Appendix 1. Data & Descriptive Statistics

Note on ARVIG Hate Crime Data: Arsons, Assaults, and Demonstrations categories include a small number (N=16 for the 2014-2017 period; none for just 2016) of mixed categories (e.g. arson and attack on refugee housing). Attacks on Refugee Housing also includes events categorized as suspicion (N=211 for the 2014-2017 period; 174 for just 2016) and other (N=25 for the 2014-2017 period; none for just 2016).

Table A 1. Difference of Means Tests For Treated/Untreated Municipalities

Municipal Characteristics	Mean (C)	N (C)	Mean (T)	N (T)	Dif.	T-Stat	p
Log Population (2015)	7.16	9213	9.15	1823	-1.99	-68.44	0.00
U25 Share (2011)	0.25	9213	0.24	1823	0.01	5.99	0.00
O65 Share (2011)	0.20	9213	0.21	1823	-0.01	-14.76	0.00
Share Male (2015)	0.50	9213	0.50	1823	0.01	21.18	0.00
Foreign Born Share (2011)	0.03	9213	0.05	1823	-0.02	-16.97	0.00
Population Density (2015)	137.50	9213	343.43	1823	-205.93	-22.81	0.00
Size in Km^2	24.86	9213	61.70	1823	-36.84	-27.85	0.00
Unemployment as % of Population (2015)	2.32	9213	2.98	1823	-0.65	-14.91	0.00
Voter Turnout (2013)	68.75	9213	67.26	1823	1.49	6.90	0.00
% AfD (2013)	0.05	9213	0.05	1823	-0.00	-6.48	0.00
% CDU/CSU (2013)	0.47	9213	0.45	1823	0.02	10.39	0.00
% FDP (2013)	0.04	9213	0.04	1823	0.00	2.15	0.03
% Green Party (2013)	0.06	9213	0.06	1823	-0.00	-3.76	0.00
% Linke (2013)	0.08	9213	0.09	1823	-0.01	-5.71	0.00
% SPD (2013)	0.22	9213	0.23	1823	-0.01	-5.26	0.00

Data Source: German Federal Statistical Office. Year of data determined by data availability, with 2011 the year of the census and 2013 the year of most most proximate national elections prior to 2016.

Table A 2. Descriptive: Outcomes

Variable	Index Component	N	Mean	SD	Min	Max	α
Attitude Towards Refugees		23,885	0.412	0.224	0	1	0.909
	Economic Impact	24,245	5.655	2.669	1	11	
	Cultural Impact	24,244	5.581	2.756	1	11	
	Impact on Germany	24,227	5.097	2.476	1	11	
	Short-Term Risks	24,244	3.859	2.270	1	11	
	Long-Term Risks	24,190	5.394	2.874	1	11	
Plans to Help Refugees		23,797	0.185	0.266	0	1	0.522
	Donate	24,057	0.346	0.476	0	1	
	Direct Support	24,014	0.129	0.335	0	1	
	Political Support	23,995	0.089	0.284	0	1	

Data source: SOEP v37. Sample: all 2016 respondents.

Appendix 1.1 Additional Variables Derived from the SOEP

In addition to the main outcomes described and summarized in the main text, we constructed a number of additional individual-level variables from the SOEP. Descriptive statistics for these additional variables are shown in Tables A3 and A4.

First, respondents were asked whether they provided help to refugees in the past, using the same three survey items that they were used to ask about future support. This is a useful compliment to our measure of planned/future support. We use this measure in a subsequent continuity test, to explore whether past refugee-facing behavior shifts at the cutoff.

Similarly, we utilize the 2015 wave of the SOEP to measure respondent political characteristics. These political characteristics include general attitudes towards immigration as well as partisanship. We again use these pre-treatment measures to test the continuity assumption in a more credible way, as well as to explore effect heterogeneity across baseline political dispositions.

Third, we derive a number of respondent-level characteristics, including sex, partner status migration status, employment status, and health. These characteristics are either time invariant or unlikely to be affected by exposure to an anti-refugee hate crime, rendering them appropriate for inclusion in a test of the continuity assumption.

Table A 3. Descriptive: Additional SOEP-derived, respondent-level characteristics Measures

Variable	Index Component	N	Mean	SD	Min	Max	α
Past Help for Refugees		24,273	0.139	0.217	0	1	0.367
	Donate	24,399	0.300	0.458	0	1	
	Direct Support	24,345	0.070	0.255	0	1	
	Political Support	24,324	0.051	0.220	0	1	
Partisanship (2015)							
SPD Partisan	22,832	0.117	0.322	0	1		
CDU/CSU Partisan	22,832	0.161	0.368	0	1		
FDP Partisan	22,832	0.011	0.103	0	1		
Green Partisan	22,832	0.058	0.233	0	1		
Left Partisan	22,832	0.032	0.175	0	1		
AfD Partisan	22,832	0.011	0.102	0	1		
Radical Right Partisan	22,832	0.004	0.061	0	1		
Other Partisan	22,832	0.006	0.077	0	1		
Anxiety about Crime and Order		24,262	0.661	0.249	0	1	0.573
	Crime	24,419	0.672	0.343	0	1	
	Social Cohesion	24,340	0.631	0.319	0	1	
	Hostility towards Foreigners	28,815	0.603	0.378	0	1	
Other Characteristics							
Female	24,651	0.542	0.498	0	1		
Birth Year	24,651	1967.432	17.644	1914	1998		
Direct Migration Background	24,651	0.200	0.400	0	1		
Direct or Indirect Migration Background	24,651	0.265	0.441	0	1		
Cohabiting with a Partner	24,649	0.687	0.464	0	1		
Employed Full Time	24,651	0.375	0.484	0	1		
Self-Reported Health	24,651	0.599	0.245	0	1.5		

Data source: SOEP v37. Sample: all 2016 respondents.

Table A 4. Descriptive: Lagged Worries about Immigration

Immigration Attitude (2015)	N
Very Concerned	6852
Somewhat Concerned	9219
Not Concerned At All	5720
Total	21791

Data source: SOEP v37. Sample: all 2016 respondents who previously responded to this question in their 2015 interview.

Appendix 2. Validating the Research Design

Regression discontinuity designs require several identifying assumptions. The first is that a known and defined threshold determines assignment to treatment status. In our case, this assumption is trivial: all residents of a municipality where a hate crime occurs are by definition treated.[10] Individuals cannot select into or out of the geographic treatment, so assignment should be perfect.

A second assumption is that individuals cannot manipulate treatment status. Our main concern is that hate crimes may alter the propensity of individuals to participate in the survey outright, due to differences in interview uptake on the part of respondents or differences in contact rates on the part of enumerators. In either case, we would expect to see a discontinuity in the daily number of observations at the cutoff. Selective behavior at the cutoff is unlikely to be a problem. First, SOEP interviews are typically scheduled well in advance of actual meetings, so enumerators and respondents are unlikely to be able to respond to hate crimes as they schedule interviews (Graeber and Schikora 2021). Second, respondents and enumerators cannot anticipate the occurrence of a hate crime in their municipality, so strategic manipulation *prior* to the cutoff is impossible. One possibility is that certain treated individuals might be more or less likely to participate in the survey – or to answer refugee-facing questions – *after* a hate crime, which could lead to imbalance across treatment status. We formally test for manipulation using a non-parametric density estimator (Cattaneo, Jansson, and Ma 2018). Figure A1 shows the empirical distribution of cases across the running variable. The distribution appears symmetrical around the cutoff for all three event types. The test does find evidence of bunching for attacks on refugee housing. This bunching for the attacks on refugee housing treatment is surprising, especially given the absence of bunching for the other two types of exposure. We believe this bunching is likely a product of chance rather than an increase in survey refusal in the aftermath of attacks on refugee housing, but we acknowledge that there is a possibility of bias.

A third assumption is that units close to the cutoff have similar potential outcomes. An empirical implication of this assumption is that predetermined characteristics do not co-vary with treatment status. Figure A2 presents the results of placebo RD models that test for effects of hate crime exposure on a set of predetermined individual and municipal characteristics known to explain

10. There could be variation in the extent to which people know about the hate crime that occurs in their municipality, but this is not directly observable in our data. Given the null result in our main analysis, we conduct an additional test of whether treatment status affects other attitudes. We find that it does increase anxiety about crime, suggesting that respondents are aware of hate crimes in their municipality.

support for refugees. For the municipal-level variables, none of the tests reveal statistically significant discontinuities across the treatment threshold (14 municipal variables). The individual-level tests reveal a small number of imbalances for each of the three event types. Due to these imbalances (and the bunching described above) we replicate our main analysis with relevant control variables. Figure A3 replicates this test with county-level characteristics. Again, none of the 23 county-level characteristics shift at the threshold.

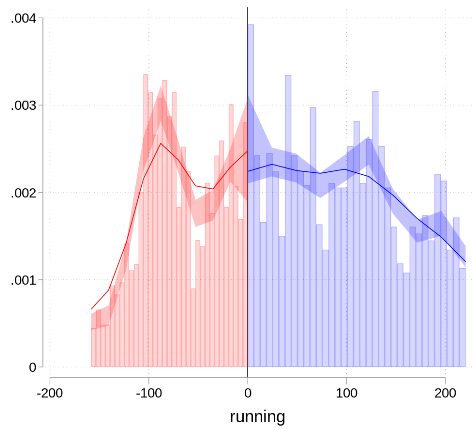

(a) Arson

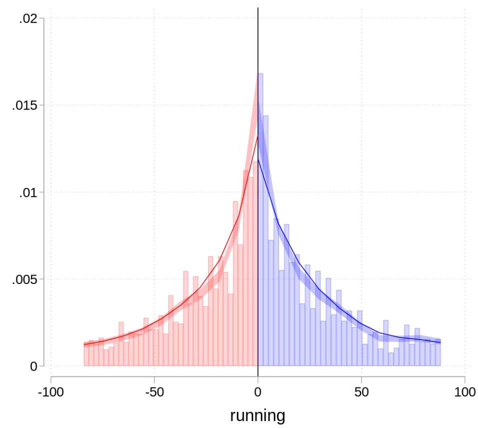

(b) Assaults

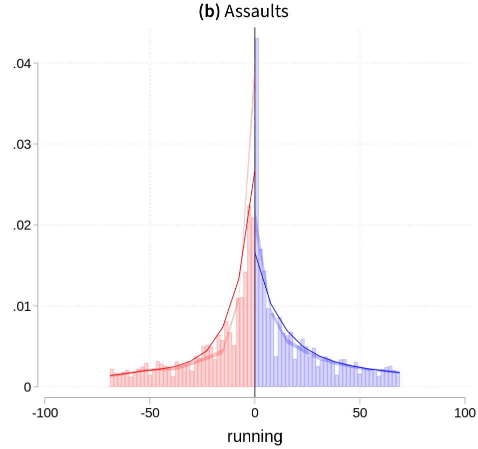

(c) Attacks on Refugee Housing

Figure A 1. Density Tests for the Running Variables

Test of manipulation at the cutoff, based on non-parametric density estimator developed by Cattaneo, Jansson, and Ma (2018). Plots depict the distribution of the running variable across the cutoff for exposure to anti-refugee hate crimes at the municipal level. Bandwidths around cutoff chosen empirically. Data Source: ARVIG (Benček and Strasheim 2016) & SOEP v37.

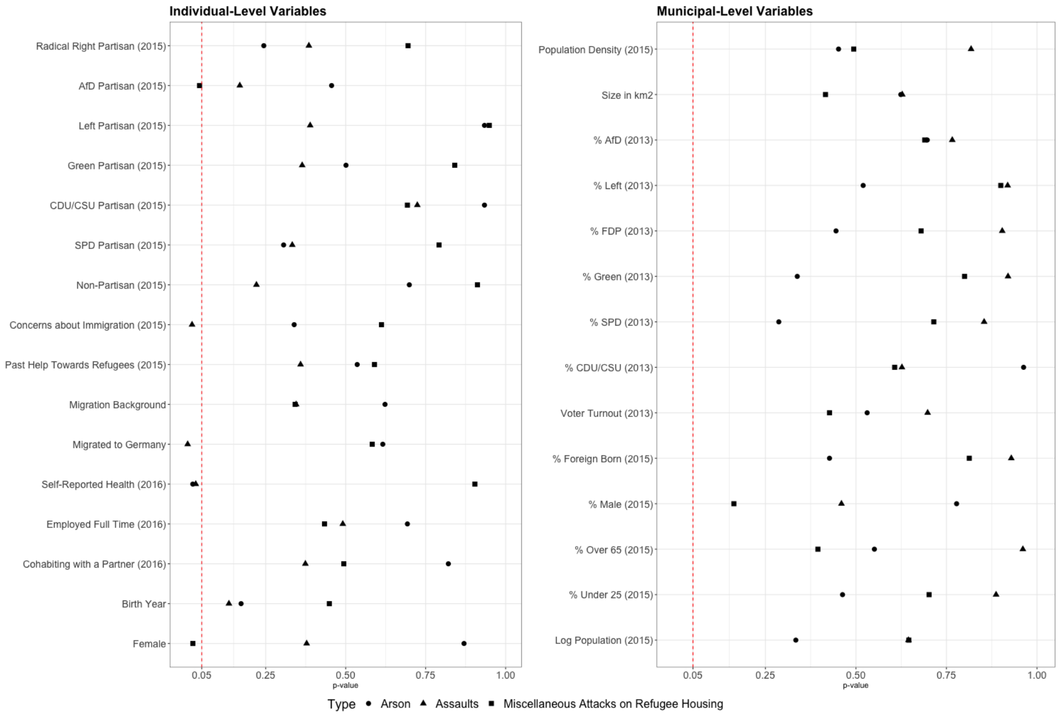

Figure A 2. Test of Continuity Assumption

Data Source: SOEP v37 for individual-level variables; German Federal Statistics Oficce for municipal-level variables; ARVIG for hate crime data. Each coefficient represents a separate test of the continuity assumption for a set of individual and municipal characteristics across the cutoff, i.e. of municipal-level exposure to a hate crime. Shape of the coefficient estimate reflects hate crime incident specification, i.e. all events, arson, assault, or attack against refugee housing. All variables are either time invariant or pre-determined, i.e. measured in the year before treatment. All models use a non-parametric RD regression using a triangular kernel and the same bandwidth for both sides of the cutoff selected using the CER-optimal bandwidth selector. The CER-optimal BW selector is preferred over the MSE-optimal BW selector here because we are are only interested in testing the null hypothesis of no effect (Cattaneo, Idrobo, and Titiunik 2019, p.94). Coefficients and standard errors reflect output from "Robust" method from *rdrobust*. As in the main analysis below, standard errors are clustered on the federal state level and respondents interviewed on the day of hate crime occurrence and respondents for whom treatment-designating hate crime occurs within 30 days of an earlier hate crime are omitted.

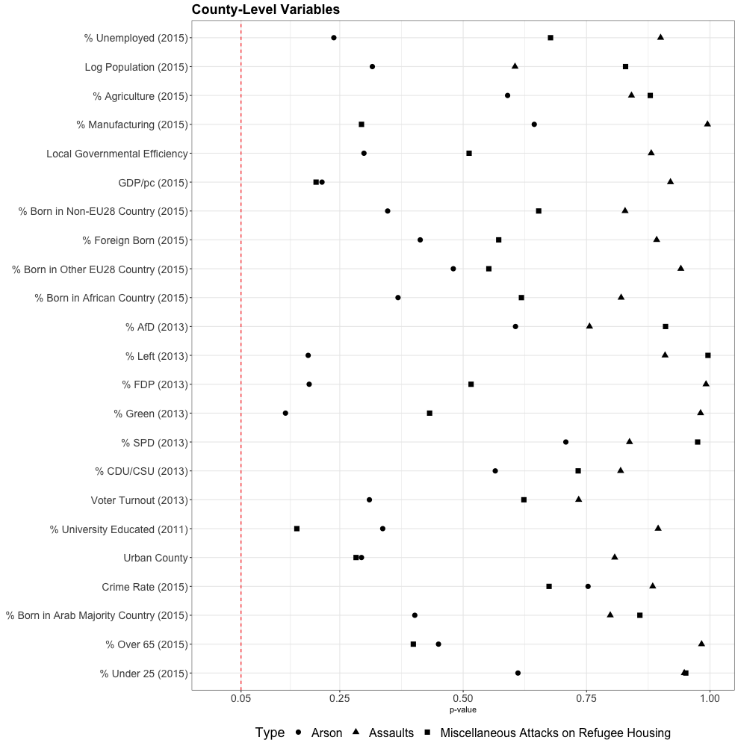

Figure A 3. Test of Continuity Assumption: County-Level Characteristics

Data Source: SOEP v37 for individual-level variables; German Federal Statistical Office, German Federal Criminal Police Office, and Ziller and Goodman (2019) for county-level variables; ARVIG for hate crime data. Each coefficient represents a separate test of the continuity assumption for a set of county characteristics across the cutoff, i.e. of municipal-level exposure to a hate crime. Shape of the coefficient estimate reflects hate crime incident specification, i.e. all events, arson, assault, or attack against refugee housing. All variables are either time invariant or pre-determined, i.e. measured in the year before treatment. All models use a non-parametric RD regression using a triangular kernel and the same bandwidth for both sides of the cutoff selected using the CER-optimal bandwidth selector. The CER-optimal BW selector is preferred over the MSE-optimal BW selector here because we are are only interested in testing the null hypothesis of no effect (Cattaneo, Idrobo, and Titiunik 2019, p.94). Coefficients and standard errors reflect output from "Robust" method from *rdrobust*. As in the main analysis below, standard errors are clustered on the federal state level and respondents interviewed on the day of hate crime occurrence and respondents for whom treatment-designating hate crime occurs within 30 days of an earlier hate crime are omitted.

Appendix 3. Analysis

Appendix 3.1 Estimation Strategy

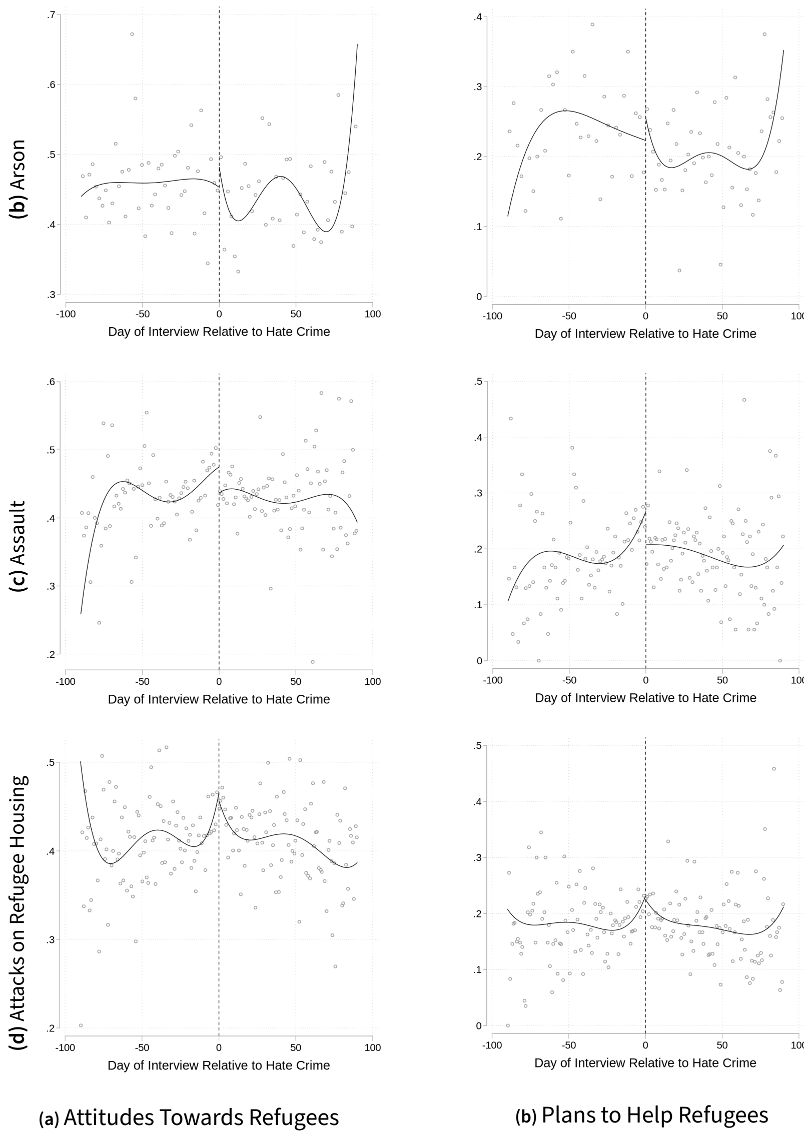

Figure A 4. Relationship Between Anti-Refugee Hate Crime Exposure and Refugee-Facing Attitudes

Flexible RD plots for the relationship between timing of anti-refugee hate crime exposure and individual attitudes towards refugees (Column 1) and plans to help refugees (Column 2). Rows represent different hate crime specifications: Arson (Row 1), Assaults (Row 2), and Attacks on Refugee Housing (Row 3). Per (Cattaneo, Idrobo, and Titiunik 2019), these plots flexibly depict the relationship between our outcomes and the running variables with a fourth-order polynomial, evenly-spaced, mimicking-variance bins. Bandwidths are ninety days on either side of the cutoff.

Appendix 3.2 Main Results

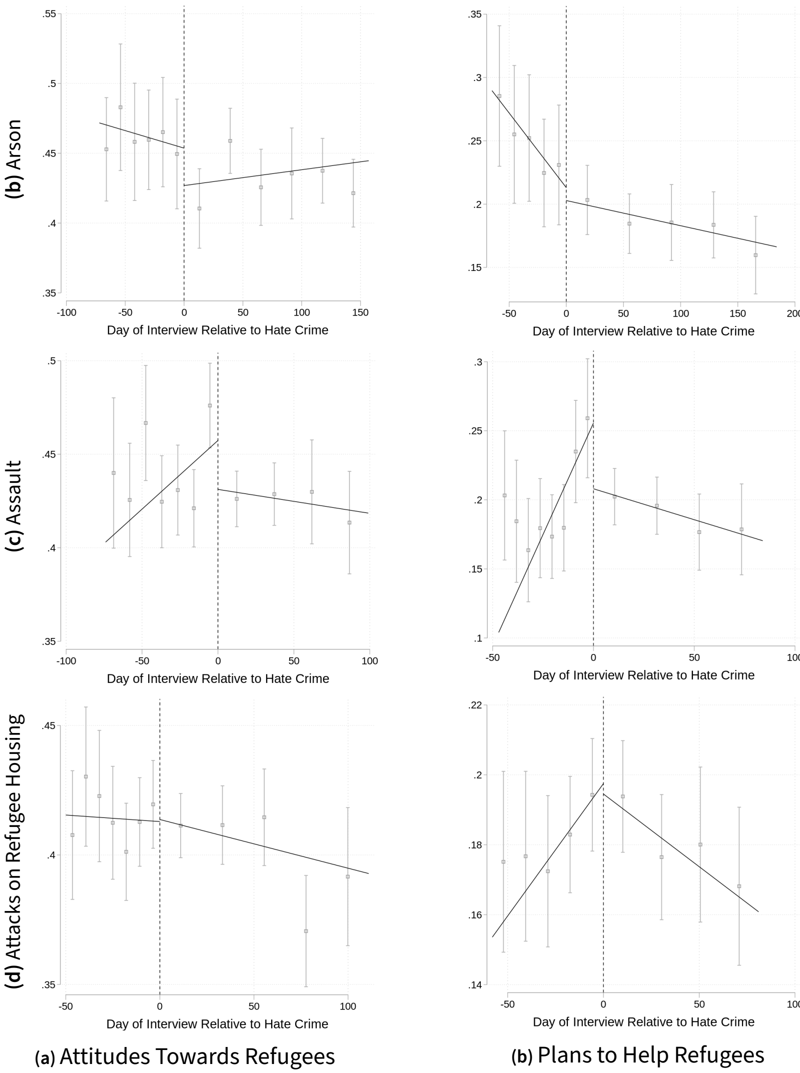

Figure A 5. Regression Discontinuity Plots

RD plots for the effect of anti-refugee hate crimes on SOEP's respondents attitudes towards refugees (Column 1) and plans to help refugees (Column 2). Rows represent different hate crime specifications: Arson (Row 1), Assaults (Row 2), and Attacks on Refugee Housing (Row 3). All models use MSE-optimal bandwidths that are permitted to differ in width across the cutoff, local linear regression, triangular kernels, and errors clustered at the federal state level. Respondents interviewed on the day of treatment are excluded, as are respondents for whom the treatment-designating hate crime occurs within 30 days of an earlier hate crime. Plots depict IMSE-optimal evenly spaced bins with standard errors clustered on the federal state level.

Appendix 3.3 Robustness Tests: LATE of Exposure to Hate Crime

In this section, we describe a number of robustness tests to probe the sensitivity of our main results, i.e. of exposure to hate crimes on attitudes towards refugees and plans to help refugees in the future. We also present the same set of tests for the analysis exploring the LATE of hate crime exposure on anxieties about crime. Discussion of results focus on the two primary findings highlighted in the main text, namely 1) the negative and significant effect of exposure to assaults on plans to help refugees, and 2) the positive and significant effect of exposure to assaults on anxieties about crime and order. More details are provided below, but the takeaway from these tests is that the first finding (LATE of assaults on plans to help refugees) is likely spurious on individual-level characteristics that are imbalanced across the threshold, as shown in Table A6. Meanwhile, the second finding (LATE of assaults on anxieties about crime) appears broadly robust.

First, we replicate the main analysis using an alternative measure of hate crime event exposure. Specifically, we bundle all three types of hate crime events together – i.e. arson, assaults, and attacks on refugee housing – so as to consider the LATE of exposure to anti-refugee hate crimes, writ large. Results are presented in Table A5. This analysis does not find a significant effect of municipal exposure to hate crimes on any outcome.

Next, we replicate our analyses with control variables for individual-level characteristics that are imbalanced at the threshold (see Table A2 above for reference). For exposure to arson events, this involves controlling for self-reported health. For assault events, this means controlling for self-reported health, lagged immigration attitudes, and migration background. For attacks on refugee housing, this requires adding a control for biological sex and lagged AfD partisanship. Results are presented in Figure A6. For exposure to arsons and attacks on refugee housing, null results do not change for any outcome when adding these controls. For exposure to assaults, including these controls attenuates the previously negative and significant effect on plans to help refugees towards zero, suggesting these imbalances may explain the discontinuity at the threshold, rather than an effect of exposure per se. However, exposure to assaults continues to exercise a positive and significant (at $p < 0.1$) effect on anxiety about crime even when adjusting for these imbalanced individual-level covariates.

Finally, we conduct a series of standard tests considered to be best practices in regression discontinuity designs (Valentim, Ruipérez Núñez, and Dinas 2021). First, we explore whether deviations from the optimal bandwidth choice alter our results. We replicate our analyses for each outcome and event type combination with manually-selected bandwidths ranging from 0.25x to 2x the optimal bandwidths on either side of the cutoff. As the bandwidth of the regression discontinuity increases, effect magnitudes on plans to help shrink slightly, along with standard errors due to increasing sample size. This is inline with the bias-variance trade-off that governs selecting bandwidth size, in which bias (variance) may increase (decrease) with bandwidth size. Results, presented in Figure A6, suggest results are robust to different bandwidth specifications. As in the main analysis, most event type and outcome combinations do not present consistent evidence of a relationship between hate crime events and refugee-facing attitudes and behavior. Two exceptions relate to the occurrence of assaults, which 1) reduce plans to help refugees and 2) increase anxieties about crime at most multiples of the optimal bandwidth (at 0.1 or below between 0.5 and 1.75 the optimal bandwidth).

Next, results could be sensitive to the fact that local regressions are specified linearly. We replicate our analysis with local polynomial regression of order two and three. Results, presented in Figure A7 are substantively comparable to the main models that use local linear regression. Concretely, these models replicate the finding that exposure to assaults reduce plans to help refugees when using second-order (at $p < 0.10$) and third-order (at $p < 0.05$) polynomials for the local regression. Moreover, exposure to assaults increase anxieties about crimes (at $p < 0.05$ for both second- and third-order polynomials).

In another test, we utilize alternative standard error specification, including clustering on the running variable per Lee and Card (2008), leaving errors unclustered, clustering on the municipality itself, and clustering on the county (kreis) in which a municipality is nested. Results are generally comparable, though estimates are significantly less precise when clustering standard errors at lower levels of regional aggregation (municipal and kreis) and are no longer singificant at conventional levels in some cases. Results are presented in Figure A8 and provide consistent evidence that exposure to assaults is associated with higher levels of anxiety about crime and order (at $p < 0.05$ or $p < 0.10$). Increased imprecision renders effects of assault exposure on plans to help refugees less clear.

We also consider sensitivity of our results to multiple treatments by replicating our models under different exclusion criteria relating to earlier exposure to hate crimes. These criteria range from not excluding any cases to excluding cases in which the treatment-designating hate crime falls within sixty days of an earlier hate crime. Results of this test are presented in Figure A9. There is some variation in how exclusion criteria shape effect magnitudes and levels of certainty. As expected, wider exclusion criteria – and therefore smaller samples – lead to less precise estimates. For the effects of exposure to assaults on plans to help refugees, estimates are significant at conventional levels (at $p < 0.10$ or below) when not excluding any case, as well as when excluding cases treated within a 30, 40, or 50 day window. For effects of exposure to assaults on anxieties about crime, effects are only significant at conventional levels when cases previously treated within one month or less are excluded.

We address the issue of multiple treatments a second way, namely by splitting the sample by previous treatment status. We categorize all individual respondents according to whether the focal hate crime to which they have been matched is the first anti-refugee hate crime in their jurisdiction recorded in ARVIG. We subsequently conduct separate analyses across previously treated and untreated units. The results are not suggestive of compounding effects (i.e. effects becoming stronger with additional treatments) or of tapering down of effects (i.e. of effects becoming weaker with each additional treatment). As in the main analysis, most event type are not associated with shifts in attitudes towards refugees, regardless of whether a municipality is experiencing its first anti-refugee hate crime. For the effect of assaults on plans to help refugees, only subsequent hate crimes are associated with a reduction in planned solidarity (at $p < 0.05$), but the effect of a first treatment is of a similar magnitude but simply estimated less precisely. Similar patterns hold for the LATE of exposure to assaults on anxieties about crime. Results are presented in Figure A10.

Next, in Figure A11 we replicate our main analysis with added controls for the day-of-week in which an interview occurred. This test allows us to account for the fact that anti-refugee hate crimes likely do not occur with similar regularity on all week days, which could confound our findings if refugee-facing attitudes also fluctuate over time. Again, results are comparable, with the effect of an assault on plans to help refugees (anxieties about crimes) significant at the 0.1 (0.05) level.

Our results could be driven by respondents nested within certain regions. To consider this possibility we conduct a Leave-One-Out analysis in which we sequentially drop respondents nested within specific federal states. Results, presented in Figure A12, suggest our findings are stable even when excluding specific federal states.

Next, we retain only respondent residing in municipalities containing ten or more (non-refugee) SOEP respondents in 2016. Results, presented in Figure A13, are similar to our main findings:

exposure to an anti-refugee assault is associated with a reduction in the planned help for refugees and an increase in anxiety about crime (at $p < 0.05$).

In a final test, we explore results on the individual response items that make up our outcome indices, for 1) the five items underpinning the attitudinal index, 2) the three items that constitute the index on plans to help refugees, as well as 3) the three items that make up the anxiety about crime outcome. Results for all three analyses are presented in Figure A14, Figure A15, and Figure A16. These findings are discussed in the main text, in Section ??.

Table A 5. LATE of Municipal Hate Crime Exposure on Refugee Attitudes, Plans to Help, and Anxiety about Crime: All Hate Crime Events

	Attitudes Toward Refugees	Plans to Help Refugees	Anxiety about Crime
All Events			
Conventional	-0.002 (0.014)	-0.010 (0.018)	0.016 (0.019)
Bias-Corrected	-0.010 (0.014)	-0.012 (0.018)	0.018 (0.019)
Robust Bias-Corrected	-0.010 (0.015)	-0.012 (0.02)	0.018 (0.019)
N	10767	10787	10964
Effective N	7762	6643	7843
BW L	100.50	54.57	72.19
BW R	93.06	78.63	126.70

$^{**}p < 0.05$, $^{*}p < 0.1$

Table 1 displays the effect of anti-refugee hate crimes on SOEP respondents' attitudes towards refugees (column 2) and plans to help refugees (column 3). The Table provides point estimates and standard errors clustered on the federal state level for both outcomes for all events in ARVIG (except demonstrations). All models use MSE-optimal bandwidths that are permitted to differ in width across the cutoff, local linear regression, triangular kernels, and errors clustered at the federal state level. Respondents interviewed on the day of treatment are excluded, as are respondents for whom the treatment-designating hate crime occurs within 30 days of an earlier hate crime.

Table A 6. LATE of Municipal Hate Crime Exposure: With Controls for Individual-Level Imbalance at the Threshold

	Attitudes Toward Refugees	Plans to Help Refugees	Anxiety about Crime
Arson			
Conventional	-0.024 (0.041)	-0.010 (0.048)	0.026 (0.028)
Bias-Corrected	-0.029 (0.041)	-0.004 (0.048)	0.032 (0.028)
Robust Bias-Corrected	-0.029 (0.051)	-0.004 (0.061)	0.032 (0.032)
N	5389	5387	5461
Effective N	2689	2702	2756
BW L	77.16	65.93	91.15
BW R	158.54	184.98	135.92
Assault			
Conventional	-0.006 (0.022)	0.001 (0.024)	0.040* (0.020)
Bias-Corrected	-0.004 (0.022)	0.000 (0.024)	0.040* (0.020)
Robust Bias-Corrected	-0.004 (0.024)	0.000 (0.027)	0.040* (0.023)
N	6053	6039	6152
Effective N	3911	3780	4114
BW L	99.65	98.32	101.65
BW R	93.62	78.90	113.75
Attacks on Refugee Housing			
Conventional	-0.005 (0.014)	-0.006 (0.017)	0.031 (0.022)
Bias-Corrected	-0.004 (0.014)	-0.006 (0.017)	0.035 (0.022)
Robust Bias-Corrected	-0.004 (0.015)	-0.006 (0.019)	0.035 (0.024)
N	9830	9828	10039
Effective N	6161	6279	6696
BW L	50.29	67.97	71.34
BW R	109.95	80.46	93.24

$^{**}p < 0.05$, $^{*}p < 0.1$

Table 1 displays the effect of anti-refugee hate crimes on SOEP respondents' attitudes towards refugees (column 2) and plans to help refugees (column 3). The Table provides point estimates and standard errors clustered on the federal state level for both outcomes, across three hate crime types, namely arson, assaults, and attacks on refugee housing. All models use MSE-optimal bandwidths that are permitted to differ in width across the cutoff, local linear regression, and triangular kernels. Respondents interviewed on the day of treatment are excluded, as are respondents for whom the treatment-designating hate crime occurs within 30 days of an earlier hate crime. Controls for arson: self-reported health. Controls for assaults: self-reported health, lagged immigration attitudes, migration background. Controls for attacks on refugee housing: sex and lagged AfD partisanship.

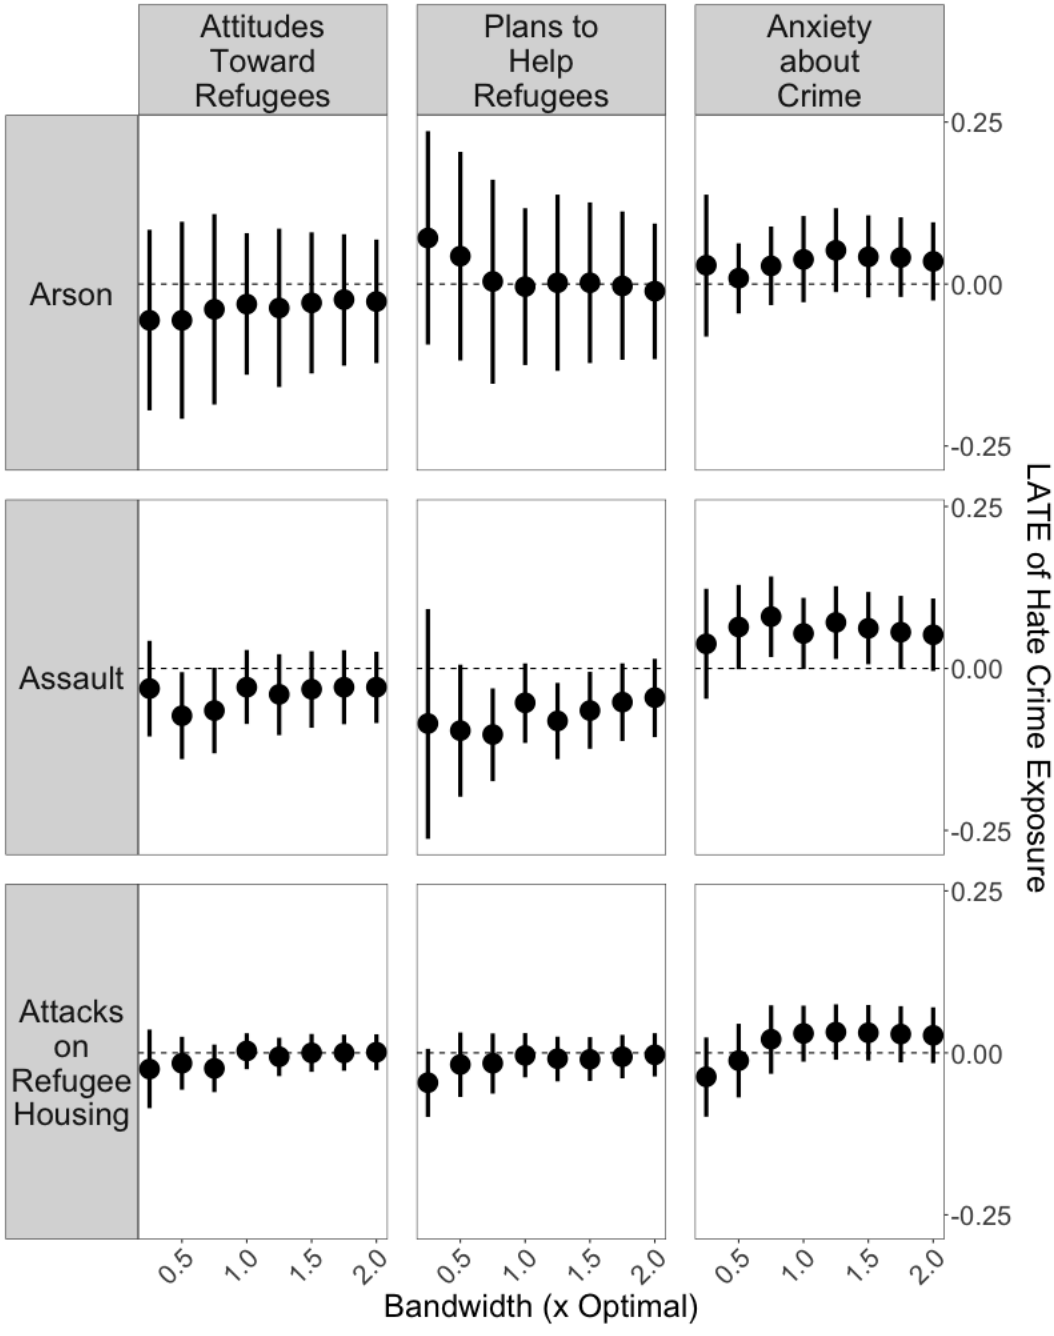

Figure A 6. Effect of Regional Hate Crime Exposure on Refugee Attitudes and Plans to Help: Alternative Bandwidths

Coefficients with 95% confidence intervals. Coefficients represent the difference in intercept below and above the cutoff, i.e. the LATE of hate crime exposure. Rows reflect different outcomes while columns reflect different hate crime event types. Position on the x-axis denotes bandwidth as a function of the MSE-optimal bandwidths used in the main model, which are permitted to differ on either side of the cutoff. Effective N varies by outcome, hate crime type, and bandwidth. All point estimates and errors based on robust estimation and inference method from *rdrobust* package. All models use local linear regression, triangular kernels, and errors clustered at the level of the German federal state. Respondents interviewed on the day of treatment are excluded, as are respondents for whom the treatment-designating hate crime occurs within 30 days of an earlier hate crime.

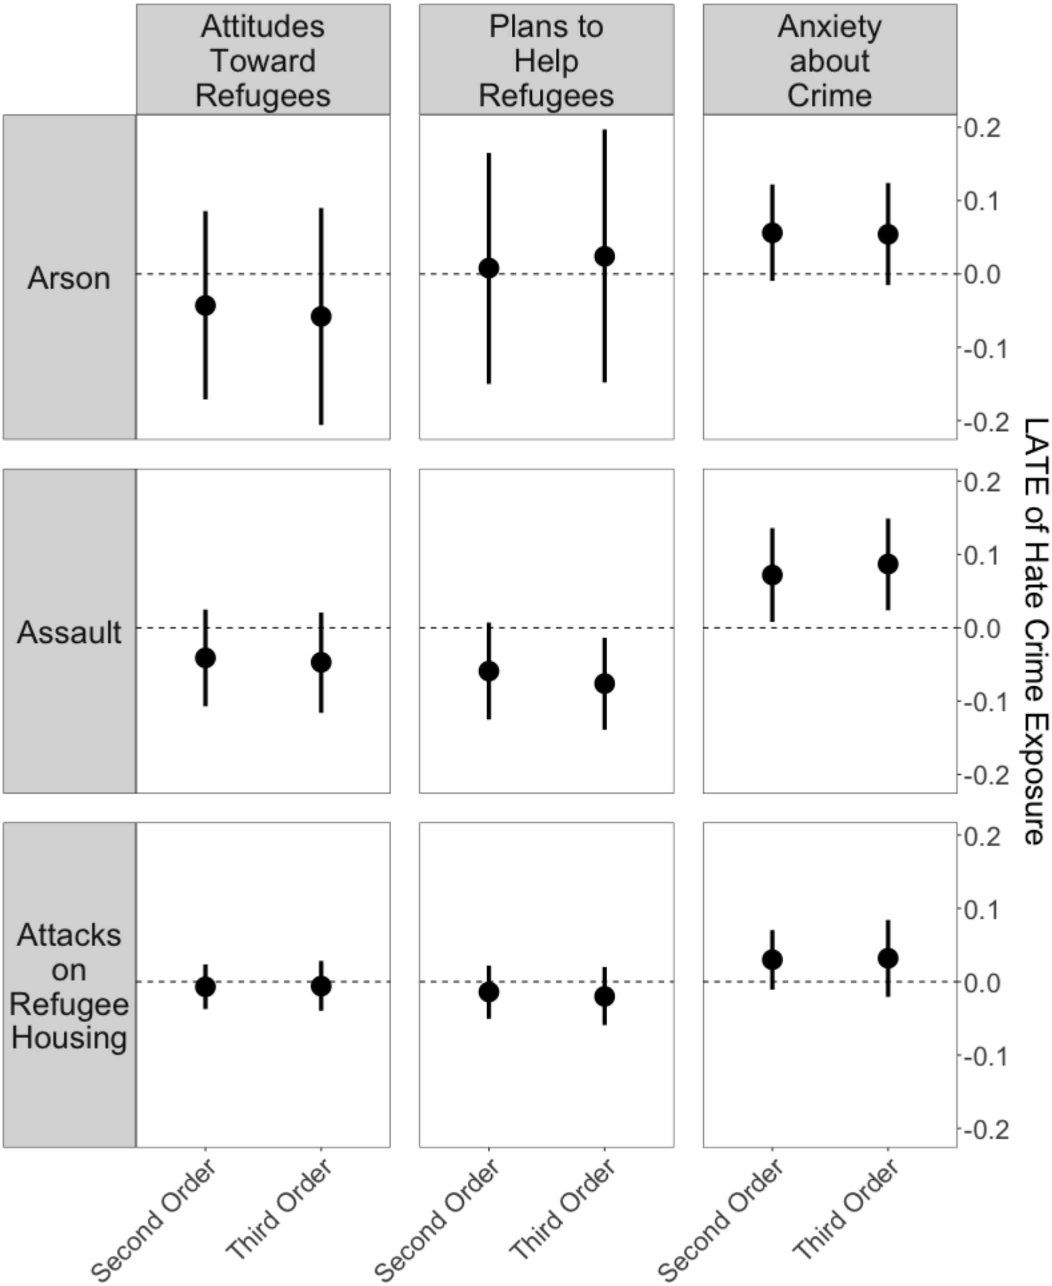

Figure A 7. Effect of Regional Hate Crime Exposure on Refugee Attitudes and Plans to Help: Local Regression with a Polynomial of Order Two and Three

Coefficients with 95% confidence intervals. Coefficients represent the difference in intercept below and above the cutoff, i.e. the LATE of hate crime exposure. Rows reflect different outcomes while columns reflect different hate crime event types. Effective N varies by outcome and hate crime type. All point estimates and errors based on robust estimation and inference method from *rdrobust* package. All models use MSE-optimal bandwidths that are permitted to differ in width across the cutoff, triangular kernels, and errors clustered at the regional unit of treatment. Respondents interviewed on the day of treatment are excluded, as are respondents for whom the treatment-designating hate crime occurs within 30 days of an earlier hate crime.

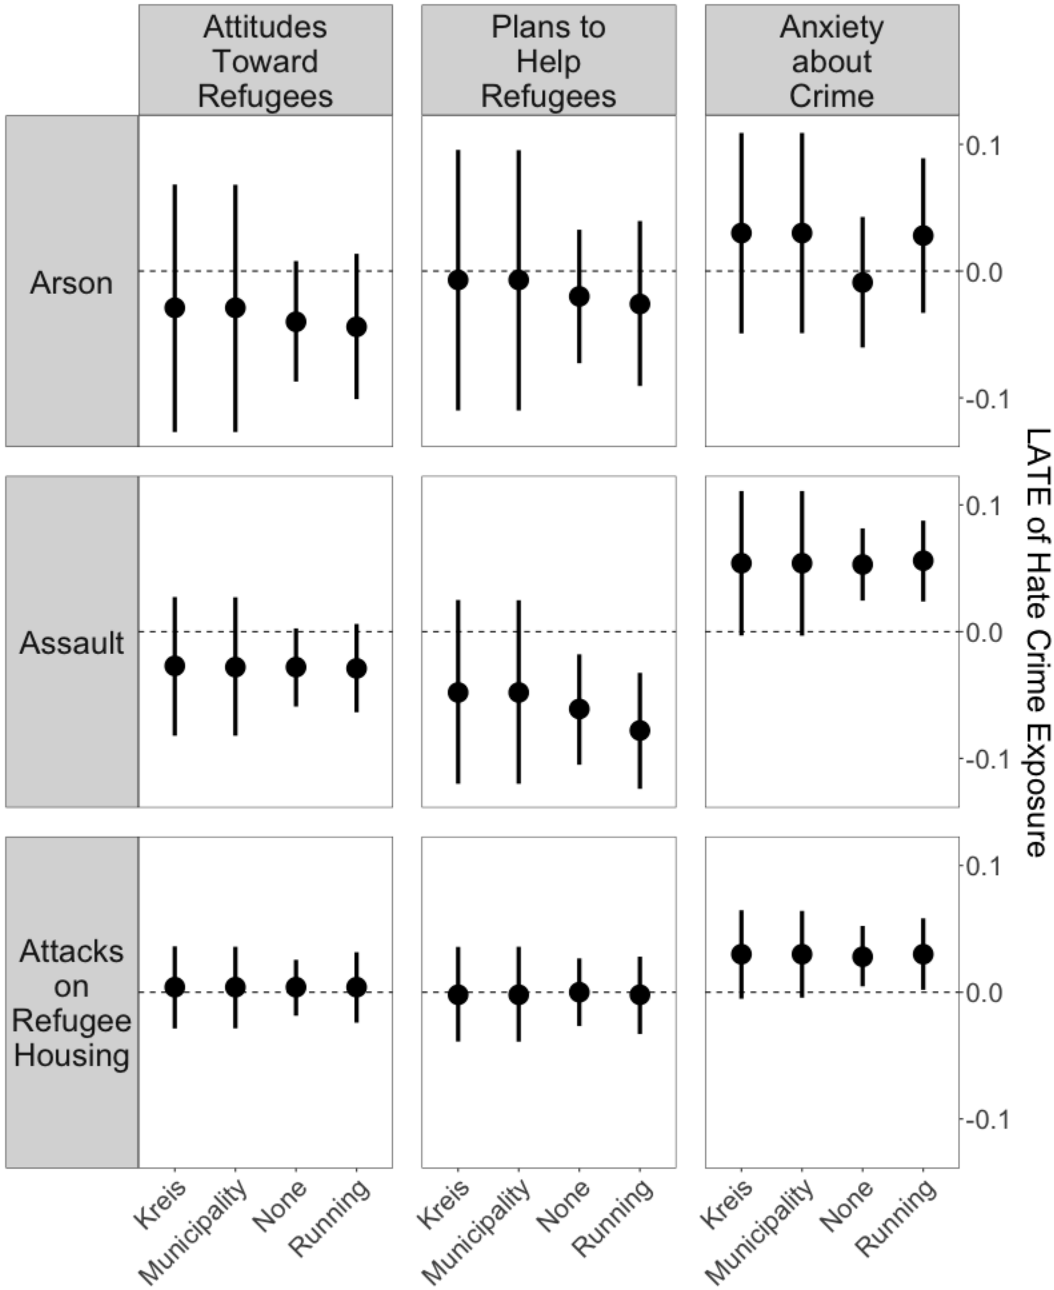

Figure A 8. Effect of Regional Hate Crime Exposure on Refugee Attitudes and Plans to Help: Alternative Standard Errors

Coefficients with 95% confidence intervals. Coefficients represent the difference in intercept below and above the cutoff, i.e. the LATE of hate crime exposure. Rows reflect different outcomes while columns reflect different hate crime event types. Effective N varies by outcome, standard error specification, and hate crime type. All point estimates and errors based on robust estimation and inference method from *rdrobust* package, local linear regression, and triangular kernels. Bandwidth specifications vary by mode. All models use MSE-optimal bandwidths. Respondents interviewed on the day of treatment are excluded, as are respondents for whom the treatment-designating hate crime occurs within 30 days of an earlier hate crime.

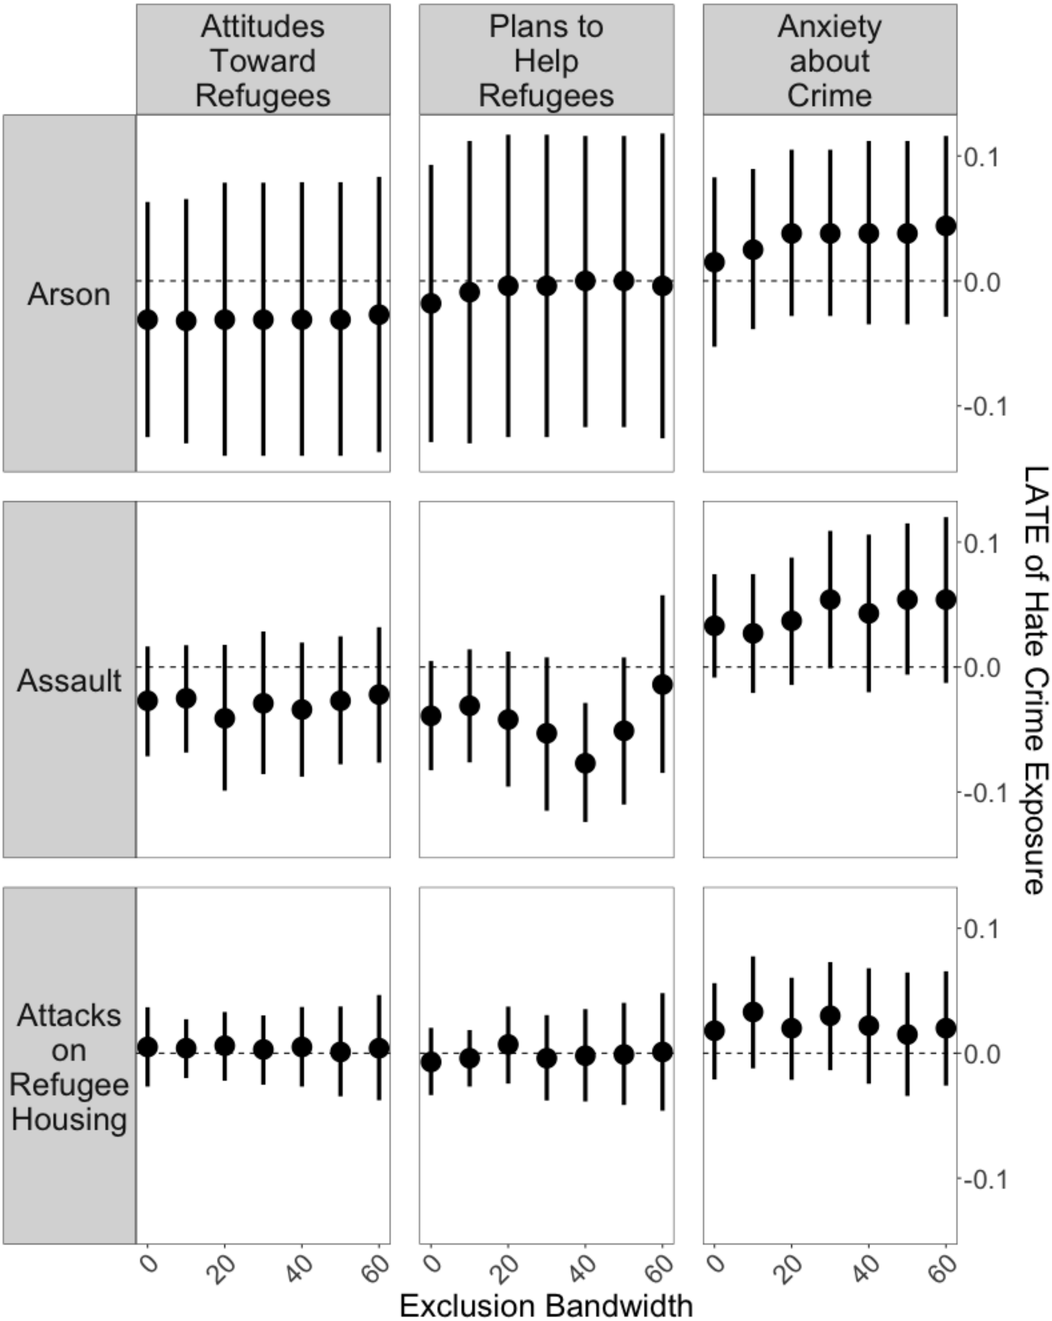

Figure A 9. Effect of Regional Hate Crime Exposure on Refugee Attitudes and Plans to Help: Alternative Exclusion Criteria for Multiply-Treated Units

Coefficients with 95% confidence intervals. Coefficients represent the difference in intercept below and above the cutoff, i.e. the LATE of hate crime exposure. Rows reflect different outcomes while columns reflect different hate crime event types. Position on the x-axis denotes exclusion criteria relating to previous exposure to hate crimes, prior to the treatment-designating hate crime. Effective N varies by outcome, hate crime type, and exclusion criteria. All models use MSE-optimal bandwidths that are permitted to differ in width across the cutoff, local linear regression, triangular kernels, and errors clustered at the regional unit of treatment. Respondents interviewed on the day of treatment are excluded.

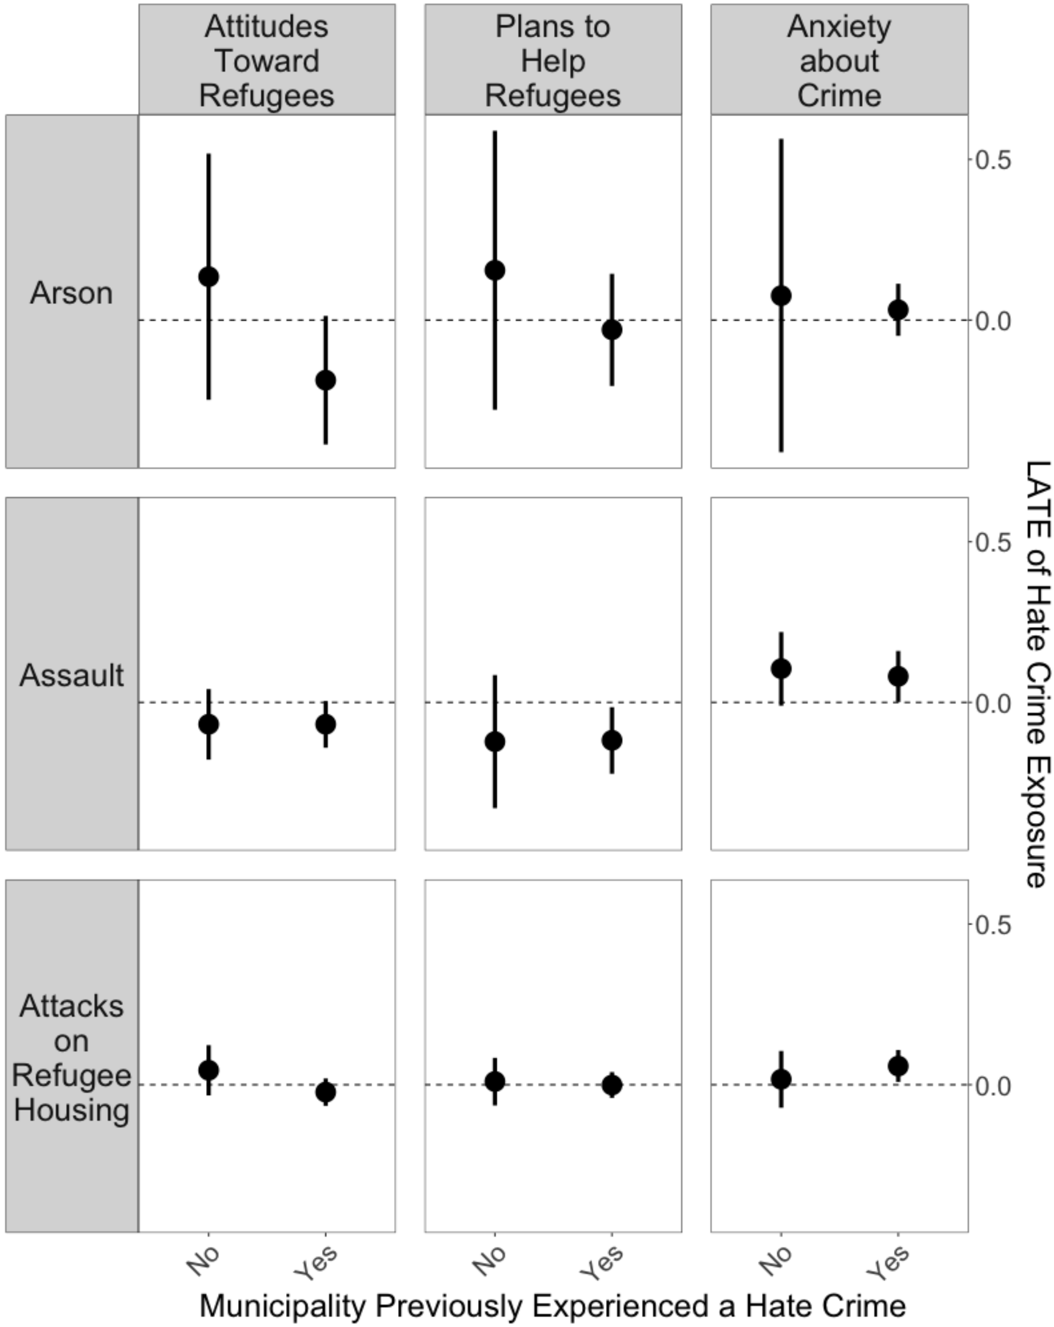

Figure A 10. Effect of Regional Hate Crime Exposure on Refugee Attitudes and Plans to Help: Sample Split by Previous Municipal Exposure

Coefficients with 95% confidence intervals. Coefficients represent the difference in intercept below and above the cutoff, i.e. the LATE of hate crime exposure. Rows reflect different outcomes while columns reflect different hate crime event types. Position on the x-axis denotes exclusion criteria relating to previous exposure to hate crimes, prior to the treatment-designating hate crime. Effective N varies by outcome and hate crime type. All models use MSE-optimal bandwidths that are permitted to differ in width across the cutoff, local linear regression, triangular kernels, and errors clustered at the regional unit of treatment. Respondents interviewed on the day of treatment are excluded, as are respondents for whom the treatment-designating hate crime occurs within 30 days of an earlier hate crime. Sample is split according to whether an individual resides in a municipality that previously experienced an anti-refugee hate crime recorded in ARVIG, i.e. since 2014.

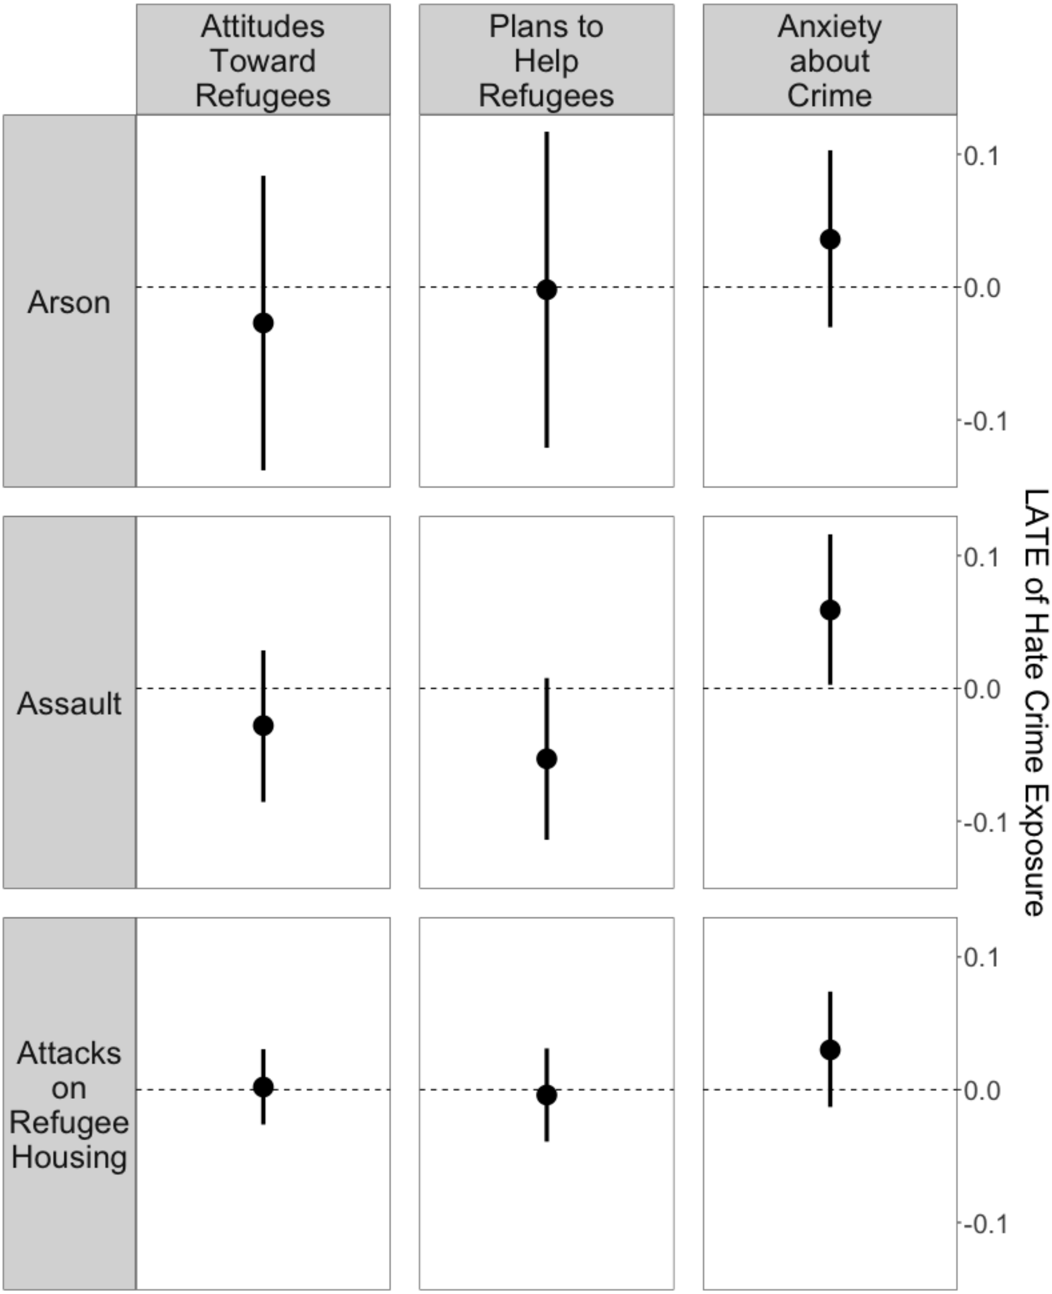

Figure A 11. Effect of Regional Hate Crime Exposure on Refugee Attitudes and Plans to Help: Controlling for Interview Day-of-Week

Coefficients with 95% confidence intervals. Coefficients represent the difference in intercept below and above the cutoff, i.e. the LATE of hate crime exposure. Rows reflect different outcomes while columns reflect different hate crime event types. Position on the x-axis denotes exclusion criteria relating to previous exposure to hate crimes, prior to the treatment-designating hate crime. Effective N varies by outcome and hate crime type. All models use MSE-optimal bandwidths that are permitted to differ in width across the cutoff, local linear regression, triangular kernels, and errors clustered at the regional unit of treatment. Respondents interviewed on the day of treatment are excluded, as are respondents for whom the treatment-designating hate crime occurs within 30 days of an earlier hate crime. This analysis adds a control for day of week in which a respondent is interviewed.

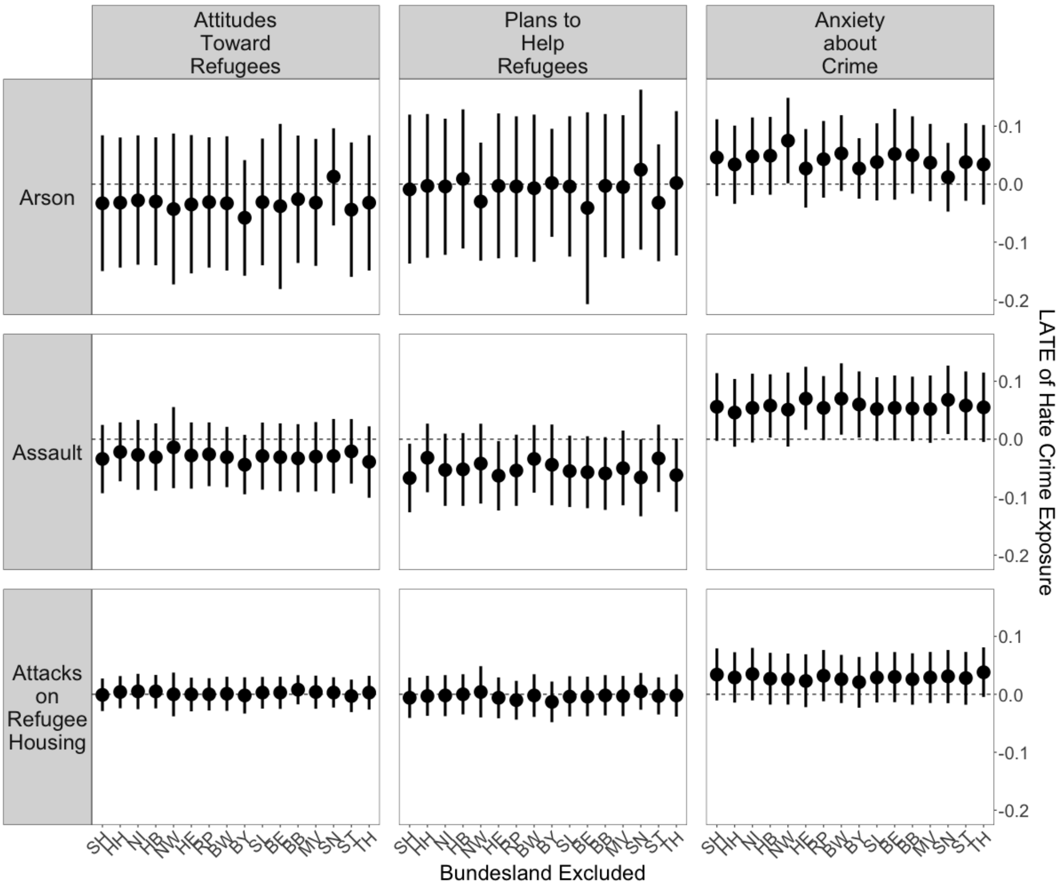

Figure A 12. Effect of Regional Hate Crime Exposure on Refugee Attitudes and Plans to Help: Sequentially Leaving Out Federal States

Coefficients with 95% confidence intervals. Coefficients represent the difference in intercept below and above the cutoff, i.e. the LATE of hate crime exposure. Rows reflect different outcomes while columns reflect different hate crime event types. Position on the x-axis denotes exclusion criteria relating to previous exposure to hate crimes, prior to the treatment-designating hate crime. Effective N varies by outcome and hate crime type. All models use MSE-optimal bandwidths that are permitted to differ in width across the cutoff, local linear regression, triangular kernels, and errors clustered at the regional unit of treatment. Respondents interviewed on the day of treatment are excluded, as are respondents for whom the treatment-designating hate crime occurs within 30 days of an earlier hate crime. This analysis sequentially drops respondents nested within each of the sixteen German federal states.

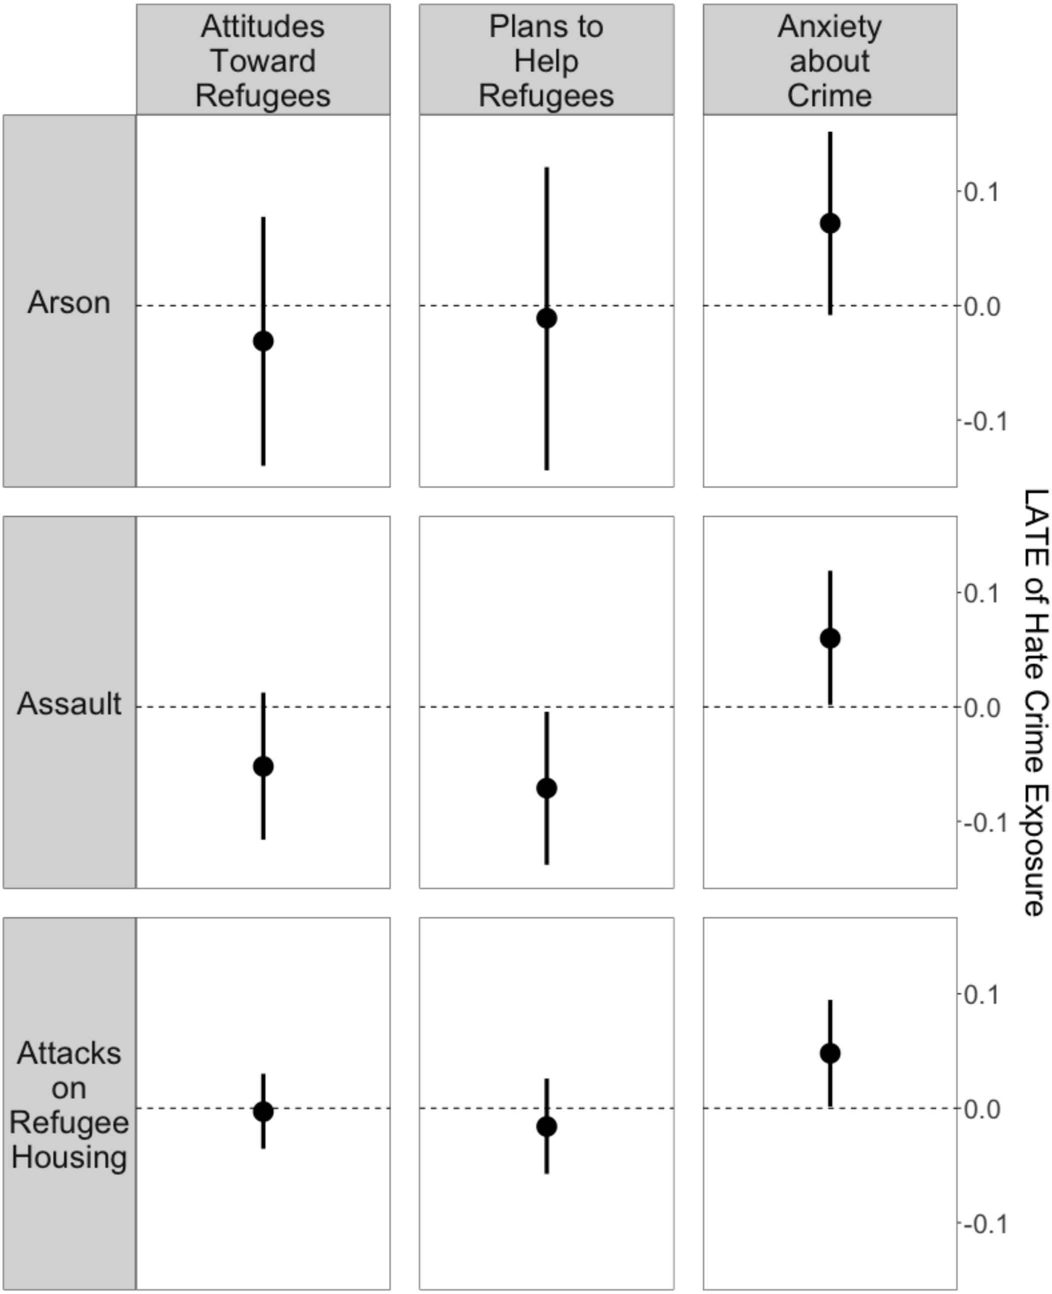

Figure A 13. Effect of Regional Hate Crime Exposure on Refugee Attitudes and Plans to Help: Including Only Municipalities with more than 10 Respondents

Coefficients with 95% confidence intervals. Coefficients represent the difference in intercept below and above the cutoff, i.e. the LATE of hate crime exposure. Rows reflect different outcomes while columns reflect different hate crime event types. Position on the x-axis denotes exclusion criteria relating to previous exposure to hate crimes, prior to the treatment-designating hate crime. Effective N varies by outcome and hate crime type. All models use MSE-optimal bandwidths that are permitted to differ in width across the cutoff, local linear regression, triangular kernels, and errors clustered at the regional unit of treatment. Respondents interviewed on the day of treatment are excluded, as are respondents for whom the treatment-designating hate crime occurs within 30 days of an earlier hate crime. This analysis includes respondents nested within municipalities that contain ten or more (non-refugee) respondents in 2016.

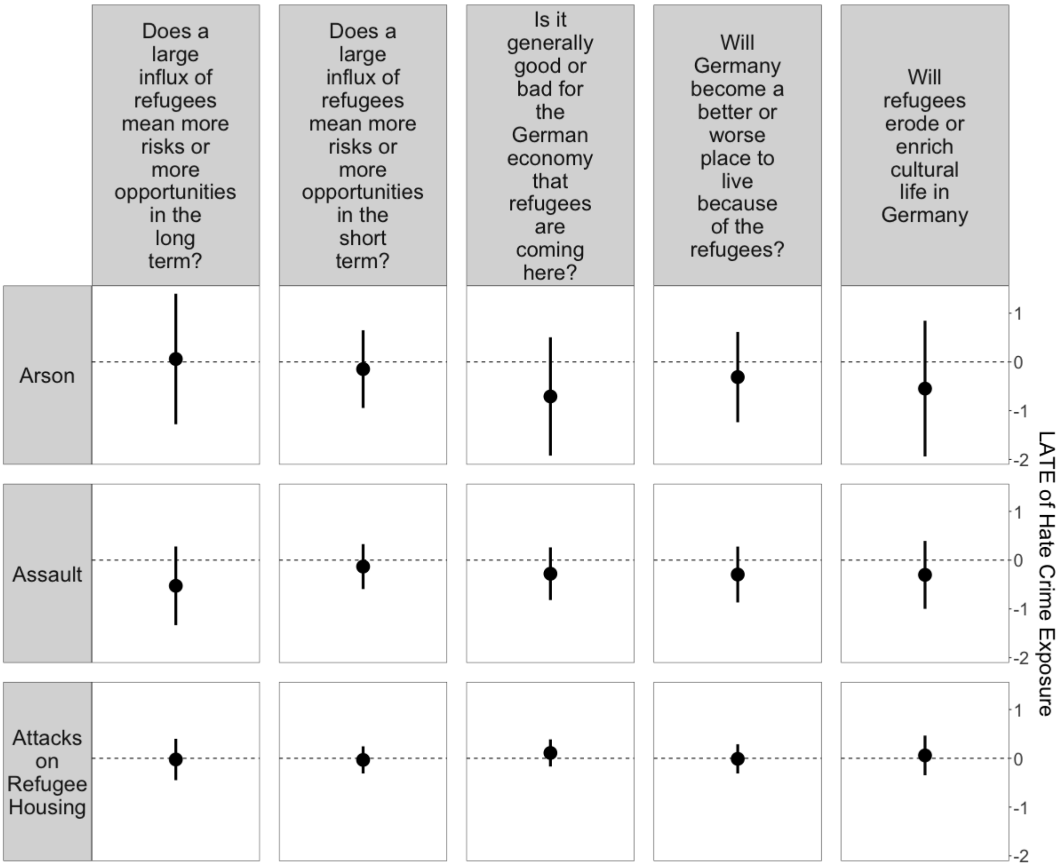

Figure A 14. Effect of Regional Hate Crime Exposure on Survey Item Components: Attitudes Towards Refugees

Coefficients with 95% confidence intervals. Coefficients represent the difference in intercept below and above the cutoff, i.e. the LATE of hate crime exposure. Rows reflect different survey items while columns reflect different hate crime event types. Effective N varies by outcome and hate crime type. All point estimates and errors based on robust estimation and inference method from *rdrobust* package. All models use MSE-optimal bandwidths that are permitted to differ in width across the cutoff, local linear regression, triangular kernels, and errors clustered at the regional unit of treatment. Respondents interviewed on the day of treatment are excluded, as are respondents for whom the treatment-designating hate crime occurs within 30 days of an earlier hate crime.

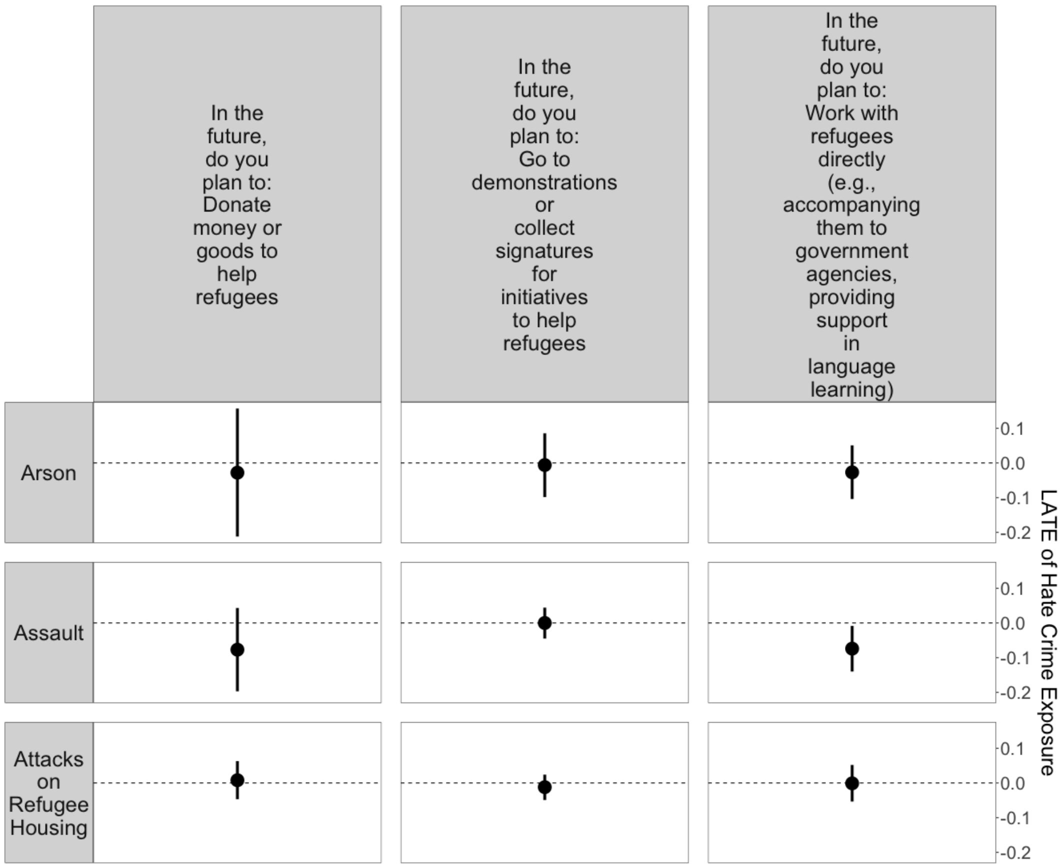

Figure A 15. Effect of Regional Hate Crime Exposure on Survey Item Components: Plans to Help Refugees

Coefficients with 95% confidence intervals. Coefficients represent the difference in intercept below and above the cutoff, i.e. the LATE of hate crime exposure. Rows reflect different survey items while columns reflect different hate crime event types. Shape of point estimate represents the level at which individuals have been treated. Effective N varies by outcome and hate crime type. All point estimates and errors based on robust estimation and inference method from *rdrobust* package. All models use MSE-optimal bandwidths that are permitted to differ in width across the cutoff, local linear regression, triangular kernels, and errors clustered at the regional unit of treatment. Respondents interviewed on the day of treatment are excluded, as are respondents for whom the treatment-designating hate crime occurs within 30 days of an earlier hate crime.

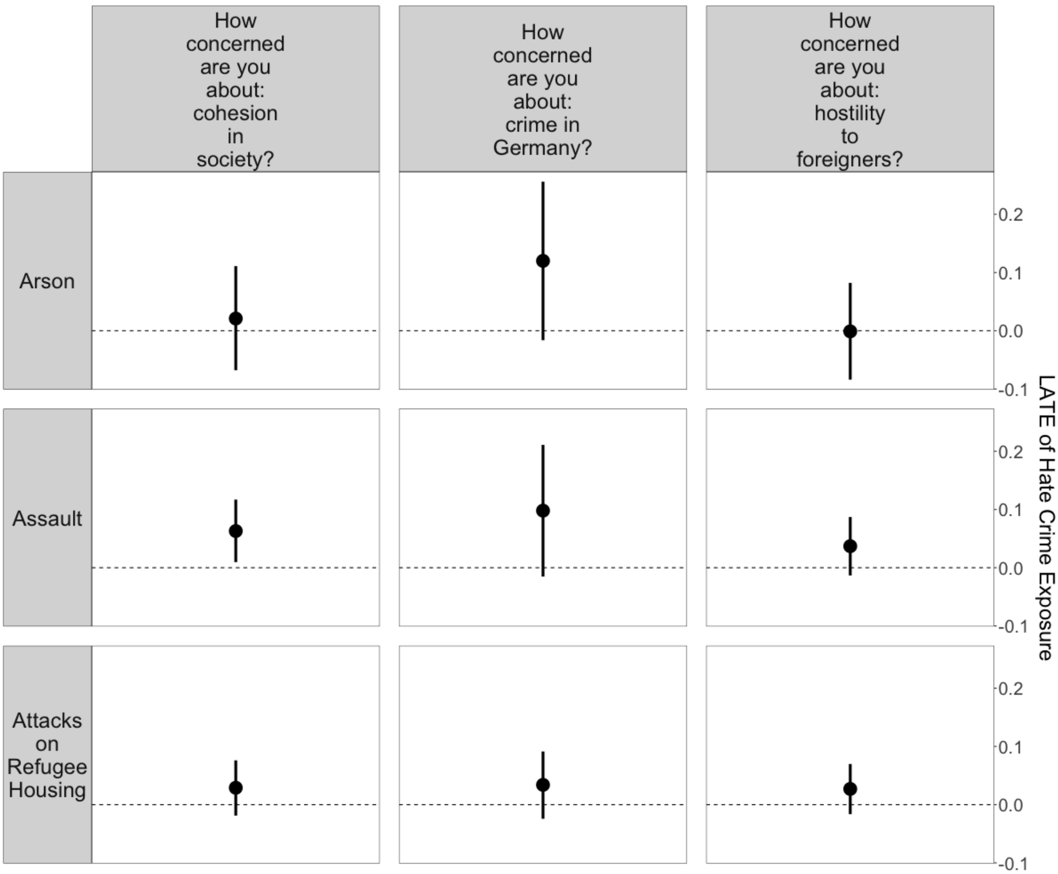

Figure A 16. Effect of Regional Hate Crime Exposure on Survey Item Components: Anxiety about Crime

Coefficients with 95% confidence intervals. Coefficients represent the difference in intercept below and above the cutoff, i.e. the LATE of hate crime exposure. Rows reflect different survey items while columns reflect different hate crime event types. Shape of point estimate represents the level at which individuals have been treated. Effective N varies by outcome and hate crime type. All point estimates and errors based on robust estimation and inference method from *rdrobust* package. All models use MSE-optimal bandwidths that are permitted to differ in width across the cutoff, local linear regression, triangular kernels, and errors clustered at the regional unit of treatment. Respondents interviewed on the day of treatment are excluded, as are respondents for whom the treatment-designating hate crime occurs within 30 days of an earlier hate crime.

Appendix 3.4 Heterogeneity in Effect Across Municipal and Individual Characteristics

Whether anti-refugee hate crime events shape attitudes towards refugees and plans to help refugees may depend on who is exposed. One group that is likely to be affected by anti-refugee hate crime are those who are skeptical of immigrants and immigration at baseline. Krause and Matsunaga (2023) show that right-wing terrorist violence in Germany increases support for radical right parties, and that this effect is driven by individuals with a conservative disposition realigning themselves with the AfD in the aftermath of right wing violence. Eger and Olzak (2023) come to similar conclusions about levels of violence against refugees: individuals negatively disposed towards immigrants may turn towards radical right parties when exposed to higher levels of violence against refugees. Here we describe a set of tests that seek to unpack whether effects of hate crime exposure on refugee-facing attitudes are conditional on individual or municipal characteristic.

We first explore how individual-level characteristics of non-target group members moderate the effect of exposure to anti-refugee hate crimes. We estimate two sets of split models. In the first analysis, we split 2016 respondents along their pre-treatment (2015) immigration attitudes and compare the LATE of hate crime exposure among those who expressed high levels of concern about immigration, some level of concern about immigration, and no concern about immigration.

In the second analysis, we separate respondents into right-wing, left-wing, and non-partisan respondents, again using their pre-treatment (2015) partisanship. Respondents are classified as right-wing if they report feeling close to the AfD, the CDU/CSU, small radical right parties, or the FDP. Respondents are classified as left-wing if they report feeling close to the Green Party, the Left Party (Linke), or the SPD. We aggregate because the number of individual partisans for specific parties is too small for reliable analysis. Next, we run separate regression models that estimate the effects of exposure for right-wing, left-wing, and non-partisan respondents.

Results for these two tests are presented in Figure A17 (immigration attitudes) and Figure A18 (partisanship). Interestingly, individuals who are very concerned about immigration (support right-wing parties) do not exhibit more negative changes in their attitudes towards refugees than individuals who are not concerned about immigration (individuals who support left-wing parties). The LATE of exposure to anti-refugee hate crimes is not significant in any sample split, regardless of hate crime type.

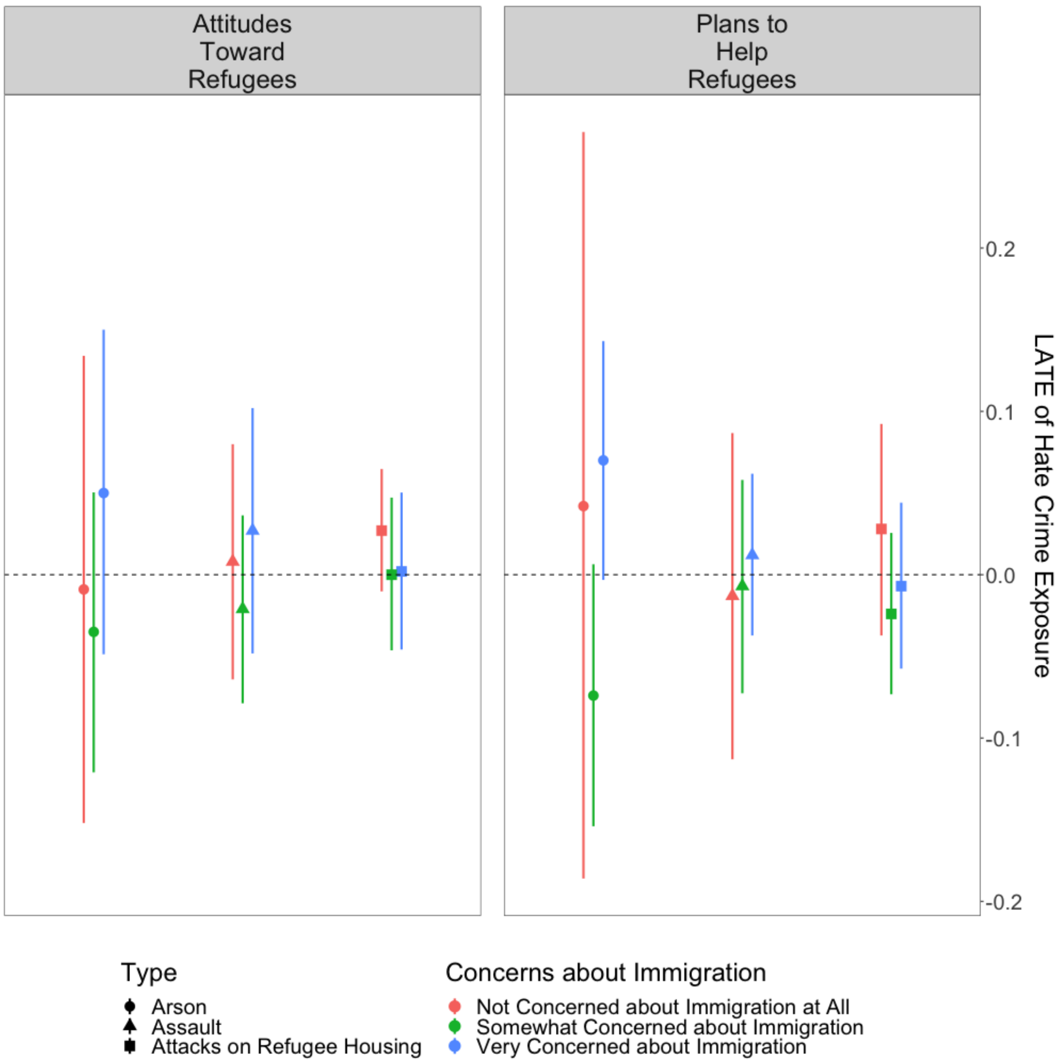

Figure A 17. Effect of Regional Hate Crime Exposure on Refugee Attitudes and Plans to Help, Across (Pre-treatment) Immigration Attitude

Coefficients with 95% confidence intervals. Coefficients represent the difference in intercept below and above the cutoff, i.e. the LATE of hate crime exposure. Columns reflect different outcomes while rows reflect 2015 partisanship (top two rows) and 2015 level of concern about immigration (bottom two rows) and. Shape of coefficient reflects type of hate crime event. Color reflects value of the individual moderator. Effective N varies by outcome, hate crime type, and pre-treatment level of concern. All point estimates and errors based on robust estimation and inference method from *rdrobust* package. All models use MSE-optimal bandwidths that are permitted to differ in width across the cutoff, local linear regression, triangular kernels, and errors clustered at the regional unit of treatment. Respondents interviewed on the day of treatment are excluded, as are respondents for whom the treatment-designating hate crime occurs within 30 days of an earlier hate crime.

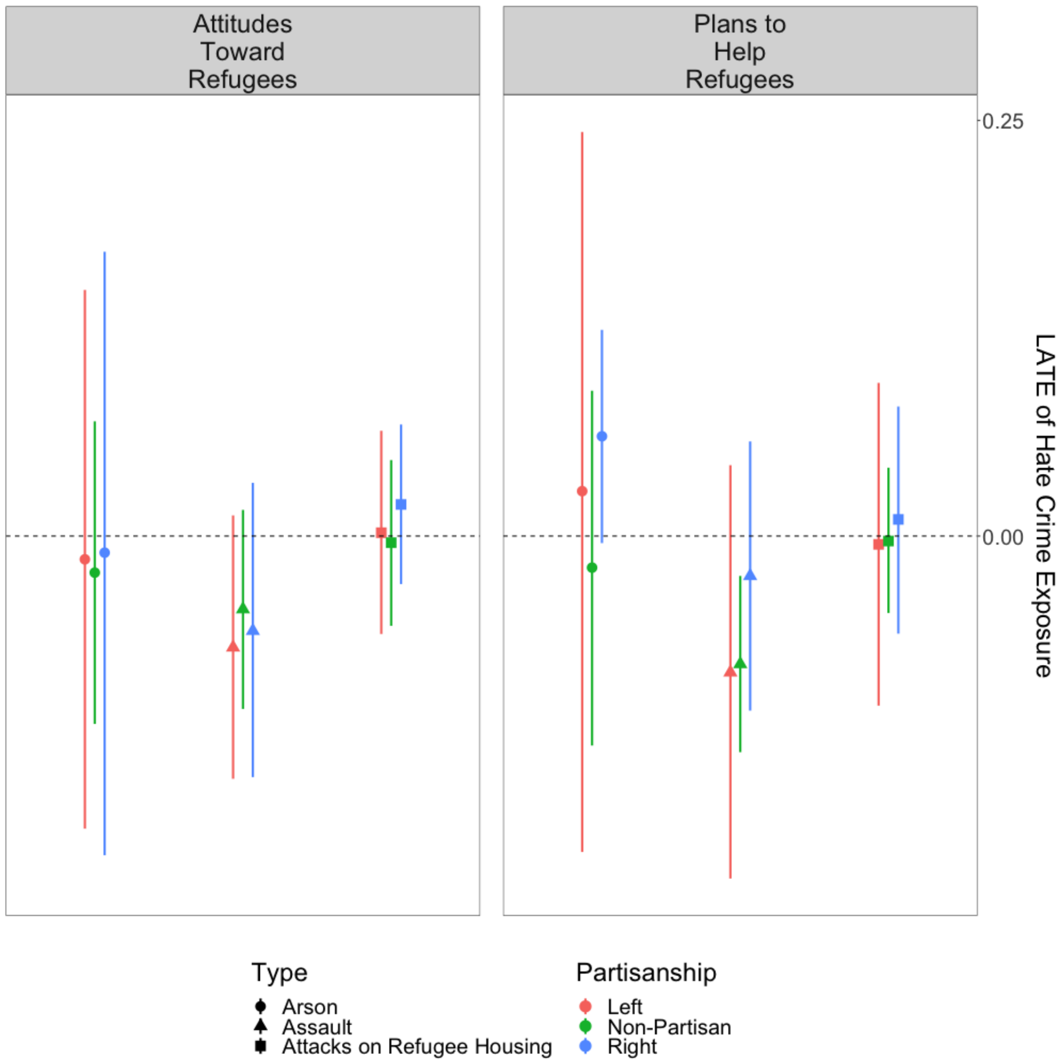

Figure A 18. Effect of Regional Hate Crime Exposure on Refugee Attitudes and Plans to Help, Across (Pre-treatment) Partisanship

Coefficients with 95% confidence intervals. Coefficients represent the difference in intercept below and above the cutoff, i.e. the LATE of hate crime exposure. Columns reflect different outcomes while rows reflect 2015 partisanship (top two rows) and 2015 level of concern about immigration (bottom two rows) and. Shape of coefficient reflects type of hate crime event. Color reflects value of the individual moderator. Effective N varies by outcome, hate crime type, and pre-treatment level of concern. All point estimates and errors based on robust estimation and inference method from *rdrobust* package. All models use MSE-optimal bandwidths that are permitted to differ in width across the cutoff, local linear regression, triangular kernels, and errors clustered at the regional unit of treatment. Respondents interviewed on the day of treatment are excluded, as are respondents for whom the treatment-designating hate crime occurs within 30 days of an earlier hate crime.

Another possibility is that effects of exposure may vary across municipal characteristics (as opposed to individual characteristics). Certain regions are more welcoming of newcomers due to the composition of local populations. Negative effects of exposure to anti-refugee hate crimes may be most pronounced in areas that are less welcoming towards newcomers at baseline. To test this possibility we split our sample along characteristics that differentiate municipalities in terms of baseline sentiment towards refugees. Specifically, we compare effects across urban municipalities and rural municipalities, across municipalities in the former west and municipalities in the former east, across municipalities in the highest and lowest tertiles of 2013 vote shares for right-wing and left-wing parties, across municipalities in the highest and lowest tertiles of 2013 vote shares for *specific* political parties, across municipalities in the highest and lowest tertiles of foreign-born population share, as well as across municipalities in the highest and lowest tertiles of unemployment rate. As shown in Figures A19 - A24, these municipal characteristics are highly predictive of refugee attitudes. Respondents in urban municipalities, municipalities in the former west, municipalities in the upper tertile of FDP and Green Party vote share and foreign-born population share are significantly more likely to voice positive refugee-facing attitudes. So too are municipalities in the lower tertiles of right-wing vote share, AfD vote share, CDU/CSU voteshare, Linke vote share, and unemployment.[11]

Despite these stark differences in the baseline level of sentiment towards refugees, our results are not indicative of heterogeneity in effects of exposure across municipal characteristics. The LATE of exposure to hate crime events typically does not differ significantly across urban or rural areas (Figure A25), across municipalities in the Former East or Former West (Figure A26, across municipalities in the upper and lower tertiles of left- or right-wing vote shares (Figure A27 and A28), across municipalities in the upper and lower tertiles of all six major German parties (Figures A29-A34), across municipalities in the upper and lower tertiles of foreign-born population share (Figure A35), or across municipalities in the upper and lower tertiles of unemployment (Figure A36). For certain moderators, the effect of exposure is significant for one value but not for another. For example, exposure to assaults in urban municipalities is associated with a negative effect on plans to help refugees (at $p < 0.05$), but not in rural areas. However, the two coefficients - i.e. the coefficient for exposure to assaults in urban areas and the coefficient for exposure in rural areas – are not distinct

11. To be clear, disentangling which if any of these contextual factors act as causal drivers of welcoming attitudes towards refugees is beyond the scope of this study. Our goal is merely to compare effects of anti-refugee hate crimes across municipalities that are more/less welcoming at baseline. Because the political, economic, and demographic characteristics of more/less welcoming contexts are highly correlated, we split our sample along a number of municipal covariates.

from one another in any case explored here.

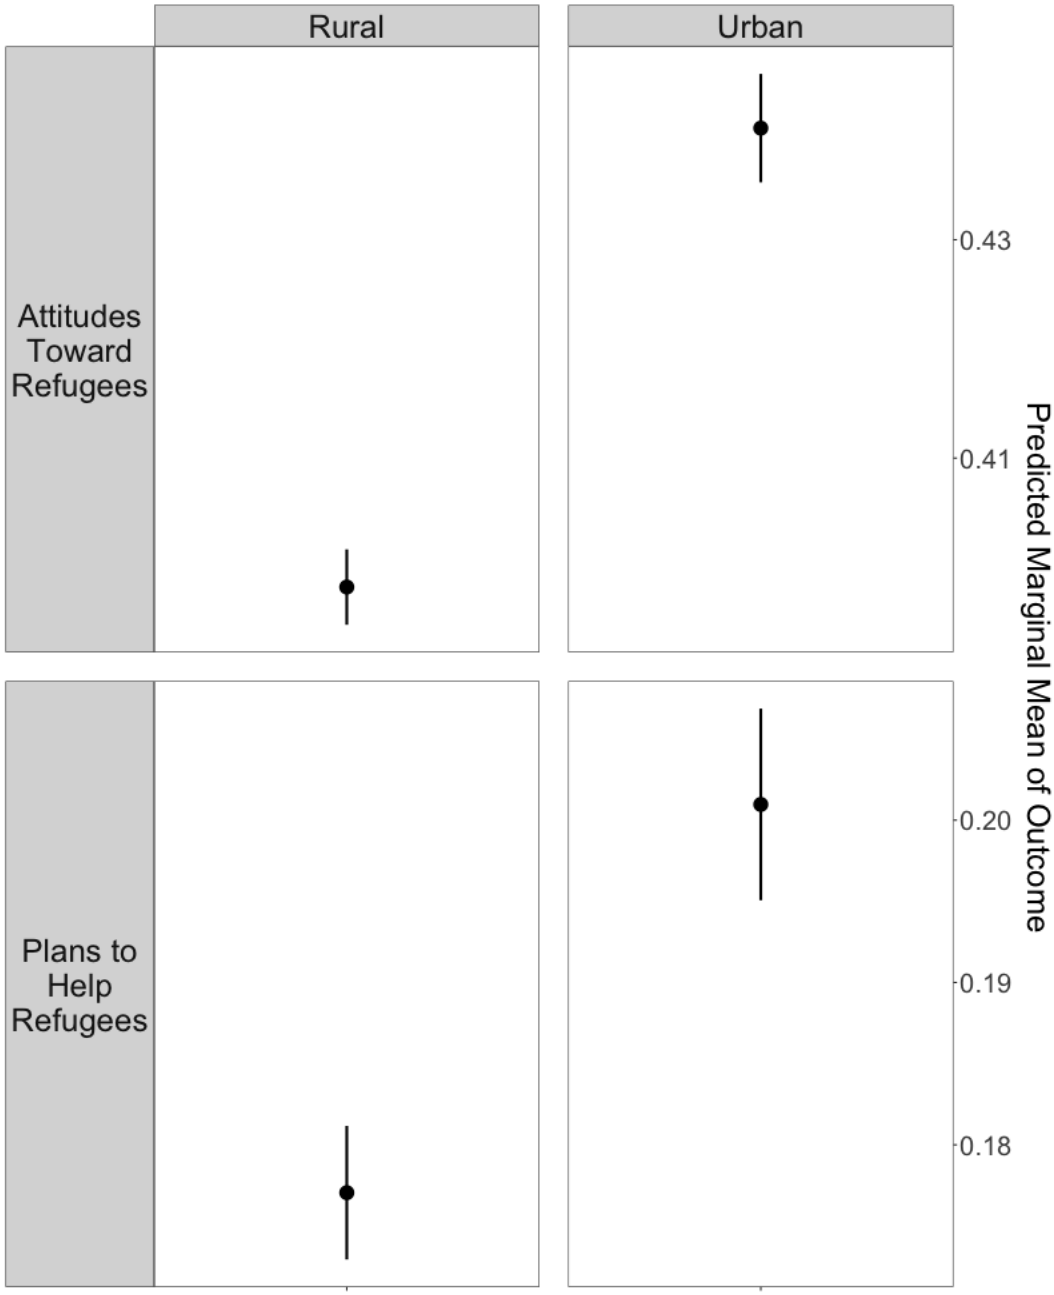

Figure A 19. Predicted Marginal Means of Outcomes, Across Type of Municipality

Predicted values with 95% confidence intervals, derived from linear regression models. Rows reflect different outcomes while columns reflect whether a municipality is urban (Kreisfreiestadt) or a rural (a municipality nested within a Landeskreis).

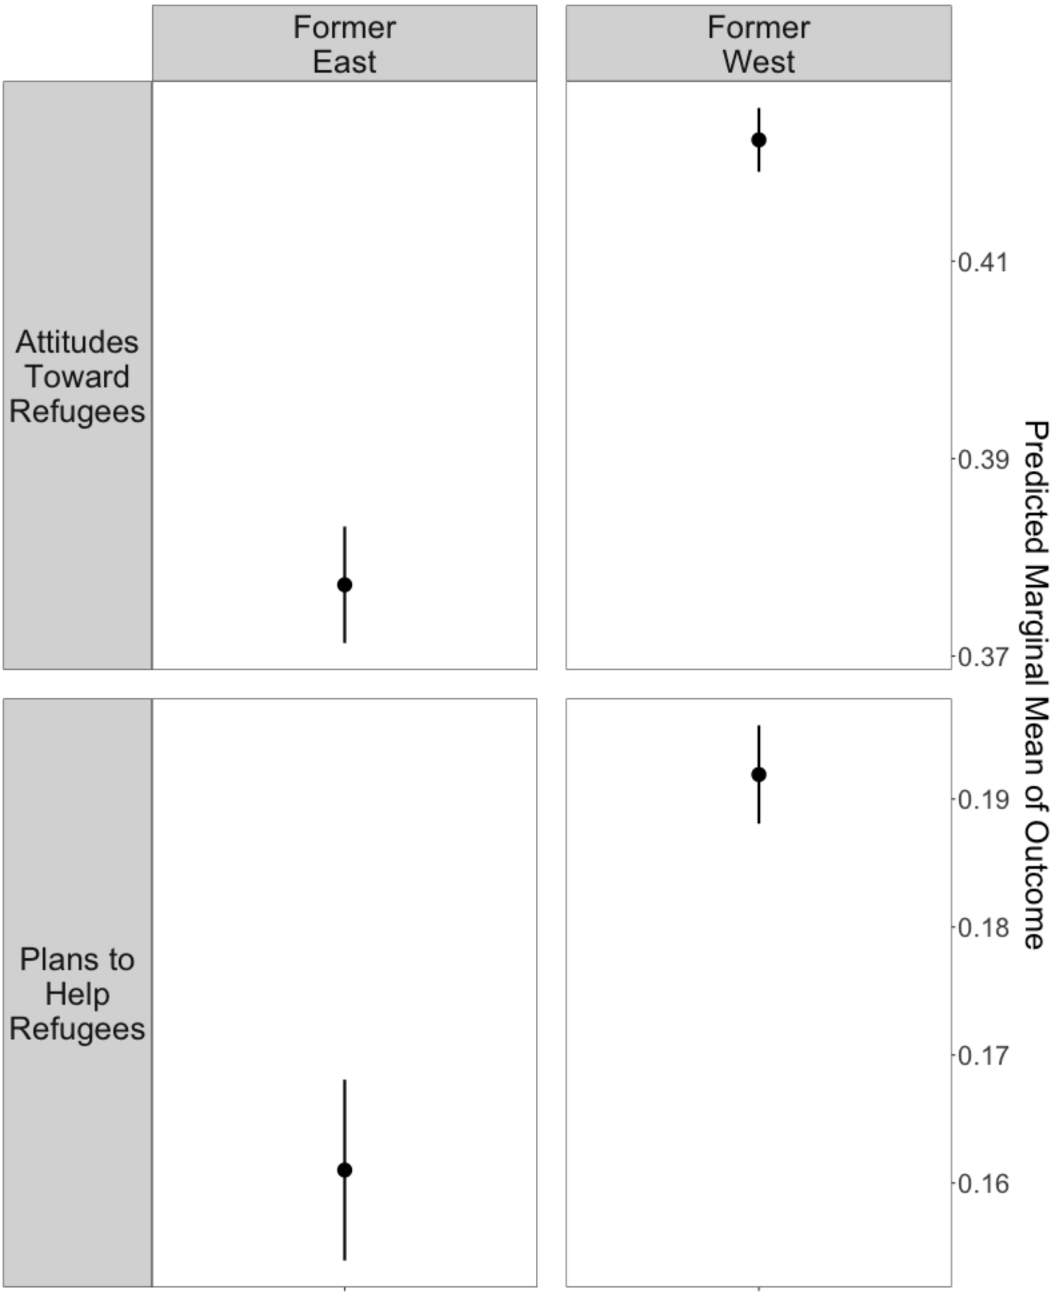

Figure A 20. Predicted Marginal Means of Outcomes, Across Region

Predicted values with 95% confidence intervals, derived from linear regression models. Rows reflect different outcomes while columns reflect whether a municipality is in the former East or West.

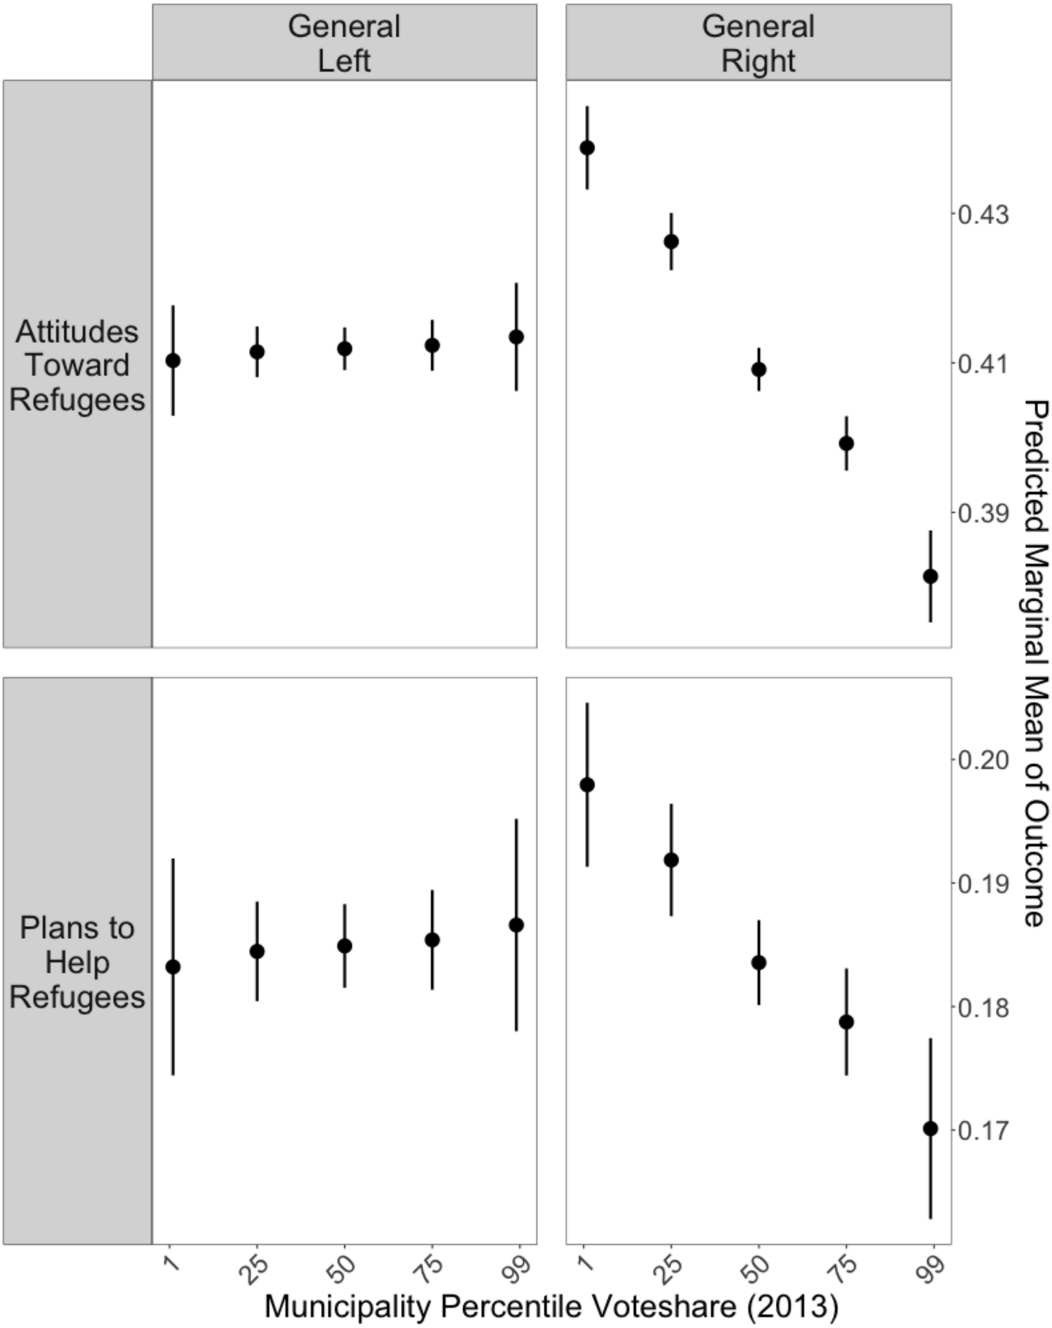

Figure A 21. Predicted Marginal Means of Outcomes, Across 2013 Municipal Vote shares for Left- and Right-Wing Parties

Predicted values with 95% confidence intervals, derived from linear regression models. Rows reflect different outcomes while columns reflect left-wing (Green, Linke, and SPD) and right-wing (AfD, CDU/CSU, and FDP) parties in the 2013 German election.

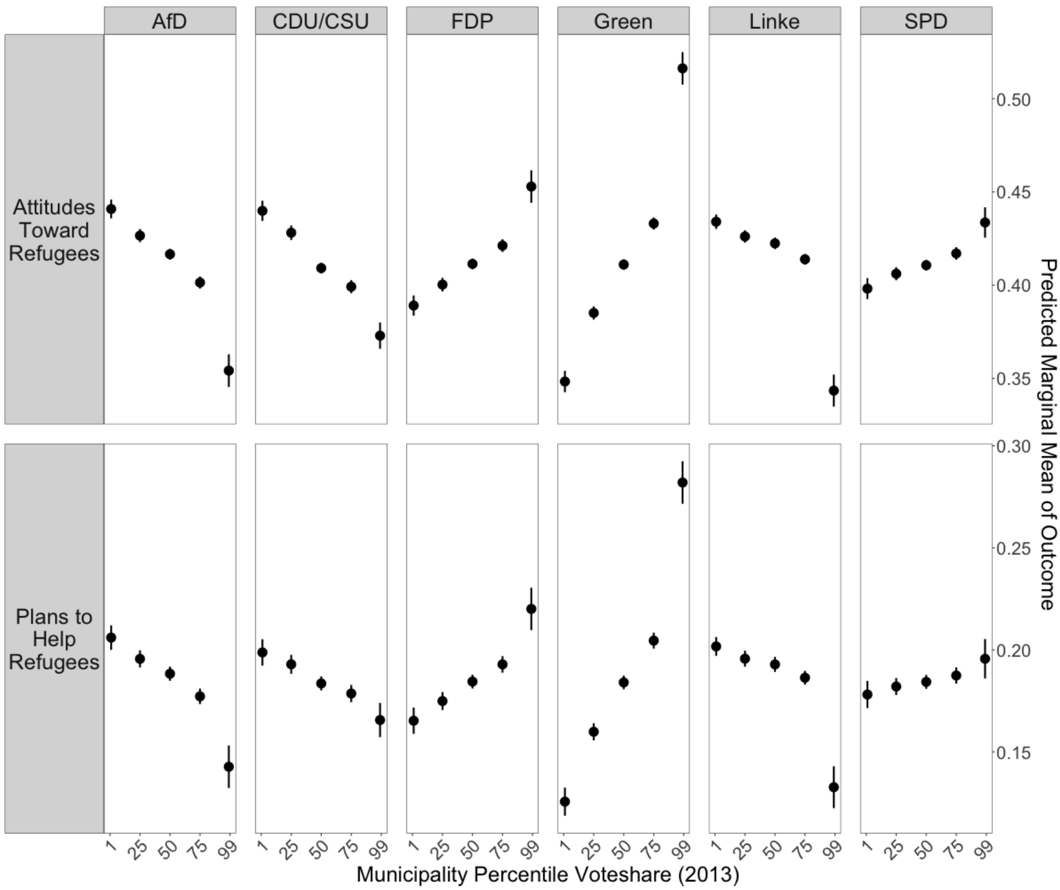

Figure A 22. Predicted Marginal Means of Outcomes, Across 2013 Municipal Vote shares for Specific Political Parties

Predicted values with 95% confidence intervals, derived from linear regression models. Rows reflect different outcomes while columns reflect the municipal vote shares of specific parties in the 2013 German election.

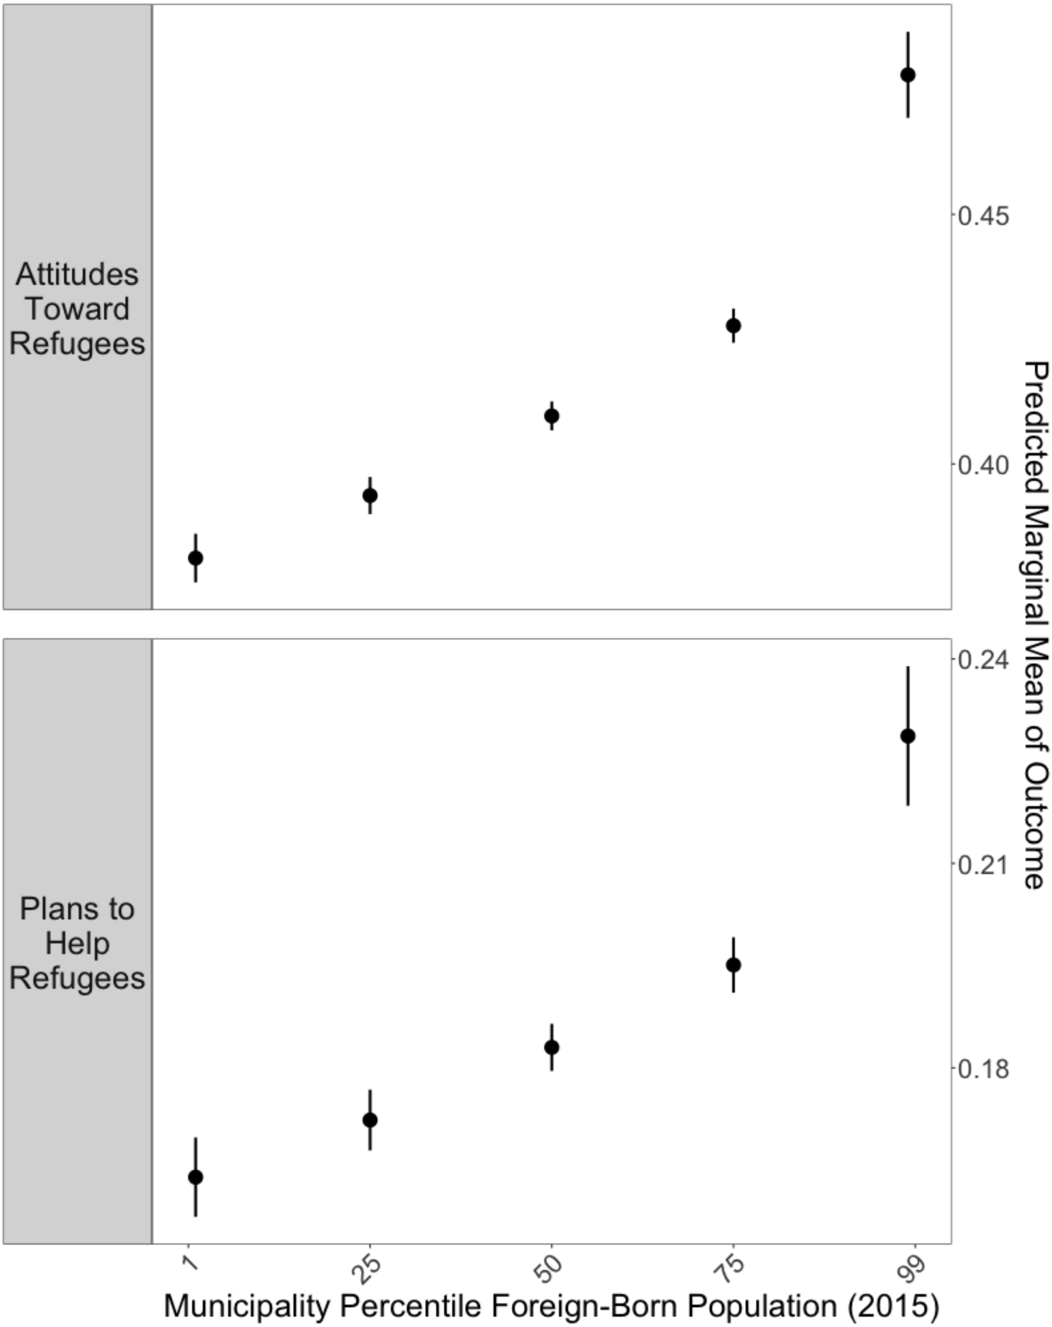

Figure A 23. Predicted Marginal Means of Outcomes, Across Municipal Foreign-Born Population Share

Predicted values with 95% confidence intervals, derived from linear regression models. Rows reflect different outcomes.

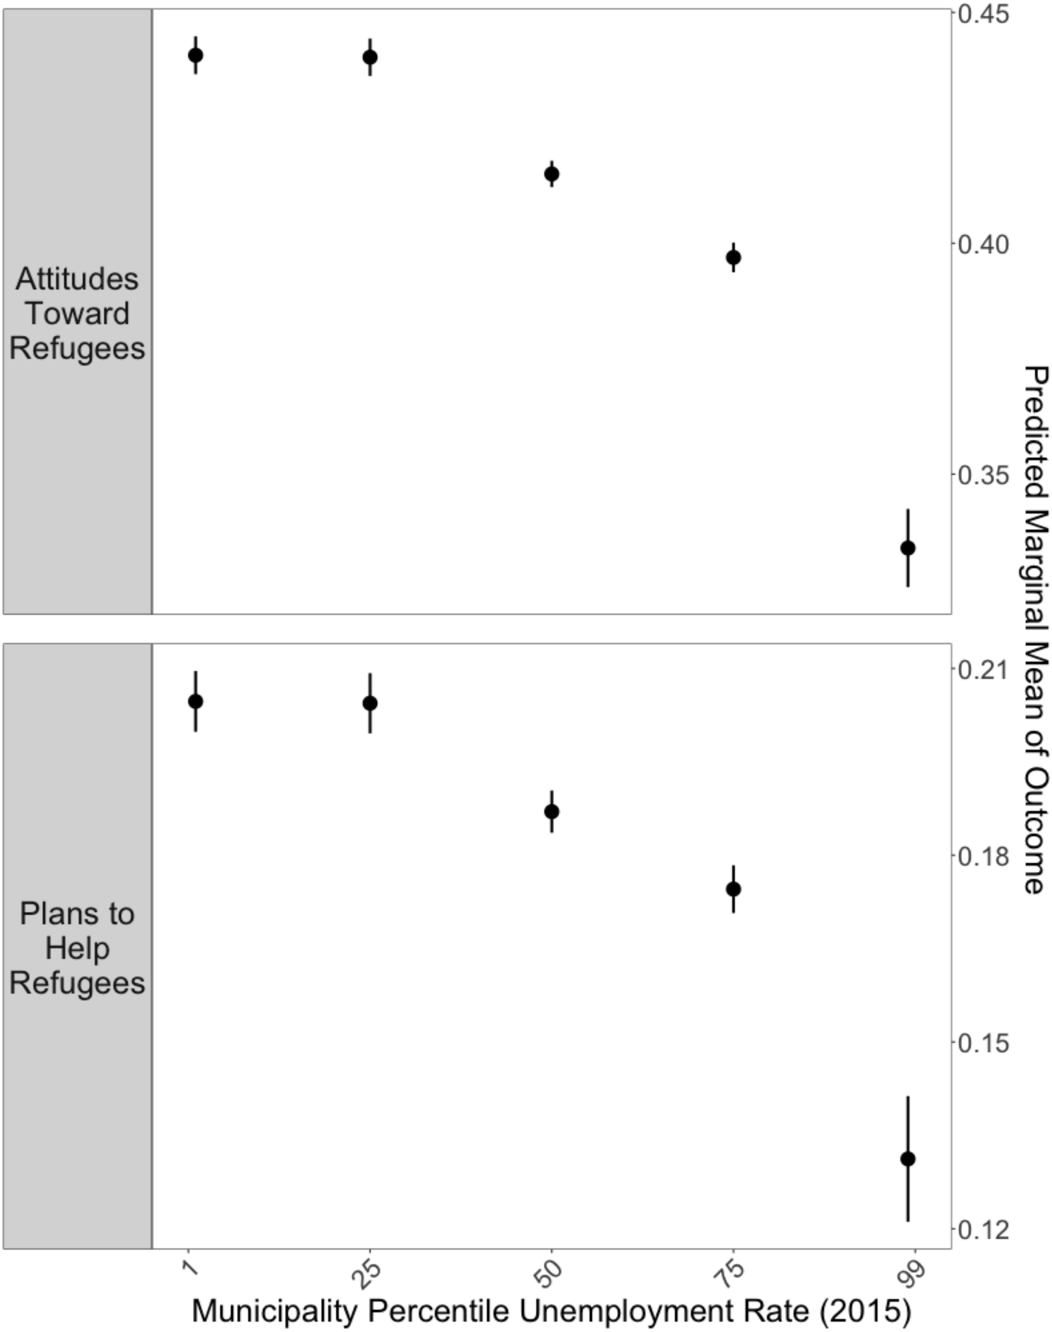

Figure A 24. Predicted Marginal Means of Outcomes, Across Municipal Unemployment Rate

Predicted values with 95% confidence intervals, derived from linear regression models. Rows reflect different outcomes.

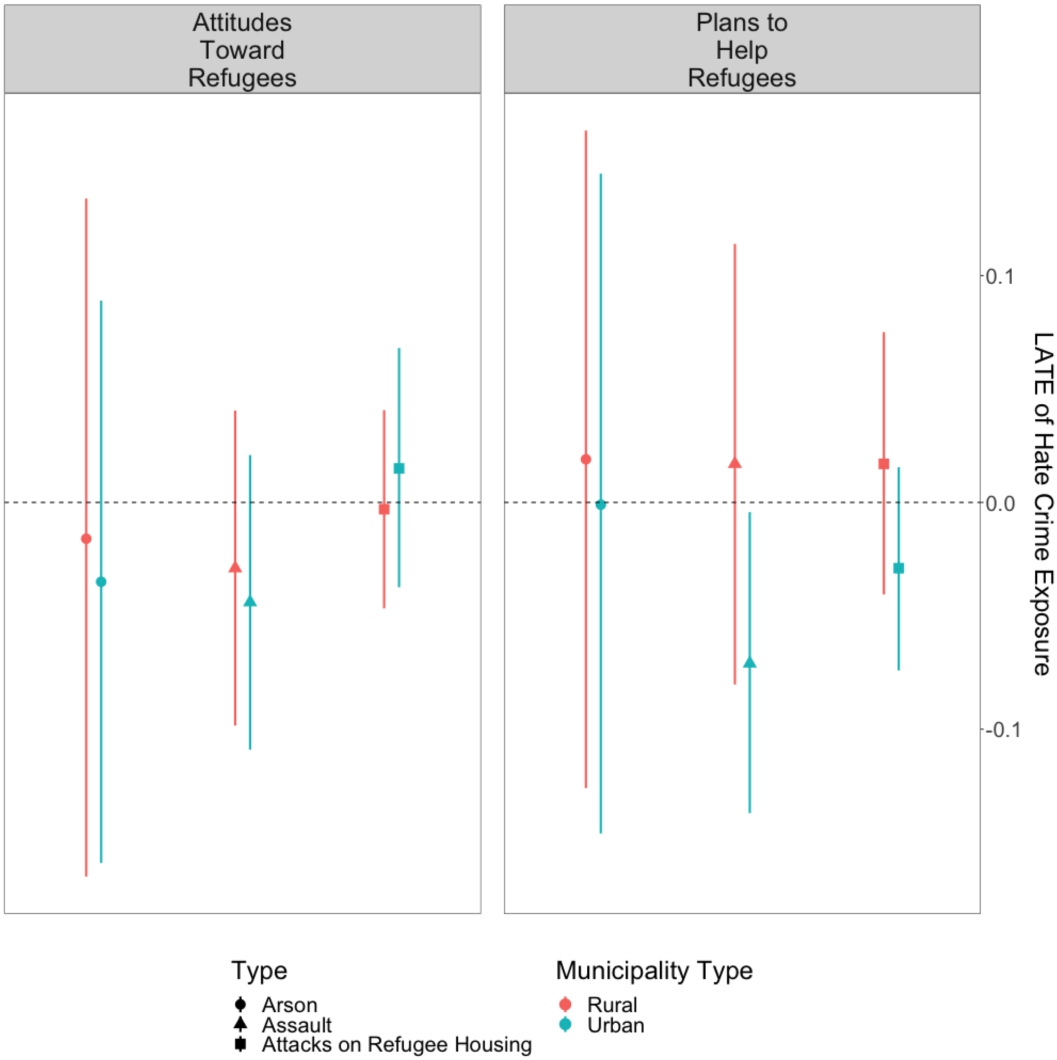

Figure A 25. Effect of Regional Hate Crime Exposure on Refugee Attitudes and Plans to Help, Across Type of Municipality

Coefficients with 95% confidence intervals. Coefficients represent the difference in intercept below and above the cutoff, i.e. the LATE of hate crime exposure. Rows reflect different outcomes while columns reflect whether a municipality is urban (Kreisfreiestadt) or a rural (a municipality nested within a Landeskreis). Shape of coefficient reflects type of hate crime event. Effective N varies by outcome, hate crime type, and pre-treatment level of concern. All point estimates and errors based on robust estimation and inference method from *rdrobust* package. All models use MSE-optimal bandwidths that are permitted to differ in width across the cutoff, local linear regression, triangular kernels, and errors clustered at the regional unit of treatment. Respondents interviewed on the day of treatment are excluded, as are respondents for whom the treatment-designating hate crime occurs within 30 days of an earlier hate crime.

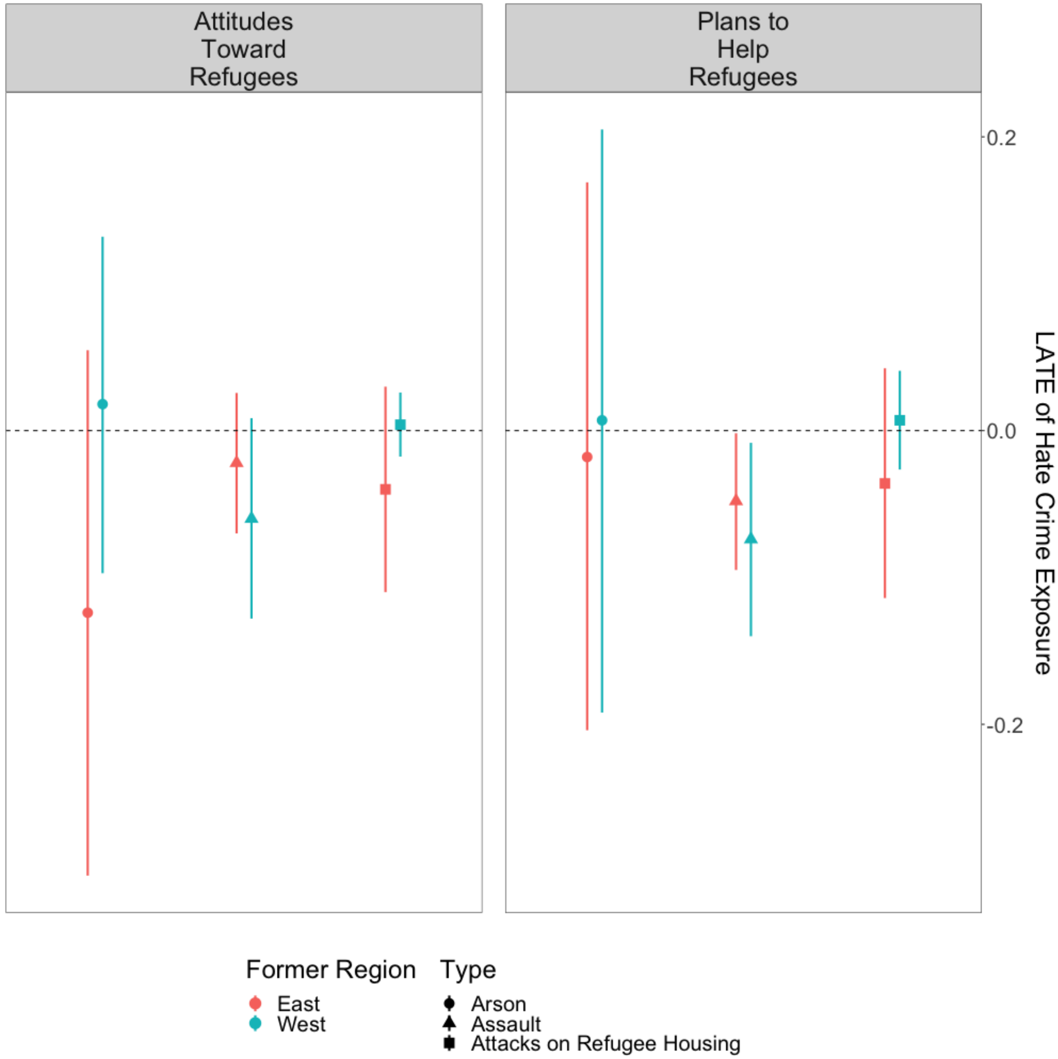

Figure A 26. Effect of Regional Hate Crime Exposure on Refugee Attitudes and Plans to Help, Across Region

Coefficients with 95% confidence intervals. Coefficients represent the difference in intercept below and above the cutoff, i.e. the LATE of hate crime exposure. Rows reflect different outcomes while columns reflect whether a municipality is in the former East or West. Shape of coefficient reflects type of hate crime event. Effective N varies by outcome, hate crime type, and pre-treatment level of concern. All point estimates and errors based on robust estimation and inference method from *rdrobust* package. All models use MSE-optimal bandwidths that are permitted to differ in width across the cutoff, local linear regression, triangular kernels, and errors clustered at the regional unit of treatment. Respondents interviewed on the day of treatment are excluded, as are respondents for whom the treatment-designating hate crime occurs within 30 days of an earlier hate crime.

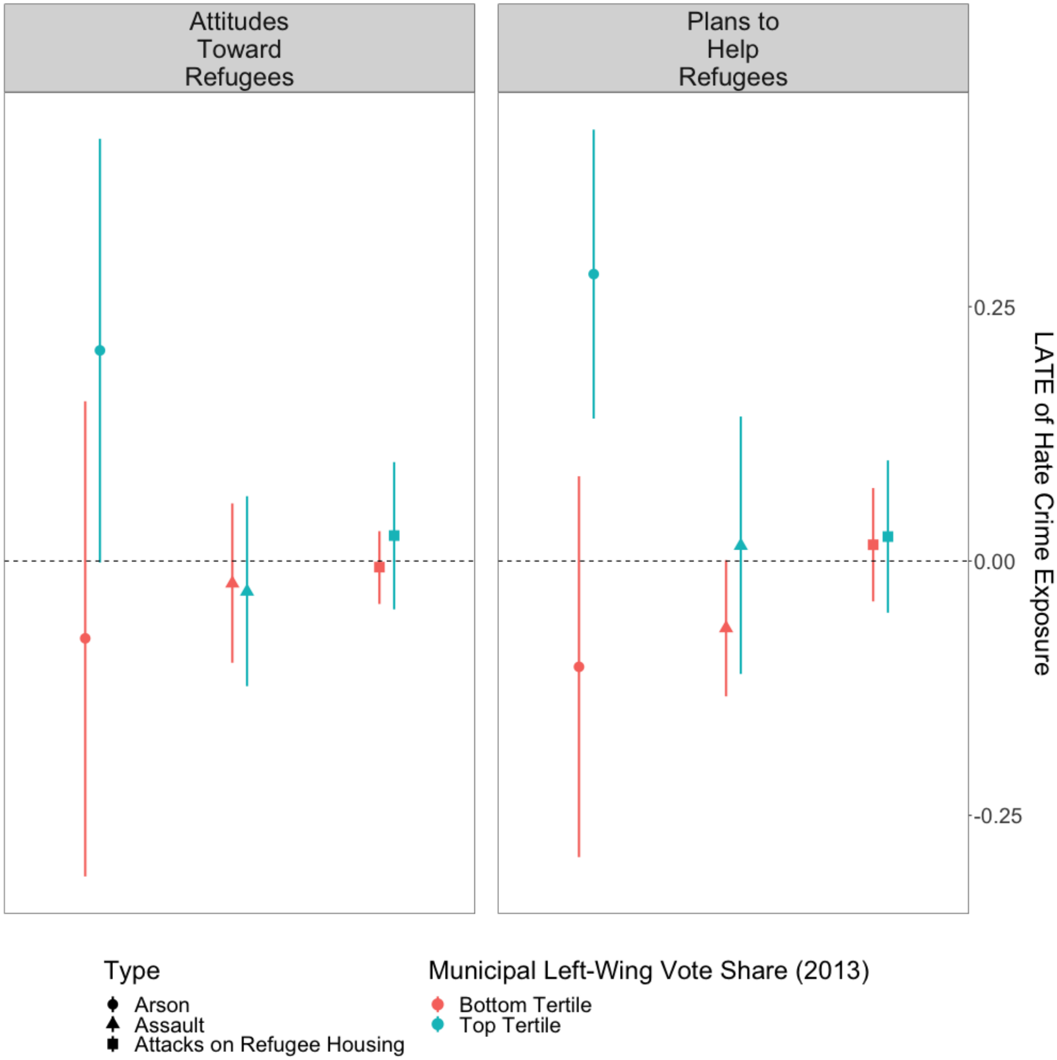

Figure A 27. Effect of Regional Hate Crime Exposure on Refugee Attitudes and Plans to Help, Across Municipal Left-leaning Party Vote Share

Coefficients with 95% confidence intervals. Coefficients represent the difference in intercept below and above the cutoff, i.e. the LATE of hate crime exposure. Rows reflect different outcomes while columns reflect the bottom and top tertiles of left-party vote share in 2013 (Green, SPD, and Linke). Shape of coefficient reflects type of hate crime event. Effective N varies by outcome, hate crime type, and pre-treatment level of concern. All point estimates and errors based on robust estimation and inference method from *rdrobust* package. All models use MSE-optimal bandwidths that are permitted to differ in width across the cutoff, local linear regression, triangular kernels, and errors clustered at the regional unit of treatment. Respondents interviewed on the day of treatment are excluded, as are respondents for whom the treatment-designating hate crime occurs within 30 days of an earlier hate crime.

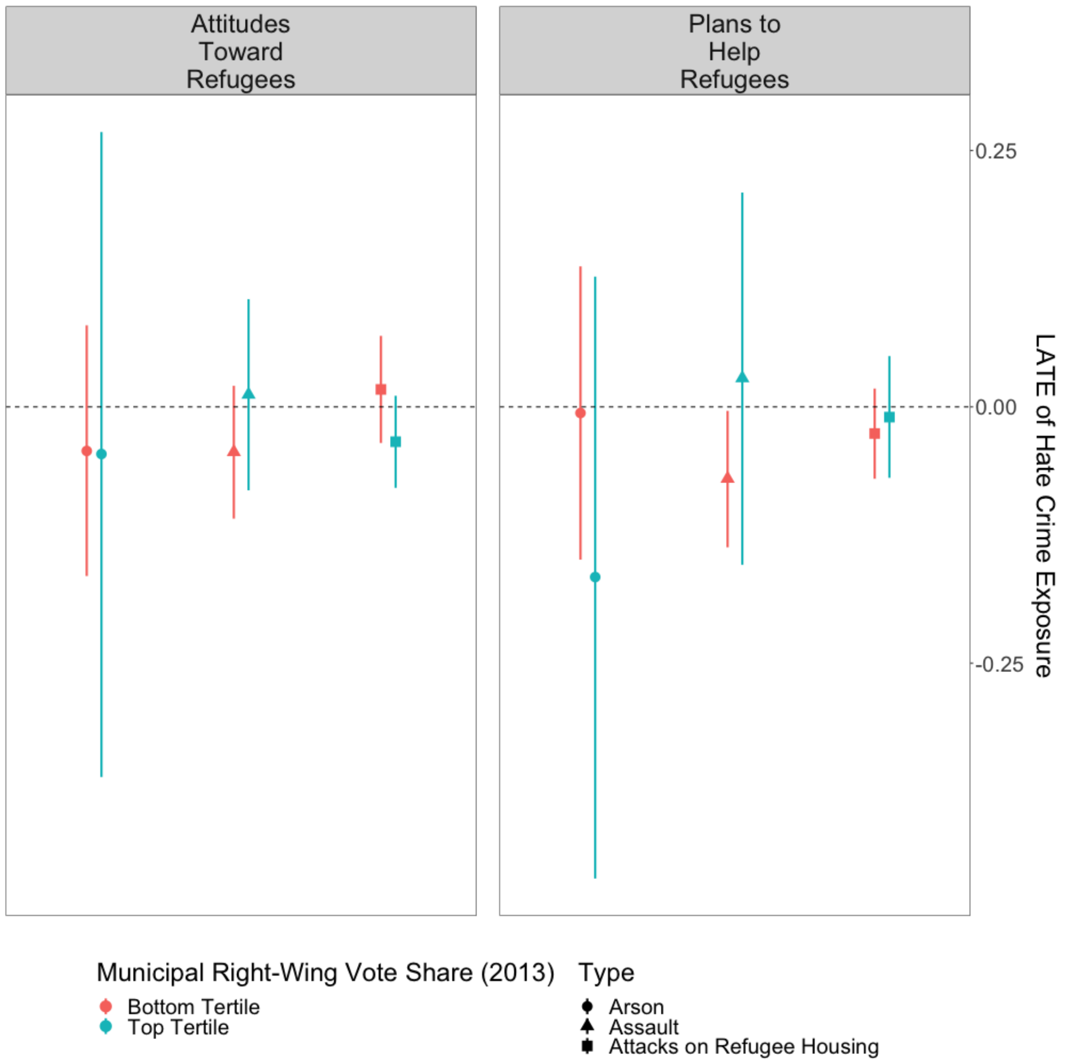

Figure A 28. Effect of Regional Hate Crime Exposure on Refugee Attitudes and Plans to Help, Across Municipal Right-Leaning Party Vote Share

Coefficients with 95% confidence intervals. Coefficients represent the difference in intercept below and above the cutoff, i.e. the LATE of hate crime exposure. Rows reflect different outcomes while columns reflect the bottom and top tertiles of left-party vote share in 2013 (FDP, CDU, AfD). Shape of coefficient reflects type of hate crime event. Effective N varies by outcome, hate crime type, and pre-treatment level of concern. All point estimates and errors based on robust estimation and inference method from *rdrobust* package. All models use MSE-optimal bandwidths that are permitted to differ in width across the cutoff, local linear regression, triangular kernels, and errors clustered at the regional unit of treatment. Respondents interviewed on the day of treatment are excluded, as are respondents for whom the treatment-designating hate crime occurs within 30 days of an earlier hate crime.

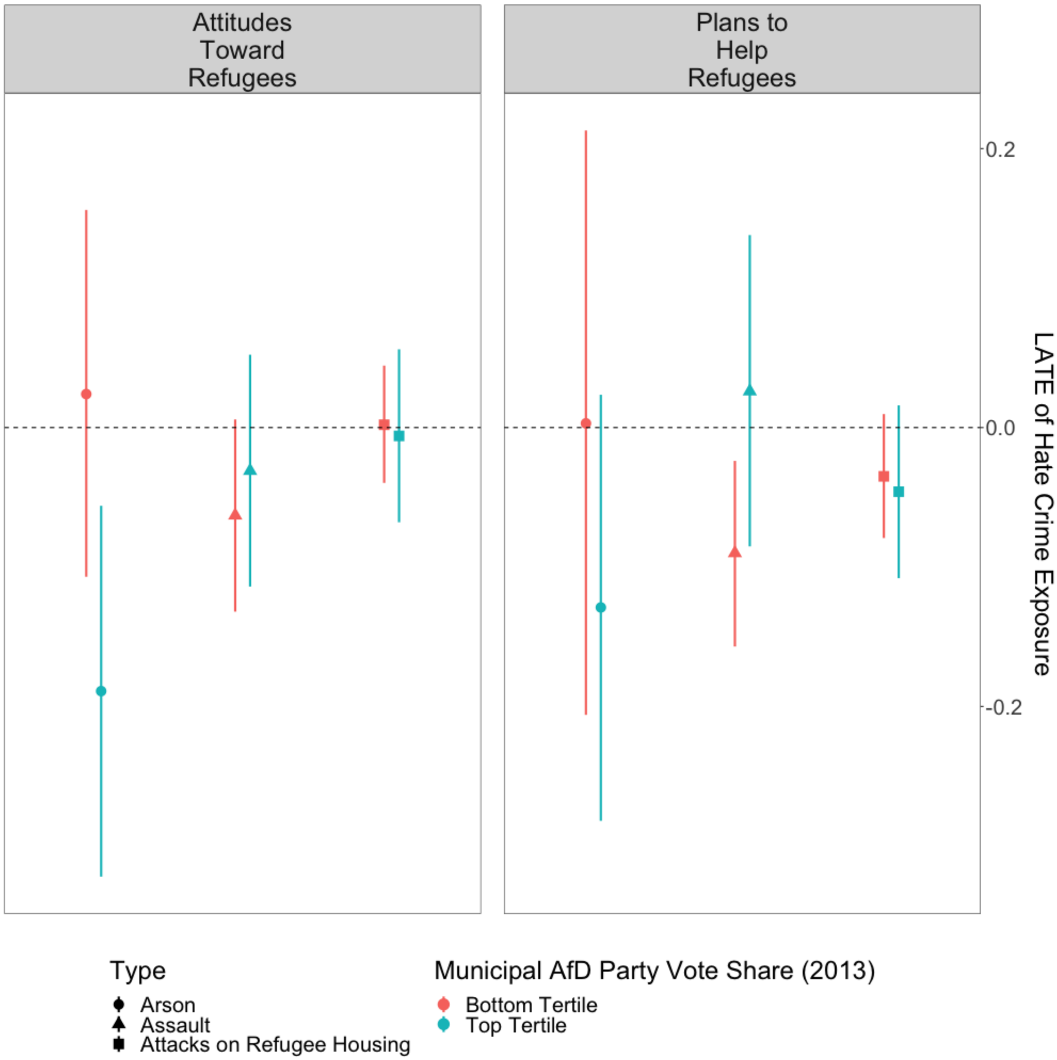

Figure A 29. Effect of Regional Hate Crime Exposure on Refugee Attitudes and Plans to Help, Across Municipal AfD Vote Share

Coefficients with 95% confidence intervals. Coefficients represent the difference in intercept below and above the cutoff, i.e. the LATE of hate crime exposure. Rows reflect different outcomes while columns reflect the bottom and top tertiles of AfD vote share in 2013. Shape of coefficient reflects type of hate crime event. Effective N varies by outcome, hate crime type, and pre-treatment level of concern. All point estimates and errors based on robust estimation and inference method from *rdrobust* package. All models use MSE-optimal bandwidths that are permitted to differ in width across the cutoff, local linear regression, triangular kernels, and errors clustered at the regional unit of treatment. Respondents interviewed on the day of treatment are excluded, as are respondents for whom the treatment-designating hate crime occurs within 30 days of an earlier hate crime.

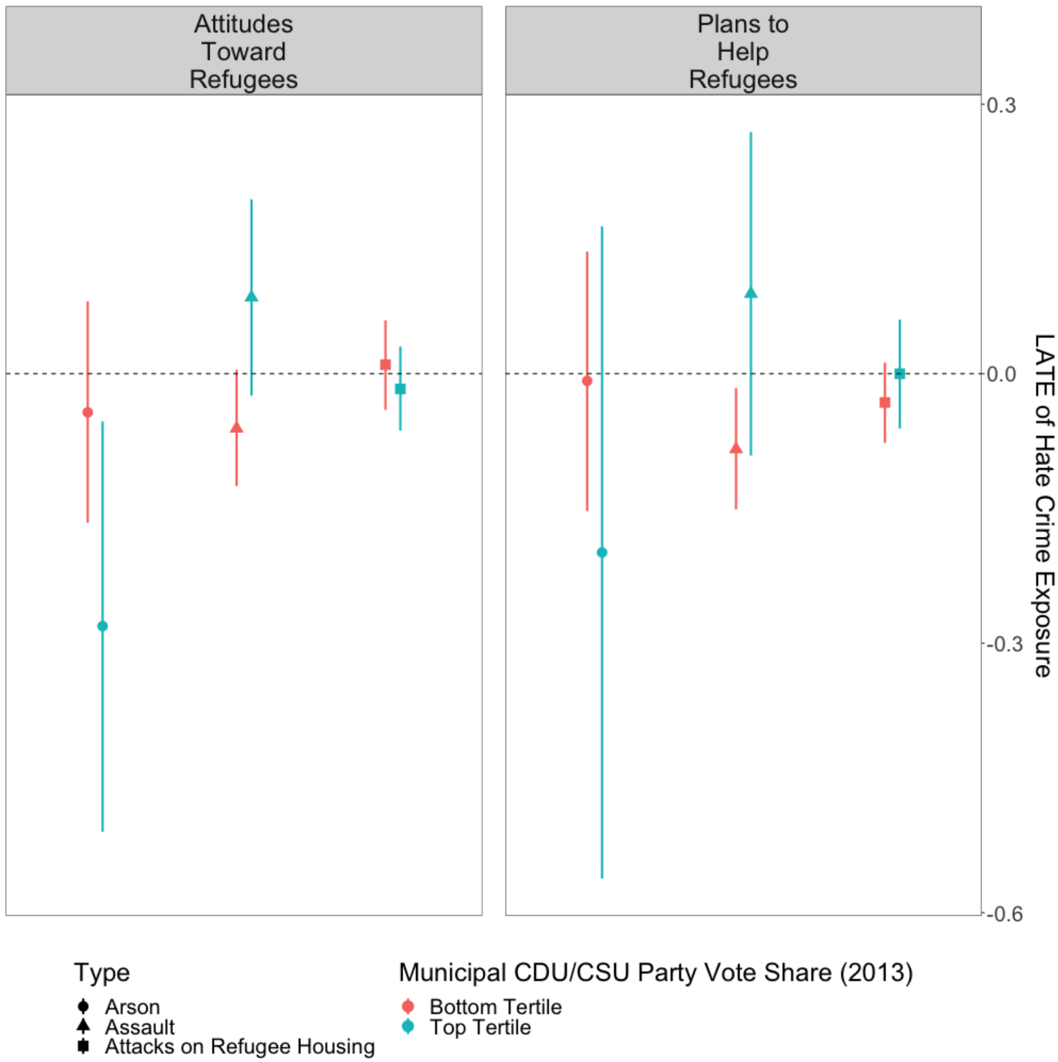

Figure A 30. Effect of Regional Hate Crime Exposure on Refugee Attitudes and Plans to Help, Across Municipal CDU/CSU Vote Share

Coefficients with 95% confidence intervals. Coefficients represent the difference in intercept below and above the cutoff, i.e. the LATE of hate crime exposure. Rows reflect different outcomes while columns reflect the bottom and top tertiles of CDU/CSU vote share in 2013. Shape of coefficient reflects type of hate crime event. Effective N varies by outcome, hate crime type, and pre-treatment level of concern. All point estimates and errors based on robust estimation and inference method from *rdrobust* package. All models use MSE-optimal bandwidths that are permitted to differ in width across the cutoff, local linear regression, triangular kernels, and errors clustered at the regional unit of treatment. Respondents interviewed on the day of treatment are excluded, as are respondents for whom the treatment-designating hate crime occurs within 30 days of an earlier hate crime.

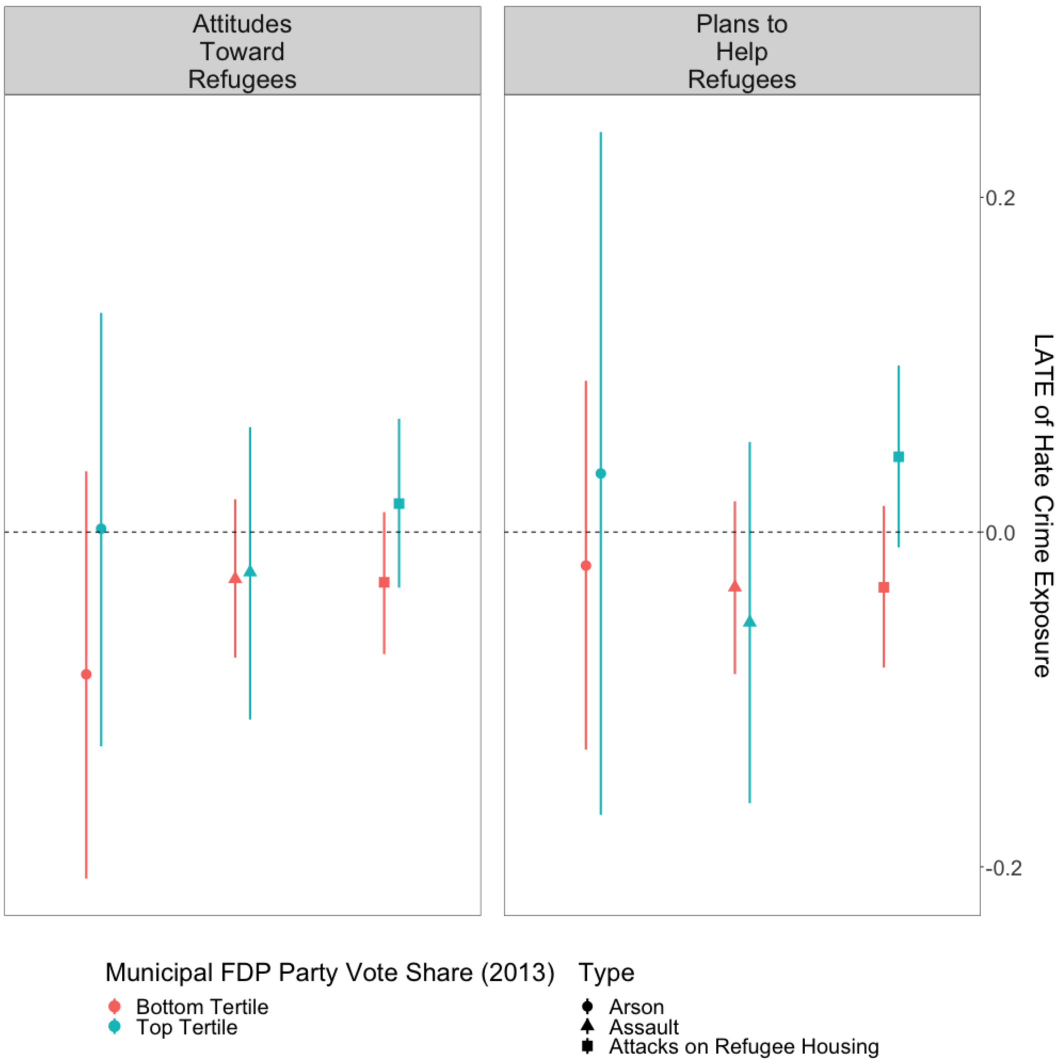

Figure A 31. Effect of Regional Hate Crime Exposure on Refugee Attitudes and Plans to Help, Across Municipal FDP Vote Share

Coefficients with 95% confidence intervals. Coefficients represent the difference in intercept below and above the cutoff, i.e. the LATE of hate crime exposure. Rows reflect different outcomes while columns reflect the bottom and top tertiles of FDP vote share in 2013. Shape of coefficient reflects type of hate crime event. Effective N varies by outcome, hate crime type, and pre-treatment level of concern. All point estimates and errors based on robust estimation and inference method from *rdrobust* package. All models use MSE-optimal bandwidths that are permitted to differ in width across the cutoff, local linear regression, triangular kernels, and errors clustered at the regional unit of treatment. Respondents interviewed on the day of treatment are excluded, as are respondents for whom the treatment-designating hate crime occurs within 30 days of an earlier hate crime.

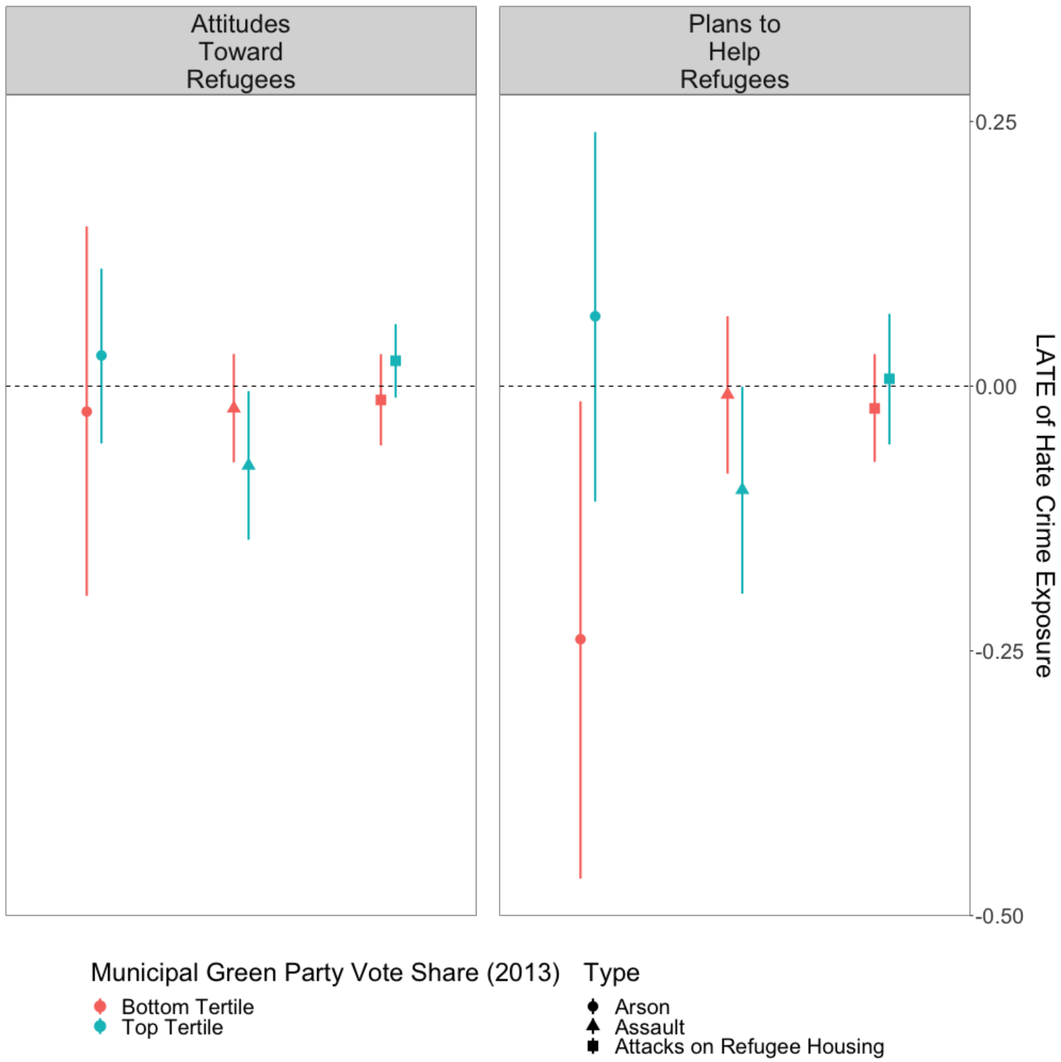

Figure A 32. Effect of Regional Hate Crime Exposure on Refugee Attitudes and Plans to Help, Across Municipal Green Party Vote Share

Coefficients with 95% confidence intervals. Coefficients represent the difference in intercept below and above the cutoff, i.e. the LATE of hate crime exposure. Rows reflect different outcomes while columns reflect the bottom and top tertiles of Green Party vote share in 2013. Shape of coefficient reflects type of hate crime event. Effective N varies by outcome, hate crime type, and pre-treatment level of concern. All point estimates and errors based on robust estimation and inference method from *rdrobust* package. All models use MSE-optimal bandwidths that are permitted to differ in width across the cutoff, local linear regression, triangular kernels, and errors clustered at the regional unit of treatment. Respondents interviewed on the day of treatment are excluded, as are respondents for whom the treatment-designating hate crime occurs within 30 days of an earlier hate crime.

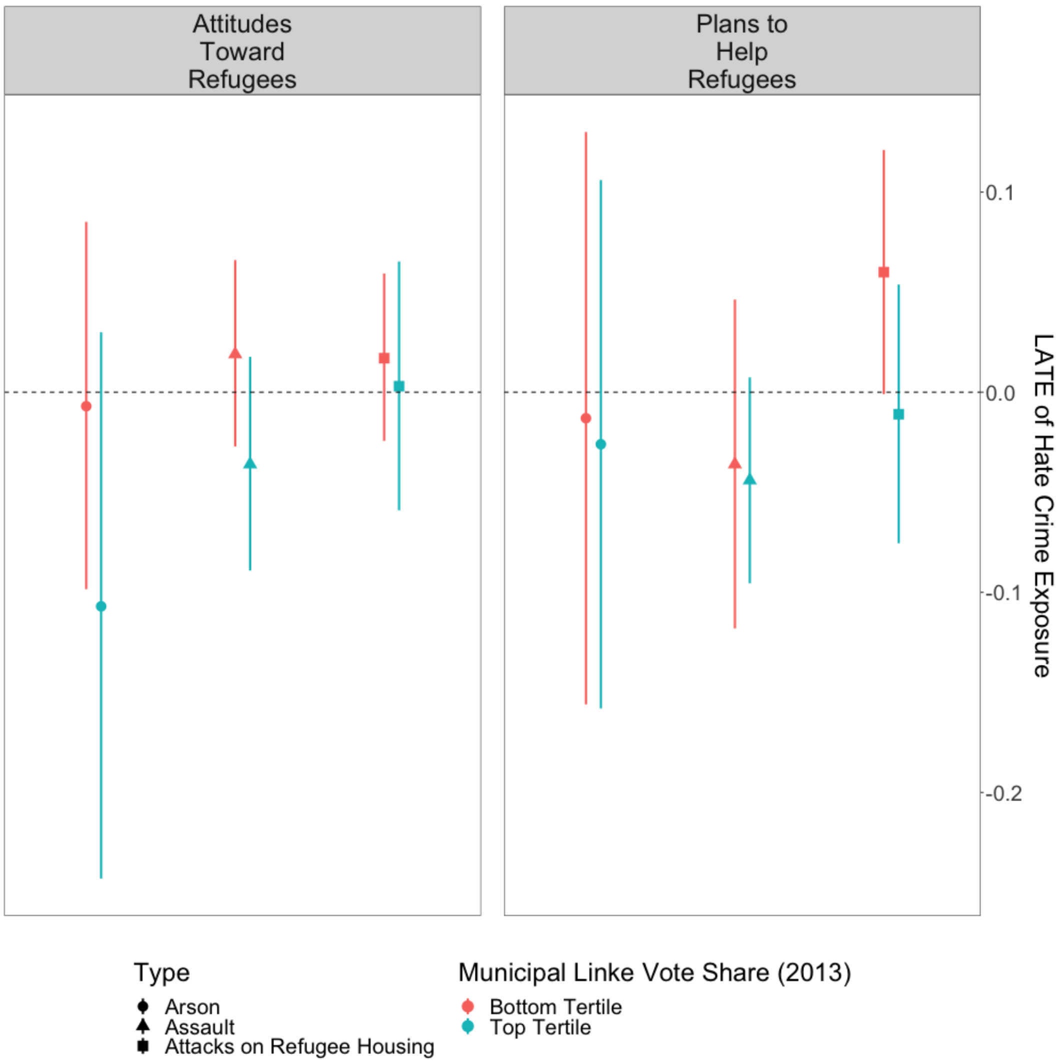

Figure A 33. Effect of Regional Hate Crime Exposure on Refugee Attitudes and Plans to Help, Across Municipal Linke Vote Share

Coefficients with 95% confidence intervals. Coefficients represent the difference in intercept below and above the cutoff, i.e. the LATE of hate crime exposure. Rows reflect different outcomes while columns reflect the bottom and top tertiles of Linke vote share in 2013. Shape of coefficient reflects type of hate crime event. Effective N varies by outcome, hate crime type, and pre-treatment level of concern. All point estimates and errors based on robust estimation and inference method from *rdrobust* package. All models use MSE-optimal bandwidths that are permitted to differ in width across the cutoff, local linear regression, triangular kernels, and errors clustered at the regional unit of treatment. Respondents interviewed on the day of treatment are excluded, as are respondents for whom the treatment-designating hate crime occurs within 30 days of an earlier hate crime.

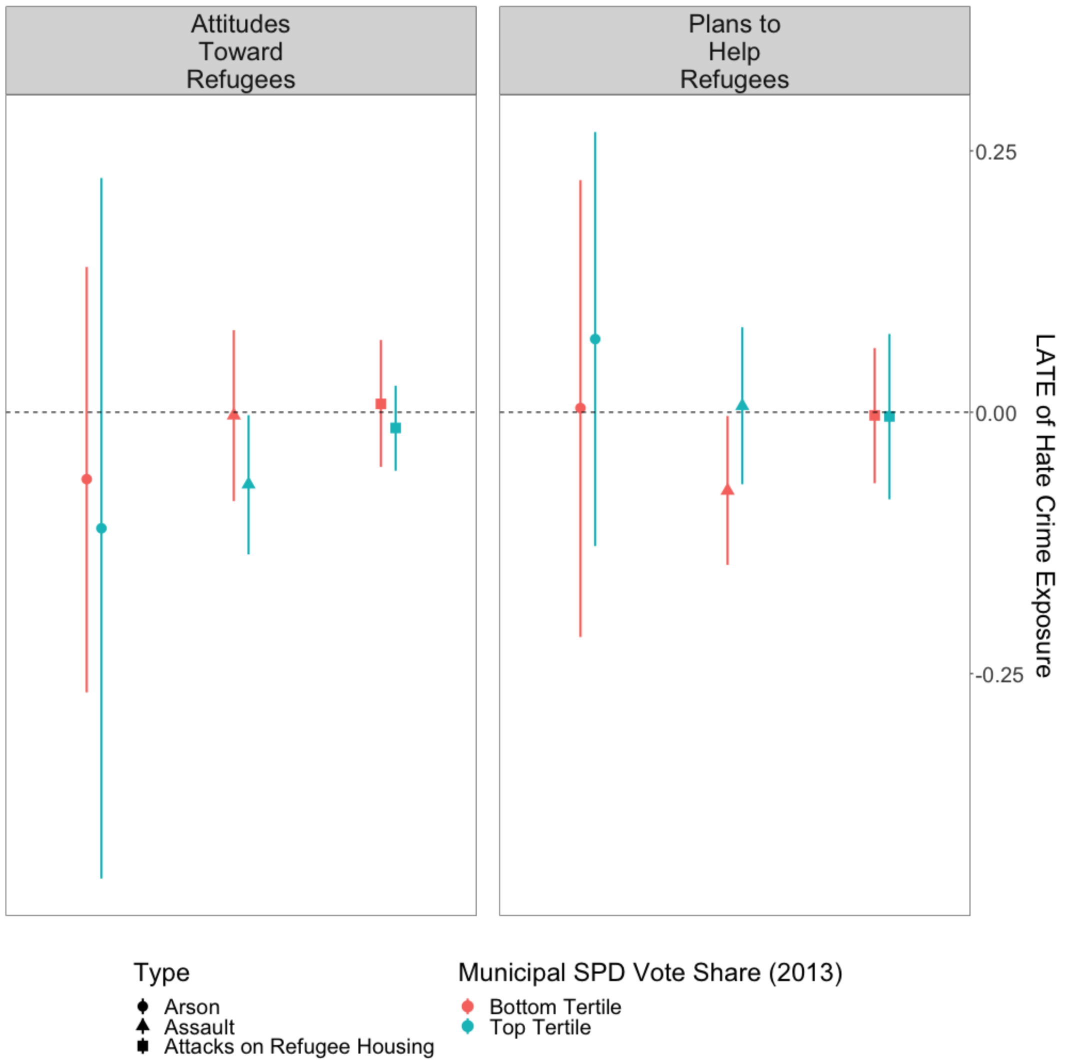

Figure A 34. Effect of Regional Hate Crime Exposure on Refugee Attitudes and Plans to Help, Across Municipal SPD Vote Share

Coefficients with 95% confidence intervals. Coefficients represent the difference in intercept below and above the cutoff, i.e. the LATE of hate crime exposure. Rows reflect different outcomes while columns reflect the bottom and top tertiles of SPD vote share in 2013. Shape of coefficient reflects type of hate crime event. Effective N varies by outcome, hate crime type, and pre-treatment level of concern. All point estimates and errors based on robust estimation and inference method from *rdrobust* package. All models use MSE-optimal bandwidths that are permitted to differ in width across the cutoff, local linear regression, triangular kernels, and errors clustered at the regional unit of treatment. Respondents interviewed on the day of treatment are excluded, as are respondents for whom the treatment-designating hate crime occurs within 30 days of an earlier hate crime.

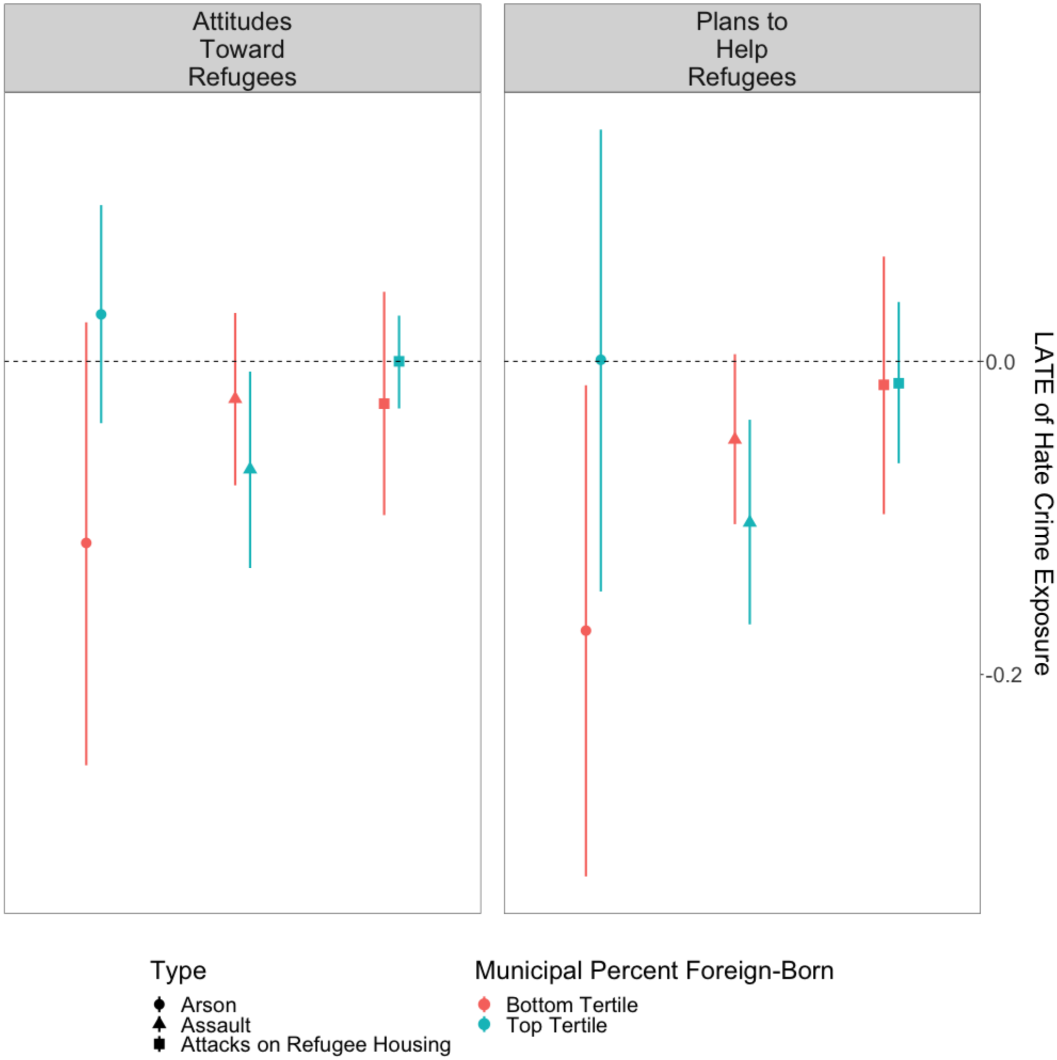

Figure A 35. Effect of Regional Hate Crime Exposure on Refugee Attitudes and Plans to Help, Across Municipal Foreign-Born Population Share

Coefficients with 95% confidence intervals. Coefficients represent the difference in intercept below and above the cutoff, i.e. the LATE of hate crime exposure. Rows reflect different outcomes while columns reflect the bottom and top tertiles of municipal foreign-born population share. Shape of coefficient reflects type of hate crime event. Effective N varies by outcome, hate crime type, and pre-treatment level of concern. All point estimates and errors based on robust estimation and inference method from *rdrobust* package. All models use MSE-optimal bandwidths that are permitted to differ in width across the cutoff, local linear regression, triangular kernels, and errors clustered at the regional unit of treatment. Respondents interviewed on the day of treatment are excluded, as are respondents for whom the treatment-designating hate crime occurs within 30 days of an earlier hate crime.

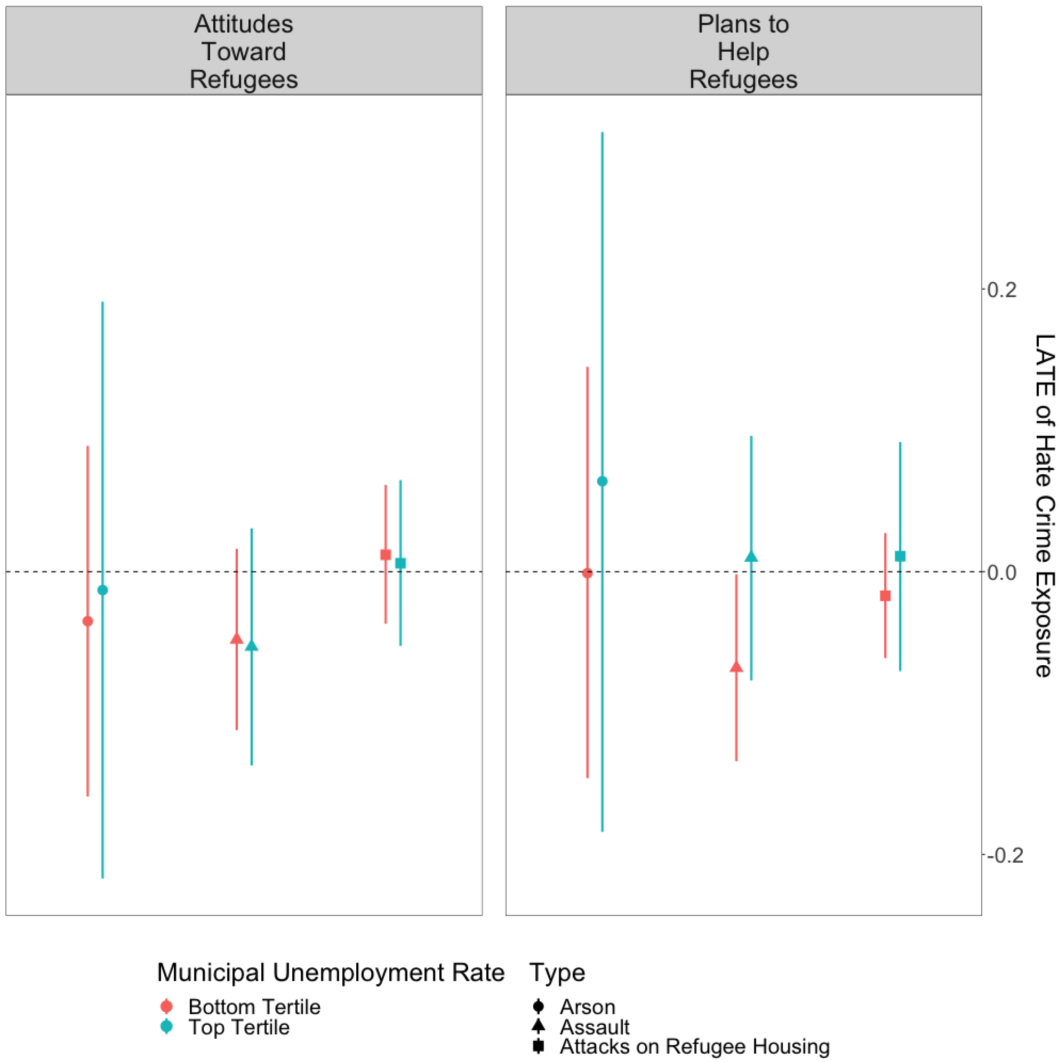

Figure A 36. Effect of Regional Hate Crime Exposure on Refugee Attitudes and Plans to Help, Across Municipal Unemployment Rate

Coefficients with 95% confidence intervals. Coefficients represent the difference in intercept below and above the cutoff, i.e. the LATE of hate crime exposure. Rows reflect different outcomes while columns reflect the bottom and top tertiles of municipal unemployment rate. Shape of coefficient reflects type of hate crime event. Effective N varies by outcome, hate crime type, and pre-treatment level of concern. All point estimates and errors based on robust estimation and inference method from *rdrobust* package. All models use MSE-optimal bandwidths that are permitted to differ in width across the cutoff, local linear regression, triangular kernels, and errors clustered at the regional unit of treatment. Respondents interviewed on the day of treatment are excluded, as are respondents for whom the treatment-designating hate crime occurs within 30 days of an earlier hate crime.

Appendix 3.5 Coverage of Hate Crime Events in German-Language Newspapers

To address whether and how information about hate crime events is disseminated in the news media, we explore German-language newspaper coverage of hate crime events. This analysis has two goals. First, we establish the baseline rate of newspaper coverage across event types, so as to establish whether individuals are sample are feasibly treated. Second, we show what types of newspapers cover events, with the goal of establishing whether treatment is indeed local (i.e. if individuals living in area that experience an event are more likely to know about said event than a German resident living in a different German region).

First, we drew random samples of events for each of the three hate crime types in ARVIG. Specifically, we sampled N=250 events for arson, assault, and attack on refugee housing, respectively. These N=750 total events amount to 12.8% of all events in ARVIG. However, given baseline differences in frequency across event type, our random sample includes 85.6% of arson events, 26% of assault events, and 5.4% of attacks on refugee housing. We excluded Berlin-based events from this sampling process for practical reasons.[12] As a result, we must caveat the findings of this analysis to highlight that news coverage of hate crime events in the rest of Germany may not reflect dynamics in Berlin.

Second, we linked individual events to German-language newspaper coverage using Factiva. Specifically, we conducted N=750 searches using the following search criteria:

- Search constrained to week after an event (from t+1, i.e. day after an event, to t+8).
- We used only one keyword: the name of the municipality in which an event occurred.
- German-language news sources only
- Search results constrained to Germany as a location
- Similar duplicated results omitted
- We specified the following "subjects" as search criteria:
  - Discrimination
  - Human Migration
  - Racism/Xenophobia
  - Hate Crime

12. German-language newspaper articles containing the word "Berlin" were so numerous that parsing through these articles to identify specific events proved unworkable.

- Civil Unrest
- Homicide
- Ethnic Minorities
- Arson (not used for assaults)
- Vandalism/Trespass (not used for assaults)
- Assault (not used for arson)

These search criteria produced highly variable search results. While many searches found zero articles matching these criteria, searches for some events found hundreds of articles. All articles found using this search method were downloaded as packets corresponding to a specific event.

Next, a research assistant was tasked with hand-coding these news packets, with the goal of ascertaining whether an event in question is covered in any of the downloaded articles. Specifically, the research assistant was tasked with identifying whether the event is mentioned in any of the newspaper articles, their level of confidence that the mentioned event corresponds to the event in ARVIG, the number of articles in which the event is mentioned (up to 5), and the specific newspaper outlets in which an event is mentioned (first 5 mentions, going in chronological order from the day after an event occurred).[13]

We subsequently classified news sources that covered at least one hate crime event in our data. Specifically, we classified sources according to whether they are local/regional, national/news agencies, or international (non-German).Given ambiguity in the distinction between regional and local newspapers, we combine these two categories in subsequent analyses. For regional/local newspapers, we also establish a coverage zone for the paper, which includes the federal state in which the newspaper is headquartered as well as all directly-neighboring federal states. Newspaper classifications are presented in Table A7.

Outlet	Type	Also Includes	Regional To
B.Z.	Local	B.Z. am Sonntag	Berlin
Berliner Morgenpost	Local	Berliner Morgenpost Online	Berlin
Göttinger Tageblatt	Local		Niedersachsen
Hamburger Abendblatt	Local	Hamburger Abendblatt Online	Hamburg
Kölner Stadtanzeiger	Local		Nordrhein-Westfalen

13. 98.7% of event mentions were recorded with high confidence.

Leipziger Volkszeitung	Local	Leipziger Volkszeitung Online	Sachsen
Neue Presse	Local	Neue Presse Online	Niedersachsen
Südthüringer Zeitung	Local		Thüringen
Aller-Zeitung	Local		Niedersachsen
Aachener Zeitung	Local		Nordrhein-Westfalen
Abendzeitung München	Local		Bayern
Peiner Allgemeine Zeitung	Local		Niedersachsen
Reutlinger Nachrichten	Local		Baden-Württemberg
Hamburger Morgenpost	Local		Hamburg
Kieler Nachrichten	Local		Schleswig-Holstein
Schaumburger Nachrichten Online	Local		Niedersachsen
Wolfsburger Allgemeine	Local		Niedersachsen
Dresdner Neueste Nachrichten	Regional	Dresdner Neueste Nachrichten Online	Sachsen
General Anzeiger	Regional		Nordrhein-Westfalen
Hannoversche Allgemeine Zeitung	Regional	Hannoversche Allgemeine Zeitung Online	Niedersachsen
Lübecker Nachrichten Online	Regional	Lübecker Nachrichten Online	Schleswig-Holstein
Märkische Allgemeine	Regional	Märkische Allgemeine Zeitung Online	Brandenburg
Ostthüringer Zeitung	Regional		Thüringen
Ostsee-Zeitung	Regional	Ostsee-Zeitung Online	Mecklenburg-Vorpommern
Sächsische Zeitung	Regional		Sachsen
Thüringer Allgemeine	Regional		Thüringen
Weser Kurier	Regional		Bremen
Berliner Kurier	Regional		Berlin
Express	Regional		Nordrhein-Westfalen
Frankenpost	Regional		Bayern
Kölnische Rundschau	Regional		Nordrhein-Westfalen
Kurier am Sonntag	Regional		Bremen
Mitteldeutsche Zeitung	Regional		Sachsen-Anhalt
Neue Westfälische	Regional		Nordrhein-Westfalen
Nürnberger Nachrichten	Regional	Nordbayerische Nachrichten	Bayern
Nordwest-Zeitung	Regional		Niedersachsen
Passauer Neue Presse	Regional		Bayern
Rheinische Post	Regional		Nordrhein-Westfalen
Stuttgarter Nachrichten	Regional		Baden-Württemberg
Stuttgarter Zeitung	Regional		Baden-Württemberg
Thüringische Landeszeitung	Regional		Thüringen
Trierischer Volksfreund	Regional		Rheinland-Pfalz
Allgemeine Zeitung Mainz	Regional		Rheinland-Pfalz

Bayerische Staatszeitung	Regional		Bayern
Bergedorfer Zeitung	Regional		Hanburg
Gieflener Anzeiger	Regional		Hesse
Meininger Tagblatt	Regional		Thüringen
Metzinger Uracher Volksblatt	Regional		Baden-Württemberg
Südwest Presse	Regional		Baden-Württemberg
BILD	National	BILD am Sonntag; Bild.de	
Der Tagesspiegel	National	Der Tagesspiegel Online	
Die Welt	National		
Handelsblatt	National	Handelsblatt Online	
Süddeutsche Zeitung	National	Süddeutsche Zeitung Online	
WELT	National	WELT kompakt; WELT online	
Berliner Zeitung	National		
Focus	National	Focus Online	
ZEIT online	National		
Spiegel	National	Spiegel Online	
taz	National		
Bremer Nachrichten	National		
Der Spiegel	National		
dpa	News Agency	dpa-AFX ProFeed; dpa-InfoLine	
epd Basisdienst	News Agency		
Deutsche Welle	News Agency		
Agence France Presse	International		
Die Südostschweiz	International		
Neue Zürcher Zeitung	International		
Sputnik German News Service	International		
Aargauer Zeitung	International		
Berner Zeitung	International		
Bündner Tagblatt	International		
Der Bund	International		
Der Standard	International		
Die Presse	International		
Reuters	International		
Salzburger Nachrichten	International		
Zofinger Tagblatt	International		
manager magazin Online	Unclear		
news aktuell OTS	Unclear		

Table A 7. Newspapers

Next, we calculate the rate of coverage across different types of events and newspapers. These results are presented in Table A8. The second column highlights total coverage across event types, which range from 14.4% for attacks on refugee housing to 60.4% for arson events. The third column demonstrates the rate of event coverage in local and regional news source. In parentheses is the share of this local/regional news coverage that occurs in newspapers that cover the region in which an event occurred. For all three event types, most regional/local newspapers are indeed regional to the event in question. The penultimate and final columns show that national and international newspapers cover hate crime events less regularly than local/regional news, which lends credence to the argument that treatment is a local or regional phenomenon - rather than a national one.

Table A 8. Newspaper Coverage of Hate Crime Events

Type	% Coverage	% Local or Regional Coverage (% of which is Regional to the Event)	% National Coverage	% International Coverage
Arson	60.40	51.20 (94.62)	33.60	7.60
Assault	21.60	16.00 (85.71)	8.40	3.20
Housing	14.40	12.80 (96.88)	5.60	1.20

Finally, we explored the extent of coverage across types of municipalities, with a focus on level of urbanization. Given the potential for non-linearity and thresholds, we use a categorical measure of urbanization provided by the German statistical service. This measure allows us to explore mean rates of coverage for all three events in "sparsely populated", "medium population density", and "densely populated" municipalities. Results of this analysis are displayed in Table A9. Again, clear difference in coverage patterns across types of municipalities do not emerge from this analysis.

Table A 9. Newspaper Coverage of Hate Crime Events, by Level of Urbanization

	% News Coverage
Arson	
Sparsely Populated	63.16
Medium Population Density	60.80
Densely Populated	56.45
Assault	
Sparsely Populated	26.32
Medium Population Density	23.64
Densely Populated	17.02
Housing	
Sparsely Populated	6.38
Medium Population Density	18.58
Densely Populated	14.12

Appendix 3.6 Anxieties about Crime and Refugee Attitudes

To explore whether anxieties about crime shape refugee attitudes in the longer term, we leverage the SOEP's panel structure. As mentioned above, the SOEP has included a battery of refugee-facing questions on a biannual basis since 2016 (up to 2020). Each of the aforementioned survey years also includes the three questions used to construct our anxiety about crime measure.

To explore whether and how changes in anxiety about crime might shape refugee attitudes, we estimate a set of regressions for both refugee-attitude outcomes. First we run models with individual-level fixed effects models, with and without year and federal state fixed effects. These fixed effects models cover all native respondents in the 2016, 2018, and 2020 waves. Next, we replicate this analysis with a first difference model (with a two-year delta in both the dependent and independent variable, due to the biannual inclusion of the refugee attitudes survey items that constitute our main outcomes). This first difference models ensures that we capture the relationship of between-wave change in anxieties about crime and refugee attitudes. Again, we replicate this analysis with and without year and federal state fixed effects for both outcomes. The first difference models cover the 2018 and 2020 waves, with the 2016 wave operating at the baseline period for 2018.

Presented in Table A10, results are mixed. Whereas increased anxieties about crime are associated with more negative assessments of refugees in general, such increased anxiety is also associated with higher self-reported plans to help refugees in the future. These findings are significant at the 0.001 level across all model specifications.

Table A 10. Do changes in anxiety about crime predict refugee attitudes?

	Attitudes Towards Refugees				Plans to Help Refugees			
	M1	M2	M3	M4	M1	M2	M3	M4
Anxiety about Crime	-0.063***	-0.047***	-0.035***	-0.047***	0.071***	0.027***	0.035***	0.028***
	0.004	0.004	0.004	0.004	0.005	0.005	0.005	0.005
Individual Fixed Effects	✓	✓			✓	✓		
First Differencing			✓	✓			✓	✓
Year FE		✓		✓		✓		✓
Bundesland FE		✓		✓		✓		✓
N	73,318	73,318	36,737	36,737	73,025	73,025	36,110	36,110

***$p < 0.001$

Point estimates and standard errors clustered on the level of the individual respondent. Data: 2016, 2018, and 2020 waves of German Socio-Economic Panel.

Response Memo for *Political Science Research and Methods*

We thank the editor and reviewer for this opportunity to revise and resubmit our manuscript at *Political Science Research and Methods*. We agree that the issues raised during the last revision merit further attention, and have addressed them in the main text and memo to the best of our abilities. Below we respond to all six areas highlighted by the editor.

Mechanism

> **R2:** The negative effect in A.15 of assaults on planning to work with refugees, and the positive effect in A.16 of assaults on concern about cohesion in society, suggest that perhaps Germans' willingness to help refugees after an assault is attenuated because they fear for their own security. I'd like the authors to engage with this possible interpretation and tell us whether this could be consistent with their data.

> **Editor:** At several junctures in your analysis it appears that survey respondents react to assaults on refugees by indicating that they will be less likely to get involved to help refugees. As the reviewer suggests, this may be owing to security/safety concerns, which would be logical and yet not motivated by animus toward refugees. "I'm scared to get involved." If this is what is going on, you would expect to see that assaults lead to a decreased likelihood of getting directly involved to help refugees and to attend demonstrations, both of which may subject the individual to danger, but *not* a decreased likelihood of donating money or goods, if this can be done remotely or with minimal possible physical danger. I would be interested in you testing whether this is true as it has implications for how you interpret this finding.

Both the Editor and R2 suggest that exposure to anti-refugee hate crimes could affect plans to help refugee for egocentric reasons. Specifically, violence against refugees could make natives less willing to associate with refugees because contact with a population that is regularly targeted for violence may be seen as risky.

We think this point is well taken and acknowledge that aspects of our analysis corroborate this story. As suggested by the editor, under this scenario we would expect to see an association between exposure to assault events and greater anxiety about risks to native respondents (the "Anxiety about Cohesion in Society" and "Anxiety about Crime in Germany" survey items). Moreover, we would not expect to see an association between exposure to assault events and anxieties about risks to refugees (the "Anxiety about Hostility to Foreigners" survey item). Indeed, Figure A16 shows that coefficients for exposure to assaults are of a similar magnitude for the"Anxiety about Cohesion in Society" item ($\beta = 0.063$ at $p < 0.1$) and for the "Anxiety about Crime in Germany" item ($\beta = 0.098$ at $p < 0.05$). Though the coefficient for effects of assaults on "Anxiety about Hostility Towards Foreigners" item is also positive, it is 2-3x smaller in magnitude than the the coefficient for the other two items and not significant at conventional levels ($\beta = 0.037$ at $p = 0.15$).

Moreover, we do find evidence that assaults change planned behavior in ways that are consistent with anxieties about greater risks to themselves. As highlighted by R2, Figure A15 shows that the only the coefficient for exposure to assaults on "working directly with refugees" is negative and significant ($\beta = -0.077, p < 0.05$), unlike the coefficient for the "plans to donate money or goods" and "attend demonstrations or collect signatures" items. To summarize, this set of results is compatible with the aforementioned mechanism: natives appear more anxious about crime/disorder that may affect themselves, and not about crime/disorder that will affect the target population, i.e. refugees.

However, other aspects of our analysis complicate the picture. The longitudinal analysis suggests that generalized anxiety about crime is *positively* associated with plans to help refugees (and negatively associated with attitudes towards refugees). Rather than an egocentric mechanism, these findings suggest a sociotropic story in which violence against refugees induce greater concern about the societal implications of humanitarian migration, which in encourages natives to change their behavior in ways that will help refugees integrate.

In short, we agree that an egocentric mechanism is theoretically and empirically feasible, but also believe that our analysis does not bear out such a mechanism in an unambiguous way. That being said, we have fleshed out our mechanism subsection to highlight the aforementioned trends relating to the individual survey items. In the conclusion we now also highlight the need for further exploration of *how and why* local violence towards out-groups might shape behavior towards out-group members.

Interpretation of Results

> **R2:** I'm not entirely convinced that this effect of exposure to assault on willingness to help refugees is null. When we look at the many robustness tests the authors conduct in the Appendix, this effect keeps popping up in some specifications (Figure A.6, A.7, A.8, A.9, A.10, A.12, A.13).

> **R2** The authors claim that the "negative and significant effect of exposure to assaults on plans to help refugees attenuates to zero and is no longer statistically significant" (p. 9). I think this is a critical part of the paper, since it supports the main take-away that exposure has null effects. Therefore, the authors needs to be very clear about where we can see this effect. They reference Appendix section 3.4. However, I believe the evidence of this attenuation is actually in appendix section 3.3, and specifically table A.6. The authors should specify this in the main text.

> **Editor:** The holdout reviewer (and I) continue to have some reservations about the unequivocal way in which you state your findings. You might consider couching them differently, such as that you rarely find effects, that they are not consistent by type of hate crime, that some of the effects appear to be driven by personal safety, that some are mediated by generalized fear of crime, that they may be driven by a failure of the media to cover the events or to identify the victims, and so on. I know that this provides a more complex story but I think it may be more consistent with your evidence.

A major theme in the request for revisions was our interpretation of the effects of assault events on plans to help refugees. Whereas R2 asked for a clearer interpretation of one specific analysis (the effect of assault events on plans to help refugees), the editor asked for more nuance in general, including about the implications of our results. We agree that the previous version of the manuscript did not provide sufficient clarity or nuance in interpreting results. The new version of our manuscript has worked to integrate and address these concerns.

First, R2 is correct in saying that Table A6 is at the crux of our null results interpretation. Our interpretation of why the attenuation in Table A6 is so significant was not sufficiently clear in the previous version of the manuscript. We have added language in the main text to clarify why we interpret our results for assault events as null results, and point unambiguously to the analysis in Table A6.

Second, we have worked to add substantial nuance to the interpretation of our findings in the main results, mechanism section, and conclusion. These additions focus on two main areas

1. We have highlighted the existence of some evidence for a mix of different mechanisms underpinning our findings. This language highlights the analysis of separate survey items, per our discussion above. We also mention directions for future research in this vein.
2. We pay more attention to the role of media, especially to how the absence of media coverage might generate informational biases and help explain our null results.

Municipalities

> **R2:** If I understand the design correctly, all analysis is done within-municipality, in other words the comparison is between individuals in municipality A who were interviewed just after a hate crime to individuals in municipality A who were interviewed just before a hate crime. If I am correct, this should be clearly stated in order to reassure the reader that any municipality-level differences (which abound per Table A.1) are neutralized.

> **Editor:** The reviewer asks whether you are making comparisons within municipalities. To my understanding you are not – you are comparing respondents who answered a survey just before an incident to those who responded just after an incident, where the two groups are pooled across municipalities. Is that correct? If so, are you able to make your comparisons within municipalities perhaps by employing municipal fixed effects? And can you cluster your standard errors at the municipal level rather than the state level?

The editor is correct in saying that our research design is not a within-municipality comparison. We are comparing individual respondents on the basis of whether they were interviewed just before or just after an incident. We make an assumption that the timing of interviews relative to the timing of hate crime events is as good as random, and therefore that the characteristics of municipalities that these individuals live in will be balanced across the treatment threshold. We test this assumption in the right panel of Figure A2 (and indirectly in Figure A3, where we compare the counties in which municipalities are nested). These tests of the continuity assumption validate our research design: the municipalities in which respondents reside are highly comparable across the treatment threshold. They do not exhibit statistically significant differences for any of the socio-demographic, economic, political, or crime variables that we were able to collect. In that sense, despite the fact that we do not compare individuals within municipalities, the treated and untreated respondents are living in highly comparable municipalities.

The addition of municipal fixed effects would help us control for unobserved municipal differences across the threshold by comparing individual respondents within municipalities only. However, we do not think an RDD with municipal fixed effects is an appropriate test for a number of reasons. First, *ex ante*, we do not have good reasons to expect the aforementioned bias given the quasi-random nature of our treatment assignment mechanism and the lack of evidence of imbalance in general.

Second, adding municipal fixed effects is problematic in the context of our study. As shown in Figure 1, many of our municipalities are included only once – especially for the comparatively rare arson and assault events (of which there were a few hundred events across 13,000 municipalities). In that sense, adding municipal fixed effects would change the composition (and decrease the size) of our sample by *de facto* dropping individuals nested within municipalities that experienced only one hate crime.

Third, in reviewing the literature on regression discontinuity designs, a number of studies have addressed inclusion of covariates, for purposes including efficiency and identification (e.g. Calonico et al. 2019; Frölich and Huber 2019). However, Cattaneo et al. (2023; p.164) caution against including fixed effects to account for potential imbalances across the threshold because this is unlikely to fix identification and may introduce new issues.

We do agree, however, that an analysis with standard errors clustered at the municipal level is an important test. This test is depicted in Figure A8. Estimates are somewhat less precise when standard errors are clustered on the municipal level (and county-level, which we include too) as opposed to when they are clustered at the state level as in the main analysis in Table 1. However, the interpretation of our results does not change substantively. For the effect of assault events on plans to help refugees, the lack of an effect that is significant at conventional levels is in line with other robustness tests. For the effect of assault events on anxieties about crime, the effect remains significant at the 0.1 level ($p = 0.064$).

Media Coverage of Events

> **Editor:** What can you tell us about the frequency with which the media identified the victims of the crimes in their stories? Are their not protections in place restraining the media from naming victims? Or do the reports just tend to identify the victims as "refugees?" If the media are not sufficiently clear in who these crimes are targeting we can't expect citizens to learn that hate crimes against refugees have taken place.

In the newspaper articles that we coded, victims of hate crimes are rarely if ever identified by name – likely due to legal protections for victims. Rather, victims are often identified on the basis of a victim's migration status (i.e. refugee, asylum seeker, migrant, etc), origin (national identity, country of citizenship, etc), or by the location of an event (at a refugee center, asylum accommodation, etc).

During data collection, we considered this challenge of linking newspaper articles to specific events from a measurement perspective. *Ex ante*, we were not sure whether we would be able to link specific events (of which we had only minimalist descriptions in the ARVIG data) to newspaper articles published in the week after an event occurred. As mentioned in our manuscript, we accounted for this potential uncertainty in our data collection. Our research assistant was instructed to indicate their level of confidence that an event in the newspaper corresponds with the event in the data. To make this assessment, the research assistant could use all information included in the relevant newspaper article. In 98.7% of cases, event mentions were recorded with high confidence. We take these results to mean that most events should be relatively easily identifiable as attacks on refugees to readers of newspapers.

Countervailing Effects

> **Editor:** Are you able to distinguish between simultaneous positive and negative effects that cancel each other out versus both effects being null? That is, across treated individuals some may be moved in a positive direction and others a negative direction. Your

heterogeneous treatment analyses begin to address this but can you go further?

If exposure to hate crimes has different effects for different groups of natives, these results may offset one another, leading to null results overall. In the previous version of our manuscript, we explore this possibility through a suite of tests that explore heterogeneity in effects of exposure across groups with notable differences in the baseline level of support for refugees . We compare across this baseline level of support because we have good reasons to expect exposure to anti-refugee violence will affect those with anti-immigrant views differently from those who are pro-immigrant at baseline (Krause and Matsunaga 2023; Eger and Olzak 2023). Our analyses compare effects across a wide variety of individual characteristic and municipal characteristics that split the German population along immigration attitudes (see Appendix 3.4 for details).

To extend this analysis, we conducted additional tests for effect heterogeneity across a set of municipal moderators. Concretely, we explored effects across municipal vote-shares for *specific* political parties. Our previous manuscript included only aggregations of vote-shares into left-leaning and right-leaning parties, and we compared effects across the bottom and top tertiles of support for both categories. Though baseline attitudes and behaviors towards refugees do differ for individuals living in areas with high/low right vote share and left vote share (see Figure A21), this aggregation scheme masks substantial variation in baseline refugee attitudes. As seen in Figure A22, municipal vote shares for the the AfD and Green Party – the two parties with likely the strongest (and most antithetical) immigration platforms – are significantly more predictive of support for refugees at baseline than the aggregated categories shown in Figure 21.

The effects of exposure to hate crime events across the bottom/top tertiles of each party's municipal vote share are presented in Figures A29 - A34. These analyses are largely in line with our previous findings. We do not find compelling evidence that effects differ by the electoral success of different parties at the municipal level. Concretely, the effects of exposure for all three event types do not differ across the high/low tertiles of municipal vote share for any of the major German parties.

In summary, we remain open to the idea that exposure to hate crimes may induce countervailing effects across individual or municipal characteristics, but our analyses have not detected such dynamics across the "most-likely" moderators, i.e. moderators that split our sample along differences correlated with baseline immigration attitudes.

Appendix Numbering

> **Editor:** In your revision please double check that all your internal references have updated numbering to figures and tables.

Figure labels have been fixed and double checked. Appendix figures and tables now consistently feature an "A" prefix, regardless of whether they are mentioned in the main text of the Appendix itself. We've also added some mention of specific appendix figures/tables in the main text, as per R2's suggestion.

Appendix 3.7 Memo Works Cited

- Cattaneo, Matias D., Luke Keele, and Rocio Titiunik. 2023. "Covariate Adjustment in Regression Discontinuity Designs." In Handbook of Matching and Weighting Adjustments for Causal Inference, edited by José R. Zubizarreta. CRC Press.
- Calonico, Sebastian, Matias D. Cattaneo, Max H. Farrell, and Rocío Titiunik. 2019. "Regression Discontinuity Designs Using Covariates." *The Review of Economics and Statistics* 101 (3): 442–51.
- Eger, Maureen A., and Susan Olzak. 2022. "The Polarizing Effect of Anti-Immigrant Violence on Radical Right Sympathies in Germany." *International Migration Review*.
- Frölich, Markus, and Martin Huber. 2019. "Including Covariates in the Regression Discontinuity Design." *Journal of Business & Economic Statistics* 37 (4): 736–48.
- Krause, Werner, and Miku Matsunaga. 2023. "Does Right-Wing Violence Affect Public Support for Radical Right Parties? Evidence from Germany." *Comparative Political Studies.*