Understanding the Influence of Government Controlled Media: Evidence from Air Pollution in China (Preliminary)*

Jie Bai (MIT), Mikhail Golosov (Princeton), Nancy Qian (Yale), Yan Kai (Yale)†

March 26, 2015

This project establishes several empirical facts to better understand the influence of government-controlled media in autocratic regimes. We document that government control and bias on the controversial topic of air pollution are salient; that individuals do not discount repeated information, even when they know the information comes from one biased government source; that individuals have difficulty interpreting conflicting pieces of information; and that the population is heterogeneous in the degree to which they update. These patterns are difficult to reconcile with Bayesian updating and frameworks that assume homogeneous populations.

*We thank Stefano DellaVigna, Matthew Gentzkow, Brian Knight and David Stromberg for many useful insights; Shuaizhang Feng and Min Lu for sharing their valuable resources; Pengtu Ni, Jinwen Wang and Mengzhu Wu for managing, training and assembling a team of over sixty surveyors. We acknowledge the NSF 0748815 grant for funding.

†Corresponding authors: golosov@princeton.edu, nancy.qian@yale.edu.

1 Introduction

This study aims to make progress in understanding the role of media in a context where government control and censorship (media bias) is a fact of life: for example, China, North Korea, Russia. The canonical model used in the media literature assumes Bayesian updating for a homogenous population. This framework stipulates that such media should have no influence on rational individuals. If citizens know that the government controls the media and that the government is biased, then citizens should update and discount government news. Government-biased news will thus have little influence on the opinions of citizens, and therefore, in equilibrium, the government will have little incentive to distort the news. This contradicts starkly with reality, where governments extensively and effectively use biased media to influence the beliefs of citizens. A growing number of studies have provided evidence on this point in the contexts of Nazi Germany (Adena, Enikolopov, Petrova, Santarosa, and Zhuravskaya, 2013), Rwanda during the 1990s (Yanagizawa-Drott, 2014), Serbia (DellaVigna, Enikolopov, Mironova, Petrova, and Zhuravskaya, 2011) and Russia (Enikolopov, Petrova, and Zhuravskaya, 2011). Past works have also provided theoretical and/or empirical evidence on non-Bayesian behavior from the United States or laboratory experiments; for examples, individuals may not discount repeated information (DellaVigna and Kaplan, 2007; Demarzo, Vayanos, and Zwiebel, 2003) or may neglect to account for the incentives of the sender (Eyster and Rabin, 2010).[1] However, there is little direct evidence on how applicable these explanations are to autocratic regimes with a high degree of government censorship.

The general difficulty in explaining the influence of government-controlled information with an alternative framework is that there are multiple explanations for any given set of

[1]DellaVigna and Kaplan (2007) finds that biased media reports from Fox News influences behavior in the United States. U.S. media is not controlled by the government. However, the question of why a news source known to be consistently biased has any influence is conceptually similar to the question we explore in this paper.

behavior. We therefore need additional empirical evidence to discipline and narrow the scope of theoretical explanations. The main goal of this study is to provide such evidence in a systematic way and use experiments to document several important departures from standard frameworks. Specifically, we question two assumptions common in the literature. The first is that updating is Bayesian. The second is that the population is homogenous.

Our study takes place in the context of China, where almost all media has been directly controlled by the government since 1949, and where new sources of non-government information has begun to gradually increase in recent years. We focus on air pollution, which is a subject of significant media attention and highly controversial due to the extremely high levels in recent years.[2] Since we are interested in the perceived risks of air pollution to health, our main outcome variable is the Air Quality Index (AQI). Average AQI in Shanghai for 2013 was 173 (which is categorized as "very unhealthy"). In contrast, Los Angeles, known as one of the most polluted cities in the United States, had an average AQI of 91 (which is categorized as "moderate"). As air pollution increased, so did controversy over the fact that the Chinese government consistently reported AQI levels lower than environmental agencies and other non-governmental sources. Eventually, in 2012, the Chinese government banned non-governmental agencies from measuring AQI.[3] The one exception to the ban is AQI reported by the U.S. embassy, which had been measuring from the roofs of the embassy for scientific purposes. Following the ban, U.S. measurements expanded to several other consulates in China, including Shanghai. We will use U.S. reports in our experiments as a second source of news. This is discussed in more detail later.

The goal of the empirical analysis is to first document that government control and bias are truly salient and some key departures from the standard model of updating. Second, we ask whether individuals discount repeated information. Third, we investigate individuals' ability to interpret conflicting pieces of information. Finally, we examine whether the

[2]For a study of the political importance of air pollution, see for example Jia (2014).
[3]For example, see and Levs (2011) and Na (2012).

population is homogenous.

To answer these questions, we conduct an experiment where volunteer respondents at three elite universities in Shanghai are randomly assigned into seven treatment groups.

For the first question, we assigned a group of students to respond to a survey about their perceptions of government control and bias in Chinese and U.S. news. Almost all respondents reported that they believe news on air pollution from Chinese media outlets to be controlled by the Chinese government, where as only half reported that the U.S. analogue is controlled by the U.S. government. They also believe that Chinese media outlets are much more likely to systematically misreport air pollution than U.S. news. In the case of misreporting, respondents believe that Chinese news understates air pollution, while U.S. news overstates it. In contrast, when asked about a non-controversial topic – temperature, very few respondents believed that either Chinese nor U.S. media outlets systematically misreport temperature. These survey responses suggest that both the government's control over news and its biases are very salient.

For the second and third questions, we develop a simple Bayesian model to derive empirically testable hypotheses on the extent to which individuals are rational and if they are not, where they depart from the rational framework. Specifically, the model predicts that rational individuals should fully discount repeated information and that conflicting information from two sources of news should be interpreted similarly regardless of whether the news are shown separately or simultaneously. The model also shows how we can infer individuals' perceived biases and accuracy of news sources from their reported priors and posteriors.

To test the first prediction, we randomly assign subjects to three groups. Each group is asked to report their prior, shown news, and then asked to report their posterior. Group 1 is shown 1 piece of news from a nominally independent Chinese website. Group 2 is shown two pieces of identical AQI reports from two different such websites. Group 3 is shown the same news as Group 2, except that we also show the sources of each piece of news, which are the same government ministry. We find that the perceived biases of Chinese news is smaller in both Groups 2 and 3 than Group 1. Thus, individuals do not discount repeated information. This is inconsistent with Bayesian updating, particularly since survey respondents (in another randomly assigned group) report that they believe Chinese news on air pollution to be controlled by the government and biased. These results are consistent with the arguments of DellaVigna and Kaplan (2007) and Demarzo, Vayanos, and Zwiebel (2003) that individuals do not discount repeated information, and Eyster and Rabin (2010) that individuals neglect to take into account the incentives of their sender.

To test the second prediction, we randomly assign subjects into another three groups. These groups are shown two pieces of news, one for temperature and one for AQI. Each group is asked to report their prior on temperature, shown the news report of temperature, report their posterior on temperature. We repeat this procedure for AQI. Group 4 is shown news from Chinese media outlets. Group 5 is shown news from U.S. outlets. Group 6 is simultaneously shown news from Chinese and U.S. outlets. Temperature reports from the two outlets are similar, while AQI reports from the U.S. is significantly higher (i.e., the U.S. outlet reports higher pollution levels than Chinese outlets). We estimate the relationship between posteriors and priors and then apply the estimates from Groups 4 and 5 to our simple Bayesian model where individuals are shown U.S. and Chinese news. We then compare the model implied responses to the responses of subjects in Group 6, who are shown these two pieces of news. We find that they are similar for temperature, but significantly different for AQI. From this we conclude that individuals have difficulty interpreting conflicting pieces of information.

Finally, we examine the assumption that the population is homogenous (this assumption is made by both Bayesian and non-Bayesian frameworks). There is substantial heterogeneity. Visually, we observe three groups: approximately one-third of the subjects update to the average of their prior and the reported AQI, while the rest update very little or very close to the reported AQI.[4] We use a FMM model to formally show that there is more than one group. Interestingly, only updating for the controversial subject, AQI, is heterogeneous; updating for the non-controversial subject, temperature, is statistically homogenous.

There are two important caveats for interpreting our results. First, we are testing the short run effects of giving conflicting information. The impact of conflicting independent information over the medium and long run may be very different. Second, our experiment takes place amongst elite university students, who have more exposure and access to non-Chinese media than the average citizen, and at the same time may be more ideological. An interesting avenue for future research is to repeat this experiment with other segments of the population.

The policy implications our the findings are discussed in the conclusion.

This study adds to the large and rapidly growing literature on the influence of the media that is reviewed by DellaVigna and Gentzkow (2010). The most closely related studies are those cited earlier. In addition, we add to the very recent studies on the influence of media in China (Bei, Stromberg, and Wu, 2013; Bei, Larsson, Stromberg, and Wu, 2014).

The results are also related to studies on environmental policy. Pollution is one of the most often discussed and controversial issues in China today. This is not surprising when, for example, air pollution in the United States, which is much lower than in China, has been shown to have detrimental effects on health outcomes such as infant mortality (e.g., Chay and Greenstone, 2003). One naturally wonders whether citizen demand to reduce pollution will increase with the amount of independent news from the internet and Weibo (Twitter in China). Our results suggest that additional news sources will not immediately cause citizens to revise their assessment of pollution upwards.

The paper is organized as follows. Section 2 presents a simple model of Bayesian

[4]Our other results carry through if we restrict the sample to the first group of partial updaters.

updating. Section 3 presents the empirical strategy and describes the experiment. Section 4 discusses the descriptive evidence on salience. Section 5 presents the experimental results. Section 6 documents the heterogeneity in updating. Section 7 discusses the preliminary conclusions.

2 Conceptual framework

The goal of the model is to develop empirically testable hypotheses that allow us to make progress on how individuals interpret news from media outlets when they are known to be controlled by the state and to report news with bias. In particular, we aim to understand the extent to which individuals are rational and if they are not, where they depart from the rational framework.

In our theoretical framework, the individuals observe noisy signals about the state of nature $\bar{\theta} \in R$ and receive potentially biased news reports.

First we describe how priors are formed. We assumed that each individual starts with a flat prior about $\bar{\theta}$ and observes noisy signals about it. We allow for a possibility that individuals' signals are systematically correlated or biased. This captures a situation when individuals may incorrectly perceive that pollution is low during a particularly nice and sunny day. Formally we model it as follows. Each individual i receives a signal consisting of two components: an aggregate component which is observed by all agents and an idiosyncratic component which observed only by agent i. The aggregate component is normally distributed with mean $\bar{\theta}$. Let s be the realization of that aggregate component. The idiosyncratic component is normally distributed with mean s. Let θ_i^{prior} be the belief of agent i after observing the aggregate and idiosyncratic signals. Since signals are normal and the initial priors are flat, θ_i^{prior} is normally distributed with mean $\bar{\theta} + s$. We use τ to denote the precision of this distribution.

A media outlet m observes an unbiased signal s_m about $\bar{\theta}$, which is normally distributed with mean $\bar{\theta}$ and precision τ_m. The media outlet reports this signal with bias b_m, that is the report is $r_m = s_m + b_m$. The value of b_m is common knowledge.

Posterior beliefs of all agents are formed by Bayes rule. Posterior belief of agent i after reading report r_m, θ_i^{post}, should satisfy

$$\theta_i^{post} = \frac{\tau_m}{\tau_m + \tau}(\theta_m - b_m) + \frac{\tau}{\tau_m + \tau}\theta_i^{prior}. \tag{1}$$

If agents are rational, reading *the same* report several times should not affect their posterior beliefs.

The same arguments extend if agents read reports from several different news sources. In particular, supposed that there are two media outlets m and m', and that each outlet independently observes $\bar{\theta}$ with precisions τ_m and $\tau_{m'}$ respectively and reports its observation with biases b_m and $b_{m'}$. The posterior of a Bayesian agent who reads the two reports is given by

$$\begin{aligned} \theta_i^{post} &= \frac{\tau_m/\tau}{\tau_m/\tau + \tau_{m'}/\tau + 1}(\theta_m - b_m) \\ &\quad + \frac{\tau_{m'}/\tau}{\tau_m/\tau + \tau_{m'}/\tau + 1}(\theta_{m'} - b_{m'}) + \frac{1}{\tau_m/\tau + \tau_{m'}/\tau + 1}\theta_i^{prior}. \end{aligned} \tag{2}$$

3 Empirical strategy

Equations (1) and (2) impose testable predictions. We first describe the implications of the theory of agents receive reports from only one news sources. Suppose we have data on agents' prior beliefs θ_i^{prior} before reading a news report r_m and their posterior beliefs θ_i^{post} immediately after reading r_m. Then equation (1) predicts that in the regression

$$\theta_i^{post} = \beta_0 + \beta_1 \theta_i^{prior} \tag{3}$$

the coefficient $\beta_1 \in [0,1]$, so that the variance of posterior beliefs is smaller than the variance of the priors. Moreover, the coefficient of this regression should be unchanged if individuals read several version of the same report r_m as they should rationally take into account that those reports carry no additional information.

Regression (3) can also be used to estimate perceived precision and bias of the media report. Let $\hat{\beta}_0$ and $\hat{\beta}_1$ be the estimated values of β_0 and β_1. Then the relative precision can be estimated as

$$\frac{\widehat{\tau_m}}{\tau} = 1/\hat{\beta}_1 - 1 \tag{4}$$

and bias as

$$\hat{b}_m = \theta_m - \frac{\hat{\beta}_0}{1-\hat{\beta}_1}. \tag{5}$$

If we have the estimates of precision and bias of several media sources, equation (2) can be used to check if individuals update their beliefs consistently in the presences of several conflicting reports. In particular, if we have estimates of bias and precision from two different media outlets, $\left(\frac{\widehat{\tau_m}}{\tau}, \hat{b}_m\right)$ and $\left(\frac{\widehat{\tau_{m'}}}{\tau}, \hat{b}_{m'}\right)$ and data on individuals priors and posteriors after reading two reports, one from each media outlet, the theoretically predicted coefficients β_0 and β_1 of the regression (3) can be deduced from (2). We can then test econometrically whether those predicted coefficients differ from the ones estimated in the data.

3.1 Experiment

We conduct an experiment in Shanghai with two types of news. The first is news about air pollution. There are several features of this news in our context to note. Air quality in China is currently reported using the Air Quality Index (AQI). This is a subjective index of health risk based on PM 2.5 readings. AQI reports both a numerical score as well as a risk category for each level of PM 2; the AQI will say whether the risk to health is “low” (1-3), moderate (4-6), high (7-10) or very high (10+).

The poor quality of air is a salient subject in China. Average AQI in Shanghai for 2013 was for example 173 (very unhealthy). In contrast, Los Angeles, known as one of the most polluted cities in the United States has an AQI of 91 (moderate). Air pollution in China receives a significant amount of attention from the official state media, the Chinese government, as well as blogs and other independent media.

Air pollution is also a politically sensitive issue. The government has vowed to reduce pollution levels. At the same time, it has banned non-government organizations from measuring and reporting air pollution in China. The one exception is AQI reported by the United States. The measurements are taken at U.S. embassies and consulates in several Chinese cities and posted on the websites of these respective organizations. The U.S. AQI is significantly higher than what is reported by the Chinese. As we discussed in our model, part of the difference is due to differences in accuracy and part of it is due to differences in biases. Note that the U.S. does not necessarily need to observe different underlying PM2.5 data to have different AQI reports since the U.S. threshold for health risks is lower than the Chinese (i.e., for any PM2.5, the U.S. AQI will assign a higher score than the Chinese AQI).

In our experiment, we will use both Chinese and U.S. reports of AQI in Shanghai. The Chinese reports are reported by two nominally independent environmental websites (these are the only two that report daily AQI in Shanghai) and the U.S. reports are reported by the U.S. consulate in Shanghai.[5] We measure AQI as the average AQI over a five day period prior to our experiment (December 9th to 15th, 1013). We use the same time period for both sources of news.

The second type of news is about temperature. This is not politically sensitive. Again,

[5]The Chinese sources are http://t.eastday.com/shhjbh and http://weibo.com/shanghaipm. The U.S. source is https://twitter.com/cgshanghaiair. These websites report hourly or every four hours (with occasionally missing hours).

we obtain this news from a U.S. source and a Chinese source.[6]

For the period of our study, Chinese news reported AQI to be 143. U.S. news reported it to be 173. Both Chinese and U.S. news reported temperature to be 6 celsius.

For the experiment, we advertised for student volunteers from three elite Shanghai universities. Participants in each university were then randomized into nine treatment groups. Bad weather caused participation to be low for the last two treatment groups. For all of the other groups, almost all of those assigned to treatment showed up. Thus, we will only discuss the first seven treatment groups.

Each group is asked to record their prior (a number), shown a piece of news on a screen, and then asked to record their posterior (a number). For groups that are asked about AQI, surveyors will first describe and explain what AQI is and show a scale with the categories, numbers, etc. on screen (see Online Appendix). For groups that are asked about temperature, surveyors will first show a scale of the thermometer. This is done to avoid respondents giving extreme values in their responses. These explanations are also shown together with the news reports. And they are left on the powerpoint when respondents are asked to report their posterior. Thus, there is no element of recall or memory test in this experiment. Respondents are give a few minutes to write down each number – they are not rushed.

Also note that each question is answered on a different sheet of paper. Surveyors make sure that respondents do not go back and change their answers once they are written down.

If a group is asked about both temperature and AQI, they are given the relevant explanations before each section of the questions. For these groups, respondents are always asked about temperature first so as to not anchor their beliefs towards the view that news is biased.

1. Chinese news: One piece on AQI

2. Chinese news: Two pieces on AQI

[6] Temperature reports from U.S. outlets come from Accuweather http://www.accuweather.com/zh/cn/shanghai/106577/november-weather/106577. The website reports daily maximum and minimums; we first take the daily average and average that across the week.

3. Chinese news: Two pieces (simultaneously) on AQI with references to the source (the Chinese government)

4. Chinese news: One piece of temperature, one piece on AQI

5. U.S. news: One piece of temperature, one piece on AQI

6. Chinese and U.S. news: One piece of temperature, one piece on AQI

7. Chinese news: One piece of temperature, one piece on AQI, additional survey questions (e.g., beliefs in bias).

In group 7, we asked respondents additional questions regarding their confidence in their accuracy, background information, as well as their perceptions of Chinese and U.S. media.

The survey is administered by students we hired and trained from the same universities. Participants are offered a small monetary reward (around 3 USD). Each treatment lasted less than 30 minutes.

Table 1 shows the balance across groups. The reported priors for AQI is statistically similar across groups 1-6. The same is true for the reported prior of temperature for Groups 5-7.

Note that in the data, the logarithm of AQI priors and posteriors are normally distributed, and the levels of temperatures priors and posteriors are normally distributed. Since the model uses the normal distribution, the empirical analysis will measure AQI in logs and temperature in levels. See Appendix Figure A.3.

We note that average priors of AQI are higher than both U.S. and Chinese reports. This is most likely because our experiment occurred immediately after a period of extremely high air pollution levels. This should not affect the interpretation of our experiments.

We we were concerned that subject responses to the experiment would be affected by the salience of pollution and other factors such as personal stress from exams, etc. We therefore minimized the number of days of the experiment and conducted it across all three universities within one week.

4 Reported Beliefs

First, we examine the descriptive statistics provided by Group 7. Several facts emerge in Table 2. Students spend approximately an hour a day watching/reading the news. Most of the time is spent on Chinese language sources. A significant amount of time is spent reading blogs.

Amongst elite university students, almost all are fluent in at least one foreign language. 26% have travelled to a foreign country. Thus, relative to the average Chinese citizen, our experimental subjects have had significantly more exposure to foreign cultures, and have much better ability to access non-government news.

We then examine their perceptions of news. Table 3 shows that students on average believe that the Chinese news is more accurate about temperature than about AQI, but the U.S. news is similarly accurate about both types of news. They also believe that the U.S. is more accurate about AQI than the Chinese news, where as the two are similar in reported temperatures. We note that these differences are not statistically significant. Thus, we interpret them as suggestive.

A more striking difference emerges in the perceptions about the source and systematic bias in news. 92% of respondents believe that the government is the source of Chinese news reports on AQI temperature. Only 50% of respondents believe that this is true for U.S. reports. 61% believe that Chinese news systematically misreports news for AQI, where as only 22% believe that U.S. news does so. In contrast, for the non-controversial topic of temperature, 16% and 11% believe that Chinese and U.S. news systematically misreports.

Amongst those who answer that there is systematic misreporting, the average belief is that the Chinese understate AQI and the U.S. overstates AQI.

5 Experimental Results

5.1 Discounting Repeated News

First, we compare Groups 1 (one piece of Chinese news on AQI), 2 (two pieces of Chinese news from two media outlets with identical AQIs) and 3 (the same two pieces of news on AQI as group 2 with referenced sources, which are the same government ministry). For each group, we estimate equation (3) and back out the precision and bias using the formulas in equations (4) and (5).

A comparison of the results in Tables 4 columns (1) and (2) show that Group 2 has a smaller estimated bias – i.e., respondents do not discount for the fact that the two pieces of news come from the same source. The p-value at the bottom of the table shows that this is significant at the 20% level. This could be because the fact that the source of news is the same is not salient – naive (not Bayesian), or because respondents have had limited prior exposure to Chinese AQI news in the past. The latter is unlikely *a priori* since 92% in Group 7 believes that the government is the source of Chinese AQI news.

However, we can investigate this by examining column (3), where the respondents are explicitly shown that the two pieces of news come from the same government source. A comparison of the estimated bias show that Group 3 has a similar bias to Group 2, which is much smaller than Group 1. This means that respondents do not discount the second piece of news, even if they see the news comes from one government source. This can still be reconciled with the Bayesian model if agents believe that the government provides (partially) independent information. However, again, this is inconsistent with reported beliefs that the government is the source of Chinese AQI news.

5.2 Conflicting Information

Next, we investigate how respondents interpret pieces of conflicting information. We estimate equation (3) for temperature and AQI reports for Groups 4 (Chinese news), 5 (U.S. news) and 6 (both news). The results are presented in Table 5. We then compare the priors and constants for temperature and AQI that we compute from Groups 4 and 5 to those that we estimate from Group 6. Recall that if respondents are Bayesian, then they should interpret the two news sources similarly when they receive them independently as when they receive them together. This follows from equation (2).

We find that for temperature, the estimated prior and constant of Group 6 is similar to those computed from Groups 4 and 5 (see column 5). The magnitudes are similar and the p-values at the bottom of the table reject the possibility that they are statistically significantly different. However, for AQI, the estimated prior and constant fro Group 6 is very different from those computed from Groups 4 and 5 (see column 6). The magnitudes are different and the p-values at the bottom of the table show that the differences are statistically significant at the 5% and 10% levels.

6 Heterogeneity

To investigate the degree of heterogeneity in updating, we first conduct a visual examination. We have two considerations in choosing the sample. First, we want to maximize sample size for power. Second, we need the subjects to be comparable. Therefore, we pool the subjects in Group 4, who are shown Chinese news on AQI and temperature, and Group 5, who are shown U.S. news on AQI and temperature.[7] Figures 1b and 1a plot the posteriors against priors for each group. The horizontal red line indicates the AQI or temperature in

[7] We alternatively pool subjects from many different combinations of the samples. The result that there are heterogeneous populations are always present, which is not surprising given the plots of the posteriors against priors. These additional results are available upon request.

the news. Points that lie along the horizontal line are those that update completely from their priors to the reported level. Points along the 45-degree line are those who do not update at all such that their posteriors equal their priors. Points that lie in between are those that take an approximate average of their priors and the reported news.

The figures show substantial heterogeneity. For formality, we use a Finite Mixture Model (FMM) to examine the bivariate relationship between priors and posteriors.[8] FMM is a Maximum Likelihood Estimation that optimally assigns individuals into groups. We allow the constants and the coefficients of the prior to vary across groups so that the assignment depends on the degree of updating and the presumed average level of AQI across subjects. The number of groups are decided by the econometrician. We alternatively estimate the model with one, two and three groups.

Table 6 Column (1) shows the OLS estimate of a regression of posteriors on priors. Column (2) shows the analogous results using MLE. Note that conceptually, these are FMM estimates where we assign all individuals to one group. They are similar to the OLS estimates.

Columns (3) and (4) show the FMM estimates where we assign individuals into two and three groups. Column (3) shows that when assigned to two groups, one groups updates partially (i.e., the coefficient for the prior is 0.225), while the other group updates very little (i.e., the coefficient for the prior is 0.953). Both coefficients are statistically significant at the 1% level. The constants are also different between the two groups.

At the bottom of the table, we present the fraction of the sample assigned to each group (π_1, π_2). We see that 68.2% of the sample are assigned to the first group of partial updaters and 31.8% are assigned to the second group of non-updaters. Note that not surprisingly, the coefficient for priors when there is one group (column (1)) is a weighted average of the coefficients when there are two groups (column (2)).

[8]See McLachlan and Peel (2004).

The estimates in column (4) show that when individuals are assigned to three groups, the first and third groups update partially (i.e., the coefficients for AQI priors are 0.211 and 0.202) and the second group again updates little (i.e., the coefficient for AQI prior is 0.968). The coefficients are statistically significant for all groups. However, the coefficients in the first and third groups are similar in magnitude and statistically indistinguishable, and the model only assigns 2.2% of the sample into Group 3.[9] Thus, the FMM model findings only two statistically distinct groups in the data.

We repeat the exercise for temperature. Figures 2a and 2b plot the posteriors for temperature against the priors. As before, we examine subjects from Groups 4 and 5. In Table 7, we pool the subjects. Columns (1) and (2) states the OLS and MLE estimates for one group. They are similar. In column (3), we assign subjects to two groups. The coefficients for the two groups are 0.396 and 0.281. The estimates are statistically significant at the 1% and 5% levels. However, they are statistically indistinguishable from one another. Thus, we conclude that the degree of updating for temperature is homogenous.

That the degree of updating is heterogenous for AQI, but not for temperature is interesting since it suggests that heterogeneity may not be an issue for non-controversial news.

One potential concern is that subjects may behave differently in an experiment setting than in the real world. Specifically, experimental subjects may be "lazy" and simply restate their prior or copy the AQI shown to them as their posterior. This is unlikely to drive our findings for two reasons. First, if heterogeneity comes from the experimental setting, then we should find similar heterogeneity for temperature. Second, we find similar results for heterogeneity fro AQI when we drop individuals whose posteriors are identical to their priors and whose posteriors are identical to the news.

[9] Note that since we specify the number of groups ex ante, the model necessarily assigns a positive number of observations into each group.

7 Conclusion

This study takes a step towards understanding how individuals interpret information in a regime where most media is historically censored by the government, and this is widely known to the public. We first document that government control of the news and its biases are very salient. We then ask whether in this context, individuals discount repeated information given by multiple government-controlled media outlets, and how they interpret conflicting information reported by non-government sources.

We find that individuals do not discount repeated information. Strikingly, this is true even when they are explicitly shown that the multiple pieces of news come from the same government sources. As such, this departs from the Bayesian framework. Second we find that they have difficulty interpreting conflicting pieces of information from government and non-government sources. In contrast, when shown non-conflicting information from the same sources, individuals behave rationally.

Finally, we document that the degree of updating is heterogeneous for news about air pollution, which is controversial in China, while it is homogenous for news about temperature, which is not controversial. These results suggests that allowing for heterogeneity is one promising avenue for future theories of how individuals interpret government-controlled news.

The results of this study only brushes the surface for understanding the role of media in autocratic regimes where government control and bias is widely known. Nevertheless, there are important policy implications. The results imply that increasing the number of state-controlled media outlets can be an effective method for persuading the population. Moreover, they show that introducing additional sources of news with conflicting information will not necessarily lead people closer to the truth, at least not in the short run, even if people have priors that these sources are independent and relatively accurate. This may explain why the increased flow of non-government information from the expansion of the internet has not obviously shifted views in countries such as China (or perhaps Russia) towards the views of the West.

The preliminary findings suggest several avenues for future research. First, the finding that individuals have difficulty interpreting conflicting information in the short run suggests that it is important to examine the effects of medium and long run exposure to conflicting information. In other words, how long does it take for individuals to learn to interpret conflicting information? This requires a natural or field experiment that can track individuals after they are exposed. Second, the finding of heterogeneity requires further exploration – e.g., what are the determinants of updating? what are the theoretical implications of heterogeneity? Related to this, it would be interesting to repeat the experiment for different segments of the population, in particular, for those that have less exposure or access to non-Chinese media, or for older generations that have been exposed to more government-driven political upheavals.

References

ADENA, M., R. ENIKOLOPOV, M. PETROVA, V. A. SANTAROSA, AND E. ZHURAVSKAYA (2013): "Radio and the Rise of the Nazis in Prewar Germany," CEPR Discussion Papers 9453, C.E.P.R. Discussion Papers.

BEI, J., T. LARSSON, D. STROMBERG, AND Y. WU (2014): "Microblogs and Newspapers in China," Working paper, Stockholm University IIES.

BEI, J., D. STROMBERG, AND Y. WU (2013): "The Determinants of Media Bias in China," Working paper, Stockholm University IIES.

CHAY, K. Y., AND M. GREENSTONE (2003): "The Impact Of Air Pollution On Infant Mortality: Evidence From Geographic Variation In Pollution Shocks Induced By A Recession," *The Quarterly Journal of Economics*, 118(3), 1121–1167.

DELLAVIGNA, S., R. ENIKOLOPOV, V. MIRONOVA, M. PETROVA, AND E. ZHURAVSKAYA (2011): "Cross-border media and nationalism: Evidence from Serbian radio in Croatia," NBER Working Papers 16989, National Bureau of Economic Research, Inc.

DELLAVIGNA, S., AND M. GENTZKOW (2010): "Persuasion: Empirical Evidence," *Annual Review of Economics*, 2(1), 643–669.

DELLAVIGNA, S., AND E. KAPLAN (2007): "The Fox News Effect: Media Bias and Voting," *The Quarterly Journal of Economics*, 122(3), 1187–1234.

DEMARZO, P. M., D. VAYANOS, AND J. ZWIEBEL (2003): "Persuasion Bias, Social Influence, And Unidimensional Opinions," *The Quarterly Journal of Economics*, 118(3), 909–968.

ENIKOLOPOV, R., M. PETROVA, AND E. ZHURAVSKAYA (2011): "Media and Political Persuasion: Evidence from Russia," *American Economic Review*, 101(7), 3253–85.

EYSTER, E., AND M. RABIN (2010): "Na?ve Herding in Rich-Information Settings," *American Economic Journal: Microeconomics*, 2(4), 221–43.

JIA, R. (2014): "Pollution for Promotion," Working paper, University of California at San Diego.

LEVS, J. (2011): "U.S. readings challenge China's smog claims," *CNN.com*.

MCLACHLAN, G., AND D. PEEL (2004): *Finite Mixture Models*, Wiley series in probability and statistics: Applied probability and statistics. Wiley.

NA, W. (2012): "No independent air quality data ban," *Global Times*.

YANAGIZAWA-DROTT, D. (2014): "Propaganda and Conflict: Theory and Evidence from Rwanda," Working paper, Harvard Kennedy School.

Table 1: Balance

Dependent Variable: Prior Beliefs		
	(1)	(2)
	AQI	Temperature
Group 2	0.0597	
	(0.0392)	
Group 3	-0.0318	
	(0.0399)	
Group 4	-0.00861	
	(0.0452)	
Group 5	-0.0442	0.0743
	(0.0460)	(0.309)
Group 6	-0.0294	-0.205
	(0.0467)	(0.315)
Group 7		-0.394
		(0.323)
Observations	1,128	533
R-squared	0.007	0.005

Table 2: Mean Exposure to News and Non-Chinese Environments

	Obs	Mean	Std. Dev.
In the last week, how many hours did you spend reading/watching any news ….			
In total	119	6.98	(8.17)
Chinese language sources	119	4.13	(5.36)
Foreign language sources	119	0.71	(1.49)
Blogs (any language)	119	2.14	(3.50)
			0.00
Fluent in a foreign language	119	0.99	(0.09)
Have travelled to a foreign country	119	0.26	(0.44)
Have travelled to Japan or Korea	119	0.10	(0.30)
Have travelled to Australia, Canada or the U.S.	119	0.11	(0.31)
Have travelled to Europe	119	0.03	(0.18)

Table 3: Perceptions of News Content

	Obs.	Mean	Std. Dev
How accurate do you think Chinese news are is on a scale 1 to 5?			
About Temperature	119	3.59	(0.775)
About AQI	119	3.19	(0.932)
How accurate do you think U.S. news is on a scale 1 to 5?			
About Temperature	119	3.34	(0.704)
About AQI	119	3.41	(0.682)
Think that the government is the source			
Chinese news on temperature	119	0.92	(0.266)
Chinese news on AQI	119	0.92	(0.279)
U.S. news on temperature	118	0.50	(0.502)
U.S. news on AQI	118	0.50	(0.502)
Do you think the news systematically misreports information?			
Chinese news on temperature	119	0.16	(0.368)
Chinese news on AQI	119	0.61	(0.491)
U.S. news on temperature	118	0.11	(0.314)
U.S. news on AQI	119	0.22	(0.415)
If yes, how much do you think news systematcially misreports information (percentages)?			
Chinese news on temperature	19	-0.03	(0.269)
Chinese news on AQI	72	-0.34	(0.356)
U.S. news on temperature	13	0.14	(0.299)
U.S. news on AQI	26	0.32	(0.197)

Table 4: Updating with Multiple Pieces of AQI News

	Dependent Variable: AQI Posterior		
	(1)	(2)	(3)
	1 News	2 News	2 News + Source
AQI Prior	0.549	0.585	0.491
	(0.0429)	(0.0366)	(0.0390)
Constant	2.311	2.086	2.548
	(0.234)	(0.202)	(0.211)
Observations	250	240	223
R-squared	0.398	0.517	0.419
Bias = [Constant/(1-AQI Prior)] - 4.96	**-0.165**	**-0.0634**	**-0.0501**
Std. Error	**(0.051)**	**(0.057)**	**(0.046)**
Precision = [(1/AQI Prior)] - 1	**0.821**	**0.710**	**1.035**
Std. Error	**(0.142)**	**(0.107)**	**(0.161)**
Precision 1 vs. 3 (p-val)	**0.511**		
Precision 1 vs. 2 (p-val)	**0.682**		
Bias 1 vs. 3 (p-val)	**0.104**		
Bias 1 vs. 2 (p-val)	**0.173**		

Table 5: Updating with Chinese and U.S. News for AQI and Temperature

	Dependent Variables: Posterior Beliefs					
	Chinese News		U.S. News		Both News	
	(1)	(2)	(3)	(4)	(5)	(6)
	Temp	AQI	Temp	AQI	Temp	AQI
Temperature Prior	0.297		0.447		0.191	
	(0.0455)		(0.0457)		(0.0315)	
AQI Prior		0.528		0.438		0.217
		(0.0489)		(0.0367)		(0.0355)
Constant	4.902	2.364	3.806	2.931	5.516	3.978
	(0.416)	(0.266)	(0.418)	(0.199)	(0.284)	(0.193)
Observations	144	145	140	138	130	131
R-squared	0.230	0.450	0.410	0.511	0.223	0.225
Bias = [Constant/(1 - Prior)] - News	-0.969	-0.0476	-0.888	-0.0638		
Std. Error	(0.206)	(0.062)	(0.256)	(0.029)		
Precision = [(1/Prior)] - 1	2.371	0.894	1.235	1.284		
Std. Error	(0.517)	(0.175)	(0.228)	(0.192)		
Prior Computed from Groups 4,5					**0.217**	**0.315**
Constant Computed from Groups 4,5					**6.039**	**3.285**
P-Value of Difference between Computed (4,5) and Estimated (6)						
AQI Prior						**0.0448**
AQI Constant						**0.0904**
Temperature Prior					**0.524**	
Temperature Constant					**0.687**	

Table 6: Test of Heterogeneity in AQI Updating

	OLS	MLE	2-component FMM	3-component FMM
	(1)	(2)	(3)	(4)
Prior on AQI (β_1)	0.481***	0.375***	0.225***	0.211***
	(0.031)	(0.030)	(0.040)	(0.034)
Constant (α_1)	2.658***	3.176***	3.962***	4.027***
	(0.170)	(0.160)	(0.206)	(0.177)
Std Dev. (σ_1)			0.156***	0.145***
			(0.011)	(0.010)
Prior on AQI (β_2)			0.953***	0.968***
			(0.056)	(0.045)
Constant (α_2)			0.184	0.104
			(0.306)	(0.244)
Std Dev. (σ_2)			0.105***	0.100***
			(0.021)	(0.016)
Prior on AQI (β_3)				0.202***
				(0.070)
Constant (α_3)				4.494***
				(0.414)
Std Dev. (σ_3)				0.058***
				(0.030)
π_1			0.682	0.671
			(0.066)	(0.055)
π_2			0.318	0.307
			(0.066)	(0.058)
π_3				0.022
				(0.016)
Log likelihood		45.31	80.01	81.70
BIC		-68.04	-120.50	-101.31
Observations	283	283	283	283

Table 7: Test of Heterogeneity in Temperature Updating

	OLS	MLE	2-component FMM
	(1)	(2)	(3)
Prior on AQI (β_1)	0.365***	0.265***	0.396***
	(0.033)	(0.034)	(0.060)
Constant (α_1)	4.420***	4.977***	4.115***
	(0.298)	(0.277)	(0.457)
Std Dev. (σ_1)			0.983***
			(0.107)
Prior on AQI (β_2)			0.281**
			(0.114)
Constant (α_2)			5.398***
			(1.167)
Std Dev. (σ_2)			2.293***
			(0.401)
Prior on AQI (β_3)			
Constant (α_3)			
Std Dev. (σ_3)			
π_1			0.785
			(0.112)
π_2			0.215
			(0.112)
π_3			
Log likelihood		-482.16	-480.23
BIC		986.92	1000.01
Observations	284	284	284

Figure 1: AQI Posterior versus Prior

(a) Group 4 – Chinese News

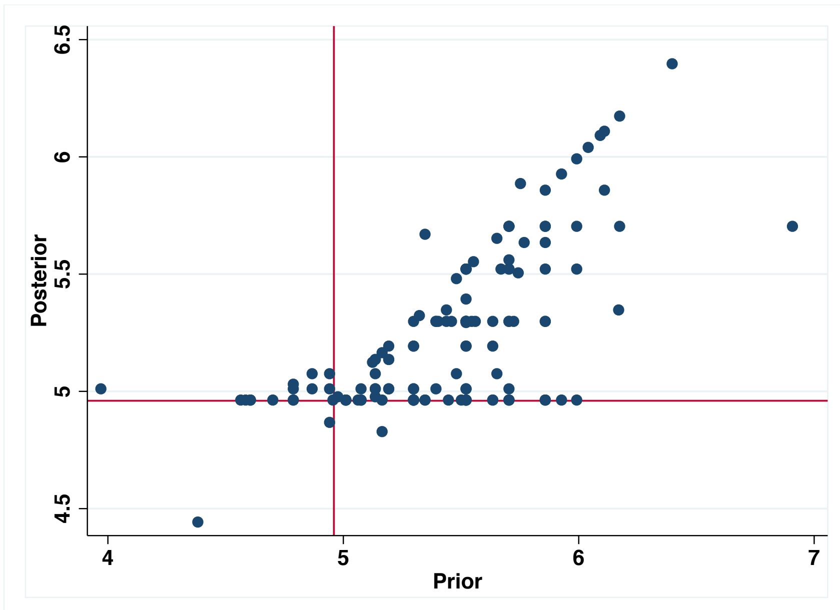

(b) Group 5 – U.S. News

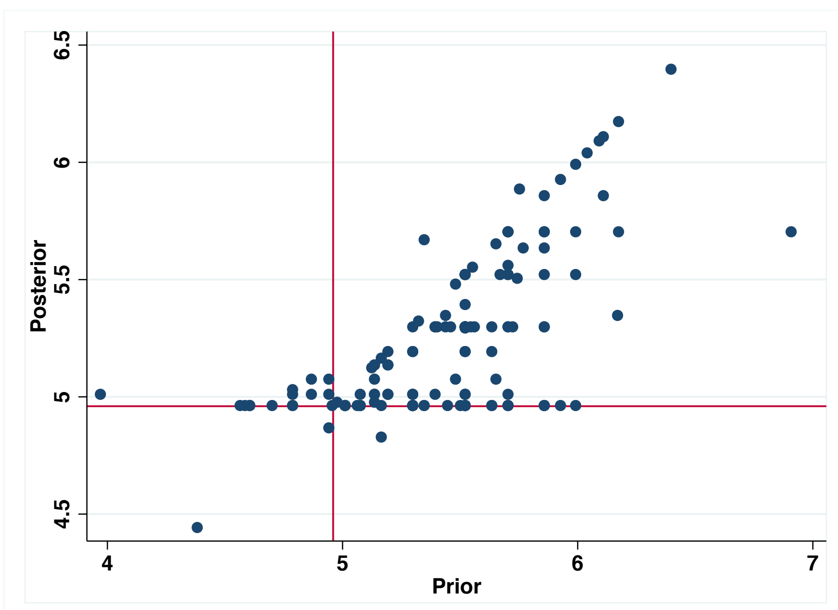

Figure 2: Temperature Posterior versus Prior

(a) Group 4 – Chinese News

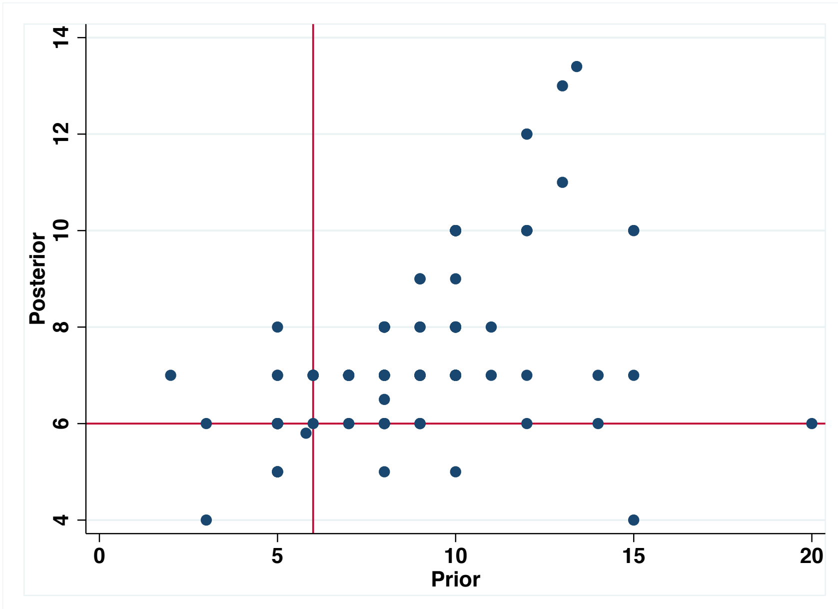

(b) Group 5 – U.S. News

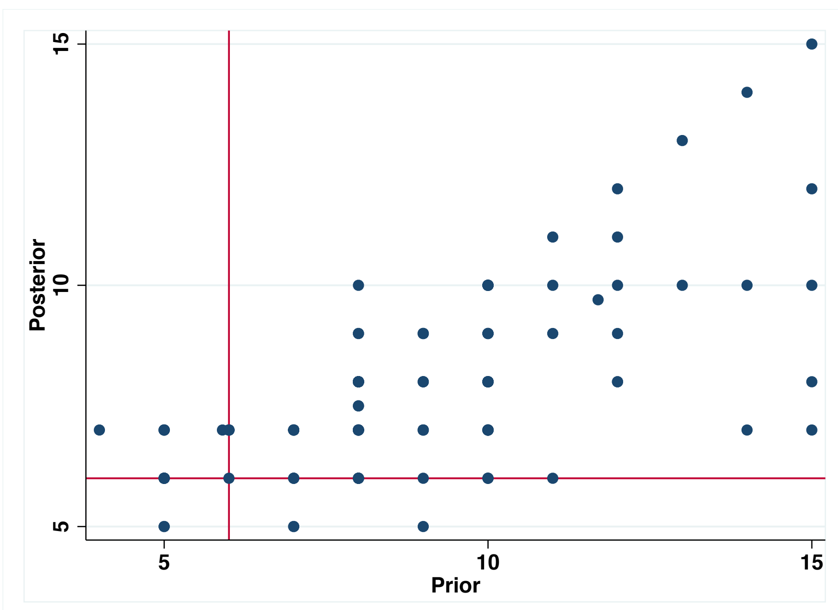

Appendix

Figure A.1: AQI Scale

什么是空气质量指数(AQI)?

AQI 数值	AQI 类别	对健康影响情况
0~50	优	空气质量好，空气污染对人体的危害很小
51~100	良	空气质量可以接受，但某些污染物可能对极少数异常敏感人群健康有较弱影响
101~150	轻度污染	易感人群症状有轻度加剧，健康人群出现刺激症状
151~200	中度污染	进一步加剧易感人群症状，可能对健康人群心脏.呼吸系统有影响
201~300	重度污染	心脏病和肺病患者症状显著加剧，运动耐受力降低，健康人群普遍出现症状
> 300	严重污染	健康人运动耐受力降低，有明显强烈症状，提前出现某些疾病

Figure A.2: Temperature Scale

Figure A.3: Distribution of Priors and Posteriors

(a) Log AQI Prior

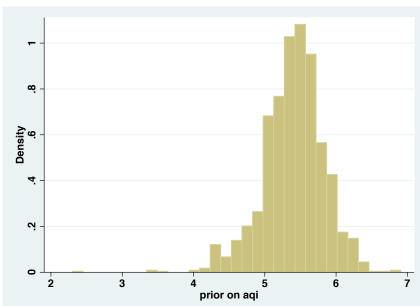

(b) Log AQI Post

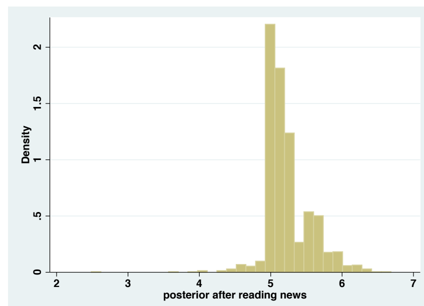

(c) Temp Prior

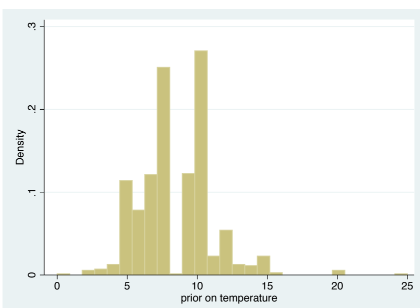

(d) Temp Posterior