

Weiss, Chagai M.; Siegel, Alexandra; Romney, David

Working Paper

How Threats of Exclusion Mobilize Palestinian Political Participation - A Response to Bochkareva, Silagadze and Stephan -

I4R Discussion Paper Series, No. 262

Provided in Cooperation with:

The Institute for Replication (I4R)

Suggested Citation: Weiss, Chagai M.; Siegel, Alexandra; Romney, David (2025) : How Threats of Exclusion Mobilize Palestinian Political Participation - A Response to Bochkareva, Silagadze and Stephan -, I4R Discussion Paper Series, No. 262, Institute for Replication (I4R), s.l.

This Version is available at:

<https://hdl.handle.net/10419/326981>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



No. 262

DISCUSSION PAPER SERIES

How Threats of Exclusion Mobilize Palestinian Political Participation – *A Response to Bochkareva, Silagadze and Stephan* –

Chagai M. Weiss

Alexandra Siegel

David Romney

This paper responds to:

Bochkareva, Y., G. Silagadze, and M. Stephan. 2025. Replication of “How Threats of Exclusion Mobilize Palestinian Political Participation. *I4R Discussion Paper Series* No. 261. Institute for Replication.

September 2025

I4R DISCUSSION PAPER SERIES

I4R DP No. 262

How Threats of Exclusion Mobilize Palestinian Political Participation – A Response to Bochkareva, Silagadze and Stephan –

Chagai M. Weiss¹, Alexandra Siegel², David Romney³

¹University of Toronto/Canada

²University of Colorado at Boulder/USA

³Brigham Young University, Provo/USA

SEPTEMBER 2025

Any opinions in this paper are those of the author(s) and not those of the Institute for Replication (I4R). Research published in this series may include views on policy, but I4R takes no institutional policy positions.

I4R Discussion Papers are research papers of the Institute for Replication which are widely circulated to promote replications and meta-scientific work in the social sciences. Provided in cooperation with EconStor, a service of the [ZBW – Leibniz Information Centre for Economics](https://www.zbw.eu/), and [RWI – Leibniz Institute for Economic Research](https://www.rwi-essen.de/), I4R Discussion Papers are among others listed in RePEc (see IDEAS, EconPapers). Complete list of all I4R DPs - downloadable for free at the I4R website.

I4R Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

Editors

Abel Brodeur
University of Ottawa

Anna Dreber
Stockholm School of Economics

Jörg Ankel-Peters
RWI – Leibniz Institute for Economic Research

E-Mail: joerg.peters@rwi-essen.de
RWI – Leibniz Institute for Economic Research

Hohenzollernstraße 1-3
45128 Essen/Germany

www.i4replication.org

How Threats of Exclusion Mobilize Palestinian Political Participation

*A Response to Bochkareva, Silagadze and Stephan (2025)**

Chagai M. Weiss[†] Alexandra Siegel[‡] David Romney[§]

August 1, 2025

Abstract

In our article, “*How Threats of Exclusion Mobilize Palestinian Political Participation*,” we argue that threats of exclusionary policies can mobilize minority citizens through both electoral and non-electoral channels. Using Donald Trump’s 2020 “Deal of the Century” (DOC)—which explicitly threatened the citizenship status of a subset of Palestinian citizens of Israel in the Triangle area—we show that the announcement increased political discourse on Facebook, voter turnout, and registrations to a Jewish-Arab social movement. Bochkareva, Silagadze and Stephan (2025) replicate our analyses and reproduce our main findings. They raise thoughtful concerns about the parallel trends assumption, and robustness to alternative treatment definitions and outcome measures. While we welcome these contributions, we clarify that our core design choices—including focusing our analysis on three recurring elections in 2019-2020, and ten localities mentioned in the DOC—were substantively motivated and pre-registered. Building on Bochkareva, Silagadze and Stephan (2025), and drawing on our substantively justified research design, we employ additional event studies that provide further support for the observable implications of the parallel trends assumption, as well as our theoretical argument that threats of exclusion can mobilize minority political participation. We also address data sharing limitations stemming from platform terms of service, data agreements, and ethical concerns, and explain our matching procedures. We conclude by commending the authors’ decomposition of mobilization patterns and agree that variation between institutionalized and non-institutionalized responses presents a promising avenue for future research.

*We sincerely thank Meret Stephan, Givi Silagadze, and Yana Bochkareva for constructively engaging with our research.

[†]Assistant Professor, University of Toronto ✉: chagai.weiss@utoronto.ca

[‡]Associate Professor, University of Colorado at Boulder. ✉: alexandra.siegel@colorado.edu

[§]Assistant Professor, Brigham Young University. ✉: david.romney@byu.edu

1 Introduction

In our article, “*How Threats of Exclusion Mobilize Palestinian Political Participation*,” we theorize and demonstrate that threats of exclusionary policies can mobilize minority citizens through both electoral and non-electoral channels. To test our theory, we leverage Donald Trump’s announcement of the Deal of the Century (DOC): a peace plan for the Israeli–Palestinian conflict—which posed a threat to the citizenship status of a subset of Palestinian citizens of Israel residing in the Triangle area adjacent to the West Bank. First, using more than 170,000 posts from public Facebook groups and pages, we show that Trump’s announcement was a more salient political event for Triangle residents. Then, employing locality-level election data as well as records detailing the origin of citizens’ joining a Jewish-Arab social movement, we use a difference-in-difference design to demonstrate that the threat to citizenship imposed by Trump’s plan increased mobilization in the Triangle area.

Bochkareva, Silagadze and Stephan (2025)—at times referred to in this writeup as BSS—replicate our analyses as part of a graduate replication seminar at the European University Institute. We appreciate their careful work and transparency. Their replication reproduces our main findings. They also raise useful questions regarding the strength of the parallel trends assumption and explore the substantive implications of alternative definitions of treatment and outcomes.

We view many of these extensions as complementary to our original argument, and we welcome the opportunity to clarify and reflect on our research design. In particular, while the assumption of parallel trends is, as in all DiD studies, ultimately untestable, we believe our design is substantively well-motivated and reasonable. The weight a reader assigns to our causal claims will necessarily depend on how credible they find this assumption in light of the empirical context and supporting evidence provided.

Importantly, we also emphasize that our decision to focus on the ten localities explicitly mentioned in Trump’s peace plan—and to limit the analysis to the three national elections occurring between 2019 and 2020—was not arbitrary. This decision was theoretically and empirically grounded choice that was pre-registered in advance of the third election under investigation. While we report a broader range of robustness checks in the paper, our core empirical strategy adheres to this pre-registered design in terms of temporal scope and treatment definitions. We note that such pre-registration is still relatively rare in the space of observational research and reflects our commitment to transparency and interpretive clarity.¹

In addition, we briefly clarify issues related to data availability and locality matching for the

¹Access to pre-registration materials is available through the following link: <https://osf.io/wvup7/>.

first component of our paper focused on analyses of Facebook posts. Due to ethical constraints surrounding the political sensitivity of this topic and platform terms of service, we are unable to share raw Facebook post data. Instead, we provide anonymized aggregated data and reproducible code. Similarly, the social movement data were shared with us under a data-use agreement that precludes public release of raw registration records.

With regard to our matching procedures, we show that our results can be replicated using the original R environment and package versions.

Finally, we commend BSS for their thoughtful decomposition of treatment effects (Section 3.2.2), which suggests variation in mobilization patterns between explicitly and implicitly threatened localities. We agree that distinguishing between institutionalized and non-institutionalized forms of political participation is a promising avenue for future research on the political consequences of exclusionary threats.

We proceed by addressing the authors' concerns in the order they were presented in the initial memo: data availability, matching, parallel trends, definition of treatment, and measurement of mobilization.

2 Data Availability

We are committed to research transparency and have made all code and processed data needed to replicate our analyses publicly available on the American Journal of Political Science Data-verse (<https://doi.org/10.7910/DVN/EGXUBU>). However, consistent with standard practices for social media research, we cannot share the raw Facebook posts used in our analysis due to restrictions imposed by Facebook's platform policies and out of respect for user privacy. As noted in the original paper, we collected these data through CrowdTangle, a Meta-owned tool that provides access to public content from Facebook pages and groups. Under CrowdTangle's terms of service and Facebook's broader data-sharing policies, redistribution of raw post-level data is not permitted. Additionally, given the politically sensitive nature of this data, there are ethical concerns associated with releasing raw post data.

Instead, we provide reproducible code and aggregated social media data allowing others to replicate our analysis. Our Facebook analysis is descriptive in nature and was not used to generate causal claims.

For the social movement mobilization data, the organization that manages the listserv used in the study requested that we do not share raw data online. The organization permitted that we use and analyze the data for academic purposes, and share it in aggregate form online for replication

purposes, but prohibited public release of information about specific registrations due to privacy concerns. We provide aggregated and anonymized versions of the data used in our analyses, along with full code to replicate our regressions.

These limitations are common in studies using sensitive or proprietary data. We have followed best practices for research ethics, data privacy, and platform compliance while maintaining a high level of transparency through detailed documentation and reproducible workflows. While in their memo, BSS characterized our paper as providing no raw data and cleaning code, we think that a more accurate characterization of our replication package is one that provides partial access to both, due to various privacy constraints.

3 Matching Response

Our descriptive analysis of Facebook posts focuses on ten Triangle localities benchmarked against ten control localities. To select the ten control localities, we employed a matching procedure that maximizes similarity across localities in light of pre-treatment covariates.

We thank BSS for noting that, when running the replication code for our matching results using current versions of R and the relevant packages, they obtain different matched localities than those we list in Table A.16 in the Appendix. We also obtained these new results when we re-ran our replication code — which also results in new reported metrics for the Mahalanobis distance matching in Table A.15.

When re-running this code, we noticed that the PSM matching procedure still produces the same matching metrics as those reported in our published Appendix, and it was only the Mahalanobis matching that had changed. This led us to suspect that updates to the Mahalanobis procedure in the `MatchIt` package might be causing the new results. It turns out that this suspicion is correct; if you re-run the replication code in an R environment replicating that used when we produced the Appendix figures, it produces the same results reported in the published paper and Appendix.

To run the minimal amount of code necessary to produce the original results (from only the matching portion of the appendix) in an R environment replicating that used when we produced the replication materials, the following steps can be followed.

1. Run the `voting_analysis.R` script in our replication materials to produce the data necessary to carry out the matching procedure and tie locality numbers to their English names. The two necessary R objects are those named `matchdat` and `yishuv_index`. These objects

should be saved so that they can be loaded into the replicated R environment. The authors used the `saveRDS()` command to do this.

2. Produce an R environment replicating the one used to produce the replication materials. The authors did this on a Mac using an R installation manager program called `rig`. This link contains instructions for how to install and use this program: <https://github.com/r-lib/rig>. Since we were working on a Mac with Homebrew, we used the following two lines of code to install this program on the command line:

```
brew tap r-lib/rig
brew install --cask rig
```

We then used the command `rig add 4.1.2` to install a version of R matching that used when we produced the replication materials. With this version of R installed, we could use Rig's menu bar app to select it for use.

3. In this new environment, we then installed the versions of the R packages used to create the replication materials. We first installed the `remotes` package. This package allows for the easy installation of older versions of other packages using the `install_version()` command. The two main packages necessary are version 1.14.2 of `data.table` and version 4.3.2 of `MatchIt`. Because of dependencies, when we load these packages, three additional packages are loaded via a namespace: version 4.1.2 of `compiler`, version 1.4.0 of `backports`, and version 1.0.7 of `Rcpp`. Additionally, we installed version 1.8.6 of `plyr`, which gets used in the call to `plyr::mapvalues()` in this part of the replication script. This is the minimal set of packages required to replicate the matching results.
4. Lastly, we run the following code to replicate our results, after reading in the `matchdat` and `yishuv_index` objects we saved in step 1:

```
library{data.table}
library{MatchIt}

# Mahalanobis matching and PSM
mymatch1 <- matchit(formula = triangle ~ pop_2018 + First + Second,
data = matchdat, method = "nearest", distance = "mahalanobis")
mymatch2 <- matchit(formula = triangle ~ pop_2018 + First + Second,
data = matchdat, method = "nearest", distance = "logit")
```

```

# Tie matched locality numbers to English names
matchloc_tri <- matchdat$lcode[as.numeric(row.names(mymatch1$match.matrix))]
matchloc_nontri <- matchdat$lcode[as.numeric(mymatch1$match.matrix[, 1])]

# English names in a table (A.16)
match_english <- data.table(
  "Triangle" = plyr::mapvalues(matchloc_tri, yishuv_index$lcode,
    yishuv_index$locality_english),
  "Non-Triangle (Match)" = plyr::mapvalues(matchloc_nontri,
    yishuv_index$lcode,
    yishuv_index$locality_english)
)

# Matching metrics (A.15)
matchtab <- data.table(
  "Variable" = c("2018 Population", "Turnout, April 2019",
    "Turnout, September 2019"),
  "Original" = summary(mymatch1)$sum.all[, 3],
  "Mahalanobis" = summary(mymatch1)$sum.matched[, 3],
  "PSM (Logit)" = summary(mymatch2)$sum.matched[2:4, 3]
)

```

5. This produces two objects, `matchtab` that corresponds to Table A.15 in the replication Appendix and `match_english` that corresponds to Table A.16.

Since we fully replicated our original R environment, we are not sure what precisely causes the change noted by BSS — this replicated environment uses a older version of R as well as of the relevant packages. However, a potential update to the Mahalanobis procedure from the `MatchIt` package seems the most likely culprit.

4 Parallel Trends

A central point raised by BSS relates to the parallel trends assumption we introduce when analyzing voting and social movement mobilization data (Tables 1-2 in the main text). In our original paper we sought to provide evidence in support of the observable implications of our parallel trends

assumptions by reporting turnout and mobilization patterns by election cycle (Figure 4-5 in the main text).² Importantly, however, we did not report a formal test for the existence of balanced pre-trends between triangle and non-triangle localities. We agree with [Bochkareva, Silagadze and Stephan \(2025\)](#) that such a test is important because it examines the observable implications of our identifying assumption.

Employing an event study approach, BSS demonstrate that when considering all 16 triangle localities, the 2015 pre-trend estimate suggests a clear violation of parallel trends (especially when considering turnout as an outcome). BSS's decision to employ an event study approach is especially valuable in this case.³ However, following our pre-registration materials that emphasize a primary focus on the ten triangle localities mentioned in the DOC, and the substantive justification for our parallel-trends assumption that leverages the recurrence of 3 national elections within a single year (which are the elections we analyze in Table 1), we believe that a more suitable analysis should focus on: 1) the ten targeted triangle localities, and 2) consider the strength of parallel trends especially during the recurring elections in 2019-2020. Indeed, as we note in the main text:

"...for our analysis of voting records, the three recurring elections in Israel, taking place over less than 1 year, provide a unique opportunity to follow PCI voting patterns in very short time intervals. The brief intervals between elections in Israel during 2019–20, along with the differential impact of Trump's proposed plan, allow us to trace and compare within-locality changes in voting behavior and then attribute any over-time differences across Triangle and non-Triangle localities to the threat to citizenship imposed by Trump's plan. Unlike other longitudinal analyses of voting behavior where several years separate each electoral cycle, in our case, time-varying factors such as changing residential patterns, economic development, or government investment are unlikely to confound our estimates."([Weiss, A. Siegel and Romney, 2023](#), p. 1087)

Building on BSS, we report additional event study regressions that consider both our original pre-registered treatment indicator (ten triangle localities mentioned in the DOC) as well as the non-preregistered indicator that considers all sixteen triangle localities as treated.⁴ Our formal

²We also examine similar patterns in voting in section 2 of our SI.

³Though, the target estimand in the event-study is subtly different from the original estimand we pre-registered and reported in Table 1-2.

⁴We also wish to emphasize that though the estimates in Table 1 of the main text are precisely estimated at conventional levels when employing the non-preregistered treatment measure, and include a degree of uncertainty when focusing on the pre-registered treatment measure, we do not report the former as our "preferred estimate" in the original paper. Moreover, the difference between both estimates is not statistically significant.

specification employed for both voting and mobilization data is presented in equation 1, where ε_{it} is clustered at the locality level:

$$Y_{it} = \sum_{k \neq k_0} \beta_k \cdot \mathbf{1}[\text{Time}_t = k] \cdot \text{Triangle}_i + \alpha_i + \gamma_t + \varepsilon_{it} \quad (1)$$

4.1 Voting Reanalysis as Event Study

In Figure 1, we report our preferred event study specification that focuses on our pre-registered measure of treatment (ten triangle localities), and considers the three recurring elections in 2019-2020. As acknowledged by BSS, it is highly plausible that in the longer-run (going back to 2015) parallel trends do not hold. However, drawing on substantive case knowledge as well as a balance tests reported in Table A3 of our SI, we believe that within the 2019-2020 time-frame, non-triangle localities present a good counterfactual for triangle localities, and their turnout trajectories should be relatively similar.

To substantiate this argument and re-examine the effects of the DOC on turnout among triangle localities, we report event-study estimates for turnout, as well as voting behavior in Figure 1. Across all panels of Figure 1 we find that prior to the DOC (in the first election cycle of the time-period we examine), the difference between triangle and non-triangle localities in terms of turnout and voting behavior is both substantively small, and imprecisely estimated. These patterns provide empirical support for the observable implication of our key identifying assumption.

Interestingly, adopting an event study approach strengthens our key results. While directionally similar, the effects of the DOC on turnout are modestly larger ($\beta = 0.025$ vs. $\beta = 0.024$) and estimated with a higher degree of precision ($p = 0.015$ vs. $p = 0.129$). In line with the results reported in our SI, we also find positive and precisely estimated treatment effects on Blue and White vote share, but no other treatment effects on support for other political parties.

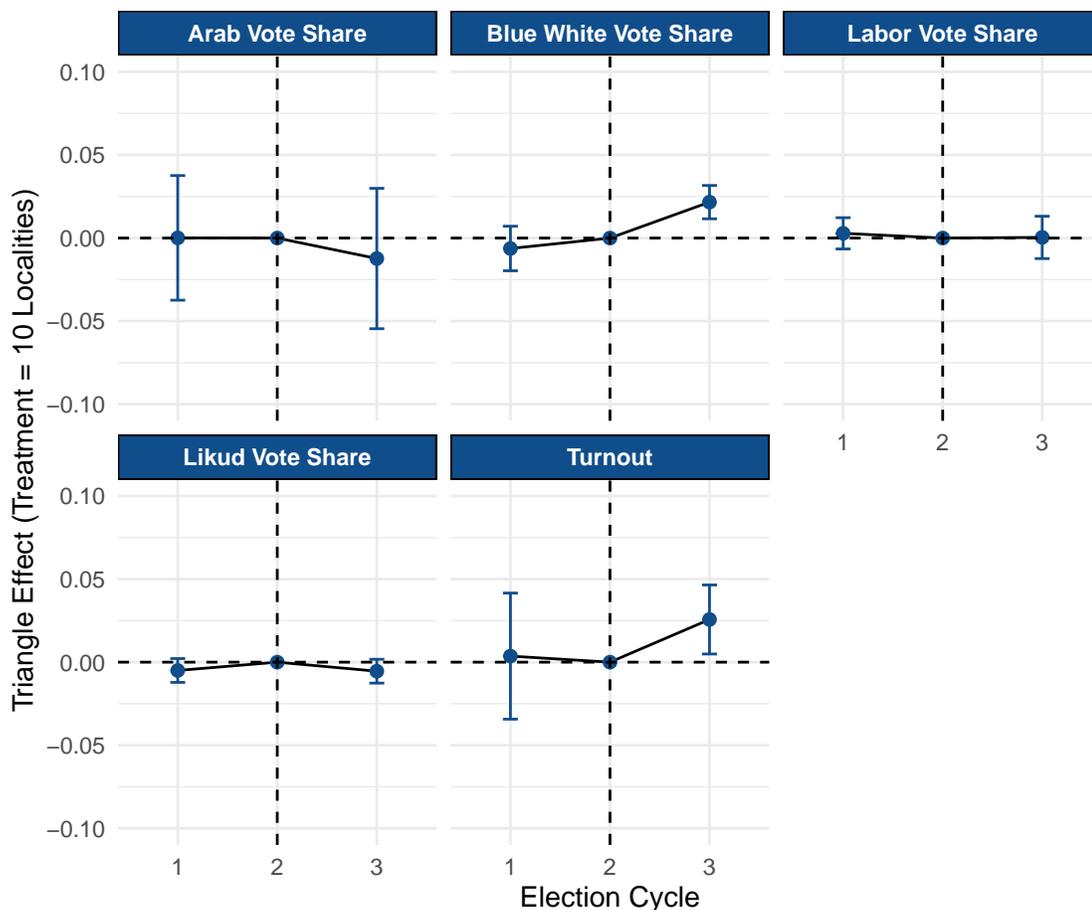


Figure 1: DOC Effect on Voting. Treatment = 10 pre-registered triangle localities.

We complement the results reported in Figure 2, with additional analyses focusing on all sixteen triangle localities. As noted above, this measurement approach of treatment was not pre-registered, but still reported given its substantive importance. Notably, as reported in Figure 2, and in line with Figure 2 in [Bochkareva, Silagadze and Stephan \(2025\)](#), the empirical evidence supporting the observable implications of the parallel trends assumption are less robust in this case. When considering all voting outcomes, the pre-treatment estimate of Triangle vs. non-Triangle localities are relatively small and imprecisely estimated. However, when focusing on turnout the pre-treatment estimate is substantially large and relatively precise ($\beta = -0.029$ and $p = 0.094$). More generally, in line with the evidence from our main text (Table 1 column 3) as well as the additional estimates reported by BSS when employing this alternative treatment measurement approach, the effects of the DOC on turnout are substantially larger. However, given the potential violation of parallel trends in this case, we advise a degree of caution when ascribing a strong causal interpretation to these supplemental estimates.

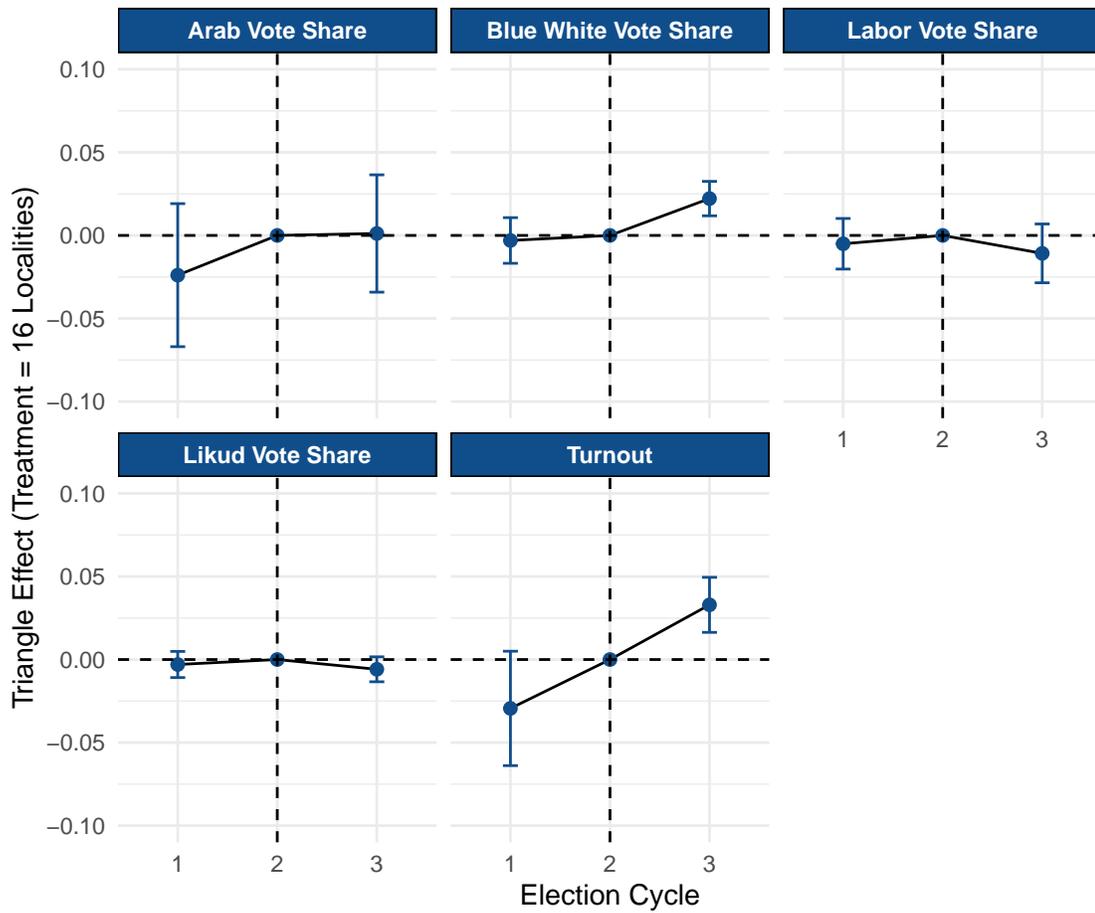


Figure 2: DOC Effect on Voting. Treatment = 16 triangle localities.

4.2 Social Movement Reanalysis as Event Study

Complementing our reanalysis of voting data, we turn to use a similar event study approach to evaluate the observable implications of our parallel trends assumption and consider the effects of the DOC on registrations to the social movement under investigation in (Weiss, A. Siegel and Romney, 2023). Our original focus in Table 2 of (Weiss, A. Siegel and Romney, 2023) was on the aggregate difference-in-difference estimate, considering the overall effect of the DOC on triangle localities between January 2017 and August 2020. The event study approach allows us to consider dynamic treatment effects over time, and examine whether there is empirical support for the observable implications of the parallel trends assumption in the pre-treatment period, and especially during in the period of three recurring elections unfolding in 2019-2020. Since we do not have specific priors regarding effects at the day-level, and since the data generating process of registrations is often bunched, likely underestimating day-by-day variation,⁵ we follow BSS and estimate event-studies where data is aggregated by month.

In Figure 3 we report our event study estimates using both measurement approaches for our treatment. Both panels reveal similar insights. At the pre-treatment period, and especially since the first 2019 election period, an overwhelming majority of estimates for triangle-non-triangle differences by month are imprecisely estimated.⁶ However, following the declaration of the DOC, we see a gradual rise in registration, that reaches its peak in March 2020. We agree with BSS that this suggests a gradual, rather than immediate response to the DOC, likely driven by both rising grievances among Triangle residents and strategic mobilization efforts by the NGO. More generally, however, we believe that these analyses provide both empirical support for the observable implications of our parallel trends assumption, as well as our primary hypothesis that threats of exclusion can mobilize minorities.

⁵A meaningful portion of registrations occur when activists register peers in the field, and log their information into a database when visiting the movements' office.

⁶Indeed, when considering the pre-registered (alternative) treatment definition 2/35 (1/35) estimates are precisely estimated ($p < 0.05$). Importantly, there is no empirical evidence for violation of parallel trends in the 2019-2020 period.

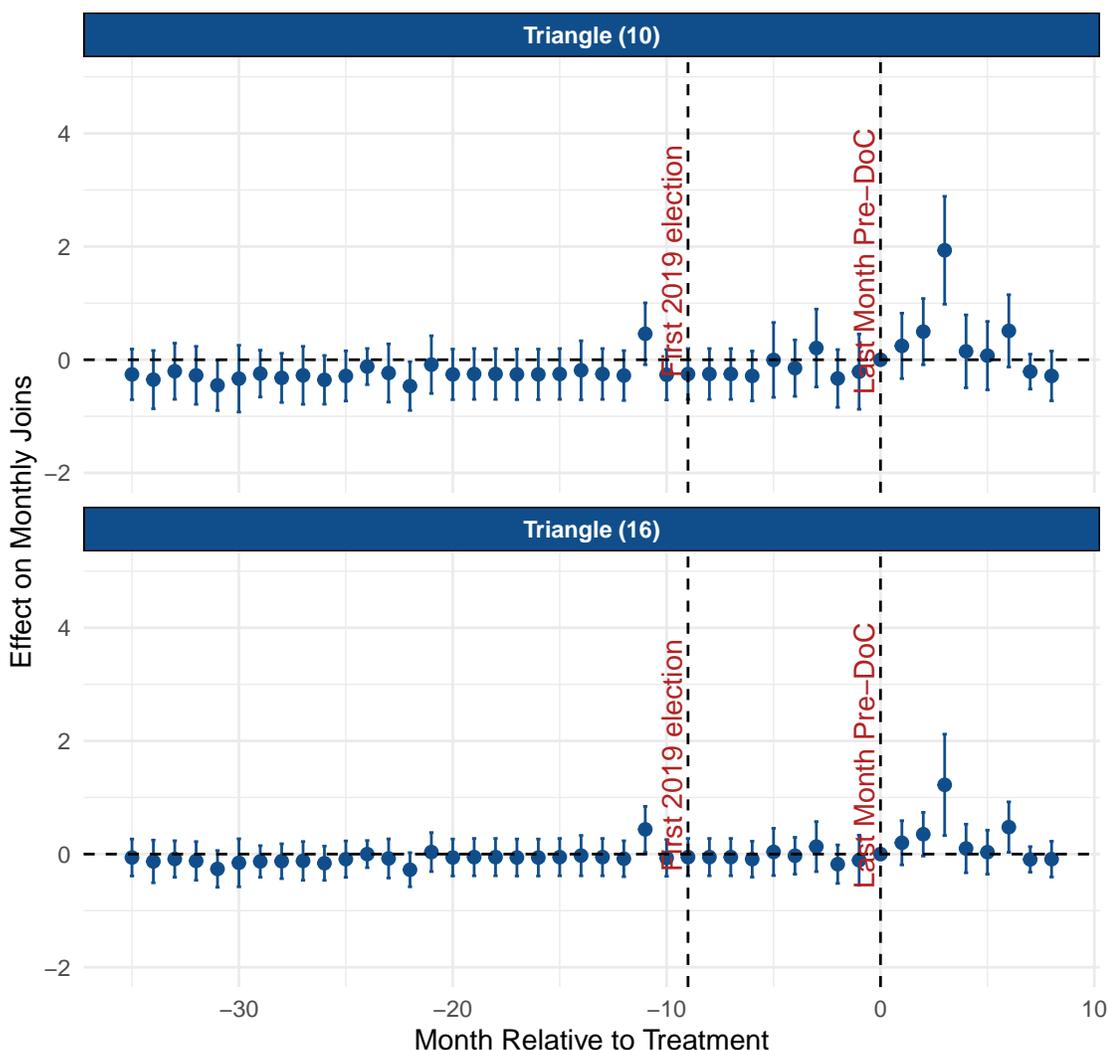


Figure 3: DOC Effect on Mobilization.

5 Definition of Treatment

A core substantive contribution of [Bochkareva, Silagadze and Stephan \(2025\)](#) is in considering the effects of the DOC among different localities which were directly (ten Triangle localities) and indirectly targeted (six triangle localities) by Trump’s plan. This approach is motivated by our estimates that were not pre-registered, and consider all sixteen triangle localities as treated. Ultimately, their analyses that decomposes the overall effect of the (sixteen) Triangle localities into those directly and indirectly targeted yields interesting insights about how explicit and implicit threats of exclusion may affect different outcomes. We believe that this extension is valuable, and

warrants further research. However, we ultimately suggest that the main conclusions from our work should be based on a comparison of the ten triangle localities with all other non-Jewish localities for three primary reasons.

First, the definition of treatment as the ten localities mentioned in the DOC, aligns with our pre-registered materials. Second, this definition aligns with our initial interest in understanding how *explicit* threats of exclusion may mobilize minority political participation. Finally, it appears (from Figures 1-3) that the empirical support for parallel trends is somewhat stronger when considering our pre-registered definition of treatment. While the consideration of all Triangle localities as treated, and the decomposition exercise by BSS is substantively illuminating, we advocate for a primary focus on the ten triangle localities that are the center of our empirical investigation.

6 Measurement of Mobilization

A final component of the BSS report focuses on our mobilization analyses, and considers the robustness of our findings to alternative outcome measurement strategies. In our initial analyses we consider mobilization at the locality level as follows:

$$\text{Mobilization}_\ell = \begin{cases} 1 & \text{if Registrations}_\ell > X \\ 0 & \text{otherwise} \end{cases} \quad \text{where } X = 1 \quad (2)$$

BSS show that these results hold when $X = 2$ but not $X = 3$. While valuable for demonstrating the robustness of our analyses, we think that this measurement approach has two limitations. First, given that most days do not include any registrations, the move from 0 to 1 which we consider in our original article is of substantive relevance. Second, while other thresholds (e.g. $X = n$) may be interesting, we do not see a theoretical reason to prefer a threshold at 2, 3, or more registrations, and therefore prefer our initial substantively motivated approach for outcome measurement. Regardless, as acknowledged by BSS, in Table A14 of our original SI we demonstrate robustness of our primary result when using a count of mobilization and a negative binomial regression suitable for count data. Taken together, our initial results, the robustness of our analyses to some aggregates (e.g., $X = 2$), the positive (albeit imprecisely estimated) estimates when using $X = 3$, and the supplementary negative binomial models, strengthen our intuition that our results are not a function of a highly specific (and biased) measurement decision.

References

- Bochkareva, Yana, Givi Silagadze and Meret Stephan. 2025. “Replication of ‘How Threats of Exclusion Mobilize Palestinian Political Participation’”. Replication study.
- Weiss, Chagai M, Alexandra A. Siegel and David Romney. 2023. “How threats of exclusion mobilize palestinian political participation.” *American Journal of Political Science* 67(4):1080–1095.