

The Impact of Media Censorship: 1984 or Brave New World?[†]

By YUYU CHEN AND DAVID Y. YANG*

Media censorship is a hallmark of authoritarian regimes. We conduct a field experiment in China to measure the effects of providing citizens with access to an uncensored internet. We track subjects' media consumption, beliefs regarding the media, economic beliefs, political attitudes, and behaviors over 18 months. We find four main results: (i) free access alone does not induce subjects to acquire politically sensitive information; (ii) temporary encouragement leads to a persistent increase in acquisition, indicating that demand is not permanently low; (iii) acquisition brings broad, substantial, and persistent changes to knowledge, beliefs, attitudes, and intended behaviors; and (iv) social transmission of information is statistically significant but small in magnitude. We calibrate a simple model to show that the combination of low demand for uncensored information and the moderate social transmission means China's censorship apparatus may remain robust to a large number of citizens receiving access to an uncensored internet. (JEL C93, D72, D83, L82, L86, L88, P36)

What Orwell feared were those who would ban books. What Huxley feared was that there would be no reason to ban a book, for there would be no one who would want to read one.

—Neil Postman, *Amusing Ourselves to Death*

*Chen: Guanghua School of Management, Peking University (email: chenyuyu@gsm.pku.edu.cn); Yang: Harvard University, and Massachusetts Institute of Technology (email: davidyang@fas.harvard.edu). Stefano DellaVigna was the coeditor for this article. Yang is deeply grateful to Ran Abramitzky, Matthew Gentzkow, and Muriel Niederle for their advice at all stages of this project, and to Doug Bernheim, Pascaline Dupas, and Al Roth for their unceasing guidance and support. Helpful and much appreciated comments and suggestions were provided by David Atkin, Abhijit Banerjee, David Broockman, Davide Cantoni, Kate Casey, Arun Chandrasekhar, Raj Chetty, Dave Donaldson, Ruben Enikolopov, Ben Enke, David Laibson, Shengwu Li, John List, Sendhil Mullainathan, Nathan Nunn, Jen Pan, Maria Petrova, Molly Roberts, Gerard Roland, Mark Rosenzweig, Andrei Shleifer, Adam Szeidl, Noam Yuchtman, Sevgi Yuksel, and many seminar and conference participants. Keping Wang, Yifei Sun, Shuai Chen, Ji Xu, Xi Yao, Yu Li, Yingzhen Zhang, Ziteng Lei, Xincheng Qiu, Min Fang, Huitian Bai, and Miqi Xiong and a dedicated survey recruitment team provided outstanding assistance in implementing the field experiment. Financial support from the Stanford Center for International Development, the Stanford Center at Peking University, the Stanford Graduate Research Opportunity Fund, the Stanford Freeman Spogli Institute for International Studies, the Stanford Institute for Research in the Social Sciences, the Stanford Gerald J. Lieberman Fellowship, the Stanford Department of Economics, the Stanford Economic Research Laboratory, the National Science Foundation of China (grant 71425006), and Al Roth is greatly appreciated. The experiment is approved by Stanford IRB (Protocol ID 34318), and is registered at AEA RCT registry (ID 0001412). The authors declare that they have no relevant or material financial interests that relate to the research described in this paper.

[†]Go to <https://doi.org/10.1257/aer.20171765> to visit the article page for additional materials and author disclosure statements.

Media censorship is a hallmark of authoritarian regimes.¹ Countries such as China spend a tremendous amount of resources to block foreign websites so that uncensored, regime-threatening information is out of citizens' reach. Scholars have long suggested that censorship is key to the popular support and stability of these regimes (Ford 1935). Nonetheless, direct empirical evidence about the effect of censorship is limited.

In this paper, we ask two questions. Does providing access to an uncensored internet lead citizens to acquire politically sensitive information? Does the acquisition of politically sensitive information change citizens' beliefs, attitudes, and behaviors? Answers to these questions are far from clear. Citizens with access to uncensored internet may not seek out politically sensitive information, due to lack of interest in politics, fear of government reprisal, and unawareness or distrust of foreign news outlets. Even if they do acquire such information and become fully informed, their attitudes and beliefs may not change.

We conduct a field experiment in China in order to answer these two questions. We randomly assign 1,800 university students in Beijing to either a control condition in which their internet use is subject to status quo censorship, or to a treatment condition in which they are given tools to bypass internet censorship for free for 18 months. A subset of the treated students also receive temporary encouragement for 4 months to visit Western news outlets otherwise blocked by China's censorship apparatus. We directly observe all browsing activities of foreign websites by the treated students. We also observe students' decisions to purchase access to uncensored internet themselves after the experiment ends. Using surveys, we repeatedly measure a wide range of outcomes, including students' knowledge of current and historical events, beliefs and attitudes toward media, economic beliefs, political attitudes, and intended behaviors.

We find four main results. First, access to uncensored internet alone has little impact on students' acquisition of politically sensitive information. Nearly one-half of the students do not use the tools to bypass censorship at all. Among those who do, almost none spend time browsing foreign news websites that are blocked. These numbers indicate that students' low demand for uncensored, politically sensitive information is an important reason why they do not consume such information, in spite of the low cost.

Second, modest and temporary incentives to visit Western news outlets lead to a persistent increase in students' acquisition of politically sensitive information. Students spend more time on foreign news websites even after the incentivized encouragement ends. This persistent increase suggests that demand is not inherently low. In particular, fear of government reprisal is unlikely the reason students do not demand sensitive information. Rather, an important factor shaping students' low demand appears to be their underestimation of the value of uncensored information. A period of exposure to foreign news outlets increases students' reported trust of these outlets, and makes them willing to pay a higher price for the access.

¹ Freedom House's *Freedom of the Press Report* shows that 86 percent of the world's population does not enjoy media free from censorship. In particular, states with "unfree" media are concentrated among regimes that are undemocratic and grant limited political rights for their citizens. Source: freedomhouse.org/report/freedom-press/2016/china (accessed December 11, 2016).

The temporary intervention, by raising demand, has resulted in a lasting increase in students' acquisition of uncensored information. By the end of the experiment, about 23 percent of the newly exposed students pay to continue their uncensored internet access.²

Third, acquisition of politically sensitive information brings broad, substantial, and persistent changes to students' knowledge, beliefs, attitudes, and intended behaviors. Acquisition, as a result of free access and temporary incentives, makes students (i) more knowledgeable of current events censored on domestic media, as well as politically sensitive events in the recent past; (ii) more pessimistic about Chinese economic growth and stock market performance in the near future, revealed in an incentive-compatible manner; (iii) more skeptical of the Chinese government, less satisfied with its performance, and more likely to demand changes in Chinese institutions; and (iv) more likely to plan on exiting through foreign graduate schools (albeit no change in direct engagement with the regime such as protesting), and more likely to report having pulled out investments in the Chinese stock market (among the small number of investors). If we rank students across all these dimensions, the access and encouragement combined have moved the median student from the forty-seventh percentile of the distribution before the experiment to the fifty-sixth percentile by the end of the experiment.

Fourth, students who acquire politically sensitive information transmit some of their knowledge to their peers, but the magnitude of such spillovers is small. Exploiting the variation in treatment saturation across dorm rooms, we find that if a student actively browses foreign news websites and is informed of a sensitive news event, her roommate is on average 12.7 percentage points more likely to correctly answer a quiz on that same event. This rate implies that direct comparisons between treatment and control groups are not substantially downward biased. Moreover, a simple calibration exercise suggests that the social transmission is too moderate to qualitatively affect the knowledge level among the entire student population, given the proportion of students who have had access to uncensored information prior to our experiment. We speculate factors that dampen social transmission of politically sensitive knowledge in Section IVB. We note, however, that our data only allow us to observe transmission among roommates, and hence can underestimate the overall social transmission of information.

Taken together, our findings suggest that censorship in China is effective not only because the regime makes it difficult to access sensitive information, but also because it fosters an environment in which citizens do not demand such information in the first place. Several limitations regarding the external validity and general equilibrium effects are worth stressing when one interprets our experimental results. In the final section of the paper, we take the partial equilibrium effects estimated from the experiment and calibrate a simple model to show that (i) the share of students who have access to uncensored internet prior to the experiment is too low for sensitive information to spread throughout the population; and (ii) the porous censorship apparatus would be robust even if the (unencouraged) access were

² Similar results are found regarding other unfamiliar but beneficial technology. For example, Dupas (2014) finds that a one-time subsidy on antimalarial bed nets has a positive impact on Kenyan villagers' willingness to pay a year later, which is predominately driven by villagers learning about the value of bed nets.

provided to a substantially larger share of students. The robustness is driven by the low demand for, and the moderate social transmission of, uncensored information, even among the young and educated population. Importantly, we note that various general equilibrium effects that could undermine the robustness are held fixed given the relatively small scale of the experiment.

Our model simulation demonstrates that the censorship apparatus can also be fragile, precisely because its effectiveness depends on citizens' suppressed demand for uncensored information. Exposure to foreign media can change citizens' beliefs regarding its value, and hence persistently raise their demand for uncensored information. If we were to provide encouragement at the level of this experiment to all students in addition to free access, enough students would begin to actively acquire sensitive information, so that the entire student population would become informed, taking into account the social transmission of information. These students could destabilize the censorship apparatus and impose substantial pressure on the regime to tighten its grip.

Our paper contributes to the large body of literature on the political economy of mass media. Our study adds an important data point to our understanding of how mass media influences citizens' political preferences and shapes aggregate outcomes.³ We investigate the case of China, the largest country engaging in state-led information control, and we provide the first causal evidence via a field experiment to identify how internet censorship shapes citizens' knowledge, economic beliefs, political attitudes, and behaviors. We show that state control of information is effective, making citizens more supportive of the regime. This relates to the broad literature on media capture (see Prat 2015, Enikolopov and Petrova 2015, and Strömberg 2015, among others). The extent to which a censored internet affects citizens also complements our knowledge on the operation and underlying objectives of China's censorship apparatus.⁴

In particular, this paper relates to the strand of the literature on mass media that emphasizes the importance of demand-side factors (e.g., Mullainathan and Shleifer 2005, Gentzkow and Shapiro 2006; see Gentzkow, Shapiro, and Stone 2015 for a survey). In the domain of censorship, while many model censorship primarily as the government obstructing access to valuable information (see Schedler 2010, Shadmehr and Bernhardt 2015, and Guriev and Treisman 2018, among others), Gehlbach and Sonin (2014) build on the framework and endogenize citizens' media consumption. Our findings show that one cannot understand the impact of media censorship without taking into account of citizens' demand for uncensored information. Specifically, such demand is inelastic if citizens are uninformed about the extent of

³ For example, DellaVigna and Kaplan (2007) and Gerber, Karlan, and Bergan (2009) on the United States; Yanagizawa-Drott (2014) on Rwanda Genocide; Adena et al. (2015) on Nazi Germany; and Enikolopov, Petrova, and Zhuravskaya (2011) on contemporary Russia. DellaVigna and Gentzkow (2010) review the empirical literature on persuasion across broader domains, and Prat and Strömberg (2009) provide a more recent survey of this literature, particularly in the domain of politics.

⁴ King, Pan, and Roberts (2013, 2014) show that the censorship algorithm prioritizes to eliminate information related to collective actions. Lorentzen (2013) and Huang, Boranbay-Akan, and Huang (2019) argue that the Chinese government strategically allows a limited amount of sensitive information to flow on domestic social media in order to facilitate the central government addressing popular discontent more effectively. More recently, Qin, Strömberg, and Wu (2017) argue that social media content can be deployed to enhance state surveillance, effectively detecting and predicting offline protests.

ensorship, but demand becomes elastic once they have been exposed to uncensored outlets. This is consistent with Abramitzky and Sin (2014) who find that inflow of Western knowledge into Eastern Europe after 1989 reflected underlying demand differences across these countries; Simonov and Rao (2018) who find that news demand responds to the bias in state-controlled media in Russia; Hobbs and Roberts (2018) who document the increased demand for censorship circumvention tools due to the sudden block of Instagram; and Roberts (2018) who shows that the Chinese censorship apparatus deploys frictions such as slowing down connections to achieve information control.

These findings also contribute to the growing empirical literature on the endogenous formation of beliefs and preferences when authoritarian regimes have a direct incentive to intervene. Among others, state indoctrination (Voigtländer and Voth 2015; Cantoni et al. 2017) and historical experiences (Alesina and Fuchs-Schündeln 2007; Fuchs-Schündeln and Schündeln 2015; Chen and Yang 2018a) have been identified as generating lasting impacts on citizens' political attitudes. We show that censorship can effectively manipulate citizens' beliefs, attitudes, and preferences along the direction of the regimes' intentions. In particular, despite citizens' moderate level of awareness and sophistication regarding media censorship and the biases in censored information, they cannot fully debias themselves from the distorted information environment.⁵

In what follows, we provide a brief overview of internet censorship in China in Section I. In Section II, we describe the experimental design, outcome variables of interest, and other empirical setups of the field experiment. In Section III, we present results on whether providing access increases acquisition of sensitive information, and in Section IV, we present results on whether acquiring sensitive information affects knowledge, beliefs, attitudes, and behaviors. In Section V, we simulate the counterfactual scenarios of media censorship in China. Finally, in Section VI, we discuss lessons from our experimental results and speculate on the external validity of this study on other authoritarian regimes that deploy internet censorship.

I. Internet Censorship in China

The media landscape in China is among the most regulated and restricted in the world, and China's media freedom is ranked consistently toward the bottom. In particular, China's information control over the internet, primarily through censorship, is second to none in terms of its scale and technological sophistication.⁶ In this section, we briefly describe the infrastructure of the Great Firewall that serves as the building block of censorship, and the market for tools to circumvent

⁵ Some recent studies investigate how people update beliefs based on censored information. In an abstract setting, Enke (2018) documents that people form biased beliefs by neglecting absence and non-occurrence, failing to take into account the selection underlying the data-generating process. In political contexts, Bai et al. (2015) show that Chinese citizens have difficulties interpreting information on air pollution when the government-controlled media conflicts with uncensored sources; and Huang and Yeh (2019) find that exposing Chinese citizens to selected news articles from foreign media that report on foreign societies may induce, in the short run, more favorable attitudes toward China.

⁶ China is home to the world's most sophisticated internet censorship apparatus. The Freedom House's *Freedom of the Net Report* in 2017 labels China's "Net Freedom Status" as *not free*, and rates its "Internet Freedom Score" as 87 (out of 100, where 100 indicates the most unfree); the "world's worst abuser of Internet freedom." Source: <https://freedomhouse.org/report/freedom-net/2017/china> (accessed November 26, 2017).

internet censorship in China. We briefly outline the administrative and legal framework of internet regulations in China in online Appendix A.

A. *The Great Firewall*

Media outlets based domestically would incur severe business and political costs from publishing content that the state deems threatening and objectionable. As a result, content on domestic media is either routinely self-censored during the editorial process, or censored and filtered according to orders from the Propaganda Department of the Communist Party of China. Among the most heavily censored topics in 2016 are government corruption, media censorship, civil society activism, ethnic tensions, health, and safety scandals.⁷ Transmission of politically sensitive information on domestic social media such as Weibo and WeChat is also limited due to platform-wide keyword filters and ex post content deletion.

Since the Chinese government does not have the jurisdiction to directly control foreign media outlets, an important aspect of China's internet regulation is its effort to block internet users in China from accessing specific foreign websites. The Great Firewall, a major part of the umbrella *Golden Shield Project* directed by China's Ministry of Public Security, has operated since 2003 and serves as the main infrastructure blocking access to potentially unfavorable incoming data from foreign media outlets.

The Great Firewall deploys several technologies to block entire websites or specific web pages from being accessed by IP addresses located in China. During the time frame of our field experiment, 12 of the 100 most trafficked websites in the world (and 161 of the Alexa top 1000 global websites) have been blocked by the Great Firewall. Some prominent examples are: Google, YouTube, Facebook, Twitter, Instagram, Blogspot, Tumblr, Dropbox, Blogger, Vimeo, Soundcloud, and Flickr. In particular, 9 of the top 20 news websites ranked by Alexa are blocked by the Great Firewall: for example, *CNN*, *The New York Times*, *The Guardian*, *BBC*, *Bloomberg*, *The Wall Street Journal*, and *Reuters*.⁸

Our project focuses on the blocked foreign news websites, a primary source of politically sensitive information and could potentially shape Chinese readers' knowledge, beliefs, and attitudes. In sharp contrast to their domestic counterparts, foreign news websites report politically sensitive news events and often feature uncensored investigative journalism on sensitive topics in China (Qin, Strömberg, and Wu 2018).

⁷ Sarah Cook, "All the News Unfit to Print: What Beijing Quashed in 2016," *Foreign Policy*, December 16, 2016.

⁸ Information on website blockage is provided by greatfire.org, an organization that monitors the activities of the Great Firewall. The full Alexa ranking of global news websites can be found at <http://www.alexa.com/topsites/category/News> (accessed December 11, 2016). Not all foreign websites are blocked by the Great Firewall, and not all blockages start at the same time. For example, while Microsoft Bing services remain unblocked by the Great Firewall as of today, IP addresses located in China have been unable to access almost all Google services (including Google search, Gmail, Google Scholar, etc.) since 2011.

B. Tools to Bypass Censorship

Access blockage introduced by the Great Firewall can be bypassed through proxy servers or traffic data encryption (e.g., the virtual proxy network, or VPN). This has led to a range of censorship circumvention tools and services to aid internet users in China in gaining access to websites blocked by the Great Firewall.

There are more than a dozen tools for bypassing censorship available to Chinese internet users. As of beginning of our field experiment, there was *no* law in China that explicitly regulates the use of VPN and similar services in China. In fact, as Roberts (2018, p. 110) describes, “[b]ecause the government focuses control on gatekeepers of information, rather than individuals, from the perspective of an ordinary citizen in China the information control system poses very few explicit constraints.” However, the enactment of the Cybersecurity Law in late 2016 indicates that the Chinese government may begin to take measures to regulate the VPN market in the near future.

The prices of censorship circumvention tools range from free of charge to US\$25 per month, fairly inexpensive even from a college student’s perspective. Faster and more stable tools typically charge a premium.⁹ For example, the Chinese government *temporarily* shuts off the connection of some lower-priced VPN services during periods such as the annual March meeting of the People’s Congress. This suggests that while the government is technologically capable of thoroughly disrupting censorship circumvention tools, it chooses to neither do so during majority of the days in the year nor target all the tools, presumably because many businesses operating in China rely on these tools to ensure a global internet connection.

Approximately 3 percent of internet users in China regularly purchase tools to bypass censorship (Roberts et al. 2010).¹⁰ As a result, all 10 of the top 10 most trafficked websites in China, as of 2017, are domestic, a much higher ratio compared to that in Hong Kong (4 out of 10), Taiwan (5 out of 10), and South Korea (3 out of 10).¹¹ The low usage of censorship circumvention tools, albeit their relatively inexpensive availability, provides *prima facie* support that citizens may not demand the access to uncensored internet.

We may expect Chinese citizens to be afraid of using tools to bypass censorship and browsing politically sensitive information, regardless of the legal conditions. To the extent that evidence exists, it suggests that Chinese internet users exhibit limited fear in browsing or posting sensitive content (Qin, Strömberg, and Wu 2017). In fact, explicit censorship and salient actions to block information access can backfire: information consumption may actually increase, similar to the so-called “forbidden fruit effect” (Hobbs and Roberts 2018). When a subgroup of students in our study recognized that *The Economist* magazine was blocked by the Great

⁹ “Circumvention Central,” from greatfire.org, provides reviews of some popular tools. Similar reviews can be found in “Leap Over the Firewall: A Review of Censorship Circumvention Tools” published by *Freedom House*.

¹⁰ The US Congressional Research Report “China, Internet Freedom, and US Policy,” published in 2012, estimates that the usage rate ranges from 1 to 8 percent. Source: <https://fas.org/sgp/crs/row/R42601.pdf> (accessed June 22, 2017). More recently, China Urban Governance Survey (2015) reports that approximately 5 percent of urban residents actively circumvent internet censorship (Roberts 2018). The total number of internet users in China was estimated to be more than 721 million in 2016. Source: *Internet Live Stats*, www.InternetLiveStats.com (accessed December 11, 2016).

¹¹ Source: Alexa top websites by country, <http://www.alexa.com/topsites/countries> (accessed July 15, 2017).

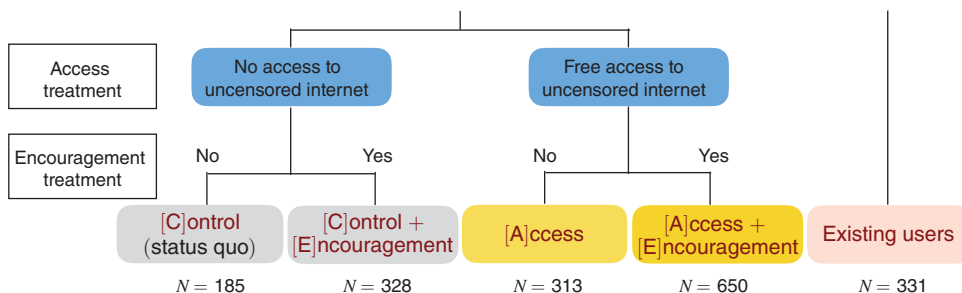
Firewall in April 2016 as a result of its coverage of top Chinese leaders' repressive policies, they spent significantly *more* time on *The Economist*. Importantly, to the extent that there is fear, this is be an important source of the low demand for sensitive information we test for in the experiment.

II. Experimental Design

In Section IIA, we describe our experimental treatments: providing free access to uncensored internet and encouraging visits to foreign news outlets. We describe outcomes related to students' acquisition of uncensored content in Section IIB, and those related to knowledge, economic beliefs, political attitudes, and intended behaviors in Section IIC. Finally, in Section IID, we describe the logistics of the field experiment, discussing the time line, recruitment, treatment assignment, panel surveys, and sample retention.

A. Access and Encouragement Treatments

The experimental design can be summarized in the figure below. Prior to treatment assignment, we identify in a baseline survey those students who have already purchased and are currently using tools to bypass censorship. We exclude these *existing users* in our subsequent treatment assignment, but we follow them throughout the study since they serve as useful benchmarks to interpret the treatment effects. For those who are not existing users of censorship circumvention tools, we randomly assign them to either a *control* condition in which they are subject to censorship as in the status quo, or an *access* (*A*) treatment in which they receive free access to uncensored internet for 18 months. Among a random subgroup of the students who receive the access treatment, we also assign an *encouragement* (*E*) treatment, where we encourage them to visit foreign news websites blocked by the Great Firewall.



To address concerns that the encouragement treatment alone generates an experimenter demand effect (e.g., explicit endorsement of specific foreign websites by the researchers) or changes students' perception of government suppression, we also provide the encouragement treatment to a random subgroup of students in the control group. These students are presented with the same encouragement

treatment material, although they are not able to visit the blocked websites mentioned in the material.

Overall, our experiment creates five groups of students: (i) the control group students [C]; (ii) the control group students who are encouraged to visit foreign news websites [CE]; (iii) students who receive only the access treatment [A]; (iv) students who receive both the access and the encouragement treatment [AE]; and (v) the existing users.

Free Access to Uncensored Internet.—The access treatment provides students with a free 18-month subscription to a censorship circumvention tool. The tool establishes a fast and stable connection to the internet unrestricted by the Great Firewall, enabling students to access websites that are otherwise blocked. The tool does not affect connections to websites that are not blocked by the Great Firewall.

We choose one of the most premium censorship circumvention tools available in China, so that unfavorable features such as slow connection speed that may prevent students from using the tool are kept to a minimum. The tool requires less than one minute to set up, and students do *not* need to sign on each time they wish to browse uncensored internet: the tool operates in each browsing session by default. The tool works on both computers and mobile devices. We provide a full list of the tool's features in online Appendix B.1. An individual account costs US\$25 per month. Although this does not exceed a typical college student's budget, almost all existing users chose cheaper censorship circumvention tools prior to our experiment.

The access treatment is distributed to assigned students in the form of a lottery after they complete the baseline survey. We inform treated students that this tool, while provided for free, is valued at US\$25 per month and is a "professional and secure internet service that allows [one] to browse internet websites around the world without restrictions, access information in a speedy manner; and it is a service adopted by many business enterprises and professionals in China."¹² Treated students are given personal accounts for the tool, and they can activate the service and start the setup process right away, following detailed instructions on the service's website. We limit each account to simultaneously operate on a maximum of two devices in order to prevent multiple students from sharing one account. We, however, cannot rule out the possibility that students lend or sell the entire account to another student.

At the same time, we also randomly draw 100 students to win a one-year VIP-account of Youku (worth approximately US\$30), a Chinese video streaming service similar to Netflix. This serves as a placebo and obfuscates the study's explicit focus on censorship circumvention tools.

¹² Students can learn more about the censorship circumvention tool itself and track their usage status on the service website. We intentionally keep the language introducing the tool vague to avoid political pressure from the school administration. Almost all participants in our study understand what censorship circumvention tools are and know what they are used for: according to our baseline survey, most of them have heard of the tools, or have people in their immediate social circle who have been using them. We communicate with study participants simultaneously, via email and WeChat (equivalent to WhatsApp) messages. In online Appendix B.2, we present the translated email script in which we inform the treated students of the access treatment.

Temporary Encouragement to Visit Foreign News Outlets.—In addition to access, we randomly provide temporary encouragement to visit foreign news outlets in order to examine students' demand for uncensored information. The encouragement consists of a variety of materials mimicking advertisement campaigns to promote foreign news outlets, and it is distributed in the format of biweekly "newsletters" to students' email and WeChat accounts. Students are told that we curate the content of these newsletters to help students stay informed.

We design two phases of the encouragement treatment, in order to distinguish whether demand for foreign news may be raised simply by informing students that foreign news outlets exist, or students need to actually consume content from these outlets. The first phase is purely informational. It consists of four newsletters introducing students to a variety of foreign websites that are blocked by the Great Firewall that students may have never heard of (e.g., the Chinese language edition of the *New York Times*). Moreover, it highlights that politically sensitive news events are often reported differently in domestic news outlets than in their foreign counterparts.

The second phase involves news quizzes with modest monetary rewards. These quizzes aim to encourage students to actually visit the Chinese edition of the *New York Times*, a website on which we focus exclusively to distill the impact of encouragement on students' foreign news consumption. We choose the *New York Times* Chinese edition because it provides extensive coverage on China-related news, offering politically sensitive content unavailable on the domestic media. The website represents one of the highest quality foreign news outlets in Chinese language that are blocked by the Great Firewall, and unlike its English counterpart, the Chinese edition does not impose a paywall. We design the quizzes so that students can locate the answer within a couple of minutes if and only if they visit the front page of the *New York Times* Chinese edition. We implement four rounds of the quizzes, and students earn US\$2.5 if they answer correctly in each round. We set the monetary incentives at a modest level, so that they are unlikely to overcome political fear, potentially a dominant factor that prevents students from browsing foreign news outlets.

The encouragement materials cover many news stories. For example, in one quiz, we ask students what percent of underground water is polluted as reported in an article on the *New York Times* Chinese edition front page. In addition to the underground water pollution, other topics covered in the quizzes include China's wealth inequality, censorship on key economic indicators, and labor unrest. In order to capture students' broad level of knowledge, we quiz students on both the news events explicitly covered in the encouragement materials, and those that are never mentioned. For instance, we measure students' knowledge of the Panama Papers episode, which is never covered in the encouragement treatment.

The encouragement to visit foreign news outlets started in December 2015, simultaneous with the distribution of the access treatment. There is no between-subjects randomization in the order we implement the two phases. Each phase of the encouragement lasts for two months, and we infer how students respond to each phase of the encouragement by changes in their behaviors over time. The temporary encouragement treatment ended in March 2016, six weeks

prior to the midline survey.¹³ Hence, there are two distinct periods during the experiment: (i) from December 2015 to March 2016, the encouragement treatment is in place and the value of visiting foreign news websites is boosted (especially during the second phase of encouragement); and (ii) from March 2016 until the end of the study in April 2017, the encouragement treatment is no longer in place. Online Appendix C provides more details of the encouragement treatment, and online Appendix Figure A.1, A.2, and A.3 present screenshots of the encouragement newsletters.

B. Outcomes: Media

We measure the impact of uncensored internet access on students' media consumption according to the following outcome categories.

Browsing Foreign Websites.—For students in the treatment group who have activated their censorship circumvention tool accounts, we directly observe all of their online activities that route toward websites hosted *outside* of China. We inform students that their online activities are logged as part of the censorship circumvention tool user agreement. Based on approximately 1.5 billion click-level activity logs recorded by the server, we construct the following four key outcome variables: (i) whether a student activates the tool; (ii) whether a student actively uses the service after activation; (iii) total time spent on browsing foreign websites each day once a student has activated the tool; and (iv) total time spent on each foreign website category, such as the Big 4 (Google, Facebook, YouTube, Twitter), news, entertainment, etc.¹⁴

We ensure study participants that the activities are recorded anonymously: rather than students' real identities, participants are only linked to their participant IDs. Approximately 77 percent (62 percent) of the study's male (female) participants use the censorship circumvention tool to browse pornography at least once: many pornographic websites are blocked by the Great Firewall. The prevalence of browsing pornography suggests that the perception of being monitored does not shy students away from content that may be considered socially undesirable.

Importantly, this measurement of information acquisition has two limitations. First, we do not observe the online activities of existing users of the censorship circumvention tools, students in the control group, or those treated students who do not activate the tool. In order to compare the exposure to foreign news websites across all study participants, we repeatedly ask all students to report the frequency with which they visit foreign websites to obtain information throughout the experiment

¹³ We do not explicitly inform students that the last round of the incentivized quiz would be in mid-March. It is unlikely that students continue to regularly visit foreign news websites out of anticipation of future quizzes; they could always visit the *New York Times* to look for answers after they are presented with the specific quiz questions.

¹⁴ We remove "passive" online activities such as those generated by automatic background refreshes, and we remove "inactive" browsing sessions where participants spend more than 30 minutes on a particular web page without any additional activities. To categorize websites, we use the Alexa categorization of domain names, and we manually categorize domains that are not covered by the Alexa database. The category of top foreign news websites consists of top 20 news sites based on Alexa Top Websites rankings, excluding news portals, such as Google News. Source: <http://www.alexa.com/topsites/category/Top/News> (accessed January 8, 2017).

(see online Appendix D, panel A for actual wording of these questions). Second, we do not observe browsing activities on websites hosted domestically, since the censorship circumvention tool need not reroute this online traffic. To assess whether the increased foreign news consumption affects time spent on its domestic counterparts, we ask students to rank the importance of various domestic and foreign media sources to their knowledge of news events.

Decisions to Purchase Access after Experiment Ends.—The free subscription to the censorship circumvention tool expires just before the endline survey, 18 months after the experiment starts. During the endline survey, we offer students who receive the access treatment an opportunity to renew their subscription (out of their own pockets) at a discounted price of US\$4.50 per month. Interested students can subscribe to the service for a minimum of three months, and the service resumes immediately after they pay online. The average monthly price is set at the median level of treated students' elicited willingness to pay for the censorship circumvention tool at the time of the midline survey. In addition, using the same language we use to describe the access treatment, we introduce the censorship circumvention tool to students who have not received the access treatment and to the existing users of similar tools. We offer these students an opportunity to purchase the subscription at the same discounted price. We directly observe the subscription renewal decisions among treated students, and the new accounts created by the students in the control group and those who are existing users of similar tools. Finally, we ask all students to report whether they would purchase tools other than the one that we provide during the endline survey. This setup allows us to compare the access purchase and renewal behaviors across *all* study participants who completed the endline survey, regardless of their treatment status or whether they adopt the access tool we provide.

Beliefs and Attitudes Regarding Media.—To assess the mechanisms behind the treatment effects on browsing foreign media outlets and acquiring uncensored information, we measure participants' attitudes and beliefs regarding media and censorship. Our questions cover (i) valuation of the access to uncensored internet and foreign news outlets (including a Becker-DeGroot-Marschak (BDM) elicitation of willingness to pay); (ii) trust in domestic and foreign news outlets; (iii) belief in the actual media censorship level and its drivers; and (iv) justification for media censorship. In addition, both the baseline and midline surveys elicit students' calibration of reporting bias and censorship depending on the news events' nature.

C. Outcomes: Knowledge, Beliefs, Attitudes, and Behaviors

We examine the effects once students are exposed to politically sensitive information, on a comprehensive set of outcomes. We repeatedly measure four broad groups of outcomes in our panel survey (see online Appendix D, panels B–E, for detailed descriptions). These survey questions are elicited in a private manner, removing various social incentives (such as signaling, coordinating, conforming) that may affect students' answers.

First, we assess a range of knowledge from contemporary to historical and from politically sensitive to non-sensitive. For example, students are quizzed about

current news events not covered in the encouragement newsletters, and are asked about their awareness of major political protests around the world. Second, we elicit students' economic beliefs in an incentive-compatible manner. Students are asked to guess, for example, China's GDP growth rate and its stock market performance by the end of 2017. They are rewarded with an additional bonus payment of up to US\$3 if their guesses are sufficiently close to the true performance indicators published at the end of the year. Third, we measure a wide range of attitudes that students hold with respect to politics, broadly defined. For example, students are asked to report their trust in various institutions, and to evaluate the Chinese government's performance over the past year. Fourth, we ask students to self-report on a range of past behaviors and intended behaviors for the near future, such as their social interactions discussing political topics, political participation, investment in the Chinese stock market, and plans after graduation.

Whether foreign news exposure affects any or all of the four dimensions hinges on the nature of foreign news outlets. For example, do they convey information, sentiment, or both? In addition, which specific aspects of knowledge, beliefs, attitudes, and behaviors may be affected depend on the particular events that take place during the experiment.

Self-Censorship in Answering Sensitive Questions.—Students may not provide honest answers to survey questions on sensitive political attitudes.¹⁵ Several facts suggest that self-censorship may not be a significant concern. First, more politically sensitive modules appear toward the latter half of the survey. We find that conditional on starting the survey and completing it through the politically non-sensitive module, less than 2 percent of the participants drop out upon seeing these sensitive questions.

Second, we use a modified “list experiment” (or “Item Count Technique”) to explicitly measure participants' degree of self-censorship in expressing distrust toward China's central government.¹⁶ The list experiment provides “cover” for expressing potentially sensitive and stigmatized attitudes (by removing individual-level identification from each answer) and allows one to estimate the attitude's prevalence only at the population level. Hence, we are able to compare estimate of adherence to such attitude from our list experiment elicitation (among a random half of the study participants) to that based on the direct question (among the other random half of the study participants) about the same

¹⁵ We emphasize to students that we are independent of the government. We promise study participants that we are committed to a high level of security, anonymity, and confidentiality with respect to the data that we collect. We assure participants that we will erase all survey data if we are faced with political pressure to share the data with government or school officials.

¹⁶ We adopt a modified version of the standard list experiment. The modification is first introduced by Coffman, Coffman, and Ericson (2017) and subsequently adopted by Cantoni et al. (forthcoming). A random half of the study participants (the “control group”) are asked the total number of statements they agree with among a list of four non-sensitive statements. The other half of the participants (the “covered group”) are asked the total number of statements they agree with, among the same list of four statements plus a sensitive statement of interest. We then append “covered” elicitation with the traditional survey method (namely, the direct question): the control group students in our list experiment setup (those who see four statements instead of five) are asked the politically sensitive questions *directly* in the form of “yes” or “no.” The sensitive political attitude of interest is: “I completely trust the central government of China,” and we report the percentage of “no” as indicators of those who do not completely trust the central government.

attitude, to determine whether any self-censorship exists (due to stigma, fear, or social desirability biases). When respondents are provided with “cover,” we estimate that 69.7 percent of the participants indicate that they do *not* completely trust China’s central government. When asked directly, 68.9 percent of the participants indicate such distrust. These two estimates are statistically indistinguishable from one another (p -value = 0.841), suggesting that the magnitude of self-censorship bias in this domain is small at the time of the baseline survey.

Concerns of Multiple Hypotheses Testing.—Given the large number of survey outcomes we examine, the threat of multiple hypotheses testing and the possibility of false positives could be prominent. We do three things to address such concerns. First, we ask *every* outcome of interest elicited in the baseline survey in the subsequent midline survey. We ask a subset of questions in the endline survey, because of space constraints and the need for other experimental modules. We report the estimated treatment effects based on every question covered in the endline survey, and in online Appendix E, we report treatment effects estimated from every question asked in the midline survey. Second, in order to reduce the number of hypotheses we test, we construct a z -score index variable for each category of outcomes we pre-registered to examine.¹⁷ We standardize each component of the index and sum respondents’ standardized outcomes, weighting each item by the inverse of the covariance matrix of the standardized outcomes (following Anderson 2008). Finally, when we examine individual survey outcomes, we adjust p -values using the multiple hypotheses testing correction procedure with multiple outcomes and treatments (following List, Shaikh, and Xu forthcoming, Remark 3.7) and the false discovery rate (FDR) procedure (following Anderson 2008).

D. Timing and Logistical Details

Recruitment and Baseline Survey (November 2015).—We recruit experiment participants from undergraduate students at two universities in Beijing: one is top ranked and considered the most liberal university in China, the other is ranked slightly lower. We believe that the group of students we study are of particular interest since elite students are core participants of anti-authoritarian movements to challenge the incumbent regime, not only in China but around the globe, and their views are likely to shape Chinese political discourse in the future. Nonetheless, one should be cautious when generalizing our results to other demographic groups in China. On one hand, elite (and often liberally-minded) college students in China are selected to be technologically savvy and intellectually curious. This may lead our estimated treatment effects to be larger than that for other Chinese citizens. On the other hand, many of our study participants come from advantaged backgrounds. They may benefit more from the regime at its status quo, and they are already fairly informed even prior to the experimental intervention.

¹⁷The category-level z -score indices are constructed from z -score indices of all corresponding subcategories of outcomes. We do not construct z -score index for the subcategories of outcomes where all variables are derived from a single survey question (e.g., indicators of top ranks in a single ranking question, as in Category A.7). The index also captures broad changes that are only imperfectly measured by any single survey question.

Elite students could also face tighter time constraints due to school work. These imply that our estimated effects may actually be smaller than the average effect among all Chinese citizens.

Recruitment is implemented via email and WeChat messages, and we end the recruitment process once the goal of 1,800 eligible study participants has been reached.¹⁸ Potential participants are informed that this is an academic research project that aims to understand Chinese college students' beliefs, attitudes, and behaviors during the age of globalization, and it involves repeated surveys over the course of 18 months. The provision of the censorship circumvention tool (or, the internet more broadly speaking) is never mentioned during the recruitment process, which assuages concerns about sample selection based on students' interest in uncensored information *a priori*. The baseline survey takes about 90 minutes to complete, and students are paid US\$15 for participating, and an additional US\$10 bonus payment, on average, depending on their survey answers. In addition to pre-treatment levels of outcomes of interest, the baseline survey collects a rich set of demographics and background characteristics, as well as participants' fundamental preferences including risk preferences, time preferences, and social preferences (see online Appendix D, panel F for details). They serve as the basis for experimental balance checks and the criteria for heterogeneity analyses.

In total, we successfully recruit 1,807 study participants who complete the baseline survey (see online Appendix Table A.1 for summary statistics). Among them, 1,490 are from the elite university (or 15 percent of its undergraduate population), and 317 from the lower ranked university (or 3 percent of its undergraduate body).

Treatment Assignment (December 2015).—After we conclude the baseline survey, we distribute the access treatment and the first encouragement newsletter simultaneously. This treatment assignment stage excludes 331 study participants, 22.0 percent of the students at the elite university and 3.4 percent at the lower ranked university, who have been using (any) censorship circumvention tools. We randomly assign two-thirds of the 1,476 non-existing users to the access treatment, and cross-randomize another two-thirds to receive the encouragement treatment. The one-to-two ratio is chosen to maximize power, accounting for the potential low take-up rate for the access and encouragement treatments. We stratify the randomization at the university-gender-cohort level.

¹⁸ We restrict participation eligibility to full-time registered undergraduate students, and who are citizens of the People's Republic of China. Both universities offer a comprehensive set of undergraduate majors and academic programs. In order to protect the study participants, we conceal the identity of these two universities per IRB arrangement. Since the universities' administration prohibits campus-wide mass email, we deploy a combination of department level mass email and informal social recruitment via class heads. Our recruitment message reaches all undergraduate students at the elite university, which constitutes our sampling frame. We face severe political pressure when implementing the study at the lower ranked university, and we terminate the planned recruitment effort before it has been fully rolled out. Hence, it is difficult to estimate the exact number of students from this university eventually reached by our recruitment message. As a result, our sample is *not* representative of the university population by cohort and by major. This does not threaten the internal validity of our findings, but should be kept in mind nonetheless. We analyze the treatment effects heterogeneity along characteristics such as students' gender, age, and college major in Section IVA.

Midline (May 2016) and Endline Surveys (April 2017).—We invite all participants to a midline survey and an endline survey 6 months and 18 months after we distribute the treatment, respectively. Each survey takes approximately 60 minutes to complete, and students are rewarded US\$20 for participation, with an additional US\$10 bonus payment, on average, depending on their survey answers. A total of 1,617 students complete the midline survey, and 1,372 students completed the endline survey (see online Appendix Tables A.2 and A.3 for summary statistics). These 1,372 students constitute the paper's main sample throughout.

Attrition and Balance Checks.—The attrition rate is 24.1 percent between the baseline and endline surveys. The relatively high attrition is likely due to the fact that we could not collect students' dorm room address and hence communication with study participants is restricted to online methods. We do not think attrition severely biases treatment effects estimations. First, survey attrition does not affect tracking of treated students' online activities, so long as they do not actively uninstall the censorship circumvention tool. Second, there is little evidence of selective attrition: (i) attrition rate does not differ by treatment status (p -value = 0.782); (ii) we cannot reject the null hypotheses that the baseline sample and endline samples are identical across the main demographics, background characteristics, and fundamental preferences examined in Table 1 (column 1 presents summary statistics of baseline participants; column 2 presents those for endline participants; column 3 reports the p -values of t -tests comparing their means; and online Appendix Table A.4 provides comparisons of all baseline survey variables); and (iii) none of the five outcome categories measured at baseline significantly predicts attrition, and importantly, the magnitudes of these associations are small (online Appendix Table A.5 reports results where we predict whether a participant attrits from the endline survey with a z -score index summarizing all variables in a given outcome category, her treatment status, and their interaction). Finally, we present three robustness exercises: (i) results on acquisition of sensitive information is robust if we exclude subjects attrited in the endline survey; (ii) bounds of estimated treatment effects taking into account of attrition are constructed (following Lee 2009); and (iii) majority of the results are robust if we re-estimate the treatment effects using midline survey with a less attrited panel sample, although we acknowledge that outcomes of interests such as renewal of censorship circumvention tool are not elicited in the midline survey.

Among the 1,372 study participants who have completed the endline survey, Table 1 presents the summary statistics of those who are existing users (column 4), and those assigned with each of the four treatment groups separately (columns 5–8). For each characteristic, we conduct an analysis of variance (ANOVA) test against the null hypothesis that students across the four experimental treatment groups are not jointly different from each other, and we report the p -value in column 9. Existing users are from households significantly richer and more politically connected than those who have not purchased such tools prior to the treatment assignment. By contrast, members of the four experimental treatment groups are statistically indistinguishable from one another, in regards to 13 out of 17 characteristics examined. The imbalance is driven by imperfect randomization rather than

TABLE 1—SUMMARY STATISTICS, ATTRITION, BALANCE TESTS

Variables	Sample and attrition			Treatment balance					
	Baseline Mean (SD) (1)	Endline Mean (SD) (2)	<i>t</i> -test <i>p</i> -value (3)	Exg. users Mean (4)	<i>C</i> Mean (5)	<i>CE</i> Mean (6)	<i>A</i> Mean (7)	<i>AE</i> Mean (8)	ANOVA test <i>p</i> -value (9)
Male	0.559 (0.497)	0.562 (0.496)	0.843	0.630	0.512	0.579	0.597	0.517	0.119
Birth year	1995.8 (1.278)	1995.8 (1.262)	0.696	1995.9	1995.9	1995.8	1995.7	1995.8	0.252
Height	170.1 (9.311)	169.9 (8.951)	0.576	170.6	168.0	170.3	170.4	169.6	0.063
Han ethnicity	0.912 (0.283)	0.914 (0.280)	0.844	0.921	0.898	0.896	0.922	0.920	0.585
Born in coastal province	0.417 (0.493)	0.415 (0.492)	0.918	0.438	0.372	0.398	0.471	0.398	0.168
Resided in coastal province	0.444 (0.497)	0.439 (0.496)	0.753	0.474	0.358	0.413	0.512	0.420	0.017
Urban hukou prior to college	0.784 (0.412)	0.771 (0.420)	0.402	0.835	0.752	0.764	0.705	0.781	0.150
Religious	0.066 (0.248)	0.066 (0.249)	0.958	0.050	0.058	0.093	0.061	0.065	0.430
Member of CCP [at baseline]	0.068 (0.252)	0.064 (0.245)	0.660	0.058	0.044	0.066	0.057	0.076	0.549
Educational background	0.000 (1.000)	−0.028 (0.980)	0.441	0.086	−0.140	−0.112	0.157	−0.109	0.002
English ability	0.000 (1.000)	−0.023 (0.986)	0.512	0.008	−0.094	0.045	−0.089	−0.019	0.383
Oversea travel experiences	0.000 (1.000)	−0.029 (0.988)	0.416	0.208	−0.158	−0.048	−0.152	−0.035	0.284
Household characteristics	0.000 (1.000)	−0.018 (0.967)	0.604	0.010	−0.067	−0.077	0.068	−0.035	0.284
Risk preferences	0.000 (1.000)	−0.017 (1.008)	0.639	0.109	−0.011	0.031	−0.023	−0.101	0.381
Time preferences	0.000 (1.000)	−0.003 (0.995)	0.930	0.057	−0.018	−0.044	−0.025	0.004	0.939
Altruism	0.000 (1.000)	0.006 (0.995)	0.874	−0.012	0.011	−0.006	0.118	−0.041	0.223
Reciprocity	0.000 (1.000)	0.014 (1.007)	0.702	−0.071	0.246	−0.007	0.005	0.006	0.076
Number of observations	1,807	1,372	—	242	137	259	244	490	—

Notes: Mean level of each characteristics are reported in column 1 for all participants who have completed baseline survey in November 2015 (and corresponding standard deviation in parentheses), column 2 for participants who completed endline survey in April 2017 (and corresponding standard deviation in parentheses), column 4 for endline participants who use censorship circumvention tools prior to the baseline survey, column 5 for endline participants in the control group (*C*), column 6 for endline participants in the control + encouragement group (*CE*), column 7 for endline participants in the access group (*A*), and column 8 for endline participants in the access + encouragement group (*AE*). All characteristics in the “personal characteristics” (category F.1 in survey) are presented. Characteristics in “educational background” (category F.2), “English ability and overseas travel experiences [at baseline]” (category F.3), “household characteristics” (category F.4), and “fundamental preferences” (category F.5) are summarized by *z*-scores. For each characteristic, a *t*-test is conducted against the null hypothesis that students who have completed baseline survey and those who completed endline are not different from each other in term of this characteristic; column 3 reports the corresponding *p*-value for each test. For each characteristic, an ANOVA test is conducted against the null hypothesis that students in the control, control + encouragement, access, and access + encouragement groups are not jointly different from each other in term of this characteristic; column 9 reports the corresponding *p*-value for each test.

differential attrition across treatment groups, and the estimated treatment effects are robust to controlling for all the imbalanced characteristics.

III. Does Access Increase Acquisition of Sensitive Information?

A. Free Provision of Access Alone Does Not Increase Acquisition

Only 55 percent of the students who receive the free access to uncensored internet actually activate the tool, despite repeated reminders (online Appendix Figure A.4 shows the cumulative activation rate during the first six months of the experiment). The low activation rate is unlikely to be an artifact of the treatment distribution modes, because 86 percent of the students who randomly receive the free Youku VIP (similar to Netflix) account via email and WeChat messages at the same time choose to activate that account within a week. Furthermore, 27 percent of students who activate the tool are not actively using the tool (defined as a student using the tool on more than 40 days after the encouragement treatment ends; robust to alternative definitions; see column 1 in panel A of Table 2). This is very likely the result of deliberately choosing to uninstall the tool.

These *Group-A* students spend virtually no time on browsing foreign news websites throughout the experiment (column 1 of panel B; assuming that students without activated accounts spend zero minute on these websites). This is true even among the positively selected subgroup who actively use the tool to bypass censorship (column 1 in panel C; online Appendix Table A.6 reports predictors of activating the censorship circumvention tool, and online Appendix Table A.7 reports predictors of active usage of the tool). The pattern is very similar if we restrict attention to students who completed endline survey (online Appendix Table A.8). Overall, less than 5 percent of the students who actively use censorship circumvention tool regularly browse foreign news websites (if they visit these websites more than twice per week on average; online Appendix Figure A.5 shows the cumulative density plot of the average number of days and total minutes per week students visit the *New York Times*). Moreover, we do not observe a trend in which active users gradually start to browse any foreign news websites even months after the tool is distributed, as we trace the weekly time they spend on the *New York Times* (Figure 1, solid blue line) and all top foreign news websites (online Appendix Figure A.6). Finally, there is little evidence that these *Group-A* students acquire politically sensitive information from other foreign sources. The total time they spend browsing foreign websites is uncorrelated with the occurrences of politically sensitive events (p -value = 0.552), measured as the share of articles published on the *New York Times* Chinese edition that report such events.¹⁹

Taken together, these results demonstrate that the free access to uncensored internet alone has little effect on students' acquisition of politically sensitive information from foreign news outlets. Students' demand for sensitive information may be low to begin with, which we investigate next.

¹⁹ On average, *Group-A* students who actively use the tool spend 79.2 minutes per day browsing foreign websites. The top four websites they spend time on are: Google and related services such as Google Maps and Gmail (17.5 minutes per day, or 22 percent of daily browsing time), YouTube (9.1 minutes per day, or 11 percent of daily browsing time), Facebook (7.7 minutes per day, or 10 percent of daily browsing time), and Twitter (7.1 minutes per day, or 9 percent of daily browsing time). Since these websites are encrypted, we observe neither the search inquiries on Google, nor the specific URLs that students click through.

TABLE 2—BROWSING ACTIVITIES ON FOREIGN WEBSITES

	Access		Access + encouragement		<i>p</i> -value (5)
	Mean (1)	SD (2)	Mean (3)	SD (4)	
<i>Panel A. Extensive margins (percent of students), among all students</i>					
Activated accounts	54.6%	49.9%	68.2%	46.6%	<0.001
Active users	39.6%	49.0%	45.5%	49.8%	0.083
<i>Panel B. Intensive margins (minutes per day), among all students</i>					
Total daily browsing time	31.45	64.99	31.87	59.14	0.922
Google and related services	6.96	13.76	7.09	13.29	0.885
YouTube	3.64	8.79	4.46	11.44	0.263
Facebook	3.09	7.28	3.27	7.83	0.732
Twitter	2.79	7.29	2.96	7.70	0.742
Top foreign news websites	0.10	0.23	0.59	0.65	<0.001
<i>New York Times</i>	0.07	0.18	0.56	0.61	<0.001
Informational websites	2.98	5.84	3.17	5.51	0.612
Wikipedia	0.05	0.19	0.54	1.78	<0.001
Entertainment websites	9.07	15.42	8.90	14.53	0.871
Pornographic websites	2.44	8.13	2.52	8.74	0.898
<i>Panel C. Intensive margins (minutes per day), among active users</i>					
Total daily browsing time	79.17	83.13	69.59	71.20	0.232
Google and related services	17.50	17.16	15.22	16.31	0.198
YouTube	9.12	12.06	9.75	15.37	0.686
Facebook	7.74	9.92	7.14	10.36	0.587
Twitter	7.05	10.22	6.51	10.36	0.625
Top foreign news websites	0.25	0.30	1.19	0.43	<0.001
<i>New York Times</i>	0.18	0.25	1.13	0.38	<0.001
Informational websites	7.50	7.24	6.72	6.59	0.284
Wikipedia	0.14	0.29	1.19	2.48	<0.001
Entertainment websites	22.79	16.97	19.22	16.37	0.044
Pornographic websites	6.16	12.03	5.49	12.32	0.613

Notes: Panel A shows the composition among students who received only the access treatment (*Group-A*) and those who received both access and encouragement treatments (*Group-AE*). They are divided into two nested categories: (i) *Activated accounts*, students who have activated the censorship circumvention tool provided during the experiment, as of April 10, 2017 (the last day of the experiment); and (ii) *Active users*, students who have activated the tool and were actively using the tool (if the student's account records at least one browsing activity per day for more than 40 days after the encouragement treatment ends). Panel B shows the average daily browsing time in total and on various categories of websites throughout the experiment after the encouragement treatment ends, among all students (assuming students without activated accounts spend zero minutes on these websites). Panel C replicates panel B, but among students who actively used the tool. Top foreign news websites, informational, entertainment, and pornographic websites are defined primarily based on Alexa Top Websites categorization. Column 5 shows *p*-values of two-sided *t*-tests on the extensive margins and the intensive margins between the *Group-A* and *Group-AE* students.




FIGURE 1

Notes: Average total browsing time (minutes) on the *New York Times* per week, among students received only the access treatment (*Group-A*) and those who received both access and encouragement treatments (*Group-AE*). *New York Times* browsing time includes both its English and Chinese websites. Dotted line (y-axis on the right-hand side) indicates the proportion of articles published on the *New York Times* that are politically sensitive during that corresponding week.

B. Temporary Encouragement Boosts Immediate Information Acquisition

When we encourage students to consume uncensored information from foreign news outlets, they respond. Students in *Group-AE* are 14 percentage points (25 percent) more likely to activate the censorship circumvention tool, compared to those who are only given the access treatment (panel A of Table 2; see online Appendix Figure A.4 for divergence in cumulative activation rates).

More importantly, the financial incentives to visit the Chinese edition of the *New York Times* increases the time students spend on this outlet during the encouragement period. The solid red line in Figure 1 traces the average weekly time spent on the *New York Times* among *Group-AE* students who activate the tool. Small monetary incentives during the second phase of the encouragement treatment increases these students' time spent on the *New York Times* to 5.6 minutes/week during that period. Similar pattern is observed in the extensive margin as we trace the percentage of students who regularly browse the *New York Times* over time (see online Appendix Figure A.7).

Students' lack of response during the first phase of the encouragement treatment, which is purely informational, suggests that ignorance of foreign news outlets and their whereabouts is unlikely to be the primary reason that students do not demand access to uncensored information. In addition, their responsiveness to financial incentives of modest magnitudes during the second phase suggests that political fear is unlikely a dominant reason they choose not to browse sensitive information. Note that the contrasting responses to first and second phases of the encouragement are unlikely to be driven by general shifts in students' demand for foreign news that precisely coincide with the timing we switch phases, because we do not observe a sharp change among *Group-A* students.

C. Increase in Information Acquisition Persists after Encouragement Ends

When the four-month encouragement treatment ends, the increase in students' information acquisition from foreign news websites persists (Figure 1). The encouragement treatment increases the browsing time on the *New York Times* by 3.4 minutes/week, among all students receiving the access treatment (assuming that those without activated accounts spend zero minutes on the *New York Times*). Similar pattern is evident among active users only, despite the negative selection of the marginal active users in *Group-AE* (increased by 6.7 minutes/week; see panels B and C of Table 2; also see online Appendix Table A.8 for similar comparisons among those who completed endline survey). This is corroborated by students' self-report, as *Group-AE* students are significantly more likely to state in the endline survey that they visit foreign news outlets more frequently to obtain information (online Appendix Figure A.8 and online Appendix Table A.9).

This increase is not driven by the encouragement treatment changing the underlying selection of who browse foreign news websites, since essentially no students spend time on the foreign news websites without the encouragement. Neither is the increase driven by a small number of students, and the comparison of the median student who activate the tool demonstrates an even sharper difference (online Appendix Figure A.5 compares the overall distribution between *Group-A* and *Group-AE* students). Finally, the increase in the *New York Times* browsing time is unlikely a result of students switching away from other foreign news websites, since the total time they spend on top foreign news websites other than the *New York Times* remains very close to zero minutes throughout the experiment.

Although the absolute minutes increased seem low, the raised *New York Times* consumption represents a substantial change. Even the 0.9 million paid subscribers of the *New York Times* in the United States only spend 12.9 minutes/week on average on the website.²⁰ Moreover, the encouragement treatment leads students to seek out information from blocked websites beyond the one we encourage them to visit. In particular, *Group-AE* students begin to spend more time on Wikipedia, as shown in panels B and C of Table 2. We speculate that sensitive news events reported on the *New York Times* prompt students to explore similarly sensitive, censored events in history, of which Wikipedia is a primary source of information. There could be other increases that we do not explicitly categorize. However, interestingly, *Group-AE* students do not begin to regularly visit other foreign news websites, presumably because the value added of browsing a second foreign news website (e.g., the *Wall Street Journal*) is limited after having already visited the *New York Times*. Finally, the increased browsing time in the *New York Times* can be complementary to acquisition of sensitive information from social media platforms. While students in *A* and *AE* groups spend similar amounts of time on Facebook, Twitter, and YouTube, it is likely that *Group-AE* students begin to consume different types of information on these platforms. We unfortunately cannot observe what students browse on these websites since traffic toward them is encrypted.

²⁰ Source: "Social, Search and Direct: Pathways to Digital News" by Pew Research Center 2014, which is based on data collected by ComScore (<http://www.journalism.org/2014/03/13/social-search-direct/>, accessed January 8, 2017). Assuming that an average Chinese reader can read 700 characters per minute, the increased *New York Times* browsing time is approximately equivalent to the time needed to read 17 headlines and news excerpts, or to skim through one medium-length article every weekday.

Group-AE Students Respond to News Shocks.—To test whether *Group-AE* students acquire politically sensitive information from the *New York Times*, we examine the extent to which their browsing time responds to sensitive news shocks. We first measure “news peaks” as the weekly share of articles published on the *New York Times* Chinese edition that report politically sensitive events not covered by domestic Chinese news outlets.²¹ This measure ranges from 26 percent during the fiftieth anniversary of the *Cultural Revolution* in May 2016 to 4 percent during the 2017 Chinese New Year.

We then superimpose this measure on students’ *New York Times* browsing time during the corresponding week (dotted line in Figure 1). They closely track each other. For example, when news on the Panama Papers broke (the week of April 4, 2016) and when President Trump called the President of Taiwan (the week of December 5, 2016), two of the highest news peaks during the experiment, *Group-AE* students increased their weekly browsing time on the *New York Times* by 157 percent and 180 percent compared to their average consumption, respectively. Overall, a 10 percent increase in the share of politically sensitive articles published corresponds to students spending 1.8 more minutes on the website during that week (see online Appendix Table A.10 for regression results). Nonetheless, the extensive margin, measured as the percentage of *Group-AE* students who regularly visit the *New York Times*, does not vary nearly as much as the browsing time. This suggests that *Group-AE* students visit the website at a fairly stable frequency, but spend additional time browsing during the weeks when there are more articles they have not yet seen on domestic news websites, which remain as most students’ primary source of information (see online Appendix Table A.9, panel A).

Encouragement Raises Willingness to Pay for Uncensored Internet Access.—The persistent increase in students’ acquisition of sensitive information may reflect their raised demand for such information, and for uncensored internet access, more broadly. To test this hypothesis, we compare the average level of willingness to pay for any kind of censorship circumvention tools across different groups of students. Figure 2 plots the willingness to pay in US\$/month, repeatedly elicited using a BDM method. As one would expect, at the time of the baseline survey (prior to treatment assignment), existing users are willing to pay 70 percent more for the access to uncensored internet, compared to those students who have not purchased a censorship circumvention tool yet. Students in the *AE* group increase their willingness to pay by US\$1.05/month, or 34 percent, considerably closing the gap with that of existing users by the endline survey.

Group-AE students’ increased demand is also captured by their decisions to purchase uncensored internet access. At the endline survey, we provide all study participants with an opportunity to purchase or renew their subscription to access uncensored internet. Approximately 23 percent of the *Group-AE* students renew their access, and

²¹ For each article published on the *New York Times* Chinese edition, we categorize it as politically sensitive either if it covers the topics explicitly mentioned in censorship commands issued by the Chinese Communist Party’s Propaganda Department (source: *China Digital Times*, collected by the Berkeley Counter-Power Lab), or if a Baidu query of the article title fails to return a relevant news story among the first five pages of the query outcomes.




FIGURE 2

Notes: Average level of willingness to pay for accessing censorship circumvention tools (US\$ per month), elicited using a BDM method, among students in control group (*Group-C*, pooling *C* and *CE* students together), those who received only the access treatment (*Group-A*), those who received both access and encouragement treatments (*Group-AE*), and the existing users, across the baseline survey (November 2015), midline survey (May 2016), and endline survey (April 2017). Sample is restricted to 1,372 students who have completed the endline survey.

they pay on average US\$21.50 up front for a seasonal subscription (darker bars in Figure 3).²² If we count the students who intend to purchase censorship circumvention tools other than the one we provide (lighter bars in Figure 3), then 52 percent of the *Group-AE* students are likely to continue having access to uncensored internet after the experiment, in contrast to their lack of interest 18 months earlier.

While *Group-AE* students' desire to acquire politically sensitive information plays a crucial role in explaining their raised demand for uncensored internet access, sensitive information is not the exclusive reason they decide to continue the access. Nearly 21 percent of the *Group-A* students also intend to renew their access. Since almost no *Group-A* students use the censorship circumvention tool to browse foreign news websites, this suggests that access to Google, social media, and entertainment websites may be a nonnegligible component of *Group-AE* students' raised demand.

D. Why Is Demand for Uncensored Information Low?

Taken together, the evidence presented above suggests that students' low demand for uncensored information is unlikely to be caused by inherent or fixed factors, such as an intrinsic lack of interest in politics or fear of government reprisal. In

²²This is a significantly higher renewal rate compared to that among the *Group-A* students. Less than 1 percent of *Group-C* students purchase the subscription of the tool that we offer. Approximately 8 percent of the existing users take up the offer. They switch from the service that they are currently using, suggesting a combination of search friction (e.g., they may not have heard of this particular premium tool before) and price discounting (e.g., the offered subscription may be cheaper than their current options).




FIGURE 3

Notes: Percentage of students who indicated that they would purchase a censorship circumvention tool at the time of endline survey (April 2017), among students in control group (*Group-C*, pooling *C* and *CE* students together), those who receive only the access treatment (*Group-A*), those who receive both access and encouragement treatments (*Group-AE*), and the existing users. Darker bars indicate the percentage of students who have actually renewed or purchased account of the premium censorship circumvention tool that we offer at a discounted price; lighter bars indicate the percentage of students who report that they would purchase any censorship circumvention tools. Dashed lines indicate percent active users among *Group-A* and *Group-AE* students. This is the same tool that we provide as the access treatment, and the free subscription expires before the endline survey. Sample is restricted to 1,372 students who have completed the endline survey.

fact, students in our experiment are politically engaged and not afraid to consume politically sensitive information. Once students become familiar with a reliable foreign news outlet where they can look for uncensored information, they are willing to spend time browsing articles reporting heavily censored news events.

An important, although not necessarily exclusive, reason students exhibit low demand for uncensored information is their belief that such information is not valuable. While less than 1 percent of the students state, in the baseline survey, that they are unaware of internet censorship in China, they hold considerably diverse beliefs regarding the extent to which content on domestic media is censored. Specifically, a key dimension of this belief is students' assessment of the value difference between foreign and domestic news outlets, and whether the value-added of foreign outlets justifies the cost of access. Among other questions, we ask students the following:

Suppose you have already read about a particular piece of news from a domestic news outlet (e.g. *Xinjin Paper*; *Caijin*; the *Southern Weekend*), how much extra information will you learn if you read news stories from foreign news outlets (e.g., the *New York Times*; the *Wall Street Journal*; the *Financial Times*) in addition?

- 0 = no extra information will be learned;
- 10 = I will learn almost everything from the foreign news outlet.

This dimension of belief on foreign news outlets strongly predicts the usage of censorship circumvention tools prior to the experiment: during the baseline

survey, existing users are more likely to believe that foreign news outlets are high value-added compared to non-users (p -value < 0.001). Moreover, while students in the control group continue to believe that foreign news outlets are not particularly valuable, as Group-AE students become exposed to reputable foreign news outlets, they significantly raise assessment of the value of these outlets (online Appendix Figure A.9, panel A, second graph).

We observe similar patterns across many dimensions of media-related beliefs (online Appendix Figure A.9, panel A summarizes belief subcategories into z -score indices; online Appendix Figure A.8 presents individual dimension of beliefs; and online Appendix Table A.9, panel A presents regression results). For example, Group-AE students become more likely to believe that content on domestic media outlets is heavily censored. Relatedly, they become less likely to trust domestic outlets, and more likely to trust the foreign counterparts. Note that there can be other belief changes such as perceiving content on foreign outlets as more entertaining that we do not explicitly capture in the survey.

If we assume that exposure to foreign news outlets makes beliefs about their quality and value more accurate (as we would expect if students are Bayesian), the patterns above imply that students' beliefs about the quality and value of foreign news outlets are biased downward at the baseline survey.²³ In online Appendix F, we develop a formal model of students' consumption of foreign news outlets using the one-armed bandit problem framework.²⁴ Following this framework, if students sufficiently underestimate the value of foreign news outlets, they may never choose to acquire information from these outlets. However, consumption of information on foreign news outlets would increase *both* during and after the period during which we encourage students to visit these outlets. In particular, one would expect that acquiring information from foreign news outlets during the encouragement period allows students to upwardly update their beliefs regarding the value of these outlets, which would result in a persistent increase in consumption of such outlets. This is precisely what we observe among Group-AE students.

It is important to emphasize that the evidence presented here does not rule out mechanisms other than learning that may explain the persistent increase in the demand for uncensored, sensitive information. For example, students may hold accurate prior beliefs about foreign media, but exposure induces upwardly biased belief updating. Beyond the belief-related mechanisms, visiting foreign news websites may be associated with a one-time, substantial mental cost. In addition, students may procrastinate in setting up the censorship circumvention tools. Another prominent candidate is habit formation that features intertemporal

²³ Downwardly biased beliefs may be an outcome of underexposure to blocked foreign news outlets. Another potential factor is the propaganda campaigns launched by the Chinese state regarding Western news media. In fact, *Foreign Policy* notes that while China ranks among the lowest in terms of media freedom, intriguingly, the conversation among Chinese citizens "regularly centers around perceived media bias elsewhere." Source: <https://foreignpolicy.com/2016/03/04/china-won-war-western-media-censorship-propaganda-communist-party>, (accessed June 20, 2017). Yet another consistent hypothesis is that citizens in authoritarian regimes do not discuss alternative information sources with each other, because they do not know what others believe (Kuran 1997).

²⁴ Armed bandit problems have been extensively used to study technological adoption decisions in many development contexts, and they highlight the process of people learning the value of new and unfamiliar technology. See Foster and Rosenzweig (2010) for a survey of the literature. Lack of consumption of uncensored, costly media outlets is also consistent with rational inattention (e.g., Caplin and Dean 2015).

complementarity in consumption (e.g., Charness and Gneezy 2009 and Hussam et al. 2016). While our experiment is not designed to distinguish between belief-driven and preference-driven models of media consumption, it is nevertheless worth noting that habit formation alone does not necessarily generate the same pattern of belief updating that we document (see online Appendix G for a discussion on habit formation and rational addiction). Yet another hypothesis is gift exchange, since the *Group-AE* students receive additional bonus payment from the experimenter. While we cannot fully rule out the income effect, it is important to note that control group students are indistinguishable from those whom in addition receive the encouragement treatment (*Group-CE*) in terms of their media-related beliefs during midline and endline surveys. This assuages the concern that the belief updating and persistent increase in media consumption among *Group-AE* students is entirely driven by experimenter demand effects (online Appendix Figure A.10).

Finally, the mechanism that citizens learn about unfamiliar media outlets or build habits of browsing could be a generic one. Even in societies absent of explicit censorship, a temporary encouragement may induce persistent increase in news consumption, and our experimental structure can be readily ported to contexts beyond authoritarian regimes.

IV. What Is the Impact of Students Acquiring Sensitive Information?

So far, we have shown that when the access treatment is combined with temporary encouragement, it effectively induces students to acquire uncensored, sensitive information. Does the acquisition of sensitive information affect students' knowledge, economic beliefs, political attitudes, and behaviors? If so, does the impact spill over to others in the social network?

A. Impact on Students Directly Exposed

To measure the impact of uncensored information on those students who are directly exposed, we exploit the variation in their acquisition of such information generated by the experimental treatment. Table 3 presents regression results where we summarize all endline questions in each outcome category pre-registered with an *z*-score index. Panel A represents the *intent-to-treat* effects, as well as bounds on *Group-AE* effect taking into account of attrition; and panel B shows two stage estimates, representing *treatment-on-the-treated* effects, where we regress treatment status on being an active user of the censorship circumvention tool as the first stage. The baseline results are robust to a range of alternative specifications (online Appendix Table A.11).

We report the regression estimates on all individual endline questions in online Appendix Table A.9, where we control for the demographic and background characteristics that are imbalanced across treatment groups, as well as students' answers to these questions in the baseline survey when applicable (see similar regression estimates on all midline survey outcomes in online Appendix E). Online Appendix Figure A.9 presents the results graphically, broken down by subcategories. Online Appendix Figures A.11, A.12, and A.13 show comparisons across all individual questions in the endline survey. For simplicity, we pool control group students with

TABLE 3—TREATMENT EFFECTS ON KNOWLEDGE, ATTITUDES, BELIEFS, AND BEHAVIORS

	A. Media-related behaviors, beliefs, and attitudes (1)	B. Knowledge (2)	C. Economic beliefs (3)	D. Political attitudes (4)	E. Behaviors and planned behaviors (5)
<i>Panel A. Reduced form</i>					
Encouragement only (CE)	0.126 [0.075]	0.066 [0.096]	0.065 [0.103]	0.054 [0.095]	0.105 [0.092]
Access only (A)	0.215 [0.080]	0.119 [0.102]	0.136 [0.100]	0.164 [0.096]	0.176 [0.096]
Encouragement + access (AE)	1.268 [0.067] (1.232–) (1.298+)	0.412 [0.088] (0.371–) (0.446+)	0.573 [0.089] (0.527–) (0.615+)	0.853 [0.086] (0.819–) (0.895+)	0.328 [0.084] (0.275–) (0.356+)
<i>Panel B. Two-stage estimates</i>					
Active user of access tool	1.977 [0.133]	0.630 [0.124]	0.901 [0.138]	1.380 [0.130]	0.469 [0.125]
Mean (all non-existing users)	–0.186	–0.108	–0.125	–0.139	–0.090
Standard deviation (all non-existing users)	0.959	0.982	0.997	0.955	0.947
Mean (control group)	–0.811	–0.327	–0.418	–0.556	–0.294
Standard deviation (control group)	0.681	0.896	0.923	0.906	0.850
Mean (existing users)	0.867	0.503	0.584	0.647	0.419
Standard deviation (existing users)	0.681	0.931	0.783	0.952	1.129

Notes: Survey outcomes in each of the A–E categories are summarized by a z-score index, weighting by the inverse covariance of the standardized variables, following Anderson (2008). Panel A shows regression coefficient estimates and robust standard errors of the *Group-CE*, *Group-A*, and *Group-AE* indicators, where *Group-C* is the omitted group. Treatment lower (–) and upper (+) bounds are calculated for the estimated encouragement + access treatment effects, following Lee (2009). Panel B shows two-stage estimates where the we use *Group-CE*, *Group-A*, and *Group-AE* indicators to estimate a first stage on whether students are active users of access tool to browse uncensored internet, defined as those who have activated the tool and were actively using the tool (if the student's account records at least one browsing activity per day for more than 40 days after the encouragement treatment ends). Coefficients are estimated using 1,130 completed endline surveys from students who have not been using censorship circumvention product at the time of baseline survey (November 2015).

and without the encouragement treatment together (labeled as *Group-C*) in these figures and in our discussions below, since these students do not differ in almost any dimension. Online Appendix Figures A.14, A.15, and A.16 present results comparing *C* (unpooled) and *CE* students.

More Informed of Sensitive Events.—Treated students are more informed of current events that are politically sensitive. We administer a set of seven quizzes on such events that occurred within three months of the endline survey (*Category B.2 of the survey*). These events range from President Trump's business in China to the Xinjiang government's surveillance effort of automobiles. None of these events are explicitly covered in the encouragement material. Students in the *AE* group can answer 0.902 more quizzes correctly. Importantly, the quizzes are able to capture knowledge stock: treated students' access subscription terminates just before the

endline survey, making them unable to look up answers on Google or the *New York Times* until they are given the chance to renew the tool in the later part of the survey. In contrast, the acquisition of sensitive information does not enable *Group-AE* students to correctly answer more quizzes for events covered by the domestic media during the same period. Students in all groups are equally likely to know events such as China stopped importing coal from North Korea in response to the newly enacted sanction. This not only indicates that acquiring uncensored information increases knowledge specifically in domains that are otherwise unavailable on domestic news outlets, but also suggests that foreign news consumption does not substantially crowd out attention on domestic news websites.

As newly exposed students realize that sensitive contemporary events remain unreported on domestic news outlets, they may suspect that censored events exist throughout history. Indeed, *Group-AE* students become 42.4 percent more likely to have heard of protest events in Greater China during the past decade (e.g., the Hong Kong Umbrella Movement in 2014), and 13.7 percent more likely to have heard of foreign protests and independence movements (e.g., the Arab Spring in 2011), all of which are highly politically sensitive and tightly censored (*Category B.3*). A likely source of such information is Wikipedia, of which we observe an increased consumption among *Group-AE* students. As a placebo, we ask students whether they have heard of the “Tomorrow Revolution,” a fake protest we create. The overall proportion of students indicating that they have heard of this event is indistinguishable from zero.

Finally, acquiring uncensored information also affects students’ assessments of their own informedness (*Category B.5*). We find that *Group-AE* students become more likely to consider themselves better informed of political issues in China in absolute terms. Interestingly, when comparing themselves to their peers, *Group-AE* students also become more optimistic about other students’ level of informedness.²⁵

More Pessimistic about China’s Economic Performance.—When students are asked to guess China’s GDP growth rate in 2017 in an incentive-compatible manner, those in the *AE* group believe that the actual growth rate would be 5.92 percent (0.90 percentage points lower than that of the *Group-C* and *Group-A* students). This is a substantial decrease in optimism, and it falls below the government’s explicit target (6.50 percent) and predictions by the Chinese Academy of Social Sciences (6.60 percent).²⁶ Moreover, *Group-AE* students lower their predictions of the closing level of the Shanghai Stock Composite Index at the end of 2017 by 317.3 index points (to 3,046.2; the closing level was 3,154.7 at the time when students made predictions). They actually become slightly too pessimistic given that the realized year-end closing level was 3,307.2, albeit the precise welfare implication remains difficult to assess.

²⁵ We explicitly study beliefs regarding others in a companion paper (Chen and Yang 2018b). In particular, we find that *Group-C* students believe that students who have never used the censorship circumvention tool are equally likely to correctly answer news quizzes compared to those who have been using the tool before the experiment starts. This suggests that an important reason for low demand for uncensored internet access is that students do not realize that uncensored information can make a difference.

²⁶ Source: “Outlook of the 2017 Chinese Economy” by the Chinese Academy of Social Sciences, <http://world.people.com.cn/n1/2017/0502/c190967-29248328.html> (accessed May 10, 2017).

Contrary to the increased pessimism on China, exposure makes *Group-AE* students more optimistic about US economic performance, again elicited in an incentive-compatible manner (*Category C.3*). For example, their prediction about the US GDP growth rate during 2017 is 1.19 percentage points higher than that of *Group-C* students. Many students anchor the United States' growth rate by halving their guess on China, resulting in an average guess (2.69 percent) considerably higher than the past growth rate in the United States (1.60 percent in 2016). Hence, optimism about the United States moves *Group-AE* students further away from the realized growth rate.

Interestingly, we find that while uncensored information significantly affects students' economic beliefs, it barely changes their levels of confidence regarding their own predictions (*Category C.2 and C.4*).

More Skeptical of China's Governance.—Uncensored information changes students' political attitudes. For example, when we ask students to evaluate the government's performance in the realm of economics (and politics) during the past year (on a scale of 0–10, where 10 indicates full satisfaction), newly exposed students in *Group-AE* report a rating 1.254 (and 1.308) lower than that of the students in the *C* and *A* groups. Moreover, the newly exposed students report lower trust toward China's central government by 1.58 (on a scale of 0–10, where 10 indicates complete trust), representing a decrease in political trust of 21.3 percent, compared to that of the unexposed students.

In fact, treatment effects are observed across a broad range of political attitudes. Compared to unexposed students, *Group-AE* students become more likely to believe that both the economic and political systems in China need fundamental changes (*Category D.1*); more likely to express distrust of China's central, provincial, and local governments, and domestic financial institutions, while more likely to state a higher trust of Japanese and the US governments (*Category D.2*); more likely to be unsatisfied with the Chinese government's performance in economic development and domestic politics (while their level of satisfaction in the domain of diplomatic affairs is unchanged) (*Category D.3*); more likely to consider living in a democratic society important (*Category D.6*); and slightly more likely to state that they are willing to battle illegal actions conducted by the government and to stand up to fight for the weak (although unchanged in terms of their willingness to report the government's misconduct) (*Category D.8*).

Changes in Behaviors and Planned Behaviors.—Finally, acquiring uncensored information leads to changes in some of the self-reported behaviors. Compared to unexposed students, *Group-AE* students become more likely to report that they discuss political topics with other students, an increase by 0.67 on a scale of 10 (*Category E.1*); and more likely to report, among the 4 percent of students who were invested in the Chinese stock market, that they have pulled investments out (*Category E.3*).²⁷ *Group-AE* students, however, are *no* more likely to report

²⁷ We do not know what portfolio students hold prior to the experiment. If we assume that the average student holds a portfolio that tracks the Shanghai Stock Index, then pulling out of stock market makes students better

participating in various political activities, such as protests concerning social issues, and voting for the local People's Congress Representatives (*Category E.2*).

Uncensored information also affects newly exposed students' future plans. *Group-AE* students are 13.5 percentage points more likely to plan on applying to overseas graduate schools and hence leaving China in the near future, a substantial increase compared to the 21.1 percent of students in *C* and *A* groups who report having such plans (*Category E.4*). Exposure to uncensored information also makes students more likely to prefer foreign cities for future work and residence, although they do not change the sectoral preferences of their careers (*Category E.5*). These results suggest that uncensored information primarily leads to plans on exiting, rather than actively engaging with the regime.

Magnitude of the Effects.—The magnitude of the treatment effects is specific to the study sample (e.g., the elite student population), the time frame (e.g., 2016 and 2017 are two unusually eventful years), the specific foreign news outlet we encourage students to visit, and the relatively small scale of the experiment. With this in mind, the *local* effects of acquiring politically sensitive information that we identify are substantial in magnitude. The two stage estimates, shown in Table 3, panel B, suggest that actively browsing foreign uncensored websites increases students' knowledge on politically sensitive events by 0.63 of a standard deviation, for instance. Such changes significantly close the gaps between students newly exposed to uncensored information and the existing users in terms of their knowledge, economic beliefs, political attitudes, and behaviors. Nonetheless, convergence with existing users does not necessarily mean converging toward truth.

Another way to quantify the magnitude of the treatment effects is to measure the quantile movement of a median *Group-AE* student (online Appendix Table A.12). If we rank students across all dimensions of the outcomes of interest, we find that the median *Group-AE* student is ranked at the forty-seventh percentile of the distribution of all study participants at the baseline survey, before the experiment starts. The treatment has moved these students to the fifty-sixth percentile of the distribution by the endline survey.

Yet another way to benchmark the effects is to compute the “persuasion rate” (following DellaVigna and Kaplan 2007), which indicates the estimated percentage of students who do not initially hold, say, skepticism toward the Chinese government (“uncensored attitude”) but change their attitudes once they are treated. Note that the name of this measure by no means suggests we could distinguish between foreign media “persuading” students or providing students with objective information, since we are unable to benchmark truth. For each outcome of interest, we calculate this as the *treatment-on-the-treated* effect of the access plus encouragement treatments, divided by the share of *Group-AE* students who do not hold “uncensored attitudes” at the time of the baseline survey. We find that the median persuasion rate across all outcomes of interest is 40.1 percent (standard error = 9.17 percent; online Appendix Table A.13). This is considerably larger than the persuasion rates estimated with respect to media in democratic societies, but of a similar magnitude

off. The Chinese stock market was among the world's worst performers in 2016: the Shanghai Composite Index decreased by 12.5 percent in 2016, compared to the Hong Kong Hang Seng Index's 0.6 percent fall.

to those found in authoritarian regimes that typically have highly regulated media markets.²⁸

Treatment Effects Heterogeneity.—To shed light on the mechanisms of exposure effect, we investigate who is more affected by politically sensitive information. We examine treatment effect heterogeneity across all baseline outcome categories, demographic characteristics, and fundamental preferences (online Appendix Figure A.17 plots results from split-sample regressions; online Appendix Table A.14 presents regression results, interacting treatment status with all sub-sample indicators). Due to the large number of heterogeneous effects simultaneously tested and the right-censoring of survey measures, one should interpret these results with caution.

We find four broad patterns. First, acquisition of uncensored information more heavily affects students with low level of baseline knowledge. This suggests that the treatment effects are unlikely to be driven by a “news junkie” type who simply shifts news consumption from domestic to foreign sources once they have access to uncensored internet. Second, treatment generates larger effects among students who hold more optimistic economic beliefs or more favorable political attitudes prior to the experiment. In other words, students predisposed toward China’s governance are actually more affected by information contrary to their priors, although there could be short-run backlash that we are unable to capture. Third, treatment effects are much larger among students whose parents are not members of the Communist Party or from households with lower incomes. These patterns are consistent with the hypothesis that students from relatively more disadvantaged backgrounds potentially lack alternative access to uncensored information, and hence the exposure induced by the experiment could lead to more dramatic shocks to their knowledge, beliefs, attitudes, and behaviors. Note that differential selection of students from distinct background admitted to elite colleges could also drive treatment heterogeneity. Finally, treatment effects are positive even among students who are, before the experiment, already fairly informed, less predisposed toward China’s governance, or from more advantaged background. This indicates that direct access and exposure to uncensored internet cannot be fully substituted by alternative access to information.

B. Social Spillover of Politically Sensitive Knowledge

Does the acquisition of uncensored information affect students beyond those who are directly exposed? The rate of information transmission allows us to adjust the naïve estimates of direct effects, which are downwardly biased when there is a social spillover. More importantly, it enables us to assess whether a small number of

²⁸ For example, DellaVigna and Kaplan (2007) estimate a persuasion rate from Fox News of approximately 3–8 percent. Enikolopov, Petrova, and Zhuravskaya (2011) find a persuasion rate of 65 percent regarding the impact of opposition messages from Russian TV stations on voting to the pro-government party. The persuasion rate we find in this paper is also larger than that documented with regard to school curriculum among a very similar demographic group (Cantoni et al. 2017). A noticeable exception concerns the impact of East Germans watching TV from the West, as Kern and Hainmueller (2009) document that they became more supportive of the East German regime.

informed students are sufficient to spread uncensored information to a majority of the student population.

We focus on the network of college dorm roommates.²⁹ The dorm room network, albeit by no means a complete mapping of social network, is closely aligned with the “conversation networks,” which play a dominant role in information transmission among university students as demonstrated by Mobius, Phan, and Szeidl (2015). Overall, 57 percent of students without the access and encouragement treatments reside with at least one treated roommate, and 42 percent of treated students reside with at least one other treated roommate.

Simple reduced-form analyses suggest that the social transmission of information is indeed present. We regress students’ likelihood of correctly answering sensitive news quizzes on (i) whether they have access to and actively browse uncensored internet (either as existing users or students in *Group-AE*); (ii) whether they have roommates with access; and the interaction of (i) and (ii). The coefficient on roommates’ access indicate the differential likelihood of correctly answering news quizzes if a control group student has one treated roommate, relative to none. Panel A of Table 4 presents the results on a sample of 3 quizzes as well as the overall correct rate across 11 quizzes; online Appendix Table A.15 shows results on each of the 11 quizzes. Take as an example students’ knowledge of the Panama Papers. Among students who neither receive the access and encouragement treatments nor residing with treated roommates, 56 percent can correctly answer the quiz on the Panama Papers. If the students have one treated roommate, the proportion increases to 78 percent. The increase is more modest among other news events, but still statistically significant when we aggregate all news quizzes. Patterns of social spillover are mixed regarding outcome categories beyond knowledge (online Appendix Table A.16).

To assess whether the magnitude of social spillover is economically meaningful, we next estimate a simple social learning model to quantify the social transmission rate of sensitive information. We consider the probability that a student correctly answer a sensitive news quiz as the sum of (i) the probability that she learns the event from browsing foreign news outlets herself (*direct learning*); and (ii) the probability that she learns about the event from her roommates who have learned about the event (*social learning*). We allow the social transmission rates to differ across news, and across students with and without direct access themselves. Importantly, we assume that there is no information transmission from students who do not have direct access to uncensored internet, since we cannot separately identify the transmission from students with and without access. This assumption is conservative when we evaluate the marginal contribution of having one additional student to receive access to uncensored internet. We use the subsample of students with less than two treated roommates for parameter estimation (panel B of Table 4), and those with two or more treated roommates for an “out-of-sample test” (panel C). Online Appendix H presents the full model and estimation details.

²⁹ A university dorm room in our experimental setting consists of four students from the same gender and cohort, assigned by the school administration. The exact algorithm of dorm roommates assignments is unknown, but they are likely to be randomly assigned within the university-gender-cohort-major cell.

TABLE 4—ESTIMATION OF SOCIAL LEARNING MODEL

Individual sensitive news events					
Direct learning rates:	Lowest	Median	Highest	Overall percentage of quizzes correctly answered	
Poli. sensitive news events:	Steel production reduction reaches target	Carrie Lam becomes HK Chief Executive	Foreign leaders involved in Panama Papers		
	(1)	(2)	(3)		(4)
Panel A. Reduced-form analyses					
Access and active	0.094 [0.056]	0.185 [0.061]	0.332 [0.040]	0.204 [0.022]	
Roommate with access	0.023 [0.056]	0.047 [0.062]	0.222 [0.040]	0.076 [0.022]	
Access and active × roommate with access	−0.001 [0.071]	0.002 [0.078]	−0.114 [0.052]	−0.032 [0.028]	
Panel B. Implied social transmission rates					
Transmission rate (receiver with access)	0.086 [0.206]	0.113 [0.150]	0.248 [0.045]	0.118 [0.034]	
Transmission rate (receiver without access)	0.084 [0.167]	0.117 [0.124]	0.121 [0.038]	0.069 [0.027]	
Panel C. Predictions and out-of-sample tests					
Actual: percent correct (receiver without access)	0.188	0.328	0.918	0.614	
Predicted: percent correct (receiver without access)	0.225 [0.070]	0.317 [0.071]	0.955 [0.053]	0.589 [0.028]	
Actual: percent correct (receiver with access)	0.314	0.490	1.000	0.743	
Predicted: percent correct (receiver with access)	0.317 [0.076]	0.505 [0.081]	1.000 [0.018]	0.733 [0.025]	

Notes: Access and active indicates whether students have access to uncensored internet and actively browse its content; the indicator takes value 1 if the student is an existing user of the censorship circumvention tool prior to the baseline survey (November 2015), or is assigned with both the access and encouragement treatments (*Group-AE*). Roommate with access indicates whether there is one college dorm roommate who actively uses censorship circumvention tool as a result of the experimental treatment. Overall percentage of quizzes correctly answered aggregates all 11 news quizzes together, and use whether roommate receives access by the endline survey (April 2017) in the baseline specifications. Reduced-form analyses and social transmission rates estimation are conducted among students who have completed the corresponding wave of the survey, have no roommates who were existing users of the censorship circumvention tool prior to the baseline survey, and have either zero or one roommates who actively use censorship circumvention tool as a result of the experimental treatment. Out-of-sample tests are conducted among students who have at least two roommates who actively use censorship circumvention tool as a result of the experimental treatment; bootstrapped standard errors are shown in brackets. See online Appendix H for more details.

Three patterns emerge. First, across all 11 news dimensions, the transmission rate from an informed student to her roommates without the access is 11.8 percent. Interestingly, this is a very similar social transmission rate of politically neutral information among Harvard undergraduates (10.3 percent; Mobius, Phan, and Szeidl 2015). The calibrated model performs well when we use it to predict the knowledge of sensitive news events among students with two or more treated roommates. Given the number of students who have access to uncensored internet prior to the experiment, this transmission rate is substantially lower than what is needed to induce the entire student population to be informed. Even if we assume that these existing users are randomly distributed across university dorms, the

social transmission rate needs to be close to 100 percent in order to saturate the population. Second, the social transmission rates monotonically increase with the direct learning rates (see online Appendix Figure A.18, where we plot the estimated direct learning rates against social transmission rates). If a student is more likely to learn about particular sensitive news event herself, she is also more likely to transmit that knowledge to her roommates. Third, the social transmission rate approximately doubles if the recipient is a fellow roommate who has no direct access herself, suggesting social substitutability in learning with respect to direct access to uncensored information.

Several factors may contribute to the relatively low social transmission rate of politically sensitive information. First and foremost, transmission of censored information is asymmetric: the uninformed students do not know what and when to ask. As a result, the informed students need to take the initiative to discuss sensitive topics. Informed students may not take such initiative, because spreading sensitive information, unlike consuming such information privately, may be perceived as politically risky. In addition, the treated and hence informed students demonstrate the “curse of knowledge”: they tend to (mistakenly) believe that other students are equally informed of politically sensitive events (Chen and Yang 2018b). This may further prevent them from taking initiative to transmit the information. Even if informed students decide to spread the information, they face additional constraints. The transmission largely relies on word-of-mouth, since sensitive messages are censored on domestic social media and messaging platforms. Finally, existing users of censorship circumvention tools are highly clustered among a small number of dorm rooms, potentially due to social complementarity.³⁰ Such a high degree of clustering limits the scope of social transmission.

V. What Does It Take to Undo Censorship?

Our experimental results allow us to get a glimpse into what does it take to undo censorship, and what makes China’s censorship apparatus effective, despite its porousness. Would most students become informed of sensitive events if a large share of them receive access to uncensored internet? Would students’ low demand for uncensored information be partially offset by the social transmission of information?

Importantly, our experiment captures partial equilibrium effects, since it covers only 15 percent of the student population of two universities and is not a permanent intervention. On one hand, the general equilibrium treatment effects may be substantially larger because political fear could be eliminated, foreign news consumption becomes more socially acceptable and desirable, or the knowledge of uncensored internet access is disseminated socially. Supply of sensitive information on domestic and foreign outlets could also change in response to the demand shifts due to a large number of citizens receiving access. On the other hand, the general

³⁰ At the time of the baseline survey, 74.8 percent of the existing users have at least one other roommate (31.1 percent have two or more) who are also currently using censorship circumvention tools. In contrast, among students who have not used censorship circumvention tools, only 26.7 percent of them have at least one other roommate (and 9.3 percent with two or more) who are existing users.

equilibrium effects may be smaller in magnitude, as the censorship apparatus itself may respond. For example, providing access to the entire population is likely to trigger a crackdown from the regime.

To illustrate, we simulate the share of students who would be able to correctly answer quiz question on the Panama Papers, a particularly high-profile, sensitive news event, if the number of students who actively acquire uncensored information increases. In addition to learning directly from foreign outlets, we incorporate information transmission among dorm mates. Specifically, we first simulate the number of information-acquiring roommates with whom a particular student may reside. We predict whether each student could correctly answer the quiz if there is no social transmission of knowledge. We then predict the number of roommates each student resides with who have acquired the knowledge about the Panama Papers. Finally, we predict the chance each student could correctly answer the quiz if her informed roommates transmit the knowledge. We calibrate the simulation using three sets of previously estimated parameters: (i) the degree to which access to uncensored internet and the temporary encouragement lead students to acquire sensitive information; (ii) the degree to which direct exposure affects students' knowledge; and (iii) the rate of social transmission of such knowledge. Online Appendix I provides details on the simulation.

We simulate three scenarios, presented in Figure 4. First, if we were to provide *all* students with free access to uncensored internet, we could increase the share of students who answer the quiz question correctly by 3 percentage points, a change that is nearly negligible. Most students would not avail themselves of the cracks of the Great Firewall to actively seek out uncensored information. Second, if we were to provide all students with both access and the same temporary encouragement in our experiment, the share of students who could correctly answer the quiz question would jump by another 30 percentage points to 98 percent, close to the full-saturation. Finally, this large increase could be sustained even if we stop fully subsidizing the access to uncensored internet. After students receive the encouragement, their raised demand for access would likely persist, and we expect 72 percent of the newly exposed students would pay to continue accessing uncensored internet. As a result, the overall share of students who would be able to correctly answer the quiz question would be retained above 90 percent. The latter two scenarios demonstrate that while the current censorship apparatus is robust due to lack of demand for uncensored information, its effectiveness could be substantially diminished if demand were raised through encouragement and exposure.

To quantify the role played by the social transmission of information, we simulate a scenario with no social spillover (shown in the dotted line in Figure 4). On average, approximately 50 percent of the increase in knowledge among the student population results from direct learning, while the remaining 50 percent is contributed by social transmission among roommates. It is worth highlighting that social transmission plays a particularly small role when only existing users are acquiring uncensored information, due to their high degree of clustering among a small number of dorm rooms.

These results are robust to relaxing the key assumptions we make in the simulation procedure. In particular, when we allow for second-degree social transmission, we do not find quantitatively differences in students' simulated knowledge, given the size




FIGURE 4

Notes: Simulation of the proportion of students who can answer quiz on the Panama Papers correctly across the entire student population, as the proportion of students who actively visit foreign news websites grows from 0 to 100 percent. Details of the simulation procedure is described in online Appendix I.

of dorm rooms and the estimated transmission rates (consistent with Mobius, Phan, and Szeidl 2015). In addition, if we use the conversation networks mapped among Harvard undergraduates to simulate the information acquisition and spillovers in our context, we find quantitatively similar results, presumably because the number of roommates in our context (3.00) is close to the average number of conversational links (3.19) a Harvard undergraduate processes. In fact, our result is robust even if we double the number of conversational links a student has.

VI. Conclusion

Media censorship is prevalent in authoritarian regimes. Little is known, however, regarding whether censorship is effective at restricting citizens' information acquisition and changing their beliefs and attitudes. In particular, one might have speculated that censorship becomes irrelevant in the age of the internet where access blockage becomes increasingly costly and technologically challenging. In this project, we conduct a field experiment among college students in China to examine the impact of accessing uncensored internet. We find that even among students in one of China's most elite and liberal universities, lack of access is not the entire story: low demand for uncensored information is a crucial reason why students don't acquire such information. Beliefs that foreign news outlets are not valuable contribute to the low demand. Importantly, a period of exposure can change these beliefs and result in a lasting increase in their acquisition of uncensored information.

These findings suggest that demand-side factors are important for comprehending how internet censorship works in China today. Depending on citizens' demand for uncensored information, the censorship apparatus in China can be either robust or fragile. After years of censorship and active propaganda campaigns, both the current level of demand and its elasticity with respect to the degree of censorship are low. Coupled with a moderate rate of socially transmitted knowledge, the regime could be highly robust. As a result, policies such as the Lantern Project that passively supply access to uncensored internet to citizens in authoritarian regimes are unlikely to be as effective as some might imagine.³¹ In fact, the Chinese government may not need to bear the extremely high costs of fully "sealing" its internet, as it can afford to leave some holes open. The masses may not begin to respond to negative news shocks, information-demanding elites may not be irritated, and business interests relying on global internet connections may not be sacrificed. Moreover, so long as the Chinese citizens perceive domestic media to be more informative, more enjoyable, or more reputable when feedback is weak due to the Great Firewall, domestic media outlets are incentivized to maintain their censored news reporting in equilibrium (Gentzkow and Shapiro 2006).

Nevertheless, our findings do *not* imply that the Chinese regime can safely eliminate the Great Firewall. The current cost of circumventing censorship imposes a huge campaign cost on foreign news outlets. Without such costs, outlets such as the *New York Times* might begin to campaign and effectively raise demand among Chinese readers. The demand for uncensored information, once raised, is likely to persist and can generate substantial pressure on the censorship apparatus. Perhaps more importantly, as demand for uncensored information shifts, domestic news outlets are likely to react to such shift, which could fundamentally change China's media landscape and further suppress demand for censored information (see Qin, Strömberg, and Wu 2018 for an investigation of China's newspaper industry).

This demand-driven censorship is not unique to contemporary China. The current Russian regime enforces repressive censorship over TV, while leaving the internet, and in particular the social media landscape, largely uncensored. Similarly, during the Cold War, the East German government employed heavy propaganda and censorship campaigns, while simultaneously allowing its citizens to purchase, de facto, antennae to access West German TV if they were sufficiently interested. What is an authoritarian regime's optimal strategy for controlling information given our findings? Would such a strategy work even in small regimes where it is unprofitable for a media company to only serve the small domestic market (e.g., Pan 2017)? If not, what would be the alternative strategy? Do information technology innovations, in particular the internet and social media, change the effectiveness of state censorship? These are fascinating questions for future research.

Finally, we find that uncensored information persistently and substantially changes students' knowledge, economic beliefs, and political attitudes. Do students realize that consuming uncensored information has made a difference? What

³¹ The Lantern Project, funded by the US State Department, aims to provide a stable internet connection that bypasses censorship at a relatively low cost. Policymakers in the West hypothesize that combating censorship boils down to ensuring that citizens have access to uncensored information. For example, former Secretary of State Hillary Clinton declares that the United States "stand[s] for a single Internet where all of humanity has equal access to knowledge and ideas."

happens to their beliefs regarding fellow students? Answers to these questions have important implications on whether coordinated and collective actions may arise, which we explicitly investigate in a companion paper (Chen and Yang 2018b).

REFERENCES

- Abramitzky, Ran, and Isabelle Sin. 2014. "Book Translations as Idea Flows: The Effects of the Collapse of Communism on the Diffusion of Knowledge." *Journal of the European Economic Association* 12 (6):1453–520.
- Adena, Maja, Ruben Enikolopov, Maria Petrova, Veronica Santarosa, and Ekaterina Zhuravskaya. 2015. "Radio and the Rise of the Nazis in Prewar Germany." *Quarterly Journal of Economics* 130 (4): 1885–939.
- Alesina, Alberto, and Nicola Fuchs-Schündeln. 2007. "Good-Bye Lenin (or Not?): The Effect of Communism on People's Preferences." *American Economic Review* 97 (4): 1507–28.
- Anderson, Michael L. 2008. "Multiple Inference and Gender Differences in the Effects of Early Intervention: A Reevaluation of the Abecedarian, Perry Preschool, and Early Training Projects." *Journal of the American Statistical Association* 103 (484): 1481–95.
- Bai, Jie, Mikhail Golosov, Nancy Qian, and Yan Kai. 2015. "Understanding the Influence of Government Controlled Media: Evidence from Air Pollution in China." Unpublished.
- Cantoni, Davide, Yuyu Chen, David Y. Yang, Noam Yuchtman, and Y. Jane Zhang. 2017. "Curriculum and Ideology." *Journal of Political Economy* 125 (2): 338–92.
- Cantoni, Davide, David Y. Yang, Noam Yuchtman, and Y. Jane Zhang. 2019. "Protests as Strategic Games: Experimental Evidence from Hong Kong's Antiauthoritarian Movement." *Quarterly Journal of Economics* 134 (2): 1021–77.
- Caplin, Andrew, and Mark Dean. 2015. "Revealed Preference, Rational Inattention, and Costly Information Acquisition." *American Economic Review* 105 (7): 2183–203.
- Charness, Gary, and Uri Gneezy. 2009. "Incentives to Exercise." *Econometrica* 77 (3): 909–31.
- Chen, Yuyu, and David Y. Yang. 2018a. "Historical Traumas and the Roots of Political Distrust: Political Inference from the Great Chinese Famine." Unpublished.
- Chen, Yuyu, and David Y. Yang. 2018b. "The Coordination Consequences of Media Censorship: Experimental Evidence from China." Unpublished.
- Chen, Yuyu, and David Y. Yang. 2019. "The Impact of Media Censorship: 1984 or Brave New World?: Dataset." *American Economic Review*. <https://doi.org/10.1257/aer.20171765>.
- Coffman, Katherine B., Lucas C. Coffman, and Keith M. Marzilli Ericson. 2017. "The Size of the LGBT Population and the Magnitude of Antigay Sentiment Are Substantially Underestimated." *Management Science* 63 (10): 3168–86.
- DellaVigna, Stefano, and Matthew Gentzkow. 2010. "Persuasion: Empirical Evidence." *Annual Review of Economics* 2: 643–69.
- DellaVigna, Stefano, and Ethan Kaplan. 2007. "The Fox News Effect: Media Bias and Voting." *Quarterly Journal of Economics* 122 (3): 1187–234.
- Dupas, Pascaline. 2014. "Short-Run Subsidies and Long-Run Adoption of New Health Products: Evidence from a Field Experiment." *Econometrica* 82 (1): 197–228.
- Enikolopov, Ruben, and Maria Petrova. 2015. "Media Capture: Empirical Evidence." In *Handbook of Media Economics*, edited by Simon P. Anderson, Joel Waldfogel, and David Strömberg, 687–700. Amsterdam: Elsevier.
- Enikolopov, Ruben, Maria Petrova, and Ekaterina Zhuravskaya. 2011. "Media and Political Persuasion: Evidence from Russia." *American Economic Review* 101 (7): 3253–85.
- Enke, Benjamin. 2018. "What You See Is All There Is." Unpublished.
- Ford, Guy Stanton, ed. 1935. *Dictatorship in the Modern World*. Minneapolis, MN: University of Minnesota Press.
- Foster, Andrew D., and Mark R. Rosenzweig. 2010. "Microeconomics of Technology Adoption." *Annual Review of Economics* 2: 395–424.
- Fuchs-Schündeln, Nicola, and Matthias Schündeln. 2015. "On the Endogeneity of Political Preferences: Evidence From Individual Experience with Democracy." *Science* 347 (6226): 1145–48.
- Gehlbach, Scott, and Konstantin Sonin. 2014. "Government Control of the Media." *Journal of Public Economics* 118: 163–71.
- Gentzkow, Matthew, and Jesse M. Shapiro. 2006. "Media Bias and Reputation." *Journal of Political Economy* 114 (2): 280–316.

- Gentzkow, Matthew, Jesse M. Shapiro, and Daniel F. Stone. 2015. "Media Bias in the Marketplace: Theory." In *Handbook of Media Economics*, edited by Simon P. Anderson, Joel Waldfogel, and David Strömberg, 623–45. Amsterdam: Elsevier.
- Gerber, Alan S., Dean Karlan, and Daniel Bergan. 2009. "Does the Media Matter? A Field Experiment Measuring the Effect of Newspapers on Voting Behavior and Political Opinions." *American Economic Journal: Applied Economics* 1 (2): 35–52.
- Guriey, Sergei, and Daniel Treisman. 2018. "Informational Autocracy: Theory and Empirics of Modern Authoritarianism." Unpublished.
- Hobbs, William R. and Margaret E. Roberts. 2018. "How Sudden Censorship Can Increase Access to Information." *American Political Science Review* 112 (3): 621–36.
- Huang, Haifeng, Serra Boranbay-Akan, and Ling Huang. 2019. "Media, Protest Diffusion, and Authoritarian Resilience." *Political Science Research and Methods* 7 (1): 23–42.
- Huang, Haifeng, and Yao-Yuan Yeh. 2019. "Information from Abroad: Foreign Media, Selective Exposure and Political Support in China." *British Journal of Political Science* 49 (2): 611–36.
- Hussam, Reshmaan, Atonu Rabbani, Giovanni Reggiani, and Natalia Rigol. 2016. "Habit Formation and Rational Addiction: A Field Experiment in Handwashing." Unpublished.
- Kern, Holger L., and Jens Hainmueller. 2009. "Opium for the Masses: How Foreign Media Can Stabilize Authoritarian Regimes." *Political Analysis* 17 (4): 377–99.
- King, Gary, Jennifer Pan, and Margaret E. Roberts. 2013. "How Censorship in China Allows Government Criticism but Silences Collective Expression." *American Political Science Review* 107 (2): 326–43.
- King, Gary, Jennifer Pan, and Margaret E. Roberts. 2014. "Reverse-Engineering Censorship in China: Randomized Experimentation and Participant Observation." *Science* 345 (6199).
- Kuran, Timur. 1997. *Private Truths, Public Lies: The Social Consequences of Preference Falsification*. Cambridge, MA: Harvard University Press.
- List, John A., Azeem M. Shaikh, and Yang Xu. Forthcoming. "Multiple Hypothesis Testing in Experimental Economics." *Experimental Economics*.
- Lorentzen, Peter. 2013. "China's Strategic Censorship." *American Journal of Political Science* 58 (2): 402–14.
- Mobius, Markus, Tuan Phan, and Adam Szeidl. 2015. "Treasure Hunt: Social Learning in the Field." Unpublished.
- Mullainathan, Sendhil, and Andrei Shleifer. 2005. "The Market for News." *American Economic Review* 95 (4): 1031–53.
- Pan, Jennifer. 2017. "How Market Dynamics of Domestic and Foreign Social Media Firms Shape Strategies of Internet Censorship." *Problems of Post-Communism* 64 (3–4): 167–88.
- Prat, Andrea. 2015. "Media Capture and Media Power." In *Handbook of Media Economics*, edited by Simon P. Anderson, Joel Waldfogel and David Strömberg, 669–86. Amsterdam: Elsevier.
- Prat, Andrea, and David Strömberg. 2009. "The Political Economy of Mass Media." In *Advances in Economics and Econometrics, Tenth World Congress*, edited by Daron Acemoglu, Manuel Arellano, and Eddie Dekel, 135–87. Cambridge: Cambridge University Press.
- Qin, Bei, David Strömberg, and Yanhui Wu. 2017. "Why Does China Allow Freer Social Media? Protests Versus Surveillance and Propaganda." *Journal of Economic Perspectives* 31 (1): 117–40.
- Qin, Bei, David Strömberg, and Yanhui Wu. 2018. "Media Bias in China." *American Economic Review* 108 (9): 2442–76.
- Roberts, Hal, Ethan Zuckerman, Jillian York, Robert Faris, and John Palfrey. 2010. *2010 Circumvention Tool Usage Report*. Cambridge, MA: Berkman Klein Center for Internet & Society.
- Roberts, Margaret E. 2018. *Censored: Distraction and Diversion Inside China's Great Firewall*. Princeton, NJ: Princeton University Press.
- Schedler, Andreas. 2010. "Democracy's Past and Future: Authoritarianism's Last Line of Defense." *Journal of Democracy* 21 (1): 69–80.
- Shadmehr, Mehdi, and Dan Bernhardt. 2015. "State Censorship." *American Economic Journal: Microeconomics* 7 (2): 280–307.
- Simonov, Andrey, and Justin Rao. 2018. "What Drives Demand for Government-Controlled News in Russia?" Unpublished.
- Strömberg, David. 2015. "Media and Politics." *Annual Review of Economics* 7: 173–205.
- Voigtländer, Nico, and Hans-Joachim Voth. 2015. "Nazi Indoctrination and Anti-Semitic Beliefs in Germany." *Proceedings of the National Academy of Sciences* 112 (26): 7931–36.
- Yanagizawa-Drott, David. 2014. "Propaganda and Conflict: Evidence from the Rwandan Genocide." *Quarterly Journal of Economics* 129 (4): 1947–94.