

SUPPORTING INFORMATION

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

Elias Dinas

*Department of Political and Social Sciences and RSCAS, European University Institute
Department of Politics and International Relations, Oxford University,*

Konstantinos Matakos

Department of Political Economy, King's College London

Dimitrios Xefteris

Department of Economics, University of Cyprus

Dominik Hangartner

*Center of Comparative and International Studies, ETH Zurich
Immigration Policy Lab, Stanford University and ETH Zurich
Department of Government, London School of Economics and Political Science
e-mail: d.hangartner@lse.ac.uk (corresponding author)*

Contents

S1 Introduction	S2
S2 The refugee crisis in context	S2
S2.1 Timing of events	S2
S2.2 Institutional set-up	S2
S2.3 External validity	S3
S3 Data	S3
S3.1 Electoral data	S3
S3.2 Refugee arrivals	S3
S3.3 Distance to Turkish border	S4
S4 Descriptive statistics and Variable coding	S4
S4.1 Descriptive statistics	S4
S4.2 Top coding	S4
S5 Estimation strategies	S5
S5.1 DID analysis	S5
S5.2 IV analysis	S6
S6 Additional results	S7
S6.1 Proximity to refugee hotspots and changes in GD vote shares	S7
S6.2 Visualization of DID estimates	S7
S6.3 Intent-to-treat effect	S7
S6.4 Additional placebo tests	S9
S6.5 Impact on turnout and vote shares of other parties	S10
S6.6 Sensitivity Analysis: Distance From Turkish coast	S14
S6.7 Alternative inference strategies for the DID estimates	S14

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-differences (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 summer 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 advise 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](https://doi.org/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 Sep 2015	500.06	176.39	85	1564
Registered voters Jan 2015	487.93	210.17	70	1838
Valid votes Sep 2015	236.23	84.03	41	439
Valid votes Jan 2015	244.14	104.72	45	470
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

Table S2: Descriptive statistics at the township level.

	Mean	SD	Min	Max
Binary treatment	0.09	0.28	0	1
Registered voters Sep 2015	465.07	171.14	58	1564
Registered voters Jan 2015	462.07	185.91	21	1838
Valid votes Sep 2015	226.98	96.74	31	429
Valid votes Jan 2015	248.15	118.97	13	528
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.

(*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 arrivals. By the end of our 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.

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. To 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-

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).

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

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 (0.031)		-0.568 (0.102)		-0.203 (0.354)
Instrumented refugee arrivals		2.080 (0.478)		0.739 (0.173)	
2012 June					-1.074 (0.869)
2012 January					-0.893 (0.894)
2015 September					2.769 (0.657)
2012 Jun \times Log Distance					0.190 (0.171)
2015 Jan \times Log Distance					0.048 (0.181)
2015 Sep \times Log Distance					-0.481 (0.143)
Constant	1.136 (0.168)	1.295 (0.104)	3.103 (0.552)	1.290 (0.173)	6.574 (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.

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 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 (0.031)		-0.568 (0.102)	
Instrumented Refugee Arrivals		2.083 (0.529)		0.739 (0.202)
Constant	1.136 (0.168)	1.295 (0.103)	3.103 (0.552)	1.290 (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.

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 Aegean 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.

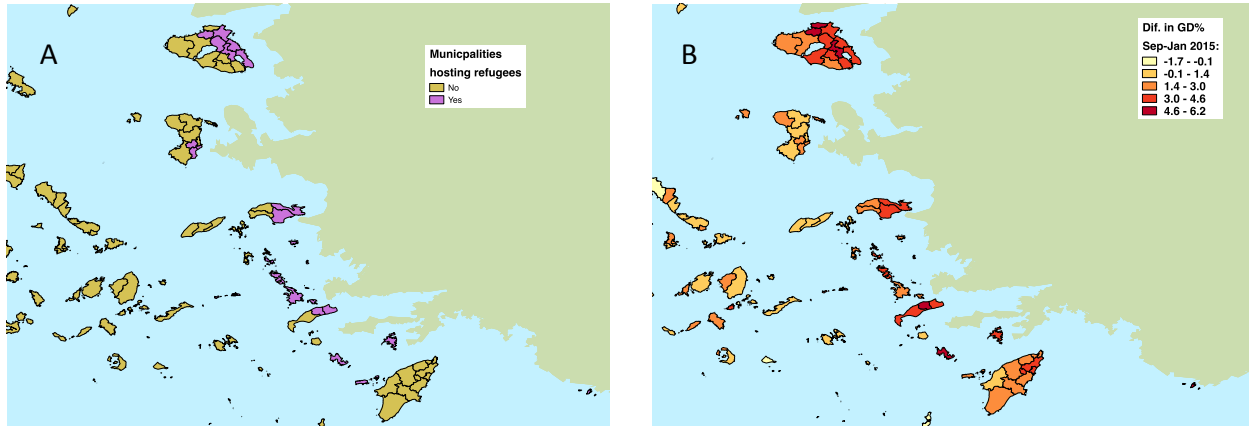



Figure S3: DID estimates of the impact of refugee arrivals on GD vote shares and placebo tests




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.

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 $p < 0.001$); 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.

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.


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 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

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).

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)											
Unit of Analysis	Municipality				Township				Arrivals per capita Municipality			
Outcome	$Party_t$	$Party_t$	$Party_{t-1}$	$Party_t$	$Party_t$	$Party_t$	$Party_{t-1}$	$Party_t$	$Party_t$	$Party_t$	$Party_{t-1}$	$Party_{t-1}$
Nea Dimokratia	-1.643 (1.141)	-4.245 (1.421)	2.294 (1.036)	0.833 (1.971)	-1.273 (0.782)	-4.498 (1.189)	3.322 (0.795)	3.106 (1.515)	-0.428 (0.436)	-1.080 (0.482)	0.567 (0.313)	0.178 (0.753)
SYRIZA	0.597 (1.575)	2.940 (2.313)	-2.132 (1.862)	-1.034 (1.583)	-1.062 (1.092)	2.040 (1.576)	-2.592 (0.868)	-0.377 (1.018)	-0.463 (0.349)	0.588 (0.687)	-0.976 (0.490)	-0.552 (0.610)
PASOK	-0.228 (0.633)	-2.891 (2.471)	2.385 (2.047)	1.012 (2.053)	0.162 (0.423)	-1.376 (1.580)	1.387 (1.311)	0.626 (1.441)	0.045 (0.201)	-1.338 (1.105)	1.315 (0.926)	0.858 (0.723)
ANEL	-1.328 (0.834)	2.005 (1.045)	-2.286 (0.928)	1.759 (1.236)	-0.708 (0.625)	1.598 (0.795)	-1.805 (0.684)	0.248 (0.904)	-0.118 (0.181)	0.668 (0.374)	-0.583 (0.272)	0.227 (0.301)
KKE	-1.355 (0.514)	-1.244 (1.071)	-0.034 (0.811)	0.239 (1.088)	-1.574 (0.380)	-1.043 (0.635)	-0.312 (0.541)	0.508 (0.703)	-0.340 (0.146)	-0.130 (0.265)	-0.231 (0.206)	-0.266 (0.277)
Turnout	2.134 (1.027)	1.420 (1.885)	0.862 (1.224)	1.236 (1.555)	2.082 (0.627)	1.429 (1.236)	0.514 (0.793)	-0.058 (1.025)	0.881 (0.326)	0.680 (0.528)	0.215 (0.300)	0.228 (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
<i>Fixed Effects</i>												
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	<i>Nea Dimokratia</i>		<i>SYRIZA</i>		<i>PASOK</i>		<i>ANEL</i>		<i>KKE</i>		<i>Turnout</i>	
	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 (1.968)	-2.532 (0.776)	5.065 (2.342)	1.736 (0.908)	-1.327 (1.002)	-0.426 (0.352)	0.732 (1.158)	0.294 (0.425)	-1.845 (0.753)	-0.547 (0.310)	5.496 (2.053)	2.054 (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.

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 way to further assess the role of distance to the coast, we use it as a predictor of a series of socioeconomic indicators. These 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)
<i>n</i>	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 (?): 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 1 and 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 shown 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:	GD _(t)	GD _(t-1)	GD _(t)	GD _(t-1)	GD _(t)	GD _(t-1)
Treatment:	Binary treatment		Binary treatment		Arrivals per capita	
Unit:	Municipality		Township		Municipality	
Exposure	2.079 (0.348)	-0.040 (0.340)	2.272 (0.282)	0.093 (0.292)	0.604 (0.179)	-0.004 (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.

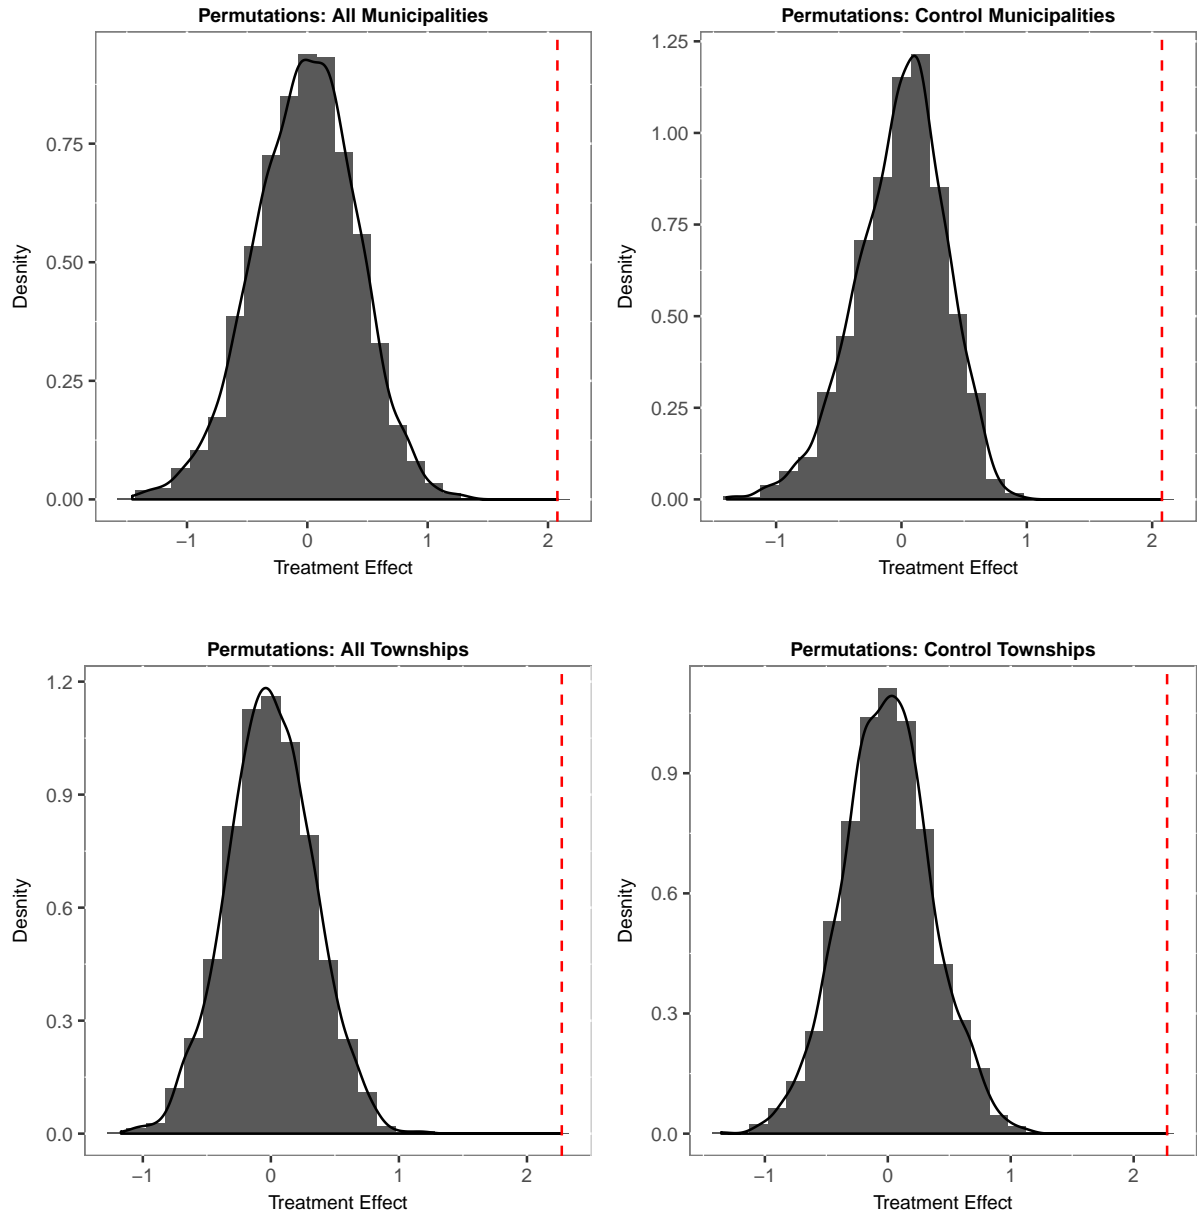
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Outcome:	GD _(t)	GD _(t)	GD _(t)	GD _(t)	GD _(t)	GD _(t)
Treatment:	Binary treatment		Arrivals p.c.	Binary treatment		Arrivals p.c.
Unit:	Municipality	Township	Municipality	Municipality	Township	Municipality
Exposure	2.079 (0.351)	2.272 (0.263)	0.604 (0.178)	2.105 (0.385)	2.210 (0.290)	0.606 (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 exercise 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.