

Placebo-Augmented PICA Design (PICA-2): Assessing the Influence of Foreign Propaganda*

Naoki Egami[†] Donald P. Green[‡]

Eunji Kim[§] Daniel C. Mattingly[¶]

Abstract

Most media experiments analyze the impact of messages chosen by researchers, often on topics participants might avoid. Addressing this limitation, the innovative “preference-incorporating choice and assignment” (PICA) design allows researchers to estimate causal effects among viewer subgroups with different content preferences. Building on this, we propose a novel placebo-augmented PICA design that more efficiently estimates these causal effects, along with a falsification test that helps detect violations of core assumptions. We use this new method to study whether Chinese and American propaganda cultivate support for their competing political models in Africa, a subject of broad scholarly and policy interest. A survey experiment conducted in five African countries shows that propaganda’s effects are especially large for those who would ordinarily choose not to watch it. Our results highlight both the value of the experimental design and the challenge faced by the two global powers in influencing foreign public opinion.

*The authors thank the Graduate School of Arts and Sciences at Columbia University, which funded this research. This study was reviewed and approved by the Institutional Review Board at Columbia University, protocol AAAU7051 and reviewed and determined exempt by the Institutional Review Board at Yale University, protocol 2000035589.

[†]Assistant Professor of Political Science, Columbia University. Email: naoki.egami@columbia.edu. URL: <https://naokiegami.com>

[‡]J.W. Burgess Professor of Political Science, Columbia University. Email: dpg2110@columbia.edu. URL: <https://donaldgreen.com>

[§]Assistant Professor of Political Science, Columbia University. Email: eunji.kim@columbia.edu. URL: <https://www.eunjikim.pub>

[¶]Associate Professor of Political Science, Yale University. Email: daniel.mattingly@yale.edu. URL: <https://daniel-mattingly.com>

1 Introduction

The study of propaganda’s impact on public opinion has captivated scholars for decades, with roots stretching back to the 1930s. This field of research has attracted renewed attention as global powers attempt to sway public opinion across the globe, with China in particular investing billions in an attempt to mold public opinion in its favor (Weiss, 2019; Blair *et al.*, 2022). Recent studies have examined whether and how Chinese and American external propaganda build mass support for authoritarian or democratic political systems among international audiences (Mattingly *et al.*, 2024; Fan *et al.*, 2024).

Similar to the broader experimental literature on political communication, recent studies often use a “forced exposure” design that randomly exposes participants to predetermined media content. A notable shortcoming of this approach is its potential to overlook the varying impacts of political communication on individuals with different media preferences. Those who would ordinarily avoid certain political topics or news sources might have a distinctive reaction to such messages, if exposed.

One important methodological advance to address this limitation is the Preference-Incorporating Choice and Assignment (PICA) design (de Benedictis-Kessner *et al.*, 2019; Knox *et al.*, 2019). This design allocates respondents to two arms: a classic “forced exposure” condition in which they are exposed to a randomly assigned piece of media or a “free choice” condition in which they select which media they consume. This design, which allows researchers to investigate how communication effects vary according to subjects’ existing media preferences, has been adopted by a growing number of studies (e.g., Kraft *et al.* 2022; Balcells *et al.* Forthcoming) with some methodological extensions (Testa *et al.*, 2021).

The PICA design represents an important innovation, but it has one noteworthy limitation: it tends to lack the statistical power to detect treatment effect heterogeneity because the method makes weak assumptions. To overcome potentially large standard errors, recent papers using the PICA design employed large sample sizes, with $N > 7,000$ in Balcells *et al.* (Forthcoming), de Benedictis-Kessner *et al.* (2019), and Markovich *et al.* (2020).

In this letter, we propose an important extension — the placebo-augmented PICA design (PICA-2) — that can greatly improve precision while retaining the ben-

efits of the original PICA design. The PICA-2 design is therefore a feasible option for political scientists operating under tight budget constraints. PICA-2 adds one feature to the original PICA design: in the free choice condition, PICA-2 randomly assigns respondents to a treatment or placebo arm (see Figure 1 for an overview). For example, in our study of foreign propaganda, for those in the free choice condition who choose to watch something about China, the treatment arm presents Chinese propaganda about the desirability of its political model, whereas the placebo arm presents content on China’s natural landscape instead. The proposed design offers substantial improvement in estimation accuracy when the placebo is valid, i.e., when videos about China’s natural landscape elicit the same outcomes as control videos about nature in general. We provide a falsification test that enables researchers to assess this assumption empirically.

Using the PICA-2 design in survey experiments conducted in five African countries, we affirm that Chinese propaganda can be quite effective in attracting support for its political model (Mattingly *et al.*, 2024; Goldsmith *et al.*, 2021; Green-Riley, 2022; Blair *et al.*, 2022). Importantly, the PICA-2 design reveals for the first time that these media effects are far more pronounced among those who would choose *not* to watch such content if given a choice — highlighting an underappreciated challenge that global powers face as they endeavor to shift foreign attitudes.

2 Placebo-Augmented PICA (PICA-2) design

To understand how causal effects vary across subgroups with different media consumption preferences, the original PICA design (de Benedictis-Kessner *et al.*, 2019; Knox *et al.*, 2019) randomly assigns respondents to a free choice condition and a forced exposure condition. In the traditional forced exposure condition, respondents are randomly exposed to certain media content. In the free choice condition, respondents consume the media they choose. Researchers can use bounds estimators¹ to assess the average choice-specific treatment effect (ACTE), which represents the conditional average treatment effect for the subpopulation of subjects who would choose a particular treatment option. An attractive feature of the original PICA design is that it allows for possible discrepancies between stated media preference

¹The quantity of interest for bounds estimators is the range of possible values for the causal effect (Duarte *et al.*, 2023), unlike conventional point estimators that attempt to recover a specific parameter value.

and actual media choice.

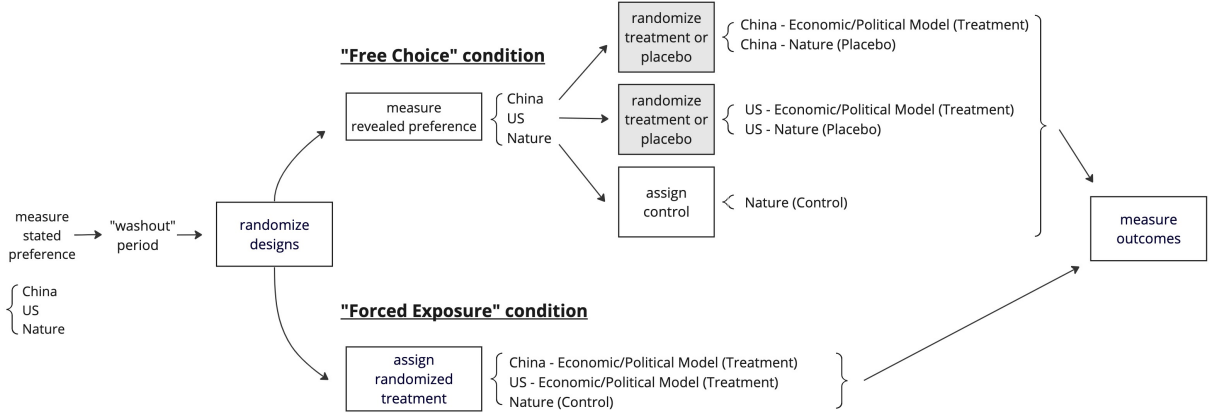


Figure 1: PICA-2 Experimental Design

Our innovation is to build a placebo design into the free choice condition *after the choice is made*. In the free choice condition, we randomly assign respondents to a treatment or placebo arm (See Figure 1). This placebo-augmented design requires the *valid placebo assumption*: that placebo videos in the free choice arm and control videos in the forced exposure arm produce the same average outcomes. This assumption is testable; when it holds, our proposed design can sharply reduce standard errors. In our application, we find that the length of confidence intervals based on our estimator is less than half the length of the original PICA bound estimator.²

When the valid placebo assumption is falsified, users can analyze the data as they would under an original PICA design. Importantly, however, invalid placebo arms in the free choice condition cannot contribute to estimation. In such a scenario, we expect to incur some increase in standard errors, but this efficiency loss is often small. Simulations suggest that the length of confidence intervals increases by only about 5% under our design compared to the original PICA design (see Appendix B). Overall, the proposed PICA-2 design offers a robust approach to improve the efficiency of the original PICA design, while retaining its benefits.

²Appendix B applies the PICA-2 design to the simulations of Knox *et al.* (2019). We find that, when the valid placebo assumption holds, the length of confidence intervals based on our estimator can be less than 40% of those based on the original PICA bound estimator under various realistic scenarios.

2.1 Design, Assumption, and Estimators

Imagine three types of media options $\{0, 1, 2\}$ where 0 represents a nature video, 1 Chinese propaganda, and 2 American propaganda. We follow the notation of de Benedictis-Kessner *et al.* (2019) and Knox *et al.* (2019) and introduce different notation only when necessary. We use the design indicator $D_i \in \{0, 1\}$ to denote whether subject i is in the forced exposure condition ($D_i = 1$) or in the free choice condition ($D_i = 0$). In the forced exposure condition, respondents are randomly assigned to treatments $T_i \in \{0, 1, 2\}$. For subjects in the free choice condition, we allow them to select which type of videos they want to watch, and this actual choice is denoted by $C_i \in \{0, 1, 2\}$.³ For those who choose to watch something about China or the US ($C_i \in \{1, 2\}$), we further randomly assign respondents to the treatment or placebo videos, $T_i \in \{1, 1_P\}$ for subjects with $C_i = 1$ and $T_i \in \{2, 2_P\}$ for subjects with $C_i = 2$, where we use subscript P to denote a placebo version of each video. For those who choose to watch something about rest of the world ($C_i = 0$), they are assigned to $T_i = 0$. Finally, we use $Y_i(t)$ to define the potential outcome where $t \in \{0, 1, 1_P, 2, 2_P\}$. As Knox *et al.* (2019) point out, this notation assumes D has no direct effect on outcomes, which is also known as the exclusion restriction.

Formally, the following conditional independence conditions define our proposed design.

Definition 1 (Placebo-Augmented PICA design)

$$\{Y(t), C\} \perp\!\!\!\perp D, \quad Y(t) \perp\!\!\!\perp T_1 \mid D = 1, \quad Y(t) \perp\!\!\!\perp T_0 \mid C = \tilde{t}, D = 0$$

where $T_1 \in \{0, 1, 2\}$, $T_0 = \{\tilde{t}, \tilde{t}_p\}$, and $\tilde{t} \in \{1, 2\}$.

The average treatment effect (ATE) is defined as $\mathbb{E}(Y(t) - Y(t'))$, which represents the causal effect of receiving treatment t rather than t' . As in the original PICA design, researchers can identify and estimate the ATEs just using the forced exposure condition.

Following de Benedictis-Kessner *et al.* (2019) and Knox *et al.* (2019), we also define the average choice-specific treatment effect (ACTE) to be $\mathbb{E}(Y(t) - Y(t') \mid C)$, which represents the causal effect of receiving treatment t rather than t' among

³Note that C_i is not observed for subjects in the forced exposure condition.

subjects who choose to watch C . In our application, we are interested in decomposing the ATE of watching Chinese propaganda $\mathbb{E}(Y(1) - Y(0))$ into two ACTEs, $\mathbb{E}(Y(1) - Y(0) \mid C = 1)$ and $\mathbb{E}(Y(1) - Y(0) \mid C \neq 1)$, which represent the causal effects of watching Chinese propaganda among those who choose to watch something about China and among those who do not choose to watch something about China. Similarly, we may decompose the ATE of watching US propaganda into two ACTEs. As emphasized in de Benedictis-Kessner *et al.* (2019), because there may be a discrepancy between the stated media preference and the actual media choice, researchers cannot estimate these quantities with simple subgroup analysis based on the stated media preferences measured in surveys.

The *valid placebo* assumption discussed above may be formalized as follows: placebo videos ($T \in \{1_p, 2_p\}$) and control videos ($T = 0$) have identical average potential outcomes within each corresponding subgroup.

Assumption 1 (Valid Placebo)

$$\mathbb{E}(Y(0) \mid C = 1) = \mathbb{E}(Y(1_p) \mid C = 1) \text{ and } \mathbb{E}(Y(0) \mid C = 2) = \mathbb{E}(Y(2_p) \mid C = 2)$$

For example, among those who choose to watch content about Chinese, the average potential outcome of a nature video from a non-Chinese source assumed to be equal to the average potential outcome of watching a nature video produced by the Chinese government. Researchers can make this assumption more plausible by carefully selecting and pre-testing placebo videos for the free choice condition. See Appendix C for practical guidance about designing valid placebos, incorporating some of the lessons learned from our mixed success in designing the trial presented below.

When the valid placebo assumption holds, we can identify the ACTE of receiving treatment t among those who choose to watch t as follows.

$$\mathbb{E}(Y(t) - Y(0) \mid C = t) = \mathbb{E}(Y \mid T = t, C = t, D = 0) - \mathbb{E}(Y \mid T = t_p, C = t, D = 0),$$

where $t \in \{1, 2\}$. Researchers can estimate this quantity using difference-in-means or linear regression estimators, comparing average outcomes of the treatment group and placebo group among subjects who choose to watch $t \in \{1, 2\}$ in the free choice condition.

We can also identify the ACTE of receiving treatment t among those who choose *not* to watch t by comparing the ATE and the ACTE among those who choose to

watch t . Researchers can estimate this quantity using plug-in estimators and use bootstrap to construct confidence intervals. We provide details in Appendix A.

One key benefit of the PICA-2 design is that researchers can also subject the valid placebo assumption to a falsification test. When this assumption holds

$$\begin{aligned} \mathbb{E}(Y | T = 0, D = 1) &= \mathbb{E}(Y | T = 1_p, C = 1, D = 0) \Pr(C = 1 | D = 0) \\ &+ \mathbb{E}(Y | T = 2_p, C = 2, D = 0) \Pr(C = 2 | D = 0) + \mathbb{E}(Y | C = 0, D = 0) \Pr(C = 0 | D = 0). \end{aligned}$$

This equality can be tested using mean outcomes from the PICA-2 design. Note that the proposed falsification test assesses the valid placebo assumption under the assumption of no design effect, so more formally, the test can be seen as an omnibus test of the valid placebo assumption and the no design effect assumption.

Finally, when the valid placebo assumption is violated, researchers can analyze the data using the original PICA bound estimator (Knox *et al.*, 2019) by ignoring the outcomes of subjects assigned to placebo videos in the free choice condition.

When the key assumptions hold, our PICA-2 design provides efficiency gains by focusing on a subset of the ACTEs, $\mathbb{E}(Y(t) - Y(0) | C = t)$ and $\mathbb{E}(Y(t) - Y(0) | C \neq t)$ where $t \in \{1, 2\}$. The former captures the causal effect of treatment t relative to the control among those who choose to watch t ; the latter captures the causal effect among those with the opposite viewing preference. While these cover many substantively important ACTEs of interest, there are other types of ACTEs, e.g., the ACTE among those who choose to watch the control video (units with $C = 0$). If these quantities are of primary interest, researchers can rely on the original PICA design, which allows users to put bounds on all possible ACTEs.

3 Implementing PICA-2

Using the PICA-2 design described above, we test the extent to which consuming Chinese or American propaganda can cultivate support for their respective political or economic models, and importantly, whether such effects are more pronounced among those who would not typically select such content, compared to those who gravitate toward it.⁴ For a discussion of ethical standards see Appendix D.

⁴The OSF preregistration is available via <http://tinyurl.com/pica2preanalysis>.

3.1 Subjects and Context

In recent years, the Chinese government has stepped up efforts to shape global public opinion in its favor. These efforts are especially visible in Africa, the site of an escalating rivalry between China and the United States (Blair *et al.*, 2022). Yet scholars remain skeptical about the extent to which China can sell its political and economic system to a foreign audience (Weiss, 2019). Scholars argue that the “China model” mixes “liberal economic policy” and “a ruling party [with a] firm grip” on society (Zhao, 2010, p. 419). Conversely, it is also an open question whether the United States can effectively maintain its own global influence and continue to promote its model of democracy, as it has long sought to do (with mixed success) around the world (Bush, 2015). We, therefore, apply our PICA-2 design to the study of Chinese and U.S. propaganda and recruited respondents from Africa.

African countries, of course, vary widely in their social, economic, and political characteristics. In order to improve external validity, we used multi-country experiments and selected diverse sites via synthetic purposive sampling. As the population of sites, we use the following three eligibility criteria due to our data and practical constraints: (a) use English as one of the official languages, (b) population size of at least 1 million, and (c) data on the UN’s ideal point alignment with China and the US exist. From this population, we selected five diverse sites—Botswana, Cameroon, Nigeria, Uganda, and Zambia—that differ in key contextual factors (GDP per capita, polity score, ethnic fractionalization, relationships with China and US measured by the UN’s ideal point alignment, and subregions in Africa). This diverse site selection allows us to evaluate whether experimental results vary across different contexts. We provide details in Appendix F.

3.2 Intervention

Our core media treatments come from two state-sponsored sources: the China Global Television Network (CGTN) and public diplomacy messaging from the U.S. Department of State. The videos highlighted the ostensible merits of each country’s economic and political model. The videos were broadly representative of the broader corpus of videos produced by each institution. These stimuli have been used in Mattingly *et al.* (2024), which finds that Chinese messaging increased the pro-

portion of respondents in nineteen countries who prefer the Chinese political model to the American model from 16 to 54 percent. In addition, we obtained placebo videos from the Chinese and American state-run media arms. These videos focused on descriptions of the flora, fauna, and landscape of each country without explicit reference to politics or economics.

3.3 Treatment Assignment

Respondents were randomly assigned to one of two conditions: a forced exposure condition (40 percent probability of assignment) or a free choice condition (60 percent probability of assignment).⁵ Based on pilot results, we elected to over-weight the probability of assignment to the free choice condition to ensure adequate power in each of the arms.

In the forced exposure condition, participants were randomly assigned to one of three scenarios, each with an equal probability of selection. They could view two videos produced by the CGTN or the U.S. State Department (these two videos were chosen from a pool of four, two about the respective country’s economic model and the other two about the political model) or two videos about nature. In the free-choice condition, participants were allowed to choose one of three categories. As a feature of our PICA-2 design, those who chose to watch something about China (or the U.S.) were further randomly assigned to one of two conditions: viewing two Chinese (or American) government-produced videos (selected from the same pool of videos as in the forced exposure condition) or viewing videos about China’s (or America’s) natural landscape. Those who elected to watch something about nature in the free choice condition watched the same videos as those in a forced condition. See Appendix E for details.

3.4 Outcomes

Our investigation centers on two key outcomes. We examine preference for a political or economic model, asking participants whether they would prefer their country

⁵To elicit media preference, we asked the following question: “We are interested in learning what people can remember from what they watch. We would now like you to watch two videos, and then answer some questions about them. Which of these categories of videos would you like to watch now? Please click on the description of the video you want to watch.” Respondents were able to choose among 1) Something about China, 2) Something about the United States, and 3) Something about the world.

to follow the Chinese or American version. We ask respondents: “If you were to choose, which one would you like your country to adopt: the Chinese political [economic] model or the United States political [economic] model?” Respondents then chose on a six-point scale that ranges between “Strongly prefer the Chinese political [economic] model” and “Strongly prefer the U.S. political [economic] model.”

4 Results

The results from our experiment highlight the advantages of the PICA-2 design. We show how the design operates both under ideal circumstances — when we have strong evidence of the validity of the placebo assumption and can reap gains in efficiency from our placebo-augmented design — and how the design fares under less favorable circumstances, when a core design assumption is violated.⁶

4.1 Average Treatment Effects

We use linear regression to estimate the ATEs. Per our pre-analysis plan, we include the following pre-treatment covariates in the regression specification: gender, age, education, national pride, left-right political orientation, and country.

Overall, the estimated ATE results from the forced exposure condition, presented as black circles in each panel of Figure 2, show that Chinese media has very strong effects on attitudes towards the Chinese political model, whereas American media has weaker effects. Viewing the Chinese treatments moves attitudes on the desirability of the Chinese political model a full two points on a six point scale, moving the median respondent from “Somewhat prefer the U.S. political model” to “Prefer the Chinese political model.” Viewing American propaganda moves respondents by one-third of a point along a six point scale, with the median respondent indicating that they “Prefer the U.S. political model.” Average treatment effects of Chinese and American propaganda were comparable (with the sign reversed) when the question focused on respondents’ preferred economic model. Appendix F reports that these causal effects seem to be relatively homogeneous across countries.

⁶See Appendix G for checks on other core assumptions. See Appendix H for details on the full results.

4.2 When the Valid Placebo Assumption Holds: Political Model

We now investigate the ACTEs, which are estimates of the effect of viewing media on respondents with different viewing preferences. The PICA-2 design allows us to do so with much greater precision than prior designs.

We find no evidence to challenge the validity of the placebo assumption for the political model outcome. Our falsification test shows that the mean outcome under the forced control condition (2.72) is almost identical to the weighted mean of outcomes under the free-choice placebo conditions (2.71), resulting in $p\text{-value} = 0.92$.

We then move to the core results based on our PICA-2 estimators, presented with red triangles in Figure 2-(A). We find that Chinese and American propaganda is most effective among those who prefer *not* to watch such messages — a substantively important finding about the conditions under which foreign propaganda is influential. The estimated effect of Chinese propaganda on preferences for China’s political model is twice as large for non-choosers of Chinese content as it is for choosers. In other words, those who are not disposed to view Chinese propaganda are most persuaded by it. Similarly, the effect of American messaging among those who choose to watch it is close to zero, while non-choosers show effects of roughly half a scale point.

To illustrate the efficiency gain, we also report the bound estimator proposed by de Benedictis-Kessner *et al.* (2019) and Knox *et al.* (2019).⁷ The valid placebo assumption ensures that our PICA-2 and the original PICA bound estimator give similar results. However, the PICA-2 estimator’s confidence intervals are notably shorter, averaging 45% the length of the bounding estimator’s intervals. In practice, this means that researchers can use the PICA-2 design and achieve the same precision with a much smaller sample size than before.

⁷Our implementation of these bounds makes minimal assumptions. In principle, as Knox *et al.* (2019) point out, one could tighten the bounds by invoking domain knowledge to specify sensitivity parameters that capture the maximum absolute difference between the average potential outcome among those who state a particular treatment preference and the average potential outcome among those who actually choose that treatment.

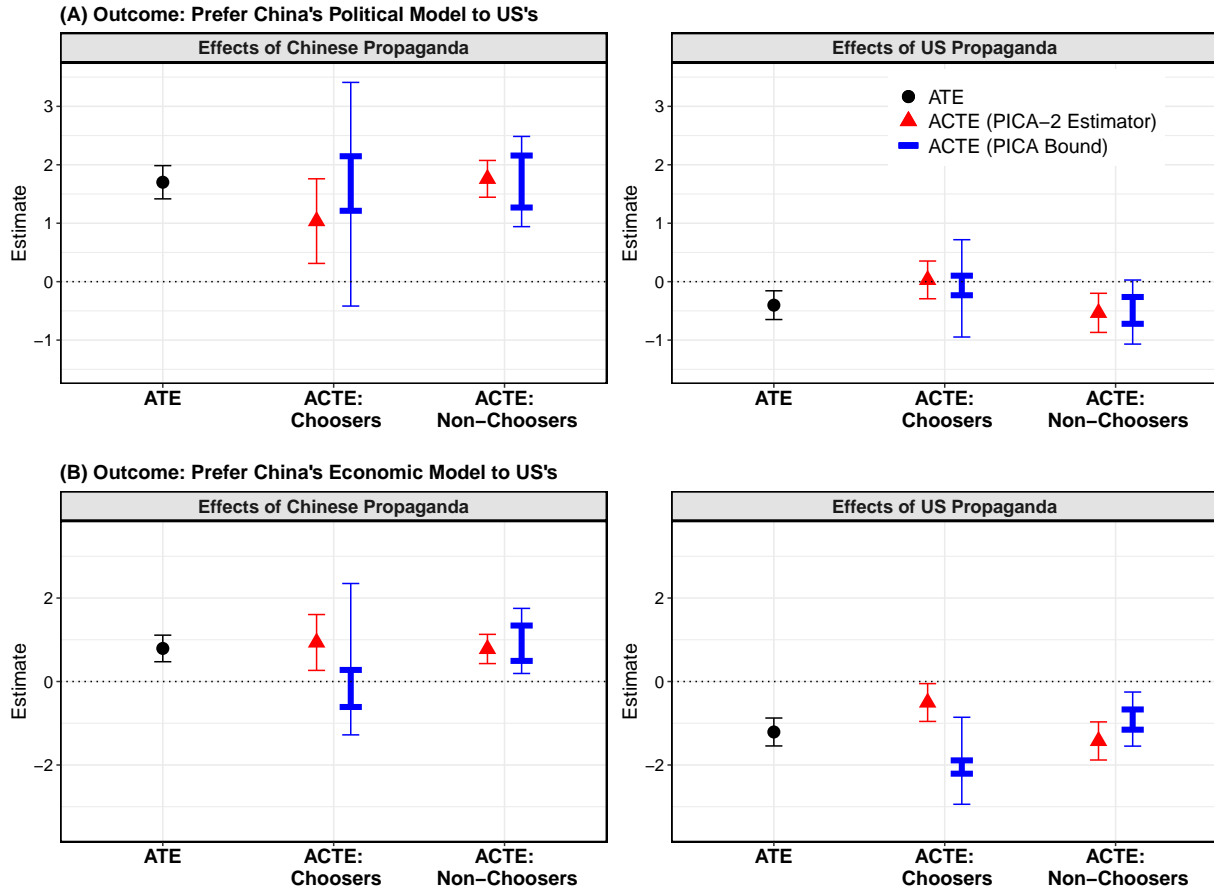


Figure 2: Effects of Foreign Propaganda on African Public Opinion

Note: For the ATEs, we report the point estimates and 95% confidence intervals (black circles and error bars). For the ACTEs, we report point estimates and 95% confidence intervals based on our PICA-2 estimators (red triangles and error bars) as well as bounds and 95% confidence intervals around the bounds based on the estimators of de Benedictis-Kessner *et al.* (2019) (blue thick and thin error bars).

4.3 When the Falsification Test Fails: Economic Model

On the contrary, we reject the validity of the placebo assumption for the economic model outcome, providing an interesting (if unintended) opportunity to assess the design when a core assumption fails. In this case, the mean outcome under the forced control condition (4.06) and the weighted mean of outcomes under the free-choice placebo conditions (3.58) differ. This difference is statistically significant (p-value < 0.01).

The failure of the placebo assumption could stem from several sources. One possibility is a violation of the exclusion restriction; being assigned to the forced or choice arms may have had a direct effect on the outcome. Another possibility is that the selection of the placebo videos in the free choice arm — the Chinese and US nature videos — moved attitudes towards the Chinese and American economic models.⁸ Ex ante, we did not expect this assumption to fail for this particular outcome measure, but the fact that it did fail exposes a risk that other researchers should be aware of as they consider the PICA-2 design. See Appendix C for a discussion of how to minimize this risk.

When the valid placebo assumption is violated, the proposed PICA-2 estimator has a bias of unknown direction. In this scenario, researchers can use the original bounding approach of de Benedictis-Kessner *et al.* (2019). In Figure 2-(B), we report both PICA-2 estimator (red triangles) and the bounding estimator (blue thick error bars). Due to the violation of the valid placebo assumption, estimates from the two estimators diverge (estimates from the PICA-2 estimator fall outside the estimated bounds), and the bounding estimator is more credible here.

5 Conclusion

These findings underscore the importance of experimental designs that efficiently detect treatment effect heterogeneity. In this application, a crucial substantive insight relies on our ability to estimate the average choice-specific treatment effects (ACTE), which are the conditional average treatment effect for subjects drawn to each type of media treatment. Our proposed PICA-2 design can significantly improve efficiency (that is, achieve the same standard errors with a much smaller sample size) compared to the original PICA design when placebos are valid; researchers can still use the original PICA bounding approach even if placebos prove to be invalid. Thus, our design offers a safe strategy to improve the original design, without losing its substantial benefits.

This placebo control approach seems especially well suited to the study of pro-

⁸A reader of an earlier draft of this paper suggested that the nature videos, by extolling the physical assets of the host country, may have convinced some viewers about the country’s economic vitality. This hypothesis could be tested by varying the content of the placebo. To hedge one’s bets against invalid placebos, it may be wise to assign subjects to alternative placebo conditions, especially during pilot testing.

paganda and other media insofar as persuasive effects have long been thought to vary according to the media tastes of the audience (Arceneaux and Johnson, 2013). Additional analyses reported in Appendix H.3, for instance, suggest that the treatment-by-preference interactions that we found are stronger substantively than the interactions between treatment and the usual demographic covariates such as age, education, or gender. When applied to the great powers' competition to curry public support in Africa, experimental studies that only involve forced exposure may overlook important dynamics of public opinion, notably who seeks exposure to particular sources or topics and, if exposed, who updates their prior views the most. Applying the PICA-2 design, we discover the remarkable effectiveness of Chinese content among those who would not ordinarily watch it. This insight aligns with China's recent strategy of embedding propaganda within entertainment media to engage a broader and more receptive audience (Yao, 2023). Our findings also highlight an opportunity to uncover substantial media effects that might remain hidden when only average treatment effects are considered.

References

- Arceneaux, K. and Johnson, M. (2013). *Changing minds or changing channels?: Partisan news in an age of choice*. University of Chicago Press.
- Balcells, L., Tellez, J. F., and Villamil, F. (Forthcoming). The Wars of Others: The Effect of the Russian Invasion of Ukraine on Spanish Nationalism. *Journal of Politics* <https://doi.org/10.1086/726939>.
- Blair, R. A., Marty, R., and Roessler, P. (2022). Foreign Aid and Soft Power: Great Power Competition in Africa in the Early Twenty-First Century. *British Journal of Political Science* **52**, 3, 1355–1376.
- Bush, S. S. (2015). *The Taming of Democracy Assistance: Why Democracy Promotion Does Not Confront Dictators*. Cambridge University Press.
- de Benedictis-Kessner, J., Baum, M. A., Berinsky, A. J., and Yamamoto, T. (2019). Persuading the Enemy: Estimating the Persuasive Effects of Partisan Media with the Preference-Incorporating Choice and Assignment Design. *American Political Science Review* **113**, 4, 902–916.

- Duarte, G., Finkelstein, N., Knox, D., Mummolo, J., and Shpitser, I. (2023). An Automated Approach to Causal Inference in Discrete Settings. *Journal of the American Statistical Association* , just-accepted, 1–25.
- Fan, Y., Pan, J., and Sheng, J. (2024). Strategies of chinese state media on twitter. *Political Communication* **41**, 1, 4–25.
- Goldsmith, B. E., Horiuchi, Y., and Matush, K. (2021). Does Public Diplomacy Sway Foreign Public Opinion? Identifying the Effect of High-Level Visits. *American Political Science Review* **115**, 4, 1342–1357.
- Green-Riley, N. (2022). *How States Win Friends and Influence People Overseas: The Micro-foundations of U.S. and Chinese Public Diplomacy*. Ph.D. thesis, Harvard Department of Government.
- Knox, D., Yamamoto, T., Baum, M. A., and Berinsky, A. J. (2019). Design, identification, and sensitivity analysis for patient preference trials. *Journal of the American Statistical Association* **114**, 528, 1532–1546.
- Kraft, P. W., Davis, N. R., Davis, T., Heideman, A., Neumeyer, J. T., and Park, S. Y. (2022). Reliable Sources? Correcting Misinformation in Polarized Media Environments. *American Politics Research* **50**, 1, 17–29.
- Markovich, Z., Baum, M. A., Berinsky, A. J., de Benedictis-Kessner, J., and Yamamoto, T. (2020). Dynamic Persuasion: Decay and Accumulation of Partisan Media Persuasion Presented at the Annual Meeting of the Southern Political Science Association, San Juan, Puerto Rico. https://scholar.harvard.edu/files/jdbk/files/multiwave_200111.pdf.
- Mattingly, D., Incerti, T., Ju, C., Moreshead, C., Tanaka, S., and Yamagishi, H. (2024). Chinese Propaganda Persuades a Global Audience That the “China Model” is Superior: Evidence From A 19-Country Experiment. *Amerian Journal of Political Science* .
- Testa, P. F., Williams, T., Britzman, K., and Hibbing, M. V. (2021). Getting the Message? Choice, Self-Selection, and the Efficacy of Social Movement Arguments. *Journal of Experimental Political Science* **8**, 3, 296–309.

Weiss, J. C. (2019). A World Safe for Autocracy? China's Rise and the Future of Global Politics. *Foreign Affairs* **98**, 4, 92–102.

Yao, L. (2023). *Popular Propaganda in Pop Culture: How China Sells Its Ideology*. Ph.D. thesis, Columbia University.

Zhao, S. (2010). The china model: Can it replace the western model of modernization? *Journal of contemporary China* **19**, 65, 419–436.

Online Appendix

Placebo-Augmented PICA Design (PICA-2): Assessing the Influence of Foreign Propaganda

Contents

A Formal Results	1
A.1 Identification of the ACTEs under PICA-2	1
A.2 Falsification Test of the Valid Placebo Treatment	2
B Simulation Study	3
B.1 When the Valid Placebo Assumption Holds.	4
B.2 When the Valid Placebo Assumption is Violated.	5
C Designing Placebo Videos and Choice Questions	7
C.1 General Guide	7
C.2 How PICA-2 Can Be Applied to a Wide Range of Contexts	8
D Ethical Standards	12
E Treatment Assignment	13
F Site Selection Using Synthetic Purposive Sampling	14
F.1 Site Selection	14
F.2 From Site Selection to External Validity Analysis	16
G Validity of the Experiment	18
G.1 Attrition	18
G.2 Attention Checks	18
G.3 Manipulation Check	20
H Full Results	21
H.1 Main Results	21
H.2 Stated vs Revealed Preferences	23
H.3 Treatment Effect Heterogeneity	24

A Formal Results

A.1 Identification of the ACTEs under PICA-2

In this section, we discuss identification results for the ACTE under the placebo-augmented PICA design. To begin with, the ATE is identified using the forced exposure condition (in which subjects are randomly assigned to either treatment or placebo arms) using the usual identification formula.

$$\begin{aligned}\mathbb{E}(Y(t) - Y(t')) &= \mathbb{E}(Y(t)|D = 1) - \mathbb{E}(Y(t')|D = 1) \\ &= \mathbb{E}(Y|T = t, D = 1) - \mathbb{E}(Y|T = t', D = 1),\end{aligned}$$

where the first line follows from $Y(t) \perp\!\!\!\perp D$ and the second from $Y(t) \perp\!\!\!\perp T_1 \mid D = 1$.

We next identify the ACTE of receiving treatment $T = t$ rather than $T = 0$ among those who choose to watch t . First, under the PICA-2 design, we have

$$\begin{aligned}\mathbb{E}(Y(0) \mid C = t) &= \mathbb{E}(Y(0) \mid C = t, D = 0), \\ \mathbb{E}(Y(t_P) \mid C = t) &= \mathbb{E}(Y(t_P) \mid C = t, D = 0).\end{aligned}$$

Therefore, under the valid placebo assumption, we have

$$E(Y(0) \mid C = t, D = 0) = \mathbb{E}(Y(t_P) \mid C = t, D = 0). \quad (1)$$

Using this equality, we can now identify the ACTE.

$$\begin{aligned}\mathbb{E}(Y(t) - Y(0) \mid C = t) &= \mathbb{E}(Y(t) - Y(0) \mid C = t, D = 0) \\ &= \mathbb{E}(Y(t) \mid C = t, D = 0) - \mathbb{E}(Y(0) \mid C = t, D = 0) \\ &= \mathbb{E}(Y \mid T = t, C = t, D = 0) - \mathbb{E}(Y(t_P) \mid C = t, D = 0) \\ &= \mathbb{E}(Y \mid T = t, C = t, D = 0) - \mathbb{E}(Y \mid T = t_P, C = t, D = 0)\end{aligned}$$

where the first equality follows from $Y(t) \perp\!\!\!\perp D \mid C$, the second from the linearity of expectations, and the third from the consistency of potential outcomes and equation (1). The final line follows from $Y(t_P) \perp\!\!\!\perp \mathbf{1}\{T = t_P\} \mid C = t, D = 0$ and the consistency of potential outcomes.

Finally, we identify the ACTE of receiving treatment $T = t$ rather than $T = 0$ among those who do *not* choose to watch t . First, we have the following decomposition based on the law of total expectation.

$$\mathbb{E}(Y(t) - Y(0)) = \mathbb{E}(Y(t) - Y(0) \mid C = t) \Pr(C = t) + \mathbb{E}(Y(t) - Y(0) \mid C \neq t) \Pr(C \neq t),$$

and thus, we have

$$\mathbb{E}(Y(t) - Y(0) \mid C \neq t) = \frac{\mathbb{E}(Y(t) - Y(0)) - \mathbb{E}(Y(t) - Y(0) \mid C = t) \Pr(C = t)}{\Pr(C \neq t)}.$$

We have already identified $\mathbb{E}(Y(t) - Y(0))$ and $\mathbb{E}(Y(t) - Y(0) \mid C = t)$ above. We also have

$$\begin{aligned}\Pr(C = t) &= \Pr(C = t \mid D = 0) \\ \Pr(C \neq t) &= \Pr(C \neq t \mid D = 0).\end{aligned}$$

Therefore, we combine everything;

$$\begin{aligned} & \mathbb{E}(Y(t) - Y(0) \mid C \neq t) \\ = & \frac{\{\mathbb{E}(Y \mid T = t, D = 1) - \mathbb{E}(Y \mid T = 0, D = 1)\} - \{\mathbb{E}(Y \mid T = t, C = t, D = 0) - \mathbb{E}(Y \mid T = t_p, C = t, D = 0)\} \Pr(C = t \mid D = 0)}{\Pr(C \neq t \mid D = 0)}, \end{aligned}$$

which completes the proof.

A.2 Falsification Test of the Valid Placebo Treatment

In this section, we explain our falsification test. Based on the law of total expectation, we have

$$\mathbb{E}(Y(0)) = \mathbb{E}(Y(0) \mid C = 1) \Pr(C = 1) + \mathbb{E}(Y(0) \mid C = 2) \Pr(C = 2) + \mathbb{E}(Y(0) \mid C = 0) \Pr(C = 0).$$

Because of randomization in the PICA-2 design, we also have

$$\begin{aligned} \mathbb{E}(Y(0)) &= \mathbb{E}(Y \mid T = 0, D = 1) \\ \mathbb{E}(Y(0) \mid C = 1) &= \mathbb{E}(Y(0) \mid C = 1, D = 0) \\ \mathbb{E}(Y(0) \mid C = 2) &= \mathbb{E}(Y(0) \mid C = 2, D = 0) \\ \mathbb{E}(Y(0) \mid C = 0) &= \mathbb{E}(Y \mid C = 0, D = 0) \\ \Pr(C = t) &= \Pr(C = t \mid D = 0) \text{ for } t \in \{0, 1, 2\}. \end{aligned}$$

Therefore, under the PICA-2 design, without additional assumptions, we have the following equality, while the right-hand side is not identified.

$$\begin{aligned} & \mathbb{E}(Y \mid T = 0, D = 1) \\ = & \mathbb{E}(Y(0) \mid C = 1, D = 0) \Pr(C = 1 \mid D = 0) + \mathbb{E}(Y(0) \mid C = 2, D = 0) \Pr(C = 2 \mid D = 0) \\ & + \mathbb{E}(Y \mid C = 0, D = 0) \Pr(C = 0 \mid D = 0) \end{aligned}$$

Finally, when the valid placebo assumption holds, the following equality holds.

$$\begin{aligned} & \mathbb{E}(Y \mid T = 0, D = 1) \\ = & \mathbb{E}(Y \mid T = 1_P, C = 1, D = 0) \Pr(C = 1 \mid D = 0) + \mathbb{E}(Y \mid T = 2_P, C = 2, D = 0) \Pr(C = 2 \mid D = 0) \\ & + \mathbb{E}(Y \mid C = 0, D = 0) \Pr(C = 0 \mid D = 0) \end{aligned} \tag{2}$$

where we use equation (1) under the valid placebo assumption and randomization of the treatment and placebo conditional on C in the free choice condition.

We test this equality (equation (2)) as a falsification test of the valid placebo assumption.

Finally, it is important to note that this equality can be violated when the assumption of no design effect (D has no direct effect on outcomes) is violated as well. Therefore, in general, the proposed falsification test can be seen as an omnibus test of the valid placebo and the no design effect assumption.

B Simulation Study

In this section, we rely on the simulation study of Knox et al. (2019) to evaluate the performance of the proposed placebo-augmented PICA design (PICA-2) and its corresponding placebo-augmented (PA) estimator.¹ Please see Section 8 of Knox et al. (2019) for details about the simulation setting.

We compare the original PICA design (its data-generating process is exactly the same as the simulation study in the original PICA paper) and our PICA-2 design. our PICA-2 data generating process is the same as the simulation study in the original PICA paper, except for how we generate outcomes under placebo arms in the free choice condition. We consider two scenarios: (i) when the valid placebo assumption holds and (ii) when the valid placebo assumption is violated. When the valid placebo assumption holds, the data-generating process for placebo arms in the free-choice condition is the same as for placebo arms in the forced exposure condition. Importantly, when the valid placebo assumption is violated, the data-generating process for placebo arms in the free-choice condition deviates from the data-generating process for placebo arms in the forced exposure condition. We first draw the error indicator $R_i \sim \text{Binomial}(v)$. When $R_i = 1$, we observe biased outcomes $Y_i = Y_i^*/2$ where Y_i^* is drawn from the data-generating process for placebo arms in the forced exposure condition. When $R_i = 0$, we observe unbiased outcomes $Y_i = Y_i^*$. This simple setup captures the violation of the valid placebo assumption, and v captures the extent of the violation.

¹We thank Dean Knox for his help in replicating their simulation study.

B.1 When the Valid Placebo Assumption Holds.

We consider settings where the valid placebo assumption holds (in our simulation setting, we set $v = 0$). In particular, we compare our placebo-augmented (PA) estimators under the new placebo-augmented design and the bound estimator under the original PICA design. In this scenario, our PA estimator and the bounding estimator (de Benedictis-Kessner et al., 2019; Knox et al., 2019) are both valid (unbiased and have valid confidence intervals), but our PA estimators have much shorter confidence intervals.

Following Knox et al. (2019), we consider seven different simulation scenarios where we vary the outcome model divergence ($OD \in \{0, 1/3, 2/3, 1\}$) and the divergence between stated and actual media ($CD \in \{0, 1/3, 2/3, 1\}$). Please see Knox et al. (2019) for the details of the simulation setting. We also vary sample size $n \in \{1500, 4500, 9000\}$. For coverage, we average over quantities over four ACTEs, i.e., $\mathbb{E}(Y(1) - Y(0) \mid C = 1)$, $\mathbb{E}(Y(1) - Y(0) \mid C \neq 1)$, $\mathbb{E}(Y(2) - Y(0) \mid C = 2)$, and $\mathbb{E}(Y(2) - Y(0) \mid C \neq 2)$. Figure B1 reports the length of the proposed PA estimators' confidence intervals as a proportion of the length of the bound estimators' confidence intervals. We find that, across simulation settings, the PA estimator can reduce standard errors substantially. The length of confidence intervals ranges from 10% to 40% of those from the bound estimator.




Figure B1: When the Valid Placebo Assumption Holds.

B.2 When the Valid Placebo Assumption is Violated.

We consider settings where the valid placebo assumption is violated (in our simulation setting, we set $v = 0.30$; the main findings do not differ depending on this parameter as long as $v > 0$). In this scenario, the placebo-augmented (PA) estimator is biased, and thus, we recommend using the bound estimator (de Benedictis-Kessner et al., 2019; Knox et al., 2019). We compare the bound estimator under the new placebo-augmented design and the same bound estimator under the original PICA design. An important point is that placebo arms under our new placebo-augmented design cannot be used when the valid placebo assumption is violated, so the bound estimator under the new placebo-augmented design is expected to have larger standard errors and longer confidence intervals.

Again, following Knox et al. (2019), we consider seven different simulation scenarios where we vary the outcome model divergence ($OD \in \{0, 1/3, 2/3, 1\}$) and the divergence between stated and actual media ($CD \in \{0, 1/3, 2/3, 1\}$). Please see Knox et al. (2019) for the details of the simulation setting. We also vary sample size $n \in \{1500, 4500, 9000\}$. For coverages, we average over quantities over four ACTEs, i.e., $\mathbb{E}(Y(1) - Y(0) \mid C = 1)$, $\mathbb{E}(Y(1) - Y(0) \mid C \neq 1)$, $\mathbb{E}(Y(2) - Y(0) \mid C = 2)$, and $\mathbb{E}(Y(2) - Y(0) \mid C \neq 2)$.

Figure B2 reports the length of the bound estimators' confidence intervals under our new PICA-2 design as a proportion of the length of the bound estimators' confidence intervals under the original PICA design. We find that, across simulation settings, the bound estimator under the PICA-2 design increases standard errors only slightly. The length of confidence intervals ranges from 100% to 105% of those from the bound estimator under the original PICA design. In practice, this means that, even when the valid placebo assumption is violated, efficiency loss from using the proposed design is relatively small.

This small loss in efficiency comes from the fact that the placebo-augmented design enables researchers to estimate most of the key ingredients of the ACTE bounds as well as the original PICA design. Suppose we are interested in $\mathbb{E}(Y(a = 1) - Y(a' = 0) \mid C = 1)$. Looking at the bound in Proposition 1 of Knox et al. (2019), researchers can estimate the following quantities under the placebo-augmented design with the same precision as under the original PICA design. $F(y \mid s, a', 1)$, $\Pr(a' \mid s, 0)$, $\Pr(A_i = c \mid S_i = s, D_i = 0)$, $F(y \mid s, a, 1)$, $F(y \mid s, a', 0)$, $\Pr(a \mid s, 0)$, and $\Pr(S_i = s \mid A_i = c, D_i = 0)$. The only change is the estimation of $F(y \mid s, a = 1, 0)$. Under the augmented-placebo design, we have the half-sample size to estimate this quantity.

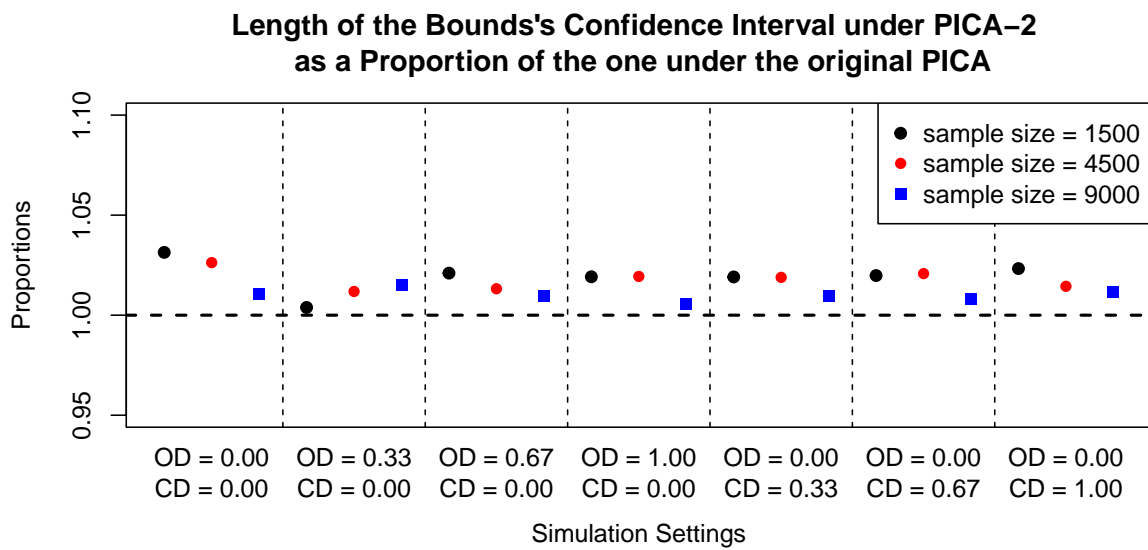


Figure B2: When the Valid Placebo Assumption is Violated.

C Designing Placebo Videos and Choice Questions

Here, we provide practical guidance about how to design placebo videos and choice questions. We provide a general logic below and offer several concrete examples in the next subsection.

C.1 General Guide

When designing placebo videos and choice questions, researchers have to consider two assumptions — the exclusion restriction and valid placebo assumptions. Recall the valid placebo assumption states that respondents exhibit the same potential outcomes under a placebo video and a control video. Intuitively, this means that placebo videos should be as similar as possible to the control video we show in the forced exposure condition. On the other hand, the exclusion restriction assumption requires that respondents exhibit the same potential outcomes regardless of whether they are in the free choice condition or in the forced exposure condition. Intuitively, this means that watching a certain video will lead to the same outcomes regardless of whether they choose to watch it or they are forced to watch it. For example, this assumption is more likely to be violated when respondents become disappointed because a video they choose to watch is less interesting than what they expected from the choice.

Let us think about these requirements by first considering extreme scenarios. If researchers just want to make the valid placebo assumption plausible, they can use the control video in the forced exposure condition as the placebo video. For example, we could show a general nature video as a placebo even for those who choose to watch “Something about China” or “Something about the United States” in the free choice condition. This will make the valid placebo assumption true by design. However, this placebo video may violate the exclusion restriction because respondents may be dismayed when they choose to watch “Something about China” yet see something very different (a general nature video). Therefore, researchers should strive to show videos that are consistent with what respondents expect to watch. On the other hand, if we make our placebo Chinese (US) videos too similar to the treatment videos and too different from a general nature video (the control video in the forced exposure condition), the valid placebo assumption may be jeopardized.

In sum, placebo videos should be as similar as possible to the control video in the forced choice condition, and placebo videos should be consistent with what respondents expect to watch given their choice. In our application, we made our placebo videos (Chinese and US nature videos) as similar as possible to the control video (a general nature video) by making all of them about nature. At the same time, to make sure placebo videos are consistent with their actual choice, we ask whether respondents want to watch “Something about China,” “Something about the United States,” or “Something about the World.” In this way, both Chinese (US) propaganda videos and Chinese (US) nature videos are natural to respondents who choose to watch “Something about China” (“Something about the United States”). Notice that the choice of the placebo subject matter – nature, in our application – is based on testable hypotheses about what content will affect the outcome. For example, it is possible that praising the natural beauty of China or the United States suggests to respondents the merits of each country’s economic system. Had we selected other placebo themes, such as cuisine, folk art, or dinosaur fossils, perhaps the risk of inadvertent effects might have been reduced. Pilot studies may provide an early indication of whether placebos are working as expected or need adjustment.

C.2 How PICA-2 Can Be Applied to a Wide Range of Contexts

To illustrate the breadth of PICA-2 applications, we describe below how recent PICA studies could incorporate a placebo into their designs. Then we turn to studies that did not use a preference incorporating design, and for the purposes of illustration, suggest potential extensions using the PICA-2 design. These are intended only to illustrate how PICA-2 could be used in practice, and are not intended as critiques of these studies, and we alone are responsible for the content below.

How PICA Studies Could Use PICA-2

Article	Choice	Treatment	Control	Suggested Treatment Placebo
de Benedictis-Kessner, Justin, Matthew A. Baum, Adam J. Berinsky, and Teppei Yamamoto. "Persuading the enemy: Estimating the persuasive effects of partisan media with the preference-incorporating choice and assignment design." <i>American Political Science Review</i> 113, no. 4 (2019): 902-916.	Fox News, MSNBC, The Food Network	Partisan news articles (from Fox News and MSNBC) discussing either the economic, social, safety, or public health effects of legalizing marijuana	The Food Network article discussing how to save money while grocery shopping, tips for buying meat, and how grocery stores might change in the future.	Partisan news articles discussing how to save money while grocery shopping
Markovich, Zachary, Matthew A. Baum, Adam J. Berinsky, Justin de Benedictis-Kessner, and Teppei Yamamoto. "Dynamic persuasion: decay and accumulation of partisan media persuasion." (2020). Working Paper.			same as above	
Balells, Laia, Juan F. Tellez, and Francisco Vilamil. "Past conflict, media, and polarization in Spain." (2022). Working Paper.	El Pais, ABC, 20 Minutes	Partisan news articles (from El Pais or ABC) discussing transitional justice and redistribution	The article from a neutral outlet (20 Minutes) discussing the homecoming of those displaced by the eruption of the La Palma volcano	The exact same news article from 20 Minutes, with the logos from the partisan news outlets (El Pais and ABC)

Here, we illustrate three recent experimental studies that adopted the original PICA design. Both de Benedictis et al (2019) and Markovich et al. (2020) are interested in estimating the persuasive effects of partisan media and investigating the extent to which pre-existing media preferences affect the treatment effects. In the free-choice arm, both studies allow subjects to choose between three media choices—Fox News, MSNBC, and Food Network. Partisan news articles discuss the economic, social, safety, or public health effects of legalizing marijuana. The Food Network article discusses how to save money while grocery shopping. In light of these messaging features of the original study, a potential placebo article could be partisan news content that delivers exactly the same content that the Food Network article has.

We suggest this in light of two assumptions central to the PICA-2 design: the exclusion

restriction and valid placebo assumptions. Using the same news content about a neutral topic unrelated to the proposed causal mechanism (which is about the biased framing of partisan media) but only with a different outlet logo ensures the content will likely lead to the same outcome. Granted, something about money-saving tips during grocery shopping could be interpreted as a reminder of the inflation, which then invites partisan interpretation of the incumbent’s failed economic policies. Yet, we consider it a remote possibility that can be readily addressed by tweaking the content if needed. At the same time, the proposed treatment placebo wouldn’t be off-putting or disappointing for those who choose to consume Fox or MSNBC; at the end of the day, they are shown an article from the outlet of their choice.

Our proposed treatment placebo follows the exact same logic for extending Balells, Tellez, and Villamil (2022). Here, the authors are interested in the impact of partisan media on polarization in Spain. Participants are given three outlet choices (El Pais, ABC, and 20 Minutes) and can choose once assigned to the free choice condition. The control is an article from a neutral outlet (20 Minutes) discussing the homecoming of those displaced by the eruption of the La Palma volcano. The placebo can feature the exact same content, with the logos from the two partisan news outlets.

How To Extend Existing Forced-Exposure Studies Using PICA-2

We now turn to six recent experimental studies that did not use a preference incorporating design, and for illustration, suggest potential extensions using the PICA-2 design. These are intended only to illustrate how PICA-2 could be used in practice, and are not intended as critiques of these studies, and we alone are responsible for the content below.

Studies that could benefit from news vs entertainment preference-based heterogeneity:

- Aytaç, Selim Erdem. (2021). “Effectiveness of incumbent’s strategic communication during economic crisis under electoral authoritarianism: Evidence from Turkey.” *American Political Science Review*.
- Culpepper, Pepper D., Jae-Hee Jung, and Taeku Lee. (2023). “Banklash: How Media Coverage of Bank Scandals Moves Mass Preferences on Financial Regulation.” *American Journal of Political Science*.
- Grewal, Sharan, and Shadi Hamid. (2022). “Discrimination, Inclusion, and Anti-System Attitudes among Muslims in Germany.” *American Journal of Political Science*.
- Siegel, Alexandra A., and Vivienen Badann. (2020). “#No2Sectarianism: Experimental Approaches to Reducing Sectarian Hate Speech Online.” *American Political Science Review*.

Across the four studies, researchers provide news-like information treatment and measure relevant attitudinal outcomes afterward. In a high-choice media environment, there is no guarantee that citizens would voluntarily choose to consume political news in the real world. In other words, the treatment effects may differ by pre-existing news versus entertainment content preference.

For example, Aytaç hypothesizes that to avoid blame during economic downturns, authoritarian rulers use pro-government media to either divert blame for poor economic performance

or shift the agenda from the economy to another area. The original research design presents vignettes in the style of pro-government media either diverting blame or attempting to shift the agenda. A potential extension of Aytac (2021) might entail offering an option of choosing either pro-government news media or entertainment media in a free choice condition. If the treatment is a news story about economic conditions that attributes blame to foreign actors, the treatment placebo could be an article that contains only objective economic statistics from the same outlet, without any attribution.

Similarly, Culpepper et al. (2023) hypothesize that “media coverage of bank scandals leads people to express attitudes that are more supportive of regulating big banks” (p. 6). Their research design randomly exposes audiences to a (simulated) media story discussing a recent bank scandal. An extension of their study could add a free choice arm that includes entertainment media or non-business media. For those who choose to consume business news, in the free choice arm, the treatment placebo in the augmented arm could be a business story about the banking sector with bank scandal information removed.

Grewal and Hamid (2022) and Siegel and Badaan (2020) show that exposure to social media feeds from religious and political leaders can reduce support for sectarian hate and for anti-system attitudes. An extension of these studies could examine whether the effects of this media differ by whether people are likely to view social media feeds of relevant religious and political leaders. In the Grewal and Hamid (2022) study, for example, a treatment is a statement from Angela Merkel containing inclusive rhetoric; in a PICA-2 extension, a treatment placebo could be an otherwise identical social media feed of Merkel with the inclusive rhetoric post removed. A similar design could be applied to an extension of Siegel and Badaan (2020).

Studies that could benefit from within-genre/within-platform content preference-based heterogeneity:

- Mullin, Megan, and Katy Hansen. (2023). “Local News and the Electoral Incentive to Invest in Infrastructure.” *American Political Science Review*.
- Kim, Eunji. (2023). “Entertaining Beliefs in Economic Mobility.” *American Journal of Political Science*.

Another strand of media effects research could benefit from incorporating within-genre and within-platform content preferences. Mullin and Hansen (2023) examine whether local news coverage increases support for investment in infrastructure projects. The study manipulated whether local officials were told that infrastructure needs had been covered in the local press. In an extension of this study, it might be pertinent to assess whether respondents prefer reading local or national news. Given that the main focus is on individuals already predisposed to consuming news – specifically, US city and county elected officials – this distinction could be crucial in understanding average treatment effects. A logical extension might involve querying respondents about their specific news preferences, then manipulating whether the story contained information about infrastructure projects. For instance, a local news story on mental health programs (as opposed to a more visible local infrastructure topic) could serve as a treatment placebo. Meanwhile, a national news story on mental health programs might serve as a control.

Kim's study (2023) examines how rags-to-riches narratives in entertainment influence beliefs about economic mobility. The study design involved randomly varying whether respondents were exposed to a reality TV show with a rags-to-riches storyline. An extension could include a question that taps into participants' pre-existing preference for different entertainment genres (reality TV, drama, or comedy). As a potential treatment placebo, a reality TV show without a rags-to-riches narrative could be employed, while a drama or comedy clip showcasing mundane everyday life could serve as a control.

We recognize that when applying the PICA framework to studies on misinformation or conspiracy theories, challenges arise. Directly questioning respondents about their tastes for misinformation or conspiracy theories may produce unreliable results. However, a potential workaround is to inquire about their general news source preferences, such as mainstream media versus alternative sources. From there, researchers can introduce accurate information from less conventional sources as a treatment placebo.

D Ethical Standards

This project was reviewed by the Institutional Review Board (IRB) of the authors' universities. However, IRB approval determinations are only a starting point for meeting ethical standards for conducting research. We took a number of steps to protect respondents from harm and to meet, and exceed, the guidelines set forward by the Belmont Report of the National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research.

First, respondents to the survey provided informed consent. The survey platform, Cint, also provided compensation to participants for their time.

Second, the information we collected on attitudes towards international politics is not generally politically sensitive in the contexts we surveyed. We did not ask respondents potentially sensitive questions about their domestic context.

Third, the data were collected anonymously through Qualtrics, an encrypted platform, and the data were stored locally on secure machines. The data we collected do not include identifiers that would make it possible to link participants to their responses.

Fourth, we prominently labelled the source of government treatments as “produced by the Chinese government” or “produced by the American government.” In the case of the CGTN videos, this labelling goes beyond what viewers of CGTN would have seen on videos broadcast in the United States by major cable networks such as Comcast and AT&T.

Fifth, the treatments did not include misinformation or deception. Each treatment included factual information, although each was slanted in favor of the country depicted in the treatment.

Sixth, though the treatments did not include deception, we nevertheless at the conclusion of the survey debriefed participants and provided them with the opportunity to read opposing viewpoints. We provided access to reports from Amnesty International and the OECD on the human rights and economic situation of each country.

E Treatment Assignment

Forced Exposure Condition

For those who are assigned to a [**forced exposure condition**], they are randomly assigned to one of three conditions with equal probability:

- 1) Two Chinese government-produced videos (Two videos will be chosen out of four: China - Economic Model 1, 2, China - Political Model 1, 2)
- 2) Two U.S. government-produced videos (Two videos will be chosen out of four: US - Domestic Economy 1, 2, US - Domestic Economy 1, 2),
- 3) A placebo condition with two nature videos (Control - Nature, Control - Dolphins)

Free Choice Condition (Placebo Augmented)

For those who are assigned to a [**free choice condition**], they had the option of choosing one of three conditions with equal probability:

- 1) China
- 2) US
- 3) Nature

For those who are in [**free choice condition - China**] arm, they are randomly assigned to one of two conditions with equal probability:

- 1) Two Chinese government-produced videos (Two videos will be chosen out of four: China - Economic Model 1, 2, China - Political Model 1, 2)
- 2) China (Nature) (China - Nature - Nonggang Babbler, China - Nature - A Sea of Bamboo)

For those who are in [**free choice condition - US**] arm, they are randomly assigned to one of two conditions with equal probability:

- 1) Two U.S. government-produced videos (Two videos will be chosen out of four: US - Domestic Economy 1, 2, US - Domestic Economy 1, 2)
- 2) US (Nature) (US - Nature - Ancient Rocks, US - Nature - Wonder of US Parks)

For those who are in [**free choice condition - Nature**] arm, they will watch two nature videos (Control - Nature, Control - Dolphins)

To download and view the videos, please click this link. Please note that clicking the link will initiate a download of about 200 MB of files.

F Site Selection Using Synthetic Purposive Sampling

F.1 Site Selection

In this study, we use a multi-site experiment in order to improve external validity with respect to contexts. To select sites, we use synthetic purposive sampling (SPS). SPS can select diverse sites such that the non-selected observations can be well approximated by the weighted average of the selected sites, while accommodating practical constraints researchers have. By doing so, without random sampling, SPS can make the weighted average of selected sites representative of the target population of sites.

First, we define the target population of sites. In particular, respecting our practical constraints, we use the following three eligibility criteria; (a) African countries, (b) in which English is one of the top 3 official languages, (c) data on the UN’s ideal point alignment with China and the US exist, and (d) population size is equal to or greater than 1,000,000. As for Condition (a), we choose Africa as the main context as it is one of the main targets of recent Chinese communication campaigns (Mattingly et al., 2024). Conditions (b), (c), and (d) come from our practical constraints: our video treatments were created in English, and we need measures of relationships with China and the US for adequate political representativeness, and we wanted to make sure that our experiment is done in relatively large countries so that we would have ample numbers of survey respondents. In total, we have 21 African countries as the target population of sites (Botswana, Cameroon, Eritrea, Ethiopia, Ghana, The Gambia, Kenya, Liberia, Lesotho, Mauritius, Malawi, Namibia, Nigeria, Sudan, Sierra Leone, South Sudan, Eswatini, Uganda, South Africa, Zambia, and Zimbabwe).

Second, using SPS, we diversified 5 key contextual factors; GDP per capita, Polity Score, Ethnic Fractionalization, Relationships with China and US, and Sub-Regions in Africa, within this practical constraint. We stratified SPS by making sure to have at least one country from each sub-region (Eastern, Middle, South, and West Africa), making sure to have at least one country below -0.5 and above 0.5 standard deviations of each variable, and selecting countries that have Cloud Research panels (only 8 countries have Cloud Research panels). Within these constraints, SPS optimally selects diverse sites such that non-selected sites can be well approximated by the weighted average of selected sites.

SPS selected five countries: Botswana, Cameroon, Nigeria, Uganda, and Zambia. These five countries successfully cover a wide range of values in all of the five key contextual factors. See Figure F1. It is also important to emphasize that SPS (or any site selection approach) cannot perfectly diversify site selection given the set of practical constraints. For example, we did not select some seemingly useful sites (e.g., sites with "Ethnic Fractionalization" lower than -2 and sites with "Distance to China" larger than 2) because these sites do not offer Cloud Research panels and SPS took this practical constraint into account.

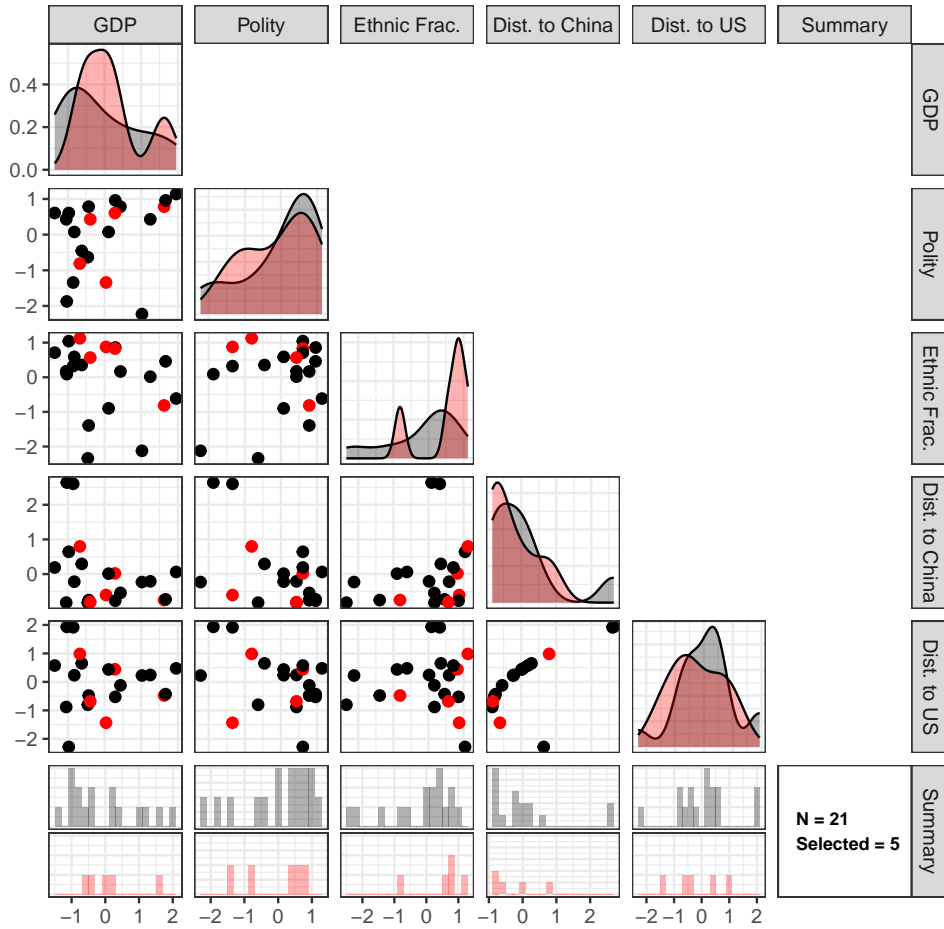


Figure F1: Distribution of Site-Level Variables in Selected Sites.

Note: Red circles represent selected sites and black circles represent non-selected sites. In the last row and the diagonal plots, we can see the marginal distribution of each variable. All the remaining figures in the middle show bivariate relationships between two variables. We can see that selected sites successfully cover a wide range of values in each site-level variable.

F.2 From Site Selection to External Validity Analysis

The first step is to report ATE estimates in each selected site. See Figure F2. Importantly, even though we optimally diversified study sites via SPS, we find that the ATEs of each treatment-outcome relationship are relatively stable across countries.

We can estimate the average-site ATEs using the SPS estimator. See Figure F3. We find that the effects of Chinese Propaganda are large and statistically significant. In particular, the effects on preferences for Chinese political models are large. The effects of US Propaganda are more ambiguous: the effects on preferences for Chinese political models are relatively small, while the effects on preferences for Chinese economic models are larger and statistically significant. Overall, the causal effects of Chinese propaganda have high external validity on both outcomes, whereas we find that the causal effects of US propaganda have high external validity only on preferences for economic models.

Finally, we can assess the influence of unobserved moderators using site-level cross-validation. In particular, we can randomly choose half of the selected sites as if they were unobserved non-selected sites and predict the average ATE of those non-selected sites based on the remaining selected sites. By repeating the same procedure many times, we can empirically check how well the SPS estimator can credibly infer the ATEs in non-selected sites. For the causal effects of US Propaganda on preferences for political models (economic models), we estimated p-values to be 0.11 (0.60), finding no evidence of significant bias. For the causal effects of Chinese Propaganda on preferences for political models (economic models), we estimated p-values to be 0.39 (0.82), finding no evidence of significant bias.

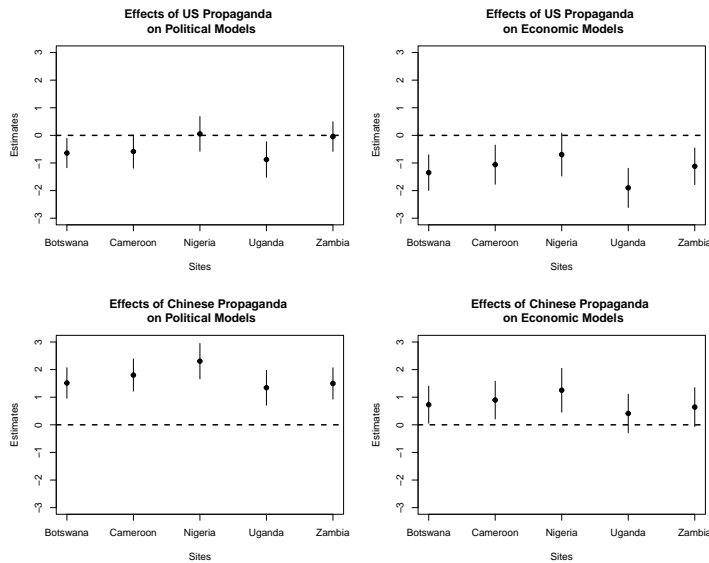


Figure F2: Site-Specific ATEs in Selected Sites. *Note:* In each plot, we report point estimates and 95% confidence intervals for site-specific ATEs.

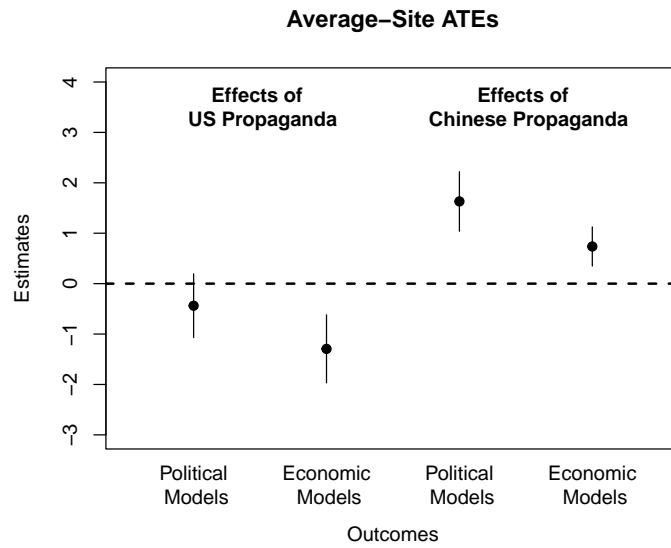


Figure F3: Average-Site ATEs. *Note:* For each treatment-outcome pair, we report point estimates and 95% confidence intervals for the average-site ATEs.

G Validity of the Experiment

G.1 Attrition

We examine whether there is differential attrition between the forced choice and free choice groups; between the placebo and treatment videos among those who choose to watch videos produced by China; and between the placebo and treatment videos among those who choose to watch videos produced by the US. Table G1 computes differences across groups using OLS with robust standard errors where the outcome is attrition. We find no evidence of differential attrition across any of these conditions.

	Outcome: Attrition				
	(1)	(2)	(3)	(4)	(5)
Forced v. Choice	0.016 (0.017)				
China Forced v. Forced Placebo		0.039 (0.033)			
US Forced v. Forced Placebo			0.022 (0.032)		
China Choice v. China Choice Placebo				0.041 (0.071)	
US Choice v. US Choice Placebo					-0.048 (0.045)
Intercept	0.178 (0.011)	0.174 (0.022)	0.174 (0.022)	0.132 (0.047)	0.220 (0.033)
N	2156	575	574	105	316
R2	0.000	0.002	0.001	0.003	0.004
R2 Adj.	0.000	0.001	-0.001	-0.006	0.001
SE Type	HC2	HC2	HC2	HC2	HC2

Table G1: No Differential Attrition by Treatment Status

G.2 Attention Checks

In light of evidence of decreased attention in online samples (Peyton, Huber and Coppock, 2022), respondents were screened according to pre-treatment attention checks and dropped from the sample if they failed the attention check. Our attention checks take the following form:

“People have different tastes in movies. For this question, however, we are not interested in your taste but want to test whether survey takers are reading questions carefully. Below, please select the options “Romance” and “Science Fiction.” The answer choices were: Action Adventure; Romance; Comedy; Science fiction; None of the above.

“For our research, careful attention to survey questions is critical! To show that you are paying attention, please select ‘I have a question.’” The answer choices were: I understand; I do not understand; I have a question.

The survey was designed in a way that if Actions adventure is selected OR comedy is selected OR None of the above is selected AND I have a question is not selected, then the survey ended before seeing any treatment.

Out of 1,737 respondents, 10.36% (N=180) didn’t pass the first attention question. In most cases, however, it was because they only chose one correct answer instead of two. 68 respondents only picked “Romance,” while 95 respondents only picked Science Fiction. 17 respondents

didn't pick any answer. Respondents performed better with the second attention check, with a passing rate of 96.09% (N=1,660). In the end, a total of 1,526 respondents passed both questions. We examined whether attention check failure is correlated with our pre-treatment variables. As displayed in Table G2, the only variable that was systematically correlated with passing the first attention check question was education. While the more educated were more likely to pass the question, such a pattern did not occur for the second question. We report the full results including those who did not pass the attention check questions in Table G3. The substantive conclusions remain the same.

	(1)	(2)	(3)
	Pass Both Q	Pass Q1	Pass Q2
Female	0.017 (0.016)	0.006 (0.015)	0.019 (0.010)
Education	0.024* (0.011)	0.023* (0.010)	0.009 (0.007)
Ideology	0.005 (0.007)	0.002 (0.006)	0.000 (0.004)
Age	-0.002 (0.001)	-0.002 (0.001)	0.001 (0.001)
National pride	0.022 (0.014)	0.009 (0.013)	0.013 (0.010)
N	1731	1731	1731
R2	0.881	0.899	0.957
R2 Adj.	0.880	0.898	0.957
SE Type	HC2	HC2	HC2

Table G2: Attention Check Results

	Political Model		Economic Model		World Leader	
	Estimate	Pr(> t)	Estimate	Pr(> t)	Estimate	Pr(> t)
US vs Nature	-0.38	0.00	-1.21	0.00	-0.61	0.00
China vs Nature	1.74	0.00	0.86	0.00	1.23	0.00
Effect of US video on US-viewer	-0.05	0.74	-0.56	0.01	0.28	0.15
Effect of China video on China-viewer	0.86	0.02	0.97	0.00	0.35	0.31
Effect of US video on non-US-viewer	-0.48	0.00	-1.42	0.00	-0.89	0.00
Effect of China video on non-China-viewer	1.82	0.00	0.85	0.00	1.31	0.00

Table G3: Full results including those who didn't pass the attention check question

G.3 Manipulation Check

We also conducted manipulation checks. After the respondents were shown videos and answered all the outcome questions, they were asked: “What were the two videos you watched about? You may select multiple answers.”

1. There was a problem and I could not watch the videos.
2. There were one or more videos about nature.
3. There were one or more videos about American society, politics, or economics.
4. There were one or more videos about Chinese society, politics, or economics.
5. There were one or more videos about Japanese society, politics, or economics.

As Table G4 shows, the overwhelming majority of respondents correctly recognized the type of videos they watched (90.6%). For those who were assigned to the placebo conditions (those who either watched videos about Chinese or American nature in a free-choice condition), we consider both answers to be correct.

	Chinese Propaganda			US Propaganda			Nature	
	Forced	Free	Nature (Free)	Forced	Free	Nature (Free)	Forced	Free
mc_problem	0.02	0.07	0.04	0.01	0.02	0.04	0.02	0.02
mc_nature	0.06	0.12	0.85	0.09	0.13	0.94	0.99	0.97
mc_american	0.06	0.05	0.00	0.99	0.97	0.15	0.00	0.01
mc_chinese	0.98	0.93	0.28	0.05	0.07	0.01	0.00	0.01
mc_japanese	0.04	0.00	0.02	0.10	0.05	0.00	0.00	0.00

Table G4: Manipulation Check

H Full Results

H.1 Main Results

	Political Model		Economic Model		World Leader	
	Estimate	Pr(> t)	Estimate	Pr(> t)	Estimate	Pr(> t)
US vs Nature	-0.40	0.00	-1.21	0.00	-0.67	0.00
China vs Nature	1.70	0.00	0.79	0.00	1.18	0.00
Effect of US video on US-viewer	0.03	0.85	-0.50	0.03	0.41	0.05
Effect of China video on China-viewer	1.04	0.01	0.94	0.01	0.55	0.13
Effect of US video on non-US-viewer	-0.53	0.00	-1.42	0.00	-1.00	0.00
Effect of China video on non-China-viewer	1.76	0.00	0.78	0.00	1.23	0.00

Table H1: Full Results

	<i>Dependent Variables:</i>		
	Political Model	Economic Model	World Leader
US Propaganda	-0.402*** (0.135)	-1.209*** (0.161)	-0.672*** (0.154)
Chinese Propaganda	1.701*** (0.137)	0.792*** (0.163)	1.176*** (0.156)
Controls:			
Female	0.137 (0.115)	0.035 (0.136)	-0.149 (0.131)
Education	-0.068 (0.072)	-0.049 (0.086)	-0.009 (0.082)
National Pride	-0.039 (0.098)	0.075 (0.117)	0.044 (0.112)
Ideology	0.021 (0.045)	0.015 (0.054)	0.016 (0.052)
Cameroon	-0.216 (0.179)	0.309 (0.213)	-0.253 (0.204)
Nigeria	-0.293 (0.180)	0.067 (0.215)	-0.469** (0.205)
Uganda	-0.129 (0.188)	0.109 (0.224)	-0.032 (0.215)
Zambia	-0.289* (0.175)	0.140 (0.208)	-0.152 (0.200)
Constant	3.110*** (0.462)	3.780*** (0.550)	3.211*** (0.529)
N	594	596	588
R ²	0.316	0.213	0.209
Adjusted R ²	0.304	0.199	0.195

*p<0.1; **p<0.05; ***p<0.01

Table H2: Full Regression Results on the ATE estimation

After asking respondents about their preferred political and economic model, we measured which country they would prefer to the foremost country in the world. We ask respondents: “If you were to choose, which country would you like to be the world leader: China or the United States?” Respondents were offered a six point scale ranging from “strongly prefer China” to “strongly prefer the U.S.” The falsification test for this outcome measure is sig-

nificant ($p = 0.0389$). The results affirm the basic pattern for the other outcome measures. The estimated ATEs suggest that, on average, Chinese propaganda moves respondents more than US propaganda. Both sources of propaganda are much more effective among viewers who would ordinarily prefer not to watch this content.

Ceiling Effects

We also investigate whether the smaller effects of Chinese propaganda on preferences toward the Chinese political model among those who choose to watch something about China compared to those who choose to watch other videos is due to ceiling effects. Figure H1 report the histogram of outcomes among those who choose to watch “Something about China” in the placebo arm and in the treatment arm. The placebo arm contains an ample number of respondents whose outcomes are less than 6, which suggests that ceiling effects cannot fully explain the smaller effects among those who choose to watch something about China.

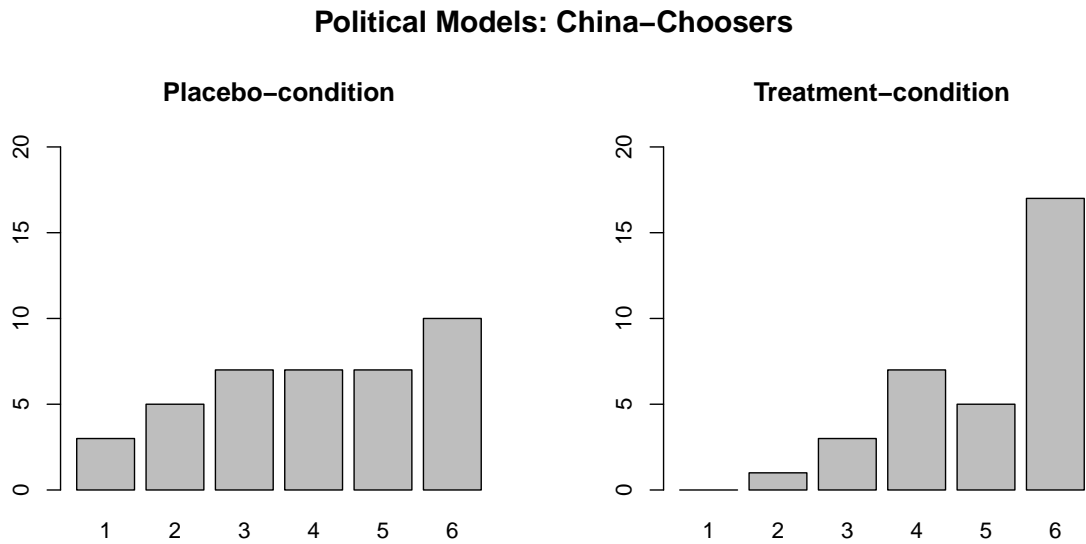


Figure H1: Histogram of Outcomes.

H.2 Stated vs Revealed Preferences

Here, we offer descriptive statistics on respondents' stated and revealed media preferences. While we provide the stated media preferences for the entire sample in Table H3, we can explore the extent to which they diverge from revealed media preferences only among respondents assigned to the free choice arm. As illustrated in Table H4, of the 895 respondents, 789 (88.2%) selected the video that matched their initial stated interest.

	China	US	World	Total N
Botswana	15	39	254	308
Cameroon	24	74	199	297
Nigeria	19	56	230	305
Uganda	24	60	199	283
Zambia	29	43	228	300
Total N	111	272	1110	1493

Table H3: Stated Media Preferences by Country (All Sample)

Revealed Preference	Stated Media Preference			Total N
	China	US	World	
Something about China	49	5	18	72
Something about the United States	5	147	57	209
Something about the World	9	12	593	614
Total N	63	164	668	895

Table H4: Difference Between Stated and Revealed Media Preferences

H.3 Treatment Effect Heterogeneity

Below, we show treatment effect heterogeneity by individual-level variables, i.e., the differences in the conditional ATEs by individual-level covariates. We find suggestive evidence that media preferences rank among the most important sources of treatment effect heterogeneity. This clarifies the importance of PICA design that can uniquely allow researchers to explore treatment effect heterogeneity by media preferences.

	Estimate	Standard Errors
Treatment Effect Heterogeneity of US Propaganda by :		
Media Preference	0.56	0.27
Gender	0.12	0.25
Education	0.09	0.29
National Pride	-0.27	0.27
Ideology	-0.95	0.38
Age	-0.10	0.25
Treatment Effect Heterogeneity of Chinese Propaganda by :		
Media Preference	-0.72	0.40
Gender	-0.14	0.28
Education	0.12	0.33
National Pride	-0.21	0.33
Ideology	-0.93	0.41
Age	-0.33	0.28

Table H5: Treatment Effect Heterogeneity by Individual-level Variables

References

- de Benedictis-Kessner, Justin, Matthew A Baum, Adam J Berinsky and Teppei Yamamoto. 2019. “Persuading the Enemy: Estimating the Persuasive Effects of Partisan Media with the Preference-Incorporating Choice and Assignment Design.” *American Political Science Review* 113(4):902–916.
- Knox, Dean, Teppei Yamamoto, Matthew A Baum and Adam J Berinsky. 2019. “Design, identification, and sensitivity analysis for patient preference trials.” *Journal of the American Statistical Association* 114(528):1532–1546.
- Mattingly, Daniel, Trevor Incerti, Changwook Ju, Colin Moreshead, Seiki Tanaka and Hikaru Yamagishi. 2024. “Chinese Propaganda Persuades a Global Audience That the “China Model” is Superior: Evidence From A 19-Country Experiment.” *American Journal of Political Science* .
- Peyton, Kyle, Gregory A Huber and Alexander Coppock. 2022. “The generalizability of online experiments conducted during the COVID-19 pandemic.” *Journal of Experimental Political Science* 9(3):379–394.