

3G INTERNET AND CONFIDENCE IN GOVERNMENT*

Sergei Guriev (r)
Nikita Melnikov (r)
Ekaterina Zhuravskaya

This draft: October 2020

First draft: June 2019

Abstract

How does mobile broadband internet affect approval of government? Using Gallup World Poll surveys of 840,537 individuals from 2,232 subnational regions in 116 countries from 2008 to 2017 and the global expansion of 3G mobile networks, we show that, on average, an increase in mobile broadband internet access reduces government approval. This effect is present only when the internet is not censored, and it is stronger when the traditional media are censored. 3G helps expose actual corruption in government: revelations of the Panama Papers and other corruption incidents translate into higher perceptions of corruption in regions covered by 3G networks. Voter disillusionment had electoral implications: In Europe, 3G expansion led to lower vote shares for incumbent parties and higher vote shares for the antiestablishment populist opposition. Vote shares for nonpopulist opposition parties were unaffected by 3G expansion. *JEL* codes: D72, D73, L86, P16.

*We thank three anonymous referees and Philippe Aghion, Nicolas Ajzenman, Oriana Bandiera, Timothy Besley, Kirill Borusyak, Filipe Campante, Mathieu Couttenier, Ruben Durante, Jeffrey Frieden, Thomas Fujiwara, Davide Furceri, Irena Grosfeld, Andy Guess, Brian Knight, Ilyana Kuziemko, John Londregan, Marco Manacorda, Alina Mungiu-Pippidi, Chris Papageorgiou, Maria Petrova, Pia Raffler, James Robinson, Seyhun Orcan Sakalli, Andrei Shleifer, Andrey Simonov, Stefanie Stantcheva, David Strömberg, Maria Micaela Sviatschi, David Yang, Alwyn Young, Andrei Zeleneev, the participants of seminars in Bocconi University, Harvard University, London School of Economics, New Economic School, Paris School of Economics, Princeton University, Sciences Po, and the Annual Globalisation Seminar at the School of Business and Management (Queen Mary University, London), Annual Workshop of CEPR RPN on Populism, Comparative Economics Webinar, IIES/SNS International Policy Talks, Kyiv Conference on Corruption, and Social Media Economics Workshop in ENS Lyon for helpful comments. We are grateful to Hites Ahir, Laura Barros, Anna Biryukova, Davide Furceri, Chris Papageorgiou, Manuel Santos Silva, and Denis Volkov for sharing their data. We also thank Etienne Madinier and Antonela Miho for excellent research assistance. The authors wish to thank the World Wide Lightning Location Network (<http://wwlln.net>), a collaboration among over 50 universities and institutions, for providing the lightning location data used in this paper. All authors contributed equally to the paper.

I. INTRODUCTION

What are the political implications of the expansion of mobile broadband internet around the world? Optimists argue that broadband internet improves access to independent political information, raising public awareness about quality of governance. Social media enables two-way information flows that help overcome collective-action problems in organizing protests against nondemocratic governments. For instance, in the wake of the Arab Spring of 2010–2012, the internet and social media were branded as “liberation technology” (Diamond and Plattner, 2010). Pessimists, in contrast, point out that social media makes it easy to disseminate fake news (Allcott and Gentzkow, 2017; Vosoughi, Roy, and Aral, 2018), empowers nondemocratic regimes by reducing the costs of propaganda and surveillance (Morozov, 2011; Mitchell et al., 2019), and helps populists connect to voters (Tufekci, 2018). These conjectures found empirical support in a number of studies that have analyzed the political implications of broadband internet expansion and social media penetration in single-country settings (for a recent survey of this literature, see Zhuravskaya, Petrova, and Enikolopov, 2020).

Our paper is the first to study the political effects of the expansion of third-generation (3G) mobile networks throughout the world. 3G was the first generation of mobile broadband internet that allowed users to freely browse the web from their smartphones and to stream or upload videos; it was a key driver of the rapid expansion of social media (Rainie and Wellman, 2012). We use Gallup World Poll (GWP) data on the attitudes and beliefs of approximately 840,000 individuals living in 2,232 subnational regions of 116 countries throughout the world from 2008 to 2017. We find that greater 3G availability, on average, decreases government approval. Citizens who gain access to mobile broadband internet show less support for their government: they become more aware of government corruption and less confident in the country’s government institutions.

This result is consistent with conjectures of many political analysts, sociologists, and psychologists, who have argued that the growth of social media, catalyzed by the expansion of mobile broadband internet, has undermined the legitimacy of governments around the world. In his recent book *The Revolt of the Public*, a former CIA analyst Martin Gurri argues that “the rise of Homo informaticus [a citizen relying on social media for information] places governments on a razor’s edge, where any mistake, any untoward event, can draw networked public into the streets... This is the situation today for authoritarian governments and liberal democracies alike. The crisis in the world [...] concerns loss of trust in government” (Gurri, 2018, p. 90). He conjectures that “the greater the diffusion of information to the public [through social media], the more illegitimate any political status quo will appear... Homo informaticus ... poses an

existential challenge to the legitimacy of every government he encounters” (Gurri, 2018, p. 91). A seminal scholar of the “network society,” Manuel Castells, argues in his recent book *Rupture: The Crisis of Liberal Democracy*, that the dissemination of images and videos through social media is a reason for this crisis of political legitimacy because “politics is fundamentally emotional” and “negative images are five times more effective in terms of influence than positive ones” (Castells, 2019, p.20). Similarly, a prominent social psychologist, Jonathan Haidt, with his coauthor, Tobias Rose-Stockwell, in their summary of recent research on the psychology of social media conclude that social media does not just serve as a spark for public outrage with the status quo, it also is especially “designed to make outrage contagious” (Haidt and Rose-Stockwell, 2019).

We find that the magnitude of the negative effect of the expansion of mobile broadband internet on government approval is substantial. An average-size increase in regional 3G coverage during the 2008–2017 decade resulted in 39% of an average subnational region’s population gaining access to mobile broadband internet, reduced the confidence in the national government of the region’s population by 2.5 percentage points (from the mean level of 51%), and increased the perception that the government is corrupt by 1.4 percentage points (from the mean of 77%).

The global setting allows us to study the heterogeneity of the effects of 3G expansion on government approval, which helps to shed light on some of the mechanisms at play. First, we show that 3G decreases government approval only when the internet is not censored. This is despite the fact that 3G networks increase internet penetration everywhere, including countries with internet censorship. This suggests that political information available online that is independent of the government makes people change their attitudes toward the government. Second, when the internet is not censored, the negative effect of 3G on government approval is stronger in countries where the government controls the traditional media, implying that mobile broadband internet becomes a major source of news when no other sources of independent political information are available. Third, we find that the effect of 3G is negative only when there is at least some corruption. The least corrupt governments (such as those of Denmark or Switzerland) suffer no drop in public approval ratings as a result of 3G expansion; in these countries, 3G expansion actually increases government approval. This evidence is consistent with Bayesian updating of public beliefs: if new information on the quality of governance made available via mobile broadband constitutes good news compared to the ex ante beliefs, 3G expansion should result in higher government approval. Fourth, we demonstrate explicitly that mobile broadband internet helps inform the public about actual corruption. Using Furceri, Papageorgiou, and Ahir (2019)’s measure of actual incidents of corruption around the world, we show that actual corruption incidents increase the public’s perception of corruption more in subnational regions

covered by 3G networks than in regions not covered by 3G. We also find that 3G affects the relationship between perceptions of corruption and actual corruption more in countries with relatively low overall corruptness than in countries with relatively high overall corruptness. This, again, is consistent with the Bayesian model, as each corruption episode constitutes bigger news in countries where such episodes are rare. We corroborate the result that 3G helps expose actual corruption using an alternative measure of actual corruption, which is based on revelations from the Panama Papers concerning offshore entities. Fifth, we explore individual, geographical, and over-time heterogeneity. We find that the effects are stronger for rural residents and respondents with lower socioeconomic status (measured by education and income), and weaker for younger respondents. 3G, on average, negatively affects government approval on all continents, but in Europe and Asia, the negative effect is only among rural residents (for whom the effects are stronger everywhere). The magnitude of the effect of 3G coverage on government approval is relatively stable over the observation period.

These results highlight one of the mechanisms behind the overall effect of 3G on government approval, namely, that mobile broadband internet helps expose actual misgovernance and corruption, suggesting that uncensored mobile broadband internet can be a powerful tool for political accountability. There may be other mechanisms as well. In particular, several observers suggest that social media is particularly well suited for the dissemination of false information.¹ For example, [Tufekci \(2018\)](#) argues that the business model of social media is likely to provide incentives to “stoke outrage, spread misinformation, and appeal to people’s existing biases.” We do not have data to systematically test whether the propagation of false news criticizing the government is also an important factor behind our main result. However, we do illustrate both mechanisms—that is, (i) the exposure of actual corruption and (ii) the dissemination of false narratives through platforms supported by mobile broadband internet—with three case studies: the exposure of corruption of Russia’s former Prime Minister Dmitry Medvedev on YouTube in 2017; the rise to power of the Romanian “Facebook President,” Klaus Iohannis, on an anti-corruption platform in 2014; and the mass dissemination of false narratives through WhatsApp by populist presidential candidate Jair Bolsonaro during Brazil’s 2018 election campaign.

Finally, we examine the electoral implications of 3G expansion. To test whether the 3G-driven disillusionment of voters in their governments translates into lower vote shares for incumbent parties, we focus on Europe. Using subnational-level data on 102 parliamentary elections in 33 European democracies between 2007 and 2018, we

¹In a recent survey of the literature, [Zhuravskaya, Petrova, and Enikolopov \(2020\)](#) discuss well-documented evidence of the massive spread of false stories on social media. Yet, they note there is no systematic study of whether false information is more prevalent in social networks than in the traditional media.

find that incumbent governments lost electoral support after the arrival of mobile 3G networks, corroborating our results on attitudes toward governments. The expansion of 3G coverage in an average subnational region in Europe during the 2008-2017 decade led to a 53-percentage-point increase in the share of the region’s population with access to mobile broadband internet, from 37% to 90%. We show that this regional 3G expansion led to a 4.7-percentage-point decrease in the incumbent party’s vote share. We then investigate what kinds of parties gained from the expansion of 3G networks. We find empirical support for the increasingly prominent hypothesis (see, e.g., [Tufekci, 2018](#)) that—in the age of social media—broadband internet empowers antiestablishment populist politicians. The decade-long expansion of 3G coverage in an average subnational region in Europe increased the vote share for right-wing populists by 4.6 percentage points and of left-wing populists by 3.6 percentage points. We also find that among opposition parties, only populist parties benefited from the expansion of 3G networks—there were no electoral gains for nonpopulist opposition parties, in general, and for Green (environmentalist) parties, in particular. Electoral support for incumbents also decreased with the expansion of 3G networks when populists were in government. We find that turnout decreased by 2 percentage points in an average region as a result of the decade of 3G expansion, which partly explains the effects on vote shares of incumbents and populists. The results, however, are statistically significant when votes cast are expressed as a share of registered voters and not of those who participated in the elections, implying that some voters did change their allegiance.

Our results suggest that, in part, the fall in incumbent governments’ political approval and the rise in popularity of populist parties in Europe are two sides of the same coin. Testing for the exact mechanisms of 3G’s effect on populists’ vote share is beyond the scope of this paper. Why populists—but not other opposition parties—benefit politically from voter disillusionment with incumbent political elites is a promising subject for future research. Overall, we find that the existence of mobile broadband internet enables voters to become more informed about their governments, leading to a fall in government approval, particularly when other sources of independent political information are scarce or nonexistent. However, in European democracies, it also helps antiestablishment populist politicians connect to voters, an effect that cannot be fully explained by the information channel, as nonpopulist opposition parties (so far) have not benefited from 3G expansion.

Our empirical strategy relies both on difference-in-differences and instrumental-variable analyses. We use the variation in the timing of 3G expansion across different subnational regions within countries, controlling for subnational region fixed effects, year fixed effects, and a large set of potential confounders, including measures of economic development, unemployment, and democracy, as well as individual sociodemo-

graphic characteristics. We document the absence of pretrends: the future availability of mobile networks has no effect on government approval, but the effect of past 3G expansion is significant. We show that our results are robust to including country-by-year fixed effects. These results are also confirmed by an event study, in which we focus on the dynamics of government approval around sharp increases in 3G coverage. We find that such sharp increases are associated with a significant drop in government approval, with a magnitude similar to the baseline specification, and that there are no changes in government approval preceding 3G expansion into a region. We also use the techniques developed by [Altonji, Elder, and Taber \(2005\)](#) and [Oster \(2017\)](#) to show that our results are highly unlikely to be driven by omitted-variable bias. Furthermore, we apply to the 3G expansion the instrumental-variable identification strategy designed by [Manacorda and Tesei \(2020\)](#) for the previous generation of mobile networks (2G). The strategy relies on exogenous variation in the regional frequency of lightning strikes per area to predict the speed of expansion of regional mobile broadband internet coverage. Frequent lightning strikes hinder the rollout of telecommunication technologies because—by causing power surges—they substantially increase the costs of providing service and maintaining the infrastructure. This approach confirms the results of the difference-in-differences OLS analysis.

We also present the results for a number of placebo outcomes. In particular, we show that the relationship between mobile broadband internet and government approval is not driven by the effect of the internet on general life satisfaction or pessimism about the future. In addition, we find no impact of 3G expansion on confidence in the local police, which we consider as a placebo outcome because the performance of the local police, in contrast to that of the national government, can be observed by voters directly, without the internet.

The only other multicountry study of the political effects of expansion of telecommunications infrastructure is [Manacorda and Tesei \(2020\)](#), which shows that 2G mobile networks facilitated political protests during economic downturns across Africa between 1998 and 2012. Our paper differs from this important work in two fundamental ways. First, our focus is mobile broadband internet (3G), which is superior to 2G in terms of possibilities for disseminating political information. While 3G enables users to browse the internet freely and seamlessly transfer images and videos—both crucial for the growth of social media—previous-generation networks allowed only texting and very limited internet connectivity. We highlight the effect of this difference by studying 2G expansion as a placebo treatment. We find that, if anything, 2G expansion, on average, is positively correlated with government approval. Furthermore, controlling for the availability of a 2G signal does not affect our results on the effect of 3G. The results of [Manacorda and Tesei \(2020\)](#) on the relationship between 2G and protests in

Africa and our result on the relationship between 2G and overall government approval are not contradictory. This is because our outcome variable reflects the opinion of the majority, whereas protests are often organized by an interested minority that has more incentives than the general public to actively seek political information and self-organize. Our results suggest that it took a new generation of mobile technology for the discontent with government to spread to the general public. Second, we make use of the global coverage of the GWP data, which allows us to shed light on some of the mechanisms by showing heterogeneity with respect to internet censorship, censorship of the traditional media, overall corruptness, and actual corruption incidents.

Broadly speaking, our paper also contributes to the growing literature on the political effects of the internet and social media. Several studies (mostly focusing on single countries) have shown that access to broadband internet hurts the incumbents' political position. For example, the expansion of high-speed cable internet in Malaysia was shown to have contributed to ending the corrupt ruling coalition's 40-year monopoly on power (Miner, 2015). In South Africa, the spread of mobile internet has also shifted votes away from the ruling political party (Donati, 2019). Social media has helped to coordinate protest activity in Russia (Enikolopov, Makarin, and Petrova, 2020). Fergusson and Molina (2019) show that the addition of a new language to the Facebook interface is associated with an increase in protests in countries where this language is spoken. In Europe, the literature has focused on political participation and the rise of populists, showing the change in the effect at the time when social media emerged. Evidence from Germany (Falck, Gold, and Heblich, 2014), the United Kingdom (Gavazza, Nardotto, and Valletti, 2019), and Italy (Campante, Durante, and Sobbrío, 2018) suggests that, initially—that is, before the emergence of social media—in Europe, broadband internet crowded out political awareness with entertainment content, reducing electoral participation, without significant gains for any specific political force. Yet, beginning in 2008—when social media was born—Campante, Durante, and Sobbrío (2018) show that broadband cable internet has contributed to the rise of Italy's populist Five-Star Movement (*Movimento 5 Stelle*). This result was confirmed by Schaub and Morisi (2020) using survey data on the electoral support for populists in Italy in 2013 (Five-Star Movement) and Germany in 2017 (*Alternative für Deutschland*, AfD).

Our contribution to this literature is threefold. First, we document the effects of the expansion of mobile broadband internet on government approval across the world and show that these effects are different from those of earlier mobile technology. Second, we use our global setting to conduct comparative analyses that identify an important mechanism at play. Third, we use election data for 33 European countries over a decade to demonstrate the electoral implications of the mobile broadband internet expansion.

The rest of the paper is organized as follows. Section II presents the data and

the empirical strategy. In Section III, we present the average effect of 3G expansion on government approval for the whole world and discuss the validity of our identification assumptions. Section IV presents comparative analyses. Section V explores the electoral implications of mobile broadband internet expansion in Europe. In Section VI, we illustrate our results with three country case studies. Section VII concludes.

II. DATA AND THE EMPIRICAL STRATEGY

II.A. Main variables

In this section, we briefly describe the main variables of interest. (For details about these measures, as well as descriptions of all the control variables, see Appendix Section A.)

The data on government approval come from the GWP and cover the period from 2008 to 2017. Approximately 80% of the data were collected via face-to-face interviews. The other 20% of the interviews were conducted over the telephone.² The exact questions about government performance in the GWP are: *“Do you have confidence in each of the following, or not: How about the national government? How about the judicial system and courts? How about the honesty of elections? Is corruption widespread throughout the government in (country), or not?”* The respondents could answer *“Yes”* or *“No.”*³ We use the responses to these four questions, also aggregating them using their first principal component and the share of positive attitudes toward the government across these four dimensions. The GWP also includes a question on individuals’ internet access at home: *“Does your home have access to the internet?”*

Because we are interested in estimating the effect of mobile broadband internet availability on attitudes and beliefs, we exploit the variation in the timing of 3G expansion. (The identification strategy is discussed below.) 3G was the first generation of mobile networks that allowed users to actively browse the web on their phones, making online content, including social media, more accessible. The technology was first introduced to the public in 2001, but it took several years for most countries to adopt it. According to the International Telecommunication Union (ITU, 2019), in 2007 there were only 0.04 active mobile broadband subscriptions per capita in the world. By 2018, the figure had jumped to 0.70. Importantly, ITU data show that most of the growth in individual broadband subscriptions over the past decade, in developing

²Telephone interviews were conducted only in countries with at least 80% telephone coverage. This sample consists primarily of high-income OECD countries and the Arab states of the Persian Gulf. Most telephone interviews were conducted via landline telephone. In Section III.B, we show that our results are robust to limiting the sample to face-to-face interviews only.

³Respondents could also choose *“Refuse to answer”* or *“Do not know.”* For the four questions about government performance, the share of respondents choosing these two options varies from 6% to 11%. We have verified that the likelihood of choosing these answers is unrelated to 3G expansion.

and developed countries alike, was due to the expansion of mobile broadband internet access rather than fixed broadband (ADSL or fiber-optic cable) access. We illustrate the global growth of fixed and mobile broadband subscriptions per capita in Appendix Figure A.1.⁴

We use digital maps of global 3G network coverage from 2007 to 2018 provided by Collins Bartholomew’s Mobile Coverage Explorer. These maps put together coverage data submitted by mobile network operators from around the world to the GSM Association, which represents the interests of mobile network operators worldwide. The data consist of 1×1 -kilometer binary grid cells. If a grid cell is covered by 4G, it is also covered by 3G, by definition.⁵ Figure I illustrates the expansion of 3G networks over the entire period of observation. It presents maps of 3G coverage in 2007 and 2018 by grid cells and the corresponding increase in the share of the subnational regions’ territory covered by 3G mobile internet for countries in the GWP sample. Subnational regions are defined by the level of geolocalization provided in the GWP data.

To combine data on mobile network coverage with the GWP surveys, which have region-level localization, we calculate regional 3G coverage in each region and year defined as the weighted average across all grid-cells in each region’s polygon of the value of 3G availability weighted by the population density in each grid cell. (The weights are normalized to sum up to one.)

The resulting dataset covers 840,537 individual respondents in 13,004 subnational region \times year cells, from 2,232 subnational regions of 116 countries. The mean number of times the same region appears in the data is six. Over 75% of the subnational regions appear in the data for at least four years. The mean number of subnational regions per country is 16. On average, 65 respondents are surveyed in a subnational region in any particular year.

To understand the drivers and consequences of the effect of mobile broadband internet on government approval, we use independent measures of corruption, censorship of the internet, and censorship of the traditional media. We use two data sources to measure actual corruption. The first one is the International Monetary Fund’s (IMF’s) Global Incidents of Corruption Index (GICI) from Furceri, Papageorgiou, and Ahir (2019), which is based on text analysis of country reports, prepared by the Economist Intelligence Unit (EIU) and made available to investors on a subscription basis. The index quantifies the intensity of actual corruption by country-year. It is the result of analysis by external (EIU) experts and is distinct from the public’s perception of cor-

⁴The ITU data are available at <https://www.itu.int/en/ITU-D/Statistics/Pages/stat/default.aspx> (accessed on July 25, 2020).

⁵These data are available for all years except 2011. The 2011 data are unavailable due to a change in the company administering the data collection that year. We use the mean of 2010 and 2012 as a proxy for 2011 coverage. All our results are robust to excluding 2011 from the sample.

ruption. This index covers 104 countries in our sample. We use both the time-variant GICI and a measure of overall country corruptness equal to the country mean of the GICI between 2000 and 2017. The second source of data on actual corruption is based on the Panama Papers Database made available by the International Consortium of Investigative Journalists (ICIJ). For each country, we calculate the number of entities featured in the Panama Papers.⁶ For the few countries that are not mentioned in the Panama Papers, we impute this number to be zero. As a baseline, we use the number of entities featured in the Panama Papers scaled by the country’s population size and establish robustness to using the total number of Panama Papers entities. Then, we examine how these two measures of actual corruption—the GICI and the number of entities in the Panama Papers—interact with regional 3G coverage in explaining perceptions of corruption.

We measure censorship of the internet using Freedom House’s Limits on Content score, a component of the Freedom on the Net (FOTN) index. It is available for 46 countries in our sample and ranges from 0 to 35, with higher values implying higher censorship. We use both the time-variant (contemporaneous) censorship measure available by country and year and the time-invariant country-level measure, which is calculated as the mean value of time-variant internet censorship in each country from 2015 to 2017, that is, the years with maximum cross-country variation in time-variant internet censorship. In addition to using these continuous measures of internet censorship, we also create dummy variables for a high level of censorship by using thresholds that indicate natural breaks in the distributions of the respective continuous measures (22 for the time-variant measure and 20 for the time-invariant measure). When we use the binary definition of internet censorship, we extend the sample by including observations with missing FOTN data from countries that one can be sure did not censor the internet. In particular, we assign a zero value to the time-varying dummy for high internet censorship when the FOTN data are missing and the country in that year is a democracy according to the Polity IV dataset (i.e., if the Polity2 score is 6 or above). Similarly, we assign a zero value to the time-invariant dummy when it is missing and the country is a democracy during our sample period (if the over-time mean of the Polity2 score in this country is 6 or above).⁷

The measure of censorship of the traditional media comes from Freedom House’s

⁶We follow [Louis-Sidois and Mougín \(2020\)](#), who used the Panama Papers revelations as a shock to corruption perceptions around the world.

⁷Below, we document that our results are robust to using alternative thresholds for the definitions of the binary measures of internet censorship. We also show that the results do not depend on the imputation of zeros for democracies. The imputation is, however, reasonable because in the sample with nonmissing FOTN data, a dummy for democracy predicts the Limits on Content score to be below 22 with 99.5% probability; and all the countries with the mean Limits on Content score in 2015–2017 above 19 have an over-time mean of a Polity2 score below 6.

Freedom of the Press (FOTP) index. It is available for all 116 countries in our sample and ranges from 0 to 100, with higher values representing higher censorship. As above, we use both a contemporaneous measure and its over-time country mean.

To single out the exogenous source of variation in the speed of regional 3G expansion, we calculate the population-weighted frequency of lightning strikes per subnational region's area using the World Wide Lightning Location Network (WWLLN) dataset. This dataset provides the exact coordinates and time of all cloud-to-ground lightning strikes across the globe. We calculate the average number of lightning strikes in a subnational region per year between 2005 and 2011, weighting each of the lightning strikes by its local population density, measured using a NASA map of population density per square kilometer for each 1×1 -kilometer grid cell. Then, we divide this number by the area of the subnational region. Thus, the resulting measure represents the number of people potentially affected by the lightning strikes per square kilometer in each subnational region. We deem a subnational region to have a high frequency of lightning strikes per area if the region was in the top half of the global distribution of population-weighted lightning strikes per square kilometer.

Finally, we use parliamentary election data from European democracies. Figure A.2 in the Appendix presents maps illustrating the growth in 3G-network coverage between 2007 and 2018 in Europe and the boundaries of the districts, that is, the spatial unit of observation in our European elections data. (The figure is organized similarly to Figure I.) To study the effect of 3G mobile internet expansion on the performance of incumbents and of establishment parties, we use the vote share for the party of the country's top executive at the time of the elections, as well as the combined vote share for the two parties that finished first and second in the first electoral race that occurred in each country since 2007. To analyze the performance of populist parties, we extend the panel dataset on the vote shares of populist parties in Europe from [Algan et al. \(2017\)](#). We classify the parties as populist or nonpopulist based on the Chapel Hill Expert Survey and on text analysis of online sources (see [Guriev and Papaioannou, forthcoming](#), for a discussion of available classifications of populist parties). The data cover 102 elections in 33 European countries from 2007 to 2018 at the level of 398 subnational districts, for a total of 1,250 district-election observations. The mean number of elections per district is 3.25 (the median is 3), and all districts appear in the data at least twice. The data on Green parties cover 97 of the 102 considered elections because, in five elections, the Greens formed joint lists with mainstream nonenvironmentalist parties, making it impossible to measure the vote share for the Greens separately. In the Appendix, we describe these data, present the lists of populist and Green parties, and outline the methodology used to classify parties into populist and nonpopulist.

Details about the exact measures used in the analysis, summary statistics, and

sources of all data are presented in Appendix Section A.

II.B. Main specifications

We estimate the effect of getting access to mobile broadband internet on individuals’ beliefs. As described above, we gauge 3G availability in each subnational region (defined by the GWP localization) of each country in each year by calculating the share of the region’s territory covered by 3G networks in that year, weighted by population density at each point on the map. Then, we relate attitudes toward government to 3G availability using a difference-in-differences model with region and year fixed effects:

$$(1) \quad Gov_approval_{irt} = \gamma_1 3G_{rt} + \gamma_2 Development_{rt} + \mathbf{X}'_{irt} \lambda + \varphi_r + \tau_t + \epsilon_{irt},$$

where i , r , and t index individuals, regions, and years, respectively. *Gov_approval* is a dummy indicating whether the survey respondent has confidence in government. As mentioned above, we use four GWP questions to measure confidence in government. *3G* represents the share of the population in the subnational region with potential access to 3G, our main explanatory variable. φ_r and τ_t are region and year fixed effects, which control for all regional time-invariant characteristics and global time-specific shocks. *Development* represents a measure of regional economic development—an important control as 3G expansion was potentially faster in regions with high economic growth. In the baseline specification, we proxy regional economic development with the log of mean household income among GWP respondents in the region and establish robustness to using nighttime light density as an alternative measure (following Henderson, Storeygard, and Weil, 2011, 2012).⁸ \mathbf{X} is a vector of additional controls: age, age squared, gender, education, marital status, employment status, indicators for urban/rural place of residence, the log of the country’s GDP per capita, the country’s unemployment rate, and dummies for democracy and for advanced democracy.⁹ In the baseline specification, standard errors are corrected for two-way clusters at the level of the subnational regions (to account for correlation over time) and at the level of the country in each year (to account for within-country-year correlation). We establish robustness of the results to using alternative assumptions about the variance-covariance matrix: in particular, the results are robust to correcting for spatial and over-time correlation following Conley (1999), Hsiang (2010), and Collela et al. (2018), and for

⁸In the few region-years where GWP income data are not available (less than 7% of the sample), we use nighttime light density and the country’s GDP per capita to predict regional income. As discussed in Appendix Section B, the results are robust to controlling for nighttime light density. We do not use this variable in the baseline specification, because it is not comparable before and after 2014.

⁹The summary statistics are presented in Table A.1 in the Appendix.

clustering at the country level.

3G mobile service allows users to freely browse the internet from a smartphone and to use social media applications. 3G coverage affects internet use (i) on the extensive margin—by affecting the probability of getting a connection, (ii) on the intensive margin—by affecting the number of hours spent online, and (iii) qualitatively—by changing what people do online. The qualitative difference that a mobile broadband connection makes comes from the fact that a number of social media, such as WhatsApp and Telegram, are particularly well-suited for users with mobile broadband access. The ease of connection also makes a qualitative difference by engaging users in social networks (Rainie and Wellman, 2012). The vast majority of active social media users access social media applications via mobile phones, even when these applications can be accessed through a fixed internet connection.¹⁰ All three of these margins are important for the overall effect of 3G, estimated by Specification (1). The GWP does not have data on the amount of time spent surfing the web and on social media. We can test only for the effect of 3G expansion on having access to the internet at home, as there is a question about this in the GWP. This is a very partial test of the extensive margin because (i) the GWP question does not specify whether home internet access is broadband or a slow connection, and (ii) mobile broadband internet enables people to access the internet outside their homes (e.g., if there is 3G coverage at their workplace but not at home). Nonetheless, we verify that 3G availability predicts internet access at home by estimating a difference-in-differences relationship between the respondent’s internet access at home and 3G coverage in the subnational region of the respondent’s residence:

$$(2) \quad \text{Internet_at_home}_{irt} = \alpha_1 3G_{rt} + \alpha_2 \text{Development}_{rt} + \mathbf{X}'_{irt} \lambda + \varphi_r + \tau_t + \epsilon_{irt},$$

where *Internet_at_home* denotes a dummy variable for self-reported access to the internet at home.

The two main identification assumptions for interpreting the estimation of Specification (1) of the effect of regional 3G coverage on confidence in government as causal are (i) the timing of 3G expansion is uncorrelated with other factors that may affect

¹⁰In 2017, out of 3.196 billion active social media users, 2.958 billion (i.e., 93%) accessed social media via mobile devices (Kemp, 2018). In 2014, this share was slightly lower, but still represented an overwhelming 81% majority (Kemp, 2015). According to YouTube, more than 70% of YouTube watch time comes from mobile devices (<https://www.youtube.com/intl/en-GB/about/press/>, accessed July 19, 2020). Twitter reports that already by 2012, two-thirds of its users were mobile, and by 2015 the share of mobile users had reached 80% (<https://www.statista.com/chart/1520/number-of-monthly-active-twitter-users/>, accessed July 19, 2020). In contrast, the growth of mobile internet use outside social media was much slower: the average share of mobile traffic was only 16% in 2013 and 50% in 2017 (<https://www.broadbandsearch.net/blog/mobile-desktop-internet-usage-statistics>, accessed July 19, 2020).

public attitudes toward government, and (ii) 3G expansion is not itself driven by the expectation of changes in government approval or by any unobserved factor that could generate a spurious correlation between government approval and 3G network coverage. These assumptions are not directly testable. However, below in Section III.A, we present a number of robustness and placebo exercises, as well as tests in the spirit of Altonji, Elder, and Taber (2005) and Oster (2017), which do suggest that the differences-in-differences results can be interpreted as causal.

To address the remaining concerns that the identification assumptions in our baseline difference-in-differences specification could be violated, we use variation in the frequency of lightning strikes per square kilometer in each subnational region to predict the speed of regional 3G expansion—the identification strategy first used by Manacorda and Tesei (2020) for 2G-network expansion in Africa. The frequency of lightning strikes has been shown to affect the diffusion of digital technologies due to an increase in the expected costs associated with voltage spikes and dips (e.g., Andersen et al., 2012). The equipment needed for mobile-phone infrastructure, including the mobile broadband networks infrastructure, is particularly sensitive to electrical surges caused by lightning strikes, which can lead both to immediate damage and to quicker depreciation of the equipment over time (Zeddiam and Day, 2014; Martin, 2016). Power-surge protection can partially alleviate the problem, but it is expensive, not always effective, and less readily available outside developed countries. We predict slower 3G expansion in regions with a high frequency of population-weighted lightning strikes per square kilometer. As both the endogenous regressor (regional 3G coverage, $3G_{rt}$) and the exogenous source of variation (lightning-strike frequency per square kilometer) vary at the regional level, as a baseline, we estimate the following first-stage equation at the region-year level:

$$(3) \quad 3G_{rt} = \delta_1[\textit{Lightning}_r \times t \times \textit{Rich}_{c_r}] + \delta_2[\textit{Lightning}_r \times t \times \textit{Poor}_{c_r}] + \mathbf{Z}'_{rt}\mu + \varphi_r + \tau_t + \epsilon_{rt},$$

where $\textit{Lightning}_r$ denotes a dummy indicating subnational regions with an above-median population-weighted frequency of lightning strikes per square kilometer; \textit{Rich}_{c_r} and \textit{Poor}_{c_r} are dummies indicating the countries with above- and below-median per capita income; and \mathbf{Z} stands for all the other controls. We include all the region-level and country-level baseline controls described above. In addition, we control for other potential determinants of 3G expansion that can correlate with lightning-strike frequency. In particular, we extend the list of covariates to include linear time trends interacted with the subnational regions' share of territory covered by deserts, share of territory covered by mountains, maximum elevation, and dummies for each quintile of population density. To control for the fact that the initial expansion of 3G networks affects the speed of subsequent expansion, we also control for linear time trends inter-

acted with regional 3G coverage in 2008, a dummy for whether the region had any 3G coverage in 2008, and a dummy for whether the country had any 3G coverage in 2008. We, then, estimate the second stage using predicted regional 3G coverage.

The identification assumption behind this approach is that the frequency of lightning strikes per square kilometer affects trends in government approval only through its effect on 3G expansion conditional on all other covariates. We also establish robustness to using individual-level data instead of region-year averages, an approach that places higher weight on more populous regions, as there are more observations per region in the GWP in regions with a larger population. As we show below, the results of the IV and OLS specifications are qualitatively similar; the magnitudes are somewhat larger in the IV estimation.

III. MOBILE BROADBAND INTERNET AND GOVERNMENT APPROVAL

Table I presents the results of estimating the effects of mobile broadband internet availability with the baseline difference-in-differences specification. Panel A presents the results for the full sample; Panel B, for the subsample of rural residents. Different columns of the table consider different measures of government approval as the outcome variable. The expansion of 3G networks, on average, is associated with individuals becoming more aware of government corruption and less confident in their country’s government and institutions. The results are statistically significant for all four measures of government approval (Columns 1–4) and for the two aggregate measures, that is, the share of positive answers and the first-principal component of the four measures (Columns 5–6), both for the full sample and for the subsample of rural residents (Panels A and B, respectively).

In Column 1 of Appendix Table A.2, we illustrate how 3G expansion affects internet access at home. We find that 3G expansion within the respondent’s region of residence significantly predicts internet availability at home. This is consistent with the observation that access to mobile broadband networks increases the extensive margin of internet use. However, 3G mobile networks have an effect on government approval above and beyond their effect on internet access at home. We show this in Columns 2 to 5 of Appendix Table A.2. The average effect of regional 3G coverage is not affected by including a dummy for having internet access at home in the list of covariates (Column 2). The effect of 3G expansion on government approval is significantly negative, both when there is and when there isn’t an internet connection at home. The effect is twice as large in magnitude for individuals *without* access to the internet at home than for individuals *with* access to the internet at home (Columns 3 to 5 of Table A.2). These estimates suggest that even when people have access to the internet, getting

access to *mobile* broadband internet significantly affects the way they use it.¹¹

The magnitude of the effect of 3G coverage on government approval, documented in Table I, is substantial; it is particularly large for residents of rural areas. The average increase in regional 3G coverage between 2008 and 2017 across the regions in the GWP sample is 0.39. As discussed in more detail in Appendix Section C, we use this increase as the basis to understand the magnitude of the effect. For example, the estimates in Column 1 of Table I imply that in an average region 3G expansion in the last decade led to a decrease in the confidence of respondents in their country’s government by 2.5 ($= -0.063 \times 0.39 \times 100$) percentage points in the full sample and by 3.5 percentage points among rural residents (from the mean levels of 51% and 54%, respectively). Similarly, as reported in Column 4, it led to a decrease in the share of people who think that the government is not corrupt by 1.4 percentage points in the full sample and 2.1 percentage points among rural residents (from the mean of approximately 22%). The results for the other measures of attitudes toward government institutions are very similar. According to the aggregate measure (Column 6, Panel A), a decade-long expansion of 3G networks in an average region led to a 2.2-percentage-point decline in government approval. (We normalize the first-principal component of the government approval variables to vary between zero and one for the ease of interpreting the magnitude of the effect.) The coefficient on the unemployment rate (measured in percentages) in the same regression is -0.010 , implying that the effect of a decade-long 3G expansion has the same-size effect on government approval as a 2.2 ($= \frac{0.057 \times 0.39}{0.010}$) percentage-point rise in the national unemployment rate.

To calculate persuasion rates for the hypothetical message “do not approve of your government,” one needs to make a number of important assumptions. We describe these assumptions in detail in Appendix Section C. Furthermore, one needs an estimate of the size of potential spillovers in exposure to the anti-government message from those connected to mobile internet to those who are not connected. Specifically, persuasion rates are inversely proportional to the number of people who, on average, get exposed to anti-government messages for each mobile device that is connected to the internet, which we denote by N . One could argue that, particularly in develop-

¹¹As mentioned above, the estimates presented in Table I take into account both the extensive and the intensive margins of the effect of the telecommunications infrastructure on internet use, which, in turn, affects attitudes. They also take into account the qualitatively different experience of using social media on a mobile phone compared to a fixed-line connection. Therefore, a hypothetical 2SLS estimation, in which one predicts internet access at home with regional 3G coverage and then uses this prediction for estimating the effect of internet access at home on government approval would lead to a gross overestimation of the effect of the internet on government approval. Such a specification would incorrectly imply that 3G only affects the probability of having a connection to the internet at home. In reality, with the arrival of the 3G technology, people who have already been using the internet started using it more because the broadband connection is more convenient and started using it differently because 3G technology is particularly conducive to social media use.

ing countries, such spillovers are substantial. A case study that we present below in Section VI.A, about a YouTube documentary exposing corruption of Russia’s Prime Minister, suggests that, indeed, people without access to the internet also get exposed to content that is available only online. As we lack the data necessary to estimate such spillovers, we can only calculate persuasion rates up to a factor of $\frac{1}{N}$. Assuming $N = 1$, that is, that there are no spillovers, we calculate the upper bound for the persuasion rates implied by the estimates for the first-principal component of the government approval variables (Column 6) to be 17.6% in the full sample and 24.2% in the sample of rural residents.

Panel A of Figure II illustrates the main result from Table I. On the horizontal axis, the figure plots the increase in regional 3G coverage in year t since 2008. The outcome variable is the residual of the first-principal component of the government approval variables in year t (after subtracting the effects of all the controls, including region and year fixed effects). Panel B provides a similar graph for the relationship between the residuals of having internet access at home in year t (again, after subtracting the effects of all the controls) and the growth in 3G coverage.¹² The graphs present the nonparametric relationship between the increase in 3G coverage and the outcome variables, along with their confidence intervals, constructed using a block bootstrap at the level of the clusters, and the data averages by equal-size bins.¹³ The figure shows that, on average, 3G expansion led to a drop in government approval (Panel A) and an increase in internet access at home (Panel B).¹⁴

III.A. Addressing identification challenges

Can these results be interpreted as causal? In this section, we present evidence suggesting that variation in 3G coverage is plausibly exogenous. We corroborate this evidence by performing an instrumental-variable analysis, in which we use the frequency of lightning strikes per area in the subnational regions as an exogenous source

¹²To generate the outcome variables net of controls, we first regress the variable of interest on the change in regional 3G coverage since 2008 and all the controls. We then take the residuals and add to them the estimated effect of the change in regional 3G coverage since 2008. This strategy accounts for the correlation between our main explanatory variable and the other controls.

¹³To construct the confidence intervals, we first generate 55 equal-size bins for the change in regional 3G coverage since 2008. We then perform 1,000 block-bootstrap iterations, sampling at the level of the clusters. In each iteration, we save the average of the outcome variable for each of the bins and the number of observations used to construct that average. After performing 1,000 iterations, we calculate the 5th and 95th percentiles of the outcome variable for each of the bins, weighting by the number of observations in each of the bins in each iteration. Finally, we perform local polynomial smoothing (lpoly) to draw the confidence intervals, using the values of the 5th and 95th percentiles for each of the bins.

¹⁴Appendix Figure A.3 presents the dynamics of raw government approval and 3G coverage separately in regions with high and low average annual growth of 3G coverage, illustrating the pattern in the data behind our difference-in-differences estimates.

of variation in the speed of the expansion of 3G networks.

Country×year FEs and pretrends. To make sure that our results are not driven by differential country-level dynamics, we redo the analysis controlling for country×year fixed effects, thus, relying only on the differential expansion of 3G in different subnational regions within countries. This is a very demanding control, because it eliminates part of the relevant variation as 3G networks often expanded to all regions of a country at the same time. Nonetheless, the results (presented in Panel A of Table A.3 in the Appendix) are largely robust. After partialing out all of the country×year variation, 3G mobile internet remains an important determinant of attitudes toward government. The effect of 3G is statistically significant for five of the six measures of government approval, with the results being most precise for the two aggregate measures, which are the least noisy among the considered outcomes (Columns 5 and 6). The point estimates are smaller than in Table I, which could be explained by the fact that part of the relevant variation is not accounted for in this specification.

A major potential concern with our difference-in-differences identification strategy is that 3G networks might expand in regions with falling confidence in government. To address this concern, we examine the effects of lags and leads of regional 3G coverage. In Panel B of Table A.3 in the Appendix, we repeat the analysis presented in Panel A, but for regional 3G coverage in year $t + 1$. We find that 3G coverage next year is not significantly related to government approval this year. In Panel C of this table, we test for the equality of the magnitude of the coefficients on regional 3G coverage and its lead (presented in Panels A and B of the table, respectively). The p-values from this test indicate that we can reject equality of the effects for five of the six outcomes and, as above, the difference is most precise for the aggregate measures of government approval. This analysis suggests parallel pretrends in the specification with country-year fixed effects, i.e., when we partial out all country-level trends and shocks.

Figure III presents the point estimates along with their confidence intervals for the coefficients on several lags and leads of regional 3G coverage from the regressions with country-year fixed effects and with the first-principal component of the government approval variables as the outcome. Consistent with the parallel pretrends assumption, we find that the future availability of mobile networks has no effect on government approval, but the effect of past 3G expansions is significant for the first lag; it stays negative, but becomes insignificant, for the second lag. The p-values for the test of equality between the coefficients on the leads of 3G coverage and on 3G coverage at t presented below each point estimate show that the coefficient on 3G coverage at t is significantly larger in magnitude (in absolute value) than the coefficients on its leads.

If we do not partial out all of the country-year dynamics, a similar pretrends test

would yield negative significant coefficients on the leads of regional 3G coverage in the full sample, because in many countries 3G expansion was gradual and there is a very strong, significant autocorrelation in the level of 3G coverage. To test for pretrends in government approval without country-year fixed effects, one needs to focus on cases in which there is a sharp discontinuous increase in regional 3G coverage; we do this in the next subsection.

Event study and pretrends. To further validate our pretrends analysis, we conduct an event study focusing on sharp increases in regional 3G coverage. As an event, we consider the situation (i.e., the region-year combination) in which regional 3G coverage increased by more than 50 percentage points within the previous year. By definition, this could happen only once per region, if it happens at all, provided that regional 3G coverage never falls substantially. On average, regional 3G coverage increases by 76 percentage points during the event.¹⁵ There are 452 regions in 65 countries that experienced such a sharp increase in 3G coverage in one year between 2008 and 2018. Focusing on the sample of respondents from these regions (130,406 observations), we estimate the average dynamics of government approval around these events, that is, the sharp increases in regional 3G coverage.

The results are presented in Table II. First, we verify that our baseline relationship holds in this subsample using the first-principal component of the government approval variables as the dependent variable (Column 1). Second, instead of regional 3G coverage, we use a postevent dummy as the treatment variable (Column 2). The results are very similar to the baseline in both cases.¹⁶ In Column 3, we present the event-study specification: we regress government approval on year dummies relative to the year of the event and all the baseline controls. In Columns 4 to 6, we repeat the same exercise in the subsample of rural residents. We find that government approval falls right after a sharp increase in regional 3G coverage (see Columns 3 and 6). All the coefficients on the postevent dummies are statistically significant and their magni-

¹⁵For the vast majority of regions, 3G expands monotonically. In 95% of the region-year observations, the change in 3G coverage is positive from one year to the next. Among all the subnational regions with 3G data, only 14 regions from three countries experienced sharp drops in 3G coverage from one year to another during our observation period. We exclude these regions from the event-study analysis in order to have a clean definition of the event. These regions are included in the sample for the baseline analysis. None of our results for either the baseline analysis or the event study depend on whether we include these regions or exclude them. Figure A.4 in the Appendix presents the distribution of events across years: It shows that we detect sharp increases in 3G in all years except 2011 and 2012, which is explained by the fact that the data for 2011 are interpolated, thus, by construction, there are no sharp increases in 3G coverage between 2010 and 2012. The figure also lists the countries with the events.

¹⁶219 regions from 36 countries have variation in the postevent dummy in the resulting sample because the 2018 GWP data are not available and not all regions are present in the GWP data in all years. All results presented in Table II are robust both to restricting the sample only to regions with variation in the postevent dummy and to including in the sample all regions without events, that is, using the full GWP sample.

tudes are similar to those presented in Table I. In contrast, all the coefficients on the pre-event dummies are very small in magnitude and statistically indistinguishable from zero, thus, confirming the absence of pretrends. In the last two rows of the table, we present the p-values from the tests of equality of the coefficients on dummies indicating years t and $t - 2$ and between the average effects for the years t and $t + 1$ as compared to the average effect for years $t - 2$ and $t - 3$. One of four tests gives the p-value of 0.119; in all other cases, the difference in magnitudes of the effects before and after the event is significant.¹⁷

We illustrate these results in Panel A of Figure IV. The figure presents the coefficients on the dummies indicating the years around the event with government approval as the dependent variable (darker line, left axis). On this figure, we also illustrate the treatment in the event study by showing the coefficients on year dummies around the event with regional 3G coverage as the outcome variable (lighter line, right axis): by construction, we observe a sharp increase in 3G coverage at the event year.¹⁸

We also verify that the events in our event study are not associated with a concurrent change in government approval in nonevent regions of the same countries (i.e., in those regions that did not experience such a sharp increase in 3G coverage).¹⁹

¹⁷To understand the size of a pretrend that can be rejected, we follow Roth (2019) and perform the following test. We assume the presence of a linear time pretrend with slope ξ and that the coefficients on the three forwards of the event dummy are jointly normally distributed. We take the variance-covariance matrix for this distribution from the estimation of the three pre-event coefficients in the event study (Column 3 of Table II). By construction, the mean of this distribution is $(S\xi, 2\xi, \xi)$, where S is the average difference in the number of years between the period $t - 1$ and each of the periods before $t - 3$. Taking into account the fact that $S \geq 3$ and that the pretrend is more easily rejected for larger S , for simplicity, we set $S = 3$. Then, we search for the smallest absolute value of ξ , such that, in 90% of all realizations of the multivariate normal distribution, at least one of the pretrend coefficients is statistically significant at the 10% significance level. In particular, we take hypothetical ξ from a grid between 0 and -0.05 , and for each value of ξ , we perform 100,000 random draws from the corresponding multivariate normal distribution to calculate the percentage of draws with at least one of the pretrend coefficients significant at the 10% significance level. The smallest $|\xi|$, such that in 90% of draws at least one of the pretrend coefficients is significant, is 0.0188. Thus, we are able to reject a pretrend with a slope that is larger than 0.0188 in absolute value, which is approximately equal to the absolute value of one half of the estimated treatment effect from Column 2 of Table II.

¹⁸Appendix Figure A.5 illustrates the dynamics of raw government approval around the event in the sample of regions for which we observe government approval both before and after the event. The figure presents the mean of government approval net of region fixed effects to account for changes in the sample composition across years.

¹⁹To do this, we restrict the sample to those countries where at most 60% of all GWP respondents are located in regions where the event occurred. Then, we randomly draw placebo-event regions among those that did not have an event from the country-years, in which other regions had an event. We repeat this exercise 500 times, comparing the distributions of the point estimates and their t-statistics for the effect of such placebo treatments with those for the actual treatment in the same sample of countries. The results are presented in Figure A.6 in the Appendix. We find that both the coefficient and its t-statistics from the estimation of the effect of the true event are outside of the corresponding distributions for the placebo events. We also verify that our results are not driven by influential observations. In Appendix Figure A.7, we present the residual scatterplot from the regression at the region-year level in the event-study sample. This regression is similar to the one presented in Column 2

A number of recent studies show that, in the presence of heterogeneous treatment effects, the coefficients on the leads and lags of the treatment variable in an event study might place negative weights on the average treatment effects for certain groups and periods (e.g., see [Borusyak and Jaravel, 2018](#); [Goodman-Bacon, 2018](#); [Sun and Abraham, 2020](#); [De Chaisemartin and D’Haultfœuille, 2020](#)). To address this concern, following [De Chaisemartin and D’Haultfœuille \(2020\)](#) we use an alternative estimator that solves this issue by calculating the average of all these treatment effects.²⁰ The results are presented in Panel B of Figure IV. Appendix Table A.4 provides the underlying regression table. Similarly to the OLS estimation of the event study, these results indicate that government approval decreases sharply after a sharp increase in regional 3G coverage, whereas before the event, the effects are not distinguishable from zero.

2G as a placebo treatment. A potential concern is that 3G availability may affect individuals’ beliefs through other mechanisms than providing access to mobile broadband internet. To address this concern, we consider the effect of the expansion of 2G networks, which allow making phone calls and sending text messages, but provide very limited internet capabilities and, in particular, do not allow browsing the internet freely or watching online videos. The key difference between 2G and 3G mobile networks is that, unlike 2G, 3G facilitates the immediate dissemination of photos and videos, which can invoke substantially stronger emotional reactions and therefore have more profound political implications than information in text form.²¹ For example, the Arab Spring started after a smartphone-recorded video of the self-immolation of a street vendor, Mohamed Bouazizi, went viral on social media ([Castells, 2015](#), p.22). In contrast, the self-immolation a few months earlier of another street vendor, Abdesslem Trimech, had no political implications, presumably because nobody recorded it ([Gurri, 2018](#), pp. 47-48).²²

If individuals’ beliefs were affected not by access to mobile broadband internet but rather by some other aspects of communications technology, one should expect similar effects of the expansion of 2G and 3G networks. In Table III, we show that, in

of Table II (apart from the level of aggregation). The results are robust to excluding observations that are away from the main cloud (marked on the scatterplot) or regions to which these region-year observations belong.

²⁰We use the software package DID_MULTIPLEGT developed by [De Chaisemartin and D’Haultfœuille \(2020\)](#). Other papers (e.g., [Borusyak and Jaravel, 2018](#); [Sun and Abraham, 2020](#)) propose similar estimators.

²¹Manuel Castells makes this point in several of his books; see, e.g., [Castells \(2015, p.15\)](#) and [Castells \(2019, p.20\)](#). The fact that videos are more powerful than text has also been shown in other contexts (e.g., [Durante and Zhuravskaya, 2018](#)).

²²Observers also argue that the spread of information about the events in Tunisia in 2011 across the Arab states was also driven by mobile broadband internet and social media: “Most of Al Jazeera’s Tunisia footage came from cell phone videos, taken by the public on the spot and communicated via Facebook. They were then re-posted online—on Al Jazeera’s website, on YouTube, and on thousands of niche sites.” ([Gurri, 2018, p.48](#)).

contrast to the effect of 3G presented above, the expansion of 2G networks, if anything, is associated with an *increase* in government approval (Columns 1 to 6 of Panel A), suggesting that the populace may credit the government—justifiably or not—for the construction of new infrastructure that improves its quality of life.

In Panel B of the table, we also show that controlling for 2G availability does not affect the estimates of the effect of 3G. In addition, as we show in Column 7 of Table III, unlike 3G coverage, regional 2G coverage is not related to respondents’ internet access at home. These findings suggest that the negative effect of 3G on government approval is driven by its effect on mobile broadband internet access rather than by other features of the expansion of mobile networks. As we noted in the introduction, the fact that we find no negative effect of 2G on overall government approval does not contradict the findings of [Manacorda and Tesei \(2020\)](#), who show that 2G mobile networks facilitated protests in Africa during recessions. This is because protests are often organized by interested minorities that have more incentives to seek political information than the general public, and therefore, the expansion of 2G, which allows texting, did help these minorities to get informed and to coordinate and organize the street protests, while having no effect on the majority’s opinion about the government.²³

Variation in observables as a proxy for unobserved variation. We follow the methodologies of [Altonji, Elder, and Taber \(2005\)](#) and [Oster \(2017\)](#) to understand whether unobserved variation is likely to explain our results. First, we construct the index of observables that is the best predictor of 3G availability, by taking the fitted value from a regression of 3G on all the controls. Then, we regress our outcome variables on this index of observables, controlling for region and year fixed effects. The results are reported in Panel A of Table A.5 in the Appendix. We find that the predicted-from-observables 3G availability is not significantly related to government approval, and the point estimates have the opposite sign of the effect of 3G for four of the six outcomes, including both aggregate measures of government approval. This suggests that, at least for these four outcomes, selection on unobservables is not driving the results, under the assumption that the observables are representative of the unobservables.

Second, in Panel B of Table A.5, we report Oster’s δ statistic, indicating how much more important unobservables need to be compared to observables to fully explain our results by omitted-variable bias. Following [Oster \(2017\)](#), we set the value of R_{\max}^2 —the R-squared from a hypothetical regression of the outcome on treatment and both observed and unobserved controls—to be equal to $1.3\tilde{R}^2$, where \tilde{R}^2 is the R-squared from the baseline estimation (Table I). In the two cases where observables should be positively selected from unobservables to explain our results (Columns 2 and 4), the

²³[Enikolopov, Makarin, and Petrova \(2020\)](#) show that the expansion of the social media platform VK in Russia increased both the likelihood of protests and support for the regime.

values of δ are 5.8 and 1.6. For all the other outcomes, observables should be negatively selected from unobservables to explain our results; for these outcomes, the δ s range between -4 and $-1,000$.²⁴ Both the magnitude and the sign of these statistics suggest that it is highly unlikely that our results are spuriously driven by unobserved variation.

The frequency of lightning strikes as an IV. Finally, we use the identification strategy proposed by [Manacorda and Tesei \(2020\)](#), who show that in Africa the incidence of lightning strikes predicts local trends in the expansion of 2G mobile networks. We use differences in the regional frequency of lightning strikes per square kilometer as an exogenous source of variation in the speed of the expansion of mobile broadband internet service. During thunderstorms, the electrostatic discharges can damage mobile-phone infrastructure, increasing the cost of providing mobile service. This is the case for both 2G and 3G infrastructure. For this reason, one could expect slower 3G expansion in places with a high frequency of lightning strikes. Moreover, one should expect the adoption of mobile broadband infrastructure to be more affected by lightning strikes in lower-income countries, because providers in these countries typically have fewer resources to protect equipment from being damaged—for instance, by using power-surge protection technology—or to repair it in case of damage.

As discussed in the methodology section, we predict regional 3G coverage with a linear time trend interacted with a dummy for a high frequency of lightning strikes per square kilometer in a subnational region, separately in countries with above- and below-median GDP per capita. To control for other factors that might influence the speed of 3G expansion and that can be correlated with the frequency of lightning strikes per square kilometer, we also include linear time trends interacted with the subnational regions' share of territory covered by deserts, share of territory covered by mountains, maximum elevation, and dummies for each quintile of population density. To account for differential trends in 3G expansion depending on its initial level at the beginning of the observation period, we also control for linear time trends interacted with the regions' initial (2008) level of 3G coverage and dummies for whether the region and the country had any 3G coverage in 2008.

We illustrate the first-stage and reduced-form relationships with graphs summarizing the data in the subsample of countries with below-median GDP per capita, which provides most of the variation for the IV estimation. The importance of the frequency of lightning strikes for the expansion of 3G networks is illustrated by Figure A.9 in the Appendix. It shows the evolution of regional 3G coverage separately in subnational regions with a high and low frequency of lightning strikes per square kilometer, limit-

²⁴Figure A.8 in the Appendix reports the sensitivity of the value of Oster's δ to alternative assumptions about the size of R_{\max}^2 for the case of the aggregate government approval. It shows that even in the case of maximum possible $R_{\max}^2 = 1$, Oster's $\delta = -70$.

ing the sample to countries with within-country variation in the frequency of lightning strikes. Appendix Figure A.10 illustrates how the frequency of lightning strikes per area affects government approval. The figure shows that, on average, government approval, net of all controls, decreased in subnational regions with a low frequency of lightning strikes and increased in subnational regions with a high frequency of lightning strikes.

Since both regional 3G coverage and the frequency of lightning strikes are defined at the level of subnational regions, in the main specification, we perform the regression analysis at the region-year level, using the mean of government approval in the regions as the outcome variable. Table IV reports the regression results for this specification. Column 1 of Table IV presents the first-stage relationship for the full sample. We find that the adoption of 3G technology is significantly slower in regions with a high frequency of lightning strikes per square kilometer, and that this effect is stronger—both in terms of magnitude and in terms of statistical significance—in countries with below-median income. The overall F-statistic for the excluded instruments is 21, driven mainly by the strong relationship for the countries in the lower half of the income distribution. The second stage, presented in Column 2, confirms our main result: 3G expansion leads to a significant decline in government approval. Columns 3 and 4 show the IV results for the subsample of rural residents. Because much of the first-stage variation is driven by poorer countries, in Columns 5–8 of Table IV, we repeat the analysis, focusing on the subsample of countries with below-median GDP per capita; we find similar results.²⁵ Table A.6 in the Appendix reports results using individual-level data, which place more weight on more populous regions. The first stage is substantially weaker in this specification: in the full sample of countries, the F-statistics for all and rural respondents are 11 and 13, respectively. The relationship between the frequency of lightning strikes per square kilometer and the expansion of 3G networks at the individual level is driven solely by countries with below-median GDP per capita. In the subsample of these countries, the F-statistics fall below 10. To address this problem, we report weak-instrument-robust Anderson-Rubin confidence intervals for the effect of 3G coverage, which show that the estimates are significant. Overall, the IV results are robust to using individual-level data despite a weaker first stage.

The magnitude of the IV estimates is substantially larger than of the OLS estimates presented in Table I. However, as most of the variation in the first stage comes from countries with below-median GDP per capita, the relevant comparison of the magnitude of the OLS and IV coefficients should come from this sample. We report the corresponding OLS estimates, keeping the same sample and the same set of controls

²⁵To rule out the potential concern that the first-stage relationship is driven by a small number of outliers (Young, 2020), we verify that the results are very similar if we use bootstrap standard errors with sampling at the cluster level. The precision of the first stage is practically unaffected, and the second-stage results are slightly more precise.

as in the IV regression, in Columns 6 and 8 at the bottom of Table IV. Using these estimates as the benchmark, we find that the magnitude of the point estimates in the IV regressions is about 2.5 times as large as in the corresponding OLS regressions (e.g., -0.329 vs. -0.120 , for all respondents, as reported in Column 6).

Given the results of the analyses of the validity of the OLS difference-in-differences specification presented above, it is unlikely that this difference is driven by endogeneity of regional 3G coverage. The first likely explanation for the difference in the magnitude between the OLS and IV estimates is the Local Average Treatment Effect (LATE) in the presence of heterogeneity of the effect of 3G on government approval. If mobile broadband internet has a larger effect on government approval among complier regions (regions where 3G expansion is, potentially, constrained by the frequency of strikes) than among noncomplier regions (regions where the expansion of 3G networks is not affected by lightning-strike frequency, for instance, because power-surge protection may be available), one should expect the IV estimates to be larger than the OLS estimates. It is probable that the population of the complier regions is particularly affected by political messages on social media. This may be because the ability to get power-surge protection when needed is positively correlated with the overall level of development in the region, which, in turn, is correlated with how informed the regional population is. Therefore, one could expect the population in the complier regions to be relatively less informed, making them more receptive to new political information compared to the residents of noncomplier regions.²⁶ The second potential source of the difference between the OLS and IV estimates is measurement error. There are several sources of such measurement error: (i) Access to mobile broadband internet is subject to numerous weather shocks, as both severe rain and wind affect connectivity (Schulman and Spring, 2011). (ii) Each year, the exact timing of the measurement of 3G network coverage does not correspond to the timing of the GWP surveys, and 3G coverage does evolve throughout the year. (iii) Providers may submit inaccurate or outdated data to the GSM Association, the ultimate source of the dataset on 3G network coverage.²⁷

As with OLS, we benchmark the magnitude of the IV estimates, comparing them to the effect of unemployment. The coefficient on the unemployment rate in Column 6 of Table IV is -0.028 . Hence, a decade-long increase in 3G coverage in an average

²⁶We control for the overall level of regional development with region fixed effects. We cannot observe which regions are compliers because the definition of compliers involves a counterfactual level of 3G expansion under an unobserved alternative level of lightning-strikes frequency. In Appendix Section D, we describe which countries provide observations that drive the variation in the first stage, as highlighted in Appendix Figure A.11.

²⁷It is possible that measurement error in 3G is nonclassical, that it is correlated with other determinants of government approval, such as governance quality. Most of this potential correlation is controlled for by region fixed effects and other covariates. If the changes in the quality of the measurement of 3G are correlated with the changes in governance quality, this could also explain the difference between the magnitudes of the OLS and IV estimates.

region of 0.39 is equivalent to a 4.6 ($= \frac{0.329 \times 0.39}{0.028}$) percentage-point increase in the national unemployment rate.

The persuasion rates implied by the IV estimates make sense only if the spillovers are substantial. In the full sample of countries, the persuasion rates are equal to $\frac{73.7}{N}$ for all residents and $\frac{80.1}{N}$ for rural residents, and in the sample of countries with below-median GDP per capita, they are $\frac{115.1}{N}$ and $\frac{133.0}{N}$, respectively. N here is the average number of people exposed to the message per smartphone with a mobile internet connection. It is reasonable to conjecture that, in complier countries, where the development of ICT is constrained by lightning strikes, N is large. (Details of the calculations and all the assumptions behind them are provided in Appendix Section C.)

Overall, the results we present in this section strongly suggest that the negative effect of 3G mobile networks on government approval can be interpreted as causal.

III.B. Robustness

Alternative assumptions about the variance-covariance matrix. Table A.7 shows that the results are robust to alternative assumptions about the correlation between the error terms. We take the specification presented in Column 6 of Panel A of Table I as the baseline (also reproduced in row 1 of Table A.7) and show in row 2 that the standard errors are only slightly larger with clusters at the country level. We then proceed to test the robustness of the results to correcting standard errors for spatial correlation following Conley (1999), Hsiang (2010), and Collela et al. (2018). In rows 3 to 8, we report the standard errors corrected for spatial correlation of the error terms within 500- and 1,000-kilometer radii with autocorrelation up to 10-year temporal lags. In all cases, the estimated effect is statistically significant at the 1% level. In addition, in Appendix Table A.8, we report the regression results for an aggregated region-level panel, in which we take simple averages of the dependent variables across individuals in each subnational region and year. As in the baseline specification, we control for the region and year fixed effects, as well as the region-level and country-level covariates (namely, we include regional-level income and the country’s per capita GDP, democracy, and unemployment in the set of covariates). The results are robust.

3G coverage and population density weights. Our baseline measure of regional 3G coverage takes into account differences in population density within regions to account for the fact that mobile networks should only matter in places where people actually live. To verify that our results are not driven by any effect of population density on government approval, we conduct two exercises. First, in Panels A and B of Appendix Table A.9, we report the results of estimating Specification (1) using a measure of regional 3G coverage equal to the share of grid cells within each region and year that are covered by 3G networks (i.e., without population density weights). Second, in

Panels C and D of this table, we replicate the results presented in Table I, using the baseline measure of regional 3G coverage, but controlling for year dummies interacted with dummies for each quintile of population density. In both cases, the results are very similar to those reported in Table I, suggesting that our results are not sensitive to how we measure regional 3G coverage.

Alternative proxy for subnational economic development. In Section B of the Appendix, we show that our results are robust to using nighttime light density as an alternative proxy for regional economic development, and we discuss the properties of this control.

Robustness to excluding individual countries. We also have verified that our results are robust to excluding any one country from the sample. In particular, we conducted this exercise for the specification presented in Column 6 of Table I.

The effect over time. We explore whether the effect of 3G coverage on government approval changes over time by replacing regional 3G coverage in Specification (1) with its interaction terms with dummies for all consecutive two-year time periods in our sample. We find that the effect is stable and does not systematically change over time.²⁸ The results are reported in Appendix Table A.10 and illustrated in Figure A.12, which plots the over-time evolution of the effect of 3G coverage.

Subsample of observations from face-to-face interviews. For most country-years in the GWP, the data were collected via face-to-face interviews. However, in certain countries with at least 80% telephone coverage, the data were collected over the telephone. In Table A.11 in the Appendix, we show that the results are robust to excluding observations from telephone interviews and are, therefore, not driven by potential differences between the sample of respondents from in-person interviews and telephone interviews.

Balance in individual characteristics. We have checked whether the expansion of regional 3G coverage is correlated with the composition of individuals in the GWP surveys. Only a few of the large number of individual characteristics are unbalanced with respect to regional 3G expansion. We show that this imbalance does not drive our results. First, we replicate the results applying the methodology developed by [Hainmueller \(2012\)](#) that uses entropy balancing to reweigh observations in order to achieve balance. Second, we show that the results are robust to focusing on the subsamples without any variation in the unbalanced individual characteristics. Details of these analyses are presented in Appendix Section E.

²⁸This provides further evidence that it is unlikely that there is time-specific heterogeneity in the treatment effect that could potentially lead to the standard difference-in-differences estimand being biased, as shown by [Borusyak and Jaravel \(2018\)](#); [Goodman-Bacon \(2018\)](#); [Sun and Abraham \(2020\)](#); [De Chaisemartin and D’Haultfœuille \(2020\)](#).

IV. EVIDENCE ON THE MECHANISM: COMPARATIVE ANALYSES

IV.A. *Heterogeneity by censorship of the internet and of traditional media*

The fact that uncensored internet can significantly undermine government popularity has not gone unnoticed by politicians, especially in nondemocratic countries. According to Freedom House, many governments have taken steps to limit internet freedom, with policies ranging from the blocking of social media and messaging apps in China, Egypt, Iran, and Russia to temporary shutdowns of mobile networks in India and Sri Lanka.²⁹ Yet, observers do conjecture that the internet is harder to censor than the traditional media (e.g., [Diamond and Plattner, 2012](#)).

In this section, we study whether and how the effect of 3G-network availability on individuals' attitudes toward government depends on internet censorship and on the censorship of the traditional media, (TV, radio, magazines, and newspapers).

We start by analyzing the heterogeneity of the effect of mobile broadband internet with respect to censorship of the internet. First, we add the interaction term between 3G coverage and a dummy for contemporaneous internet censorship, controlling for the direct effect of internet censorship, to our baseline difference-in-differences Specification (1). Panel A of Table V presents the results. The coefficients on 3G, indicating the effect of 3G without internet censorship, are negative and statistically significant, whereas the coefficients on the interaction term of 3G coverage with internet censorship, indicating the difference between the effects with and without internet censorship, are positive, significant for five of the six outcomes, and of similar magnitude in absolute value to the direct effect of 3G.³⁰

As internet censorship is often introduced to prevent messages critical of the government from reaching the public, it is reasonable to assume that censorship is more likely when government approval is low. In that case, one should worry about a bias arising from reverse causality in this estimation. In Appendix Section F, we derive the formula for the probability limit of the estimator of the coefficient on the interaction term between 3G and internet censorship and show that it is biased downwards (toward zero) and against finding an effect. Thus, with the contemporaneous censorship dummy, we can interpret the *sign* of the effect as causal, but we are likely to underestimate the magnitude.

Panel B of Table V addresses the potential issue with reverse causality by using a time-invariant dummy for countries with internet censorship. In this estimation, we do not allow reverse causality by construction, but we introduce measurement error, as

²⁹See <https://freedomhouse.org/report/freedom-net/freedom-net-2018> (accessed September 7, 2019). For academic work on internet censorship, see, e.g., [King, Pan, and Roberts \(2013, 2014\)](#), [Qin, Stromberg, and Wu \(2017\)](#), [Roberts \(2018\)](#), and [Chen and Yang \(2019\)](#).

³⁰The coefficients on the direct effect of internet censorship are positive and marginally significant.

internet censorship evolves over time. The results are very similar whether we use the time-variant or the time-invariant measure. Thus, we conclude that internet censorship weakens the effect of 3G on government approval. When the internet is free, 3G coverage has a strong and statistically significant negative effect on government approval. In contrast, with internet censorship, the impact of 3G coverage on government approval is zero.³¹

Figure V illustrates these results. Panel A presents the nonparametric relationships between government approval in a region (net of all controls) and 3G expansion in this region since 2008, separately for the two groups of countries: with free internet and with censored internet, according to the time-invariant measure. The figure shows that in countries with low internet censorship (left-hand-side graph), 3G expansion is associated with lower government approval, while in countries where the internet is censored (right-hand-side graph), there is no relationship between these variables.

In Panel B of Figure V, we present the nonparametric relationships between the increase in 3G coverage since 2008 and internet access at home in the two groups of countries. Whether the internet is censored or not, the presence of 3G networks facilitates internet access at home for the population. This suggests that the difference in the effect of 3G on government approval between countries with free and with censored internet comes from the content available online rather than from internet penetration.³²

Censoring the internet is technically difficult, due to its decentralized nature. Only a few governments restrict online content, whereas censorship of the traditional media is common throughout nondemocratic regimes. All countries with internet censorship in our sample have above-median censorship of the traditional press. In Panels C and D of Table V, we explore how the effect of 3G on government approval depends on the government’s control of the traditional media. We include the interactions of 3G coverage with dummies for both internet censorship and censorship of the traditional media. As above, we use both the contemporaneous and the time-invariant measures (in Panels C and D, respectively). We define the time-variant dummy for press censor-

³¹In the Appendix, we document that these results are not driven by the choice of functional form, the threshold for defining the internet censorship dummy, or the fact that we imputed zero censorship values for democracies. In Panels A and B of Table A.12, we replicate the results of Panels A and B of Table V in the subsample of countries with nonmissing internet censorship (FOTN) data: if anything, the effects are stronger without the imputation. Panels C and D of Table A.12 show that the results are also robust to using the continuous measures of internet censorship instead of the dummies. Panel A of Figure A.13 shows that the results are robust to using alternative thresholds for the definitions of the internet censorship dummies. Panel B of this figure reports the distributions of the underlying continuous measures of internet censorship and shows that the baseline thresholds are chosen to reflect natural breaks in these distributions.

³²Figure A.14 in the Appendix presents the corresponding nonparametric relationships, in which all controls are partialled out from the explanatory variable in addition to the dependent variable.

ship as an indicator that the FOTP index in that country and year is above the median value of this index among all countries in the sample without internet censorship, and the time-invariant measure is an indicator that the over-time mean of the FOTP index in this country is above the same median value. The coefficients on the interaction terms between 3G and the internet-censorship dummy remain positive and statistically significant in this specification, whereas the coefficients on the interaction of 3G with the dummy for above-median censorship of the traditional media are always negative and significant (for all but one outcome).³³ The coefficients on 3G coverage are also always negative, but are significant only for the aggregate measures of government approval.

The results are the same whether we use the time-variant or the time-invariant measures, which is particularly important in the case of censorship of the traditional media, because reverse causality could potentially bias the coefficient on the interaction term between 3G and censorship of the traditional media downward (away from zero), in favor of finding a negative effect. The specification with time-invariant measures of censorship is not subject to this reverse-causality problem. We illustrate the heterogeneity with respect to censorship of the traditional media in Appendix Figure A.15: focusing on countries with uncensored internet, the figure shows that the relationship between regional 3G expansion since 2008 and government approval (net of controls) is steeper in countries with above-median censorship of the traditional media compared to countries with below-median censorship of the traditional media.

Table A.13 in the Appendix replicates the entire Table V for the subsample of rural residents: the results are similar to those presented in Table V.³⁴

Overall, we conclude that, with internet censorship, 3G does not affect government approval. Without internet censorship, the effect of 3G coverage on government approval is, on average, negative. The effect is stronger (more negative) when the traditional media are controlled. This evidence suggests that uncensored internet plays a particularly important role in informing the public about politics, when the traditional media do not report independent-of-the-government political information.

³³Panels E and F of Table A.12 in the Appendix show that these results are robust to using the continuous measures of censorship of the internet and of the traditional media instead of dummies.

³⁴One potential concern with the interpretation of the results about the difference in the differential effects by the censorship of the traditional media versus censorship of the internet is the potential unobserved heterogeneity between those autocratic governments that control the traditional media but not the internet and those that censor both. In particular, if the latter are more sophisticated, our results on the heterogeneity by censorship may be driven by the heterogeneity with respect to the government's sophistication. In Appendix Section G, we show that there is no correlation between the censorship-of-the-internet score and any available measure of the level of education of the political elite and their prior occupations from Gerring et al. (2019). If the sophistication of the political leadership is related to education and occupations, it is not driving our results. In the Appendix, we list countries with internet censorship.

IV.B. Is the effect of 3G always negative? Heterogeneity by country-level corruption

In theory, if the expansion of mobile broadband internet provides the public with new information about the integrity and competence of the government, the sign of the effect of 3G on government approval should depend on the relationship between the public's prior beliefs and the content received online. The expansion of 3G should decrease government approval if the new information provides a worse view of the government relative to the ex ante beliefs. However, for honest and competent governments, greater transparency may increase approval. If the new information delivers a better view of the government compared to ex ante expectations, the Bayesian public should update the assessment of the government upward. This may be the case even if online platforms disseminate predominantly negative information. For example, if social and other online media expose more damning information about governments of other countries than about one's own government, government approval may increase.

If there is no systematic bias in the information received via 3G and in the ex ante beliefs, then the negative updates by the Bayesian public should on average be balanced by the positive updates. Our results in Section III, however, indicate a statistically significant negative average impact of 3G expansion on government approval. There are two potential explanations. First, if social media is more conducive to disseminating negative information about the status quo no matter how good the government actually is and the public is unaware of this asymmetry (e.g., [Castells, 2019](#); [Haidt and Rose-Stockwell, 2019](#)), one should expect the average effect of 3G expansion on government approval to be negative. Second, if the public's ex ante views are biased upward, for example, because the mainstream media controlled by the elites overstated the benefits of the status quo before the arrival of social media (as argued by [Gurri, 2018](#)), an increased transparency due to 3G expansion, on average, should also result in a downward shift in government approval.

To explore the heterogeneity of the impact of 3G by the actual quality of government, we use a cross-country measure of corruption constructed by IMF economists that is not based on perceptions ([Furceri, Papageorgiou, and Ahir, 2019](#)): their Global Incidents of Corruption Index (GICI) quantifies the actual corruption incidents in each country and year by measuring the share of the text of the quarterly EIU country reports devoted to corruption. In the next section, we explore within-country variation in this index over time. In this section, we use the long-term mean of this index as a measure of the overall level of corruptness in the country to understand whether and how the sign of the relationship between 3G and government approval differs in countries with high and low overall actual corruption.

We measure the overall level of the country's corruption as the country's mean of

the GICI from 2000 to 2017.³⁵ There are 104 countries in our sample with GICI data. We divide them into 13 equal-size groups of eight countries according to their rank in the overall level of corruption.³⁶ We put the remaining 12 countries with missing GICI data into a separate group, denoted by “M” for missing. Then, we estimate our baseline Specification (1), but allowing the coefficient on regional 3G coverage to vary depending on which group the country is in. The results are presented in Figure VI for each of our outcome variables. The figure presents the point estimates of the coefficients on regional 3G coverage for each group, along with their 90%-confidence intervals. Even though the estimates in some of the subgroups are rather noisy, the overall picture is clear: with the exception of the least corrupt countries, the expansion of mobile broadband internet has a negative effect on government approval regardless of how corrupt the country is. In contrast, in the first group of the least corrupt countries, which consists of Denmark, Germany, Japan, the Netherlands, Norway, Sweden, Switzerland, and the United Kingdom, 3G expansion led to an increase in government approval.³⁷

To ensure that the positive effect of 3G on government approval in the countries with the lowest overall corruption is not a result of pure chance, we conduct a set of 500 placebo estimations, in which we rank countries with nonmissing GICI data randomly, rather than according to the GICI, and we estimate the same specification as in Figure VI. The distribution of the t -statistics of the coefficients on the placebo group for the least corrupt countries from these regressions is presented in Appendix Figure A.16. It shows that it is extremely unlikely that the result about the effect of 3G in the countries with the least corrupt governments is just a random realization.

The results of the two heterogeneity exercises presented above—with respect to censorship and with respect to overall corruption—are consistent with the hypothesis that the consumption of political information available online is an important channel behind the political effect of 3G. However, these results provide no details on the content of such political information, in particular, whether voters are getting access to accurate political information or to false news, which—as has been shown in a number

³⁵This cross-country measure is highly correlated with the various measures of quality of governance from the Worldwide Governance Indicators (WGI), which—unlike the GICI—are based on perceptions.

³⁶We chose 13 groups to yield an equal number of countries in each group.

³⁷The composition of this least-corrupt group is consistent with [Castells \(2019, p.18\)](#) who argues that “corruption is a systemic trait of contemporary politics ... with a handful of exceptions, such as Switzerland or Scandinavia (excluding Iceland).” Our results are robust to expanding the group of least corrupt countries to include Finland (the 9th least corrupt country in the world according to the index of overall corruption based on the mean of the GICI), New Zealand (10th), Belgium (11th), Portugal (12th), and Singapore (13th). The effect of 3G in the least-corrupt group becomes zero if one includes the United States, the world’s 14th least corrupt country according to this index. (This is consistent with [Gurri \(2018\)](#), who provides extensive anecdotal evidence on how the expansion of social media has undermined the confidence in the U.S. government.) The positive and significant effect of 3G expansion in the least corrupt countries is robust to estimating the effect separately for this group of countries, instead of a full-sample estimation that allows the effect to vary by subgroup.

of studies, for example, [Allcott and Gentzkow \(2017\)](#); [Vosoughi, Roy, and Aral \(2018\)](#); [Grinberg et al. \(2019\)](#); [Guess, Nagler, and Tucker \(2019\)](#)—does get disseminated on social media. We address this question directly in the next section.

IV.C. Does mobile broadband internet help expose actual corruption?

If mobile broadband internet helps inform the public about actual corruption in government, incidents of actual corruption should translate into higher perceptions of corruption in subnational regions with greater access to mobile broadband internet. Thus, one should expect the link between actual and perceived corruption to be stronger in areas with higher 3G coverage. To test this, one needs to measure new incidents of actual corruption in a global setting. We use two alternative measures of actual corruption. The first is based on the analysis conducted by the Economist Intelligence Unit and aggregated into the GICI by [Furceri, Papageorgiou, and Ahir \(2019\)](#), the other is based on information from the Panama Papers, a trove of leaked documents about offshore entities.

New incidents of corruption measured with the GICI.—We consider the over-time within-country variation in the GICI as a measure of actual corruption incidents. To test whether mobile broadband internet helps expose corruption, we regress the dummy indicating whether the respondent believes that the government is *not* corrupt on the measure of actual corruption incidents (GICI) and its interaction with regional 3G coverage, controlling for the direct effect of 3G as well as all the baseline controls, including region and year fixed effects.

We find strong support for the hypothesis that the internet helps expose corruption. The results are reported in Table VI. The first two columns present the results for the full sample of countries, for which the GICI is defined, that is, including observations with zero actual corruption incidents. Columns 3 and 4 consider the subsample of country-years, in which the measure of actual corruption incidents is strictly positive, so that we rely on the variation in how much focus is given to corruption incidents in the EIU country reports, provided that corruption is among the topics covered by the report. In odd columns, we present the results for all the respondents; in even columns, for the respondents from rural areas.

The results are very similar, whether we consider all respondents or only respondents from rural areas and whether observations with zero corruption incidents are included. We find that the within-country correlation between actual corruption incidents and the perceptions of corruption is significantly higher in regions with higher 3G coverage. In regions with no 3G signal, the correlation between corruption incidents and perception that the government is not corrupt is negative but small in magnitude and is (marginally) significant only if one excludes observations with zero corruption

incidents (Columns 3 and 4). In contrast, if a region has full 3G coverage, there is a large, robust, and statistically significant link between the incidence of actual corruption and its perception. According to the baseline-sample estimates (Column 1), a one-standard-deviation increase in the measure of actual corruption incidents (0.31) is associated with a 2.9-percentage-point lower perception that the government is clean in places fully covered by 3G networks, and with a nonsignificant 0.4-percentage-point lower perception that the government is clean in places without mobile broadband internet coverage. (Overall, 21.5% of respondents believe that the government is clean.) In Panel A of Figure VII, we illustrate these results by presenting the marginal effect of an increase in the index of actual corruption incidents on the respondents' perceptions that the government is not corrupt for different levels of regional 3G coverage (implied by the estimates from Column 1): the effect becomes stronger (more negative) with the increase in 3G coverage.³⁸ The effect of 3G expansion when there are no corruption incidents, measured by the coefficient on regional 3G coverage in Columns 1 and 2, is small in magnitude and not statistically significant, suggesting that information about corruption available online is an important channel behind the negative effect of 3G.

Columns 5 and 6 of Table VI show that these average effects mask important heterogeneity by the country's overall level of corruption. We allow the effect of the interaction between 3G coverage and actual corruption incidents as well as the direct effect of 3G to vary between two groups of countries: with above- and below-median overall corruption, measured by the long-term mean of the GICI (used in the previous section). We find that the coefficient on the triple-interaction term between 3G coverage, actual corruption incidents, and a dummy for countries with below-median overall corruption is much larger in magnitude and more significant compared to a similar interaction with a dummy for countries with above-median overall corruption. This suggests that any particular corruption incident that gets exposed via mobile broadband internet contains bigger news in countries with relatively low overall corruption. At the same time, when the index of corruption incidents is zero, 3G expansion does not significantly affect corruption perceptions when overall corruption is relatively low.

In contrast, in countries with relatively high overall corruption, having access to mobile broadband internet is associated with significantly lower perceptions of no corruption in government, even when the index of corruption incidents is zero. This can be explained by the fact that, in these countries, many corruption incidents are not reflected in the EIU reports but are exposed with the help of mobile broadband internet. In this group of countries, an increase in the index of actual corruption incidents leads to

³⁸The results do not depend on the functional form of the measure of actual corruption incidents. In particular, the results are very similar if one uses $\log(GICI + 0.1)$ or $\log(GICI + 1)$ instead of using the raw GICI data.

a much smaller widening of the gap in corruption perceptions between regions covered and not covered by 3G than in countries with relatively low overall corruption. (The coefficients on the interaction terms between regional 3G coverage, a measure of actual corruption incidents, and a dummy for countries with above-median overall corruption are negative, small in magnitude, and significant only for rural residents.)³⁹

In Appendix Table A.15, we show that in the subsample of European countries, mobile broadband internet also helps inform the public about corruption incidents. These results help us interpret the findings on European elections, which we present in Section V. In Columns 1 and 2, we show that—similarly to the results for the global sample presented in Table VI—in Europe, the relationship between actual and perceived corruption is stronger in those subnational regions that are covered by 3G networks compared to the subnational regions without 3G coverage. This is true both for the full sample (Column 1) and for the sample focusing on the intensive-margin variation in actual corruption that excludes country-years with zero corruption incidents (Column 2). In Column 3, we verify that 3G expansion is associated with a significant increase in internet access at home among European respondents.

New incidents of corruption measured with the Panama Papers. On April 3, 2016, the Panama Papers, 11.5 million leaked documents detailing sensitive financial information of a large number of offshore entities, were made public. These documents directly implicated many corrupt government officials around the world in tax fraud and money laundering. Although offshore accounts are not *a priori* illegal and many private individuals use them, the revelations were particularly important in exposing corruption.⁴⁰ We base our second measure of new incidents of actual corruption on the number of unique offshore entities named in the Panama Papers.

First, we estimate a specification in which we regress the respondent’s perception that the government is not corrupt on the interaction between regional 3G coverage and the number of Panama Papers entities per 1,000 people in each country (i.e., we use the cross-country variation in the number of Panama Papers entities per capita, assuming that the actual corruption that gave rise to these offshore accounts can partially be observed by independent journalists and the opposition). In addition to our standard set of controls, to address the potential confounding factor that people in rich regions are more likely to know about offshore accounts than people in poor regions, we add the interaction of 3G with regional income to the set of covariates. The results are

³⁹In Appendix Table A.14, we test for a pretrend in actual corruption incidents; we find no evidence of such a pretrend. In particular, we show that regional 3G coverage is not predicted by contemporaneous or past levels of actual corruption incidents (Columns 1 and 2), and the index of actual corruption incidents is not predicted by lagged regional 3G coverage (Column 3).

⁴⁰See, e.g., the *New York Times* Editorial Board’s article from April 5, 2016: <https://www.nytimes.com/2016/04/06/opinion/the-PanamaPapers-sprawling-web-of-corruption.html> (accessed January 19, 2020).

reported in Column 1 of Table VII. The coefficient on the interaction between regional 3G coverage and the number of Panama Papers entities per 1,000 people is negative and significant. Thus, if the revelations from the Panama Papers are a measure of the level of overall corruption, this result confirms that mobile broadband internet helps expose it. To understand the magnitude of this effect, one can compare the difference in differences between the shares of people who believe that the government is corrupt in regions covered and not covered by 3G between two hypothetical countries, such that the number of Panama Papers entities per 1,000 people differs between these countries by one standard deviation. This difference in differences is equal to 5 percentage points. Panel B of Figure VII illustrates this result by presenting the magnitude of the marginal effect of an increase in the level of actual corruption measured by the Panama Papers on the belief that the government is not corrupt by different levels of regional 3G coverage (implied by the estimates presented in Column 1 of Table VII).

Next, we factor in the date when the Panama Papers were released to the public. In particular, we estimate specifications in which we allow the effect of the interaction between regional 3G coverage and the number of Panama Papers entities per 1,000 people to vary between two time periods: before and after the Panama Papers were released. We find that the effects are negative and significant both before and after the Panama Papers were released. The effect for the period after is larger than for the period before (presented in Column 2 of Table VII), but the difference in magnitude of these coefficients is not statistically significant.

The vast majority of the entities implicated by the Panama Papers come from middle-income and rich countries. Evidently, this is not because there is less corruption in poorer countries, but instead, because corrupt officials in these countries have no access to offshore bank accounts. In addition, in many low-income countries corruption is so pervasive that people observe it directly and do not need the internet to learn about it. Thus, we exclude low-income countries from the sample in Columns 3 to 6.⁴¹ As we show in Column 3, once low-income countries are excluded, the magnitude of the coefficient on the postrelease period becomes larger, and the difference in magnitude between the preperiod and postperiod effects becomes statistically significant (the p-values for this test are presented at the bottom of the table).

These results suggest that only part of the information contained in the Panama Papers was news to the public. Even though before the release of the Panama Papers the public did not know where corrupt officials hid their wealth, some information about the corruption of these officials was already available on the internet. For this reason, the effect of the interaction of 3G coverage with the number of Panama Papers entities

⁴¹We use the standard World Bank definition of low-income countries for 2015—the year before the Panama Papers revelations. The results are robust to alternative definitions of low-income countries.

is negative and statistically significant even before the leak. The difference between the coefficients from before and after the leak illustrates both the extent of surprise from the revelations of the Panama Papers and the fact that this new information was more likely to reach the public in regions covered by 3G networks.

In Column 4, we verify that these results do not rely on a linear functional form. In particular, instead of the number of Panama Papers entities per 1,000 people, we use a dummy indicating that this number exceeds 0.1, which corresponds to the top 10% of the distribution of Panama Papers entities per capita. In this specification, only the effect for the postperiod is statistically significant; the difference between the effects in the pre- and postperiods remains statistically significant.

The country ranking of the implication in the Panama Papers differs somewhat if one considers the total number of entities rather than the number of entities per capita. In particular, some large countries such as the United States and Russia have a big number of Panama Papers entities but a relatively small number of entities per capita. In Columns 5 and 6, we show that our results are robust to using the number of entities not divided by the size of country's population (factoring in that only elites have offshore accounts). Column 5 presents the results for the number of entities and Column 6 for a dummy indicating that this number is above 2000, which corresponds to the top 10% of countries in terms of the total number of Panama Papers entities per country. In all specifications, we find that the coefficients on the triple-interaction terms between regional 3G coverage, a measure of the country's exposure to the Panama Papers, and a dummy for the period after the Panama Papers were revealed are negative and significant. They are also significantly larger in magnitude than the corresponding effect for the preperiod.

To sum up, we find robust evidence that mobile broadband internet helps expose government corruption.

IV.D. Heterogeneity with respect to other country and individual characteristics, as well as placebo outcomes

We also interact regional 3G coverage with a number of other country-level and individual-level characteristics, focusing on the first principal component of government approval as the outcome variable.

Geography, income, and democracy. The first eight columns of Table A.16 report heterogeneity by continents, OECD membership, level of per capita income, and level of democracy. As above, we present the results for the full sample and for the subsample of rural residents. Columns 1 and 2 present the effect of 3G expansion separately for each continent. In the full sample, the effect is significant for Africa and the Americas and is not significant for Asia and Europe. The magnitude of the effect in European countries

in the full sample is essentially zero. In contrast, in the rural subsample, the effect is significant for all the continents, including Europe, where the effect is the smallest in magnitude among all continents (but is still sizeable). Columns 3 and 4 present the results separately for OECD and non-OECD countries. The effect is significant in non-OECD countries in both samples, while in OECD countries, it is significant only in the subsample of rural residents. Columns 5 and 6 show heterogeneity by the countries' per capita income. The results in high-income countries are virtually identical to those for OECD countries; whereas in middle-income and low-income countries, the effect of 3G coverage is significant both in the full sample and the rural subsample. It is the largest in magnitude in the group of upper-middle-income countries. Columns 7 and 8 document the absence of heterogeneity with respect to the level of democracy.

In the last two columns of Table A.16, we show that censorship of the internet and of the traditional media—considered in Section IV.A above—are the most important determinants of the effect of 3G coverage on government approval: qualitatively, the results on the heterogeneity by level of censorship do not change if we control for the interaction of regional 3G coverage with dummies for continents, levels of income, and levels of democracy.

Individual socioeconomic status. Table A.17 in the Appendix tests for heterogeneity with respect to the individual characteristics of the respondents. As above, we present the results for the full sample and for the subsample of rural residents. Columns 1 and 2 show that the effects are one-and-a-half times larger for the unemployed than for the employed. Columns 3 and 4 show that there is no effect of 3G on government approval among respondents with tertiary education, in sharp contrast to the negative and significant effects for respondents with secondary education and for respondents with less than secondary education, for whom the magnitude of the effect is the largest. Columns 5 and 6 show that the attitudes of respondents whose income is above the median country income in that year are less affected by 3G expansion than the attitudes of respondents with below-median income. Finally, Columns 7 and 8 report heterogeneity with respect to age groups. The results indicate that government approval among respondents who are younger than 25 is less affected by the expansion of mobile broadband internet than among respondents of other age groups. The effect on elderly people (above 60) is similar in magnitude to the effect on middle-aged people (between 25 and 60). The individual-level heterogeneity results are essentially the same for the total population and for the rural subsample, as can be seen from the comparison of the estimates presented in the odd and even columns of Table A.17.

Life satisfaction and other placebo outcomes. In Table A.18 in the Appendix, we show that 3G did not affect attitudes unrelated to the government. In particular, we show that 3G availability is not related to life satisfaction today, the expectation about

life satisfaction in five years, satisfaction with current standards of living, and beliefs about whether the standards of living are getting better. 3G coverage also has no effect on the confidence in the local police, suggesting that mobile broadband internet affects individuals' opinions about the government only for those government functions that people cannot observe directly through their day-to-day experience.

V. ELECTORAL IMPLICATIONS OF THE 3G EXPANSION

The results presented above suggest that mobile broadband internet is an important source of political information for voters. Does 3G expansion also have electoral implications? The evidence from previous literature (briefly discussed above) suggests that it does, but previous studies have addressed this question only in single-country settings. We use panel data on the election results in European democracies to examine the effects of the decade-long expansion of mobile broadband internet on the vote shares of incumbent and opposition political parties, including populist ones. We focus on Europe for two reasons. First, Europe has recently experienced a significant rise of populism (Rodrik, 2018); and we are particularly interested in whether the internet facilitates the electoral success of populist parties, as has been suggested by several observers (e.g., Gurri, 2018; Tufekci, 2018; Castells, 2019) and by previous research on Italy (e.g., Campante, Durante, and Sobbrío, 2018). Second, a conventional classification of political parties into populist and nonpopulist is not available outside Europe.

We use data on 102 parliamentary elections that took place between 2007 and 2018, covering 398 subnational districts in 33 European countries (EU-28 plus Liechtenstein, Montenegro, Northern Macedonia, Norway, and Switzerland), and we estimate regression equations analogous to Specification (1) but aggregated to the level of the subnational districts at which the election data are available. In all the specifications, we control for subnational-district and year fixed effects, as well as for a proxy for subnational district income (for which we use nighttime light density), and employ the following country-level controls: log GDP per capita, the rate of unemployment, inflation, labor-force participation, and the share of population that is 65 or older.⁴²

Our aim is to test whether the relationship between 3G expansion and a decline in government approval, which we have documented above, translates into tangible electoral losses for incumbent parties. The empirical challenge is that incumbent parties change over time. We address this challenge in two ways. First, we consider how electoral support for the parties that initially were part of the establishment evolved

⁴²We cannot use the IV strategy in the analysis of elections, because the frequency of lightning strikes does not have predictive power in the sample of European countries, as all of them are in the group of countries with above-median GDP per capita.

depending on the expansion of mobile broadband internet availability. For simplicity, we focus on the two largest parties in parliament from the first election during our observation period. The reason for considering two parties is that in most European democracies, the two top parties traditionally have rotated in and out of power. The advantage of this approach is that the parties that constitute the political establishment under this definition do not change over time, and we can measure their political support throughout the period.

As a more direct alternative, we consider the vote share for the ruling party, defined as the party of the country’s top executive (e.g., the Prime Minister). Because the ruling parties change over time, we first make a list of all political parties that were the ruling party at any point during our observation period. Next, we track the vote share for these parties, starting from the election in which they became the incumbent to the election in which they lost their incumbency. We then pool these observations. To compare vote shares within the same incumbent parties, in addition to all the baseline covariates, we control for incumbent-party-by-district fixed effects.⁴³

The results are presented in Columns 1 and 2 of Table VIII. In Column 1, the outcome is the vote share for the top two parties in the first observed election; in Column 2, it is the vote share for the incumbent party. No matter the specification, we find that 3G expansion reduces incumbents’ electoral support. We illustrate this relationship in Figure VIII. The point estimates imply that the expansion of mobile broadband networks in an average subnational district over a decade resulted in a 4.7-percentage-point lower vote share for the incumbent, both when the incumbents’ vote share is proxied by the vote share for the two top parties from the first election (the sample mean is 56%), and when it is measured as the vote share for the ruling party (with the sample mean of 30%).⁴⁴

In Column 3, we reestimate the specification presented in Column 2, allowing the effect of 3G to differ between populist and nonpopulist incumbents. We find that 3G expansion leads to a decrease in the incumbents’ vote share, whether or not the incumbent is populist. (There is no statistically significant difference between the coefficients on the interaction terms between district 3G coverage and dummies for populist and nonpopulist incumbents.) In Column 4, we confirm this result by showing

⁴³In the first approach, our unit of observation is a subnational district in an election. In the second approach, it’s an incumbent party in a subnational district in an election; namely, in those elections that led to a change of an incumbent party, there are two observations in each subnational district: one for the outgoing incumbent party and the other for the incoming incumbent party. In this specification, we control for incumbent-party-by-district fixed effects to account for geographic differences in political support for different parties. The results are the same in a less conservative specification that controls separately for district fixed effects and incumbent-party fixed effects.

⁴⁴This magnitude is based on the following calculation: $-4.7 = -0.089 \times 0.53 \times 100$, where 0.53 is an average increase in 3G coverage for subnational districts in Europe from 2008 to 2017, as discussed in Appendix Section C, and -0.089 is the coefficient on district 3G coverage.

that populist parties that were one of the top two parties in the beginning of the period lost votes as a result of 3G expansion.

In Column 5, we show that electoral turnout decreased more in districts that got higher 3G network coverage. This result could be driven by voters getting discouraged from participating in the elections due to their disillusionment with electoral institutions, consistent with our findings based on the GWP. It also could be that potential voters lose interest in politics as a result of exposure to online entertainment.⁴⁵ Appendix Table A.19 presents the results for the incumbent vote as a share of the number of registered voters rather than of those who actually voted in the election. The magnitudes are smaller but remain statistically significant. This implies that 3G expansion did spur some voters to change their political preferences.

Factoring in voter turnout, the estimates from Columns 1 and 5 of Table VIII imply the upper bound for the persuasion rate of the message “do not vote for the ruling party” of 27%. This upper bound is calculated under the assumption of no spillovers, that only the smartphone-owner (and nobody else) gets exposed to the message delivered by the device. In European countries, such spillovers are likely to be smaller compared to the rest of the world and, particularly, compared to poor countries; yet, one cannot rule out positive spillovers, even in the European context.⁴⁶

Taken together, these results strongly corroborate our findings on government approval from the GWP.⁴⁷ The expansion of 3G made voters more critical of their governments and resulted in worse electoral performance by the incumbents in Europe.

Which parties gain electoral support when incumbents lose it as a result of 3G expansion? In Columns 1 to 5 of Table IX, we consider the effect on the vote shares for populist and Green (environmentalist) parties. As the definitions of populist and Green parties do not change over time, the unit of observation is a subnational district in an election. First, we consider the populists’ vote shares and find that 3G expansion has contributed to stronger electoral performance by populist parties in Europe. A decade-long increase in subnational district 3G coverage, on average, results in a 4.6-percentage-point higher vote share for right-wing populists and a 3.6-percentage-point higher vote share for left-wing populists (Columns 1 and 2). The effects are large relative to the mean vote shares for right-wing and left-wing populists, equal to 13.6% and 6.5%, respectively. As we show in Column 3, there is no effect on parties classified as “other populists” (those that are not classified as right-wing or left-wing). Not all

⁴⁵Previous literature has found that political participation may increase or decrease with access to the internet depending on the setting; see a recent review of the literature by [Zhuravskaya, Petrova, and Enikolopov \(2020\)](#).

⁴⁶For details of this calculation and the assumptions behind it, see Appendix Section C.

⁴⁷As we show above, the expansion of 3G networks helped expose actual corruption incidents to European voters (see Appendix Table A.15) and led to a significant decline in government approval among European voters residing in rural areas (see Appendix Table A.16).

observers agree with the classification of populist parties into right-wing, left-wing, and other. In Column 4, we show that the results do not depend on this classification; the effects are large and statistically significant for all populists taken together. We find a 6.1-percentage-point increase in the vote share for all populists (from the mean of 26%) as a result of an average-size 3G expansion over the 2008–2017 decade.

During our observation period, populist parties were in power during some electoral terms in Bulgaria, Hungary, Italy, Montenegro, North Macedonia, Poland, Slovakia, and Slovenia. In Column 5, we exclude these countries from the sample and find a larger point estimate of the coefficient on district 3G coverage, as one would expect given that populist incumbents suffer electoral losses due to 3G expansion (see Column 3 of Table VIII).

Appendix Table A.20 reports these results with the vote share expressed as the share of registered voters. The point estimates of 3G’s effects on the populists’ vote (total, right-wing, and left-wing) are smaller in magnitude but remain statistically significant. The average region-level 3G expansion over the decade increased the electoral support for all populists as a share of registered voters by 2.5 percentage points (see Column 4).⁴⁸

The baseline estimates imply the upper bound for the persuasion rate—under the assumption of no spillovers—of 10.7% for the message “vote for a populist party” (for details on how we arrive at this figure, see Appendix Section C).

Does the nonpopulist opposition also gain from 3G expansion? Column 6 of Table IX shows that 3G network availability has a precisely-estimated zero impact on the vote share for Green parties. In Column 7, we consider all the nonpopulist opposition. We define a party to be in opposition if it is not included in the current ruling coalition. Similarly to the specifications presented in Columns 2 and–3 of Table VIII, this outcome is defined for each ruling coalition, and we control for the ruling-coalition-by-district fixed effects. We find no significant effect of 3G on the nonpopulist opposition’s vote share; point estimate is actually negative. Figure IX illustrates the results for the opposition parties’ vote share as the outcome variable.

In the Appendix, we establish robustness of these results to excluding any single country from the sample, as reported in Figure A.17. We also present the nonparametric relationships illustrating the election results with controls partialled out from the treatment variable as well as from the outcome variables. Figure A.18 shows the results for the incumbents’ vote share; Figure A.19 shows the results for the opposition.⁴⁹

⁴⁸The effect of 3G on the share of votes cast for populists classified as “other” becomes negative and significant, but as there are very few parties like this and there is a strong positive effect on both left-wing and right-wing populists, the overall effect for all populists remains positive and significant.

⁴⁹We also verify that the results are robust to excluding countries with compulsory voting: Belgium, Liechtenstein, and Luxembourg.

Overall, we find that, in European democracies, only populist opposition parties benefit from the disillusionment of voters with incumbent governments as a result of 3G expansion. If exposure to online criticism of incumbents were the only mechanism behind the fall in government approval with 3G expansion, one would expect all opposition parties to benefit from this phenomenon. Explaining why populists are the ones who gained from 3G expansion in Europe is beyond the scope of this paper. The mechanism could be both coincidental and causal. For instance, it is possible that the timing of 3G expansion coincided with the time when populist messages began to strongly resonate with voters, so that they just turned to the opposition that was the most appealing to them. However, it could also be that populists messages are particularly suited to the format of social media. In particular, populists’ rejection of existing democratic institutions as entrenched and serving the elites implies that they should talk directly to the voters bypassing the traditional media. Such direct contact on a large scale was made possible only with the arrival of social media. Populist messages may also be simpler, and thus, better suited for a short, catchy communication than messages of other opposition parties (see, e.g., [Levy, Razin, and Young, 2020](#)).⁵⁰

VI. COUNTRY CASE STUDIES

In this section, we briefly discuss three case studies illustrating the possible mechanisms behind our findings from the Gallup World Poll and from the European elections. Appendix Section H provides a detailed discussion of these case studies, backs them up with descriptive evidence, and lists the sources.

VI.A. *Russia 2017: YouTube video on Prime Minister’s corruption*

On March 2, 2017, a leading Russian opposition politician, Alexei Navalny, posted on YouTube a 50-minute documentary, *He Is Not Dimon to You* (or *Don’t Call Him “Dimon”*), detailing the corruption of Prime Minister Dmitry Medvedev. Because the Kremlin controls the traditional media, the documentary was not mentioned, let alone shown, on any of Russia’s TV channels; it could be viewed exclusively on YouTube. By the time Vladimir Putin removed Medvedev from the Prime Minister’s job, in January 2020, the film had 35 million views on YouTube.⁵¹

Within one month of the film’s release, Medvedev’s approval rating sank to a historic low and never recovered. It was unprecedented: in ten years, Medvedev’s

⁵⁰Consider, for example, the Greens’ narrative, which is substantially more complex than that of the populists. Greens call for voters to take responsibility for the planet, which requires costly policy choices. Populists, in contrast, apportion all the blame for economic and social problems to elites and foreigners, suggesting that they are the ones who should bear the costs of change.

⁵¹The population of Russia in 2019 was 146 million (Source: the UN’s [World Population Prospects 2019](#), accessed August 5, 2020.)

popularity had never before fallen by 10 percentage points in one month. There was no news related to Medvedev or his government that month apart from the release of the film.

According to a nationally representative survey, only two weeks after the release of the documentary, 4.5% of respondents had watched it and another 15.4% had heard about it. Both having watched the documentary and having heard about it is strongly positively correlated with respondents' self-reported internet use and with Medvedev's disapproval. One cannot establish causality in these relationships, because much of this correlation is driven by (unobserved) individual and location characteristics. It is noteworthy, however, that 9.6% of respondents who had never used the internet had some exposure to the film: 2.4% of respondents had someone else show them the film and 7.2% had heard about it from others. This indicates the importance of spillovers in the effect of the release of stories about government corruption on social media platforms: these stories reach not only those who are directly exposed but also those with whom the users of social media communicate.⁵²

VI.B. Romania 2014: the election of a "Facebook President"

In democratic countries (in contrast to Russia), exposing corruption online can also have electoral consequences. In the 2014 presidential election in Romania, incumbent Prime Minister Victor Ponta lost to a former physics teacher, Klaus Iohannis, who became known as the Facebook President. The margin of victory in the second round was 8.9 percentage points, which was a major and unexpected change in the Romanian political landscape—just two years earlier, Ponta won the parliamentary election with 59% of the vote. Romania was the second-most corrupt country in the EU at that point, and the stand on anti-corruption policies was the main cleavage in Romanian politics. Iohannis won on the anti-corruption ticket.

Iohannis attributed his success to his Facebook campaign. On the election night, he wrote the following post: *"Together, we have won the battle here on Facebook! . . . For the first time, the online has made a difference."* In the last two weeks of the campaign, Iohannis published eight Facebook posts per day, criticizing the status quo for corruption and emphasizing the need for change. During the campaign, Iohannis overtook Ponta in terms of followers and also strongly outperformed Ponta in terms of the number of comments, likes, and shares. A postelection survey reported that 54% of a representative sample of Romanian voters used the internet, and 93% of those internet users had a Facebook page; 70% of those respondents who used the internet

⁵²In the presence of such spillovers, the calculation of the persuasion rates for the political effect of 3G expansion (see Appendix Section C) overstates the true persuasion rates, because it assumes that only those who get a mobile broadband subscription are exposed to the anti-government message.

said that the internet and social media influenced their decision to vote.⁵³

VI.C. Brazil 2018: the election of a “WhatsApp President”

In addition to helping inform voters about misgovernance, mobile broadband internet and social media may provide a platform for disseminating misleading and outright false narratives, which can also have electoral implications.

During election campaigns in Brazil, free TV time slots are allotted to political parties based on their size and the seats in the legislature. Therefore, as an outsider in Brazilian politics, the right-wing populist candidate, Jair Bolsonaro, got virtually no access to television during the 2018 presidential election campaign. Thus, he campaigned almost exclusively online and mostly on WhatsApp, a digital social network used by 90% of Brazilian internet users. The high penetration of WhatsApp in Brazil is related to the popularity of so called “zero-rating” plans that offer free access to a limited number of social-media applications, including WhatsApp. Zero-rating plans are popular because they are affordable. Mobile subscriptions with unlimited broadband internet access are too expensive for most Brazilians.

WhatsApp, especially when accessed via zero-rating plans, is particularly well-suited for disseminating misinformation. WhatsApp messages are sent through encrypted chat groups of up to 256 members, which makes fact-checking hard for two reasons. First, WhatsApp messages are private and, therefore, not always available to fact-checkers. Second, zero-rating-plan users cannot access fact-checking information provided on non-WhatsApp platforms because of the limitations imposed by their zero-rating plan. Several sources, detailed in Appendix Section H, provide anecdotal evidence that WhatsApp was widely used to expose voters to false political narratives during the 2018 presidential election, much of which was carried out in a coordinated campaign by a network of Bolsonaro supporters. Bolsonaro won the election with 55.13% of the second-round vote.

To provide suggestive evidence of the importance of access to WhatsApp for the election results, we use cross-sectional geographic variation in mobile broadband network coverage in Brazil in 2018. We merge our 3G network availability data with the 2018 election results for Brazil’s microregions (*Microrregião*).⁵⁴ We find a strong positive correlation between a microregion’s 3G coverage and Bolsonaro’s vote share in the second round of the election (see Appendix Figure A.24). This correlation is

⁵³By 2014, Romania was almost fully covered by mobile broadband internet. Even outside Bucharest, the average share of population with access to 3G was 94%. At the time of the previous presidential election, in 2009, Romania’s 3G coverage was only 10%.

⁵⁴Across the 558 microregions, the mean 3G coverage was 28% (with the standard deviation of 24%); 22 microregions had no coverage, 166 had coverage below 10%, 104 had coverage above 50%, and 11 had 3G more than 90% coverage.

especially striking given that 3G coverage in urban areas—where the share of educated voters, who were more likely to vote against Bolsonaro—is higher than in rural areas. The slope of the correlation between 3G coverage and the electoral outcome is steep: in microregions with 3G coverage below 10%, the average vote for Bolsonaro was only 40.7%, whereas in the microregions with 3G coverage above 50%, his electoral support was 63.4%.

Overall, the three case studies illustrate how mobile broadband networks can affect government approval and lead to a fall in the incumbents' popularity.

VII. CONCLUSIONS

In this paper, we document the political effects of the expansion of mobile broadband internet throughout the world. Our analysis yields the following conclusions. The decade-long expansion of 3G networks that we studied has, on average, led to a significant reduction in government approval around the globe. However, there is substantial heterogeneity in this effect, depending on censorship of the internet, censorship of the traditional media, and overall corruptness. Government approval falls with 3G expansion only when there is no internet censorship. It is more negatively affected by the expansion of 3G networks if the traditional media are censored but the internet is not. Expansion of 3G decreases government approval if there is at least some corruption. In very few noncorrupt countries, the effect of 3G expansion on government approval is actually positive. Overall, mobile broadband internet is an important medium for providing voters with political information that is independent of the government. Part of this information is about actual corruption in government, part could be misinformation.

In Europe, the expansion of mobile broadband networks has had electoral implications. As 3G network coverage has increased, so has voters' discontent with their governments, leading to a decline in vote shares for the incumbent parties, a decrease in turnout, and electoral gains for populist parties, both on the right and on the left. On average, 3G expansion has not helped the nonpopulist opposition in Europe, including Green parties.

SCIENCES PO AND CENTRE FOR ECONOMIC POLICY RESEARCH

PRINCETON UNIVERSITY

PARIS SCHOOL OF ECONOMICS (EHESS) AND CENTRE FOR ECONOMIC POLICY RESEARCH

References

- Algan, Yann, Sergei Guriev, Elias Papaioannou, and Evgenia Passari, “The European Trust Crisis and the Rise of Populism,” *Brookings Papers on Economic Activity*, 2017 (2017) 309–400.
- Allcott, Hunt, and Matthew Gentzkow, “Social Media and Fake News in the 2016 Election,” *Journal of Economic Perspectives*, 31 (2017) 211–236.
- Altonji, Joseph G., Todd E. Elder, and Christopher R. Taber, “Selection on Observed and Unobserved Variables: Assessing the Effectiveness of Catholic Schools,” *Journal of Political Economy*, 113 (2005) 151–184.
- Andersen, Thomas Barnebeck, Jeanet Bentzen, Carl-Johan Dalgaard, and Pablo Selaya, “Lightning, IT Diffusion, and Economic Growth Across U.S. States,” *Review of Economics and Statistics*, 94 (2012) 903–924.
- Borusyak, Kirill, and Xavier Jaravel, “Revisiting Event Study Designs, with an Application to the Estimation of the Marginal Propensity to Consume,” Mimeo, Harvard University, (2018).
- Campante, Filipe, Ruben Durante, and Francesco Sobbrío, “Politics 2.0: The Multifaceted Effect of Broadband Internet on Political Participation,” *Journal of the European Economic Association*, 16 (2018) 1094–1136.
- Castells, Manuel, *Networks of Outrage and Hope: Social Movements in the Internet Age* (Malden, MA: Polity Press, 2015).
- , *Rupture: The Crisis of Liberal Democracy* (Medford, MA: Polity Press, 2019).
- Chen, Yuyu, and David Yang, “The Impact of Media Censorship: 1984 or Brave New World?” *American Economic Review*, 109 (2019) 2294–2332.
- Collela, Fabrizio, Rafael Lalive, Seyhun Orcan Sakalli, and Mathias Thoenig, “Inference with Arbitrary Clustering,” Mimeo, University of Lausanne, (2018).
- Conley, T.G., “GMM estimation with cross sectional dependence,” *Journal of Econometrics*, 92 (1999) 1–45.
- De Chaisemartin, Clément, and Xavier D’Haultfœuille, “Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects,” *American Economic Review*, 110 (2020) 2964–96.
- Diamond, Larry, and Marc F. Plattner, “Liberation Technology,” *Journal of Democracy*, 21 (2010) 69–83.
- Diamond, Larry, and Marc F. Plattner, (eds.) *Liberation Technology: Social Media and the Struggle for Democracy (A Journal of Democracy Book)* (Baltimore, MD: Johns Hopkins University Press, 2012).
- Donati, Dante, “Mobile Internet access and political outcomes: Evidence from South Africa,” Mimeo, Universitat Pompeu Fabra, (2019).

- Durante, Ruben, and Ekaterina Zhuravskaya, “Attack When the World Is Not Watching? US News and the Israeli-Palestinian Conflict,” *Journal of Political Economy*, 126 (2018) 1085–1133.
- Enikolopov, Ruben, Alexey Makarin, and Maria Petrova, “Social media and protest participation: Evidence from Russia,” *Econometrica*, 88 (2020) 1479–1514.
- Falck, Oliver, Robert Gold, and Stephan Heblich, “E-lections: Voting Behavior and the Internet,” *American Economic Review*, 104 (2014) 2238–65.
- Fergusson, Leopoldo, and Carlos Molina, “Facebook Causes Protests,” Mimeo, Universidad de los Andes, (2019).
- Furceri, Davide, Chris Papageorgiou, and Hites Ahir, “Global Incidents of Corruption Index,” Technical report, IMF (2019).
- Gavazza, Alessandro, Mattia Nardotto, and Tommaso Valletti, “Internet and Politics: Evidence from U.K. Local Elections and Local Government Policies,” *Review of Economic Studies*, 86 (2019) 2092–2135.
- Gerring, John, Erzen Oncel, Kevin Morrison, and Daniel Pemstein, “Who rules the world? A portrait of the global leadership class,” *Perspectives on Politics*, 17 (2019) 1079–1097.
- Goodman-Bacon, Andrew, “Differences-in-Differences with Variation in Treatment Timing,” *NBER Working Paper No. 25018*, (2018).
- Grinberg, Nir, Kenneth Joseph, Lisa Friedland, Briony Swire-Thompson, and David Lazer, “Fake News on Twitter during the 2016 U.S. Presidential Election,” *Science*, 363 (2019) 374–378.
- Guess, Andrew, Jonathan Nagler, and Joshua Tucker, “Less than You Think: Prevalence and Predictors of Fake News Dissemination on Facebook,” *Science Advances*, 5 (2019) eaau4586.
- Guriev, Sergei, and Elias Papaioannou, “The Political Economy of Populism,” *Journal of Economic Literature*, (forthcoming).
- Gurri, Martin, *The Revolt of the Public and the Crisis of Authority in the New Millennium* (San Francisco, CA: Stripe Press, 2018).
- Haidt, Jonathan, and Tobias Rose-Stockwell, “The Dark Psychology of Social Networks: Why it feels like everything is going haywire,” *The Atlantic*, (2019), <https://www.theatlantic.com/magazine/archive/2019/12/social-media-democracy/600763/> accessed on July 22, 2020.
- Hainmueller, Jens, “Entropy Balancing for Causal Effects: A Multivariate Reweighting Method to Produce Balanced Samples in Observational Studies,” *Political Analysis*, 20 (2012) 25–46.
- Henderson, Vernon, Adam Storeygard, and David Weil, “A Bright Idea for Measuring Economic Growth,” *American Economic Review*, 101 (2011) 194–199.

- , “Measuring Economic Growth from Outer Space,” *American Economic Review*, 102 (2012) 994–1028.
- Hsiang, Solomon, “Temperatures and cyclones strongly associated with economic production in the Caribbean and Central America,” *Proceedings of the National Academy of Sciences*, 107 (2010) 15367–15372.
- ITU, “Measuring digital development. Facts and figures 2019,” Technical report, ITU Publications, Geneva, Switzerland (2019).
- Kemp, Simon, “Digital, Social and Mobile in 2015: We Are Social’s Compendium of Global Digital Statistics,” Technical report, We Are Social (2015), <https://www.socialmediatoday.com/content/global-digital-social-media-stats-2015> (accessed on July 20, 2020).
- , “Digital in 2018: Essential Insights into Internet, Social media, Mobile and Ecommerce Use around the World,” Technical report, We Are Social and Hootsuite (2018), <https://wearesocial.com/blog/2018/01/global-digital-report-2018> (accessed on July 20, 2020).
- King, Gary, Jennifer Pan, and Margaret E. Roberts, “How Censorship in China Allows Government Criticism but Silences Collective Expression,” *American Political Science Review*, 107 (2013) 1–18.
- , “Reverse-Engineering Censorship in China: Randomized Experimentation and Participant Observation,” *Science*, 345 (2014) 1–10.
- Levy, Gilat, Ronny Razin, and Alwyn Young, “Misspecified Politics and the Recurrence of Populism,” Mimeo, London School of Economics, (2020).
- Louis-Sidois, Charles, and Elisa Mougin, “Silence the Media or the Story? Theory and Evidence of Media Capture,” Mimeo, Sciences Po, (2020).
- Manacorda, Marco, and Andrea Tesei, “Liberation Technology: Mobile Phones and Political Mobilization in Africa,” *Econometrica*, 88 (2020) 533–567.
- Martin, Al, “Effects of Lightning on ICT Circuits: Induction and GCR,” *In Compliance Magazine*, (2016), <https://incompliancemag.com/article/effects-of-lightning-on-ict-circuits-induction-and-gcr/> (accessed on July 20, 2020).
- Miner, Luke, “The unintended consequences of Internet diffusion: Evidence from Malaysia,” *Journal of Public Economics*, 132 (2015) 66–78.
- Mitchell, Amy, Jeffrey Gottfried, Sophia Fedeli, Galen Stocking, and Mason Walker, “Many Americans Say Made-Up News Is a Critical Problem That Needs To Be Fixed,” Technical report, Pew Research Center (2019).
- Morozov, Evgeny, *The Net Delusion: The Dark Side of Internet Freedom* (New York, NY: Public Affairs, 2011).
- Oster, Emily, “Unobservable Selection and Coefficient Stability: Theory and Evidence,” *Journal of Business & Economic Statistics*, 37 (2017) 187–204.

- Qin, Bei, David Stromberg, and Yanhui Wu, “Why Does China Allow Freer Social Media? Protests versus Surveillance and Propaganda,” *Journal of Economic Perspectives*, 31 (2017) 117–140.
- Rainie, Lee, and Barry Wellman, *Networked – The New Social Operating System* (Cambridge, MA: The MIT Press, 2012).
- Roberts, Margaret E., *Censored: Distraction and Diversion Inside China’s Great Firewall* (Princeton, NJ: Princeton University Press, 2018).
- Rodrik, Dani, “Populism and the Economics of Globalization,” *Journal of International Business Policy*, 1 (2018) 12–33.
- Roth, Jonathan, “Pre-test with Caution: Event-study Estimates After Testing for Parallel Trends,” Mimeo, Harvard University, (2019).
- Schaub, Max, and Davide Morisi, “Voter mobilization in the echo chamber: Broadband internet and the rise of populism in Europe,” *European Journal of Political Research*, 59 (2020) 752–773.
- Schulman, Aaron, and Neil Spring, “Pingin’ in the Rain,” *IMC ’11: Proceedings of the 2011 ACM SIGCOMM conference on Internet measurement conference*, (2011).
- Sun, Liyang, and Sarah Abraham, “Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects,” Mimeo, MIT, (2020).
- Tufekci, Zeynep, “How social media took us from Tahrir Square to Donald Trump,” Technical report, MIT Technology Review (2018).
- Vosoughi, Soroush, Deb Roy, and Sinan Aral, “The spread of true and false news online,” *Science*, 359 (2018) 1146–1151.
- Young, Alwyn, “Consistency without Inference: Instrumental Variables in Practical Application,” Mimeo, London School of Economics, (2020).
- Zeddani, Ahmed, and Phil Day, “Improving the protection of ICT equipment against lightning strikes,” Technical report, ITU News (2014), <https://news.itu.int/improving-protection-ict-equipment-lightning-strikes/> (accessed on July 20, 2020).
- Zhuravskaya, Ekaterina, Maria Petrova, and Ruben Enikolopov, “Political Effects of the Internet and Social Media,” *Annual Review of Economics*, 12 (2020) 415–438.

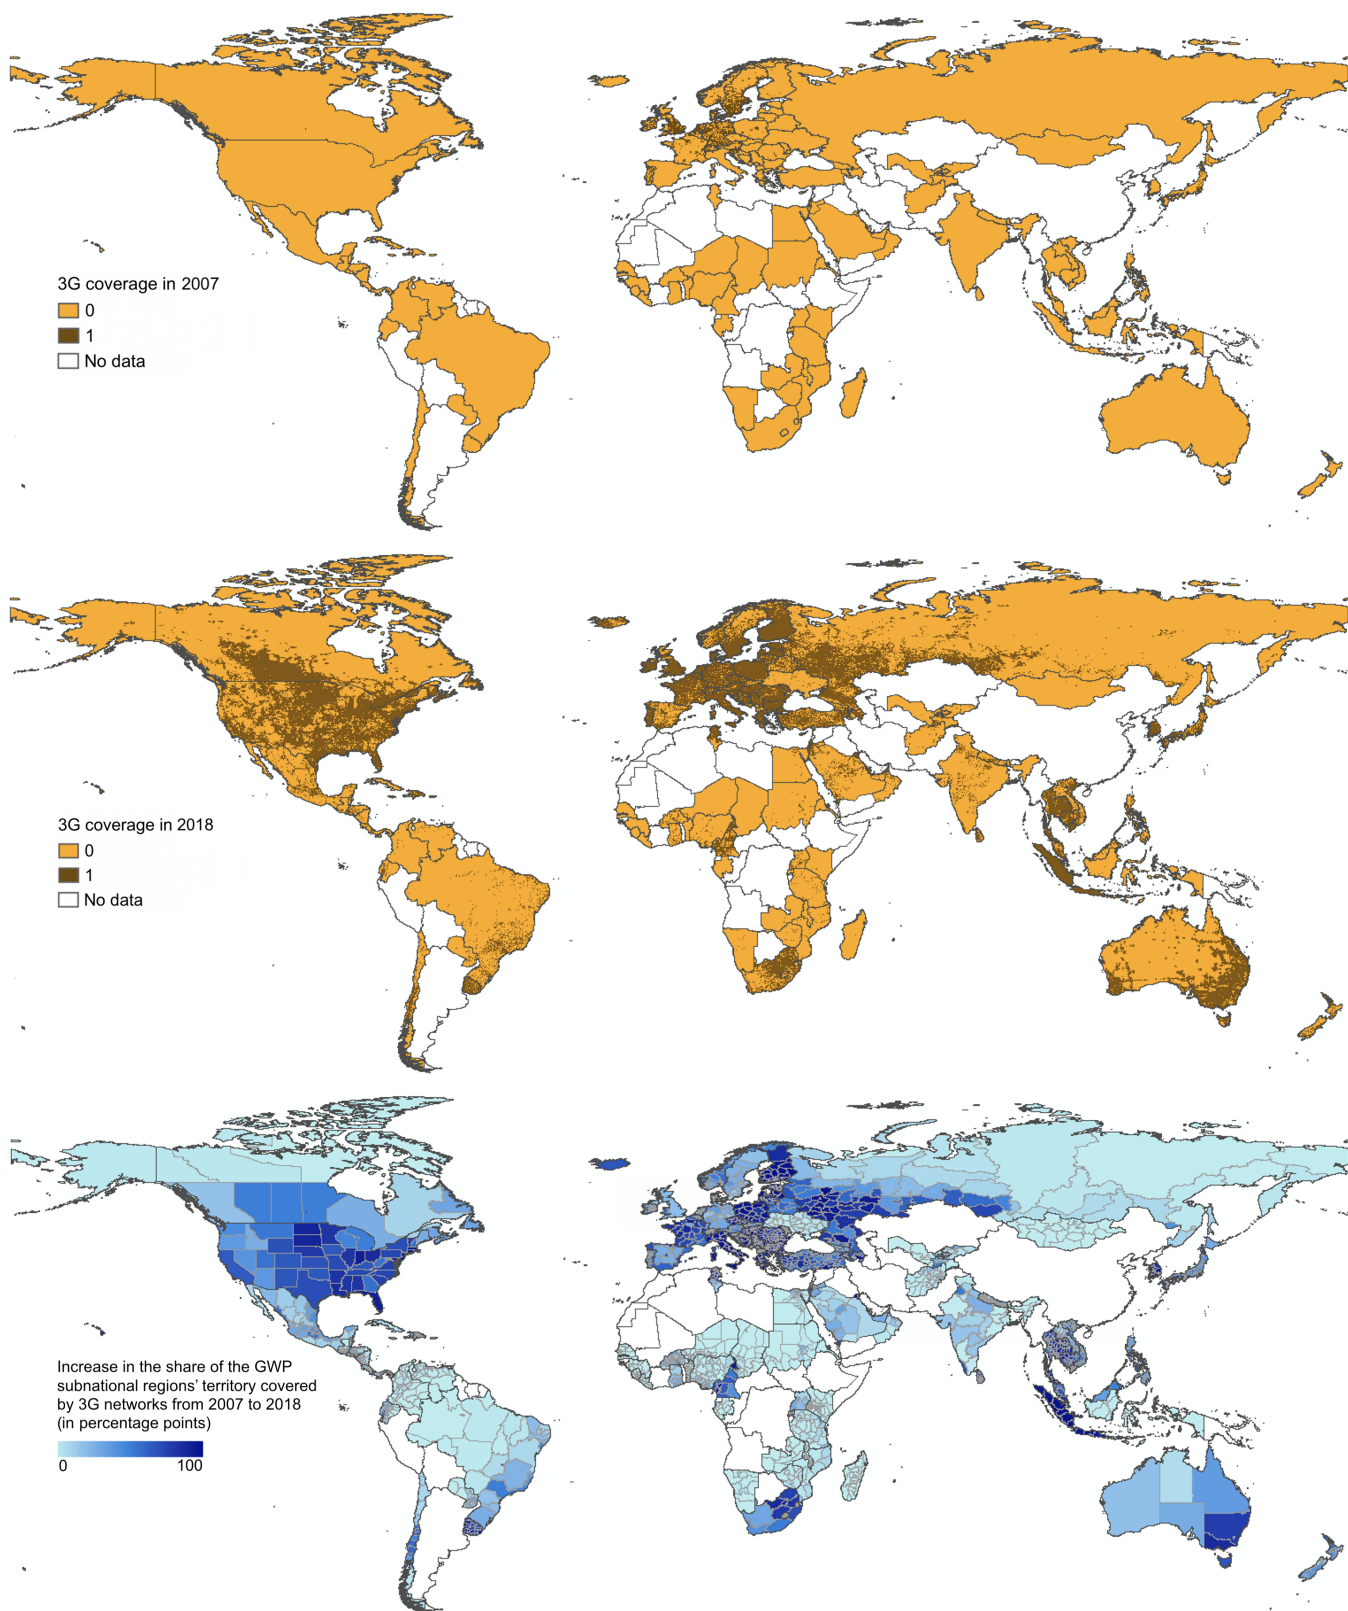


Figure I

The growth of 3G network coverage between 2007 and 2018

Note: The first two maps present 3G network coverage by grid cell in 2007 and 2018. The third map presents: (1) the boundaries of the subnational regions (the unit of localization in the GWP data) and (2) the increase in the share of each subnational region's population covered by 3G networks from 2007 to 2018. The sample consists of all countries covered by the GWP data. There are 2,232 subnational regions in the sample.

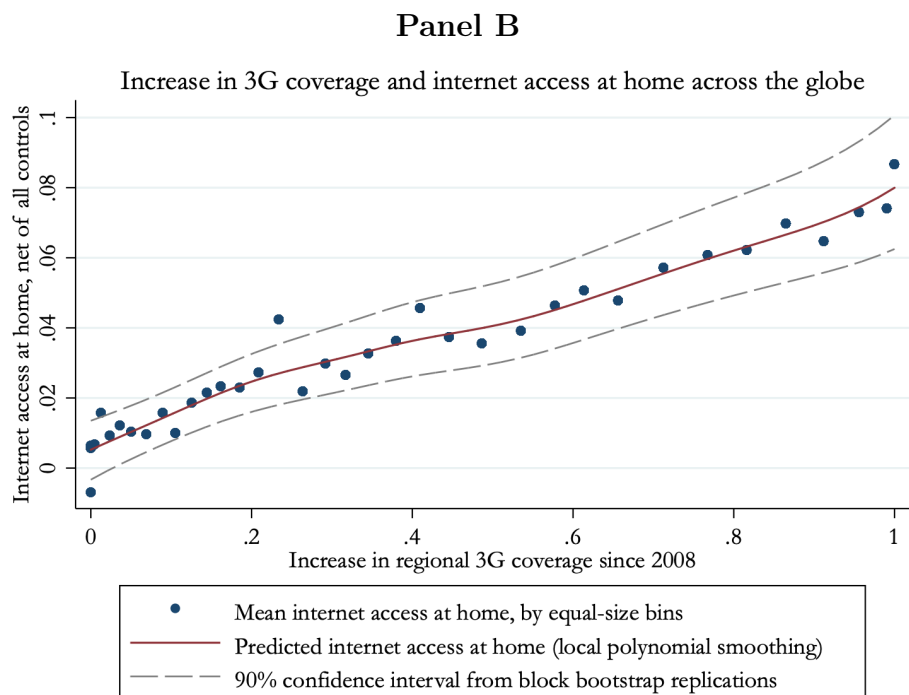
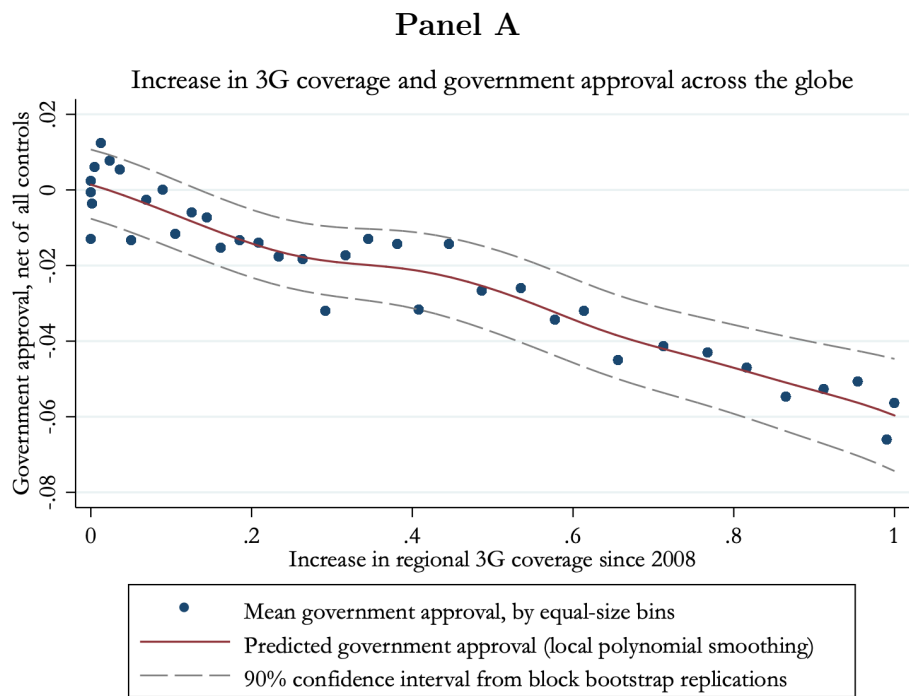


Figure II

Increase in 3G coverage and confidence in government

Note: Panel A of the figure illustrates the relationship between regional 3G coverage and government approval (Column 6 of Panel A of Table I). Panel B of the figure illustrates the relationship between the increase in regional 3G coverage and access to the internet at home (Column 1 of Panel A of Appendix Table A.2). The dots show the means of the respective outcome variables, net of all controls, by equal-size bins. The lines on the graphs show the predicted outcomes (Gaussian kernel, local polynomial smoothing). The confidence intervals are constructed by performing a block bootstrap at the level of the clusters.

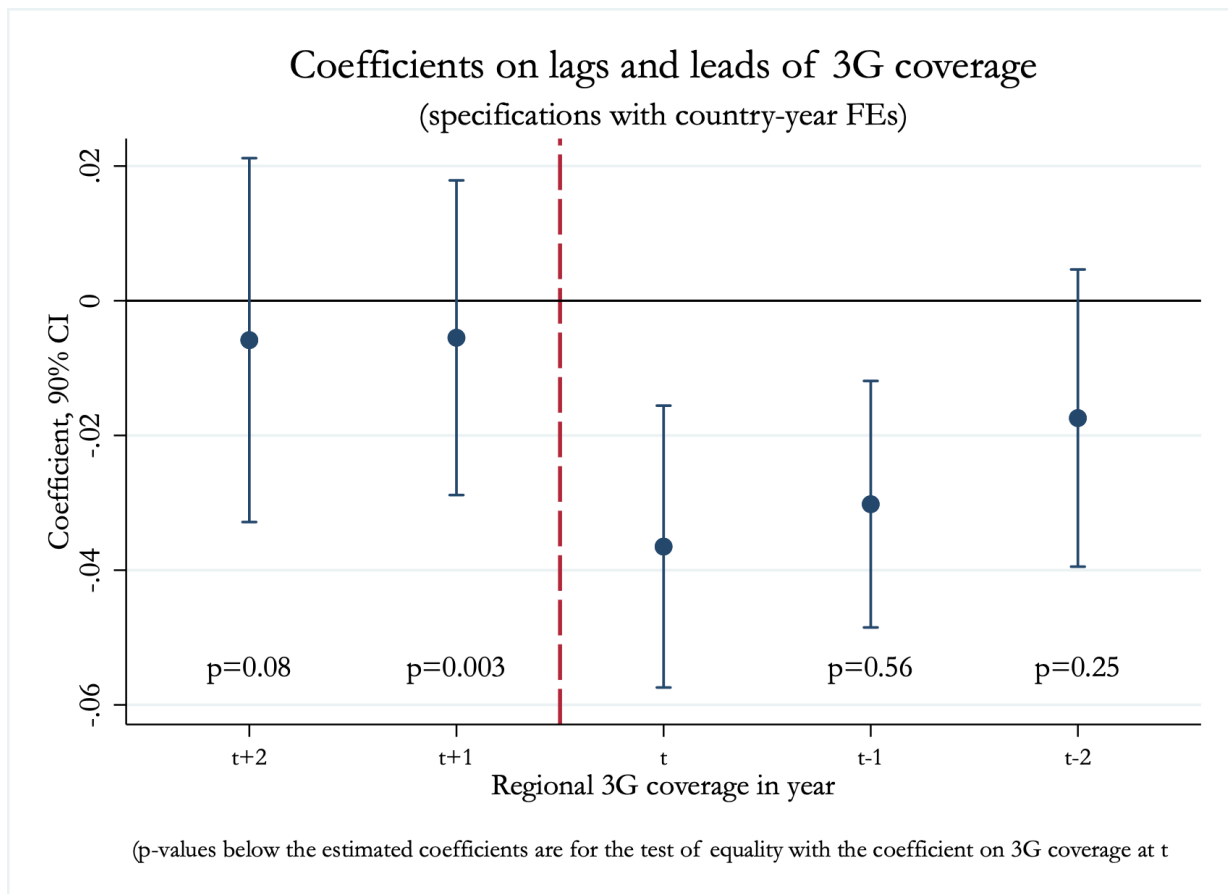


Figure III
Pretrend analysis with country \times year FEs

Note: The figure presents the coefficients from the regressions of government approval on the lags and leads of 3G coverage in the full sample, controlling for country-year fixed effects and all the baseline controls. Each coefficient is from a separate regression. The results suggest that future expansions of 3G networks are not associated with current changes in government approval. The p-values below the estimates are for the test of equality of magnitudes between the respective coefficient and the coefficient on regional 3G coverage at t . The coefficients on the leads of 3G coverage are significantly smaller in absolute value than on 3G coverage at t , confirming the parallel pretrends assumption required for identification.

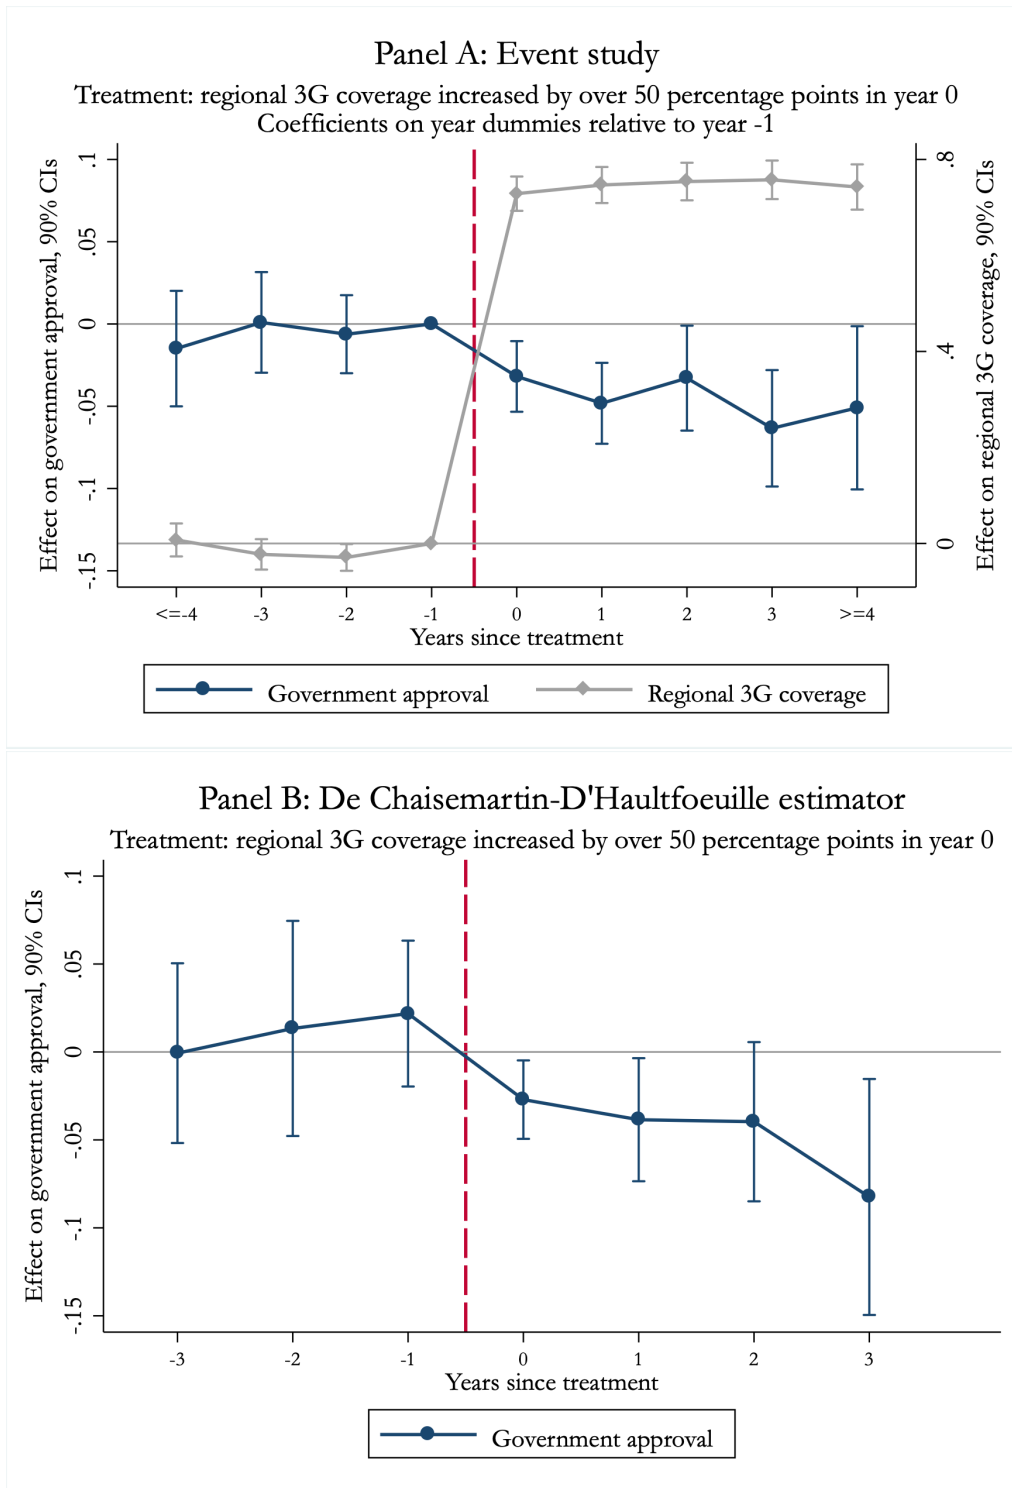
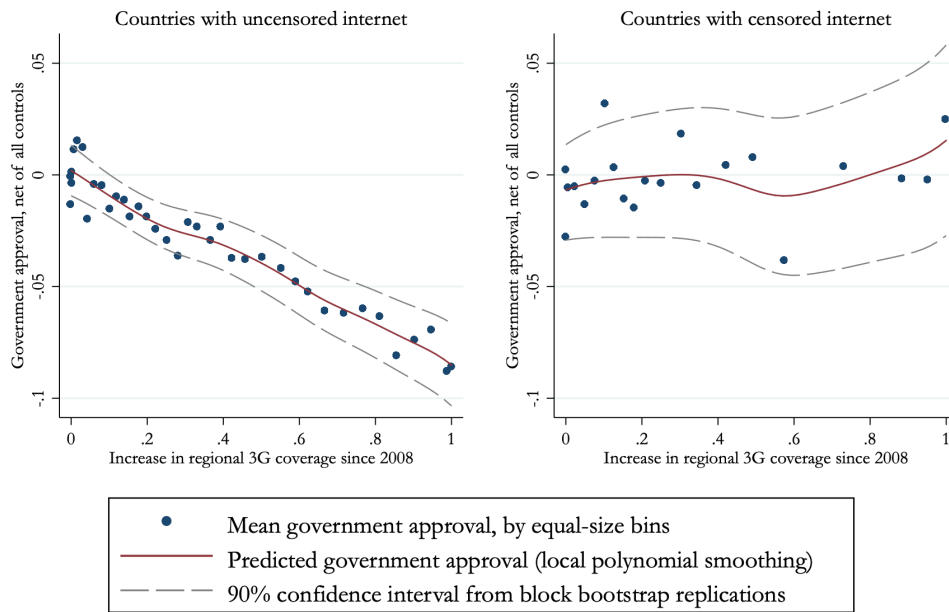


Figure IV
Event study analysis

Note: Panel A presents an event study in which government approval (left axis) and 3G coverage (right axis) are regressed on a set of year dummies around the event defined as an annual increase in regional 3G coverage of more than 50 percentage points. The regressions are run on the subsample of 452 regions in 65 countries where 3G did increase sharply in a single year between 2007 and 2018. The results of the underlying regression for government approval as outcome are presented in Column 3 of Table II. For each outcome variable, all the coefficients come from the same regression, which includes all the baseline controls and the freedom-of-the-press score in the list of covariates. Panel B presents the estimates based on the estimator proposed in [De Chaisemartin and D'Haultfoeuille \(2020\)](#), which ensures that the average treatment effects in each group and period do not have negative weights. Both panels of the figure show that the decrease in government approval occurred after the significant expansion of 3G networks.

Panel A

Increase in 3G coverage and government approval across the globe



Panel B

Increase in 3G coverage and internet access at home across the globe

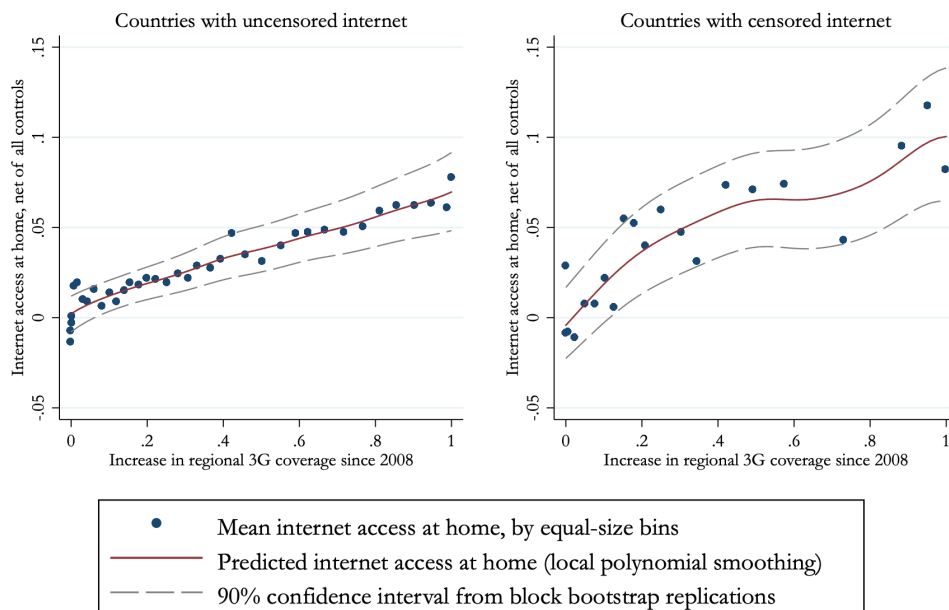


Figure V

Increase in 3G coverage and confidence in government, depending on internet censorship

Note: Panel A of the figure illustrates the results presented in Column 6 of Panel B of Table V, showing the relationship between the increase in regional 3G coverage and government approval separately for countries with and without internet censorship. Panel B of the figure illustrates the relationship between the increase in regional 3G coverage and access to the internet at home for countries with high and low levels of internet censorship. The dots show the means of the respective outcome variables, net of all controls, by equal-size bins. The lines on the graphs show the predicted outcomes (Gaussian kernel, local polynomial smoothing). The confidence intervals are constructed by performing a block bootstrap at the level of the clusters.

The effect of 3G by country's overall corruption level on:

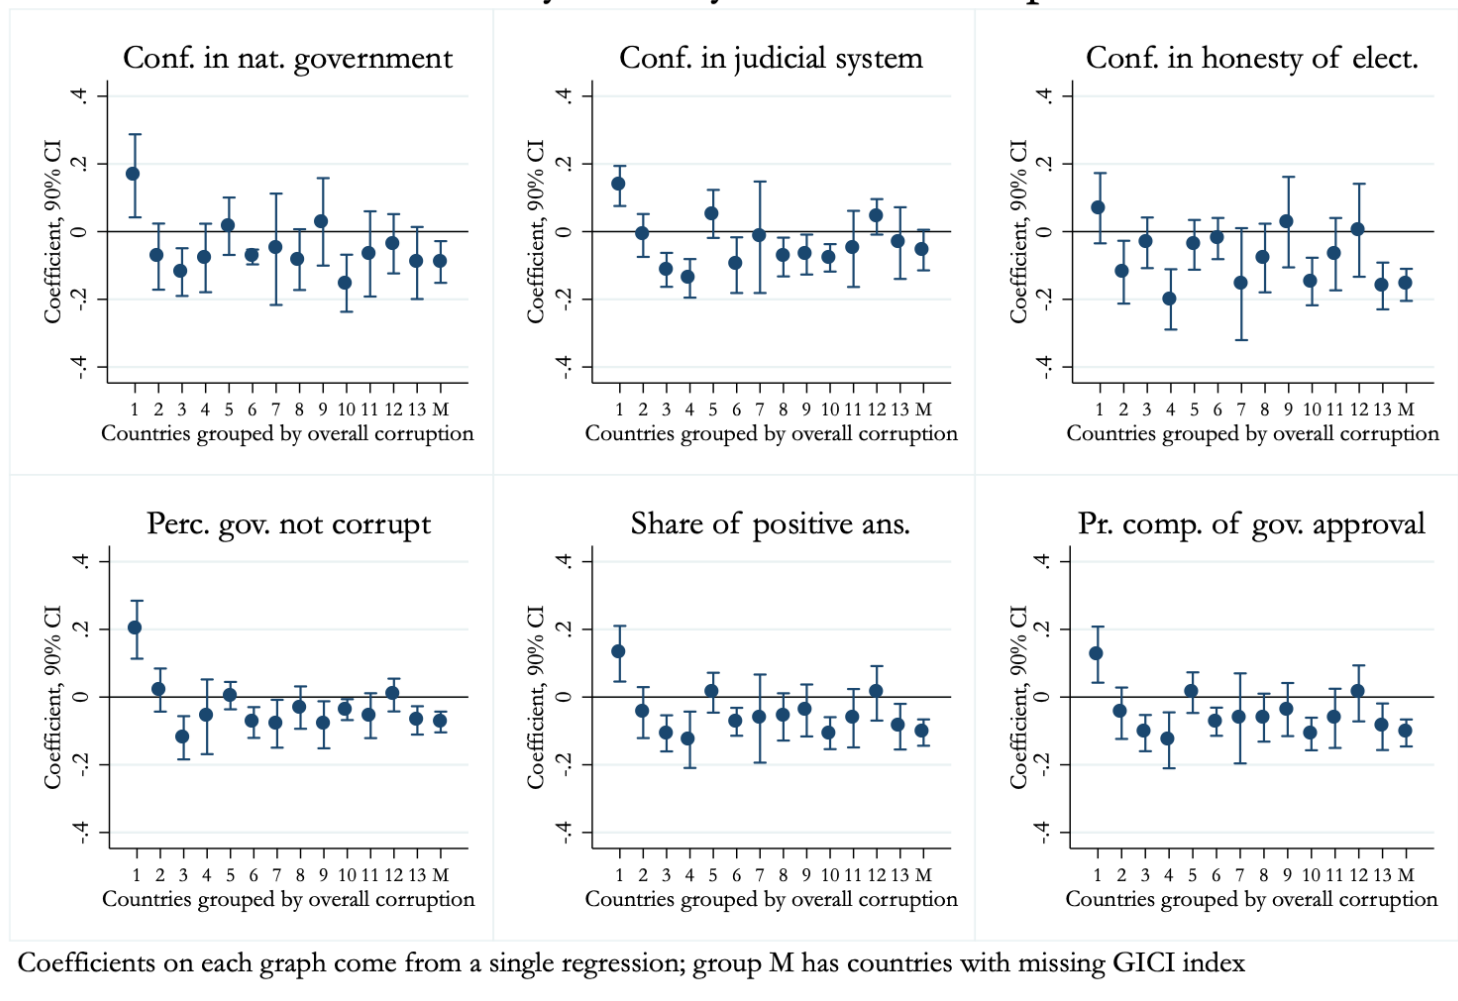
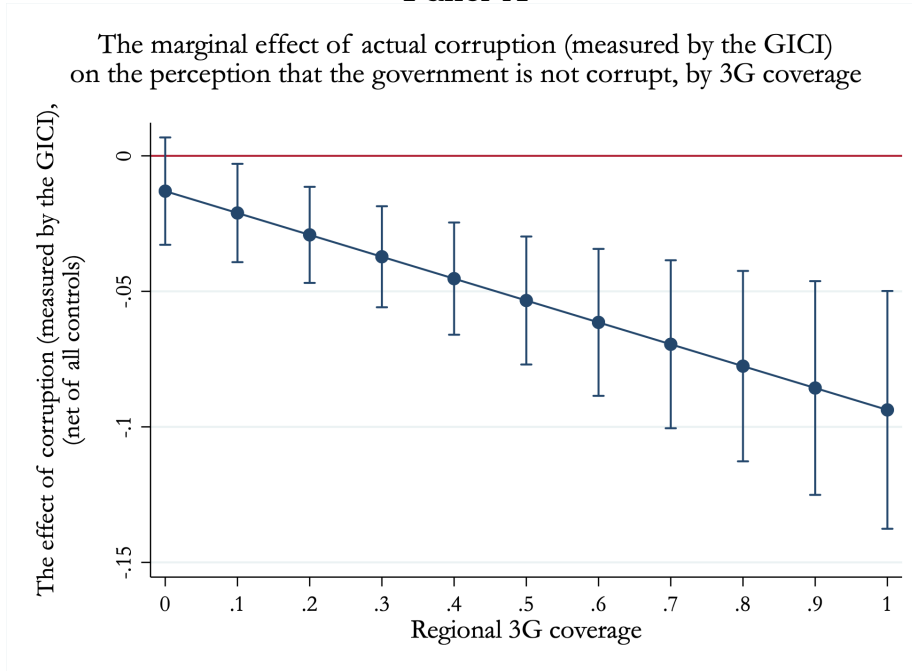


Figure VI

The effect of 3G coverage on government approval by overall corruption level

Note: The graphs present the coefficients on the interactions between regional 3G coverage and dummies for each of the 13 groups of countries, grouped by the overall level of corruption (i.e., mean GICI over 2000-2017) with 8 countries in each group. Group M has all 12 countries with missing GICI data. The graphs also present 90% confidence intervals, that are calculated from standard errors, corrected for two-way clusters at the subnational district level (to account for correlation over time) and at the level of the countries in each year (to account for within-country-year correlation).

Panel A



Panel B

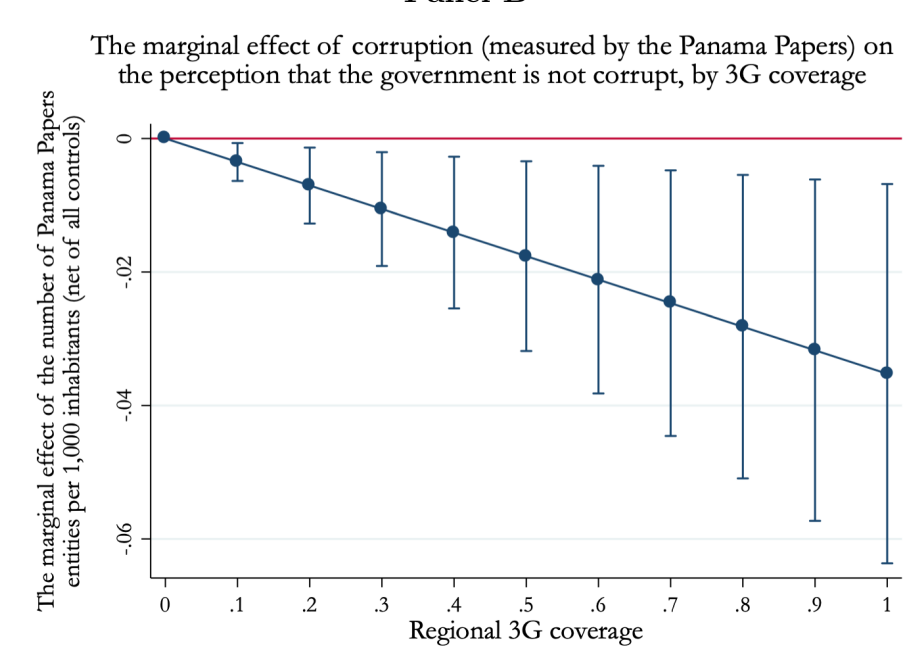


Figure VII

3G coverage and actual and perceived corruption

Note: The outcome variable is a dummy for the perception that there is no corruption in government. In Panel A, the explanatory variables are: regional 3G coverage, actual corruption incidents (GICI), their interaction term, as well as all the baseline controls, including region and year fixed effects (Column 1 of Table VI). In Panel B, the explanatory variables are: regional 3G coverage, the interaction term of regional 3G coverage and the number of entities in the Panama Papers per 1,000 people, the interaction of regional 3G coverage with regional income, as well as all the baseline controls, including region and year fixed effects (Column 1 of Table VII). The graphs present the marginal effects of an increase in actual corruption (measured by the GICI and the Panama Papers) on the perception of corruption. The graphs also present 95% confidence intervals, that are calculated from standard errors, corrected for two-way clusters at the level of the subnational districts (to account for correlation over time) and at the level of the countries in each year (to account for within-country-year correlation). The difference in the shape of the confidence intervals in the two graphs comes from the fact that the GICI varies both across countries and over time, whereas the Panama Papers provide information on countries at one point in time.

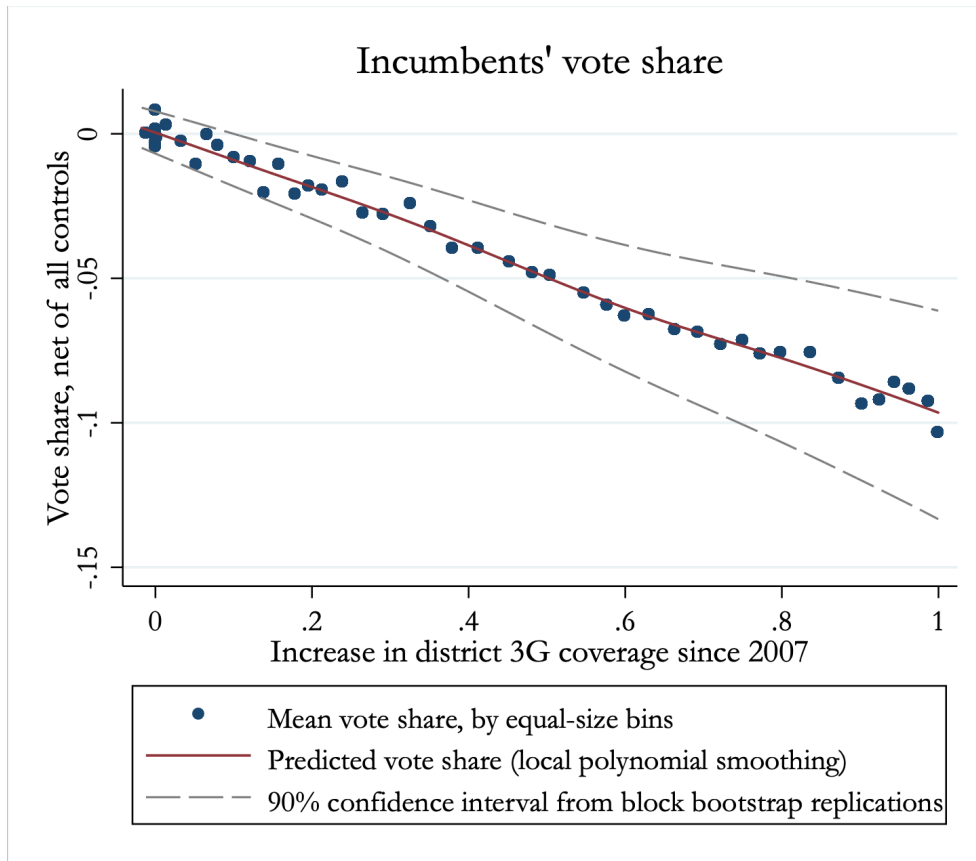


Figure VIII

Electoral implications of the expansion of 3G coverage for incumbents

Note: The figure illustrates the results presented in Column 2 of Table VIII. The dots represent the vote shares, net of all controls, by equal-size bins. The solid line on the graphs shows the predicted vote shares (Gaussian kernel, local polynomial smoothing). The 90% confidence intervals are constructed by performing a block bootstrap at the level of the clusters.

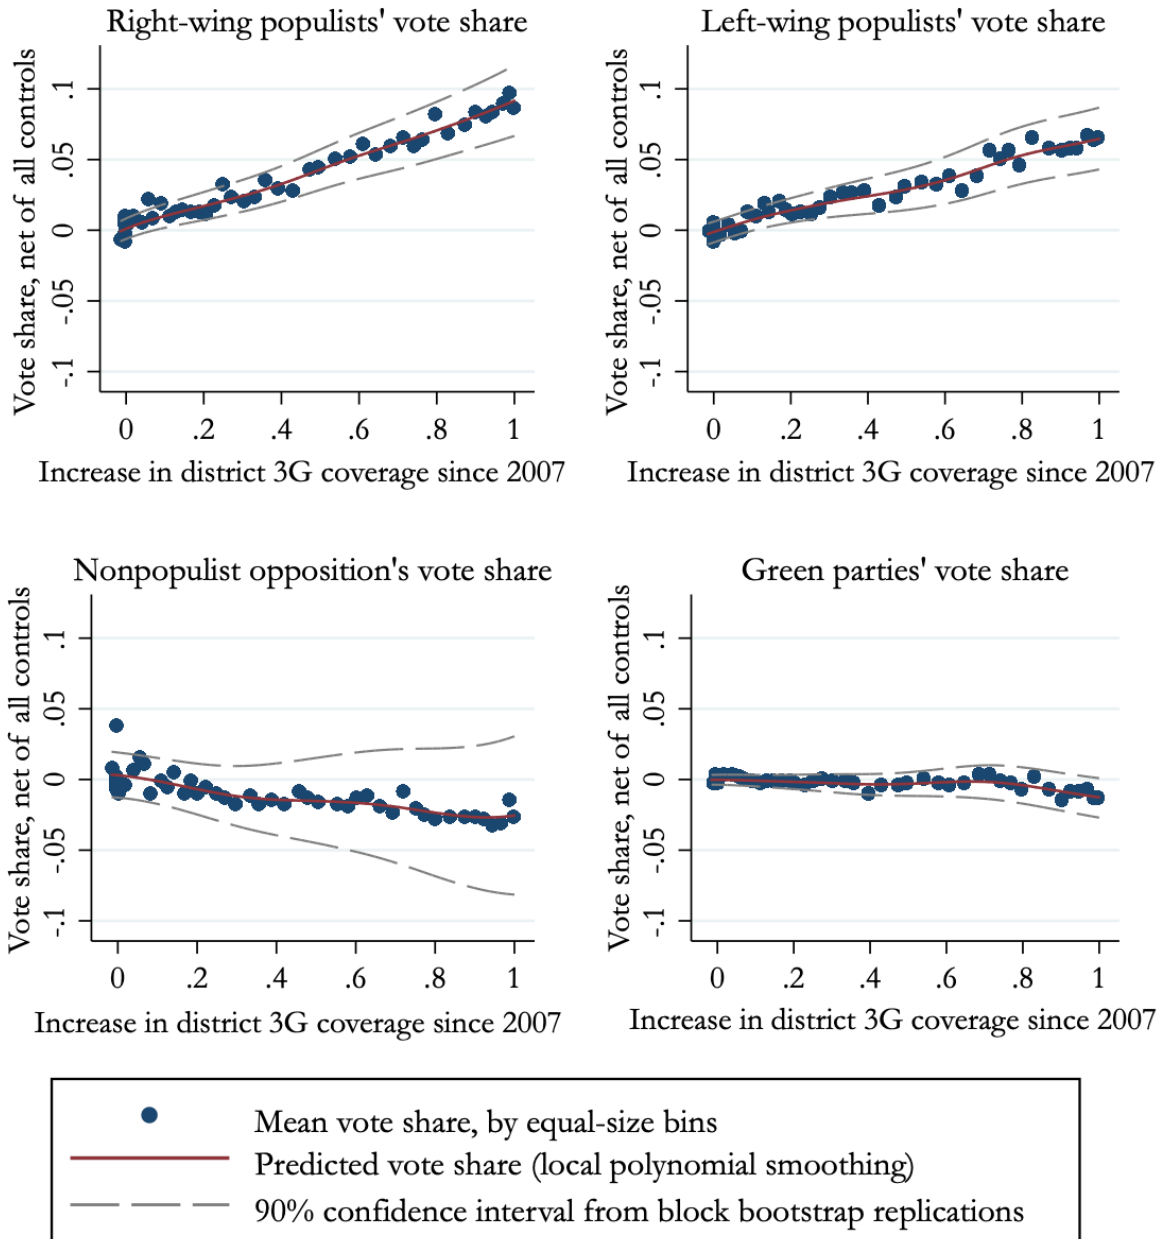


Figure IX

Electoral implications of the expansion of 3G coverage for opposition parties

Note: The plots on the first row illustrate the results presented in Columns 1 and 2 of Table IX. The plots on the second row illustrate the results presented in Columns 6 and 7 of Table IX. The dots represent the vote shares, net of all controls, by equal-size bins. The solid lines show the predicted vote shares (Gaussian kernel, local polynomial smoothing). The 90% confidence intervals are constructed by performing a block bootstrap at the level of the clusters.

Table I
The effect of mobile internet coverage on confidence in government

	(1)	(2)	(3)	(4)	(5)	(6)
<i>Dep. Var.:</i>	Confidence in national government	Confidence in judicial system	Honesty of elections	No corruption in government	Share of questions with positive responses	1st principal component of responses
Panel A: Sample of all respondents						
Regional 3G coverage	-0.063*** (0.021)	-0.040*** (0.015)	-0.079*** (0.021)	-0.036** (0.014)	-0.056*** (0.015)	-0.057*** (0.015)
Observations	772,353	748,471	732,856	722,768	617,863	617,863
R-squared	0.164	0.163	0.168	0.225	0.242	0.239
Mean dep. var	0.439	0.534	0.505	0.226	0.432	0.439
Mean 3G coverage	0.397	0.381	0.383	0.383	0.381	0.381
Number of countries	111	116	112	112	110	110
Panel B: Subsample of rural residents						
Regional 3G coverage	-0.091*** (0.024)	-0.058*** (0.017)	-0.115*** (0.026)	-0.054*** (0.016)	-0.080*** (0.018)	-0.081*** (0.018)
Observations	464,831	448,449	440,786	432,460	371,055	371,055
R-squared	0.171	0.157	0.161	0.194	0.224	0.222
Mean dep. var	0.349	0.556	0.516	0.215	0.445	0.452
Mean 3G coverage	0.329	0.314	0.316	0.316	0.311	0.311
Number of countries	110	115	111	111	109	109
Subnational region & year FEs	✓	✓	✓	✓	✓	✓
Baseline controls	✓	✓	✓	✓	✓	✓

Note: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. 3G internet reduces government approval. The unit of observation is an individual. Panel A reports the results for the full sample and Panel B for the subsample of respondents from rural areas. The table presents the results of the estimation of Specification (1). The dependent variables are individuals' perceptions of government and the country's institutions. Controls include age, age squared, gender, marital status, dummies for high school and university education, employment status, urban status, the regions' average level of income, the log of the countries' GDP per capita, the countries' unemployment rate, and dummies for democracy status. Standard errors in parentheses are corrected for two-way clusters at the level of the subnational regions (to account for correlation over time) and at the level of the countries in each year (to account for within-country-year correlation).

Table II
Event-study results

	(1)	(2)	(3)	(4)	(5)	(6)
<i>Dep. Var.:</i>	1st principal component of the government approval responses					
<i>Sample note 1:</i>	Regions with a sharp increase in 3G coverage in one year in 2007-2018					
<i>Sample note 2:</i>	All respondents			Rural respondents		
Regional 3G coverage	-0.055*** (0.014)			-0.073*** (0.018)		
Post-event dummy		-0.036*** (0.012)			-0.052*** (0.014)	
Sharp increase in regional 3G coverage occurred in:						
Year $t + 4$ or later			-0.015 (0.021)			-0.011 (0.024)
Year $t + 3$			0.001 (0.019)			-0.005 (0.021)
Year $t + 2$			-0.006 (0.014)			0.006 (0.017)
Year t			-0.032** (0.013)			-0.035** (0.015)
Year $t - 1$			-0.048*** (0.016)			-0.066*** (0.019)
Year $t - 2$			-0.033* (0.019)			-0.053** (0.021)
Year $t - 3$			-0.063*** (0.022)			-0.067*** (0.024)
Year $t - 4$ or earlier			-0.051* (0.030)			-0.061** (0.030)
Observations	130,406	130,406	130,406	66,078	66,078	66,078
R-squared	0.213	0.212	0.213	0.242	0.242	0.242
Number of countries	65	65	65	62	62	62
Number of regions	452	452	452	444	444	444
Number of countries with variation	65	36	65	62	32	62
Number of regions with variation	452	219	452	444	206	444
Subnational region & year FEs	✓	✓	✓	✓	✓	✓
Baseline controls	✓	✓	✓	✓	✓	✓
Censorship of the traditional press control	✓	✓	✓	✓	✓	✓
P-value: $\gamma[Y_t] = \gamma[Y_{t-2}]$			0.119			0.010
P-value: $(\gamma[Y_t] + \gamma[Y_{t+1}])/2 = (\gamma[Y_{t-2}] + \gamma[Y_{t-3}])/2$			0.024			0.005

Note: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. The table presents the results of the event study. The unit of observation is an individual. The sample is comprised of individuals from regions that had a sharp increase in 3G coverage, more than 50 percentage points in the share of a subnational region's population covered by 3G in a single year, between 2007 and 2018. There are 452 regions from 65 countries like this. All regions in this sample have variation in the lags and leads of the year of the event (estimated in Columns 3 and 6). However, only 219 regions out of all regions with an event have variation in the post-event dummy within the sample, due to missing region-years in GWP data. Columns 1 to 3 report results for the full sample; Column 4 to 6—for the subsample of respondents from rural areas. The unreported controls include age, age squared, gender, marital status, dummies for high school and university education, employment status, urban status, the regions' average level of income, the log of the countries' GDP per capita, the countries' unemployment rate, dummies for democracy status, and censorship-of-the-traditional-press score. Standard errors in parentheses are corrected for two-way clusters at the level of the subnational regions (to account for correlation over time) and at the level of the countries in each year (to account for within-country-year correlation).

Table III
The effect of 2G coverage on internet usage and confidence in government

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
<i>Dep. Var.:</i>	Confidence in national government	Confidence in judicial system	Honesty of elections	No corruption in government	Share of questions with positive responses	1st principal component of responses	Internet access at home
Panel A: The effect of 2G on confidence in government and internet access at home							
Regional 2G coverage	0.045 (0.029)	0.031 (0.020)	0.098*** (0.030)	0.054*** (0.019)	0.056*** (0.021)	0.056** (0.022)	-0.013 (0.020)
Observations	772,353	748,471	732,856	722,768	617,863	617,863	840,537
Mean dep. var.	0.514	0.534	0.505	0.226	0.432	0.439	0.440
Panel B: The effect of 3G and 2G on confidence in government and internet access at home							
Regional 3G coverage	-0.060*** (0.020)	-0.038*** (0.015)	-0.074*** (0.020)	-0.032** (0.014)	-0.053*** (0.015)	-0.053*** (0.015)	0.080*** (0.017)
Regional 2G coverage	0.037 (0.028)	0.026 (0.019)	0.088*** (0.030)	0.049** (0.019)	0.048** (0.021)	0.048** (0.021)	-0.002 (0.019)
Observations	772,353	748,471	732,856	722,768	617,863	617,863	840,537
Mean dep. var.	0.514	0.534	0.505	0.226	0.432	0.439	0.440
Subnational region & year FEs	✓	✓	✓	✓	✓	✓	✓
Baseline controls	✓	✓	✓	✓	✓	✓	✓

Note: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. The table presents the effects of 2G coverage on internet usage and government support. The results suggest that, as expected, the change in 2G coverage did not increase internet access at home and, on average, increased government support. The unit of observation is an individual. Panel A reports results for the effect of 2G coverage, Panel B—similar results with 3G coverage included as a control variable. Columns 1 to 6 present the results for government approval as the outcome variables; Column 7—for a dummy for having access to the internet at home. Baseline controls include age, age squared, gender, marital status, dummies for high school and university education, employment status, urban status, the regions' average level of income, the log of the countries' GDP per capita, the countries' unemployment rate, and dummies for democracy status. Standard errors in parentheses are corrected for two-way clusters at the level of the subnational regions (to account for correlation over time) and at the level of the countries in each year (to account for within-country-year correlation).

Table IV
Lightning strikes, 3G coverage, and government approval

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<i>Dep. Var.:</i>	Regional 3G coverage	1st principal component of government approval	Regional 3G coverage	1st principal component of government approval	Regional 3G coverage	1st principal component of government approval	Regional 3G coverage	1st principal component of government approval
<i>Stage, 2SLS:</i>	1	2	1	2	1	2	1	2
<i>Countries in the sample:</i>	All countries				Countries with below-median GDP per capita			
<i>Respondents in the sample:</i>	All	All	Rural	Rural	All	All	Rural	Rural
Regional 3G coverage		-0.283*** (0.086)		-0.308*** (0.093)		-0.329*** (0.099)		-0.389*** (0.108)
$\mathbb{1}[\text{High frequency of lightning strikes per sq. km}] \times \text{Year} \times \mathbb{1}[\text{GDP per capita below median}]$	-0.032*** (0.005)		-0.032*** (0.005)		-0.033*** (0.007)		-0.033*** (0.007)	
$\mathbb{1}[\text{High frequency of lightning strikes per sq. km}] \times \text{Year} \times \mathbb{1}[\text{GDP per capita above median}]$	-0.010** (0.005)		-0.009* (0.005)					
Observations	12,860	12,860	11,743	11,743	5,789	5,789	5,324	5,324
Mean dep. var.	0.373	0.432	0.369	0.439	0.134	0.433	0.124	0.440
F-stat, excluded instrument		20.74		19.13		25.15		25.47
Corresponding OLS coefficient on regional 3G coverage						-0.120*** (0.027)		-0.158*** (0.031)
Subnational region & year FEs	✓	✓	✓	✓	✓	✓	✓	✓
Extended set of controls	✓	✓	✓	✓	✓	✓	✓	✓

Note: *** p<0.01, ** p<0.05, * p<0.1. The table presents the results of an IV analysis, where the frequency of lightning strikes per sq. km. in a subnational region is used as an IV for the expansion of regional 3G coverage. The methodology follows [Manacorda and Tesei \(2020\)](#). High frequency of lightning strikes per sq. km is defined by the subnational region being in the top half of the distribution of lightning strikes per sq. km. Odd columns present the results of the first stage. Even columns—the results of the second stage. Columns 1 to 4 present the results for all the countries in the sample; Columns 5 to 8—for the subsample of countries with below-median GDP per capita. The unit of observation is a subnational region. Controls include the region’s average level of income, the log of the countries’ GDP per capita, the countries’ unemployment rate, dummies for democracy status, and linear time trends interacted with the subnational regions’ share of territory covered by deserts, share of territory covered by mountains, maximum elevation, dummies for each quintile of population density, 3G coverage in 2008, a dummy for whether the region had any 3G coverage in 2008, and a dummy for whether the country had any 3G coverage in 2008. Standard errors in parentheses are corrected for two-way clusters at the level of the subnational regions (to account for correlation over time) and at the level of the countries in each year (to account for within-country-year correlation).

Table V

The effect of 3G coverage on government approval, depending on the level of internet censorship and on the level of censorship of the traditional media

	(1)	(2)	(3)	(4)	(5)	(6)
<i>Dep. Var.:</i>	Confidence in national government	Confidence in judicial system	Honesty of elections	No corruption in government	Share of questions with positive responses	1st principal component of responses
Panel A: Time-variant dummy for internet censorship						
Regional 3G coverage	-0.100*** (0.023)	-0.057*** (0.016)	-0.117*** (0.021)	-0.054*** (0.016)	-0.081*** (0.016)	-0.082*** (0.016)
Regional 3G coverage × Internet censorship dummy	0.105** (0.041)	0.037 (0.029)	0.173*** (0.043)	0.054* (0.029)	0.093*** (0.034)	0.094*** (0.035)
Internet censorship dummy	0.068** (0.033)	0.042* (0.024)	0.053* (0.031)	0.011 (0.023)	0.045* (0.027)	0.046* (0.028)
Observations	656,015	631,606	618,480	613,737	521,632	521,632
R-squared	0.157	0.166	0.157	0.234	0.238	0.235
Panel B: Time-invariant dummy for internet censorship						
Regional 3G coverage	-0.098*** (0.025)	-0.055*** (0.018)	-0.124*** (0.023)	-0.056*** (0.017)	-0.081*** (0.018)	-0.082*** (0.018)
Regional 3G coverage × Dummy: countries with internet censorship	0.091** (0.043)	0.027 (0.028)	0.201*** (0.043)	0.056*** (0.021)	0.084*** (0.031)	0.085*** (0.032)
Observations	648,705	624,264	611,221	606,955	515,365	515,365
R-squared	0.157	0.166	0.158	0.235	0.239	0.235
Panel C: Time-variant dummies for internet censorship and above-median press censorship						
Regional 3G coverage	-0.032 (0.029)	-0.027 (0.020)	-0.089*** (0.024)	-0.026 (0.021)	-0.046** (0.020)	-0.047** (0.020)
Regional 3G coverage × Internet censorship dummy	0.157*** (0.044)	0.059** (0.030)	0.195*** (0.046)	0.078** (0.031)	0.121*** (0.035)	0.123*** (0.036)
Regional 3G coverage × Above-median press censorship dummy	-0.116*** (0.034)	-0.051** (0.023)	-0.046 (0.030)	-0.049* (0.025)	-0.059** (0.025)	-0.060** (0.025)
Internet censorship dummy	0.057* (0.032)	0.037 (0.024)	0.049 (0.031)	0.005 (0.023)	0.039 (0.027)	0.040 (0.027)
Above-median press censorship dummy	0.123*** (0.034)	0.023 (0.021)	0.070** (0.030)	0.059** (0.025)	0.068*** (0.023)	0.069*** (0.023)
Observations	656,015	631,606	618,480	613,737	521,632	521,632
R-squared	0.158	0.166	0.158	0.234	0.239	0.236
Panel D: Time-invariant dummies for internet censorship and above-median press censorship						
Regional 3G coverage	-0.040 (0.032)	-0.019 (0.024)	-0.109*** (0.026)	-0.025 (0.023)	-0.052** (0.023)	-0.053** (0.023)
Regional 3G coverage × Dummy: countries with internet censorship	0.154*** (0.050)	0.066** (0.033)	0.218*** (0.050)	0.089*** (0.025)	0.115*** (0.037)	0.117*** (0.037)
Regional 3G coverage × Dummy: countries with above-median press censorship	-0.117*** (0.043)	-0.072** (0.032)	-0.031 (0.038)	-0.061** (0.030)	-0.057* (0.031)	-0.058* (0.032)
Observations	648,705	624,264	611,221	606,955	515,365	515,365
R-squared	0.157	0.166	0.158	0.235	0.239	0.236
Subnational region & year FEs	✓	✓	✓	✓	✓	✓
Baseline controls	✓	✓	✓	✓	✓	✓

Note: *** p<0.01, ** p<0.05, * p<0.1. The unit of observation is an individual. The dependent variables are individuals' perceptions of government and the country's institutions. Panels A and C use time-variant measures of censorship, whereas Panels B and D use time-invariant measures. Unreported controls include age, age squared, gender, marital status, dummies for high school and university education, employment status, urban status, the regions' average level of income, the log of the countries' GDP per capita, the countries' unemployment rate, and dummies for democracy status. Standard errors in parentheses are corrected for two-way clusters at the level of the subnational regions (to account for correlation over time) and at the level of the countries in each year (to account for within-country-year correlation).

Table VI
The relationship between actual and perceived corruption

	(1)	(2)	(3)	(4)	(5)	(6)
<i>Dep. Var.:</i>	Perception of no corruption in government					
<i>Sample:</i>	All	Rural	All	Rural	All	Rural
Regional 3G coverage × Actual corruption incidents	-0.081*** (0.025)	-0.101*** (0.030)	-0.059** (0.024)	-0.062** (0.029)		
Regional 3G coverage × Actual corruption incidents × × Country with below-median overall corruption					-0.222*** (0.035)	-0.243*** (0.041)
Regional 3G coverage × Actual corruption incidents × × Country with above-median overall corruption					-0.030 (0.020)	-0.039* (0.024)
Regional 3G coverage	-0.014 (0.016)	-0.025 (0.019)	-0.019 (0.016)	-0.037* (0.020)		
Regional 3G coverage × × Country with below-median overall corruption					0.005 (0.021)	0.002 (0.026)
Regional 3G coverage × × Country with above-median overall corruption					-0.033* (0.019)	-0.057*** (0.020)
Actual corruption incidents	-0.013 (0.010)	-0.013 (0.012)	-0.017* (0.010)	-0.021* (0.012)	-0.015 (0.010)	-0.017 (0.012)
Observations	691,872	414,346	581,944	354,966	691,872	414,346
R-squared	0.226	0.192	0.151	0.126	0.227	0.193
Subnational region & year FEs	✓	✓	✓	✓	✓	✓
Baseline controls	✓	✓	✓	✓	✓	✓
Sample excludes observations with zero corruption incidents			✓	✓		

Note: *** p<0.01, ** p<0.05, * p<0.1. The outcome variable is a dummy for the perception that there is no corruption in government. Actual corruption incidents stand for the IMF's Global Incidents of Corruption Index (GICI). The unit of observation is an individual. Unreported controls include age, age squared, gender, marital status, dummies for high school and university education, employment status, urban status, the regions' average level of income, the log of the countries' GDP per capita, the countries' unemployment rate, and dummies for democracy status. Standard errors in parentheses are corrected for two-way clusters at the level of the subnational regions (to account for correlation over time) and at the level of the countries in each year (to account for within-country-year correlation).

Table VII

3G coverage, the number of entities in the Panama Papers, and perceived corruption

	(1)	(2)	(3)	(4)	(5)	(6)
<i>Dep. Var.:</i>	Perception of no corruption in government					
<i>Countries in the sample:</i>	All countries		Excluding low-income countries			
Regional 3G coverage ×						
× Number of Panama Papers entities per capita	-0.035** (0.014)					
× Number of Panama Papers entities per capita × Before Panama Papers		-0.031** (0.014)	-0.033** (0.014)			
× Number of Panama Papers entities per capita × After Panama Papers		-0.037** (0.018)	-0.048*** (0.017)			
× 1[Top 10% of countries by Panama Papers entities per capita] × Before Panama Papers				-0.045 (0.033)		
× 1[Top 10% of countries by Panama Papers entities per capita] × After Panama Papers				-0.100** (0.040)		
× Number of Panama Papers entities × Before Panama Papers					-0.012*** (0.004)	
× Number of Panama Papers entities × After Panama Papers					-0.017*** (0.005)	
× 1[Top 10% of countries by Panama Papers entities] × Before Panama Papers						-0.092*** (0.028)
× 1[Top 10% of countries by Panama Papers entities] × After Panama Papers						-0.174*** (0.038)
Regional 3G coverage	-0.027* (0.014)	-0.017 (0.013)	-0.015 (0.014)	-0.015 (0.014)	-0.012 (0.014)	-0.008 (0.014)
Regional 3G coverage × After Panama Papers		-0.011 (0.014)	-0.001 (0.014)	0.003 (0.014)	-0.002 (0.014)	0.000 (0.015)
Observations	722,768	722,768	620,827	620,827	620,827	620,827
R-squared	0.225	0.226	0.232	0.232	0.232	0.232
<i>p-value</i> $\beta(\text{Before Panama Papers}) = \beta(\text{After Panama Papers})$		0.490	0.055*	0.058*	0.073*	0.0095***
Subnational region & year FEs	✓	✓	✓	✓	✓	✓
Baseline controls	✓	✓	✓	✓	✓	✓
All lower-level interactions	✓	✓	✓	✓	✓	✓
Interactions of 3G and regional income	✓	✓	✓	✓	✓	✓

Note: *** p<0.01, ** p<0.05, * p<0.1. The outcome variable is a dummy for the perception that there is no corruption in government. “Number of Panama Papers entities” is the number of entities from a country in the Panama Papers. “Number of Panama Papers entities per capita” is the number of entities from a country in the Panama Papers per 1,000 inhabitants. “Before Panama Papers” and “After Panama Papers” are dummies indicating whether the GWP interview took place before or after the release of the Panama Papers to the public. The unit of observation is an individual. Controls include age, age squared, gender, marital status, dummies for high school and university education, employment status, urban status, the region’s average level of income, the log of the countries’ GDP per capita, the countries’ unemployment rate, dummies for democracy status, the Freedom of the Press score, and the interactions of regional 3G coverage with the region’s average level of income. Standard errors in parentheses are corrected for two-way clusters at the level of the subnational regions (to account for correlation over time) and at the level of the countries in each year (to account for within-country-year correlation).

Table VIII
The effect of 3G coverage on incumbents' electoral performance in Europe

	(1)	(2)	(3)	(4)	(5)
<i>Dep. Var.:</i>	Vote share of:				
	Top 2 parties from the 1st election	Ruling party (the party of the Prime Minister)	Populist parties if they are among top 2 parties from the 1st election	Turnout	
<i>Unit of observation:</i>	District-year	District-year-incumbent	District-year	District-year	
District 3G coverage	-0.089** (0.045)	-0.089*** (0.031)		-0.090** (0.036)	-0.038*** (0.012)
District 3G coverage × Populist party			-0.120** (0.050)		
District 3G coverage × Nonpopulist party			-0.084*** (0.032)		
Observations	1,234	1,536	1,536	341	1,250
R-squared	0.889	0.917	0.917	0.982	0.968
Mean dep. var.	0.561	0.304	0.304	0.329	0.656
Mean 3G coverage	0.649	0.645	0.645	0.655	0.647
District & year FEs	✓			✓	✓
Incumbent-by-district & year FEs		✓	✓		
Baseline controls	✓	✓	✓	✓	✓
Excl. countries without populists among top 2 in the 1st election				✓	

Note: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. The 3G expansion led to a decrease in the vote share for incumbent parties. This is true for both nonpopulist and populist incumbent parties. In Columns 1, 4, and 5, the unit of observation is a subnational district in an election. In Columns 2-3, the unit of observation is an incumbent party in a subnational district in an election. The data in Column 5 cover 102 parliamentary elections in 33 European countries (this is the full panel). In Columns 1, 2, and 3, Romania is excluded because, in Romania, after the first election, the top 2 parties merged with other large parties. In Columns 2-3, Switzerland is excluded because, in Switzerland, the position of the president rotates among the parties in the ruling coalition. In Column 4, the sample is restricted to countries that had populist parties among the top 2 parties in the first election. Controls include the country's unemployment rate, labor-force participation rate, inflation rate, log of GDP per capita, the share of population over 65 years old, and the subnational district's average nighttime light density. As the nighttime light density data for 2007-2013, 2014, and 2015-2018 come from different sources (DMSP-OLS, a combination of DMSP-OLS and VIIRS, and VIIRS, respectively), we interact the measure of nighttime light density with a dummy for each of those time periods. Standard errors presented in parentheses are corrected for two-way clusters at the level of the subnational districts (to account for over time correlation) and at the level of the countries in each year (to account for within-country-year correlation).

Table IX

The effect of 3G coverage on the opposition's electoral performance in Europe

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
<i>Dep. Var.:</i>	Vote share of:						
	Right-wing populists	Left-wing populists	Other populists	All populists	All populists	Green parties	Nonpopulist opposition
<i>Unit of observation:</i>	District-year	District-year	District-year	District-year	District-year	District-year	District-year- ruling coalition
District 3G coverage	0.086*** (0.024)	0.067*** (0.022)	-0.038 (0.024)	0.115*** (0.039)	0.129*** (0.042)	-0.007 (0.012)	-0.030 (0.053)
Observations	1,250	1,250	1,250	1,250	1,002	1,141	1,566
R-squared	0.961	0.876	0.934	0.924	0.813	0.870	0.904
Mean dep. var	0.136	0.065	0.060	0.260	0.189	0.039	0.431
Mean 3G coverage	0.647	0.647	0.647	0.647	0.648	0.636	0.654
District & year FEs	✓	✓	✓	✓	✓	✓	
Ruling-coalition-by-district & year FEs							✓
Baseline controls	✓	✓	✓	✓	✓	✓	✓
Excl. countries with populists in power					✓		

Note: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. The expansion of 3G networks led to an increase in both right-wing and left-wing populists' vote share, but not in the vote share for green parties or the nonpopulist opposition in general. In Columns 1 to 6, the unit of observation is a subnational district in an election. In Column 7, the unit of observation is the ruling coalition in the subnational district in an election. The data in Columns 1-5 cover 102 parliamentary elections in 33 European countries (the full panel). In Column 6, there are fewer observations than in Columns 1-5 because in five elections (Spain in 2015-2016, Croatia in 2015-2016, and Greece in 2015) Green parties formed joint lists with large non-Green parties, making it impossible to determine what share of the votes went to the Green parties and what to their partners. Column 5 excludes all countries, in which populists were a ruling party at some point during the sample period: Bulgaria, Hungary, Italy, Montenegro, North Macedonia, Poland, Slovakia, and Slovenia. In Column 7, the election results for Switzerland and Romania are excluded because, in Switzerland, all the major parties are a part of the ruling coalition, and in Romania, after the first election, the parties in the ruling coalition merged with parties outside of the ruling coalition. Controls include the country's unemployment rate, labor force participation rate, inflation rate, log of GDP per capita, the share of population over 65 years old, and the regions' average level of nighttime light density. As the nighttime light density data for 2007-2013, 2014, and 2015-2018 come from different sources (DMSP-OLS, a combination of DMSP-OLS and VIIRS, and VIIRS, respectively), we also interact the measure of nighttime light density with a dummy for each of those time periods. Standard errors presented in parentheses are corrected for two-way clusters at the level of the subnational district (to account for over time correlation) and at the level of the countries in each year (to account for within-country-year correlation).