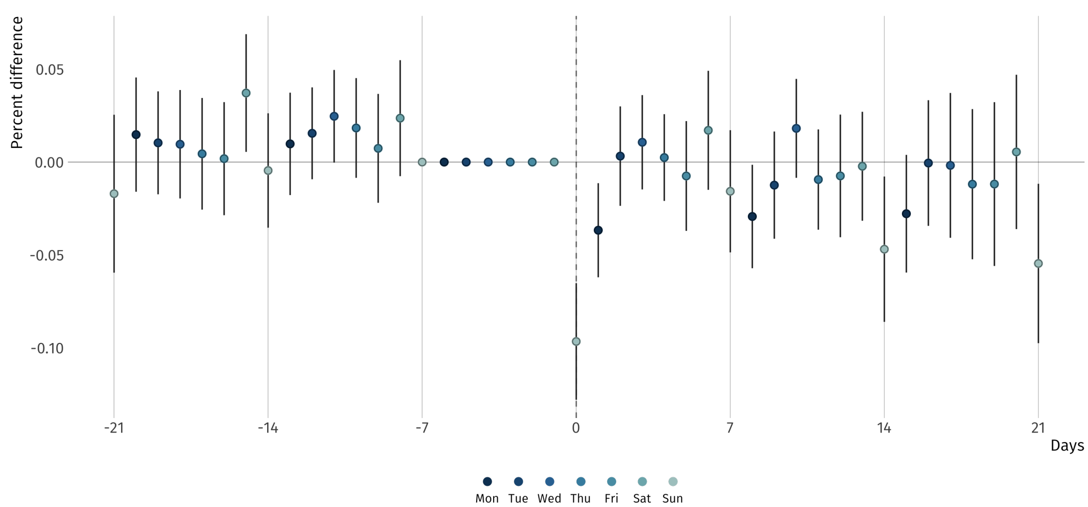

Notes: The sample includes the observations of the total number of GitHub events (logged) three weeks before and after DST ($n = 43 \times 50 = 2,150$, representing 141,899 unique users). In Panel A we estimate one model, using a pooled sample of all days and leaving the Saturday prior to treatment as the omitted category. We estimate city, day-of-week, day-of-year, year, and holiday fixed effects. With evidence that day-of-week effects vary by city (e.g., some cities exhibit more or less weekend activity), we are inclined to allow day-of-week effects to vary by city. However, as treatment always falls on a Sunday we cannot estimate separate day-of-week effects by city. In Panel B we therefore estimate seven separate models (i.e., one model for each day of the week) where the city fixed effect absorbs day-of-week differences.

Figure 3: Do we see different responses within working hours, or if we restrict events to “pushes” only?

A: Total events during work hours (9a-5p)

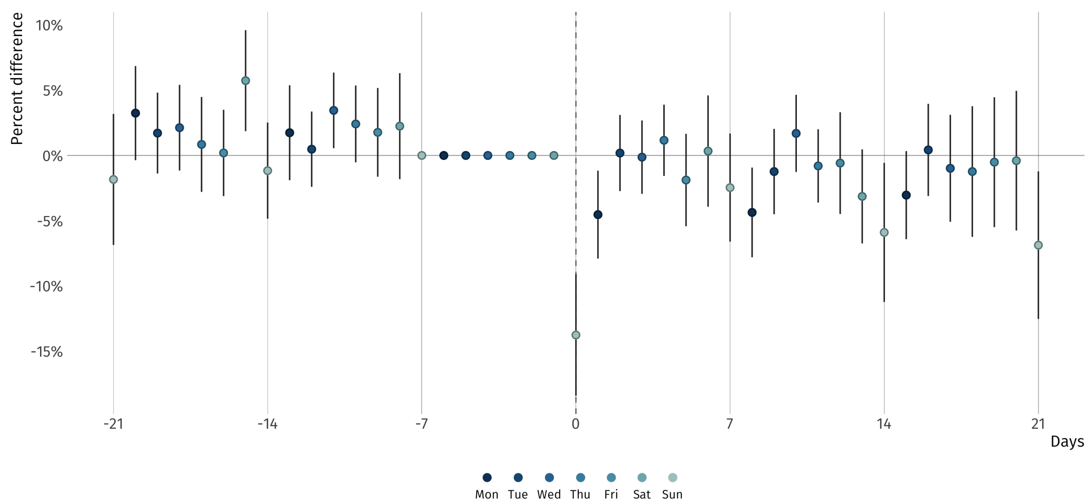

B: Pushes only (all hours)

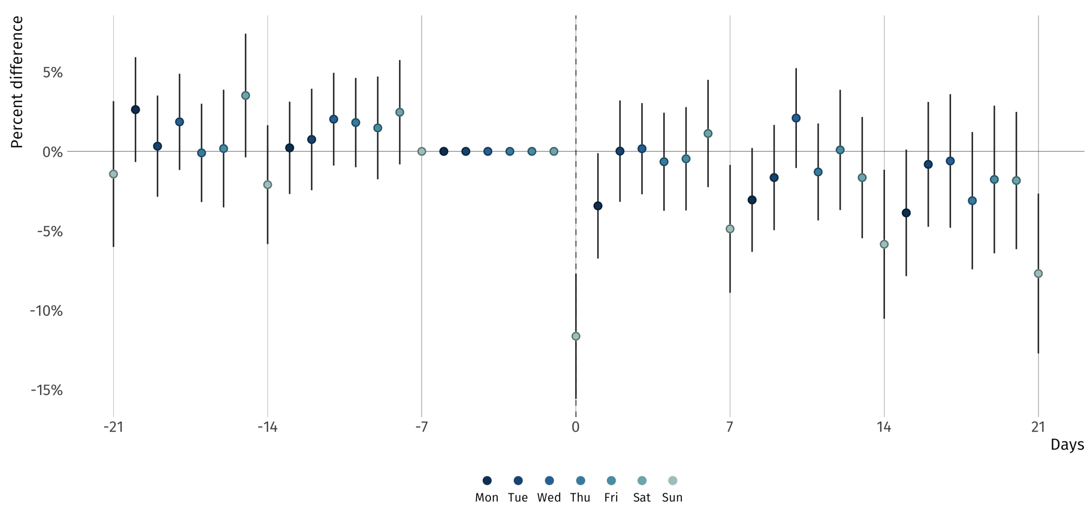

Notes: The sample includes the observations three weeks before and after DST ($n = 43 \times 50 = 2,150$). In both panels we estimate seven separate models (i.e., one model for each day of the week). In Panel A the sample includes the total number of GitHub events (logged) three weeks before and after DST, restricted to those falling within the hours of 9am to 5pm. In Panel B the sample includes the number of pushes (logged) three weeks before and after DST.

Figure 4: Constructing hourly difference-in-differences estimates

A: Mean counts of total events by treatment status, DST ±1 week

(i) North America and the never treated — Pre — Post

(ii) EU and the never treated — Pre — Post

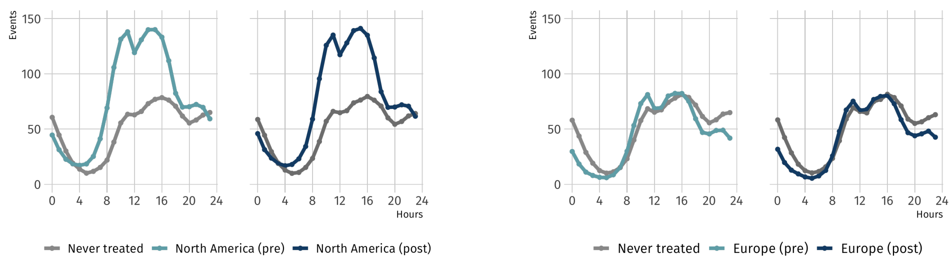

B: Post- and pre-DST differences in total number of events ("treatment – control")

(i) North America

(ii) European Union

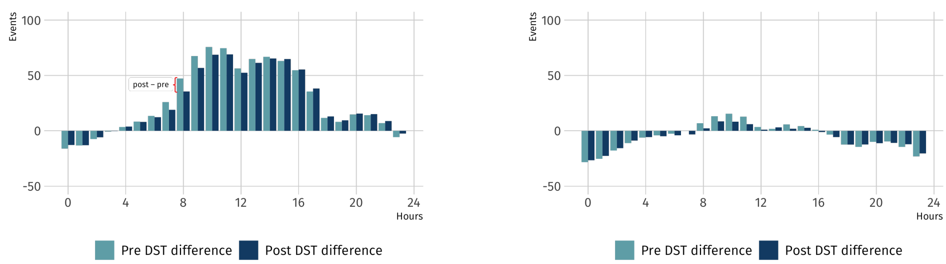

C: Difference in differences

(Post-DST "treatment – control" – Pre-DST "treatment – control")

(i) North America

(ii) European Union

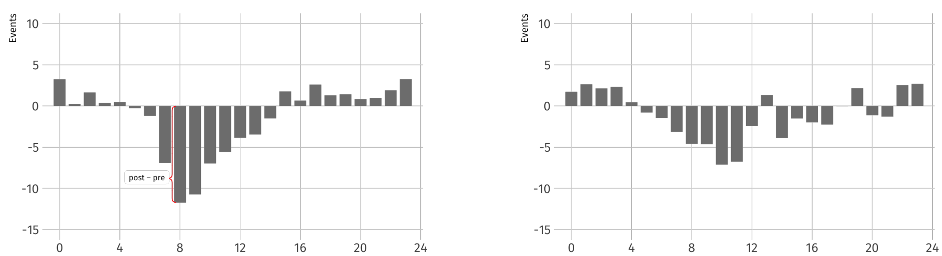

Notes: In Panel A we plot the mean number of hourly events in the one week before and one week after DST, separately for North America and Europe. In Panel B we plot the difference in means, subtracting hourly averages by treatment status separately for pre- and post-treatment period. In Panel C we plot the "second difference" in means, subtracting the post-treatment differences by pre-treatment differences

Figure 5: Hour-specific estimates of the effect of DST on GitHub events

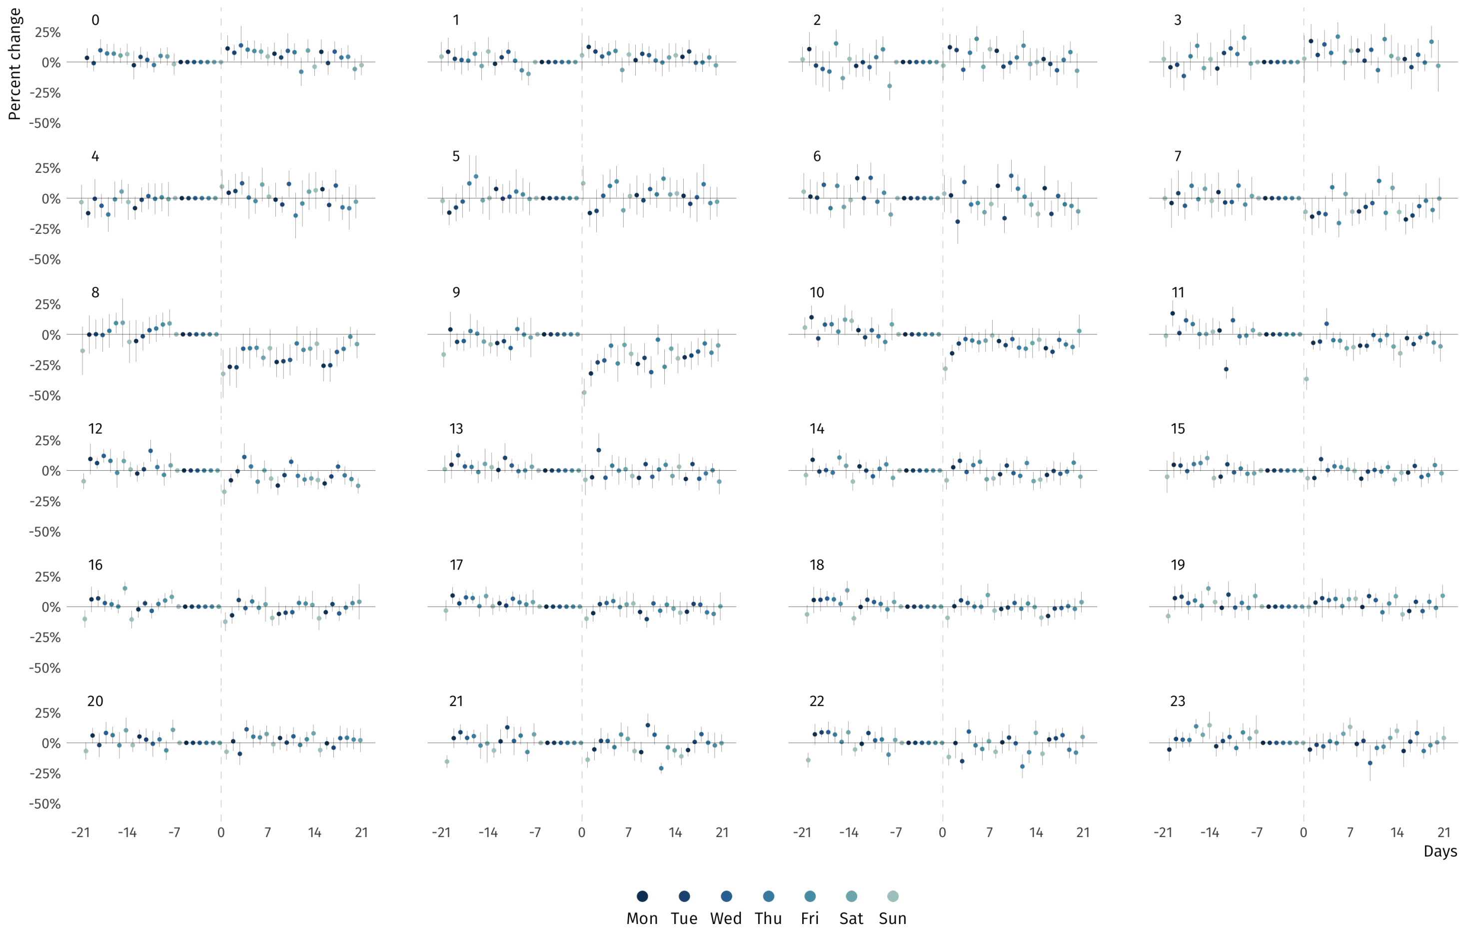

Notes: The sample includes the observations of the total number of GitHub events (logged) three weeks before and after DST ($n = 43 \times 50 \times 24 = 51,600$, representing 141,899 unique users).

Figure 6: Does the pattern of hourly events change around DST?

A: -21 to -15 days to DST

(i) Hour-specific estimates

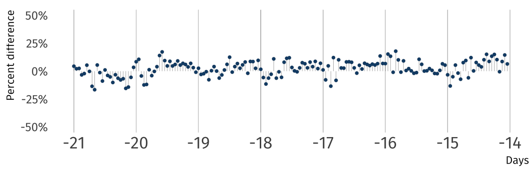

(ii) Average activity by hour

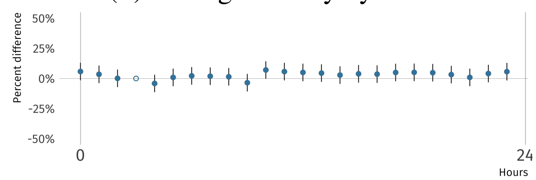

B: -14 to -8 days to DST

(i) Hour-specific estimates

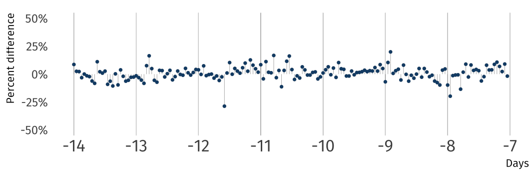

(ii) Average activity by hour

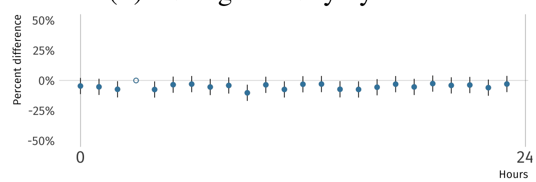

C: 0 to 6 days from DST

(i) Hour-specific estimates

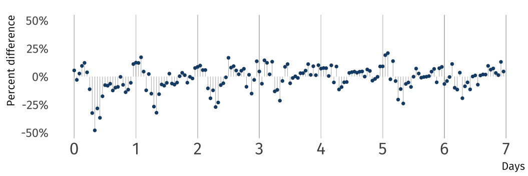

(ii) Average activity by hour

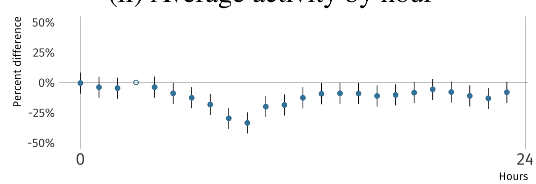

D: 7 to 13 days from DST

(i) Hour-specific estimates

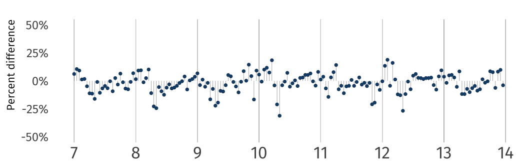

(ii) Average activity by hour

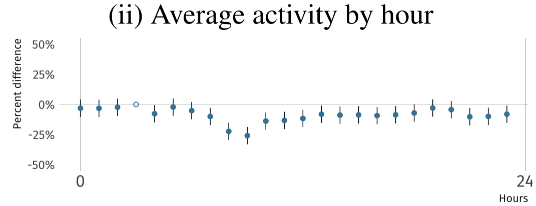

E: 14 to 20 days from DST

(i) Hour-specific estimates

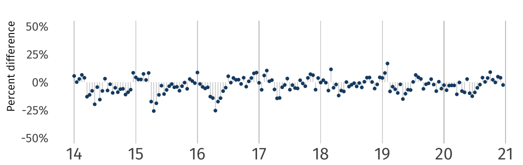

(ii) Average activity by hour

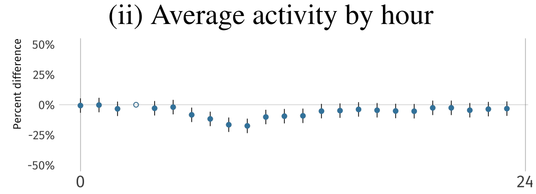

Notes: Point estimates in (i) are from corresponding hours in Figure 5. In (ii), we estimate the average hourly activity in the corresponding week.

Supplementary materials

Table S1: GitHub events with proportional breakdowns

Name	Description	Proportion
Push	One or more commits are pushed to a repository branch or tag	0.424
Issue comment	Activity related to an issue or pull request comment	0.145
Watch	When someone stars a repository	0.107
Create	A Git branch or tag is created	0.083
Pull request	Activity related to pull requests	0.083
Issue	Activity related to an issue	0.055
Pull request review	Activity related to pull request review comments in the pull request's unified diff	0.044
Delete	A Git branch or tag is deleted	0.031
Fork	A user forks a repository	0.014
Gollum	When a wiki page is created for a repository	0.006
Release	Activity related to a release	0.005
Member	Activity related to repository collaborators	0.002
Public	When a private repository is made public	0.001

Table S2: Top 50 GitHub cities, by total events in 2013–2019

Rank	City	Events	Pushes	Early / Late / Never DST
1	San Fransisco Bay, US [a]	24,661,994	9,866,982	Early
2	London, GB	11,899,803	4,921,165	Late
3	Seattle / Redmond, US	10,603,863	4,220,038	Early
4	New York, US	10,134,305	4,491,422	Early
5	Berlin, DE	7,505,479	2,965,833	Late
6	Tokyo, JP	6,725,725	3,002,795	Never
7	Paris / Paris, FR	6,440,234	2,743,032	Late
8	Boston / Cambridge, US	5,758,715	2,452,179	Early
9	Los Angeles, US	4,088,571	1,671,551	Early
10	Toronto, CA	3,785,193	1,657,969	Early
11	Washington / Baltimore, US	3,674,995	1,661,466	Early
12	Portland, US	3,472,755	1,517,178	Early
13	Moscow, RU	3,352,144	1,477,497	Never
14	Austin, US	3,244,818	1,372,864	Early
15	Denver / Boulder, US	3,055,455	1,316,222	Early
16	Chicago, US	3,047,534	1,384,378	Early
17	Amsterdam, NL	2,643,473	1,150,170	Late
18	Vancouver, CA	2,482,046	1,028,486	Early
19	Montreal, CA	2,325,379	1,015,984	Early
20	Bengaluru, IN	2,272,297	1,014,638	Never
21	Stockholm, SE	2,268,401	1,013,413	Late
22	Munich, DE	2,233,769	917,433	Late
23	Raleigh, US	2,040,004	660,386	Early
24	Madrid, ES	2,031,618	823,199	Late
25	Zürich, CH	2,018,204	862,027	Late
26	Seoul, KR	1,746,734	837,688	Never
27	Prague, CZ	1,740,558	733,109	Late
28	Philadelphia, US	1,728,980	741,104	Early
29	Hamburg, DE	1,666,523	732,504	Late
30	Kyiv, UA	1,624,576	723,113	Late
31	San Diego, US	1,617,850	823,158	Early
32	Oslo, NO	1,484,586	657,178	Late
33	Cambridge, GB	1,471,832	681,485	Late
34	Atlanta, US	1,464,569	686,143	Early
35	Saint Petersburg, RU	1,461,972	630,312	Never
36	Pittsburgh, US	1,439,968	747,827	Early
37	Copenhagen, DK	1,438,777	580,865	Late
38	Lausanne / Geneva, CH	1,333,881	552,245	Late
39	Warsaw, PL	1,274,021	556,611	Late
40	Tulsa, US	1,221,202	480,833	Early
41	Helsinki, FI	1,163,743	534,872	Late
42	Cologne, DE	1,139,073	499,945	Late
43	Vienna, AT	1,132,286	519,615	Late
44	Taipei, TW	1,113,622	492,575	Never
45	Minneapolis, US	1,097,682	466,351	Early
46	Lyon, FR	1,027,755	442,294	Late
47	Ottawa, CA	1,003,609	419,492	Early
48	Brno, CZ	973,900	379,024	Late
49	Dallas, US	953,706	419,439	Early
50	Budapest, HU	936,181	444,574	Late

[a] Mountain View, Oakland, San Jose, Palo Alto, Berkeley.

Table S3: The timing of Daylight Saving Time across treated locations

Year	North American	European Union
2013	March 10	March 31
2014	March 09	March 30
2015	March 08	March 29
2016	March 13	March 27
2017	March 12	March 26
2018	March 11	March 25
2019	March 10	March 31

Figure S1: Do we see different responses across treatment groups?

A: Total events, North America and never treated

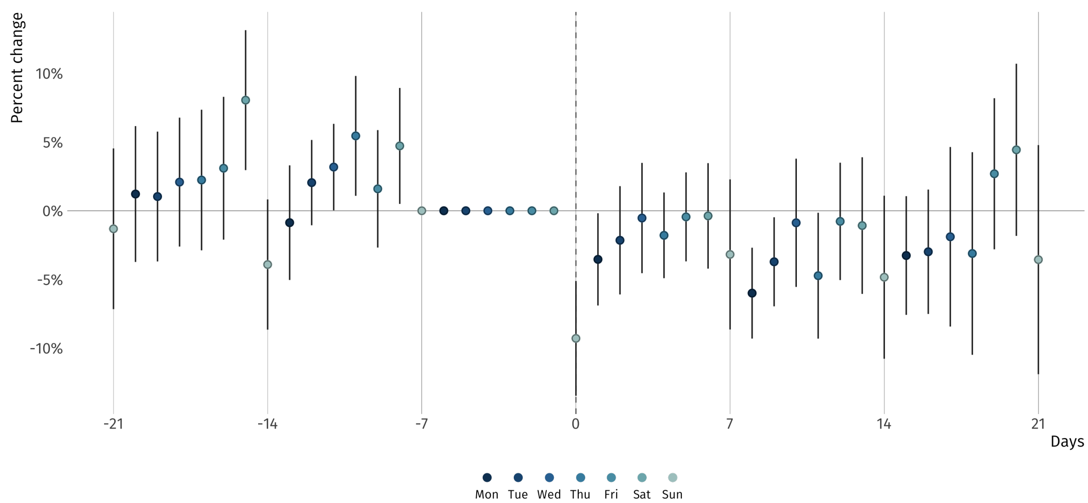

A: Total events, European Union and never treated

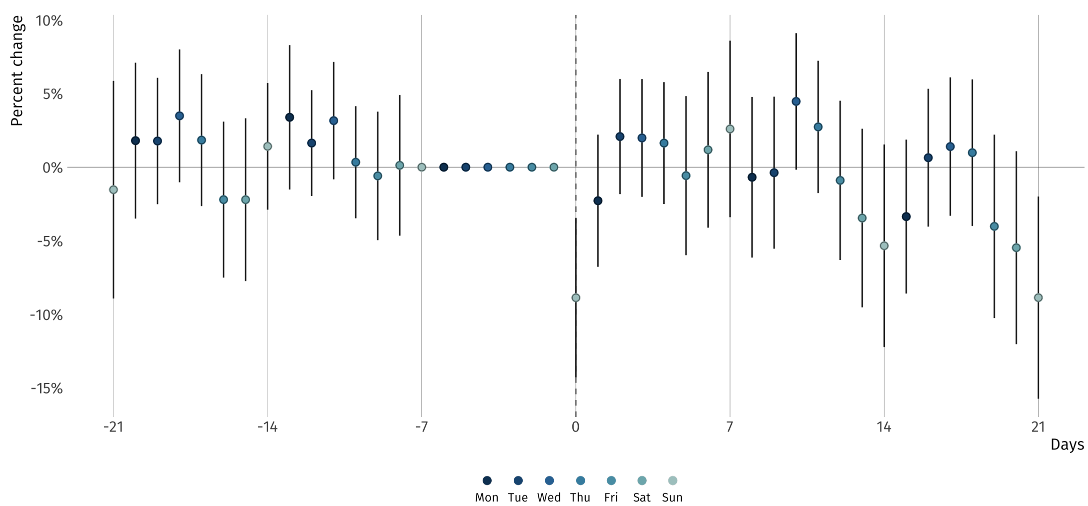

Notes: In both panels we estimate seven separate models (i.e., one model for each day of the week). In Panel A the sample includes the total (logged) number of GitHub events three weeks before and after DST, restricted to users in North American and in never-treated cities (see Table S2). In Panel B the restriction is to users in European Union cities and in the never-treated cities.