Waking Up the Golden Dawn: Does Exposure to the Refugee Crisis Increase Support for Extreme-right Parties?

Elias Dinas
Department of Politics and International Relations, Oxford University, Oxford, OX1 4AJ, U.K.

Konstantinos Matakos
Department of Political Economy, King's College London, London, WC2R 2LS, U.K.

Dimitrios Xefteris
Department of Economics, University of Cyprus, Nicosia

Dominik Hangartner
Department of Government, London School of Economics and Political Science, London, WC2A 2AE, U.K.
e-mail: d.hangartner@lse.ac.uk (corresponding author)

ABSTRACT

Does exposure to the refugee crisis fuel support for extreme-right parties? Despite heated debates about the political repercussions of the refugee crisis in Europe, there exists very little—and sometimes conflicting—evidence with which to assess the impact of a large influx of refugees on natives' political attitudes and behavior. We provide causal evidence from a natural experiment in Greece, where some Aegean islands close to the Turkish border experienced sudden and drastic increases in the number of Syrian refugees while other islands slightly farther away—but with otherwise similar institutional and socioeconomic characteristics—did not. Placebo tests suggest that pre-crisis trends in vote shares for exposed and non-exposed islands were virtually identical. This allows us to obtain unbiased estimates of the electoral consequences of the refugee crisis. Our study shows that among islands that faced a massive inflow of refugees just before the September 2015 election, vote shares for Golden Dawn, the most extreme-right party in Europe, moderately increased by 2 percentage points (a 44 percent increase at the average). Our findings have implications for the theoretical understanding of the "flash potential" of immigration politics as well as for the management of refugee flows.

INTRODUCTION

Countries across the globe are struggling to cope with the most severe refugee crisis since the aftermath of World War II. Since the spring of 2015, more than 2 million (UNHCR 2017) new asylum claims were submitted in Europe alone, most of them from Syrian refugees who crossed into Europe via the Mediterranean Sea. In addition to the direct financial costs and economic consequences (Foged and Peri 2016) that transition and receiving countries bear, the crisis has resulted in severe political repercussions in many destination countries, including protests against asylum seekers and political violence directed at refugees (Board 2015). Perhaps most importantly, the ongoing refugee crisis may also have long-lasting electoral consequences by affecting natives' political preferences and their attitudes toward asylum seekers.

Does the influx of refugees increase support for extreme-right parties? Despite heated debates about asylum policies and the allocation of refugees across Europe (Bansak et al. 2017), there exists very little evidence regarding the impact of refugee arrivals on native voters' political preferences and behavior. While there is a sizeable literature that documents a positive relationship between labor migration and vote shares for anti-immigrant parties in receiving countries (Barone et al. 2016; Becker et al. 2016; Brunner and Kuhn 2014; Halla et al. 2015; Mendez and Cutillas 2014), it is unclear how these findings translate to the context of refugee migration. Since humanitarian concerns are particularly important in structuring native attitudes towards asylum seekers (Bansak et al. 2016), we would expect that host communities are ceteris paribus more welcoming to refugees than to labor migrants. Only recently have two seminal studies attempted to isolate the effects of refugee migration on electoral outcomes in Austria (Steinmayr 2016) and Denmark (Dustmann et al. 2016), with partially contradictory conclusions. This lack of knowledge is particularly problematic with respect to countries that have experienced a large influx of refugees within a very short period of time, such as Greece, which received more than 50% of all refugees crossing into Europe in 2015 (UNHCR 2016).

When trying to isolate the effect that exposure to refugees has on the political attitudes and behavior of the host population, the key problem for causal inference is that the number of arrivals is far from randomly assigned. Instead, the process through which refugees select between different potential host countries and communities depends on factors such as religious or cultural proximity, labor market conditions, and openness toward refugees (Neumayer 2004). In other words, refugees will, as much as possible, flee to communities that they can reach and where they are welcome. This selection process severely confounds most existing comparisons of the political responses of natives to the presence of asylum seekers and refugees. If we find that communities that experienced a large influx of refugees are less (or more) supportive of anti-immigrant parties than places that have not, we cannot conclude that these differences are caused by refugee arrivals, because the selection ensures that these two groups of communities are also different on many other confounding characteristics.

By focusing on Greece, one of the European countries most affected by the current refugee crisis, we overcome this issue by taking advantage of a natural experiment to identify the electoral consequences of refugee arrivals. As the refugee crisis unfolded in the waters of the Aegean in spring 2015, Greek islands were differentially exposed to the influx of refugees merely because they were closer to or farther away from the Turkish coast where boats with refugees departed. For example, the island of *Lesvos* (with a population of about 80,000) received more than 200,000 asylum seekers between May and September 2015. Yet the neighboring island of *Lemnos*, which belongs to the same geographical and administrative unit and features similar socioeconomic and political conditions, did not receive any refugees over this period, simply because it is a few miles farther away from the Turkish coast (see Figure 1, Panel A, for a map of the Aegean Sea).

Several factors facilitate our identification strategy. First, among the 95 habitable islands in the Aegean Sea, distance to the Turkish border caused significant variation in the number of refugees arriving in each of these islands. Second, refugees tended to be housed within specific areas on each island. This allows us to exploit both between- and within-island variation in the number of refugee arrivals. Third, Greece held an election on September 20, 2015, right after the first wave of refugee arrivals that the country encountered throughout the spring and summer of 2015. The previous election had taken place only eight months prior in January 2015, before significant numbers of refugees arrived (see Figure 1, Panel B). The SI Appendix Sections S2.1 and S2.2 provide more information about the elections, which took place in the context of the Greek financial crisis. Two more elections, both in 2012, also serve as possible counterfactuals of support for extreme-right parties in the absence of the crisis. Fourth, many of these islands belong to the same electoral and administrative district, which ensures that they are identical across a plethora of observable and unobservable characteristics, such as the candidates running for office, regional government, police, judiciary, and access to EU funds.

Our study makes four main contributions. First, we provide rare causal evidence of the effects of exposure to the refugee crisis on support for extreme-right parties. For a number of reasons, the Aegean refugee crisis is an ideal case to test for the "flash potential" (Sniderman et al. 2004) of immigration politics, when widespread concerns about asylum seekers (Bansak et al. 2016) coupled with a sudden influx of immigrants are hypothesized to lead to large-scale electoral mobilization (see also Hopkins 2010). Specifically, Greece received about 400,000 new asylum seekers between April 2015 and September 2015 (UNHCR 2016), which translates to more than 2.8 refugees per Greek resident in affected islands. These sudden inflows reached their first peak just before the September 2015 election, thereby ensuring that the refugee crisis would be a highly salient issue when island residents cast their vote in the midst of the concurrent financial and economic crisis. The results suggest that this sudden and massive increase in refugee arrivals indeed fueled support for the radical anti-immigrant and anti-asylum-seeker party *Golden Dawn* (GD), but only moderately so: vote shares for GD increased about 2 percentage points (a 44 percent increase at the average) because of the exposure to the refugee crisis. This increase in vote shares, however, is statistically significant and substantively meaningful. GD became the third-largest (and minor opposition) party in the fragmented Greek multi-party system, and was awarded certain constitutionally protected privileges, such as the right to appoint the second deputy Speaker of the house and important vice-chair positions in parliamentary committees. In a companion study, we fielded a targeted survey to examine the impact of refugee arrivals on natives' political attitudes and policy preferences (reference suppressed to protect author anonymity).

Second, by employing within-island variation in proximity to refugee hotspots and measuring the number of refugee arrivals per resident, we show that the intensity of exposure further catalyzes the electoral boost of the extreme-right. Third, our study fills a gap by examining the political repercussions of the refugee crisis in Greece, the European country arguably most affected by the current refugee crisis and home to GD, the most extreme-right and neo-fascist party currently represented in a European parliament (Heinö 2016).

Fourth, Greek islands served as temporary transit destinations during our study period, and most refugees continued their journeys to Athens within 48 hours of arrival (Capon 2015). This very unusual feature of our case enables us to make two important contributions to our theoretical understanding of the drivers of anti-refugee attitudes. The refugees' very temporary presence on the island eliminated most avenues for sustained interactions between locals and refugees, a prerequisite for the contact theory to work (Allport 1979). Hence, this setting allows us to isolate the impact of exposure to (rather than contact with) refugee populations on political behavior—something that previous studies could not achieve. In addition, our findings are difficult to reconcile with realistic group conflict theory (Campbell 1965), in which competition for scarce resources is a necessary condition for conflict between the outgroup and the dominant group. Our study shows that mere exposure is sufficient to fuel prejudice and change political behavior.

THE UNFOLDING OF THE REFUGEE CRISIS IN THE AEGEAN SEA

Setting

Our study focuses on Greece, which was the main entry points to Europe, due to its proximity to the Turkish border, a long coastline that marks the EU external borders, and many difficult-to-patrol islands. Out of the 1.3 million new asylum seekers who reached European territory in 2015, more than 850,000 of them did so by arriving in one of the Greek Aegean islands (UNHCR 2016). But while some islands were strongly affected by these sudden refugee inflows, many other islands did not experience any contact with refugees. Most refugees left the islands of first arrival within a very short period, typically less than 48 hours, to continue their journeys via the ports of *Piraeus* or *Thessaloniki* to central and northern Europe.

The SI Appendix Section S.2 provides more information about how the refugee crisis unfolded in Greece.

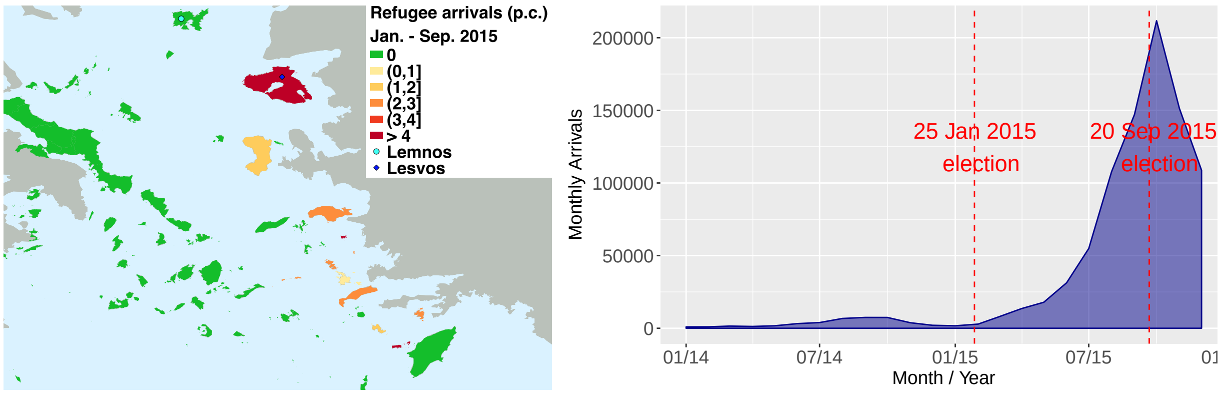

Figure 1: Panel A shows that islands close to the Turkish border received the most refugee arrivals per capita. Panel B shows the monthly number of asylum seekers arriving at Greek islands over the period from January 2014 to March 2016. During the study period, the first election took place just before the onset of the refugee crisis on January 25, 2015. A second election took place at the height of the refugee crisis on September 20, 2015.

Data

This study draws upon a new panel data set that covers all habitable Greek islands. The units of analysis are either municipalities ($N = 95$) or townships ($N = 248$). The data come from three different sources: electoral outcomes (at the municipality and township level) are provided for all four elections between 2012 and 2015 by the official record of the Greek Ministry of Interior and Public Administration; monthly data on refugee arrivals per island are provided by the UNHCR (UNHCR 2015), and geographic data on island's distance from the Turkish coast are obtained from the online mapping service Google Maps, which provides satellite imagery and geospatial data visualization and measurement. The SI Appendix Sections S3 and S4 provide more information about the data sources, descriptive statistics and coding decisions.

Empirical Strategies

We employ two complementary empirical strategies to identify the causal effect of exposure to the refugee crisis on electoral support for the GD. First, we apply a difference-in-differences (DID) estimator to obtain an estimate of the effect of refugee arrivals on the September 2015 vote share for GD. The underlying logic behind the design is that for the islands that received refugees in the summer of 2015, we construct the counterfactual change in GD vote share between January and September 2015 had they not received refugees, by using the change in GD vote share in the unaffected islands. This strategy yields causal estimates so long as the parallel trends assumption holds. In the current context, this assumption implies that the vote share for GD would follow the same trajectory from January to September 2015 among treated and non-treated islands in the absence of the refugee crisis.

To assess whether this assumption holds and thus to evaluate the validity of the design, we turn to past elections, exploring whether the over-time trend in GD vote share differs between treatment and control units. The left panel of Figure 2 uses the municipality-level analysis, whereas the right panel depicts the results from the township-level analysis. In both graphs two lines are shown, each representing the over-time change in the vote share for GD from the first election of 2012 until the last election of 2015. The red line denotes treated units, and the blue line denotes control units. We see that both groups travel in parallel until September 2015, the first election after the onset of the refugee crisis in Greece. Moreover, consistent with our expectation that the within-island variation helps to account more effectively for unobservable differences in support for GD, we see that the treated and control groups are similar not only in terms of changes, but also in levels before the refugee crisis. Taken together, the two graphs strengthen our confidence in the validity of the DID design.

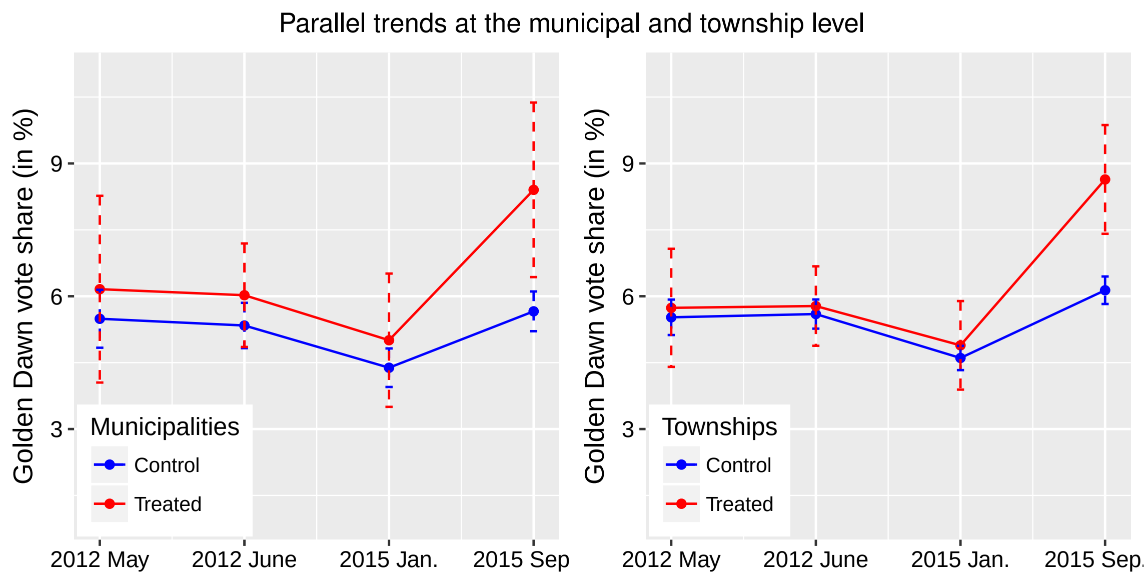

Figure 2: The analyses at the municipality (left panel) and township-level (right panel) show that treated and control islands experience highly similar changes in support for GD prior to the refugee crisis, thereby strengthening our confidence in the parallel trend assumption. The blue connected line indicates the average vote share for GD in the municipalities (left panel) and townships (right panel) that received refugees. The red connected line denotes the average GD vote share in municipalities and townships without refugee exposure.

The second empirical strategy employs an Instrumental Variables (IV) approach to identify the impact of the refugee crisis on the GD vote share. In particular, we use each island's distance to the Turkish coast as an instrument for the number of refugee arrivals. To test whether the first stage regression is sufficiently strong (Stock and Yogo 2005), we must determine whether the islands closer to the coast were indeed more likely to have received refugees than those farther away. Figure 3 shows that this the case. The left panel displays the probability of having received any refugees as a function of the distance from the coast. The blue curve represents a local regression smoother that is fitted to the data. We find a monotonically descending relationship, indicating that the farther away an island is from the coast, the less likely it is to have received any refugees during the current crisis. The right panel shows that the same monotonically decreasing relationship exists for the intensity of refugee exposure (measured as the relative number of refugee arrivals per island resident). Both graphs indicate that this relationship can be accurately modelled by a logarithmic function (red line). Accordingly, we use logged distance from the Turkish coast as an instrument of refugee exposure in the subsequent analysis.

The two identification strategies are complementary. The DID strategy is more efficient and identifies the average treatment effect for the treated islands, but we have to rely on the parallel trend assumption. In contrast, the IV strategy holds many potential confounders constant by design, but relies on the exclusion restriction (Angrist et al. 1996), namely the assumption that the only way proximity to the Turkish coast

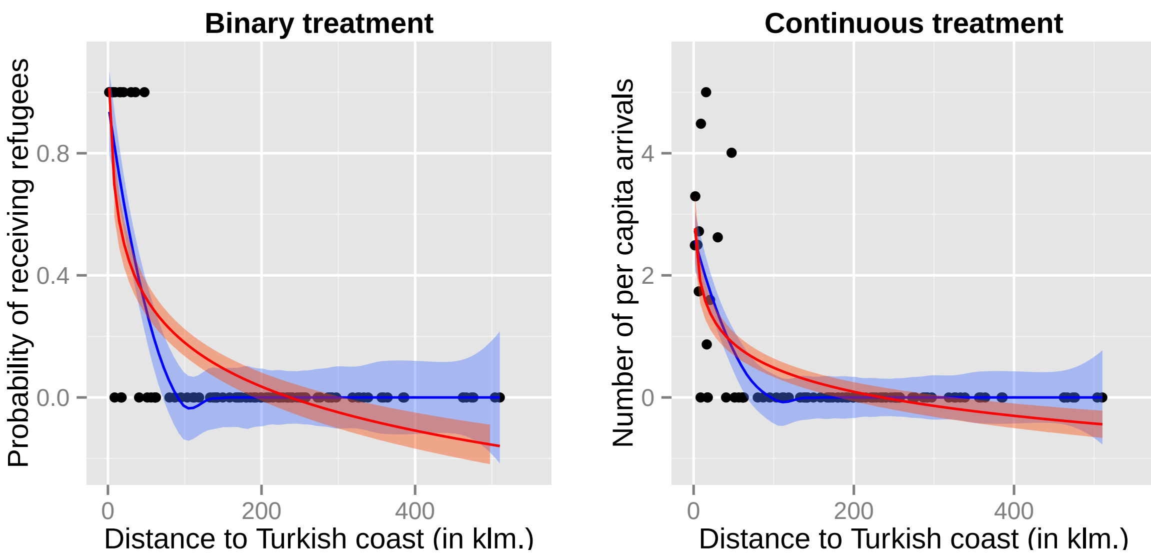

Figure 3: First stage: islands' distance from the Turkish coast predicts the number of refugee arrivals. The blue curves denote the local linear regression smoother (span=0.5), with the shaded area capturing the 95% confidence intervals. The left panel shows the propensity of receiving refugees against the distance to the Turkish coast, whereas the right panel shows the number of refugee arrivals per capita against the distance. The red curves display the predicted level of treatment exposure conditional on logged distance. In both graphs, the logarithmic function of distance from the Turkish coast provides a good approximation to the data.

affects support for GD is through refugee arrival. The SI Appendix Section S5 provides more information about the estimation strategy for the DID and IV analyses.

A placebo test confirms that control islands provide a credible counterfactual for treated islands. For this, we exploit a particular feature of the Greek election law: Registered nonresident voters cast their ballot in their town of residency, even when they do not live in the town they are registered in. We use this group to examine the behaviour of nonresidents who live on the Greek mainland but are registered to vote for islands with refugee exposure. These placebo tests, presented in the SI Appendix S6.3, reveal that electoral support for GD did not change substantially among i) nonresident voters of treated islands over time and ii) between nonresidents of treated and control islands.

RESULTS

Panel Data Estimates

Massive refugee arrivals from the Turkish coast to the Greek islands started in spring and escalated over the summer of 2015. Our panel data estimates exploit that Greece held two elections in the same year, one in January and one in September. Furthermore, refugee arrivals were highly clustered among a subset of the Aegean islands. We use the cross-island as well as within-island variation (Figure S2, Panel A) to assess the impact of refugee presence of vote share for GD (Figure S2, Panel B). Employing a DID estimator, we treat the September 2015 election as the post-treatment election and the three last elections preceding the January 2015 election as the pre-treatment elections. A detailed description of empirical strategy is provided in the SI Appendix Section S5.

Table 1 presents the main result and shows that refugee arrivals catalyze GD vote share.[1] Model 1 uses municipalities as the unit of analysis and shows that in islands that received refugees, support for GD rose by 2 percentage points (two-tailed $p < 0.001$). The average vote share for GD in 2015 among all municipalities in the estimation sample is 4.5 percentage points. Thus, the magnitude of the backlash effect amounts to a 44 percent increase in the party's vote share. Model 2 shows that the results are virtually identical when adding linear municipality-specific trends that capture smooth changes in unobserved confounders. Models 3 and 4 are the placebo analogues of Models 1 and 2, respectively. They present the same result but using GD's lagged vote share as the outcome. Reassuringly, we find no effect of exposure to the refugee crisis on GD's vote share in the *previous* election.

Models 5 and 6 use townships as the unit of analysis. Recall that townships are clustered within municipalities and thus allow for within-island variation. The results are very similar to the municipality-level analysis. We find that townships receiving refugees see an increase in the GD vote share by slightly more than 2 percentage points (two-tailed $p < 0.001$) compared to townships that were not exposed to refugee arrivals. Model 6 shows that the effect remains robust to the inclusion of township-specific linear trends. This result suggests that the increase in the party's vote share is not uniform within the affected islands, but rather more concentrated among those areas directly exposed to the refugee crisis. The placebo tests shown in Models 7 and 8 confirm that treated and control townships follow parallel trends prior to the January 2015 election. The difference in the vote share for GD between treated and untreated townships before that election is essentially zero, both with and without the inclusion of township-specific trends.

Finally, Models 9 to 12 replace the binary treatment indicator with the per capita number of refugees arriving in each island. Models 9 and 10 present the treatment effects and indicate that the arrival of one refugee per resident increases GD's vote share by 0.77 percentage points. To see more concretely what this means, take *Samos*, which received approximately 2.5 refugee arrivals per resident. As a result, the vote share for GD in September 2015 is expected to have risen by almost two percentage points because of the refugee arrivals. This estimate remains robust when adding municipality-specific linear trends. Models 11 and 12 display the placebo results, which use the lagged values of support for GD. Reassuringly, we find that exposure to refugee arrivals in 2015 plays no role in the vote share for GD in the pre-treatment elections. The average treatment effect is again very close to zero both with and without municipality-specific trends. Figure S2 in the SI Appendix provides a visual summary of those results.

Instrumental Variables Estimates

Distance to the Turkish coast plays a key role in whether and how many refugees arrived on a particular island in the Aegean in 2015. As discussed in the previous section, we exploit this natural experiment by using distance from the Turkish coast as an instrument of refugee exposure.

Figure 4 illustrates the logic of the estimation strategy. The left panel of Figure 4 shows the change in GD's vote share, from May 2012 to January 2015, as a function of the (logged) distance of each island from the Turkish coast. Since the Syrian refugee crisis impacted Europe only after the spring of 2015, we expect no relationship between distance from the coast and support for GD. This is exactly what we see in the right panel of Figure 4: The linear regression coefficient measuring the effect of distance to the Turkish coast on change in GD support between May 2012 and January 2015 is 0.048 (two-tailed,

[1] To address the intra-class correlation stemming from the panel structure of our data, we cluster standard errors at the municipality and township level, respectively. To address further concerns about inference based on DID estimates (Bertrand et al. 2004), we also adopted three alternative inference strategies: i) block-bootstrapped standard errors (Table S9 in SI Appendix); ii) collapsing all pre-treatment time points into one pre-treatment period (Table S10 in SI Appendix); and iii) placebo tests based on permutation tests (Figure S8 in SI Appendix). All these analyses are described in more detail in the SI Appendix.

Table 1: Impact of refugee arrivals on GD vote share

Model:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Outcome:	$GD_{(t)}$		$GD_{(t-1)}$		$GD_{(t)}$		$GD_{(t-1)}$		$GD_{(t)}$		$GD_{(t-1)}$	
Treatment:	Binary treatment		Binary treatment		Binary treatment		Binary treatment		Arrivals per capita		Arrivals per capita	
Unit:	Municipality		Municipality		Township		Township		Municipality		Municipality	
Exposure	2.079	2.112	-0.040	-0.055	2.272	2.193	0.093	0.127	0.604	0.600	-0.004	-0.033
	(0.351)	(0.674)	(0.392)	(0.713)	(0.382)	(0.552)	(0.284)	(0.417)	(0.178)	(0.264)	(0.119)	(0.262)
Unit FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Election FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Unit trends		✓		✓		✓		✓		✓		✓
N	380	380	285	285	992	992	744	744	379	379	284	284
Elections	4	4	3	3	4	4	3	3	4	4	3	3
Clusters	95	95	95	95	248	248	248	248	95	95	95	95

Notes: Models 1–12 display ordinary least squares (OLS) regression coefficients with clustered standard errors in parentheses. Models 1-8 use a binary treatment indicator while models 9-12 use the number of refugee arrivals per capita. Models 1, 2, 5, 6, 9 and 10 show the effect on GD vote share (in red). Models 3, 4, 7, 8, 11, and 12 use the GD vote share from the previous election as placebo outcome (in blue). All models control for election and unit of analysis (municipality or township) fixed effects. In addition, models 2, 4, 6, 8, 10 and 12 also include unit-specific linear time trends.

$p < 0.782$), i.e. close to zero and far away from conventional level of statistical significance. The right panel of Figure 5 shows the same analysis but using the change in support for GD from January to September 2015 as outcome. For this period, we estimate a linear regression coefficient of -0.433 (two-tailed $p < 0.001$), indicating that after the crisis, proximity to the coast is associated with a significantly higher increase in GD vote share.

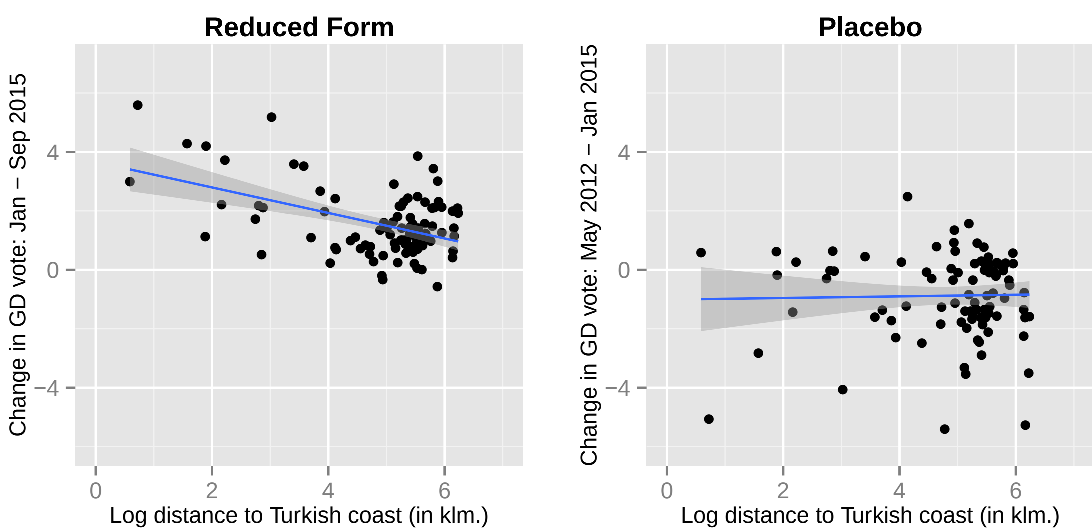

Figure 4: Intention-to-treat effect: Proximity to the Turkish coast increases GD vote share after the inflow of refugees in September 2015, but has no impact on previous elections. The blue lines indicate linear OLS regressions with 95% confidence intervals (shaded areas). The right panel shows the change in GD vote shares from January to September 2015, the left panel shows the change in GD vote shares between May 2012 and January 2015. Both graphs use the logged distance from the Turkish coast in the horizontal axis.

Table S3 in the SI Appendix presents the IV estimates, using the binary and the continuous treatment indicator. Models 1 and 3 show the first-stage regression results. Both the binary and continuous measures for refugee arrivals are negatively related to distance to the Turkish coast. The F statistic measuring the strength of the first stage is 46 and 31, respectively, well above the critical value (Stock and Yogo 2005) of 10. Model 2 shows that the presence of refugees during 2015 increases the GD vote share by more than two percentage points. Model 4 employs the continuous measure and shows that for an additional refugee per capita the vote share for GD increases by almost one percentage point. Both estimates are very close to those obtained using the DID design.[2] Model 5 presents the intent-to-treat result using all four elections since May 2012. The result shows that the effect disappears entirely when moving from an island that is 2 km away from the Turkish coast (e.g. *Samos*) to an island that is almost 100 km away (e.g. *Karpathos*).[3]

Taken together, both empirical strategies lead to the same conclusion: The refugee crisis resulted in a statistically significant and politically substantial increase in electoral support for Europe's most radical right-wing party, GD. In the SI appendix, we extend this evidence in two related dimensions. First we look at the intent-to-treat effect in greater detail, showing that distance to the Turkish coast predicts GD vote share only in September 2015, but not in any previous election (see SI Appendix Table S5). Second, we drop observations far from the Turkish coast. We use several cut-off points, ranging from 500 to 50 kilometer maximum distance to the coast. The results appear in Figure S7 of the SI Appendix and appear remarkably robust to changes in the range of distance to the Turkish coast.

Is GD 'stealing' votes from other parties or mobilizing citizens that would otherwise abstain? While a full analysis of the mechanism underlying these shifts in electoral support is beyond the scope of this study, our data provides some evidence that allows us to examine these two hypotheses. The results, shown in Tables S6 and S7 of the SI Appendix Section S6.5, show that the rise of GD caused by the refugee crises did not affect the vote shares of the governing coalition parties of the leftist *SYRIZA* and *ANEL*, but that the major opposition party *Nea Dimokratia* suffered significant electoral losses. This suggests that a significant segment of voters turned from the center-right *Nea Dimokratia*, whose electoral agenda was dominated by economic issues and the financial bail-out negotiations, to the extreme-right GD, which established itself as the fiercest anti-immigrant and anti-asylum-seeker platform. At the same time, the DID and IV results show that overall turnout increased significantly in treated islands, which suggests that the refugee crisis also enabled GD to mobilize new voters who previously had not participated in elections.

CONCLUSION

This study exploits a natural experiment in the Aegean Sea to examine the effect of exposure to the refugee crisis on natives' support for extreme-right parties. Using two complementary identification strategies and multiple outcomes and placebo tests, we find that in municipalities and townships that experienced sudden and sizeable refugee inflows, electoral support for the extreme-right party GD increased by 2 percentage points, a more than 40 percent increase at the average. These effects are further amplified by the degree of exposure, measured by the number of refugees per resident who arrived in the treated islands. However, given the size and suddenness of the refugee inflows, the ensuing chaos on affected islands just before the election, and the concurrent economic downturn caused by the financial crisis—all factors that have been

[2] Whereas the 2SLS estimator recovers the Local Average Treatment Effect, the DID estimator unpacks the Average Treatment Effect on the Treated (ATT). The similarity in the magnitude of the effects between the two estimation techniques suggests low levels of non-compliance to treatment assignment status in the IV analysis.

[3] The 95 municipalities in our analysis are nested within 20 prefectures. We account for this grouping by clustering standard errors at the prefecture level. To explore whether our inferences are fragile due to the relatively low number of higher-level prefectures, we repeat the IV analysis using bootstrapped standard errors (SI Appendix Table S4). All results are robust to these alternative specifications.

theorized to fuel electoral backlash (Sniderman et al. 2004)—we interpret the 2 percentage point increase in GD vote shares as a relatively modest effect.

These findings have important implications for our theoretical understanding of the dynamic nature of attitudes toward asylum seekers. Mirroring findings from research on how immigration to the U.S. can trigger local anti-immigrant sentiment and policies (Hopkins 2010), we find strong evidence that extreme-right parties can successfully convert prevalent negative attitudes towards Muslim asylum seekers (Bansak et al. 2016) into vote shares in times of crisis. Previous research on the electoral consequences of refugee migration has found negative effects on support for far-right parties in Austria (Steinmayr 2016), lending support to the theory that contact between natives and refugees alleviates animus (Allport 1979). Our study shows that hostility prevails in the Greek context, where asylum seekers quickly continue their journeys from the arrival island to the Greek mainland and extensive contact and interactions between island residents and asylum seekers are therefore not possible. Given the only temporary presence of the refugee population on most of these islands, our findings are also hard to reconcile with theories of realistic group conflict (Campbell 1965), which posits that conflict between the outgroup and dominant group emerges over scarce resources such as access to jobs, housing, or education. In our context, there is no specific competition between refugees and residents on treated islands over any of these resources. Our study therefore shows that mere exposure to the refugee crisis is sufficient to fuel support for extreme-right parties.

Our findings have also direct implications for policymakers regarding the management of refugee flows in Europe. As our analysis demonstrates, a substantial part of the electoral backlash can be attributed to the fact that some islands and townships received a disproportionately high number of asylum seekers in a very short period of time. This, in turn, suggests that xenophobic repercussions among natives might be mitigated if European governments were to invest more resources in supporting Mediterranean countries in processing asylum claims and by allocating asylum seekers and refugees fairly across all countries in the Common European Asylum System (Bansak et al. 2017). With more and more people fleeing civil war, protracted conflicts, poverty, and natural disaster, there is no expectation that in the foreseeable future refugee migration will decrease substantially. If an important goal is to reduce electoral backlash in EU countries of first arrival, European governments need greater cooperation and solidarity in managing the current and future refugee crises.

References

Allport, G. W. (1979). *The nature of prejudice*. Basic books.

Angrist, J. D., G. W. Imbens, and D. B. Rubin (1996). Identification of causal effects using instrumental variables. *Journal of the American statistical Association 91*(434), 444–455.

Bansak, K., J. Hainmueller, and D. Hangartner (2016). How economic, humanitarian, and religious concerns shape european attitudes toward asylum seekers. *Science 354*(6309), 217–222.

Bansak, K., J. Hainmueller, and D. Hangartner (2017). Europeans support a proportional allocation of asylum seekers. *Nature Human Behaviour 1*, 0133.

Barone, G., A. D'Ignazio, G. de Blasio, and P. Naticchioni (2016). Mr. rossi, mr. hu and politics. the role of immigration in shaping natives' voting behavior. *Journal of Public Economics 136*, 1–13.

Becker, S. O., T. Fetzer, et al. (2016). Does migration cause extreme voting.

Bertrand, M., E. Duflo, and S. Mullainathan (2004). How much should we trust differences-in-differences estimates? *The Quarterly journal of economics 119*(1), 249–275.

Board, E. (2015). Ugly attacks on refugees in europe. *New York Times*.

Brunner, B. and A. Kuhn (2014). Immigration, cultural distance and natives' attitudes towards immigrants: Evidence from swiss voting results.

Campbell, D. T. (1965). Ethnocentric and other altruistic motives. In *Nebraska symposium on motivation*, Volume 13, pp. 283–311.

Capon, F. (2015). Greek island ferries fully booked as refugees stream to mainland. *Newseeek*.

Dustmann, C., K. Vasiljeva, and A. P. Damm (2016). Refugee migration and electoral outcomes. *CReAM DP 19*, 16.

Foged, M. and G. Peri (2016). Immigrants' effect on native workers: New analysis on longitudinal data. *American Economic Journal: Applied Economics 8*(2), 1–34.

Halla, M., A. F. Wagner, and J. Zweimüller (2015). Immigration and voting for the far right.

Heinö, A. J. (2016). Timbro's authoritarian populism index 2016.

Hopkins, D. J. (2010). Politicized places: Explaining where and when immigrants provoke local opposition. *American Political Science Review 104*(01), 40–60.

Mendez, I. and I. M. Cutillas (2014). Has immigration affected spanish presidential elections results? *Journal of Population Economics 27*(1), 135–171.

Neumayer, E. (2004). Asylum destination choice what makes some west european countries more attractive than others? *European Union Politics 5*(2), 155–180.

Sniderman, P. M., L. Hagendoorn, and M. Prior (2004). Predisposing factors and situational triggers: Exclusionary reactions to immigrant minorities. *American Political Science Review 98*(01), 35–49.

Steinmayr, A. (2016). Exposure to refugees and voting for the far-right:(unexpected) results from austria.

Stock, J. H. and M. Yogo (2005). Testing for weak instruments in linear iv regression. *Identification and inference for econometric models: Essays in honor of Thomas Rothenberg*.

UNHCR (2015). Emergency information sharing web portal: http://data.unhcr.org/ mediterranean/documents.php?page=28&view=grid&country

UNHCR (2016). Emergency Information Sharing Web Portal: http://data.unhcr.org/mediterranean/regional.php.

UNHCR (2017). Mid-year trends 2016: http://www.unhcr.org/en-us/statistics/unhcrstats/58aa8f247/mid-year-trends-june-2016.html.

Supporting Information

"Waking Up the Golden Dawn: Does Exposure to the Refugee Crisis Increase Support for Extreme-right Parties?"

Contents

S1 Introduction

This Supporting Information is structured as follows: Following the introduction, the second section provides more information about the unfolding of the refugee crisis in Greece, the timing of events, the institutional set-up relevant for our empirical strategy and a discussion of the external validity of our findings. The third section discusses the various data sources for election results, refugee arrivals, and islands' distance to the Turkish border. The fourth section provides descriptive statistics and information about coding decisions. The fifth section provides more details about the difference-in-difference (DID) and instrumental variable (IV) analyses. The last section reports additional results referenced in the main paper and the SI Appendix, including further results of the DID, IV, and intention-to-treat analyses, placebo tests, and effect estimates for parties other than *Golden Dawn* (GD) as well as turnout.

S2 The refugee crisis in context

Greece is the European country that is arguably most strongly affected by the refugee crises due to its proximity to the Turkish border (with Turkey being a major route through which refugees try to reach the EU countries), a long coastline that marks the EU external borders, and many, difficult-to-patrol islands. In 2015 alone, out of the 1.3 million refugees and asylum seekers that reached EU soil for the first time, more than 850,000 of them did so by arriving in one of the Greek Aegean islands [1]. Most refugees left the islands of first arrival within a very short period, typically less than 48 hours, to continue their journeys via the port of *Piraeus* or *Thessaloniki* to central and northern Europe. But while some islands were strongly affected by these sudden refugee inflows, many other islands did not experience any contact with refugees and asylum seekers.

S2.1 Timing of events

The timing of events between the two elections (January and September 2015) is critical for our empirical strategy. The short time that elapsed between the two elections, the speed with which the crisis unfolded in the first half of 2015, the localization of the events, and the fact that the news cycle and political agenda were dominated by the financial crises and capital controls all ensure that there are limited spill-over effects (e.g. via interpersonal contacts or news coverage) across islands. For these reasons, the political and economic impact of the refugee crisis—at least during its initial stages—was mostly limited to the communities and islands that were directly affected by it and were receiving new asylum seekers and refugees. Note that the presence of such spill-over effects from exposed to unexposed islands would bias the absolute value of our estimate downwards.

S2.2 Institutional set-up

The second key element that we exploit in our identification strategy is a special feature of the Greek electoral law (Law 3636/2008) which dictates that if a new general parliamentary election takes place within a period of less than eighteen months since the last general election, then the electoral lists must be closed (as opposed to open-list in regular cases) and the order of candidates must remain unchanged from the last election. The law aims to prevent additional campaign spending by candidates in such a short time interval and to eliminate the necessity for candidates to raise more money. Thus, in the September 2015 election, the lists were closed, and the order of candidates remained unchanged from January 2015. This, in turn, implies that voters were not given the opportunity to express preferences over candidates within a given party list (only preferences over parties), and, hence, individual candidates had no incentive to campaign. But most importantly, this feature of the electoral law effectively guarantees that candidate quality (and ranking) remained constant between the two elections, thus keeping the fundamentals of political competition between the two elections almost unchanged. That is, both in treated and control islands, voters were presented in September with the same party lists as they were in January.

In addition to this, most of the islands in the Aegean Sea belong to the same electoral and administrative districts (NUTS-3). This ensures that they are identical on a plethora of observable and unobservable characteristics such as the candidates running for office, regional government, police, judiciary, and access to EU funds. Together, these institutional features lend further credibility to our identification strategy by holding constant many potential sources of variation across municipalities.

S2.3 External validity

Given that our identification strategy exploits a particular natural experiment in Greece, one could raise legitimate questions on the generalizability of our findings. While we would advice against over-claiming the external validity of our results, we believe that the effects and mechanisms that we identify also has implications for other countries. Albeit not at this scale, the refugee crisis that Greece experienced over spring and summer 2015 is not unprecedented. In the past, Greece as well as other EU countries with extensive sea borders such as Italy, faced structurally similar situations. Furthermore, the demographic and ethnic composition of the refugee population that arrived in Greece is very similar to that in other EU countries. While we believe that external validity is best addressed by replicating this study in other contexts, we may therefore expect to find a similar electoral reaction to large-scale and sudden refugee inflows in other European democracies.

S3 Data

In our analysis, we use three different sources of data: electoral outcomes at the municipality and township level, data on refugee arrivals (temporal and spatial), and geographic data (distance from the Turkish coast). Below we describe in detail how we have collected, processed, and analyzed our data. Upon publication, all data will be made publicly available at the dedicated Dataverse `doi:10.7910/DVN/XXXX`.

S3.1 Electoral data

Our electoral data cover vote outcomes for all the parties that participated in the four Greek legislative elections between May 2012 and September 2015. Our sample includes all inhabited Greek islands. In administrative terms, each island might contain one or more municipalities (large islands such as *Crete* and *Evvoia* contain more than one municipality, while smaller islands contain only one). Electoral data are collected at the municipality and at the township level (some municipalities can have multiple townships) using publicly available sources provided by the Greek Ministry of Interior and Public Administration, the office that is responsible for conducting elections and reporting the official results. The data we use is publicly available and freely accessible on the Ministry's website (`http://www.ypes.gr/el/Elections/NationalElections/Results/`). The data include a) the number of total votes cast, b) the number of votes that each party that participated in elections obtained, c) the number of blank votes, d) the number of valid votes, e) the number of invalid votes, f) the number of registered voters, and g) the number of voters who turned out to vote. Our empirical analysis is based on the vote shares of GD and all other parties. In order to compute the vote share for each party, we have divided the number of votes that each party received over the number of total valid votes cast (that is, excluding blank votes).

S3.2 Refugee arrivals

Our study population consists of all inhabited islands in Greece. The units of analysis are either at the municipality or township level. With the exceptions of *Crete* and *Evvoia*, the two largest islands of the country, each island represents a separate municipality (see above). Municipalities are further disaggregated into townships, thus offering within-island variation in refugee exposure. Data on the number of refugee arrivals are obtained through the United Nations High Commission on Refugees (UNHCR) and are publicly available on the UNHCR website (`http://data.unhcr.org/mediterranean/country.php?id=83`). Data on arrivals are disaggregated at the island (i.e. municipality) level and are available on a monthly basis. Data include aggregate information on the country of origin of refugees, the month and the location of arrival (at the island or municipality level), and other demographic characteristics (such as age, gender, etc.). We code an island (municipality) as treated if it received a positive number of refugees in the period between February and September 2015, when our analysis ends. To measure the intensity of treatment, we compute the cumulative number of refugee arrivals per island inhabitant between February and September 2015. For this time period, we compute the total number of asylum-seekers that arrived in each municipality. We divide this number over the size of the local population that resides in this particular municipality (data obtained via the Greek Ministry of Interior and Public Administration) to obtain our measure of the intensity of treatment. For our analysis at the township level, we code as treated those townships with a hotspot center (points of first reception and recording of refugees on arrival) or a refugee camp, which is a more permanent hospitality facility, or both. Information about the location of these facilities is also obtained from the UNHCR (`http://data.unhcr.org/mediterranean/country.php?id=83`). These data do not provide a detailed breakdown of refugee arrivals (on a monthly basis) at the township level, and, hence, we cannot compute a measure of the treatment intensity at the this level.

S3.3 Distance to Turkish border

Geographic data on the distance between our units of analysis (island or municipality) and the Turkish coast were computed using the web mapping service developed by Google (`https://www.google.com/maps/`) that provides satellite imagery and geospatial data visualization and measurement. For islands that contain a single municipality, we computed the Euclidian distance between the population center of this municipality and the most proximal point in the Turkish coast line as identified by Google Maps. For islands that contain more than one municipality, the distance to the Turkish coast was calculated for each individual municipality using the above algorithm.

S4 Descriptive statistics and Variable coding

S4.1 Descriptive statistics

Tables S1 and S2 display the descriptive statistics of all variables used in the analysis at the municipality and township level.

Table S1: Descriptive statistics at the municipality level.

	Mean	SD	Min	Max
Binary treatment	0.13	0.33	0	1
Arrivals per capita	0.33	1.01	0	5
Log distance	4.86	1.24	0.59	6.23
Registered voters	485.43	180.17	92	1564
Valid votes	227.80	87.08	30	408
Turnout in Sep 2015 (%)	48.41	9.76	18.03	64.04
Turnout in Jan 2015 (%)	51.57	13.35	17.79	74.32
GD vote share in Sep 2015 (%)	6.02	2.47	0	17.51
GD vote share in Jan 2015 (%)	4.46	2.13	0	11.93
GD vote share in Jun 2012 (%)	5.44	2.37	0.483	12.62
GD vote share in Jan 2012 (%)	5.59	3.15	0.794	18.84
Nea Dimokratia vote share in Sep 2015 (%)	28.51	6.62	12.32	47.50
Nea Dimokratia vote share in Jan 2015 (%)	29.84	8.33	13.09	54.14
PASOK vote share in Sep 2015 (%)	8.93	3.99	1.639	21.14
PASOK vote share in Jan 2015 (%)	6.27	3.95	0	24.41
KKE vote share in Sep 2015 (%)	5.59	4.56	1.126	33.20
KKE vote share in Jan 2015 (%)	5.44	4.58	0	31.83
SYRIZA vote share in Sep 2015 (%)	34.80	6.43	14.17	49.81
SYRIZA vote share in Jan 2012 (%)	35.34	8.96	12.30	61.82
ANEL vote share in Sep 2015 (%)	3.66	1.86	0	10.59
ANEL vote share in Jan 2012 (%)	4.70	2.58	0	19.13
Municipalities	95			

Notes: Table shows the mean, standard deviation, and minimum and maximum values for each variable used in the analyses.

S4.2 Top coding

Figure S1 shows the distribution of the refugee arrivals among all treated islands. For all but one treated islands, the number of refugee arrivals varies from fewer than one refugee for every resident (*Kalymnos*) up to four and a half refugees per resident (*Lesvos*). The only exception is *Agathonisi*, a tiny island with less than two hundred residents but a massive number of refugee arrivales. By the

Table S2: Descriptive statistics at the township level.

	Mean	SD	Min	Max
Binary treatment	0.09	0.28	0	1
Registered voters	473.56	165.98	109	1564
Valid votes	231.97	92.67	34	459
Turnout in Sep 2015 (%)	49.42	10.27	16.84	69.84
Turnout in Jan 2015 (%)	54.08	13.78	12.22	77.09
GD vote share in Sep 2015 (%)	6.36	2.54	0	17.51
GD vote share in Jan 2015 (%)	4.63	2.14	0	11.93
Nea Dimokratia vote share in Sep 2015 (%)	27.18	7.15	8.441	52.20
Nea Dimokratia vote share in Jan 2015 (%)	28.01	8.77	9.472	56.48
PASOK vote share in Sep 2015 (%)	8.89	4.08	0	22.02
PASOK vote share in Jan 2015 (%)	5.88	3.34	0	24.41
KKE vote share in Sep 2015 (%)	6.27	5.08	0.345	37.35
KKE vote share in Jan 2015 (%)	5.99	5.11	0	36.37
SYRIZA vote share in Sep 2015 (%)	35.42	6.87	12.96	52.66
SYRIZA vote share in Jan 2012 (%)	36.95	8.81	8.108	67.86
ANEL vote share in Sep 2015 (%)	3.52	1.92	0	17.62
ANEL vote share in Jan 2012 (%)	4.77	3.62	0	32.93
Townships	248			

Notes: Table shows the mean, standard deviation, and minimum and maximum values for each variable used in the analyses.

end of ours study period, *Agathonisi* received 125 times more refugees than its resident population. The huge gap between *Agathonisi* and the rest of the islands generates an extreme interpolation in our estimation. To avoid this problem, we top-code *Agathonisi*, using the value of five, still the highest in the data. The second panel of the figure shows the distribution of refugee arrivals after top-coding, which we use for further analysis. This strategy does not affect the binary treatment analysis.

Figure S1: The distribution of refugee arrivals with and without top-coding.

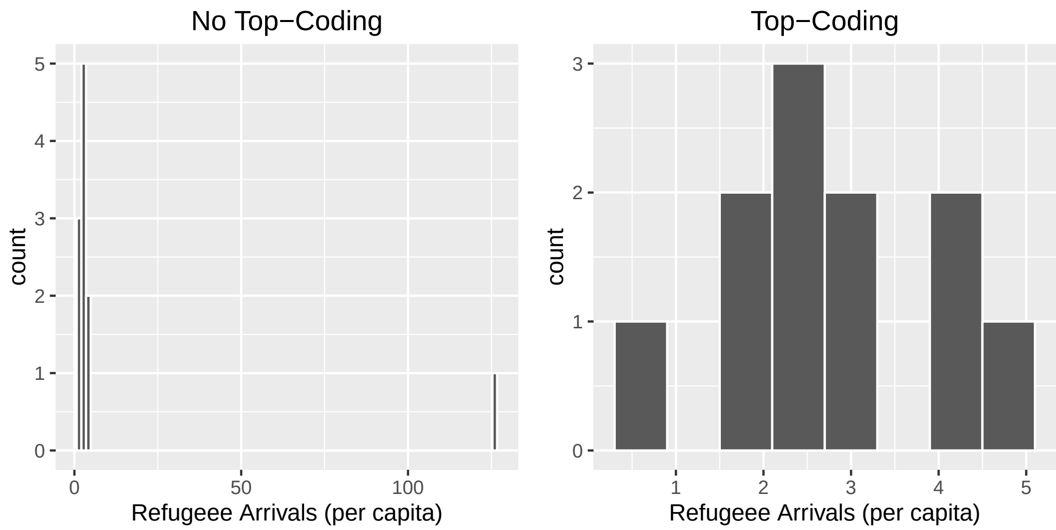

Note: The left panel shows the density of refugees per capita and illustrates the interpolation problem caused by one island, *Agathonisi*, which has received 125 times more refugees than its resident population. The right panel shows the same distribution when we top-code *Agathonisi* at five (the maximum value in the data).

S5 Estimation strategies

We use two complementary identification strategies: DID estimation that relies on changes in the number of refugee arrivals and voting behavior over time and IV analysis that additionally leverages the distance to the Turkish coast as an instrument for the number of refugee arrivals. We briefly describe each estimation strategy below.

S5.1 DID analysis

Our DID model uses municipalities or townships as the unit of analysis. Below, we present the municipality-level specification thebut one for townships is completely analogous. Two estimate the effect of refugee inflows on GD vote share, we use a two-way fixed effects regression given by

$$GD_{s,t} = \gamma_s + \lambda_t + \delta_{DID} T_{s,t} + u_{s,t} \,,$$

where $GD_{s,t}$ is the local vote share for GD in municipality s and election t; γ_s is a municipality fixed effect that rules out omitted variable bias from unobserved municipality characteristics that are invariant over our study period; λ_t is an election fixed effect to control for common factors that change nonlinearly over time, $T_{s,t}$ is the (binary or continuous) treatment indicator measuring refugee exposure, and $u_{s,t}$ is an idiosyncratic error term. The quantity of interest is δ_{DID}, which identifies the effect of refugee inflows on GD vote share based on the within-municipality variation among municipalities that have received refugees between spring and summer 2015 (the average treatment effect for the treated). As a robustness check, we further relax the model specification and add municipality-specific linear time trends. This ensures that all unobserved municipality-specific differences that vary smoothly over time (such as local trends in voter preferences) are purged from the estimate of δ_{DID}.

S5.2 IV analysis

In the IV analysis, identification relies solely on the exogenous variation in the distance to the Turkish coast, our instrument. In order to serve as a valid instrument, three assumptions have to hold. First, islands close to the Turkish coast have to have a higher propensity (or number) of refugee arrivals (first stage). Second, distance to the Turkish coast can only affect changes in GD vote share through refugee exposure (exclusion restriction). Third, we have to rule out any other time-varying confounder that affect closer islands more (less) than islands further away and simultaneously impacts changes in GD vote share (independence of the instrument). Under these assumptions, we can consistently estimate the impact of refugee exposure on changes in GD vote shares using two-stage least squares (2SLS) regression. The first stage is given by:

$$T_s = \alpha + \beta Z_s + v_s \,,$$

where T_s is the binary (or continuous) treatment indicator of refugee exposure, Z_i is the logged distance to the Turkish coast, α an estimated constant, β the coefficient measuring the strength of the first stage, and v_i an idiosyncratic error term. The second stage regression is given by

$$\Delta GD_s = \gamma + \delta_{IV} \hat{T}_s + u_s$$

where ΔGD_s is the change in GD vote share between January and September 2015, $\hat{T}_s$ is the instrumented treatment indicator; γ an estimated constant; and u_s an idiosyncratic error term assumed to be orthogonal to v_s. Here, the quantity of interest is δ_{IV}, which identifies the causal effect of refugee exposure on change in GD vote share by leveraging the distance to the Turkish coast as an instrument.

Table S1 displays the results of the IV analysis. Models 1 and 3 present the results of the first-stage estimation of the binary and continuous (arrivals per resident) treatment models, while Models 2 and 4 present the results of the corresponding second-stage estimation. Model 5 shows the intention-to-treat analysis for all four elections and confirms that distance to the Turkish coast only has an effect on GD vote share in the September 2015 election, after the onset of the refugee crisis.

With only 20 districts, the clustered standard errors in Table might be biased. In order to assess this issue, Table S4 replicates Models 1–4 above but uses the bootstrap to calculate standard errors. We find that the standard errors are virtually identical.

Table S3: 2SLS regressions of change in GD vote share on refugee exposure instrumented by distance to the Turkish coast.

Model:	(1)	(2)	(3)	(4)	(5)
Outcome:	Difference in GD vote share: January to September 2015				GD vote share
Treatment:	Binary Treatment		Arrivals per capita		Distance to coast
Stage:	First Stage	Second Stage	First Stage	Second Stage	Reduced Form
Log Distance	-0.208		-0.568		-0.203
	(0.031)		(0.102)		(0.354)
Instrumented refugee arrivals		2.080		0.739	
		(0.478)		(0.173)	
2012 June					-1.074
					(0.869)
2012 January					-0.893
					(0.894)
2015 September					2.769
					(0.657)
2012 Jun × Log Distance					0.190
					(0.171)
2015 Jan × Log Distance					0.048
					(0.181)
2015 Sep × Log Distance					-0.481
					(0.143)
Constant	1.136	1.295	3.103	1.290	6.574
	(0.168)	(0.104)	(0.552)	(0.173)	(1.785)
F statistic	46.21		31.08		
N	95	95	94	94	380

Notes: Models 1 and 3 display the coefficients of the first stage of a two stage least squares (2SLS) regression. Models 2 and 4 show the coefficients of the corresponding second stage. Model 5 shows ordinary least squares (OLS) coefficients of the reduced form regression of GD vote share on the distance from the Turkish Coast. Standard errors, shown in parentheses, are clustered at the district level in models 1–4, and at the municipal level in model 5.

Table S4: IV estimates using bootstrapped standard errors.

Model	(1)	(2)	(3)	(4)
Outcome	Difference in GD vote share: January to September 2015			
Treatment	Binary Treatment		Treatment Intensity	
Stage	First Stage	Second Stage	First Stage	Second Stage
Log Distance	-0.208		-0.568	
	(0.031)		(0.102)	
Instrumented Refugee Arrivals		2.083		0.739
		(0.529)		(0.202)
Constant	1.136	1.295	3.103	1.290
	(0.168)	(0.103)	(0.552)	(0.096)
F statistic	46.21		31.08	
N	95	95	94	94

Notes: Models 1 and 3 display the coefficients of the first stage of a two stage least squares (2SLS) regression. Models 2 and 4 show the coefficients of the corresponding second stage. Standard errors, shown in parentheses, are bootstrapped with 300 replication samples.

S6 Additional results

S6.1 Proximity to refugee hotspots and changes in GD vote shares

S6.2 Visualization of DID estimates

S6.3 Intent-to-treat effect

To further investigate the relationship between distance from the Turkish coast and GD support, we estimate a regression model in which we interact logged distance with each election-dummy (using the May 2012 election as baseline). The results are presented in Table S5. We see that distance from the Turkish coast has no effect on GD vote share in the May 2012, June 2012 and January 2015 elections. Only for the September 2015 election, after the onset of the refugee crisis, do we find that logged distance decreases the vote share for GD, and significantly more so than for the May 2012 baseline election. The difference in the slope between 2015 and May 2012 is -0.423 (std. error 0 .155, two-tailed $p < 0.007$); between September 2015 and June 2012 -0.676 (std. error 0.133, two-tailed $p < 0.001$);

Figure S2: Townships that hosted refugees experienced a higher increase in GD vote shares than townships on the same island that did not. Panel A indicates the townships on Aegan islands that received refugees during the period from January to September 2015. Panel B shows the change in the GD vote share during the same period.

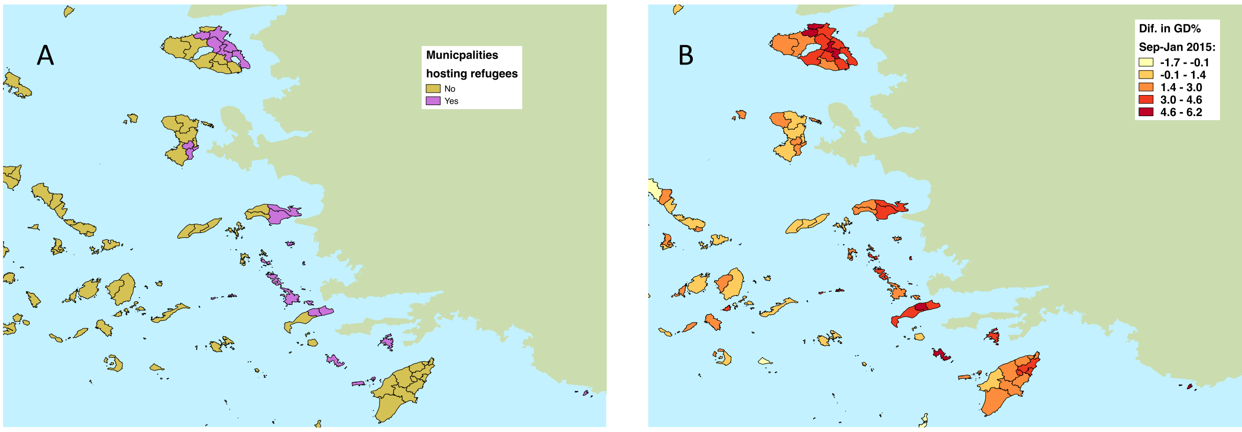

Figure S3: DID estimates of the impact of refugee arrivals on GD vote shares and placebo tests

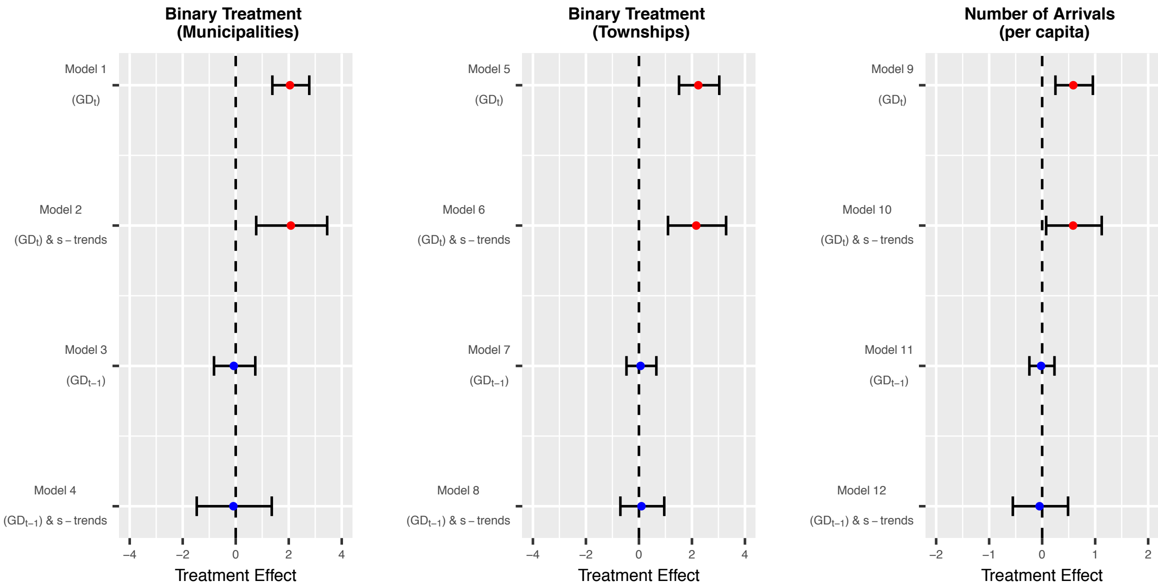

Notes: The red and blue dots denote DID regression coefficient of the average treatment effect on the treated. The horizontal bars display the 95% confidence intervals. Models 1, 5, and 9 show the treatment effects (red dots) of the baseline model for binary treatment (municipalities), binary treatment (townships), and continuous treatment (municipalities), respectively. Models 2, 6, and 10 (red dots) use the same specification but also include unit-specific trends (s-trends). Models 3, 7, and 11 (blue dots) show the estimates of baseline placebo models, while models 4, 8, and 12 (blue dots) show the estimates of the placebo models with s-trends.

and between September 2015 and January 2015 -0.463 (std. error 0.116, two-tailed $p < 0.001$). Figure S4 visualizes this pattern, showing the marginal effect of logged distance on GD vote share for each of the four elections.

S6.4 Additional placebo tests

In addition to the placebo tests using pre-crisis elections, we can also exploit a specific feature of the Greek electoral law for an additional placebo test. In Greece, in the absence of postal votes, registered

Table S5: Intention-to-treat effect of distance to the Turkish coast on GD vote share.

Intention-to-treat effect	
(Logged) Distance from the Coast	-0.203
	(0.354)
2012 June	-1.074
	(0.869)
2015 January	-0.893
	(0.894)
2015 September	2.769
	(0.657)
Logged Distance X 2012 June	0.190
	(0.171)
Logged Distance X 2015 January	0.048
	(0.181)
Logged Distance X 2015 September	-0.481
	(0.143)
Constant	6.574
	(1.785)
N	380
Clusters	95

Figure S4: Intention-to-treat effect of distance to the Turkish coast on GD vote share.

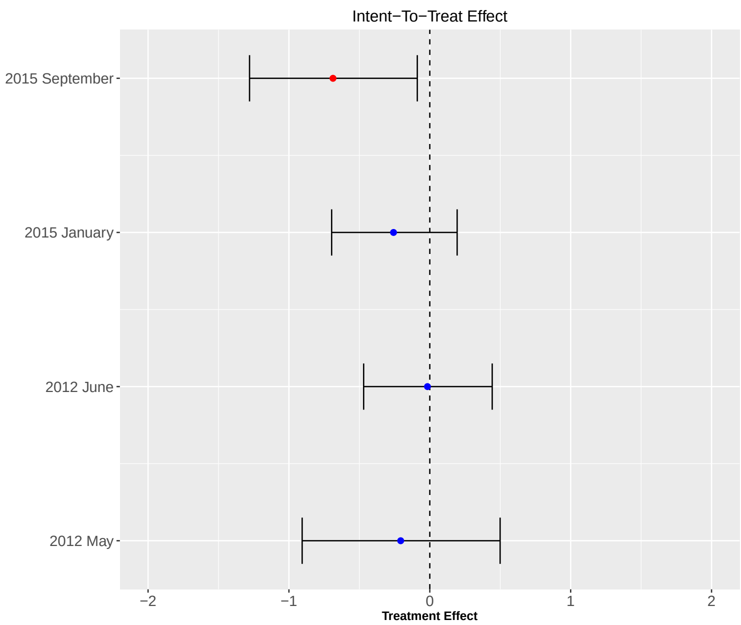

Notes: The red (blue) dots denote the intention-to-treat effect of logged distance on GD vote share for the September 2015 (pre-refugee crisis) elections. The black bars display 95% confidence intervals, with standard errors clustered at the municipality level.

non-resident voters are allowed to vote in the area that they reside, in special polling stations and separate ballot boxes for the electoral district in which they are registered to vote. Their ballots are then collected and counted separately. This means that voters who are registered in treated or control islands but reside in other parts of Greece voted in their area of current residency but used the exact same electoral lists that voters on the islands received. This allows us to examine the behavior of non-residents who are registered to vote on islands with refugee exposure. As an example, consider a non-resident voter registered in *Lesvos* who currently resides on the mainland in *Athens*, and was therefor not directly exposed to the refugee arrivals on her home island. We leverage this setting for a DID analysis that uses as a placebo treatment group the non-resident voters of the treated islands.

We conduct two sets of placebo sets. For the first test, we compare non-resident voters of treated islands to non-resident voters of control islands. In the second placebo test, we compare non-resident voters of treated islands to resident voters of control islands. For both tests, Figure S5 shows that changes in vote shares across the different groups, whereas Figure S6 displays the treatment effects with and without municipality-specific trends. The placebo tests confirm that electoral support for GD did not increase substantially among non-resident voters originating from treated islands between the two elections in 2015 when compared to i) non-resident voters of the control islands (see left panels of Figures S5 and S6) or ii) to the resident voters of the control islands (see right panels of Figures S5 and S6). If anything, they appear less prone to increase their electoral support for GD compared to resident voters from control islands, which implies that our estimates are most likely a lower bound of the impact of refugee arrivals.

Figure S5: DID placebo estimates comparing non-resident voters from treated islands with resident voters from control islands.

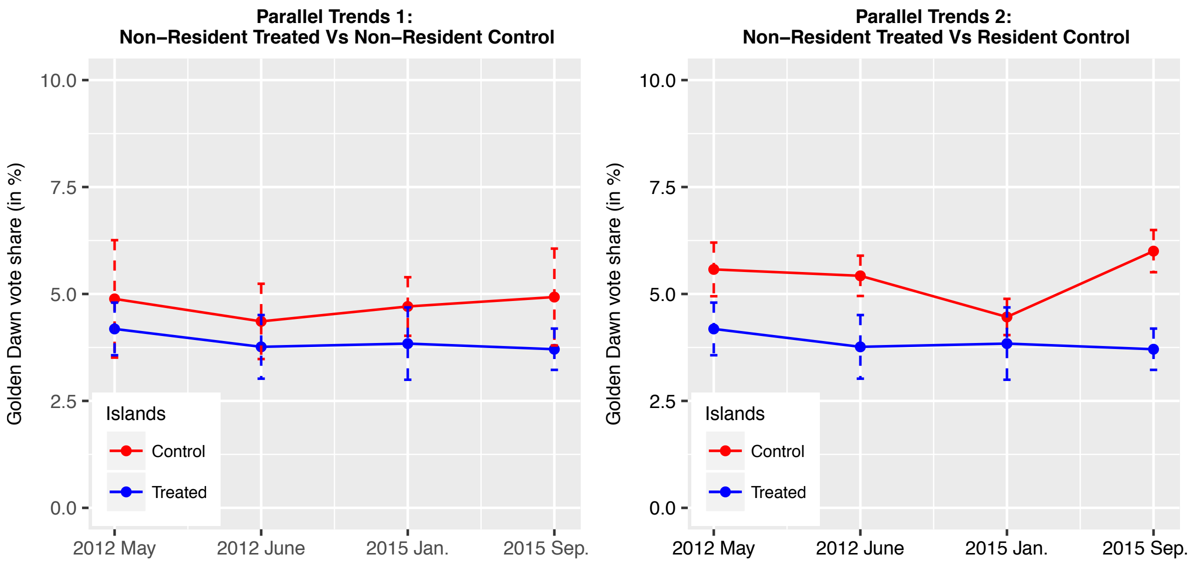

Notes: The left panel shows that non-resident voters registered to vote in treated islands do not differ from non-resident voters from control islands in their change of support for GD. The right panel shows that non-resident voters originating from treated islands are, if anything, less likely increase support for GD party compared to resident voters from control islands.

S6.5 Impact on turnout and vote shares of other parties

In this section, we explore the effect that exposure to the refugee crisis has on overall turnout and the electoral performance of the other parliamentary parties that contested the January and September 2015 elections. Employing the same DID and IV analysis used to generate the main results, we replicate them using the vote share for each of the other parliamentary parties and find no significant changes for all of them except for the center-right *Nea Dimokratia*, which incurred losses between 1 and 4 percentage points depending on the specification. *Nea Dimokratia*, whose electoral agenda was dominated by economic issues and the financial bail-out negotiations, was *SYRIZA*'s major competitor in the January and September 2015 elections. Tables S6 and S7 report the estimates. In addition, treated islands also experienced higher levels of turnout between 1 and 5 percentage points depending on the specification. Taken together, these results suggest that in treated islands, GD successfully attracted former voters of the *Nea Dimokratia* as well as mobilized additional voters that have not participated in the January 2015 election.

Table S6: *DID estimates of the impact of refugee arrivals on turnout and vote shares of the other parties.*

Model	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Treatment	Binary: Area received refugees? (YES/NO)								Arrivals per capita			
Unit of Analysis	Municipality				Township				Municipality			
Outcome	$Party_t$		$Party_{t-1}$		$Party_t$		$Party_{t-1}$		$Party_t$		$Party_{t-1}$	
Nea Dimokratia	-1.643	-4.245	2.294	0.833	-1.273	-4.498	3.322	3.106	-0.428	-1.080	0.567	0.178
	(1.141)	(1.421)	(1.036)	(1.971)	(0.847)	(1.561)	(0.987)	(1.649)	(0.436)	(0.482)	(0.313)	(0.753)
SYRIZA	0.597	2.940	-2.132	-1.034	-1.062	2.040	-2.592	-0.377	-0.463	0.588	-0.976	-0.552
	(1.575)	(2.313)	(1.862)	(1.583)	(1.438)	(2.006)	(0.943)	(0.972)	(0.349)	(0.687)	(0.490)	(0.610)
PASOK	-0.228	-2.891	2.385	1.012	0.162	-1.376	1.387	0.626	0.045	-1.338	1.315	0.858
	(0.633)	(2.471)	(2.047)	(2.053)	(0.443)	(1.819)	(1.537)	(1.795)	(0.201)	(1.105)	(0.926)	(0.723)
ANEL	-1.328	2.005	-2.286	1.759	-0.708	1.598	-1.805	0.248	-0.118	0.668	-0.583	0.227
	(0.834)	(1.045)	(0.928)	(1.236)	(0.709)	(1.028)	(0.872)	(0.894)	(0.181)	(0.374)	(0.272)	(0.301)
KKE	-1.355	-1.244	-0.034	0.239	-1.574	-1.043	-0.312	0.508	-0.340	-0.130	-0.231	-0.266
	(0.514)	(1.071)	(0.811)	(1.088)	(0.438)	(0.750)	(0.722)	(0.820)	(0.146)	(0.265)	(0.206)	(0.277)
Turnout	2.134	1.420	0.862	1.236	2.082	1.429	0.514	-0.058	0.881	0.680	0.215	0.228
	(1.027)	(1.885)	(1.224)	(1.555)	(0.713)	(1.591)	(0.912)	(1.114)	(0.326)	(0.528)	(0.300)	(0.442)
N	380	380	285	285	992	992	744	744	379	379	284	284
N of clusters	95	95	95	95	248	248	248	248	95	95	95	95
N of elections	4	4	3	3	4	4	3	3	4	4	3	3
Fixed Effects												
Municipality/Township	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Election	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Unit-specific trends		✓		✓		✓		✓		✓		✓

Notes: Models 1–12 display ordinary least squares (OLS) regression coefficients with clustered standard errors in parentheses. Models 1-8 use a binary treatment indicator while models 9-12 use the number of refugee arrivals per capita. Models 1, 2, 5, 6, 9 and 10 show the effect on contemporary elections. Models 3, 4, 7, 8, 11, and 12 use the party vote share from the previous election as placebo outcome. All models control for election and unit of analysis (municipality or township) fixed effects. In addition, models 2, 4, 6, 8, 10 and 12 also include unit-specific linear time trends.

Table S7: *IV estimates of the impact of refugee arrivals on turnout and vote shares of the other parties..*

Party	*Nea Dimokratia*		*SYRIZA*		*PASOK*		*ANEL*		*KKE*		*Turnout*	
Treatment	Binary	Arrivals p.c.	Binary	Arrivals p.c.	Binary	Arrivals p.c.	Binary	Arrivals p.c.	Binary	Arrivals p.c.	Binary	Arrivals p.c.
LATE	-6.769	-2.532	5.065	1.736	-1.327	-0.426	0.732	0.294	-1.845	-0.547	5.496	2.054
	(1.968)	(0.776)	(2.342)	(0.908)	(1.002)	(0.352)	(1.158)	(0.425)	(0.753)	(0.310)	(2.053)	(0.788)
N	95	94	95	94	95	94	95	94	95	94	95	94
N of clusters	20	20	20	20	20	20	20	20	20	20	20	20

Note: Estimates are obtained from the 2SLS regression. Robust standard errors, clustered at the municipality level, are shown in parentheses. The outcome variable is the percentage change between January and September 2015 in the vote share for each party. The first stage regression for the binary and the continuous treatments are identical to the specification used for Table S2.

Figure S6: Placebo tests: Resident vs. non-resident voters.

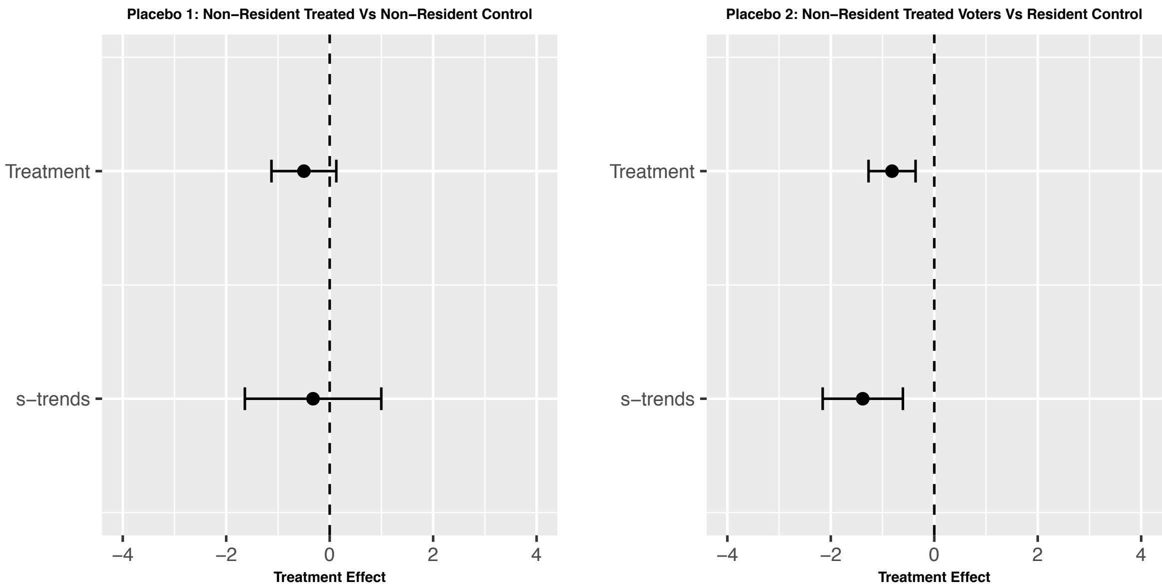

Notes: The black dots show the ATET from the DID regression. Solid black lines indicate 95% confidence intervals. The placebo tests shows that GD vote share did not increase among non-resident voters of treated islands when compared to non-resident voters of control islands (left panel) or resident voters of control islands (right panel).

S6.6 Sensitivity Analysis: Distance From Turkish coast

Distance from the Turkish coast plays a key role in our identification strategy, because it helps us predict which islands were exposed to refugee arrivals and which ones were not. Implicitly this logic assumes that distance from the coast is not related to other potential determinants of change in GD vote share between the post- and pre-treatment elections. Although this assumption can be expected to hold locally, in the part of the Aegean sea that is relatively close to the coast, it might not hold when we expand the radius to all islands in the country. To examine whether this is the case, we repeat the main analyses, focusing on islands closer to the Turkish coast. Distance ranges from one to more than 530 *klm*. We use various cut-off points, from 500 and up to 50 *klm*. Using each cut-off point as the maximum distance to the coast, we repeat both the DID and the IV analysis. The results of this exercise are shown in Figure S7. As expected the level of uncertainty increases as the maximum distance to Turkish coast decreases. Yet, throughout the range, all treatment effect estimates remain remarkably robust.[1] This is the case even when we include unit-specific linear trends in the DID analysis. The evidence seems to rule out the possibility that the effects are due to some distance-related confounder.

As a a way to further assess the role of distance to the coast, we use it as a predictor of a series of socioeconomic indicators. There results are shown in Table S8. We use population, area (in Km^2), population density, GDP (p/c), unemployment, an indicator about tourist activity in the area and rates of foreign population. Exact information about the measurement of these indicators is provided in the note of the table. *Distance* does not seem to predict to any of these outcomes. This evidence matches well the results of the main text, as well as Table S5 and Figure S4, which show that distance is unrelated to change in GD vote share in the pre-treatment period.

Table S8: Balance Tests.

	Area Km^2	# Inhabitants (2011)	Population Density (2011)	Unemployment Rate (2011)	% Non-Natives (2011)	Tourism (2014)	GDP p/c (2014)
Logged	-235.4	-16.10	6.362	0.414	-0.141	1.698	158.6
Distance	(2638.2)	(35.68)	(4.968)	(0.334)	(0.168)	(1.579)	(165.4)
n	95	95	95	95	95	95	95

Notes: Table shows the OLS coefficient of (logged) distance to Turkish coast as predictor of each of the variables shown in the first row of the table. Standard errors in parentheses. Tourism stands for the percentage of bed occupancy in hotel accommodation. Source: XXX.

S6.7 Alternative inference strategies for the DID estimates

An oft-neglected problem with difference-in-differences estimates is that conventional standard errors are inconsistent, because they ignore the serial correlation stemming from repeated measurements over time. We try to address this problem of intra-class correlation in our main analysis by clustering observations either at the municipality or the township level, depending on the unit of analysis. Alternative methods have been suggested, however. We employ those that seem to perform best with small number of units (Bertrand et al. 2004): block bootstrapping and ignoring time-series information. Moreover, we perform a set of permutation-based placebo tests, which help to shed light on the variance of the DID estimators used in the main analysis.

Block Bootstrap

[1] The only exception here is the IV estimate when including islands up to 50 klm from the Turkish coast. No island more than 50 klm far from the Turkish coast received refugees. However, not all islands within this range were treated. Some did not receive refugees (e.g. Rhodes, Nisiros). This makes distance a weak instrument within this range, resulting into a non-significant first stage (OLS coefficient -0.111 with std error 0.091), which in turn generates second-level 2SLS estimates characterized by high levels of uncertainty, as shown in the graph. Increasing the threshold to 100 klm is sufficient to turn distance into a strong instrument of refugee exposure (OLS coefficient -0.277 with std error 0.032), yielding reliable second-stage estimates.

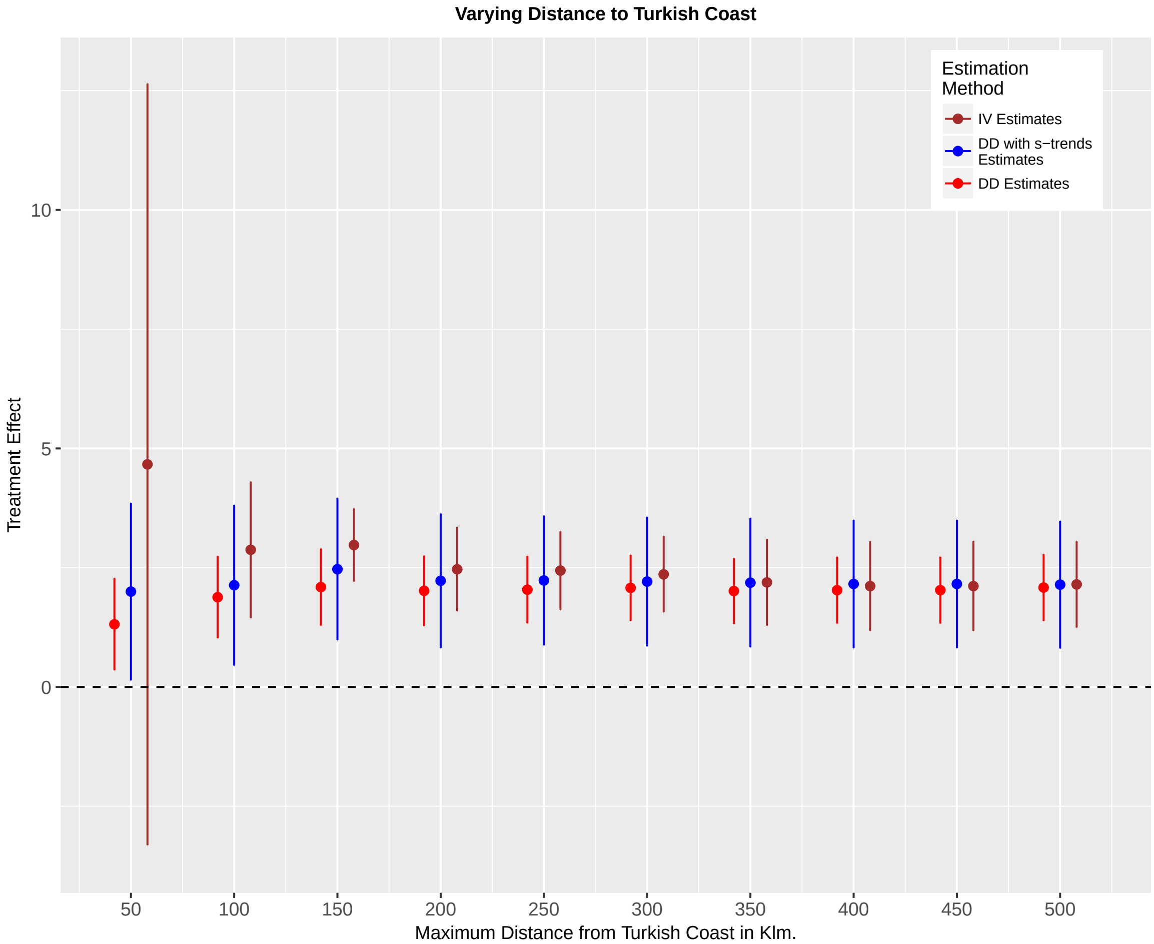

Figure S7: Sensitivity of the Effects to Distance from Turkish Coast.

Note: Each entry denotes the treatment effect of refugee exposure on GD vote share, conditional on the distance from the Turkish coast. The horizontal axis indicates the maximum distance from the Turkish coast in each analysis. The vertical spikes encapsulate the 95% confidence intervals.

The block bootstrapped analysis is shown in Table S9. Municipalities (Model 1and Model 2 for the binary treatment and Model 5 and 6 for refugee exposure) and townships (Model 3 and Model 4) are resampled with replacement (1000 iterations). The DID estimator is used in each bootstrapped sample. The variance of the resulting empirical distribution of treatment effect estimates is used to derive the standard errors. As show in the Table, block bootstrapping causes no change in our inference about the effect of refugee exposure on GD vote share.

Table S9: Impact of refugee arrivals on GD vote share, block bootstrapped standard errors.

Model:	(1)	(2)	(3)	(4)	(5)	(6)
Outcome:	$\text{GD}_{(t)}$	$\text{GD}_{(t-1)}$	$\text{GD}_{(t)}$	$\text{GD}_{(t-1)}$	$\text{GD}_{(t)}$	$\text{GD}_{(t-1)}$
Treatment:	Binary treatment		Binary treatment		Arrivals per capita	
Unit:	Municipality		Township		Municipality	
Exposure	2.079	-0.040	2.272	0.093	0.604	-0.004
	(0.348)	(0.340)	(0.281)	(0.265)	(0.179)	(0.106)
Unit FE	✓	✓	✓	✓	✓	✓
Election FE	✓	✓	✓	✓	✓	✓
N	380	285	992	744	379	284
Elections	4	3	4	3	4	3
Clusters	95	95	248	248	95	95

Notes: Models 1–6 display ordinary least squares (OLS) regression coefficients with block-bootstrapped standard errors in parentheses. Models 1-4 use a binary treatment indicator while models 5 and 6 use the number of refugee arrivals per capita. Models 1, 3 and 5 show the effect on GD vote share (in red). Models 2, 4 and 6 use the GD vote share from the previous election as placebo outcome (in blue). All models control for election and unit of analysis (municipality or township) fixed effects.

Ignoring Time-Series Information

Another, relatively conservative, approach, which however seems to perform well with even low number of units is to ignore the time dimension in the data, by collapsing all observations into one pre-treatment period. We do this in two ways. First, we take the average of the GD vote share in all pre-treatment elections (Models 1–3); second, we use only the last pre-treatment election, January 2015 (Models 4–6). We present the results from this exercise in Table S10. Again, inference remains intact to this exercise. The resulting estimates are remarkably close to those reported in the main analysis.

Table S10: Impact of refugee arrivals on GD vote share: Pre-Post Analysis.

Model:	(1)	(2)	(3)	(4)	(5)	(6)
Outcome:	$GD_{(t)}$	$GD_{(t)}$	$GD_{(t)}$	$GD_{(t)}$	$GD_{(t)}$	$GD_{(t)}$
Treatment:	Binary treatment		Arrivals p.c.	Binary treatment		Arrivals p.c.
Unit:	Municipality	Township	Municipality	Municipality	Township	Municipality
Exposure	2.079	2.272	0.604	2.105	2.210	0.606
	(0.351)	(0.263)	(0.178)	(0.385)	(0.290)	(0.171)
Unit FE	✓	✓	✓	✓	✓	✓
Post-Period	✓	✓	✓	✓	✓	✓
N	190	496	190	190	496	190
Time Points	2	2	2	2	2	2
Clusters	95	248	95	95	248	95

Notes: All models display ordinary least squares (OLS) regression coefficients with clustered standard errors in parentheses. Pre-treatment elections have been collapsed into one pre-treatment observation. Models 1–3 use the average GD vote share in all elections prior to refugee arrivals as the pre-treatment observation for each municipality. Models 4–6 use only the last election prior to refugee arrivals, January 2015, as the pre-treatment observation for each municipality. Models 1, 2, 4 and 5 use a binary treatment indicator while models 3 and 6 use the number of refugee arrivals per capita. All models include a post-period dummy and municipality (or township) fixed effects.

Randomization Inference

Finally, we also implemented two sets of permutation-based tests. The first uses all units and randomly apply the treatment status to few of them (the same number as the number originally treated). Permutations are taken at the municipality (or township) level. After each permutation we end up with a set of municipalities (or townships) being treated. We implement the DID analysis for all permutations and plot the treatment effect estimates in the first column of Figure S8. The first row presents the municipality-based empirical distribution of 2,000 such estimates, whereas the second row displays the empirical distribution from the township-based analysis. The second column repeats this exericse but using only the control islands. Again, 2,000 placebo treatment effects are estimated. Each distribution is compared to the treatment effect obtained from the original dataset. The results confirm previous analyses in that it seems quite unlikely that our original treatment effect estimates are due to sheer chance.

Figure S8: Permutation-based evaluation of the DID effects.

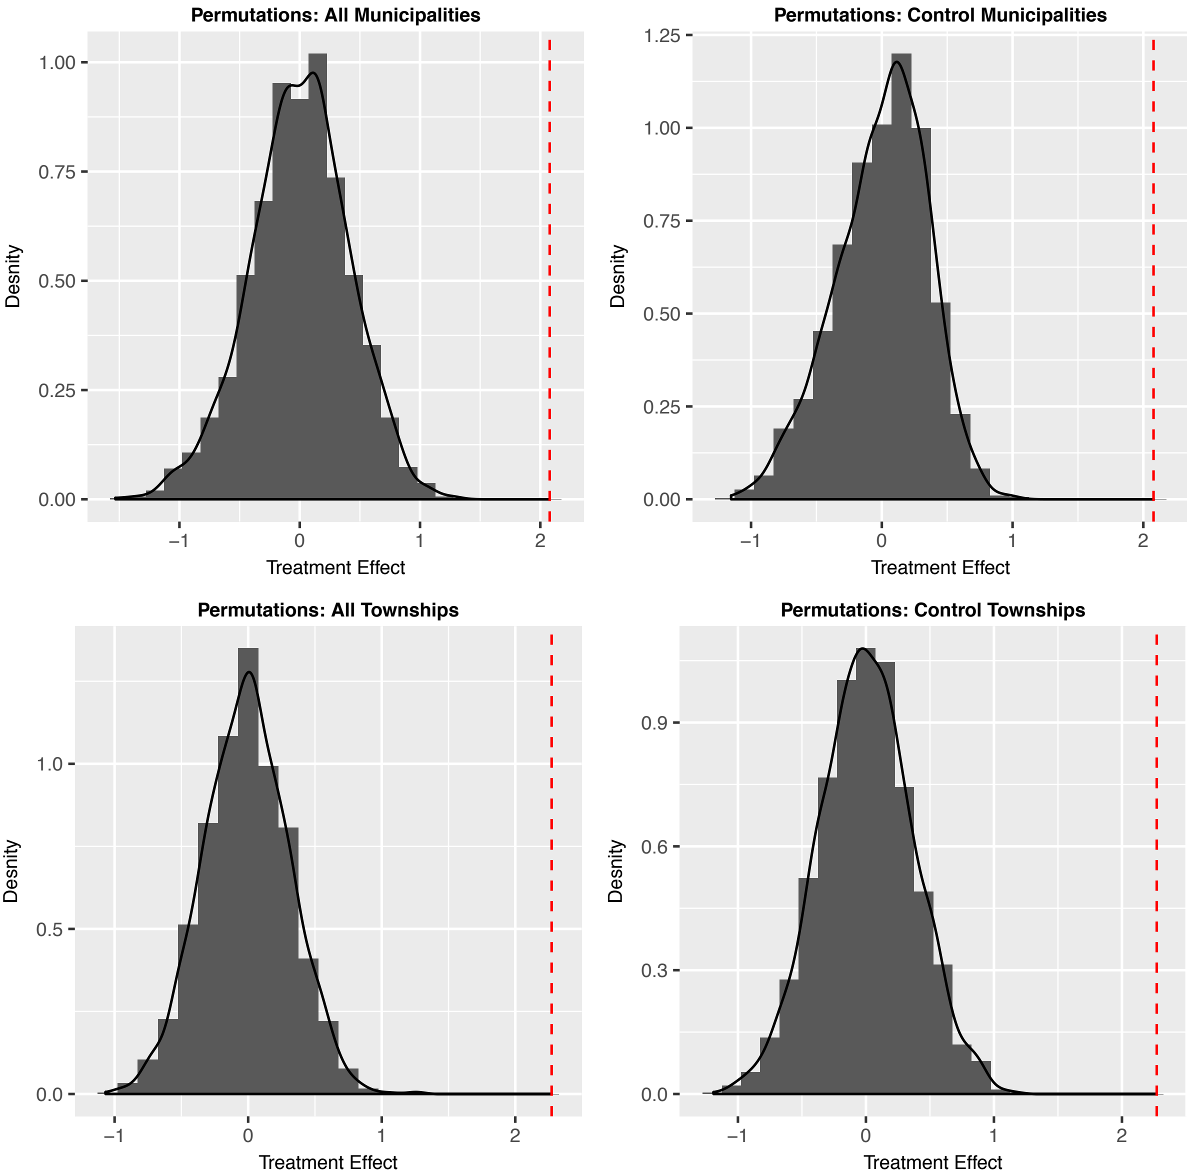

Note: Each graph displays 2,000 treatment effect estimates, based on placebo difference-in-difference analyses. The first two sets of analyses are based on randomly assigning municipalities in the post-treatment period into treatment and control condition. The last two panels follow the same procedure but the analysis is implemented at the township level and thus assigns treatment randomly to townships in the post-treatment period. In the two panels of the first column, 2,000 permutations are drawn from the full set of islands. In the second column graphs, 2,000 permutations are drawn from the set of control municipalities (townships). In both analysis permutations are clustered at the municipality (township level. The vertical red dashed line in each graph denotes the treatment effect from the original analysis.

References

Bertrand, M., E. Duflo, and S. Mullainathan (2004). How much should we trust differences-in-differences estimates? *The Quarterly journal of economics 119*(1), 249–275.