

# Taxing dissent: The impact of a social media tax in Uganda

Levi Boxell, *Stanford University*\*  
Zachary Steinert-Threlkeld, *UCLA*

September 2019

## Abstract

We examine the impact of a new tool for suppressing the expression of dissent—a daily tax on social media use. Using a synthetic control framework, we estimate that the tax reduced the number of georeferenced Twitter users in Uganda by 13 percent. The estimated treatment effects are larger for poorer and less frequent users. Despite the overall decline in Twitter use, tweets referencing collective action increased by 31 percent and observed protests increased by 47 percent. These results suggest that taxing social media use may not be an effective tool for reducing political dissent.

---

\*E-mail: lboxell@stanford.edu and zst@luskin.ucla.edu. We thank Nicholas Bloom, Matthew Gentzkow, Guy Grossman, William Hobbs, Matthew O. Jackson, Jenn Larson, Janet Lewis, Vianney Mbonigaba, Nicholas Obradovich, Chiara Pasquini, Margaret Roberts, Jesse M. Shapiro, and Chuan Yu for their suggestions and comments, along with participants at Stanford University. Funding was provided by the National Science Foundation (grant number: DGE-1656518) and the Institute for Humane Studies.

# 1 Introduction

Autocratic regimes have become increasingly interested in digital censorship, and new strategies for controlling online information are rapidly spreading. After tests of internet shutdowns in the Xinjiang province of China in 2009, governments shut down internet access nearly 190 times in 2017 (Griffiths 2019). China has also been actively exporting its digital repression strategies to other countries. In 2017, officials from Uganda met with members of a Chinese government subsidiary to build capacity to monitor and prevent social media abuse (Mugerwa 2017). Later, Uganda innovated on the set of tools available to authoritarian regimes by implementing a 200 shilling daily tax on users of social media.<sup>1</sup> Critics view the tax as an attempt to suppress political dissent ahead of upcoming elections, and even the Ugandan President has stated that reducing “gossip” was an intended outcome of the tax (Crabtree 2018).

In this paper, we examine the impact of the social media tax on social media use and levels of collective action in Uganda. Using georeferenced tweets from across sub-Saharan Africa and exploiting the unexpected implementation of the social media tax, we compare outcomes in Uganda relative to a synthetic control composed of non-treated countries with similar pre-treatment trends (Abadie et al. 2010).

There was a large and sustained decrease in the number of active users on Twitter after the tax. The tax decreased the number of unique Twitter users in each period by 13 percent on average. New account registrations also dropped by 20 percent. Poorer and less frequent users exhibited a larger change in behavior, with a 17 percent decrease in non-Apple users (a proxy for wealth) and an 18 percent decrease in users who tweet less than once per day. These estimated effects are large relative to a placebo distribution of non-treated countries.

Despite the reduction in overall social media use, tweets referencing collective action increased by 31 percent in the period following the implementation of the tax. In contrast, there was no meaningful increase in the overall level of tweets or tweets using political phrases not related to collective action.<sup>2</sup> The period after the tax’s implementation also experienced a 17 percent increase in the number of unique users discussing collective action. This increase in

---

<sup>1</sup>The tax is equivalent to \$1.60 per day for the United States (normalized by GDP per capita PPP) and applies to over 50 social media applications including Facebook, WhatsApp, and Twitter.

<sup>2</sup>Increasing the availability of information on collective action is particularly concerning for authoritarian regimes. While the impact of increasing the flow of information on grievances against the government is theoretically ambiguous depending on citizen expectations, information on the events themselves nearly always increases the likelihood of successful coordination (Little 2016). Ananyev et al. (2017) make a related argument that less competent governments will be more likely to invest in coordination censorship than content censorship. Observed behavior of Chinese censorship corroborates the focus on collective action information (King et al. 2013).

discussions of collective action online is mirrored by increases in physical protests and riots. There was a 47 percent increase in observed protests and riots in the subsequent period in Uganda relative to a synthetic control. Again, these estimated effects are large relative to the placebo distribution of non-treated countries.

The tax could increase collective action through two potential mechanisms.

First, it could alter views of or beliefs about the regime. The high salience of the tax as a form of digital repression may have decreased coordination costs for collective action by creating a focal point or providing common information (Chwe 2000; Tucker 2007; Egorov et al. 2009; Chwe 2013). Alternatively, the signal from the tax could have shifted the perceived returns to successful protest and regime change. Qualitative evidence from news coverage of the tax and the language used by oppositional leaders on Twitter corroborate the tax's role as a rallying point for collective action, and previous work has also suggested that the salience of repression is an important determinant of the overall effect on dissent (Roberts 2018).

Second, the tax could increase collective action by altering the composition or structure of the network. In network models of collective action, an individual's decision to participate often depends on the extent to which users within her *local* network participate (Chwe 2000; Watts 2002; Barbera and Jackson 2019).<sup>3</sup> If the likelihood of participating in collective action is positively correlated with willingness to pay for Twitter, then the tax would increase the density of activists on Twitter, increasing the likelihood of collective action.<sup>4</sup> We find some empirical support for this positive correlation, as users who are less sensitive to the tax (users with Apple devices) are also more likely to tweet about collective action in the pre-intervention period.

There are two important caveats of this study. First, we study the short-run effect the tax, and the long-run effects may be different. This is particularly true if the salience of the informational content of the tax decays. Second, we can not exclude the possibility that other concurrent changes are driving the observed trends rather than the tax itself. That being said, over 30 percent of collective action tweets contained the string "tax" in the 10 days after the tax's implementation, which, along with the qualitative evidence on elite behavior, is suggestive

---

<sup>3</sup>Local network participation can be important for inference about global participation, information on the event itself, and consumption value from social network participation. The importance of local network participation in determining individual turnout for protests is also highlighted in recent empirical work (Burzstyn et al. 2019).

<sup>4</sup>See the Online Appendix for a model highlighting this relationship. Previous work has also pointed out that the relationship between group size and collective action varies depending on other forms of network structure (Oliver and Marwell 1988; Gould 1993; Siegel 2009). Hassanpour (2014) makes a related argument that full connectivity can hinder collective action in threshold models. And, in models of information overload, a small tax on sending messages will filter out low-value messages—potentially increasing consumption of information on collective action (Van Zandt 2004; Anderson and de Palma 2009; Iyer and Katona 2016).

of a backlash effect of the tax.

Among autocratic regimes, China's digital censorship and repression is most studied. It has developed a sophisticated regime that focuses on censoring information on collective action and actively produces information to distract from sensitive topics (King et al. 2013, 2014, 2017). China also allows some forms of dissent on social media to assist in monitoring public opinion and the performance of local officials (Qin et al. 2017). Long-run exposure to censorship in China has reduced the incentives to acquire uncensored information (Chen and Yang 2017). However, the sudden and visible censorship of Instagram in China caused citizens to actively seek out sensitive information (Hobbs and Roberts 2018).

In other contexts, disruptions to communication networks correlate with increases in realized dissent (Hassanpour 2014), though they also appear to weaken the ability of organized opposition groups to coordinate (Ghodes 2015). Physical repression against dissenting elites in Saudi Arabia reduced criticism by those directly involved, but did not reduce overall levels of criticism (Pan and Siegel 2018). More generally, the growth of electronic and online social networks tends to increase protest activity (Manacorda and Tesei 2016; Enikopolov et al. 2017; Enikopolov et al. 2018). We show that taxing social media can effectively reduce social media use, but its high salience as digital repression and its impact on the underlying network structure can increase the likelihood of collective action.

This work also relates to the literature on measuring consumer demand and willingness to pay for social media. Previous estimates of willingness to pay for social media (or willingness to accept for quitting) have relied on experimental elicitation (Allcott et al. Forthcoming) or non-incentivized survey responses (Sunstein Forthcoming). Our study is the first to estimate part of the willingness-to-pay distribution for social media using a quasi-experiment in which users face a naturally occurring price for accessing the platforms.

## **2 Ugandan Context**

President Yoweri Museveni has ruled Uganda since 1986. In this span, the elimination of several important constitutional limits on executive rule have allowed the 74-year-old to continue to rule. In 2005, the two-term limit on the presidential office was removed. In 2017, the age limit on the presidential office (75 years) was lifted—leading to a “brawl” in the Ugandan parliament (BBC 2017). Museveni, who largely gathers electoral support from rural and aging populations, has faced increasing difficulty connecting with the youth of Uganda. In contrast,

Robert Kyagulanyi Ssentamu, a Ugandan pop-star with the stage name Bobi Wine, has quickly risen to become the face of the opposition party after becoming a member of parliament in 2017 (Kagumire 2018). Kyagulanyi frequently uses social media, in addition to his music, to speak directly about government failures and the hopes of the Ugandan people.

Within this growing discontent and opposition, Museveni and his government actively seek to control the growth of digital dissent in Uganda. They have arrested dissenting voices on social media for violating the 2011 Computer Misuse Act and shut down social media services on Election Day in 2016 (Mugerwa 2017). In 2017, Ugandan officials met with a Chinese government-owned subsidiary to provide Uganda with “technical capacity to monitor and prevent social media abuse” across the more than 13 million internet users (Mugerwa 2017). Shortly after, the Ugandan Communication Commission announced that they would launch their own versions of popular social media platforms, mirroring a prominent Chinese strategy for monitoring and controlling information flows (Observer 2017). These alternative platforms have yet to materialize.

In May 2018, the Ugandan parliament passed a bill creating a 200 shilling daily tax on the use of any of more than 50 social media applications, along with a tax on mobile money transfers. The tax was implemented on July 1, 2018, and is collected by telecommunication companies via mobile money transfers (Crabtree 2018). The Online Appendix (Figure A1) shows that there was a large spike in news articles referencing the “social media tax” after July 1, 2018 and limited discussion prior to this. In relative terms, the 200 shilling (roughly \$0.05 USD) per day tax is large. The GDP per capita (PPP) for Uganda was \$1,864 in 2017 relative to \$59,532 in the United States, per the World Bank. Annualizing the \$0.05 per day gives a tax burden of the social media tax at roughly 1 percent of Uganda’s GDP per capita. In comparison, the median willingness-to-pay estimate from Sunstein (Forthcoming) for Twitter is \$5.00 per month or 0.1 percent of the United States’ GDP per capita.

Many view the tax as a form of digital repression by Museveni who hopes to limit social media use for the upcoming 2021 presidential election (Crabtree 2018). Protests, led by Kyagulanyi, occurred shortly after the tax’s implementing, and Kyagulanyi, who announced he will also campaign in the 2021 presidential election, has been subsequently arrested multiple times (Al Jazeera 2019).

## 3 Data and Empirical Framework

### 3.1 Data

Our primary data consists of georeferenced tweets from 50 sub-Saharan Africa countries and territories (hereafter, “countries”) collected from the one percent Twitter Streaming API. For each tweet, the API returns the text, the time of the tweet, and various user characteristics at the time of the tweet. We pass a global bounding box to the streaming API, delivering an estimated 1/3 to 1/2 of all georeferenced tweets (Leetaru et al. 2013; Steinert-Threlkeld 2018). The tweets are assigned to their respective countries using the provided geolabelling from Twitter. We collapse observations into ten day periods and exclude countries in the bottom quintile of unique users (e.g., Central African Republic and Eritrea).

For data on collective action events, we merge data on riots and protests as recorded by ACLED and ICEWS (Raleigh et al. 2010; Boschee et al. 2015). Separately, for each dataset, we compute the number of events occurring in a given time period. We then take the average across both datasets for a given country in a given period. We restrict attention to countries that appear in both our ACLED and ICEWS samples.

### 3.2 Empirical Framework

To estimate the impact of the social media tax, we adopt a synthetic control approach (Abadie et al. 2010). Let  $Y_{it}(1)$ ,  $Y_{it}(0)$  be outcome  $Y$  for country  $i$  at time  $t$  with and without the social media tax respectively. Let  $i = 0$  represent Uganda,  $i \in \{1, \dots, N\}$  represent the control countries, and  $t \in \{-T_0, \dots, 0, \dots, T_1\}$  be the timespan considered where 0 indicates the first period in which the social media tax was implemented (July 1, 2018). Each time period represents a 10-day span.

Assume

$$Y_{it}(0) = \delta_t + \lambda_t \cdot \mu_i + \epsilon_{it} \quad (1)$$

where  $\delta_t$  is a constant factor across countries,  $\mu_i$  is a vector of unobserved factors for each country  $i$  with time-varying loadings  $\lambda_t$ , and  $\epsilon_{it}$  is a mean zero error term. The treatment effect of the tax  $\tau$  is assumed to be additive, but can vary across countries and time

$$Y_{it}(1) = Y_{it}(0) + \tau_{it}. \quad (2)$$

Under mild conditions,<sup>5</sup> if a vector of weights  $w$ , where  $\sum_{j=1}^N w_j = 1$  and  $w_j \geq 0$  for all  $j \in \{1, \dots, N\}$ , is chosen such that  $\sum_{j=1}^N w_j Y_{jt}(0) = Y_{0t}(0)$  holds approximately for all  $t < 0$  when  $T_0$  is relatively large, then

$$\hat{\tau}_{0t} = Y_{0t}(1) - \sum_{j=1}^N \hat{w}_j Y_{jt}(0) \quad (3)$$

provides an unbiased estimate of  $\tau_{0t}$ .<sup>6</sup>

For inference, we compare the estimated treatment effect for Uganda  $\hat{\tau}_{0t}$  to the placebo distribution of estimated effects created by repeating the synthetic control procedure for each control unit. Specifically, we estimate  $\tau_{it}$  for each control unit and compute the associated mean squared prediction error

$$\sigma_i^2 = \frac{1}{T_0} \sum_{t < 0} (Y_{it}(0) - \hat{Y}_{it}(0))^2.$$

We then compare  $\hat{\tau}_{0t}$  to the scaled placebo distribution  $\mathcal{P}_t = \{\hat{\tau}_1 \frac{\sigma_0}{\sigma_1}, \dots, \hat{\tau}_N \frac{\sigma_0}{\sigma_N}\}$ . Scaling by  $\frac{\sigma_0}{\sigma_i}$  accounts for systematic differences in how well the synthetic control performs across countries (e.g., a country on the edge of the convex hull will tend to perform poorly). An estimate of  $\hat{\tau}_{0t}$  outside the .025 and .975 quantiles of  $\mathcal{P}_t$  is indicative of a statistically meaningful effect. Note that these are point-wise distributions.

See the Online Appendix and the replication code for additional details.

## 4 Results

### 4.1 Effect of the Tax on Social Media Use

Panel A of Figure 1 plots the log of unique, active Twitter users in Uganda relative to Rwanda, Kenya, and Ghana in each ten day period surrounding the implementation of the social media tax.<sup>7</sup> Panel B uses the average of the other countries in the estimation sample. In the 100 days prior to the tax, Uganda followed a similar trend relative to the other countries. After the tax, however, Twitter use in Uganda is systematically lower. When comparing Uganda to the

<sup>5</sup>See Abadie et al. (2010).

<sup>6</sup>In estimating the weights, we choose  $w$  to minimize  $\|\mathbf{X}_0 - \mathbf{X}_1 w\| = \sqrt{(\mathbf{X}_1 - \mathbf{X}_0 w)' \mathbf{I} (\mathbf{X}_1 - \mathbf{X}_0 w)}$  where  $\mathbf{X}_0 = (Y_{0,T_0}(0), \dots, Y_{0,-1}(0))'$ ,  $\mathbf{X}_1$  is an analogous matrix over the control units, and  $\mathbf{I}$  is an identity matrix. We estimate the weights separately for each outcome of interest.

<sup>7</sup>We define an active user to be one that has an observed tweet during a given period. The countries selected in Panel A have the highest weights in the synthetic control estimates.

synthetic control, Panel C makes this post-tax distinction even clearer. The synthetic control matches Uganda closely in the entire pre-tax period, but there is sharp break in trend at the onset of the tax that persists throughout the post-treatment period.

Panel D of Figure 1 plots the estimated treatment effect of the social media tax on Twitter use. The estimates suggest that the daily willingness to pay for Twitter remains below 200 shillings during most periods for 10 to 20 percent of users. Furthermore, the estimated effect is consistently outside the .025 and .975 quantiles of the scaled placebo distribution.

Tweet and user characteristics in the Twitter dataset reveal heterogeneous effects of the tax. For each category and period, we determine the number of unique users displaying a characteristic at least once.<sup>8</sup> We then re-estimate the synthetic control treatment effects and report the average of the post-intervention treatment effects along with the corresponding placebo distributions.

Figure 2 shows the tax had a larger effect on newer, less frequent, and poorer users. There was a 20.1 percent reduction in new account creation as a result of the tax. For users that tweeted less than one tweet per day on average, there was a 17.8 percent drop. For individuals without an Apple device (an indicator of wealth), we see a 16.8 percent reduction. In contrast, there was an estimated 12.7 percent decrease across all users.

To address concerns regarding anticipatory effects and overfitting, the Online Appendix (Figure A1) examines the robustness of the main treatment effects to only using data from more than 100 days before the onset of the tax to fit the synthetic control. The Online Appendix (Figures A2 and A3) also shows the results of a falsification test using pre-tax data only and the estimated treatment effects when using 1, 7, and 28-day aggregation periods. Lastly, the Online Appendix (Figure A4) reports the full time path of all treatment effects considered here and in the section below.

## 4.2 Effect of the Tax on Collective Action

To examine the impact of the tax on collective action, we first construct an indicator for whether a tweet contains phrases associated with collective action (e.g., “protest” and “rally”).<sup>9</sup> Figure 3

---

<sup>8</sup>See the Online Appendix for additional details on how these groups are constructed. These treatment effects incorporate the combined effect on Twitter use and user characteristics conditional on use because the user characteristics are not always defined prior to the tax. However, the user characteristics considered here are unlikely to be effected by the tax in the timespan examined.

<sup>9</sup>See the Online Appendix for the complete list of collective action and political phrases. English is Uganda’s only official language, and there is limited Twitter support for local languages.

shows an estimated 31.3 percent increase in the number of tweets referencing collective action events in response to the tax. In contrast, there is no meaningful increase in the total number of tweets or the number of tweets containing political phrases unrelated to collective action (e.g., “president” and “government”).

The increase in the amount of online discussion of collective action is partly driven by an increase in the number of unique users participating in these discussions. Figure 3 shows a 17.1 percent increase in the number of users that discuss collective action in a given period. This increase is statistically significant relative to the placebo distribution.

To examine whether this increase in online discussion of collective action is associated with an increase in offline activity, we use data on protests and riots from ACLED and ICEWS. We use the averaged value for the number of events recorded across both datasets for a given period. Figure 3 suggests that collective action events increased by 46.7 percent after the implementation of the tax—well outside the .025 and .975 quantiles of the placebo distribution. The Online Appendix (Figure A4, last panel) shows that the estimated treatment effects on collective action events are highest in the first 100 days after the tax’s implementation.

The Online Appendix (Figure A5) reports the results of a falsification test where the synthetic control is estimated with data more than 100 days prior to the tax’s implementation and the estimated treatment effects are averaged over the 100 day pre-implementation period. Consistent with the tax driving the effect on collective action outcomes, the average estimated effects in the falsification test are either negative or not statistically meaningful. The Online Appendix (Figure A6) also reports sensitivity analysis to alternative specifications of the outcome variables.

## 5 Discussion and Conclusion

Censorship strategies face a trade off between decreasing access to sensitive information and creating visible evidence of government interference in the information environment that has a tendency to create backlash from those experiencing it.<sup>10</sup> Backlash is a key reason the Chinese government has shifted towards more subtle approaches to censorship, including flooding social media with benign information that distracts from sensitive topics (King et al. 2017; Roberts

---

<sup>10</sup>Streisand effects can also increase consumption of specific pieces of sensitive information by drawing attention to and updating beliefs on what the government does not want citizens to be aware of (Nabi 2014; Jansen and Martin 2015; Roberts 2018).

2018). Highly salient forms of digital repression can also reduce coordination costs for collective action by creating a focal point or providing common information (Chwe 2000; Tucker 2007; Egorov et al. 2009; Chwe 2013).

News coverage of Uganda corroborates this mechanism:

[the social media tax] backfired, and the backlash was fierce. Young Ugandans, who organised online under #NoSocialMediaTax and #ThisTaxMustGo hashtags, poured into the streets of the capital city Kampala to express their anger against the government. (Durmaz 2019)

The language used by opposition leaders on Twitter also suggests the tax acted as a focal point. Under the handle @HEBobiWine, Kyagulanyi tweeted the following on July 11, 2018,

..there's no amount of bullets, teargas or arrests that will stop us. Some of our colleagues have been beaten up and others arrested. They must be freed for they have nothing against the police but rather the terrible tax. [...] #ThisTaxMustGo.<sup>11</sup>

The Online Appendix (Figure A7) shows that over 30 percent of collective action tweets in the ten days after the tax's implementation include the string "tax," whereas this proportion is typically below 5 percent.

In addition to the signalling value of the tax, the effect of the tax on the composition or structure of the network can alter collective action outcomes. In network models of collective action, an individual's decision to participate often depends on the participation rate of her local network (Chwe 2000; Watts 2002; Barbera and Jackson 2019). If tendency for collective action and willingness to pay the tax are positively correlated, then the tax could increase the likelihood of successful coordination of collective action by increasing the density of activists. In the Online Appendix, we develop a model following Jackson and Yariv (2007), but allowing for endogenous Twitter use, to highlight the importance of this correlation. Table A2 in the Online Appendix shows that Apple users and long-time users are more likely to tweet about collective action in the pre-intervention period. These users are also users who are more likely to continue using Twitter after the tax, thus, providing some evidence for the tax altering the network composition in a way that could facilitate collective action.

While both mechanisms are consistent with aspects of the data, we are unable to differentiate which, if any, of the aforementioned mechanisms are the primary driver for the increase in

---

<sup>11</sup>On July 7, 2018, Kyagulanyi also tweeted the following call to action, "In an hour, I'll be addressing the people of Uganda on the state of our nation. I will address the boycott on mobile money and social media tax, the national citizen's protest coming up on Wednesday and other issues at hand. Join me live on facebook. #ThisTaxMustGo."

collective action in the period after the tax. Furthermore, we can not exclude the possibility that other concurrent changes are driving the observed trends rather than the tax itself. However, the anecdotal evidence and language of elites suggest that the tax was an important impetus to subsequent collective action.

Overall, the results suggest the economic forces driving selection into social media use and the signalling value of the tax likely caused this form of digital repression to backfire by increasing collective action despite the reduction in social media users.

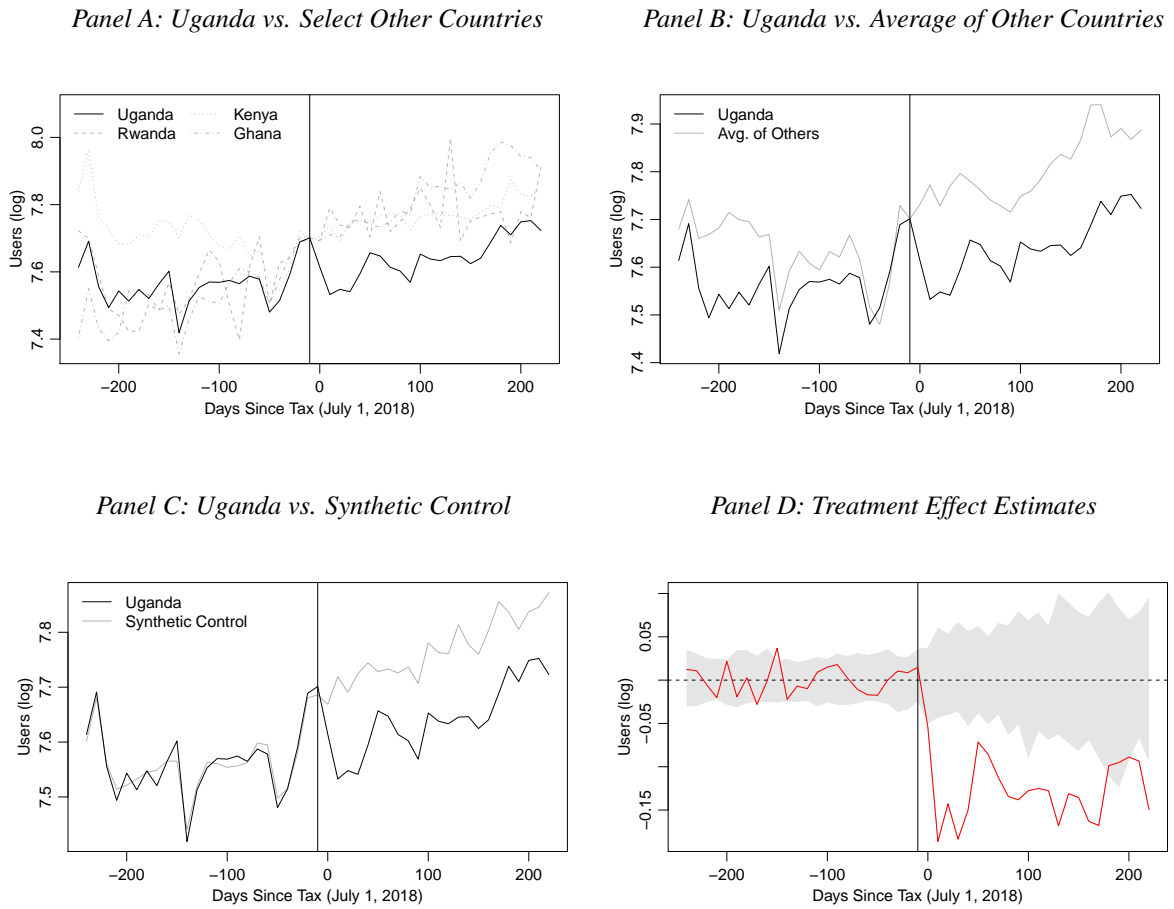
## References

- Abadie, Alberto, Alexis Diamond, and Jens Hainmueller. 2010. Synthetic control methods for comparative case studies: Estimating the effect of California's Tobacco Control Program. *Journal of the American Statistical Association*. 105:493–505.
- Al Jazeera. 2019. Ugandan pop star opposition MP Bobi Wine arrested again. *Al Jazeera*. Accessed at <https://www.aljazeera.com/news/2019/04/ugandan-pop-star-opposition-mp-bobi-wine-arrested-190429153709250.html>.
- Allcott, Hunt, Luca Braghieri, Sarah Eichmeyer, and Matthew Gentzkow. Forthcoming. The welfare effects of social media. *American Economic Review*.
- Ananyev, Maxim, Galina Zudenkova, and Maria Petrova. 2017. Information and communication technologies, protests, and censorship. *Working Paper*. Accessed at <https://ssrn.com/abstract=2978549>.
- Anderson, Simon P. and Andre de Palma. 2009. Information congestion. *RAND Journal of Economics*. 40(4): 688–709.
- Barbera, Salvador and Matthew O. Jackson. 2019. A model of protests, revolution, and information. *Working Paper*. Accessed at <https://ssrn.com/abstract=2732864>.
- BBC. 2017. Uganda MPs vote to scrap presidential age limit. *BBC*. Accessed at <https://www.bbc.com/news/world-africa-42434809>.
- Boschee, Elizabeth, Jennifer Lautenschlager, Sean O'Brien, Steve Shellman, James Starz, and Michael Ward. 2015. ICEWS Coded Event Data. Harvard Dataverse Network. <http://dx.doi.org/10.7910/DVN/28075>.
- Bursztyn, Leonardo, Davide Cantoni, David Y. Yang, Noam Yuchtman, and Y. Jane Zhang. 2019. Persistent political engagement: Social interactions and the dynamics of protest movements. *Working Paper*. Accessed at [https://home.uchicago.edu/bursztyn/Persistent\\_Political\\_Engagement\\_July2019.pdf](https://home.uchicago.edu/bursztyn/Persistent_Political_Engagement_July2019.pdf).
- Chen, Yuyu, and David Y. Yang. Forthcoming. The impact of media censorship: 1984 or brave new world? *American Economic Review*.
- Chwe, Michael Suk-Young. 2000. Communication and coordination in social networks. *Review of Economic Studies*. 67: 1–16.
- Chwe, Michael Suk-Young. 2013. *Rational Ritual: Culture, Coordination, and Common Knowledge*. Princeton University Press: Princeton, NJ.
- Crabtree, Justina. 2018. Facebook says it is committed to Uganda despite social media tax to quash 'gossip'. *cncb.com*. Accessed at <https://www.cncb.com/2018/08/10/facebook-committed-to-uganda-despite-social-media-tax.html> on April 3, 2019.
- Durmaz, Mucahid. 2019. Do African autocrats fear the internet? *TRT World*. Accessed at <https://www.trtworld.com/magazine/do-african-autocrats-fear-the-internet-23327>.

- Egorov, Georgy, Sergei Guriev, and Konstantin Sonin. 2009. Why resource-poor dictators allow freer media: A theory and evidence from panel data. *American Political Science Review*. 103(4): 645–668.
- Enikolopov, Ruben, Alexey Makarin, Maria Petrova, and Leonid Polishchuk. 2017. Social image, networks, and protest participation. *Working Paper*. Accessed at <https://ssrn.com/abstract=2940171>.
- Enikolopov, Ruben, Alexey Makarin, and Maria Petrova. 2018. Social media and protest participation: Evidence from Russia. *Working Paper*. Accessed at <https://ssrn.com/abstract=2696236>.
- Gohdes, Anita. 2015. Pulling the plug: Network disruptions and violence in civil conflict. *Journal of Peace Research*. 52(3): 352–367.
- Gould, Roger V. 1993. Collective action and network structure. *American Sociological Review*. 58(2): 182–196.
- Griffiths, James. 2019. Democratic Republic of Congo internet shutdown shows how Chinese censorship tactics are spreading. *CNN*. Accessed at <https://www.cnn.com/2019/01/02/africa/congo-internet-shutdown-china-intl/index.html>.
- Hassanpour, Navid. 2014. Media disruption and revolutionary unrest: Evidence from Mubarak’s quasi-experiment. *Political Communication*. 31(1):1–24.
- Hobbs, William and Margaret E. Roberts. 2018. How sudden censorship can increase access to information. *American Political Science Review*. 112(3):1–16.
- Iyer, Ganesh and Zsolt Katona. 2016. Competing for attention in social communication markets. *Management Science*. 62(8): 2304–2320.
- Jackson, Matthew O. and Leeat Yariv. 2007. Diffusion of Behavior and Equilibrium Properties in Network Games. *American Economic Review (Papers and Proceedings)*. 97(2): 92–98
- Jansen, Sue Curry and Brian Martin. 2015. The Streisand effect and censorship backfire. *International Journal of Communication*. 9: 656–671.
- Kagumire, Rosebell. 2018. Bobi Wine and the beginning of the end of Museveni’s power. *Al Jazeera*. Accessed at <https://www.aljazeera.com/indepth/opinion/bobi-wine-beginning-museveni-power-180828111608108.html>.
- King, Gary, Jennifer Pan, and Margaret E. Roberts. 2013. How censorship in China allows government criticism but silences collective expression. *American Political Science Review*. 107(2): 1–18.
- . 2014. Reverse-engineering censorship in China: randomized experimentation and participant observation. *Science*. 345(6199): 1–10.
- . 2017. How the Chinese government fabricates social media posts for strategic distraction, not engaged argument. *American Political Science Review*. 111(3): 484–501.
- Leetaru, Kalev H., Shaowen Wang, Guofeng Cao, Anand Padmanabhan, and Eric Shook. 2013.

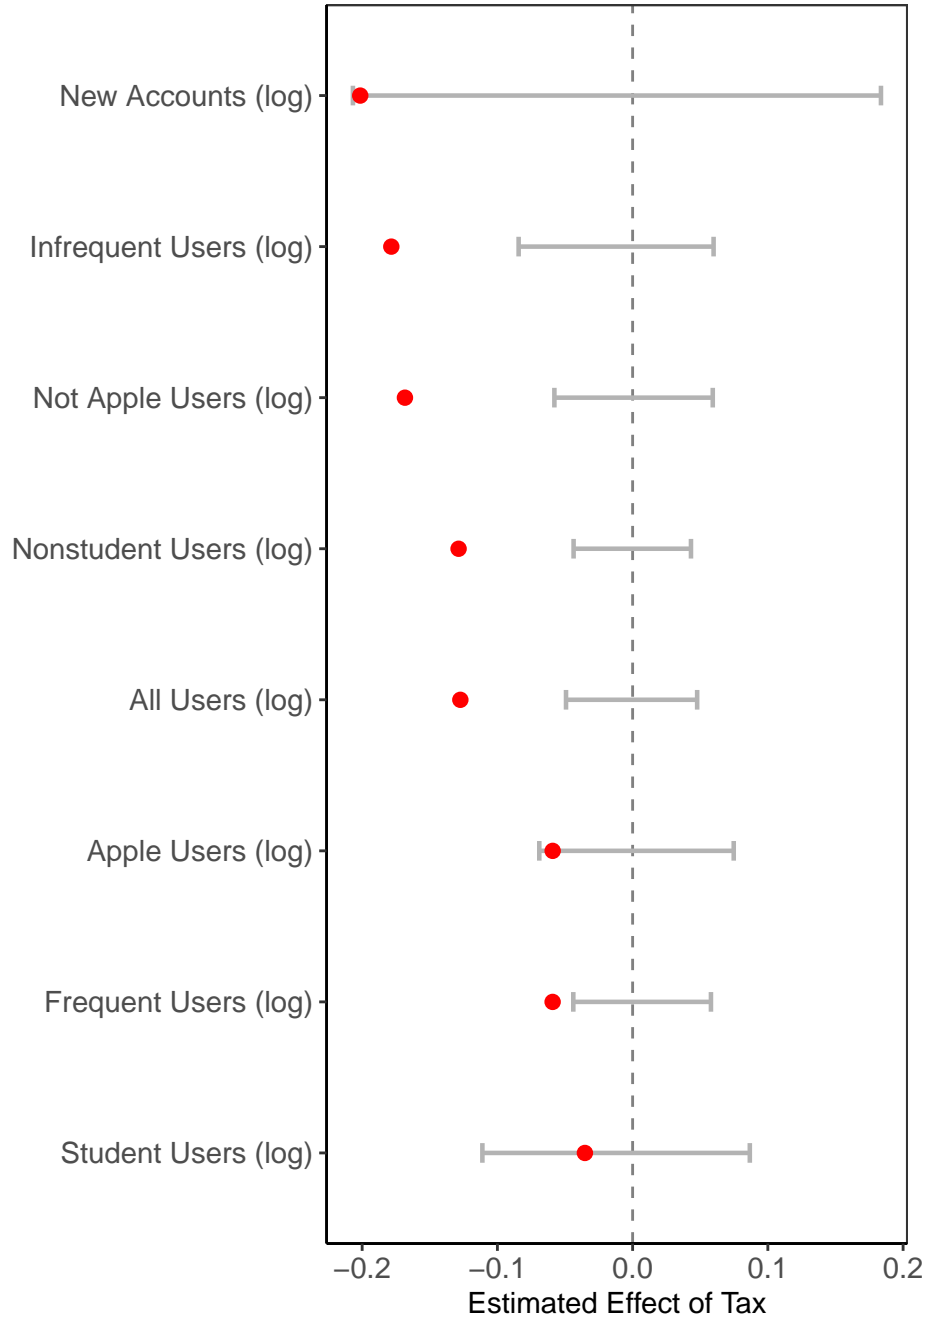
- Mapping the global Twitter heartbeat: The geography of Twitter. *First Monday*. 18(5-6).
- Little, Andrew T. 2016. Communication technology and protest. *Journal of Politics*. 78(1): 152–166.
- Manacorda, Marco and Andrea Tesei. 2016. Liberation technology: Mobile phones and political mobilization in Africa. *Working Paper*. Accessed at <https://ssrn.com/abstract=2795957>.
- Mugerwa, Yasiin. 2017. China to help Uganda fight internet abuse. *Daily Monitor*. Accessed at <https://www.monitor.co.ug/News/National/China-Uganda-Internet-Evelyn-Anite-Africa-Internet-Users/688334-4032626-u1161r/index.html>.
- Nabi, Zubair. 2014. Censorship is futile. *First Monday*. 19(11).
- Observer. 2017. Uganda to launch own version of Facebook, Twitter - UCC. *The Observer*. Accessed at <https://observer.ug/news/headlines/56506-uganda-to-launch-own-version-of-facebook-next-year-ucc.html>.
- Oliver, Pamela E. and Gerald Marwell. 1988. The paradox of group size in collective action: A theory of the critical mass. II. *American Sociological Review*. 53(1): 1–8.
- Pan, Jennifer and Alexandra A. Siegel. 2018. Physical repression and online dissent: Evidence from Saudia Arabia. *Working Paper*. Accessed at [https://alexandra-siegel.com/wp-content/uploads/2018/11/Pan\\_Siegel\\_Saudi\\_Twitter\\_Oct29.pdf](https://alexandra-siegel.com/wp-content/uploads/2018/11/Pan_Siegel_Saudi_Twitter_Oct29.pdf).
- Qin, Bei, David Stromberg, and Yanhui Wu. 2017. Why does China allow freer social media? Protests versus surveillance and propaganda. *Journal of Economic Perspectives*. 31(1): 117–140.
- Raleigh, Clionadh, Andrew Linke, Håvard Hegre and Joakim Karlsen. 2010. “Introducing ACLED—Armed conflict location and event data.” *Journal of Peace Research* 47(5): 65–660.
- Roberts, Margaret E. 2018. *Censored: Distraction and Diversion Inside China’s Great Firewall*. Princeton University Press: Princeton, NJ.
- Siegel, David A. 2009. Social networks and collective action. *American Journal of Political Science*. 53(1): 122–138.
- Steinert-Threlkeld, Zachary. 2018. *Twitter as Data*. Cambridge University Press: Cambridge, UK.
- Sunstein, Cass R. Forthcoming. Valuing Facebook. *Behavioural Public Policy*.
- Tucker, Joshua A. 2007. Enough! Electoral fraud, collective action problems, and post-communist colored revolutions. *Perspectives on Politics*. 5(3): 535–551.
- Van Zandt, Timothy. 2004. Information overload in a network of targeted communication. *RAND Journal of Economics*. 35(3): 542–560.
- Watts, Duncan J. 2002. A simple model of global cascades on random networks. *Proceedings of the National Academy of Sciences*. 99(9): 5766–5771.

Figure 1: Impact of Tax on Twitter Use



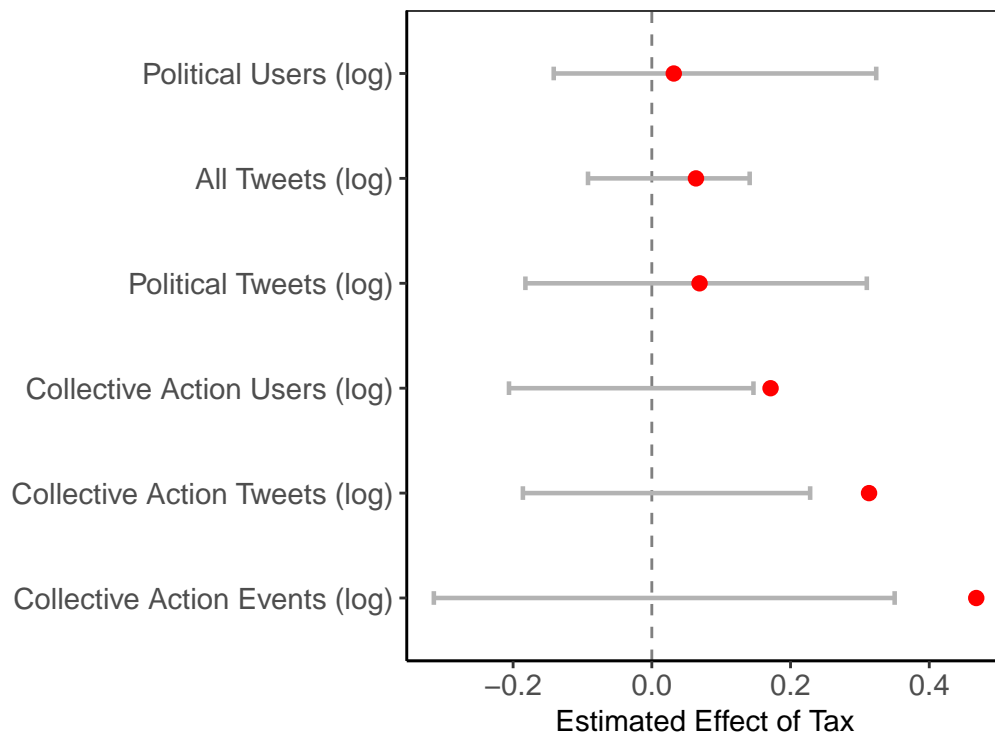
Notes: Panel A compares the log of one plus the number of unique, active users in each ten day period for Uganda (black) and select other African countries (grey). Panel B compares the log of one plus the number of unique, active users in each ten day period for Uganda (black) and the average of the African countries and territories kept after the sample restrictions (grey). The difference between the log of one plus the number of users in Uganda versus the other African countries in the period immediately prior to the intervention ( $t = -1$ ) is subtracted from the series for each individual country (Panel A) and the average (Panel B). Panel C plots the log of one plus the number of unique, active users in Uganda during each ten day period alongside the synthetic control. Panel D plots the estimated treatment effect  $\hat{\tau}_{0t}$  of the social media tax on the log of one plus the number of unique, active users in Uganda. The shaded region indicates the .025 and .975 quantiles of the scaled placebo distribution  $\mathcal{P}_t$ . The vertical line indicates the 10 day period immediately prior to when the social media tax was implemented in Uganda (July 1, 2018).

Figure 2: Effect of the Tax on User Composition



Notes: The figure plots the average estimated treatment effect  $\frac{1}{T_{1+1}} \sum_{t \geq 0} \hat{\tau}_{0t}$  (dot) of the social media tax on various outcomes along with the .025 and .975 quantiles of the scaled placebo distribution  $\mathcal{P}$  (grey bars). See the text and Online Appendix for definitions of each outcome.

Figure 3: Effect of the Tax on Collective Action



Notes: The figure plots the average estimated treatment effect  $\frac{1}{T_1+1} \sum_{t \geq 0} \hat{\tau}_{0t}$  (dot) of the social media tax on various outcomes along with the .025 and .975 quantiles of the scaled placebo distribution  $\mathcal{P}$  (grey bars). See the Online Appendix for further details on each outcome.

# A1 Online Appendix

## A1.1 Additional Data Details

See replication code for exact details on the data construction.

### A1.1.1 Twitter Outcomes

For the Twitter outcomes, we restrict the set of countries for which we use to construct the synthetic control to those in the top 80 percent of the sample of sub-Saharan African countries with respect to the average number of unique, active users per period, and we do this separately for each level of temporal aggregation. We also exclude tweets whose user description contains phrases associated with bot activity.<sup>12</sup> For all logged variables, we use the log of one plus the level.

- New accounts: The number of users in our dataset who created their Twitter account in a given period.
- Infrequent users: Users who, upon their first appearance in our dataset, averaged less than one tweet per day as measured by the number of posted statuses listed on their account and the date of account creation.
- Not Apple users: Users for which our dataset never contains a tweet posted via “ios,” “ipad,” or “iphone” in a given period.
- Student users: Users for which the user description or user location includes “university,” “college,” “universite,” “universita,” “universidade,” “faculdade” or “student” for at least one tweet in our dataset in a given period.
- Activist users: Users for which at least one tweet contains a phrase referencing collecting action in a given period.<sup>13</sup>
- Political users: Users for which at least one tweet contains a phrase referencing the government in a given period.<sup>14</sup>

---

<sup>12</sup>“weather,” “4:20,” “job,” “career,” “hire,” and “hiring.”

<sup>13</sup>‘protest’, ‘rally’, ‘rallies’, ‘freedom of assembly’, ‘freedom of expression’, ‘riot’, ‘march against’, ‘march with’, ‘inciting violence’, ‘unlawful assembly’, ‘civil unrest’, ‘demonstration’, ‘i stand with’, ‘take to the street’, ‘taking to the street’, or ‘boycott’

<sup>14</sup>‘supreme court’, ‘legislator’, ‘election’, ‘president’, ‘vote’, ‘ballot’, ‘administracao’, ‘presidente’, ‘mp’, ‘constitution’, ‘partisan’, ‘mps’, ‘administration’, ‘dictator’, ‘gouvernement’, ‘citizen’, ‘republic’, ‘political party’, ‘political parties’, ‘parlement’, ‘parliament’, ‘government’, ‘law’, ‘constituent’, ‘governo’, ‘democrac’, ‘politic’, ‘amministrazione’, ‘military’, ‘constituency’, and ‘parlamento’.

- Collective action tweets: Tweets containing at least one phrase referencing collective action as defined above for activist users.

### **A1.1.2 ACLED and ICEWS Outcomes**

The ACLED data includes all countries in Africa and is described by Raleigh et al. (2010). It was downloaded from <https://www.acleddata.com/> on March 6, 2019. We restrict attention to events labelled as “Riots/protests.” Other event categories include, among others, “Violence against citizens” and “Battle-No change of territory.”

The ICEWS data comes from Boschee et al. (2015). It is a machine-coded dataset that generates 20 event types, one of which is protests. Of the 45 events in Uganda during the main window of analysis, none are exact duplicates.<sup>15</sup>

Separately, for each dataset, we compute the number of events that occur in a given time period. We then take the average across both datasets as our baseline measure of the number of collective action events for a given country in a given period, where this measure is appropriately transformed depending on the specification (e.g., log). We restrict attention to countries that appear in both our ACLED and ICEWS samples throughout. For all logged variables, we use the log of one plus the level.

---

<sup>15</sup>Defined by source, actor, target, and location.

## A1.2 Model of Collective Action with Social Media Tax

The following discussion builds on Jackson and Yariv (2007). There is a continuum of agents. Each agent decides whether to pay a price  $q$  and join Twitter. Users have an intrinsic valuation of Twitter  $w_i \geq 0$  with c.d.f.  $F_w$ .

Each agent subsequently chooses whether to participate in a protest and has costs of participating  $c_i \geq 0$ . The utility from participating in the protest is a function of the average participation of connections on Twitter  $v(x)$  (where Twitter connections are formed randomly from the population of Twitter users), and that  $v(\cdot)$  is an increasing function in  $x$ .<sup>16</sup>

Let  $c_i$  and  $w_i$  be jointly normal with mean  $\mu$  and variance  $\Omega$ . Agents will choose to protest if  $v(x) > c_i$ . Let  $F_{c|w_i \geq q-v(x)}$  denote the conditional distribution of idiosyncratic protest costs for users on Twitter for a given rate of protest participation  $x$ . Following Jackson and Yariv (2007), we have that

$$x = \phi(x) \equiv F_{c|w_i \geq q-v(x)}[v(x)]$$

characterizes the symmetric Bayesian equilibria of the game.

This leads to the following theorem.

**Theorem 1.** *Impact of tax depends on the covariance  $\Omega_{cw}$  between Twitter valuations  $w_i$  and protest costs  $c_i$ .*

1. *If  $q' > q$  and  $\Omega_{cw} < 0$ , then  $q'$  will generate more diffusion of behavior on the Twitter network.*
2. *If  $q' > q$  and  $\Omega_{cw} > 0$ , then  $q'$  will generate less diffusion of behavior on the Twitter network.*

*Proof.* Note that  $F_{c|w_i \geq q-v(x)}[v(x)]$  is an increasing [decreasing] function in  $q$  for all  $x$  when  $\Omega_{cw} < 0$  [ $\Omega_{cw} > 0$ ]. Therefore,  $\phi(x, q') > \phi(x, q)$  [ $\phi(x, q') < \phi(x, q)$ ] for all  $x$ .

Proposition 1 of Jackson and Yariv (2007) states that, for two regular environments,<sup>17</sup> an environment  $\tilde{\phi}(\cdot)$  will generate greater diffusion of behavior<sup>18</sup> than  $\phi(\cdot)$  if  $\tilde{\phi}(x) \geq \phi(x)$  for all  $x$ . □

<sup>16</sup>Define  $x \equiv 0$  for non-Twitter users and  $v(0) = 0$ .

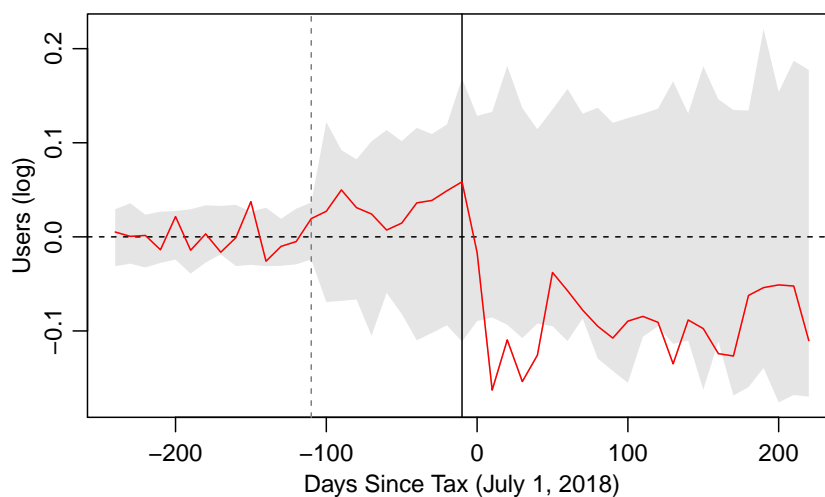
<sup>17</sup>See Jackson and Yariv (2007) for technical details.

<sup>18</sup>In the sense that tipping points are lower and stable equilibria are higher.

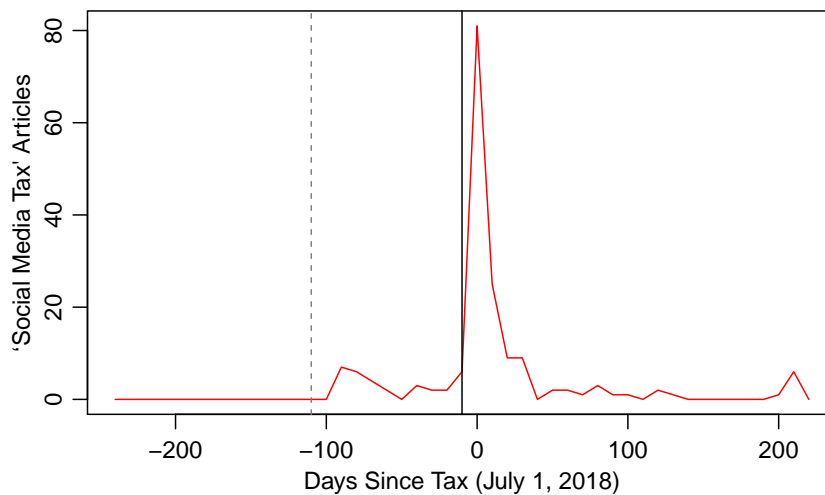
### A1.3 Online Appendix Figures and Tables

Figure A1: Robustness to Restricting Model Fitting Period

*Panel A: Estimated Treatment Effects*

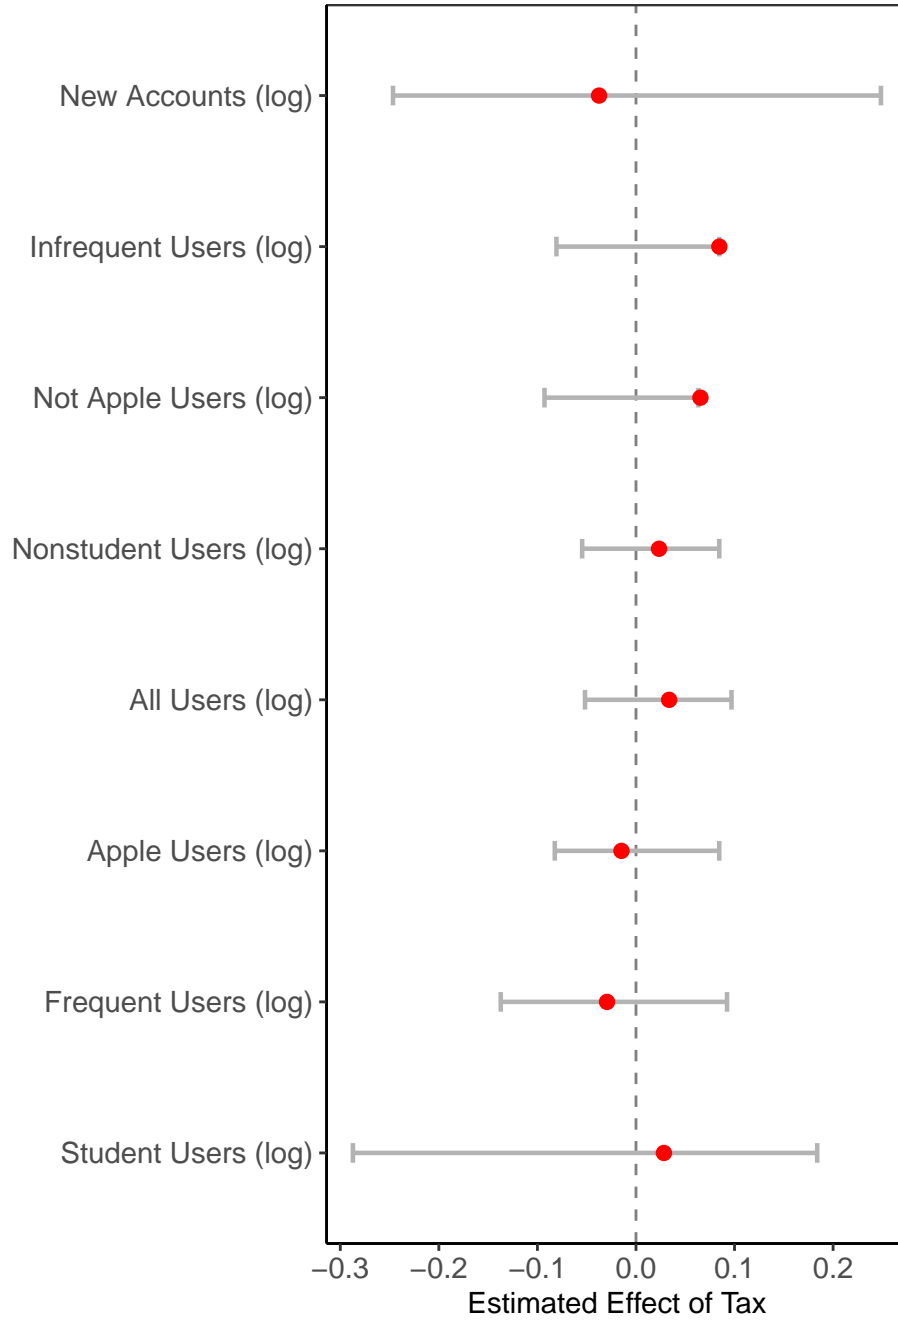


*Panel B: Discussion of Social Media Tax*



Notes: Panel A shows the log of one plus the number of users per period in each 10-day period for Uganda relative to the synthetic control. The solid vertical line indicates the 10 day period immediately prior to when the social media tax was implemented in Uganda (July 1, 2018). The dashed vertical line indicates the last period which was used to fit the synthetic control. Panel B shows the total number of news articles, as measured by data from BuzzSumo, referencing “social media tax” with a .ug domain in each 10-day period.

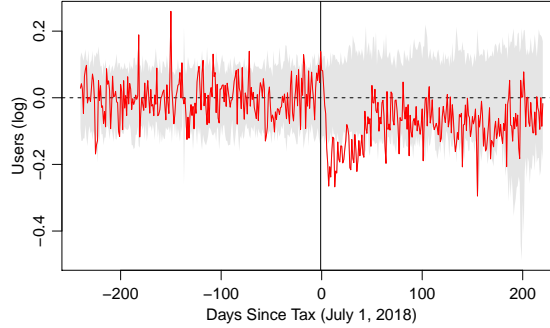
Figure A2: Effect of the Tax on User Composition – Falsification Test



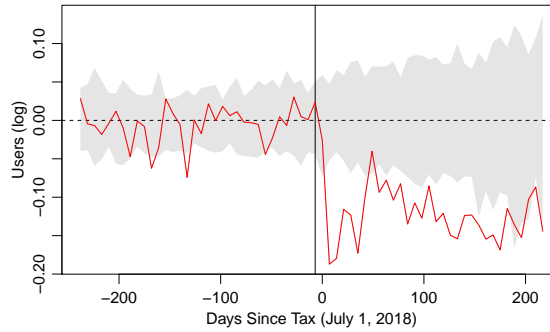
Notes: The figure plots the average estimated treatment effect (dot) of the social media tax on various outcomes along with the .025 and .975 quantiles of the scaled placebo distribution  $\mathcal{P}$  (grey bars) when restricting the model fit of the synthetic control to more than 100 days prior to the tax. The average of the estimated treatment effects are then taken over the 100 days prior to the tax. See Section A1.1 for further details on each outcome and Figure 2 to compare to the true estimated treatment effects.

Figure A3: Robustness to Other Levels of Aggregation

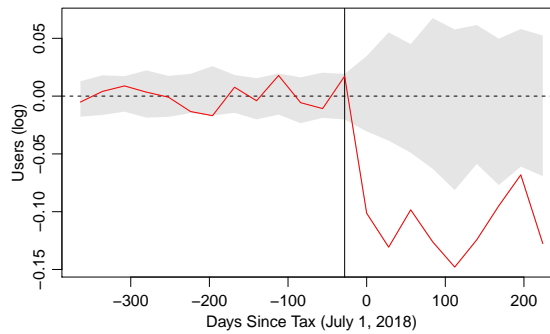
*Panel A: 1-day Aggregation*



*Panel B: 7-day Aggregation*

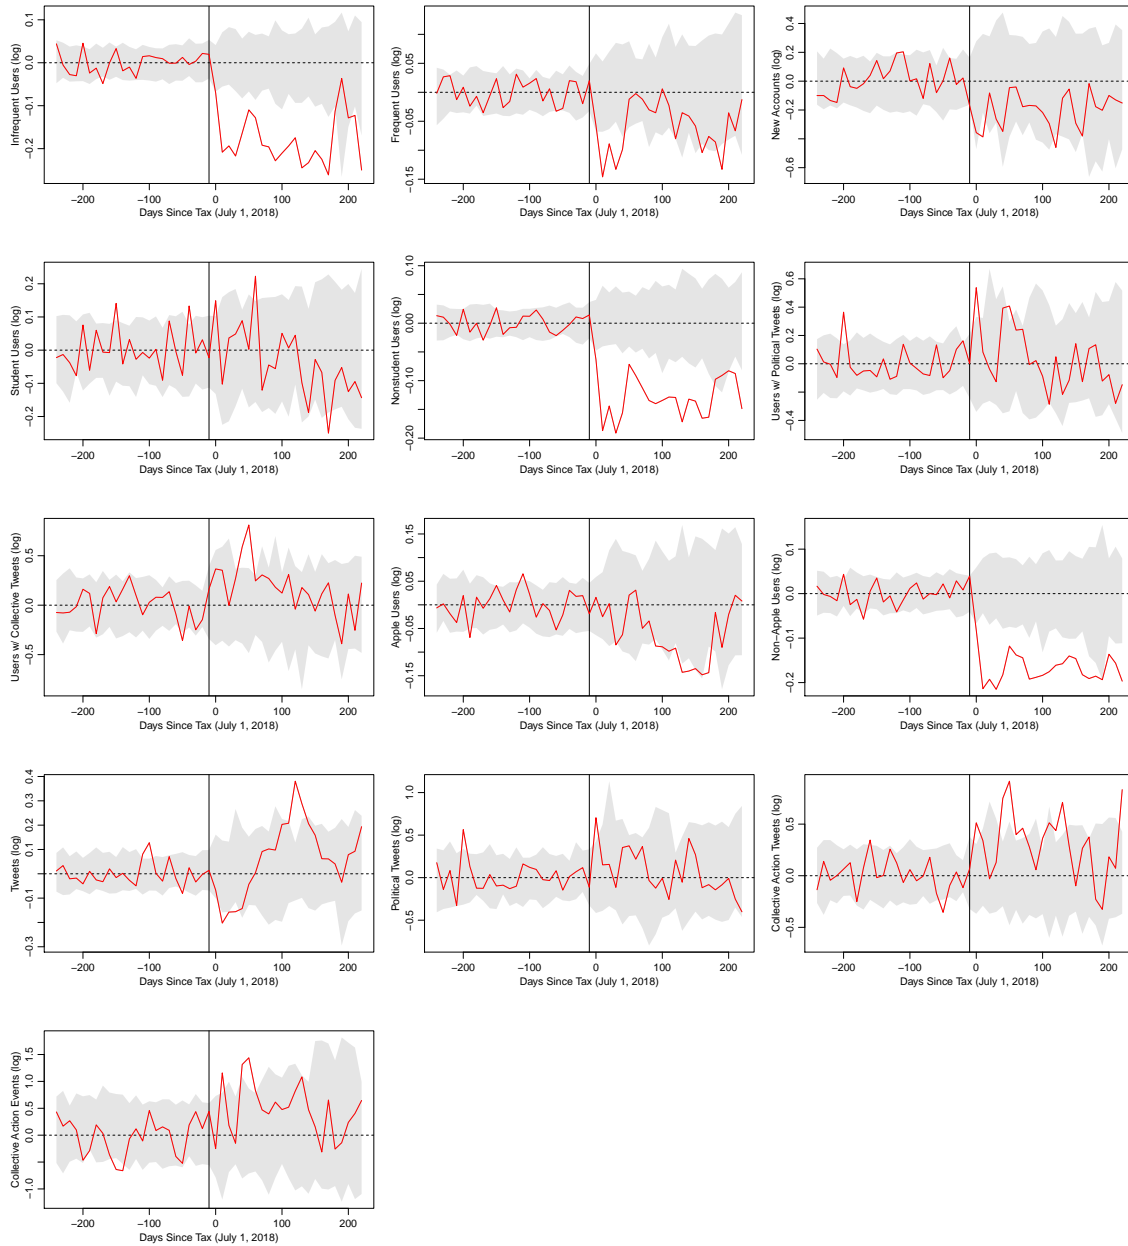


*Panel C: 28-day Aggregation*




Notes: The figure shows the estimated effect of the tax on the log of one plus the number of unique, active users in a given period using various levels of temporal aggregation. Panel A uses 1-day periods, Panel B uses 7-day periods, and Panel C uses 28-day periods. Panel A uses every 10th period in the pre-tax period for  $\mathbf{X}_0$  in estimating the weights. Panel B uses every 4th period in the pre-tax period for  $\mathbf{X}_0$  in estimating the weights. Panel C uses all periods in the pre-tax period for  $\mathbf{X}_0$  in estimating the weights. Since Panels A and B do not use all pre-intervention periods for  $\mathbf{X}_0$ , the weights are chosen to minimize  $\sqrt{(\mathbf{X}_1 - \mathbf{X}_0 w)' \mathbf{V} (\mathbf{X}_1 - \mathbf{X}_0 w)}$  where  $\mathbf{V}$  is a positive definite diagonal matrix chosen to minimize the mean squared prediction error across all pre-intervention periods. The shaded region indicates the .025 and .975 quantiles of the scaled placebo distribution  $\mathcal{P}_t$ . The vertical line indicates the 10 day period immediately prior to when the social media tax was implemented in Uganda (July 1, 2018).

Figure A4: Time Path of Estimated Treatment Effects



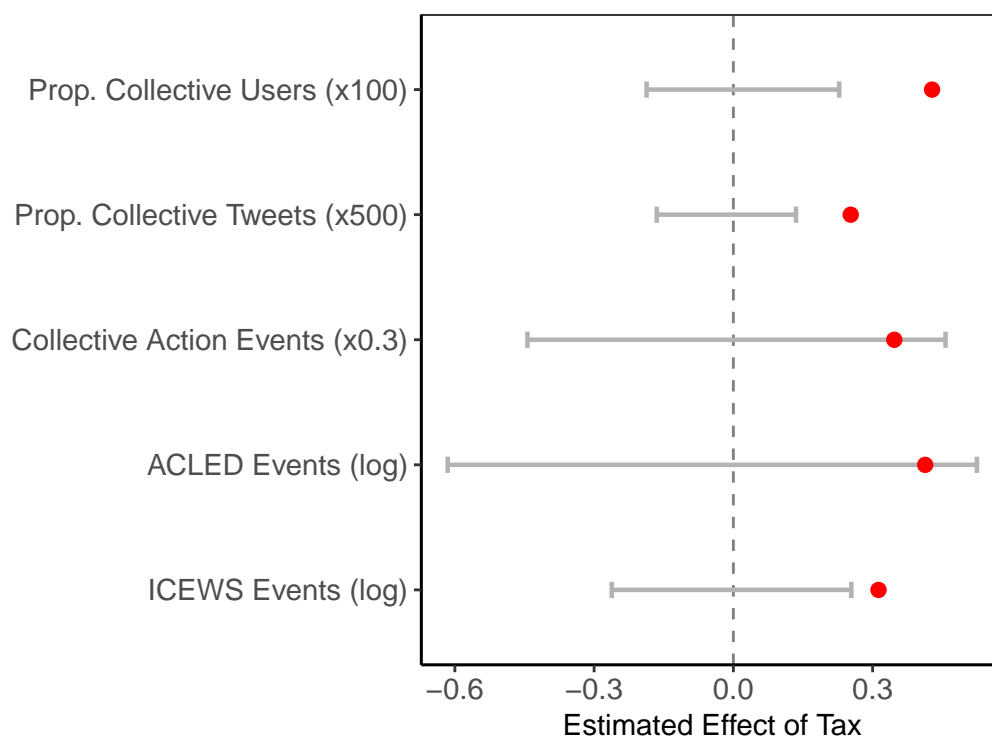
Notes: The figure reports the estimated treatment effects of the social media tax on various outcomes in Uganda during each ten day period relative to a synthetic control. The shaded region indicates the .025 and .975 quantiles of the scaled placebo distribution  $\mathcal{P}_t$ . The vertical line indicates the 10 day period immediately prior to when the social media tax was implemented in Uganda (July 1, 2018). See Section A1.1 for details on the construction of each outcome variable.

Figure A5: Effect of the Tax on Collective Action – Falsification Test



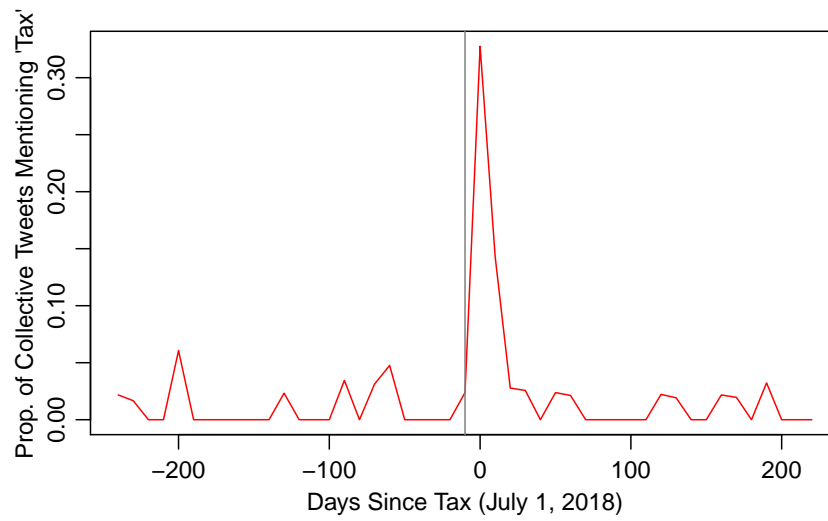
Notes: The figure plots the average estimated treatment effect (dot) of the social media tax on various outcomes along with the .025 and .975 quantiles of the scaled placebo distribution  $\mathcal{P}$  (grey bars) when restricting the model fit of the synthetic control to more than 100 days prior to the tax. The average of the estimated treatment effects are then taken over the 100 days prior to the tax. See Section A1.1 for further details on each outcome and Figure 3 to compare to the true estimated treatment effects.

Figure A6: Effect of the Tax on Collective Action – Robustness



Notes: The figure plots the average estimated treatment effect  $\frac{1}{T_1+1} \sum_{t \geq 0} \hat{\tau}_{0t}$  (dot) of the social media tax on various outcomes along with the .025 and .975 quantiles of the scaled placebo distribution  $\mathcal{P}$  (grey bars). The ‘Prop. Collective Users,’ ‘Prop. Collective Tweets,’ and ‘Collective Action Events’ outcomes are scaled by 100, 500, and 0.3 respectively. The ‘Prop.’ outcomes report the proportion of users that are collective action users or the proportion of tweets that are collective action tweets. See the main text and Section A1.1 for additional details on each outcome.

Figure A7: Collective Action Tweets Mentioning the Tax



Notes: The plot shows the proportion of collective action tweets mentioning 'tax' in Uganda in each 10 day period. The solid vertical line indicates the 10 day period immediately prior to when the social media tax was implemented in Uganda (July 1, 2018).

Table A1: Top 10 Weights for Synthetic Control

Users (log)		Infrequent Users (log)		Activist Users (log)		Tweets (log)		Collective Action Tweets (log)	
Country	Weight	Country	Weight	Country	Weight	Country	Weight	Country	Weight
GH	0.312	GH	0.376	KE	0.302	KE	0.394	KE	0.315
KE	0.288	KE	0.204	GH	0.214	CM	0.169	ZW	0.215
RW	0.152	CD	0.128	ZW	0.143	TZ	0.112	ZM	0.210
ET	0.059	ZM	0.113	AO	0.100	SN	0.064	SL	0.078
SO	0.058	RW	0.105	ZM	0.064	GN	0.052	MU	0.052
CD	0.049	ET	0.046	MZ	0.057	BJ	0.042	AO	0.050
ZM	0.032	SO	0.029	RE	0.034	GH	0.035	SZ	0.025
BI	0.024	TD	0.000	MU	0.032	BI	0.034	RW	0.021
NE	0.023	SN	0.000	CV	0.028	TD	0.026	MW	0.020
SN	0.002	ZW	0.000	SZ	0.026	ZW	0.018	BJ	0.008

Notes: The table shows the top ten countries along with their associated weights use for constructing the synthetic control for each outcome variable. Country is the two-letter ISO 3166 country code.

Table A2: Determinants of Collective Action and Political Tweets

Dependent Variable:	Collective Action Tweets			Political Tweets		
	(1)	(2)	(3)	(4)	(5)	(6)
Apple User $\times$ .01	0.077 (0.039)	25.804 (12.148)	0.001 (0.000)	-0.316 (0.249)	-10.660 (8.375)	-0.000 (0.002)
Frequent User $\times$ .01	0.017 (0.032)	5.985 (11.025)	0.038 (0.035)	-0.238 (0.264)	-7.310 (7.533)	-0.058 (0.164)
Student User $\times$ .01	-0.065 (0.047)	-25.099 (19.076)	-0.001 (0.001)	0.918 (0.419)	24.618 (10.660)	0.012 (0.004)
Longtime User $\times$ .01	0.106 (0.041)	38.093 (15.834)	0.116 (0.035)	0.372 (0.333)	10.607 (9.650)	0.001 (0.173)
Level of Observation	Tweet	Tweet	10-day Period	Tweet	Tweet	10-day Period
Estimator	OLS	Logit	OLS	OLS	Logit	OLS
Period FE	Y	Y	Y	Y	Y	Y
Location FE	Y	Y	Y	Y	Y	Y
Clusters	17016	16646	17016	17016	17007	17016
Observations	483090	476652	81393	483090	482860	81393

Notes: Table shows the results of regressing measures of collective action or political tweets on user and period characteristics. Columns (1) and (2) use an indicator for whether a tweet contains a collective action phrase. Column (3) collapses observations by user id, location, and 10-day periods by taking the mean of each variable over the associated set of tweets. Columns (4)–(6) are analogous, but use a measure of whether a tweet contains a political phrase not related to collective action. Fixed effects for 10-day periods and location name as assigned by Twitter are included throughout. Observations are restricted to the period prior to the tax, and to tweets where both the user language and tweet language is specified as English. All independent variables reported are scaled by .01. ‘Apple User’ is an indicator for whether the tweet came from an Apple device; for Columns (3) and (6), it is an indicator for whether a user tweeted from an Apple device at least once during the period. ‘Frequent User’ is an indicator for whether an individual tweeted more than once per day between joining Twitter and when they initially appear in our data. ‘Student User’ is an indicator for whether the user included a university-related phrase in their description or location at the time of the tweet; for Columns (3) and (6), it is an indicator for whether a university-related phrase appears in the description or location for at least one tweet during the period. ‘Longtime User’ is an indicator for whether an individual joined Twitter more than four years prior to the start of the tax. Coefficients are reported throughout, and standard errors reported in parentheses below are clustered at the user level.