State Repression, Trust, and Political Participation*

Francesco Capozza† Matteo Giugovaz‡ Giuseppe Musillo§

October 22, 2026

Abstract

How does violent state repression affect political participation in a mature democracy? We study the 2001 G8 summit in Genova, Italy, where police attacked authorized protests and committed abuses that courts later ruled met the international definition of torture. We measure each municipality's exposure by its proximity to the associations that organized travel to Genova. Exposure lowers turnout in national elections by 0.62 percentage points per standard deviation. The gradient is still visible in 2022, two decades after the summit. The effect is withdrawal, not realignment: it raises abstention among former voters, first on the left, and does not show up in petitions, demonstrations, or local and European elections. Two patterns point to a repression narrative that spread by word of mouth against a competing official account. The decline reaches one percentage point in rural municipalities and persists through 2013, while the urban effect fades to zero, and it shrinks where the reception of the mainstream television signal was stronger. Trust falls for the central institutions of the national state, and rises for interpersonal trust and local belonging. Repression cut the tie between citizens and the state, not the ties among citizens.

Keywords: State Repression, Political Participation, Trust, Trials, Narrative Competitions, Media
JEL Classification: C83, D72, D91, K42

*We are grateful to Eleonora Alabrese, Andreu Arenas, Mathias Bühler, Davide Cipullo, Tommaso Colussi, Simone Cremaschi, Gianmarco Daniele, Ruben Durante, Ruben Enikolopov, Elisa Facchetti, Francesco Ferlenga, Mery Ferrando, Matteo Gamalerio, Albrecht Glitz, Krzysztof Krakowski, Nicola Mastrorocco, Arianna Ornaghi, Salvo Nunnari, Maria Petrova, Mounu Prem, Massimo Pulejo, Marta Reynal-Querol, Anna Salomons, Sulin Sardoschau, Antonio Schiavone, David Schindler, Piero Stanig, Andrea Tesei, Miriam Venturini and seminar participants at CLEAN, Bologna (Workshop on "Social Economy" for Young Economists), BSE PhD 2025 Jamboree, 30th ENTER Jamboree, Tilburg, Universitat Carlos III de Madrid, Universitat Pompeu Fabra, Université Paris Dauphine-PSL, University of Essex for their helpful comments. The project has received an IRB approval by WZB with the number 293/2025. Musillo aknowledges financial support from CentER. The authors have no conflicts of interest to declare.

†UB, IEB, CESifo, and IZA@LISER. Email: francesco.capozza@ub.edu
‡Universitat Pompeu Fabra. Email: matteo.giugovaz@upf.edu
§Tilburg University. Email: g.musillo@tilburguniversity.edu

1 Introduction

How does state repression change political participation and trust in institutions? This question is especially important in a democracy, where elections, courts, and legal protections offer ways to challenge abuses by the state. These institutions may sustain participation by allowing citizens to replace a government or obtain recognition of a wrong. Yet the availability of those remedies does not ensure that citizens continue to place confidence in institutions. An abuse committed by state agents can raise doubts about institutions beyond those immediately responsible, while shared experiences may strengthen ties among citizens. To understand these responses in a democratic setting, we need to distinguish a change in political preferences from withdrawal from national institutions, and both from a broader retreat from political and social life.

V-Dem's 2026 Democracy Report finds that global levels of democracy have fallen back to where they stood in 1978 (V-Dem Institute, 2026). Repression is central to how this happens. The CIVICUS Monitor records the detention of protesters as the single most common civic-space violation in 2025, documented in at least 76 countries, with excessive force against protesters documented in at least 67 (CIVICUS, 2025). Crucially, this is no longer confined to fragile or transitional regimes. Autocratization now affects seven EU member states, and civic space has deteriorated in consolidated democracies.[1] Yet state repression is usually studied in autocracies or civil conflicts, and the large literature asking whether it deters political participation or provokes backlash, finding evidence of both, draws mostly on such settings (Davenport, 2007a; Lawrence, 2017; Young, 2019).

This paper studies police violence during the July 2001 G8 summit in Genova, Italy, and its consequences for participation and trust. Roughly 200,000 people travelled to the city to protest, with many relying on associations belonging to the Genoa Social Forum (GSF) to organize transport and mobilization. Police charged an authorized section of the demonstrations, assaulted activists and journalists at the Diaz school, and abused detainees at the Bolzaneto barracks. Subsequent judgments established serious violations, including treatment amounting to torture. We use the geography of the associations that organized participation to examine how political behavior subsequently changed in communities connected to the protests.

The empirical challenge is that only a small share of Italians were physically present in Genova. We exploit the logistics of participation to measure which places were more closely connected to the event. Travel was organized largely through GSF-affiliated associations, which chartered coaches and trains and served as local nodes for mobilization and information. This infrastructure mattered especially in 2001, when internet use remained limited, broadband penetration was negligible, and social media did not exist (Interna-

[1] France, for instance, was downgraded in 2025 after repeated use of tear gas, explosive grenades and mass arrests against peaceful demonstrations (CIVICUS, 2025).

tional Telecommunication Union, 2002). [2] We digitize the 2001 location of the GSF associations and construct a market-access measure for every Italian municipality, equal to the inverse-distance-weighted number of associations within 100 kilometers (Ahlfeldt et al., 2015). Municipalities surrounded by more GSF associations were more likely to be connected to people who travelled to Genova and to the accounts they brought back.

We combine this measure with municipal election returns from 1983 to 2013 and individual survey data from ITANES between 1994 and 2008. Our main design compares changes in national-election turnout across municipalities with different levels of exposure, using municipality fixed effects, province-by-year fixed effects, and a rich set of 2001 characteristics interacted with election-year indicators. Throughout the baseline analysis we exclude municipalities that directly hosted a GSF association, so identification comes from differential proximity to the network rather than from the locations in which the associations themselves were based.

Exposure predicts a persistent decline in national-election turnout. The pre-2001 event-study coefficients are small and close to zero. Beginning with the first election after the summit, turnout falls and remains lower through 2013. Averaged across the three post-summit elections, a one-standard-deviation increase in exposure lowers turnout by 0.62 percentage points, relative to a pre-period mean of 82.8%. The decline reaches 0.80 percentage points in 2008. Extending the panel beyond our baseline window, the gradient is still negative in 2018 and 2022. A median split gives the same qualitative result without imposing linearity in the dose: municipalities above the median of exposure lose about 2.5 percentage points of turnout relative to those below it at each post-2001 election.

The effect is withdrawal from national voting rather than a shift across political blocs. In ITANES, exposure does not systematically change the probability of voting for the right rather than the left conditional on voting. Instead, former voters become more likely to abstain. The response appears first among prior left-wing voters, whose abstention rises by 5.3 percentage points in 2006, and later more weakly among prior right-wing voters. This ordering matches the composition of the counter-summit, whose participants were overwhelmingly on the left. At the same time, exposure has no detectable effect on signing referendum petitions, attending demonstrations, or reported interest in politics, and neither municipal nor European elections display a post-2001 break comparable to the national elections. The response is therefore specific to national electoral participation rather than a general retreat from politics. When an explicitly anti-establishment party, the Five Star Movement, entered national politics in 2013, exposed municipalities gave it somewhat more support among those who still voted, most clearly in 2018 and not at all by 2022, once the party had governed.[3]

[2] Contemporary survey evidence shows the same pattern: 36% of participants first heard about the counter-summit through their own organization, compared with only 4.2% through the internet (Andretta et al., 2002).

[3] We treat this as suggestive, since the estimates are cross-sectional.

Adverse political experiences are often studied as shocks to beliefs and preferences (Alesina and Fuchs-Schündeln, 2007; Di Tella et al., 2007; Nunn and Wantchekon, 2011; Malmendier and Nagel, 2011; Fuchs-Schündeln and Schündeln, 2015), while work on violence and social capital shows that such shocks can either strengthen or destroy local cooperation depending on the social boundary the violence creates (Bauer et al., 2016; Hager et al., 2019). Our evidence points to a reallocation rather than a uniform decline. Exposure is associated with lower confidence in central institutions of the national state, particularly the President of the Republic, the judiciary, and the army, while generalized interpersonal trust, trust in trade unions, and local belonging rise. Trust in the European Union is unchanged. The pattern is close to the distinction in Scott (1985) between withdrawal from institutions of domination and continued cooperation within local social relations: what changes is not the level of trust but the institutions toward which it is directed. We treat these estimates as suggestive evidence on mechanism, because the relevant ITANES modules are available only in 2001 and 2006 and some individual items are sensitive to coding.

The second part of the paper asks how an event concentrated on a small group generated a persistent geographic pattern. Our interpretation centers on competing narratives. Protesters returned through a dense associational network with first-hand accounts of police violence, while national media and the incoming center-right government initially emphasized violent protesters and justified the police response. Much of the repression literature treats proximity to an event as equivalent to knowledge of it, which seems unlikely to hold in general and did not hold here: what reached most Italian households was not the crackdown but one of two competing accounts of it. A corpus of 569 press and wire items shows that protester blame dominated early coverage, while police blame rose during the trials and became dominant after the 2008 convictions. The turnout effect is strongest and most persistent in rural municipalities, where face-to-face ties are denser, and weaker where terrain-driven variation produced stronger reception of the mainstream Mediaset signal. The media specification follows the logic of work using topographic variation in broadcast reception to isolate differential media exposure (Olken, 2009; Enikolopov et al., 2011; Durante et al., 2019). We read this evidence as corroborating a narrative-diffusion mechanism but it suggests that the informational environment in which a repressive event is learned about may be one reason why estimated responses to repression differ so much across settings.

Several exercises probe alternative explanations. The turnout result survives alternative fixed effects, a relevance-weighted treatment, a median split of exposure, the inclusion of association-hosting municipalities, leave-one-province-out and leave-one-region-out estimation, placebo treatment dates, and the spatial-noise diagnostics of Conley and Kelly (2025). Province-level clustering substantially reduces precision. The continuous 50-kilometer treatment also becomes weak, while the corresponding median split retains a

large and persistent effect. Exposure is correlated with observable municipal characteristics, including baseline partisanship, which is why each is interacted with election-year indicators; the absence of pre-2001 trends is the main diagnostic that these differences do not drive the post-summit change.

Contribution to the Literature We contribute to five strands of research.

First, we bring the study of state repression, long treated as a feature of nondemocratic regimes (Davenport, 2007a; Davenport and Armstrong, 2004), inside a mature democracy, and show that its political legacy there has a particular shape. Most causal evidence on policing and repression comes from developing countries and young democracies (González and Prem, 2024; Bühler and Madestam, 2023; Arjona et al., 2025; Bautista et al., 2023), or from protest dynamics under authoritarian rule (Cantoni et al., 2019, 2022; Hager and Krakowski, 2022).[4] In a consolidated democracy a domestic crackdown produces persistent demobilization without any detectable left-right realignment, backlash, or prejudice contagion, and it does so at the ballot box specifically. We can draw that distinction because we observe participation at three levels of government, in several arenas, and attitudes toward ten institutions.

Second, we add to work on how lived experience shapes beliefs and preferences (Alesina and Fuchs-Schündeln, 2007; Di Tella et al., 2007; Nunn and Wantchekon, 2011; Malmendier and Nagel, 2011; Fuchs-Schündeln and Schündeln, 2015). One episode of state repression, we show, durably reshapes trust in the state and political behaviour.

Third, we speak to the determinants of social capital (Guiso et al., 2004, 2008, 2016; Nannicini et al., 2013; Bigoni et al., 2016, 2019; Campante et al., 2022; Buggle and Durante, 2021; Durante et al., 2025). Where much of that literature treats bad experiences as eroding trust broadly, we document a reallocation: vertical trust in the national state falls while generalized trust, union trust and local belonging rise. Law enforcement emerges not as a force that raises or lowers social capital as a whole, but as one that moves it between the vertical and horizontal margins.

Fourth, we contribute to a growing body of work showing that citizens' encounters with the coercive arm of the state shape political trust and participation (Ang et al., 2025; Facchetti, 2025; González and Prem, 2024; Canini et al., 2026; Mummolo, 2018). Our addition is that a single salient episode durably depressed national electoral participation, and that it did so through an erosion of trust organized along the state/non-state boundary rather than through general disengagement.

Finally, we connect to work on the long-run legacies of mass repression in authoritarian

[4]For a review of the protest literature, see Cantoni et al. (2024). Unlike Enikolopov et al. (2020), we study a mobilization that was not coordinated through social media. We also connect to work on how protest tactics and their repression shape public opinion and voting (Wasow, 2020). Our contribution there is to trace, past the contemporaneous shift in opinion, the electoral withdrawal that follows once the repression frame prevails.

and post-authoritarian settings: Stalinist terror and lower turnout decades later (Zhukov and Talibova, 2018), lower institutional trust (Nikolova et al., 2022), durable opposition to the perpetrators' heirs (Rozenas et al., 2017), loyalty contingent on the threat of retribution (Rozenas and Zhukov, 2019), and parallel findings after China's Cultural Revolution (Wang, 2021), alongside the literatures on wartime and civil-war violence (Cremaschi and Masullo, 2025; Komisarchik et al., 2022; Harada et al., 2024; Fouka and Voth, 2023; Rodon and Tormos, 2023; Balcells, 2012; Tur-Prats and Valencia Caicedo, 2025; Barceló, 2021; Vásquez-Cortés, 2024) Evidence from Pinochet's Chile points both ways at once, as ours does: local exposure to regime violence raised opposition to the dictatorship without lasting post-transition shifts in preferences (Bautista et al., 2023), and the repressed left parties and unions without changing ideology (Bautista, 2015). Our value added is to separate a channel those settings cannot. Where the state may still retaliate, withdrawal, fear and broad institutional distrust are hard to tell apart, and repression often corrodes trust across the board (Nunn and Wantchekon, 2011; Daniele et al., 2025). In a democracy with free elections, a free press, and courts that eventually convicted the perpetrators, the legacy is neither general cynicism nor ideological realignment, but a cut vertical tie beside intact and even stronger horizontal ones.

The closest papers to ours are Rodon and Guinjoan (2022) and Balcells et al. (2021), on police violence and participation in the 2017 Catalan referendum. We differ in tracing the persistent electoral legacy of an unexpected crackdown on a peaceful protest, an effect we follow through 2022, and in substantiating the mechanism with survey, electoral, media and trial data rather than the contemporaneous response alone. Moreover, we provide evidence of the channels, namely direct interactions and media, that propagate the effects.

The rest of the paper proceeds as follows. Section 2 gives the background of the G8 and, drawing on the trial record, the human imprint of the repression. Section 3 describes the data and Section 3 the empirical strategy. Section 4 presents the effects on electoral outcomes and on trust. Section 5 traces how the repression narrative diffused through the press, face-to-face networks and the competing television signal. Section 6 reports the robustness checks and Section 7 concludes.

2 The Setting

Three features of the Genova summit matter for everything that follows. The violence was the state's, not a few officers'. The people who met it had been brought to Genova by a dispersed network of associations, and they went home through the same network. And what they experienced left them with an account that was both painful to carry and easy to repeat. This section establishes each in turn, keeping to what the empirical work needs; Appendix C gives the fuller history, from the origins of the anti-globalization movement to the verdict of each trial.

For readers outside Italy, a word on how large this event looms there. Genova is not a half-remembered episode of public disorder: it is a reference point in Italian political memory, and it has been revisited on screen for a quarter of a century.[5] Our argument turns on the claim that Genova entered and stayed in collective memory.

2.1 The Summit and the Crackdown

The G8 met in Genova on 20–22 July 2001. Alongside it, a coalition of Italian and international associations, the Genoa Social Forum (GSF), organised a *counter-summit* whose demonstrations and routes had been agreed in advance with the police and the government's local representative.[6] More than 200,000 people were in the city by 19 July.

On 20 July the police charged an authorised and peaceful section of the march, and an officer shot and killed a protester, Carlo Giuliani. On the night of 21 July the police raided the Diaz school, where activists and journalists were sleeping, and assaulted its 93 occupants; many were then taken to the Bolzaneto barracks, where the abuse continued. Figure 1 shows two moments from those three days.

2.2 State Repression, Not Individual Excess

If Genova were a matter of individual officers losing control, a loss of trust in the state as an institution would be a strange response to it. If it was the state acting as an institution, that response is the natural one. Three independent sources point to the second reading.

The police accounts describe a planned operation. We coded what officers and defendants said they had anticipated (Table A.7): 85% of police hearings describe acting on orders from above and 74% invoke intelligence or radio reports of an imminent threat, while roughly four in ten anticipate organised violent infiltrators and a third expect an armed or masked crowd.

The victim accounts contradict the premise of that justification. The people in the school were unarmed and were beaten while seated, lying down or already on the ground (Table A.5), and the Bolzaneto detainees describe prolonged mistreatment in custody rather than force used to subdue a violent crowd.

[5] Francesca Comencini's *Carlo Giuliani, ragazzo* (2002) was built around the testimony of Giuliani's mother (IMDb). Carlo A. Bachschmidt's documentary *Black Block* (2011) gives the floor to seven foreign demonstrators who survived the Diaz raid and Bolzaneto and traces how they lived with it afterwards (IMDb). Bachschmidt is also one of the GSF logistical coordinators we interviewed for Section 2.3. Daniele Vicari's feature film *Diaz — Don't Clean Up This Blood* (2012) reconstructs the raid from the perspectives of protesters, journalists and police (IMDb); Franco Fracassi and Massimo Lauria's *The Summit* (2012) covers the summit and the proceedings that followed. A catalogued archive of footage from 19–22 July, organised by day and location, is available at documentafilm.com.

[6] The daily programme is archived at processig8.net. The GSF was founded in December 1999 and comprised 1,187 associations by February 2001, of which 1,009 were Italy-based: cultural associations, parties, unions, religious movements, student groups, NGOs and single-issue organisations.

Figure 1: Pictures from Genova's Countersummit



(a) Riots in Genova



(b) Raid at Díaz School

The courts reached the same conclusion. The Bolzaneto judges held that the officers' conduct met the international definition of torture and amounted to a serious suspension of the rule of law. The Diaz verdict convicted officers and senior officials of aggravated assault, injury, slander, forgery and fabrication of evidence. The European Court of Human Rights condemned Italy and required it to introduce the crime of torture into domestic law.[7]

Two features of how the courts handled Genova matter later. The trials ran for more than a decade, so the event kept returning to public attention rather than fading. And, as the ECHR noted, nobody responsible for the Diaz or Bolzaneto abuses served a custodial sentence, because the officers could not be identified and the offences lapsed; several senior officials kept their posts and were later promoted. The courts vindicated the victims without punishing anyone, which is relevant when we find in Section 4.4 that trust in the judiciary falls too.

2.3 The Network That Took People to Genova

How people reached Genova is the basis of our identification, so we describe it in some detail. The airport and port were closed, only designated stations were open, and road access was controlled by the police. Participants arrived by chartered coaches, or by special trains.

Most came on coaches and trains organised by GSF-affiliated associations. The GSF's logistical coordinators told us this was true even of participants who belonged to no association,[8] and the trial record says the same thing in the participants' own words. Victims describe travelling on coaches chartered from their province by the unions, by Rifondazione and by the Democratic Left; leaving from Imola on one of the special trains; or joining an organised group that gathered at Bologna station. These details are incidental to the proceedings, which is part of why we find them credible.

A survey of 763 participants conducted during the protest (Andretta et al., 2002) confirms the picture. 41.2% were formal members of a participating association, and most had first heard of the counter-summit through their own organisation (36%), well ahead of family and friends (23%), the mass media (22.7%) and the internet (4.2%). The network was therefore both the means of getting to Genova and the main channel of information about it. The participants were also politically homogeneous: their left–right self-placement was overwhelmingly on the left and centre-left, with negligible shares on the centre or right.[9]

[7] AFFAIRE CESTARO c. ITALIE (Requête no 6884/11). The charges against the peaceful section of the march were also ruled illegitimate. The officer who shot Giuliani was cleared on grounds of legitimate defence, and of the 25 protesters tried for devastation 15 were acquitted and 10 convicted. For a full account of the proceedings see Agnoletto and Guadagnucci (2011).

[8] Vittorio Agnoletto, Carlo Bachschmidt, Raffaella Bolini and Stefano Kovac.

[9] 39.5% came from the Attac/Fiom/Rifondazione/ARCI area, 29.9% from the Lilliput network, 24.5%

2.4 What Victims Brought Home: The Trial Record

The depositions of the four first-instance trials record what the repression did to the people who lived through it, and how they spoke about it afterwards.

The abuse was severe, and different at each site (Table A.5). The Diaz raid was an acute mass beating: batons appear in 93% of its victim hearings, and victims describe being beaten while already on the ground in 73%. Bolzaneto was custodial abuse: forced stress positions appear in 93% of its hearings, handcuffing in 87% and forced nudity in 82%. What victims carried afterwards is recorded in Table A.6. Persistent fear appears in 77% of hearings, and substantial shares describe panic and anxiety, intrusive memories, and a lasting loss of trust. What these people describe losing is confidence in the state that beat them, not in one another.

Figure 2 scores victim, police and forensic testimony on an emotion-versus-reason axis built from word embeddings trained on the corpus, following Gennaro and Ash (2022), and on a rate of hedging. Victim testimony is the most emotive (-0.18 on the standardised affect–cognition axis), forensic testimony the most cognitive (-0.64), and police testimony lies in between (-0.29). Victims also hedge far more: 9.79 uncertainty markers per thousand words, against 6.67 for police and 2.47 for forensic witnesses.[10] The point is not which side is more reliable. It is that the victims' accounts are first-person, emotional and hard to forget, which is the kind of account people repeat to those around them.

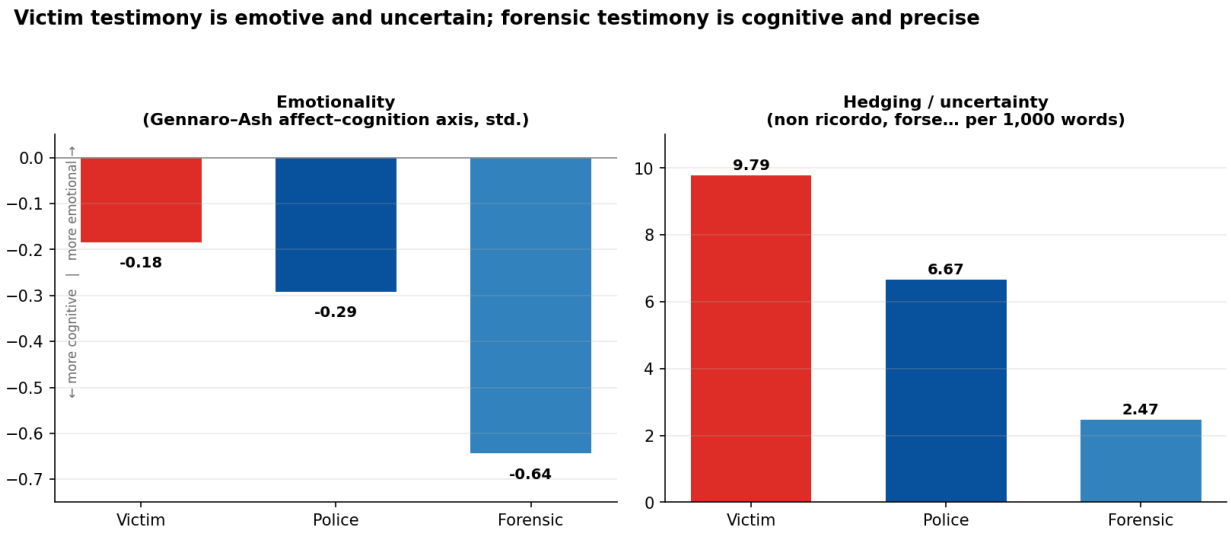
2.5 Protest After Genova

Genova did not end political contention in Italy. Figure B.1 plots Italian protest from 2000 to 2015 using event-level data from the POLDEM Protest dataset, split by the three actors most central to the counter-summit coalition and by the radicalism of each event. Student-led events all but disappear between 2002 and 2007, and when student protest returns from 2008 it is tied to the university-reform cycle rather than to the alter-globalization coalition. Aggregate protest does not collapse, but it shifts toward institutionalised labour action: unions account for 283 weighted events over the period, of which 2% are coded violent. Nor did the radical fringe grow: the violent share of student events falls to essentially zero between 2001 and 2004. These series are national and carry no variation in exposure, so we use them only as context.

from the *centri sociali*, and 3.8% from local organisations. Women made up 42.8% of the sample and 96% were Italian; 42.4% came from the North and 35.3% from the Centre. Students were the largest occupational group (49.9%), followed by salaried employees (12.6%) and the unemployed and precariously employed (10.7%).

[10] We read the hedging gap as the difference between traumatic recall and the reading of physical evidence, not as evasiveness.

Figure 2: Emotions, Cognition, and Hedging in Trial's Depositions



Note. The figure compares victim, police, and forensic testimony from the Genova trials (Diaz, Bolzaneto, and the proc-25 proceedings, while Piazza Alimonda's trial is not considered because it contains exclusively police's hearings) on two dimensions of speech, computed over speaker-attributed deponent turns (foreign witnesses' testimony is taken from the interpreter's rendering). The left panel reports *emotionality*, operationalising the emotion-versus-reason measure of Gennaro and Ash (2022): we train word embeddings on the trial corpus (1.83 million tokens), define an affect pole and a cognition pole from Italian seed words (e.g. *paura, terrore, dolore, panico* vs. *accertato, rilevato, misura, dunque*), and score each word by the difference between its cosine similarity to the two poles; a side's emotionality is the standardised mean of this score over its words. Higher values denote more affective, lower values more cognitive language; the data-driven axis places *paura, panico, dolore, urla* at the emotion end and *denominato, eseguito, risultavano, corrispondenza* at the cognition end. The right panel reports *hedging*, the rate of uncertainty markers (*non ricordo, forse, mi sembra*) per 1,000 words.

3 Data and Empirical Strategy

The empirical strategy follows directly from the setting. If the associations that organised travel were also the channels through which accounts of Genova came home, then living near many of them should predict both a higher chance of knowing someone who was there and a larger political response afterwards. This section describes the data and turns that idea into a measure.

3.1 Data

The GSF associations. We digitised the list of Italian associations that signed the GSF manifesto before the counter-summit and recovered their 2001 locations from archived web pages on the Wayback Machine. We locate 863 of the 1,009 Italian associations, or 85.5%. For each we also record whether it sat on the GSF national board, whether it belonged to a nationwide organisation, and whether it had a website in 2001, which we use for a relevance-weighted variant of the treatment.

Elections. Municipality-level turnout and party results for every national election from 1983 to 2013 come from the Italian Ministry of the Interior, as do the corresponding series for European Parliament elections (1984–2014) and municipal (*comunali*) elections

(1985–2015), which we use in Section 4.3.

Our baseline national panel ends in 2013. That election saw the entry of the Five Star Movement, which took roughly a quarter of the vote at its first national contest and reorganised the party system around a new anti-establishment pole. After 2013 a disaffected voter had a party to vote for as well as the option of staying home, so abstention means something different, and comparing turnout across the two regimes would mix two things. We keep 2013 because it completes the post-summit sequence, but read it more cautiously than 2006 and 2008. Section 6 extends the panel to 2022 and shows that the conclusion on persistence does not depend on where we stop.

Survey data. The Italian National Election Study (ITANES) provides individual-level turnout, vote choice, previous vote, trust and sense of local belonging for every national election year from 1994 to 2008 (Daniele et al., 2023). Its trust battery covers the police, the army, the President of the Republic, parliament, political parties, the judiciary, trade unions, the European Union and the Church, together with the standard generalized-trust item.

We use the two sources for different things. Turnout and vote shares are recorded for every municipality at every election, so the official returns dominate the survey on coverage and we use them wherever they can answer the question. ITANES is reserved for what no aggregate source contains: how an individual voted conditional on voting, whether a respondent who voted last time abstains this time, trust in specific institutions, and participation outside elections. The transition outcomes in particular require observing the same person's previous vote. The cost is coverage, since ITANES samples many small municipalities thinly or not at all.

Municipal characteristics. Demographic and socio-economic characteristics come from the 1991 and 2001 censuses. The Mediaset signal, free-space signal and terrain variables come from Durante et al. (2019). Tables A.1 and A.2 report summary statistics for the baseline estimation samples.

3.2 Measuring Exposure

Think of someone living in a small town in Emilia in July 2001 who wanted to go to Genova. They would almost certainly have gone the way the witnesses describe, on a coach chartered by a local union branch or on a special train out of Imola or Bologna, and they would have heard about the trip through the same organisations. Whether they went, and whether a neighbour or a relative went and came back with a story, depended on how many such organisations were within reach. Online channels were no substitute in 2001: roughly 27% of Italians used the internet at all, broadband stood

below one subscription per 100 inhabitants, and social media did not exist (International Telecommunication Union, 2002).

We capture this with a market-access index (Ahlfeldt et al., 2015). For each municipality we count the GSF associations within 100 km, each weighted by the inverse of its road distance:

$$\text{exposure}_m = \sum_{j=1}^{J} \left(\frac{1}{d_j} \text{asso}_j \right) \tag{1}$$

where d_j is the road distance between municipality m and association j, and asso_j equals one for each association. The index rises with both the number of nearby associations and their proximity. Figure 3 maps it. On the baseline sample it has a mean of 1.28, a standard deviation of 1.21, a median of 0.95 and a maximum of 9.64 (Table A.1). It is strongly right-skewed, and every municipality has at least one association within 100 km, so there is no untreated mass at zero.[11] We standardise the index on each estimation sample after applying the sample restrictions, so every coefficient reads as the effect of a one-standard-deviation increase in exposure within the sample that produced it.

Three design choices follow. First, we exclude municipalities that themselves host a GSF association. What we want to compare is proximity to the network, not the qualitatively different case of living where an organisation is based; Section 6 brings these municipalities back.

Second, the index is smooth over space by construction, and the accounts we argue carried the effect do not stop at municipal borders either. Less-exposed municipalities in the same province are therefore partly treated. Because our specifications include province-by-year fixed effects, the coefficient comes from comparing more- and less-exposed municipalities within a province. If the story also reached the less-exposed ones, this comparison understates the total effect of the crackdown.

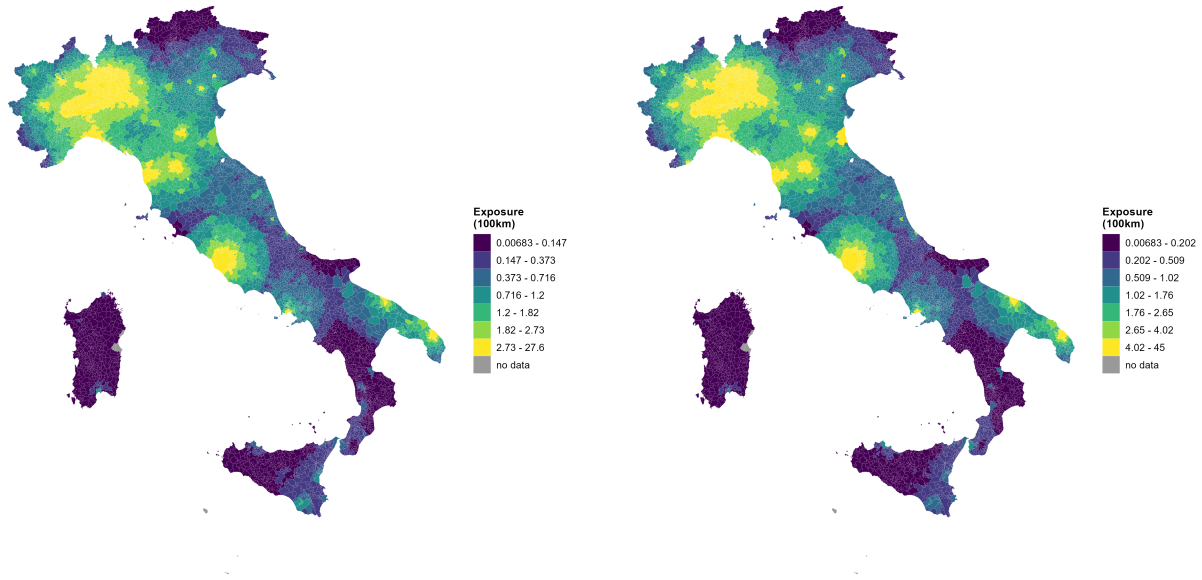
Third, the baseline treats every association as equally relevant. As a check, we weight associations by how distinctive they were within the counter-summit network, using three traits: sitting on the GSF national board, belonging to a nationwide organisation, and having a website in 2001. Rare traits are more informative about an association's prominence than common ones, so each trait t with sample frequency p_t receives weight $-\log p_t$, its Shannon self-information, and an association's score is the sum over the traits it holds:

$$\text{asso}_j = \sum_t \text{trait}_{tj} \left(-\log p_t \right).$$

An association with national reach and a website thus scores higher than one with neither. This score replaces the indicator in Equation 1 in the robustness analysis.

[11] This is a property of the 100 km radius. At 50 km, 8.2% of municipalities have no association in range, and at 20 km, 39.9% do; "no association within reach" is a different statement from "little exposure", which is one reason we take 100 km as the baseline. Section 6 reports the 50 km specification in both continuous and binarised form.

Figure 3: Market Access Exposure Index at the Municipality Level, 2001



(a) Without Association Hosts index (b) With Association Hosts index

Note. Distribution of the market-access exposure measure at the municipality level in Italy in 2001, at the 100 km radius, without association host municipalities (panel a) and with association host municipalities (panel b) variants. Lighter colours indicate higher exposure to Genova G8 state repression, darker colours lower exposure. Classes are quantiles of the raw measure.

3.3 Municipal Specification

The municipal panel covers the nine national elections from 1983 to 2013 for all Italian municipalities, roughly 8,000 per election. After excluding association-hosting municipalities and observations with missing covariates, the baseline sample contains 69,504 municipality-election observations. We estimate

$$Y_{mt} = \beta(\text{Exposure}_m \times \text{Year}_t) + \gamma(\text{Z}_m \times \text{Year}_t) + \rho_{tp} + \mu_m + \epsilon_{mt}, \tag{2}$$

where Y_{mt} is the outcome in municipality m and election t; $\text{Exposure}_m \times \text{Year}_t$ interacts standardised exposure with election-year indicators, omitting 2001, the last election before the summit; Z_m contains 2001 population density, gender balance, mean age, the share aged 75 or older, the shares with a diploma and with a degree, the shares employed and unemployed, the shares employed in agriculture and in industry, historical Pentapartito and PCI vote shares, television signal, latitude and longitude; and ρ_{tp} and μ_m are province-by-year and municipality fixed effects. Because every characteristic in Z_m is interacted with election-year indicators, municipalities that differ at baseline are allowed to follow different paths in every election.[12] Standard errors are clustered by municipality.

[12]All controls are measured in 2001 and are time-invariant, so they identify nothing until interacted with the time term. Subsample splits likewise use 2001 values, so that no municipality moves between

We report each specification in two forms. The event-study form traces the effect election by election and shows the pre-period. The pooled form replaces the year indicators with $\text{Post}_t = \mathbf{1}[t > 2001]$ and summarises the post-summit path in a single coefficient. Where useful we also report the cumulative effect, the sum of the three post-2001 coefficients, with its standard error computed from the full variance–covariance matrix.

3.4 Individual Specification

For the ITANES outcomes we use a repeated cross-section of the five national election years from 1994 to 2008, with roughly 2,000 usable respondents per year, and estimate

$$Y_{imt} = \beta(\text{Exposure}_m \times \text{Year}_t) + \gamma W_{imt} + \rho_{tp} + \mu_m + \epsilon_{imt}, \tag{3}$$

where i indexes respondents, W_{imt} contains gender, age, education, employment status and religious orientation, and the fixed effects and clustering are as in the municipal specification. The trust battery is fielded only in 2001 and 2006, so for trust we can estimate the pooled form but not the event study: the specification and identifying assumption are unchanged, but we cannot display for trust the flat pre-period we display for turnout. We return to this in Section 4.4.

4 How Citizens Responded

The introduction set out three ways citizens might have responded to the crackdown, and each makes different predictions. If exposed citizens *realigned*, turnout need not fall, but the balance between left and right should move. If they *disengaged* from politics altogether, turnout should fall together with every other form of participation, in every kind of election. If they *withdrew from the national state*, turnout in national elections should fall while the left–right vote, participation outside elections, and local and European elections stay where they were, and trust should fall for the institutions of the national state but not for others. This section takes these in turn. We first establish the decline in national turnout, then show that it is neither realignment nor general disengagement, then turn to the attitudes that accompany it, and finally look at where some of that support went once an anti-establishment party existed.

4.1 National Turnout

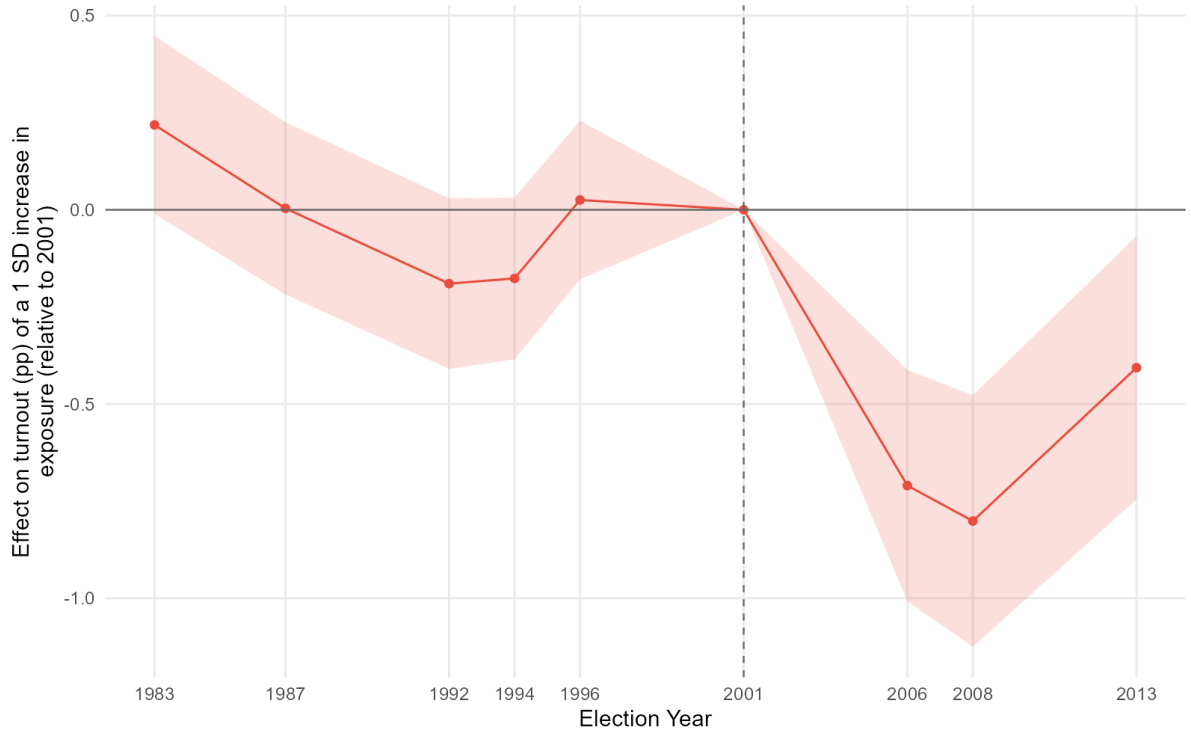
Figure 4 plots the event-study coefficients for turnout in national elections, relative to 2001, the last election before the summit. Before 2001 the coefficients are small and none is significant at the 5% level; all lie within a quarter of a percentage point of the

groups over the panel.

reference year. From the first election after the summit they turn negative and stay there: a one-standard-deviation increase in exposure lowers turnout by 0.71 percentage points in 2006, 0.80 in 2008 and 0.41 in 2013. Pooling the post-summit elections, the effect is 0.62 percentage points (s.e. 0.135). It also lasts: extending the panel beyond our baseline window, the gradient is still negative in 2018 and 2022 (Section 6).

The same pattern is visible without a model. Figure B.2 in Appendix B plots raw turnout for municipalities above and below the median of exposure. The two groups sit at different levels, since exposure is correlated with the characteristics the regression conditions on, so the levels themselves are not the point. The gap between them is: it is stable at 11 to 13 percentage points across the three elections before the summit and closes to about 5 points at both elections after it.[13]

Figure 4: Effect of Exposure on Turnout



Note. The figure plots the event study coefficients from a fixed effects regression of voter turnout on exposure. The coefficients represent the interaction terms between year indicators and the standardised exposure measure, with 2001 serving as the reference year. The shaded areas represent 95% confidence intervals with standard errors clustered at the municipality level. The regression includes municipality fixed effects, province-by-year fixed effects, and a comprehensive set of time-invariant 2001 controls interacted with each election year: density, gender balance, mean age, share age +75, share of population with diploma, share of population with bachelor, share employed, share unemployed, employment in agriculture, employment in industry, historical political leaning, TV signal, latitude and longitude. Municipalities hosting at least one GSF association are excluded. $N = 69{,}504$.

A persistent decline of this kind is what the repression literature would describe as

[13]Two caveats apply, since the figure is unconditional. The gap closes mostly because the below-median group rises in 2006 rather than because the above-median group falls, and a raw plot cannot separate that from any other post-2001 shock that hit the two groups differently. And 1994 is a level break in Italian electoral history, so we would not read the 1994–1996 segment as a trend. We treat the figure as a check that the divergence is visible in the raw data and begins when the event does, not as evidence in its own right.

deterrence: coercive policing raises the perceived cost of engaging and lowers the perceived legitimacy of the institutions one would engage with (della Porta, 1995; Davenport, 2007b; Tilly, 1978), and salient episodes that keep being recalled can have durable effects (González and Prem, 2024). The partial recovery in 2013, where the coefficient roughly halves, is worth flagging now because Section 5 explains it: in rural municipalities, where we argue the memory of Genova took hold, there is no recovery at all.

How large is the effect? Mean turnout on the estimation sample is 82.8%, with a standard deviation of 9.9 points (Table A.1), so a 0.62 point decline is about 0.75% of the mean, or 0.06 standard deviations. The number is small because the crackdown directly touched a small group of people; the question is how far a shock concentrated on that group travelled. A more intuitive comparison splits municipalities at the median of exposure: the more-exposed half loses about 2.5 points of turnout relative to the less-exposed half at every post-summit election (Figure B.12), a comparison that imposes no linearity in the dose.

Two back-of-the-envelope calculations put the magnitude in terms of people. The first translates exposure into participants. Allocating the roughly 200,000 participants (Andretta et al., 2002) across municipalities in proportion to each one's share of total exposure,

$$\text{Participants}_m \equiv 200{,}000 \times \frac{\text{Exposure}_m}{\sum_{m'} \text{Exposure}_{m'}},$$

the pooled coefficient implies that 100 additional participants from a municipality are associated with a turnout decline of roughly 2.3 percentage points.[14]

The second asks whether direct experience alone could account for the result. If the crackdown had only affected those who experienced it, the votes lost could not exceed the number of people who went to Genova. Because exposure is standardised, β identifies a gradient rather than a level, so an aggregate needs a benchmark; we take the least-exposed municipality in the country, which sits 1.05 standard deviations below the mean. Summing over municipalities and weighting by the electorate,

$$\Delta V \;=\; \frac{\beta}{100} \sum_m (z_m - z_{\min})\, E_m,$$

implies roughly 224,000 fewer votes at each post-2001 election at the pooled estimate, and about 289,000 at the 2008 trough.[15] What the calculation does rule out is an effect confined to those who were there: that would require essentially every person who went

[14] This mapping is not used for identification and is intended only to convey magnitudes.

[15] Computed on the 2006 electorate of the estimation sample, 31.4 million registered voters; the sample excludes the association-hosting municipalities and therefore most large cities. The figure depends on the benchmark: taking the 1st, 5th and 10th percentiles of exposure in place of the minimum gives 221,000, 216,000 and 209,000 at the pooled estimate, and 286,000, 279,000 and 271,000 at the trough. Weighting by the electorate matters, since exposure is positively correlated with density: the electorate-weighted mean of z_m is 0.098 against an unweighted mean of zero.

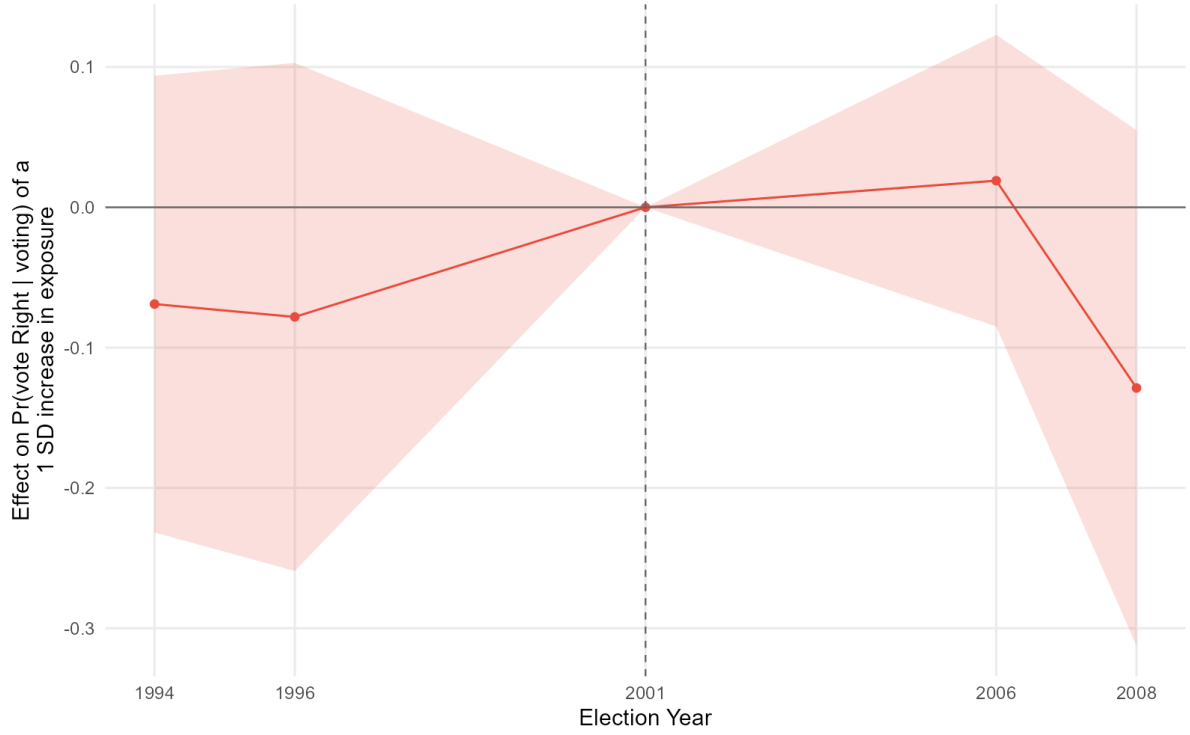
to Genova to have stopped voting at all three subsequent national elections, which is not a credible reading of who they were. Whatever happened in Genova did not stay with the people it happened to. Section 5 asks how it travelled.

4.2 Not Realignment

The first alternative is that exposure moved voters across the ideological divide, and that lower turnout is the by-product of one bloc mobilising less than the other. ITANES lets us check, because it records both how respondents voted and how they say they voted at the previous election.

Among those who voted, exposure does not change the probability of voting for the right rather than the left (Figure 5). The coefficients are small, change sign, and are never significant: +0.019 in 2006 and −0.129 in 2008, against pre-period values of −0.069 and −0.078. The pooled estimate is −0.024 (s.e. 0.052).[16]

Figure 5: Effect of Exposure on Votes for Right-Wing Parties – ITANES



Note. The figure plots the event study coefficients from a fixed effects regression of a dummy for individual voting decision (0 for left, 1 for right), estimated on respondents who voted, on exposure. The coefficients represent the interaction terms between year indicators and the standardised exposure measure, with 2001 serving as the reference year. The shaded areas represent 95% confidence intervals with standard errors clustered at the municipality level. The regression includes municipality fixed effects, province-by-year fixed effects, and individual-level time-varying controls: gender, age, education, employment status, religious orientation. Respondents in the residual third vote category are dropped. $N = 6{,}327$.

The previous-vote question lets us go further and follow individual voters out of the

[16]The confidence intervals are wide, so modest realignment cannot be ruled out; but the 2008 point estimate, if anything, runs *away* from the right, the opposite of what a backlash-to-the-right reading would predict.

electorate. We classify respondents by the bloc they voted for last time and ask whether they still vote. Figure 6 plots the probability of *still voting*, so a downward movement is an increase in abstention. Prior left-wing voters become 5.3 percentage points more likely to abstain in 2006 and 7.8 points more likely in 2008, against a baseline abstention rate of just 1.9% ($p = 0.004$ in 2006); the 2006 effect alone is nearly three times that baseline. Prior right-wing voters also drift toward abstention, but later and more weakly: 1.4 points in 2006 and 3.4 points in 2008, neither significant at the 5% level, against a baseline of 2.4%.[17]

The counter-summit's participants were overwhelmingly left-wing (Section 2.3), so the people who came home first with an account of Genova were disproportionately left-wing voters, and theirs is the bloc that moves. That the right then drifts the same way although not significantly, rather than the opposite way, is what rules out re-sorting. Whatever exposure did, it did not move voters from one bloc to the other: it moved left-wing voters out of the national electorate.

4.3 Not General Disengagement

The second alternative is that exposure made people less political in general: more apathetic, more tired, more cynical about everything. That predicts declines everywhere.

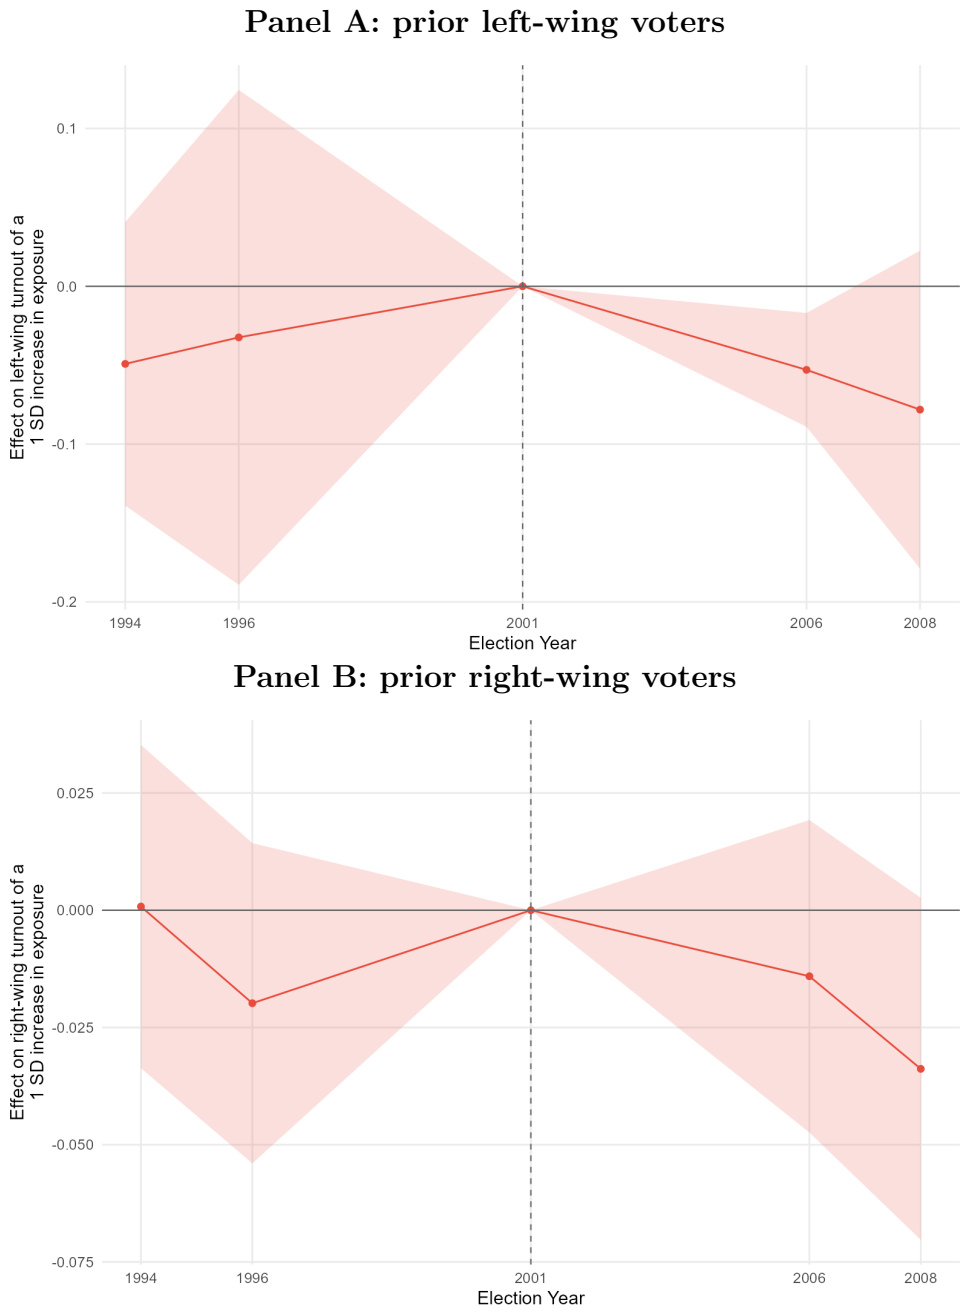
Politics outside elections. ITANES records three non-electoral forms of engagement: signing a referendum petition, attending a demonstration, and interest in politics. Exposure moves none of them (Table A.3). The estimates are +1.9 percentage points for signing a petition ($p = 0.20$), +3.3 points for attending a demonstration ($p = 0.40$), and -0.07 on the four-point interest scale ($p = 0.36$). The two participation measures are, if anything, positive. Exposed citizens kept doing politics. They stopped doing one specific thing: voting in national elections.

Other elections. If exposure made people less willing to vote in general, the effect should appear in every election. If it made them less willing to engage with the national state, it should be specific to national contests. We re-estimate the main specification on European Parliament elections (1984–2014) and on municipal elections, using the same treatment and the same controls, so that any difference from the national result comes from the election alone.

European elections have no 2001 round, so the reference year is 1999 (Figure B.3). This panel needs careful reading, because it shows a trend *before* the summit: the three earlier coefficients are large, positive and precisely estimated (+1.26 in 1984, +0.94 in 1989, +1.39 in 1994, all $p < 0.001$), meaning more-exposed municipalities were already

[17] We read the right-bloc estimates as suggestive: the transition sample is small ($N = 5{,}265$) relative to what these intervals would need to be decisive.

Figure 6: Effect of Exposure on Switching from Voting to Non-Voting – ITANES



Note. The figure plots event study coefficients from fixed effects regressions of an indicator equal to one if a respondent who voted left (Panel A) or right (Panel B) in the previous election *still votes* in the current one, on exposure. A negative coefficient is therefore an increase in the probability of switching to abstention. The coefficients represent the interaction terms between year indicators and the standardised exposure measure, with 2001 serving as the reference year. The shaded areas represent 95% confidence intervals with standard errors clustered at the municipality level. The regressions include municipality fixed effects, province-by-year fixed effects, and individual-level time-varying controls: gender, age, education, employment status, religious orientation. The post-treatment period covers the 2006 and 2008 waves only. $N = 5{,}265$.

on a relative downward path in European turnout. After 1999 the coefficients are small and flat (-0.08, -0.15, -0.09, none significant). So there is no break at the summit. The pooled coefficient of -1.00 ($p < 0.001$) simply reflects the earlier trend. This panel is also our best check that the design detects trends when they exist: it has three pre-treatment waves where the national panel has two usable ones, and where a pre-trend exists, the event study finds it.

Municipal elections are held on each municipality's own five-year cycle, so there is no single election year to pivot on. We therefore look at them on two clocks and report both. The first keeps only the six large common rounds (1990, 1995, 1999, 2004, 2009, 2014), with 1999 as the reference (Figure B.4). The second counts each municipality's elections relative to the summit, with $k = -1$ its last election before 22 July 2001 and $k = +1$ its first after (Figure B.5). Both show the same thing, and it is not the national pattern. In calendar time the coefficients decline steadily across the whole window, from $+0.49$ in 1990 and $+0.21$ in 1995 to -0.11, -0.16 and -0.30 after 1999, with no visible break and no significant post-period coefficient. In event time the same steady decline appears, from $+0.57$ at $k = -3$ to -0.48 at $k = +3$; the later coefficients are individually significant, but they continue the earlier line rather than departing from it. More-exposed municipalities were on a gently declining path in local turnout throughout, and the summit does not bend it.[18]

Only the national panel combines a flat pre-period with a sharp drop after 2001. The response is not general disengagement. Exposed citizens kept voting for their local governments and in European elections at the rates they would have anyway, and kept doing politics outside elections; what they withdrew from was voting for the national institutions of the state that had carried out the violence. Of the three possible responses, only withdrawal from the national state fits the electoral evidence.

4.4 A Reallocation of Trust: State versus Society

The electoral evidence describes what exposed citizens did. The survey lets us look at what they came to believe. If the response is withdrawal from the national state, trust should not fall across the board: it should fall for the institutions of the national state, and hold or rise for attachments outside it. Table 1 reports the full ten-item battery and two measures of local belonging.[19] Two features of the data shape how the estimates should be read. The trust modules exist only in 2001 and 2006, so the specification is the

[18] Two caveats apply to the municipal panel. $k = -3$ falls around 1986–1991, where three years are missing from the municipal archive, so some municipalities are mis-indexed there; $k = -2$ and all post-treatment periods are clean. And $k = -1$ is not a fixed distance from the summit, averaging 2.5 years before it and ranging from two months to sixteen years.

[19] The outcomes are binary high-trust indicators estimated by linear probability models, and several have low 2001 baselines, so ratios to baseline should be read as orders of magnitude rather than precise elasticities.

same difference-in-differences as elsewhere but we cannot display a pre-period path.

Trust in the national state falls. A one-standard-deviation increase in exposure is associated with a 27.2 percentage point lower probability of high trust in the President of the Republic (s.e. 0.023) on a 2001 base of 32.0%, a 14.8 point fall for the judiciary on a base of 7.1%, a 7.5 point fall for the army on a base of 15.8%, and a 2.3 point fall for political parties on a base of 1.4%. The fall in trust in the judiciary may seem surprising, since the courts eventually found for the victims; but, as Section 2 described, a judicial system that vindicates without punishing is not an obvious source of restored confidence. Trust in parliament is the one state item that moves up (+2.0 points, $p = 0.019$), from a baseline of 6.9% low enough that we read it as noise around a floor. The police item falls by 17.3 points on a base of 21.9%.

Local and horizontal attachments rise. Generalized interpersonal trust rises by 19.6 percentage points (s.e. 0.074) on a 2001 base of 23.6%. Trust in trade unions, the associational institution closest to the GSF network, rises by 4.8 points (s.e. 0.008) on a base of 3.2%. The share of people who say they belong first to their own town or region, rather than to the nation or the world, rises by 8.0 points (s.e. 0.014) on a base of 43.5%, roughly an 18% increase; on the underlying five-point scale (1 = local, 5 = world) the shift is −0.21 points, toward the local end.

Outside the national state, nothing moves. Trust in the European Union, which had no part in the crackdown and no role in local life, is unchanged (−0.019, $p = 0.44$), the attitudinal counterpart of the flat European turnout result above. Trust in the Church is also unchanged in this coding (+0.008, $p = 0.59$), but, as we show next, not in the others.

4.5 An Anti-Establishment Outlet: The Five Star Movement

From 2013 Italian voters who had lost confidence in national institutions had an alternative to staying home: the Five Star Movement (M5S), which entered national politics that year as an explicitly anti-establishment party and took roughly a quarter of the vote at its first contest. If exposure weakened people's attachment to the national state, rather than just making voting costlier, exposed municipalities that did vote should have given the new party more support.

Because the M5S did not exist before 2013, there is no pre-period, and the event-study design of Equation 2 cannot be used. We instead estimate a separate cross-sectional regression of the M5S vote share on standardised exposure for each of the three national elections the party contested, 2013, 2018 and 2022. These estimates compare munici-

Table 1: Effect of Exposure on Trust and Sense of Community Belonging (ITANES)

Outcome	Coef.	S.E.	2001 baseline	N
Panel A. The national state				
Trust: President of the Republic	-0.272^{***}	(0.023)	0.320	4,207
Trust: judiciary	-0.148^{***}	(0.018)	0.071	4,242
Trust: army	-0.075^{***}	(0.019)	0.158	4,264
Trust: political parties	-0.023^{***}	(0.003)	0.014	4,257
Trust: parliament	$+0.020^{**}$	(0.008)	0.069	4,182
Trust: police†	-0.173^{***}	(0.014)	0.219	4,340
Panel B. Outside the national state				
Trust: European Union (placebo)	-0.019	(0.024)	0.125	4,106
Trust: Church	$+0.008$	(0.014)	0.250	4,314
Panel C. Horizontal and local attachments				
Generalized (interpersonal) trust‡	$+0.196^{***}$	(0.074)	0.236	4,743
Trust: trade unions	$+0.048^{***}$	(0.008)	0.032	4,159
Belonging is local (1/0)	$+0.080^{***}$	(0.014)	0.435	4,293
Belonging, ordinal (1 local .. 5 world)	-0.206^{***}	(0.047)	2.604	4,293
Individual-Level Controls		Yes		
Municipality FE		Yes		
Province-by-Year FE		Yes		

Notes. Each row reports the coefficient on $Exposure_m \times Post_t$ from a separate regression of the stated outcome on the standardised exposure measure, following Equation 3 in pooled form. The trust battery is fielded in the 2001 and 2006 ITANES waves only, so the specification is the paper's difference-in-differences estimated on two waves, with 2001 as the pre-period; no event-study counterpart is available, so the pre-trends shown for turnout cannot be displayed here; the estimates are correlational and are read as suggestive of mechanism, not as causal. Trust outcomes are binary indicators estimated by linear probability models and, except where marked, take the value one for the *top category* of the four-point scale. ‡ Generalized trust is a *top-half* indicator, so its coefficient is not on the same scale as the top-category rows and should not be compared to them directly. "Belonging is local" equals one for respondents reporting attachment primarily to their municipal or regional community and zero for the national, European or world community; the ordinal variant is the raw five-point scale, on which *lower* means more local, so its sign is the mechanical flip of the binary. "2001 baseline" is the mean of the outcome in the 2001 wave. All regressions include municipality fixed effects, province-by-year fixed effects, and individual-level time-varying controls (gender, age, education, employment status, religious orientation). Standard errors clustered at the municipality level in parentheses. Municipalities hosting at least one GSF association are excluded. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

palities at a point in time rather than tracking changes within them, and we read them accordingly.

Figure 7 shows the results. The coefficients are positive in the two elections in which the M5S ran as an outsider: a one-standard-deviation increase in exposure is associated with a 0.22 percentage point higher M5S share in 2013, against a mean of 23.7%, and a 0.43 point higher share in 2018, against a mean of 30.0%. In 2022 the coefficient is 0.03, against a mean of 13.0%. Only the 2018 estimate is distinguishable from zero, and only when standard errors are clustered by municipality; with province-level clustering none of the three is. The magnitudes are modest, about 1.4% of the mean share in 2018.

We therefore treat this as suggestive. First, its timing follows the party's position rather than its size. The association is largest in 2018, when the M5S campaigned most explicitly against the established parties, and it disappears in 2022, after four years in which the party had governed with the League, then with the Democratic Party, and then within the Draghi national-unity government, and could no longer present itself as outside the system. Second, selection works against the result: if the voters who withdraw in exposed municipalities are disproportionately the disaffected ones, those who remain are less anti-establishment than they would otherwise be, which biases the M5S coefficient toward zero.

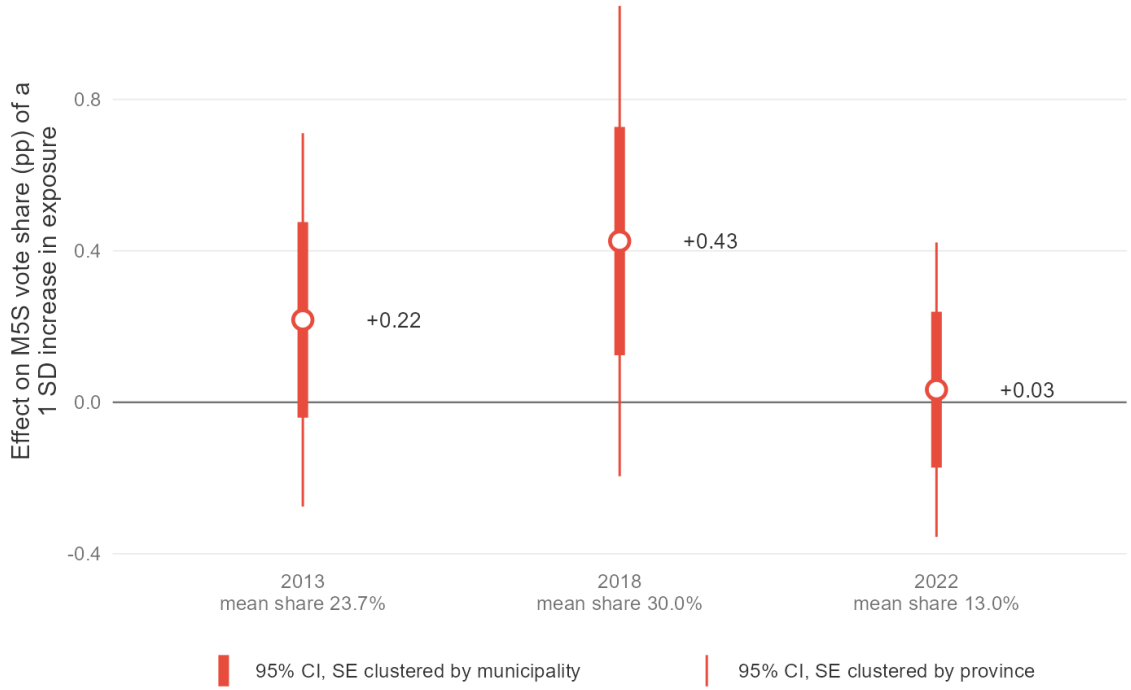
Support for the M5S is not movement along the left–right dimension: the party drew voters from both blocs and defined itself against the political class as a whole. Where exposed voters could vote against national institutions rather than abstain from them, some appear to have done so, which is the same disposition the turnout results describe, expressed through a different channel.

5 How the Response Travelled

Section 4 showed what exposed citizens did: they withdrew from national elections, without changing sides and without withdrawing from politics in general. This section asks how an event witnessed by a few hundred thousand people reached enough others to move turnout across a large part of the country, and why the effect lasted.

Our answer has two parts, and we use three terms consistently throughout. By the *account* of Genova we mean an interpretation of what happened there; there were two in competition, as we describe below. By *diffusion* we mean the spread of the protesters' account from the people who were there to the people around them, which explains where turnout fell. By *memory* we mean what keeps that account alive once it has spread, which explains why turnout stayed down. Memory in this sense is partly individual, a matter of which events come to mind when a person is prompted (Gennaioli and Shleifer, 2010), and partly collective: in the tradition running from Halbwachs (1992), what lasts is not a trace in one person's head but a story that a group keeps retelling. That second view

Figure 7: Exposure and Support for the Five Star Movement



Note. Each point is the coefficient from a separate cross-sectional regression of the M5S vote share in national elections on the standardised 100 km exposure index, for 2013, 2018 and 2022 (specification P8). The controls are analogous to those from the main specification, province FE, and the sample excludes GSF-host municipalities as in the baseline. The M5S first contested national elections in 2013, so no pre-period exists and these are cross-sectional associations rather than difference-in-differences estimates. Thick bars are 95% confidence intervals with standard errors clustered by municipality; thin bars cluster by province. The mean M5S share in the estimation sample is reported under each year.

makes a testable prediction. If the effect faded because individuals forget, it should fade at similar rates everywhere; if it lasts because groups retell, it should last longest where the ties that carry the retelling are densest.

The official account, carried by national television and endorsed by the centre-right government that had taken office in June 2001, said the demonstrators were violent rioters and the police had done their job. The other came from the people who had been there: first-person, emotional, and centred on the police. It was eventually vindicated in court, but for years it could spread mainly by word of mouth, through the same local networks that had taken people to Genova. If that is right, three things should be true. The protesters' account must have existed and been available; the effect should be strongest where face-to-face ties are densest; and it should be weakest where the official account reached households most strongly. We take these in turn.

5.1 Diffusion Across Space and Media

Two accounts in the press. We assemble 569 press and wire items on Genova from Factiva, spanning 2000 to 2011, and track how much was written and whom it blamed (Figure 8). Public attention did not end with the summit. It returned in two large

waves, around the first-instance trials of 2007–2008 and the appeals of 2010, so the event kept coming back into national discussion for most of a decade. And the blame shifted. Through the early 2000s the coverage mostly blamed the protesters, peaking in 2004 when 55% of items carried protester-blame language against 5% carrying police-blame language. The two frames converged during the trials, and from the 2008 convictions onward blame directed at the police and their commanders dominated, reaching over 80% of items by 2011.

This matters for the argument in two ways. The police-as-perpetrator account existed and kept being reactivated, so there was something to spread. And in the years right after the summit, when a memory of Genova would have had to form, the national coverage ran the other way, blaming the protesters by roughly ten to one. In those years the protesters' account could only travel locally, which is what makes its geography informative.

Where word of mouth works: rural and urban municipalities. If the protesters' account spread through personal ties, the effect should be strongest where those ties are densest. We split municipalities into rural and urban using Eurostat's Degree of Urbanisation applied to 2001 population density, which fixes the classification for the whole panel, and re-standardise exposure within each group.[20]

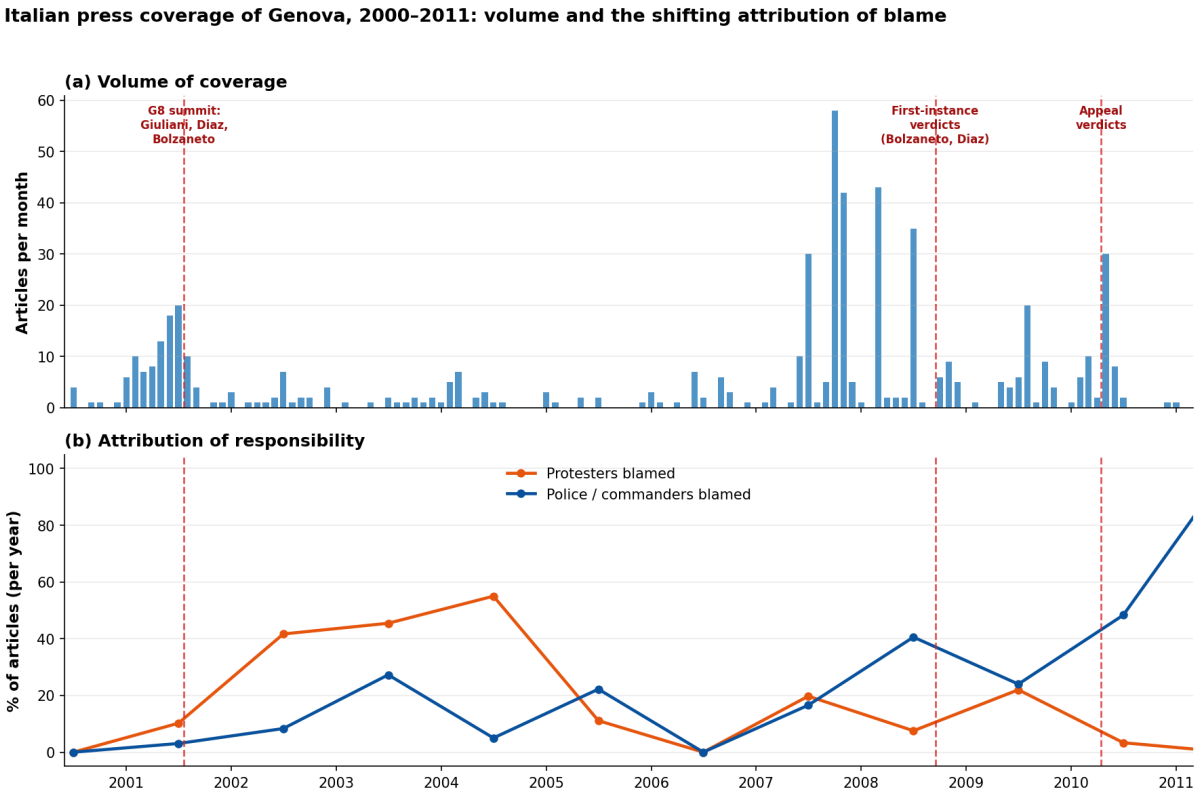
The contrast is sharp (Figure 9). In rural municipalities the pre-period is flat, and after 2001 turnout steps down and stays down: -0.92 percentage points in 2006, -1.00 in 2008 and -0.88 in 2013, all significant at $p < 0.001$, pooling to -0.90 (s.e. 0.184). In urban municipalities the effect is smaller and does not last: -0.46 in 2006 and -0.48 in 2008, then $+0.10$ in 2013, pooling to -0.13 (s.e. 0.175), which is not distinguishable from zero. The rural pooled effect is nearly seven times the urban one.

The timing matters as much as the size. Both groups react at the first election after the summit, but only the rural group is still reacting twelve years later. That is the prediction of the collective view of memory: the effect lasts where dense, overlapping ties keep an account in circulation and give it social confirmation, and fades where they do not (Coleman, 1988; Putnam, 2000). It also explains the partial recovery in the national series in 2013, which Section 4.1 left open: the rebound is the urban effect dying out, not the rural one. Rurality can stand in for other things too, so we read this contrast as consistent with the diffusion account rather than as a clean test of a single network parameter.

Where the official account arrived loudest: television. The press record shows that the protesters' account existed, and the rural–urban contrast shows that the effect lives where personal ties are dense. Neither shows directly that the two accounts *competed*.

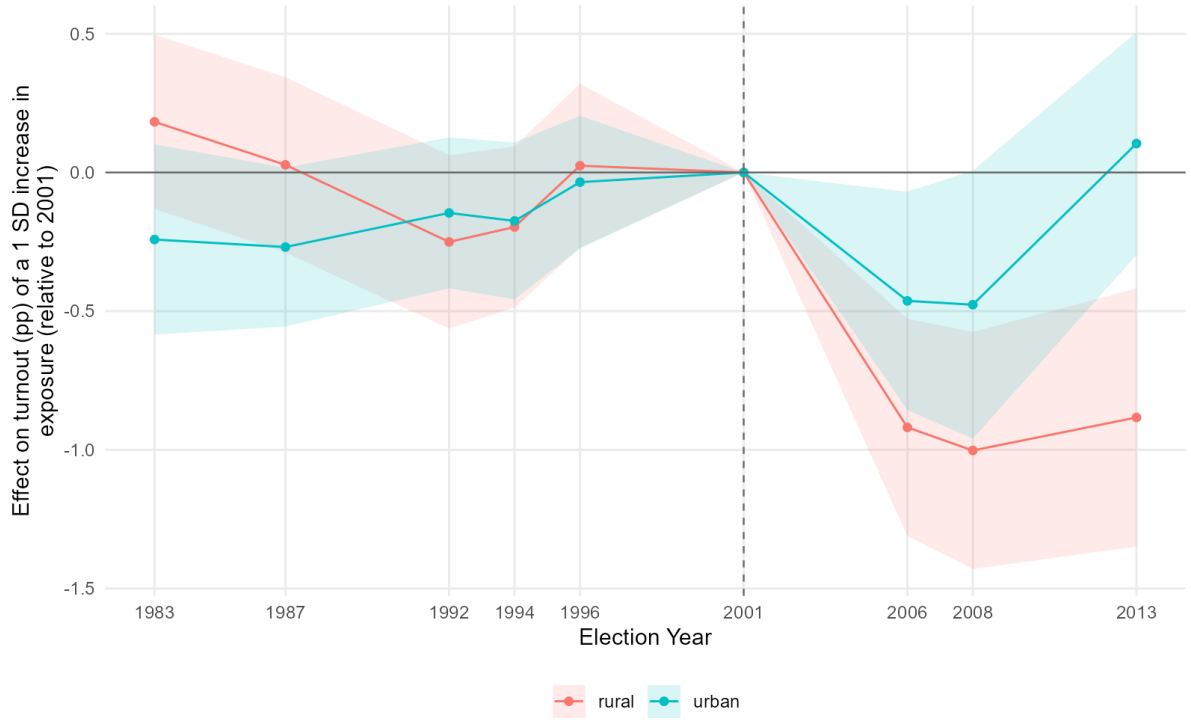
[20] Rural municipalities are those with a 2001 density below 300 inhabitants per square kilometre; 79.7% of municipalities in the estimation sample are rural. Because exposure is re-standardised within each group, each coefficient is per standard deviation of exposure *within* the group.

Figure 8: Italian Press Coverage of Genova, 2000–2011: Volume and the Attribution of Blame



Notes. The figure summarizes 569 Italian news items on the Genova G8 retrieved from Factiva and published between February 2000 and January 2011. The corpus is dominated by national wire services (AGI, ANSA, and Apcom account for the large majority of items), so the series reflects agency reporting rather than editorial-page opinion. Panel (a) plots the number of articles per month. Coverage clusters in three periods: the July 2001 summit and the nights of the Diaz raid and the Bolzaneto barracks; the first-instance trials and verdicts of 2007–2008; and the appeal verdicts of 2010. Vertical dashed lines mark these three moments. Panel (b) plots, for each year, the share of articles whose language attributes responsibility for the violence to the protesters (for example *black bloc*, *devastazioni*, *saccheggi*) and the share that attributes it to the police or their commanders (for example *pestaggi*, *tortura*, *irruzione illegittima*, and, in the trial period, convictions of officers and officials). Guilt and conviction terms are counted toward police-blame only when a police actor is named in the same article, so that the 2005 trial of protesters for devastation does not load onto the police series. The two shares are roughly balanced through the mid-2000s and then cross during the trials: from 2008 onward, as the courts convict officers and senior officials, police-and-commander blame rises steeply and becomes the dominant frame. The measure is a transparent dictionary count over headline and body text and is descriptive; "neutral" procedural items are the single largest category and are omitted from the two plotted shares, and the sparse years 2005 and 2006 (9 and 24 items) yield noisier points. Source: Factiva (AGI, ANSA, Apcom, Reuters Italia, and others).

Figure 9: Effect of Exposure on Turnout - Heterogeneity by Rural and Urban Areas



Note. The figure plots the event study coefficients from fixed effects regressions of voter turnout on exposure, estimated separately on rural and urban subsamples. The definition of urban is taken from DEGURBA applied to 2001 density: a municipality is rural if its density is lower than 300 inhabitants per square kilometre. Exposure is re-standardised within each subsample, so coefficients are per standard deviation of exposure within group. The coefficients represent the interaction terms between year indicators and the exposure measure, with 2001 serving as the reference year. The shaded areas represent 95% confidence intervals with standard errors clustered at the municipality level. The regressions include municipality fixed effects, province-by-year fixed effects, and the full set of time-invariant 2001 controls interacted with each election year. $N = 55{,}341$ (rural) and $N = 14{,}163$ (urban).

For that we use the strength of the Mediaset television signal from Durante et al. (2019). Mediaset is the natural measure of exposure to the official account: it is the commercial broadcaster Durante et al. (2019) link to support for the centre-right, and it was the centre-right government that publicly defended the police and cast the demonstrators as rioters.

Signal strength is not random, since it is stronger near cities and transmitters. Following Olken (2009), Enikolopov et al. (2011) and Durante et al. (2019), we use only the part driven by terrain. We add to the main specification the interaction of exposure, actual signal strength and year, together with the same interactions for the *free-space* signal, the reception a municipality would get from the same transmitters with no terrain in the way, and for terrain ruggedness, plus all lower-order terms.[21] Holding free-space signal and ruggedness fixed, two municipalities with the same exposure and the same distance to the transmitters differ in reception only because of the hills that happen to lie between them and the transmitter.

The moderation behaves as the argument predicts. Before 2001 the triple interaction is

[21] Television signal is removed from the control vector Z_m here, since it becomes the moderator.

flat and none of the coefficients is significant. After 2001 it turns positive and stays there: +0.31 in 2006, +0.30 in 2008 and +0.32 in 2013 (Figure 10). Holding exposure fixed, a one-standard-deviation stronger signal offsets about 0.3 percentage points of the decline at every post-summit election. The pooled estimate is +0.234 (s.e. 0.097), and adding a competing interaction with population density barely changes it (+0.204, $p = 0.035$), so this is not urbanisation under another name.

Taken together, the three pieces fit one picture. The protesters' account existed and was eventually vindicated nationally, but for years it travelled person to person. The turnout decline is largest and most lasting where personal ties are dense, and smallest where the official account arrived most strongly on television. None of these pieces on its own proves that diffusion of an account caused the turnout response. Together they explain how an event experienced directly by a few could become a lasting local memory, and why the same repressive event produced a large effect in some places, a small one in others, and none in some.

6 Robustness and Alternative Explanations

This section asks how much the turnout result depends on choices we made and on things we could not control. We group the checks by the question each answers: how long the effect lasts and whether it starts when the event does; how precise it is once spatial dependence is taken seriously; whether it depends on how exposure is defined; whether a few places drive it; and whether it could reflect something other than the crackdown. Across these checks the pattern is the same: no pre-trends, a break at the first election after the summit, and a negative gradient that persists through 2022. What varies is the magnitude, with the metric, and the precision, which falls under the most conservative treatment of spatial dependence. The figures are collected in Appendix B.

6.1 Timing and Persistence

Extending the panel to 2022. The baseline panel ends in 2013, for the reasons given in Section 3. Adding the 2018 and 2022 national elections, with the same specification and 2001 as the reference year, leaves the earlier coefficients essentially unchanged (Figure B.6). The gradient remains negative at both later elections, at about −1.06 percentage points in 2018 and −0.43 in 2022, and neither confidence interval includes zero. The 2018 coefficient is the largest in the series and coincides with the election in which the M5S was strongest (Section 4.5). We read the post-2013 coefficients cautiously, since abstention means something different once an anti-establishment party is on the ballot, but the extension speaks directly to persistence: two decades after the summit, the gap has not closed.

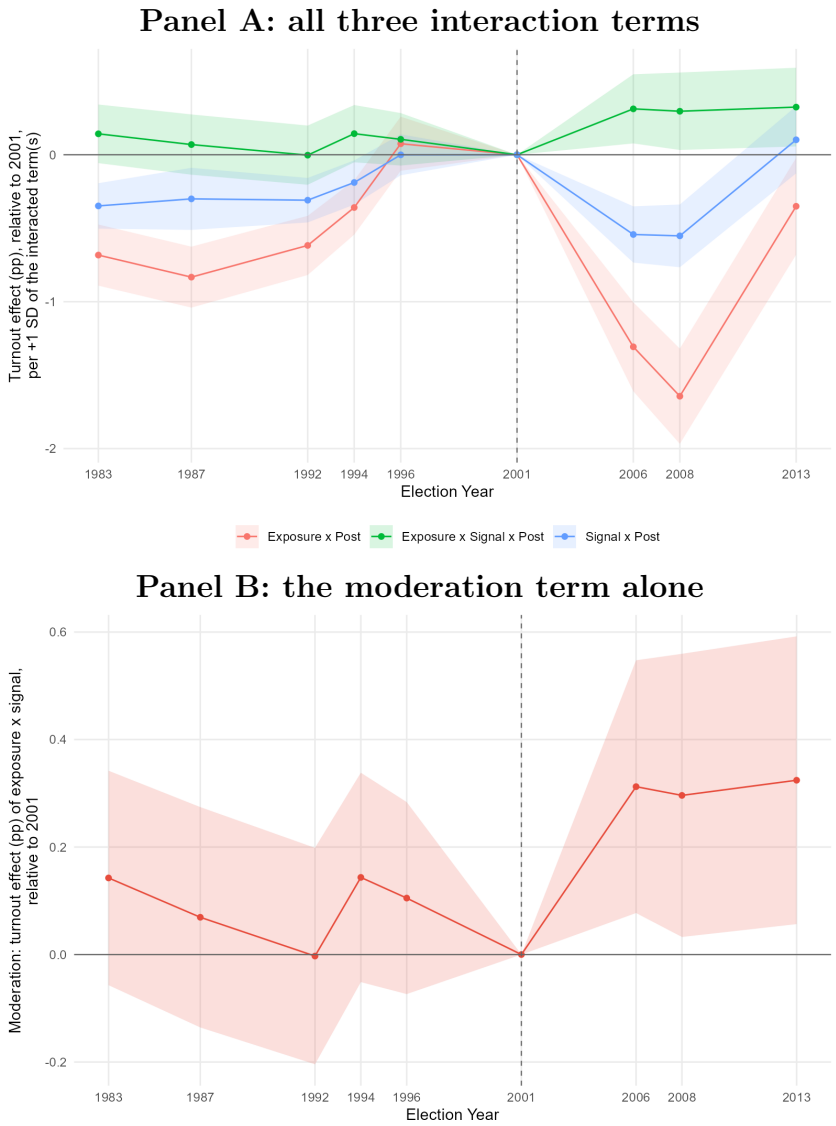


Figure 10: Media Narratives: The Turnout Effect of Exposure Weakens Where the Mainstream Signal Is Stronger

Notes. The figure plots coefficients from a fixed-effects regression of national-election turnout that augments the main specification withexposure × signal × year terms and all lower-order interactions. Panel A shows all three interaction paths: the exposure effect at mean signal (red), the signal effect at mean exposure (blue), and the moderation (green). Panel B isolates the moderation term. Following Olken (2009), Enikolopov et al. (2011), and Durante et al. (2019), the specification additionally includes the analogous interactions with the *free-space* signal (the predicted reception from the same transmitters absent any terrain obstruction) and with terrain ruggedness, so that the identifying variation in actual signal strength reflects topographic shadowing along the path between transmitter and receiver rather than proximity to transmitters or urban geography. TV signal is excluded from the control vector, since it is the moderator. Each point on the moderation path is the *differential* effect on turnout of a one-standard-deviation increase in exposure associated with a one-standard-deviation stronger actual signal, relative to 2001. Signal measures are the municipality-level Mediaset variables of Durante et al. (2019). A positive coefficient means that, at a given level of exposure, municipalities that received the mainstream television signal more strongly experienced a *smaller* post-2001 turnout decline. Shaded areas are 95% confidence intervals with standard errors clustered at the municipality level; the moderation is not significant under province-level clustering. The regression includes municipality and province-by-year fixed effects and the remaining controls interacted with each election year. $N = 69{,}414$.

Placebo in time. The identifying claim is that the break happens at 2001, not somewhere in the panel. We move the treatment date to a counterfactual summit in 1991, 1993 and 1995, each time truncating the sample at 2001 so the real event cannot enter the false post-period (Figure B.7). None of the false dates produces a sustained decline: the coefficients drift slightly downward in the early 1990s and return to zero or above by 1996 and 2001. That mild early-1990s drift is the same one visible in the baseline event study, and this exercise shows it is not large enough, at any false date, to produce the post-2001 pattern on its own.

Individual-level replication. ITANES gives the same timing at the individual level (Figure B.8). The pre-2001 coefficients are small and insignificant, and exposure lowers the probability of voting by 11.5 percentage points in 2006 ($p < 0.001$) and 6.1 points in 2008 ($p = 0.106$).[22]

6.2 Inference

Province-level clustering. Our baseline clusters standard errors by municipality. Clustering by province instead (112 clusters) leaves the point estimates unchanged but widens the confidence intervals considerably (Figure B.9): the 2008 coefficient remains significant ($p = 0.045$), the 2006 coefficient becomes marginal ($p = 0.088$), and the pooled and cumulative estimates are no longer significant at conventional levels ($p = 0.124$ and $p = 0.114$). Because exposure is smooth over space, province clustering is a natural way to allow for spatially correlated errors, and we do not dismiss it. The sign, timing and size of the effect are unaffected; how precisely it is estimated depends on how conservatively spatial dependence is modelled.

Spatial noise. Both exposure and turnout vary smoothly over space, so spatially correlated unobservables could in principle produce a gradient that looks like an effect. The two diagnostics of Conley and Kelly (2025) ask whether an estimate of our size could arise from spatial noise alone. The first re-estimates Equation 2 on placebo treatments drawn with the same spatial structure as the real exposure measure; the second keeps the real treatment and adds spatially correlated noise to the outcome. Each gives a reference distribution of t-statistics for the cumulative post-2001 effect (Figure B.10). The actual t-statistic is -4.19. None of the placebo treatments comes close to it, and most fall inside the conventional ± 1.96 band. The synthetic-outcome distribution is wider, as it should be when noise enters the outcome directly, but all but one of its draws lie to the right of

[22] The pooled coefficient is -0.032 and insignificant ($p = 0.14$), which is unsurprising: reported turnout in ITANES is 94% and the sample is two orders of magnitude smaller than the municipal panel. The direction and timing replicate; the precision does not, and the municipal panel remains the primary evidence.

the actual estimate.

6.3 How Exposure Is Defined

Relevance-weighted exposure. Weighting associations by the relevance score of Section 3 rather than counting them equally gives almost the same path as the baseline: -0.69, -0.78 and -0.40 at the three post-summit elections (Figure B.11). The result does not depend on treating every association as equally important.

Binary exposure. A continuous treatment brings two caveats from the recent difference-in-differences literature (Callaway et al., 2024; de Chaisemartin and D'Haultfœuille, 2020). It requires parallel trends at every level of exposure, not just between a treated and an untreated group, and the coefficient is a weighted average of responses across doses whose weights need not be even. We therefore also split municipalities at the median of exposure (Figure B.12). The pre-period is flat and insignificant, and the post-summit coefficients are -2.52, -2.49 and -2.74 percentage points, all $p < 0.001$, with no decay. The binary version confirms the timing and sign of the baseline without assuming a linear dose response.

A narrower radius. Our index counts associations within 100 km. Halving the radius to 50 km is the one check that does not reproduce the baseline in every form. With the continuous index, the post-2001 coefficients are small and insignificant (-0.09, -0.14 and $+0.16$) and the 1983 coefficient is positive and significant (Figure B.13). With a median split of the same index, the pattern returns: the pre-period is flat apart from a marginal 1983 coefficient, and the post-2001 coefficients are -1.35, -1.18 and -1.12, all highly significant (Figure B.14), about half the size of the 100 km binary contrast, which is the ordering one would expect if a wider radius captures more of each association's catchment. We read the divergence as a matter of functional form. At 50 km, 8.2% of municipalities have no association in range and the rest of the distribution is strongly skewed, so a single linear coefficient describes the dose poorly. The ordinal contrast between more- and less-exposed places survives the change of radius; the linear dose response does not, and a reader who prefers the continuous specification should treat that as a real limitation.

Host municipalities. Bringing back the municipalities that host a GSF association keeps the sign and timing but shrinks the effect to about a third of the baseline: -0.25 in 2006, -0.30 in 2008 and -0.07 in 2013, pooling to -0.25 (Figure B.15). This is what the design implies. Host municipalities combine proximity with the direct presence of the organisation, so they need not lie on the same gradient as their neighbours, and adding them dilutes the variation.

6.4 Where the Effect Comes From

Region-by-year fixed effects. Replacing province-by-year with region-by-year fixed effects widens the comparison set and gives a larger decline: -0.80 in 2006, -0.84 in 2008 and -0.81 in 2013, with no recovery by 2013 (Figure B.16). This fits the point made in Section 3: a wider comparison set brings in less-contaminated comparison municipalities. The cost is that two pre-period coefficients become significant (-0.40 in 1992 and -0.27 in 1994), so we do not treat this specification as stronger than the baseline.

Leave one out. The effect could in principle rest on a small part of the country, since exposure is highest in a few dense, well-organised areas of the North and Centre. We re-estimate the pooled specification dropping one province at a time, and then one region at a time (Figure B.17). No single area drives the result. Across the 106 provinces in the sample the coefficient ranges from -0.83 to -0.41, with a median of -0.62 and an interquartile range of $[-0.63, -0.61]$; across the 19 regions it ranges from -1.40 to -0.46, with a median of -0.61. All 125 estimates are significant at the 5% level. Two runs are worth naming. Dropping Liguria gives -0.58 and dropping the province of Genova gives -0.63, so the result is not a local Genova story, which is what the diffusion account implies. And the most influential exclusion strengthens the result rather than weakening it: dropping Lombardia doubles the estimate to -1.40, and dropping Milan alone takes it to -0.83. Lombardy, with the densest and most exposed urban municipalities, works against our result, consistent with the weak urban response in Section 5.1.

6.5 Alternative Explanations

Exposure and baseline characteristics. GSF associations cluster in organised, politically dense places, so exposure is not independent of observable characteristics, and it is better to show this than to assume it away. Table A.4 regresses exposure on each 2001 characteristic in turn, with and without province fixed effects. More-exposed municipalities are denser, younger, better educated, more industrial and less agricultural, with higher employment and stronger television reception. Within provinces they also lean left: exposure is positively associated with the historical PCI vote and negatively with the historical Pentapartito vote. A differential turnout trend in left-leaning, organised places could in principle look like an effect of the crackdown. The specification is designed for this: every characteristic in Table A.4, including both historical vote shares, is interacted with election-year indicators, so differential trends along each of them are absorbed election by election. Two results then cut against a partisan-trend explanation directly. The pre-2001 coefficients are flat, which they would not be if the trend predated the summit; and the left–right vote conditional on voting does not move (Section 4.2), which a differential left-area dynamic would be expected to produce.

Selective migration. Turnout could fall with no change in anyone's behaviour if politically engaged residents left high-exposure municipalities after 2001. The timing makes this unlikely, since the decline is already visible in 2006 and is strongest in rural areas, where young people were leaving for the cities regardless of exposure. More importantly, migration is a story about who is counted at the ballot box, and it makes no prediction about what those who stay believe. To reproduce the trust results through composition alone, the stayers in exposed municipalities would have to trust the President, the army and the judiciary less, trust one another and their unions more, feel more attached to their town, and feel exactly the same about the European Union. This does not rule out every compositional change, but it makes sorting an incomplete explanation for the outcomes taken together.

7 Conclusion

This paper examines how a salient episode of state repression during the 2001 G8 summit in Genova affected civic and political behavior in a mature democracy. We exploit cross-municipality variation in proximity to the Genoa Social Forum (GSF) network to construct a market-access measure of local exposure to the crackdown, and estimate difference-in-differences specifications with a continuous treatment.

Our main finding is that greater network-based exposure to repression caused a persistent decline in voter turnout in subsequent national elections. Event-study estimates reveal flat pre-trends and a discrete break at the first post-summit election, with a one-standard-deviation increase in exposure lowering turnout by 0.62 percentage points on average and by 0.80 at the peak, and it stayed negative until 2013. The result is stable across alternative fixed-effect structures, a relevance-weighted treatment, a binarised exposure definition, and the inclusion of municipalities that host associations.

Using individual-level survey evidence, we find that this response operates on the extensive margin: exposure increases transitions from voting to abstention, first and most clearly among prior left-wing voters and later, more weakly, among prior right-wing ones, while we detect no shift between left and right conditional on voting. The withdrawal is specific to the ballot: exposure has no detectable effect on petitioning, demonstrating, or interest in politics, and neither European nor municipal contests display a break at the summit.

More-exposed areas report substantially lower trust in the President of the Republic, the judiciary, the army and political parties, and at the same time higher generalized trust, higher trust in trade unions, and a stronger sense of local belonging, with trust in the European Union unchanged. The relevant survey modules exist only in the 2001 and 2006 waves, so the difference-in-differences is estimated without an event-study counterpart.

We also provide suggestive evidence on mechanisms. The turnout effects are four to

seven times larger in rural municipalities and persist there through 2013 while fading in urban ones, and they weaken where terrain-driven variation gave households stronger access to the mainstream television signal. Both are consistent with a channel operating through the diffusion and persistence of first-hand accounts of the crackdown within local social networks, in competition with an official broadcast frame that, as the Factiva record shows, dominated national coverage for years before the courts adjudicated the event the other way.

More broadly, the results highlight that even in consolidated democracies, acute episodes of state violence can have durable consequences for participation in national electoral politics, without operating through partisan realignment and without dissolving social capital. By linking localized exposure to repression to subsequent civic and electoral behavior, the paper contributes evidence on how coercive institutional performance shapes political engagement and the resilience of democratic participation.

References

Agnoletto, V. and Guadagnucci, L. (2011). *L'eclisse della democrazia: le verità nascoste sul G8 2001 a Genova*. Feltrinelli Editore.

Ahlfeldt, G. M., Redding, S. J., Sturm, D. M., and Wolf, N. (2015). The economics of density: Evidence from the berlin wall. *Econometrica*, 83(6):2127–2189.

Alesina, A. and Fuchs-Schündeln, N. (2007). Good-bye lenin (or not?): The effect of communism on people's preferences. *American Economic Review*, 97(4):1507–1528.

Andretta, M., Della Porta, D., Mosca, L., Reiter, H., et al. (2002). *Global, noglobal, newglobal. La protesta contro il G8 a Genova*. Edizioni Giuseppe Laterza.

Ang, D., Bencsik, P., Bruhn, J., and Derenoncourt, E. (2025). Community engagement with law enforcement after high-profile acts of police violence. *American Economic Review: Insights*, 7(1):124–140.

Arjona, A., Chacón, M., and García-Montoya, L. (2025). The impact of political assassinations on turnout: Evidence from colombia. *Proceedings of the National Academy of Sciences*, 122(11):e2414767122.

Ash, E. and Hansen, S. (2023). Text algorithms in economics. *Annual Review of Economics*, 15:659–688.

Balcells, L. (2012). The consequences of victimization on political identities: Evidence from Spain. *Politics & Society*, 40(3):311–347.

Balcells, L., Dorsey, S., and Tellez, J. F. (2021). Repression and dissent in contemporary catalonia. *British Journal of Political Science*, 51(4):1742–1750.

Barceló, J. (2021). The long-term effects of war exposure on civic engagement. *Proceedings of the National Academy of Sciences*.

Bauer, M., Blattman, C., Chytilová, J., Henrich, J., Miguel, E., and Mitts, T. (2016). Can war foster cooperation? *Journal of Economic Perspectives*, 30(3):249–274.

Bautista, M. A. (2015). Political effects of state-led repression: The Chilean case. Working paper, Harris School of Public Policy, University of Chicago.

Bautista, M. A., González, F., Martínez, L. R., Muñoz, P., and Prem, M. (2023). The geography of repression and opposition to autocracy. *American Journal of Political Science*, 67(1):101–118.

Bigoni, M., Bortolotti, S., Casari, M., and Gambetta, D. (2019). At the root of the north–south cooperation gap in italy: Preferences or beliefs? *The Economic Journal*, 129(619):1139–1152.

Bigoni, M., Bortolotti, S., Casari, M., Gambetta, D., and Pancotto, F. (2016). Amoral familism, social capital, or trust? the behavioural foundations of the italian north–south divide. *The Economic Journal*, 126(594):1318–1341.

Buggle, J. C. and Durante, R. (2021). Climate risk, cooperation and the co-evolution of culture and institutions. *The Economic Journal*, 131(637):1947–1987.

Bühler, M. and Madestam, A. (2023). State repression, exit, and voice: Living in the shadow of cambodia's killing fields.

Callaway, B., Goodman-Bacon, A., and Sant'Anna, P. H. C. (2024). Difference-in-differences with a continuous treatment. *Journal of Econometrics*. forthcoming; NBER Working Paper 32117.

Campante, F., Durante, R., and Tesei, A. (2022). Media and social capital. *Annual review of economics*, 14(1):69–91.

Canini, R., González, F., and Prem, M. (2026). Police-civilian fatal encounters and state trust. Technical report, GLO Discussion Paper.

Cantoni, D., Heizlsperger, L.-J., Yang, D. Y., Yuchtman, N., and Zhang, Y. J. (2022). The fundamental determinants of protest participation: Evidence from hong kong's antiauthoritarian movement. *Journal of Public Economics*, 211:104667.

Cantoni, D., Kao, A., Yang, D. Y., and Yuchtman, N. (2024). Protests. *Annual Review of Economics*, 16.

Cantoni, D., Yang, D. Y., Yuchtman, N., and Zhang, Y. J. (2019). Protests as strategic games: experimental evidence from hong kong's antiauthoritarian movement. *The Quarterly Journal of Economics*, 134(2):1021–1077.

CIVICUS (2025). People power under attack 2025. Technical report, CIVICUS Monitor.

Coleman, J. S. (1988). Social capital in the creation of human capital. *American Journal of Sociology*, 94(Supplement):S95–S120.

Conley, T. G. and Kelly, M. (2025). The standard errors of persistence. *Journal of International Economics*, 153:104027. Article 104027.

Cremaschi, S. and Masullo, J. (2025). The political legacies of wartime resistance: How local communities in italy keep anti-fascist sentiments alive. *Comparative Political Studies*, 58(14):3065–3108.

Daniele, G., Aassve, A., and Le Moglie, M. (2023). Never forget the first time: The persistent effects of corruption and the rise of populism in italy. *The Journal of Politics*, 85(2):468–483.

Daniele, G., Martinangeli, A., Passarelli, F., Sas, W., and Windsteiger, L. (2025). Externalities and the erosion of trust.

Davenport, C. (2007a). State repression and political order. *Annual Review of Political Science*, 10:1–23.

Davenport, C. (2007b). State repression and political order. *Annual Review of Political Science*, 10:1–23.

Davenport, C. and Armstrong, D. A. (2004). Democracy and the violation of human rights: A statistical analysis from 1976 to 1996. *American Journal of Political Science*, 48(3):538–554.

de Chaisemartin, C. and D'Haultfœuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, 110(9):2964–2996.

della Porta, D. (1995). *Social Movements, Political Violence, and the State: A Comparative Analysis of Italy and Germany*. Cambridge University Press, Cambridge.

Di Tella, R., Galiant, S., and Schargrodsky, E. (2007). The formation of beliefs: evidence from the allocation of land titles to squatters. *The Quarterly Journal of Economics*, 122(1):209–241.

Durante, R., Mastrorocco, N., Minale, L., and Snyder, J. M. (2025). Unpacking social capital. *The Economic Journal*, 135(667):773–807.

Durante, R., Pinotti, P., and Tesei, A. (2019). The political legacy of entertainment TV. *American Economic Review*, 109(7):2497–2530.

Enikolopov, R., Makarin, A., and Petrova, M. (2020). Social media and protest participation: Evidence from russia. *Econometrica*, 88(4):1479–1514.

Enikolopov, R., Petrova, M., and Zhuravskaya, E. (2011). Media and political persuasion: Evidence from russia. *American Economic Review*, 101(7):3253–3285.

Facchetti, E. (2025). Police infrastructure, police performance, and crime: Evidence from austerity cuts. Technical report, IFS Working Papers.

Fouka, V. and Voth, H.-J. (2023). Collective remembrance and private choice: German–greek conflict and behavior in times of crisis. *American Political Science Review*, 117(3):851–870.

Fuchs-Schündeln, N. and Schündeln, M. (2015). On the endogeneity of political preferences: Evidence from individual experience with democracy. *Science*, 347(6226):1145–1148.

Gennaioli, N. and Shleifer, A. (2010). What comes to mind. *The Quarterly journal of economics*, 125(4):1399–1433.

Gennaro, G. and Ash, E. (2022). Emotion and reason in political language. *The Economic Journal*, 132(643):1037–1059.

Gentzkow, M., Kelly, B., and Taddy, M. (2019). Text as data. *Journal of Economic Literature*, 57(3):535–574.

González, F. and Prem, M. (2024). Police violence, student protests, and educational performance. *Review of Economics and Statistics*, 106(3):712–727.

Guiso, L., Sapienza, P., and Zingales, L. (2004). The role of social capital in financial development. *American Economic Review*, 94(3):526–556.

Guiso, L., Sapienza, P., and Zingales, L. (2008). Social capital as good culture. *Journal of the European economic Association*, 6(2-3):295–320.

Guiso, L., Sapienza, P., and Zingales, L. (2016). Long-term persistence. *Journal of the European Economic Association*, 14(6):1401–1436.

Hager, A. and Krakowski, K. (2022). Does state repression spark protests? evidence from secret police surveillance in communist poland. *American Political Science Review*, 116(2):564–579.

Hager, A., Krakowski, K., and Schaub, M. (2019). Ethnic riots and prosocial behavior: Evidence from kyrgyzstan. *American Political Science Review*, 113(4):1029–1044.

Halbwachs, M. (1992). *On Collective Memory*. University of Chicago Press, Chicago.

Harada, M., Ito, G., and Smith, D. M. (2024). Destruction from above: Long-term legacies of the tokyo air raids. *The Journal of Politics*, 86(2):782–797.

International Telecommunication Union (2002). *World Telecommunication Indicators Database*. ITU, Geneva. Italy: individuals using the Internet and fixed-broadband subscriptions, 2001.

Komisarchik, M., Sen, M., and Velez, Y. R. (2022). The political consequences of ethnically targeted incarceration: Evidence from japanese american internment during world war ii. *The Journal of Politics*, 84(3):1497–1514.

Lawrence, A. K. (2017). Repression and activism among the Arab Spring's first movers: Evidence from Morocco's February 20th movement. *British Journal of Political Science*, 47(3):699–718.

Malmendier, U. and Nagel, S. (2011). Depression babies: do macroeconomic experiences affect risk taking? *The quarterly journal of economics*, 126(1):373–416.

Mummolo, J. (2018). Militarization fails to enhance police safety or reduce crime but may harm police reputation. *Proceedings of the National Academy of Sciences*, 115(37):9181–9186.

Nannicini, T., Stella, A., Tabellini, G., and Troiano, U. (2013). Social capital and political accountability. *American Economic Journal: Economic Policy*, 5(2):222–250.

Nikolova, M., Popova, O., and Otrachshenko, V. (2022). Stalin and the origins of mistrust. *Journal of Public Economics*, 208:104629.

Nunn, N. and Wantchekon, L. (2011). The slave trade and the origins of mistrust in africa. *American economic review*, 101(7):3221–3252.

Olken, B. A. (2009). Do television and radio destroy social capital? evidence from Indonesian villages. *American Economic Journal: Applied Economics*, 1(4):1–33.

Putnam, R. D. (2000). *Bowling Alone: The Collapse and Revival of American Community*. Simon and Schuster, New York.

Rodon, T. and Guinjoan, M. (2022). Beaten ballots: Political participation dynamics amidst police interventions. *Political Science Research and Methods*, 10(1):153–170.

Rodon, T. and Tormos, R. (2023). The burden of a violent past: Formative experiences of repression and support for secession in catalonia. *British Journal of Political Science*, 53(1):269–280.

Rozenas, A., Schutte, S., and Zhukov, Y. M. (2017). The political legacy of violence: The long-term impact of stalin's repression in Ukraine. *The Journal of Politics*, 79(4):1147–1161.

Rozenas, A. and Zhukov, Y. M. (2019). Mass repression and political loyalty: Evidence from Stalin's 'Terror by Hunger'. *American Political Science Review*, 113(2):569–583.

Scott, J. C. (1985). *Weapons of the Weak: Everyday Forms of Peasant Resistance*. Yale University Press, New Haven.

Tilly, C. (1978). *From Mobilization to Revolution*. Addison-Wesley, Reading, MA.

Tur-Prats, A. and Valencia Caicedo, F. (2025). The long shadow of the spanish civil war. *Available at SSRN 3659549*.

V-Dem Institute (2026). Democracy report 2026: Unraveling the democratic era? Technical report, University of Gothenburg, V-Dem Institute. 10th edition.

Vásquez-Cortés, M. (2024). Peer effects and recidivism: Wartime connections and criminality among colombian ex-combatants. *American Political Science Review*, pages 637–652.

Wang, Y. (2021). The political legacy of violence during china's cultural revolution. *British Journal of Political Science*, 51(2):463–487.

Wasow, O. (2020). Agenda seeding: How 1960s black protests moved elites, public opinion and voting. *American Political Science Review*, 114(3):638–659.

Young, L. E. (2019). The psychology of state repression: Fear and dissent decisions in Zimbabwe. *American Political Science Review*, 113(1):140–155.

Zhukov, Y. M. and Talibova, R. (2018). Stalin's terror and the long-term political effects of mass repression. *Journal of Peace Research*, 55(2):267–283.

For online publication only:

State Repression, Trust, and Political Participation

A Additional Tables

Table A.1: Summary Statistics — Municipal Panel (ELIGENDO)

	N	Mean	SD	Min	P25	P50	Max
Panel A: Outcome							
Turnout (%)	70,813	82.78	9.893	0.000	77.94	84.90	100.00
Panel B: Treatment (raw units, 2001 cross-section)							
Exposure, 100 km (baseline treatment)	7,860	1.282	1.207	0.007	0.298	0.953	9.640
Exposure, 100 km, relevance-weighted	7,860	1.874	1.787	0.007	0.414	1.379	13.62
Exposure, 50 km	7,860	0.726	0.965	0.000	0.089	0.317	9.109
Exposure, 50 km, relevance-weighted	7,860	1.059	1.420	0.000	0.115	0.464	12.80
Panel C: Controls Z_m (2001 cross-section)							
Population density, 2001	7,778	257.70	569.34	1.159	46.64	101.19	13,381.78
Share male, 2001	7,859	0.491	0.015	0.336	0.483	0.490	0.611
Mean age, 2001	7,856	42.67	3.931	30.36	40.04	42.28	63.76
Share age 75+, 2001	7,859	0.100	0.040	0.019	0.072	0.093	0.390
Share with high school diploma, 2001	7,859	35.75	8.381	2.752	30.27	35.56	78.74
Share with university degree, 2001	7,858	9.353	4.979	0.000	6.250	8.896	100.00
Employment rate, 2001	7,859	42.90	8.977	14.88	35.39	44.08	71.43
Unemployment rate, 2001	7,859	10.07	8.839	0.000	3.826	5.867	51.32
Share employed in agriculture, 2001	7,859	10.22	9.187	0.000	3.398	7.317	68.51
Share employed in industry, 2001	7,859	37.43	12.50	4.265	27.33	37.01	79.36
Historical Pentapartito vote (avg.)	7,761	60.86	13.96	0.214	52.40	61.74	96.50
Historical PCI vote (avg.)	7,761	26.19	13.26	0.227	15.93	24.92	77.07
TV signal strength	7,784	-0.397	1.006	-9.859	-0.769	-0.235	3.404
Latitude	7,859	43.41	2.576	35.59	41.28	44.54	47.05
Longitude	7,859	11.50	2.787	6.732	9.072	11.18	18.46
Panel D: Terrain (television-signal moderator)							
Free-space TV signal	7,784	-0.062	0.999	-2.566	-0.687	-0.248	6.868
Altitude	7,784	0.357	0.292	0.000	0.118	0.293	2.035
Terrain ruggedness	7,774	221.28	210.26	0.806	42.69	165.47	1,125.41
Panel E: Sample and subgroup indicators							
Hosts a GSF association (excluded from baseline)	72,876	0.028	0.165	0.000	0.000	0.000	1.000
Rural: 2001 density ≤ 300	7,778	0.797	0.403	0.000	1.000	1.000	1.000
Share in the G8 target age band	7,856	0.254	0.036	0.078	0.233	0.257	0.395
Exposure above median, 100 km	7,860	0.500	0.500	0.000	0.000	0.500	1.000
Exposure above median, 50 km	7,860	0.500	0.500	0.000	0.000	0.500	1.000

Notes. Statistics are computed on the baseline estimation sample, which excludes municipalities hosting at least one GSF association. The outcome in Panel A is reported at the municipality $\times$ election-year level. Everything in Panels B to E is a 2001 cross-section and is therefore reported once per municipality: pooling these over nine election years would describe the panel's wave structure rather than the covariate's distribution. The exception is the GSF-host indicator, reported over the full panel because it defines the sample restriction itself. Exposure is reported in *raw* units in Panel B; every regression standardises it on its own estimation sample after the sample filter, so the transform that enters has mean zero and unit standard deviation by construction and a coefficient always reads per within-sample standard deviation. N varies across rows with covariate availability. Bloc vote counts are omitted: the underlying series contains implausible values on a subset of municipality-years and no specification in the paper uses it.

Table A.2: Summary Statistics — Individual Survey Panel (ITANES)

	N	Mean	SD	Min	P25	P50	Max
Panel A: Electoral outcomes							
Voted (1/0)	9,509	0.941	0.236	0	1	1	1
Votes Right (1) vs Left (0), conditional on voting	6,350	0.520	0.500	0	0	1	1
Prior-Left voter still votes	5,289	0.981	0.135	0	1	1	1
Prior-Right voter still votes	5,289	0.976	0.154	0	1	1	1
Prior-Left voter switches to abstention	5,289	0.019	0.135	0	0	0	1
Prior-Right voter switches to abstention	5,289	0.024	0.154	0	0	0	1
Panel B: Trust battery (high-trust indicators)							
Police	4,340	0.242	0.428	0	0	0	1
Army	4,264	0.188	0.391	0	0	0	1
President of the Republic	4,207	0.302	0.459	0	0	0	1
Judiciary	4,242	0.086	0.280	0	0	0	1
Parliament	4,182	0.060	0.237	0	0	0	1
Political parties	4,257	0.017	0.129	0	0	0	1
Trade unions	4,159	0.045	0.208	0	0	0	1
European Union	4,106	0.113	0.317	0	0	0	1
Church	4,314	0.236	0.424	0	0	0	1
Generalized (interpersonal) trust, top half	4,743	0.195	0.397	0	0	0	1
Panel C: Belonging							
Community of belonging (1 local .. 5 world)	4,293	2.631	1.344	1	1	3	5
Community of belonging is local (1/0)	4,293	0.419	0.493	0	0	0	1
Panel D: Non-electoral participation							
Signed a referendum petition	4,296	0.225	0.418	0	0	0	1
Attended a demonstration	4,353	0.123	0.328	0	0	0	1
Interest in politics (1–4)	8,424	2.873	0.848	1	2	3	4
Panel E: Treatment and controls							
Exposure, 100 km (raw)	9,989	1.410	1.371	0.014	0.394	1.011	9.640
Gender	9,994	1.505	0.500	1	1	2	2
Age	9,990	42.83	18.74	1	28	42	96
Education level	9,990	4.553	1.930	1	3	4	12
Religious orientation	9,964	1.651	6.187	1	1	1	99
Employed	9,994	1.655	3.807	1	1	2	99

Notes: Summary statistics on the baseline ITANES estimation sample (municipalities hosting at least one GSF association excluded), at the respondent × wave level. The trust and belonging batteries are fielded in the 2001 and 2006 waves only, which accounts for the smaller N in Panels B and C. Codes 99 in the religion and employment variables are missing-value codes carried over from the survey and are dropped in estimation.

Table A.3: Effect of Exposure on Non-Electoral Political Engagement (ITANES)

Outcome	Coef.	S.E.	95% CI	N
Signed a referendum petition	+0.019	(0.015)	[−0.010, 0.049]	4,296
Attended a demonstration	+0.033	(0.039)	[−0.044, 0.109]	4,353
Interest in politics (1–4)	−0.066	(0.072)	[−0.208, 0.076]	8,398
Individual-Level Controls			Yes	
Municipality FE			Yes	
Province-by-Year FE			Yes	

Notes. Each row reports the coefficient on $Exposure_m \times Post_t$ from a separate regression following Equation 3 in pooled form. The first two outcomes are binary indicators estimated by linear probability models; interest in politics is a four-point scale. None of the three is distinguishable from zero, and the two participation measures are positive in sign. The event-study form gives the same picture, with the single exception of a marginally negative interest coefficient in 2006 (-0.092, $p = 0.094$) that neither precedes nor follows a pattern. These outcomes test whether the electoral withdrawal documented in Section 4 extends beyond the ballot box; it does not. All regressions include municipality fixed effects, province-by-year fixed effects, and individual-level time-varying controls (gender, age, education, employment status, religious orientation). Standard errors clustered at the municipality level in parentheses. Municipalities hosting at least one GSF association are excluded.

Table A.4: What Exposure Is Correlated With: Municipality Characteristics in 2001

	No fixed effects		Province FE		
Covariate (2001)	Coef.	(S.E.)	Coef.	(S.E.)	N
Population density	0.0005^{*}	(0.0003)	0.0003^{***}	(0.0001)	7,759
Share male	1.634	(1.990)	-1.004^{*}	(0.517)	7,834
Mean age	0.0068	(0.0185)	-0.035^{***}	(0.006)	7,831
Share age 75+	−2.965	(2.034)	-3.169^{***}	(0.541)	7,834
Share with high school diploma	0.016^{**}	(0.007)	0.0064^{***}	(0.0020)	7,834
Share with university degree	0.0077	(0.0072)	0.0071^{***}	(0.0025)	7,833
Employment rate	0.045^{***}	(0.009)	0.022^{***}	(0.005)	7,834
Unemployment rate	-0.042^{***}	(0.007)	0.0015	(0.0021)	7,834
Share employed in agriculture	-0.031^{***}	(0.007)	-0.0074^{***}	(0.0014)	7,834
Share employed in industry	0.025^{***}	(0.005)	0.0028	(0.0017)	7,834
Historical Pentapartito vote (avg.)	−0.0017	(0.0057)	-0.0060^{***}	(0.0012)	7,736
Historical PCI vote (avg.)	0.011^{**}	(0.005)	0.0069^{***}	(0.0013)	7,736
TV signal strength	0.190^{***}	(0.047)	0.094^{***}	(0.023)	7,759
Latitude	0.172^{***}	(0.030)	−0.211	(0.182)	7,834
Longitude	-0.168^{***}	(0.028)	−0.056	(0.136)	7,834

Notes: Each row reports two separate regressions of the standardised 100 km exposure index on the stated 2001 municipality characteristic, entered one at a time: the first pair of columns without fixed effects, the second with province fixed effects. The province-fixed-effects column is the relevant one, since our estimating equation identifies β from within-province contrasts. The cross-section is the 2001 municipal sample excluding municipalities that host at least one GSF association; N varies with covariate availability. Coefficients are in the natural units of each covariate and are therefore not comparable in magnitude across rows; the table is read for sign and significance, not for effect size. This table is descriptive and is *not* a balance test in the experimental sense: exposure is not randomly assigned, and we do not claim it is orthogonal to observables. Every covariate listed here enters the estimating equation interacted with election-year indicators (Equation 2), so differential trends along each of these dimensions are absorbed nonparametrically in every post-period year. Standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A.5: Methods of Violence Reported in Victim Testimony

	Bolzaneto (n=45)	Diaz (n=30)	All (n=75)
Beaten with batons	78	93	84
Kicked or punched	96	60	81
Beaten while already down	76	73	75
Insults, spitting, humiliation	93	60	80
Forced songs / fascist chants	82	73	79
Thrown or dragged	47	60	52
Forced stress positions	93	30	68
Handcuffed / restrained	87	10	56
Forced undressing / nudity	82	27	60
Verbal threats	71	27	53
Pepper spray / gas	53	17	39
Denial of medical care	20	63	37
Denial of water / toilet	33	3	21

Notes. Cell entries are the percentage of victim hearings whose testimony describes each method, by trial. Coding is a conservative keyword-based lower bound over speaker-attributed testimony (a hearing is flagged only on explicit mention); the top block covers acute assault, the lower block custodial abuse.

Table A.6: Lasting Harm Reported in Victim Testimony

	Bolzaneto (n=45)	Diaz (n=30)	All (n=75)
Persistent fear	76	80	77
Panic / anxiety (PTSD-type)	40	40	40
Intrusive memories / nightmares	31	20	27
Lasting physical injury	42	30	37
Hopelessness / resignation	31	27	29
Loss of trust	20	17	19
Lasting shame	20	3	13

Notes. Percentage of victim hearings reporting each lasting effect.

Table A.7: Threat Anticipated by the Police

	% of police hearings (n=53)
Acting on received orders	85
Intelligence / radio reports of threat	74
Perceived threat to property	49
Expected black bloc / infiltrators	43
Framed as not a peaceful protest	36
Expected armed / masked crowd	34
Expected to be outnumbered	17

Notes. Percentage of police/defendant hearings whose testimony invokes each prior expectation.

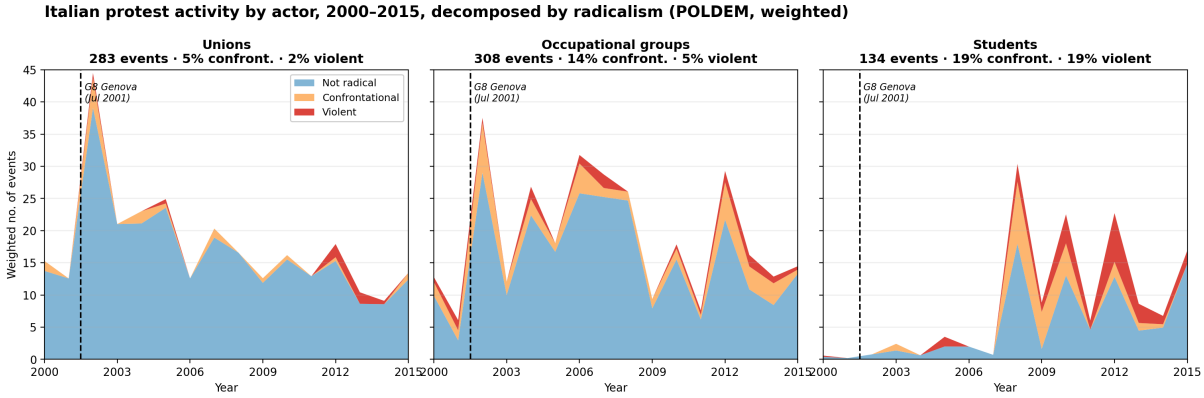
Table A.8: The Trust Battery Under Two Codings of the Same Item

Institution	(1) Top category $\mathbf{1}[molta]$ *probability*	(2) Top half $\mathbf{1}[abbastanza,\ molta]$ *probability*
Police	−0.173*** (0.014)	+0.144*** (0.034)
Church	+0.008 (0.014)	+0.174*** (0.009)
President of the Republic	−0.272*** (0.023)	−0.022** (0.011)
Army	−0.075*** (0.019)	−0.028* (0.015)
Judiciary	−0.148*** (0.018)	−0.057 (0.041)
Parliament	+0.020** (0.008)	+0.052 (0.043)
Political parties	−0.023*** (0.003)	+0.088 (0.059)
Trade unions	+0.048*** (0.008)	+0.052 (0.039)
European Union	−0.019 (0.024)	+0.004 (0.025)
Generalized trust	—	+0.196*** (0.074)

Notes. Each cell is the coefficient on $Exposure_m \times Post_t$ from a separate pooled regression following Equation 3, with municipality-clustered standard errors in parentheses. Column (1) is the coding reported in Table 1; columns (2) is a specification which holds sample, controls, fixed effects, clustering and the standardisation of exposure fixed and varies only how the four-point item is converted into an outcome. Generalized trust has no ordinal counterpart, since the 2001 instrument is a two-point forced choice and the 2006 instrument a ten-rung ladder, so it is reported under coding (2) only. Municipalities hosting at least one GSF association are excluded. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

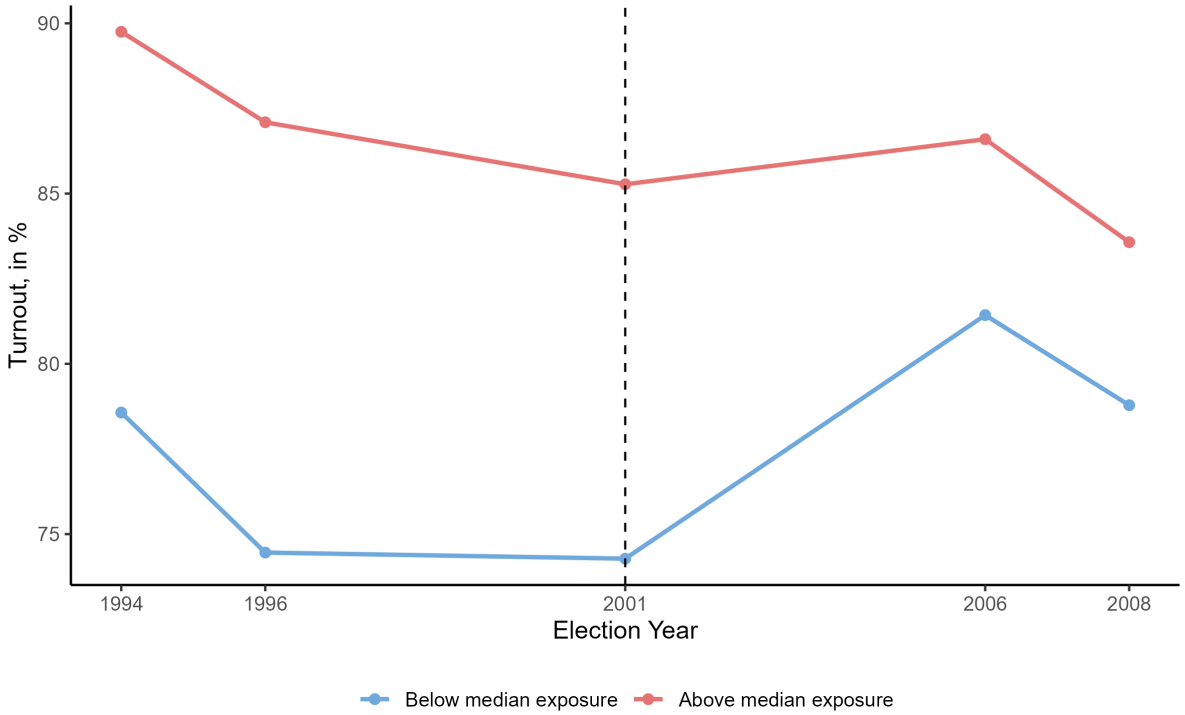
B Additional Figures

Figure B.1: Italian Protest Activity by Actor, 2000–2015



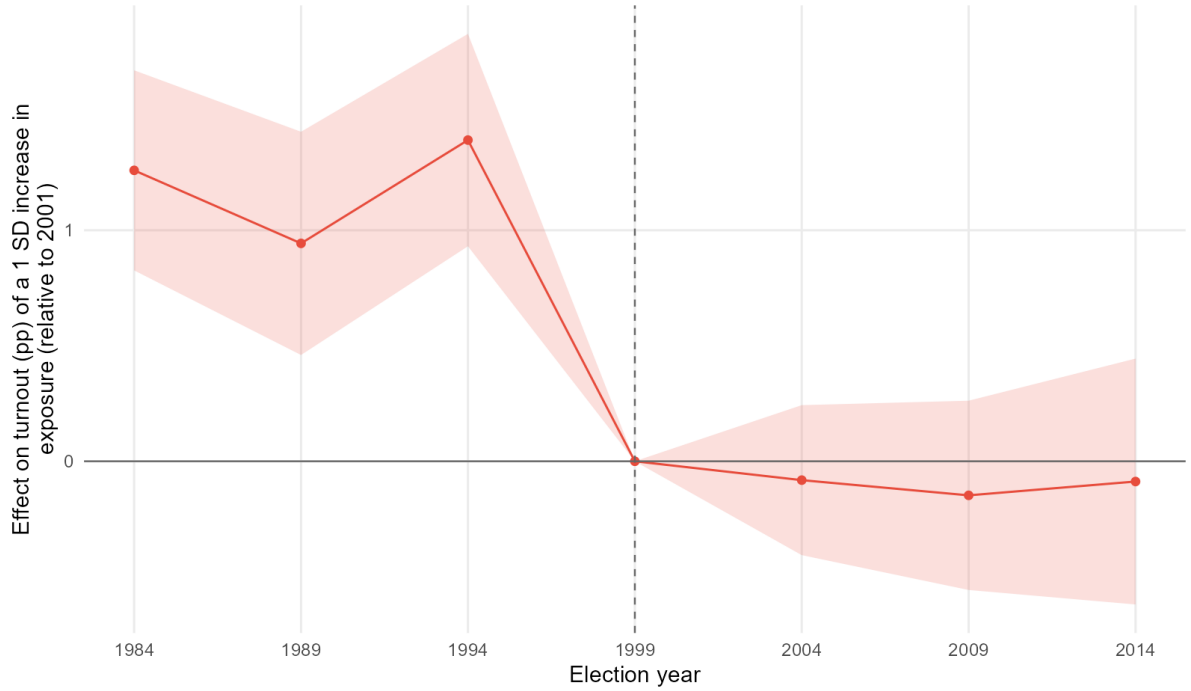
Notes. The figure plots protest activity in Italy from the POLDEM Protest dataset, which codes individual protest events reported in international newswires for 30 European countries. The sample is restricted to Italy. Each of the three upper panels corresponds to one type of protest actor: trade unions (pooling the private-sector, public-sector, mixed and unspecified union codes), occupational groups, and students. Actor codes are non-exclusive, so an event mobilising more than one actor enters more than one panel; the panels are therefore not mutually exclusive and do not sum to the Italian total. Roughly one third of Italian events carry no identified actor and are excluded from all panels. Within each upper panel, yearly activity is stacked into three mutually exclusive radicalism categories from the dataset's `radical_action` variable: not radical, confrontational (disruptive but non-violent, e.g. blockades and occupations), and violent. Heights are weighted event counts: each event is weighted by the dataset's sampling weight, which corrects for the newswire sampling design, so the vertical axis approximates the true number of events rather than the number sampled. Panel subtitles report the actor's total weighted events over 2000–2015 and the share of those events that are confrontational and violent. The lower panel plots, for each actor, the share of that actor's yearly events coded violent. Source: POLDEM.

Figure B.2: Raw Turnout Trends by Exposure Group (Median Split)



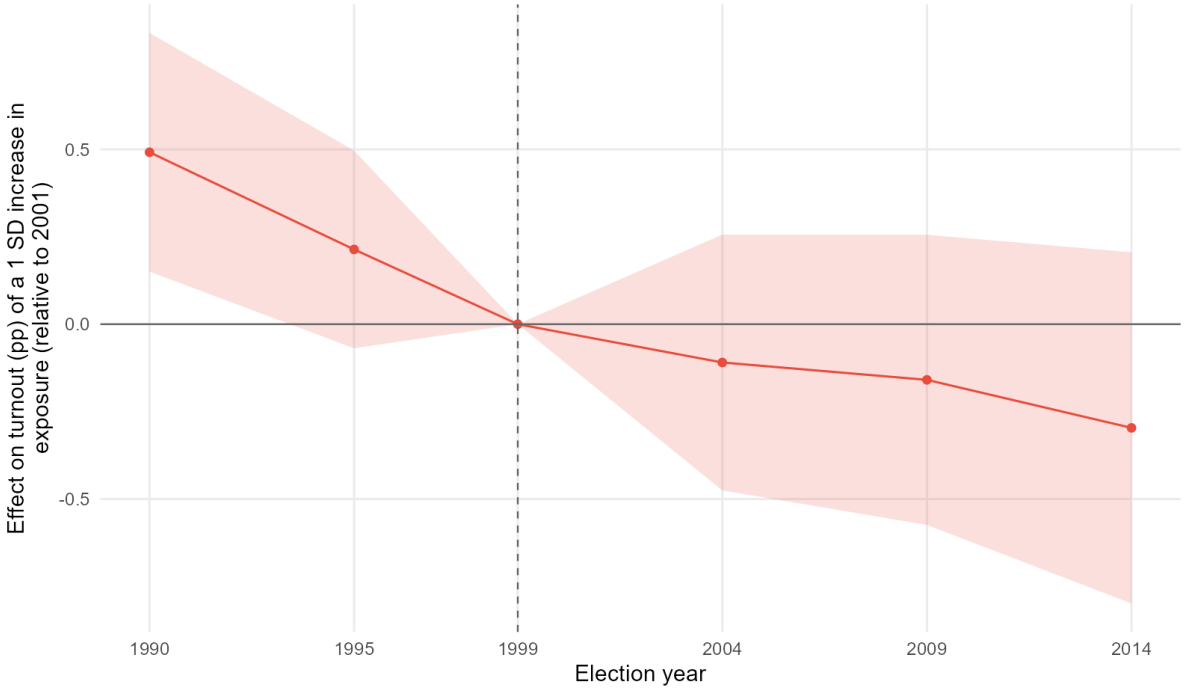
Note. The figure plots average raw national-election turnout (in %) over the 1994–2008 elections, separately for municipalities above and below the median of the 100 km exposure index. No controls, fixed effects or weights are applied. The vertical dashed line marks the July 2001 G8 summit. The between-group gap is 11.3 percentage points in 1994, 12.8 in 1996 and 10.9 in 2001, and then narrows to 5.3 in 2006 and 4.7 in 2008. This is the model-free counterpart to the event-study evidence in Figure 4. Because the series are unconditional, the levels reflect the covariate differences that the regression conditions on, and the post-2001 narrowing is driven more by a rise in the below-median group than by a fall in the above-median one; the figure is read for the timing of the divergence rather than for its magnitude.

Figure B.3: Turnout in European Parliament Elections, 1984–2014



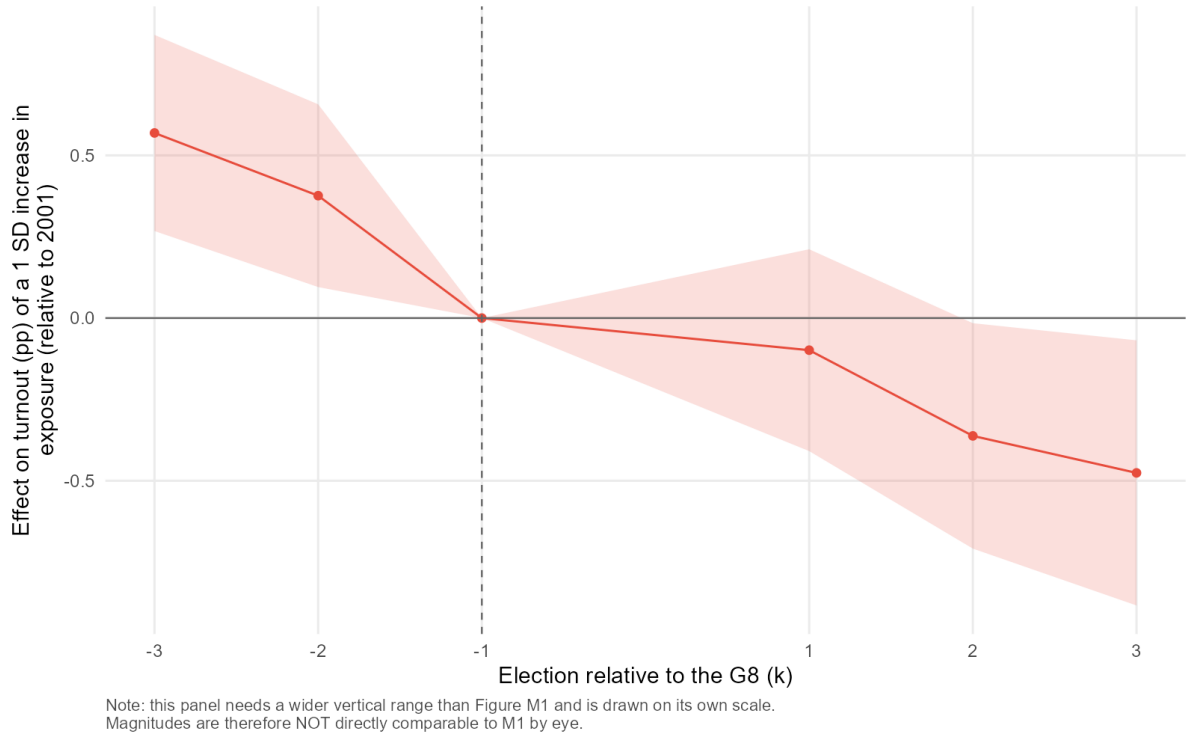
Note. Event-study coefficients from the baseline specification estimated on European Parliament election turnout, seven waves 1984–2014. **The omitted year is 1999**, the last pre-treatment wave, since there is no 2001 European election. The covariate block is lifted from the national panel, so treatment and controls are identical across the two datasets and any difference in β is attributable to the outcome alone. Coefficients are interactions between year indicators and the standardised exposure measure; shaded areas are 95% confidence intervals with municipality-clustered standard errors. Note the vertical scale differs from Figure 4 and magnitudes are not comparable by eye. The three pre-1999 coefficients are large, positive and significant, and the three post-1999 coefficients are small and insignificant: there is a differential pre-trend and no break at the summit. $N = 54{,}166$.

Figure B.4: Turnout in Municipal Elections: Synchronous Calendar Years



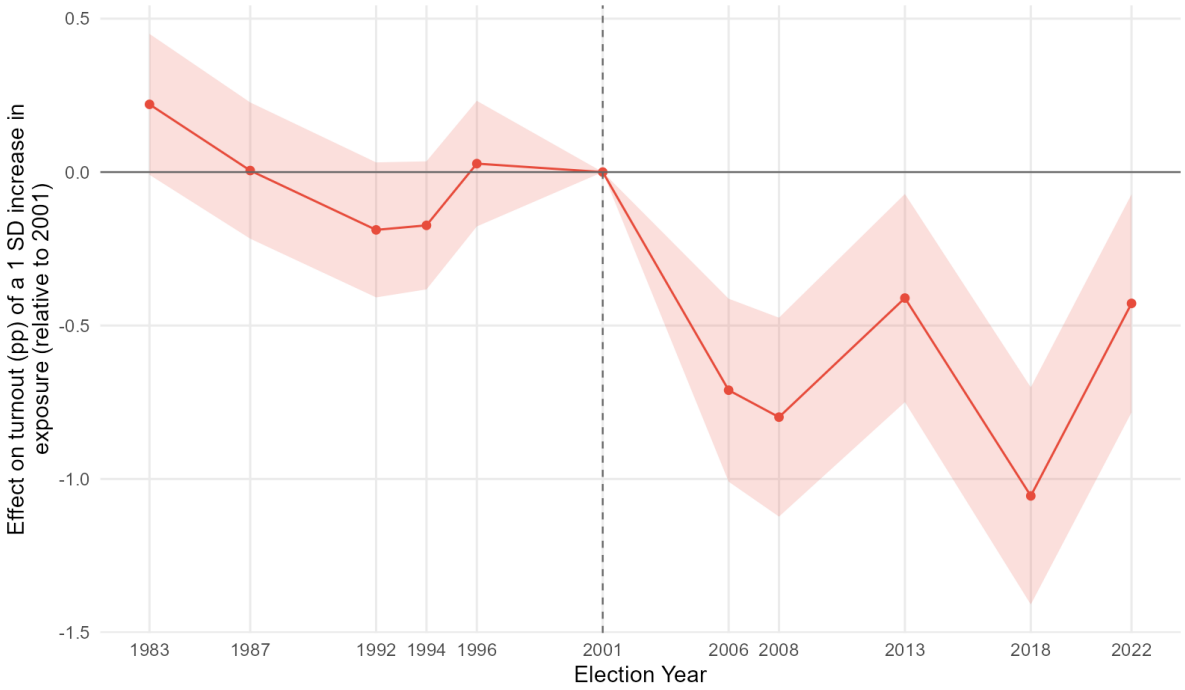
Note. Event-study coefficients from the baseline specification estimated on municipal (*comunali*) election turnout, restricted to the six main election days — 1990, 1995, 1999, 2004, 2009 and 2014 — with **1999 omitted**. Restricting to these six days makes the panel synchronous, so it has the same shape as the national and European event studies and the three sit on one clock; Italy holds one large municipal round a year, but the years differ enormously in size, and a year-by-year event study on the full panel would be dominated by the small years. Controls are lifted from the national dataset. Shaded areas are 95% confidence intervals with municipality-clustered standard errors. Vertical scale differs from Figure 4. $N = 27{,}090$.

Figure B.5: Turnout in Municipal Elections: Event Time



Note. Event-study coefficients from the baseline specification estimated on municipal election turnout indexed in **event time** rather than calendar time: $k = -1$ is the municipality's last election before 22 July 2001 and $k = +1$ its first after, with no zero, windowed to $k \in [-3, +3]$ over all 130 election dates. Municipal elections are asynchronous, so there is no common election year to pivot on. Fixed effects are province-by-*calendar*-year, deliberately not province-by-k, so that national shocks are absorbed while the treatment path is indexed in event time. Two caveats: $k = -3$ lands around 1986–1991, where three years are missing from the municipal archive, so some municipalities are mis-indexed there ($k = -2$ and all post-treatment periods are clean); and $k = -1$ is not a uniform distance from treatment, averaging 2.5 years before the summit and ranging from 0.2 to 16.2 years. Figures B.4 and B.5 are the same data on two clocks and are reported together, never one alone. $N = 42{,}396$.

Figure B.6: Robustness: Effect of Exposure on Turnout, Panel Extended to 2022



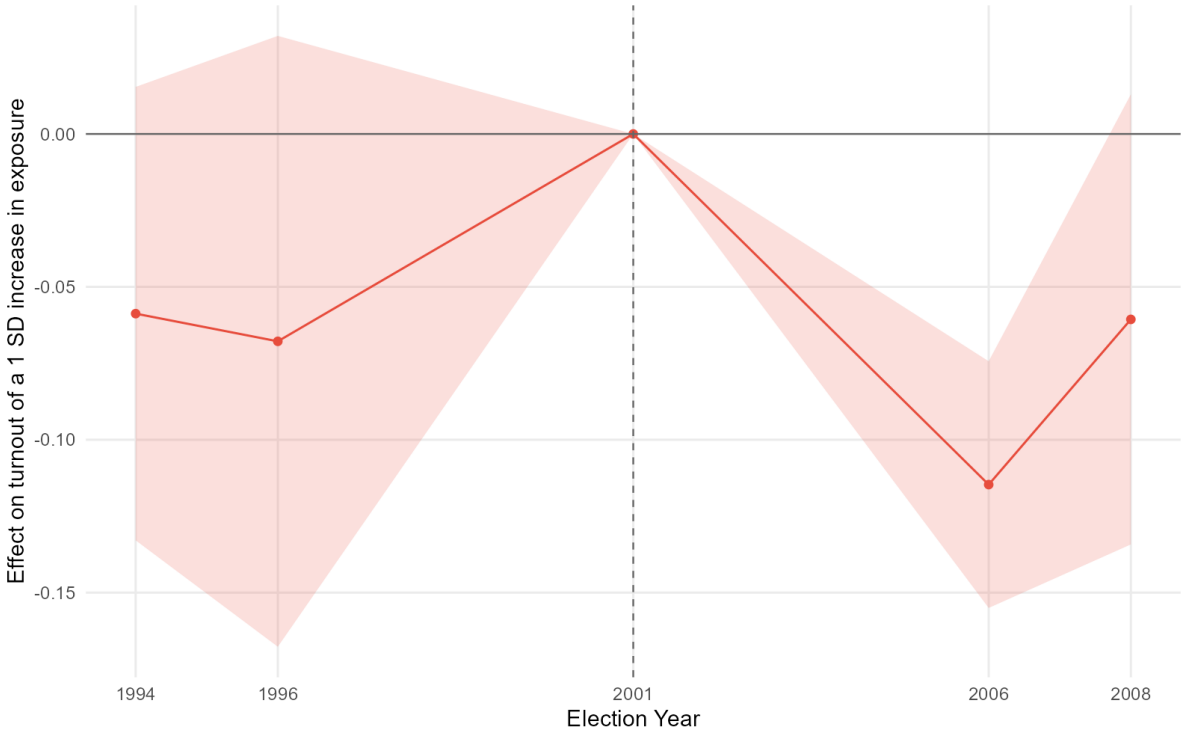
Note. Event-study coefficients from the baseline specification (Equation 2) estimated on national elections from 1983 to 2022, adding the 2018 and 2022 elections to the baseline panel. Coefficients are interactions between year indicators and the standardised exposure measure, with 2001 as the reference year. Shaded areas are 95% confidence intervals with municipality-clustered standard errors; controls and fixed effects as in Equation 2. Note the vertical scale differs from Figure 4. The M5S entered national politics in 2013, so the 2018 and 2022 coefficients are estimated under a changed party system and are read with the caution discussed in Section 3.

Figure B.7: Placebo in Time: Counterfactual Summit Dates



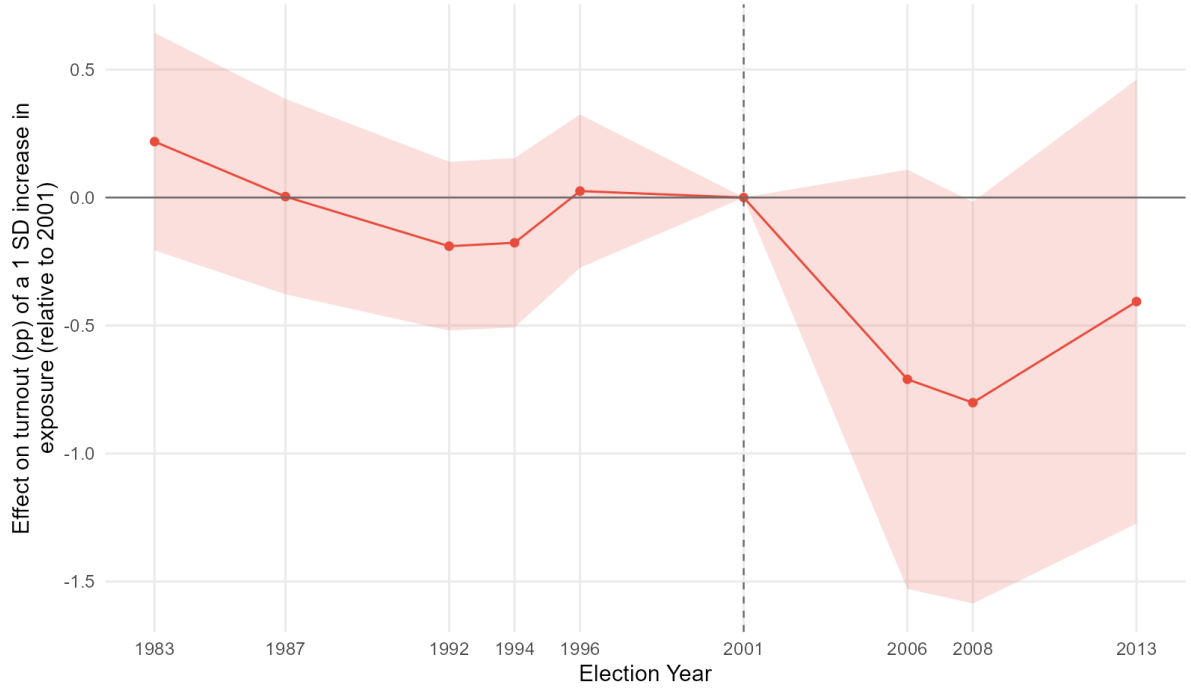
Note. Event-study coefficients from the baseline specification (Equation 2) estimated with the treatment date moved to a counterfactual summit in 1991, 1993 and 1995. For each false date the sample is truncated at 2001, so the false post-period contains only elections that precede the real event and the true treatment cannot contaminate the placebo. The omitted category is the last election before the false date. The fourth panel repeats the exercise at the true date on the full sample and is the baseline of Figure 4 drawn on the same axes for comparison. Shaded areas are 95% confidence intervals with municipality-clustered standard errors; controls and fixed effects as in Equation 2. A false date that reproduced the post-2001 step down would indicate that the design detects a break wherever it is asked to look; none does.

Figure B.8: Individual-Level Replication: Effect of Exposure on Voting (ITANES)



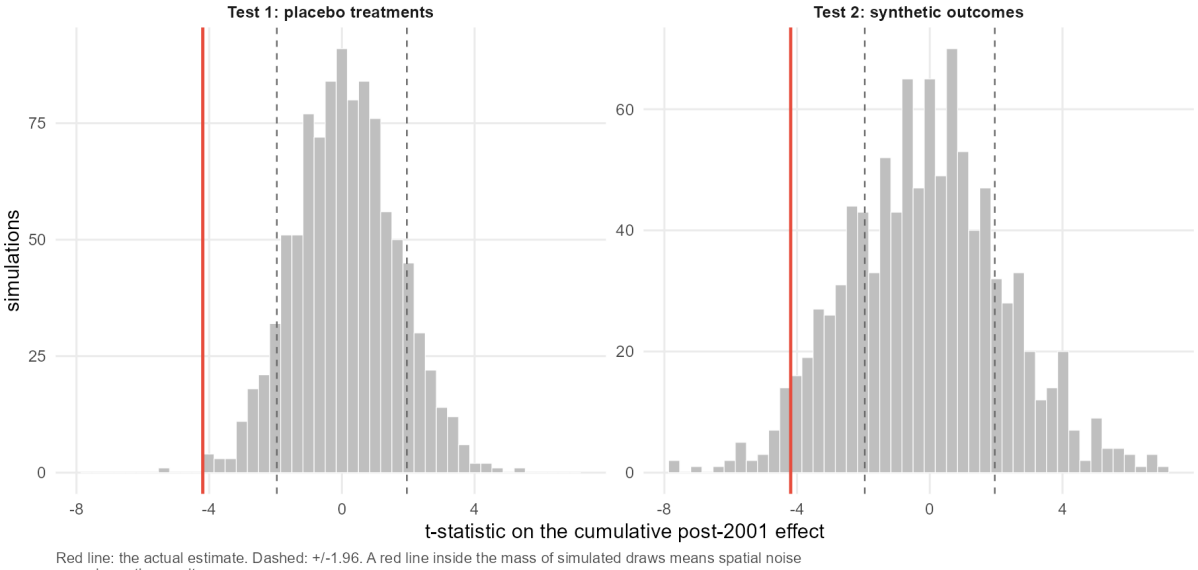
Note. Event-study coefficients from a fixed effects regression of an individual-level indicator for having voted on exposure, estimated on the ITANES repeated cross-section. The vertical axis is in probability units, so −0.10 is a ten percentage point fall in the probability of voting. Coefficients are interactions between year indicators and the standardised exposure measure, with 2001 as the reference year. Shaded areas are 95% confidence intervals with standard errors clustered at the municipality level. The regression includes municipality fixed effects, province-by-year fixed effects, and individual-level time-varying controls (gender, age, education, employment status, religious orientation). Municipalities hosting at least one GSF association are excluded. The pre-2001 coefficients are imprecise, since only two pre-treatment waves are available and the sample is two orders of magnitude smaller than the municipal panel; individual turnout in ITANES also averages 94%, leaving little variation to explain. Direction and timing replicate the municipal result of Figure 4; precision does not. $N = 9{,}468$.

Figure B.9: Robustness: Effect of Exposure on Turnout (Province-Level Clustering)



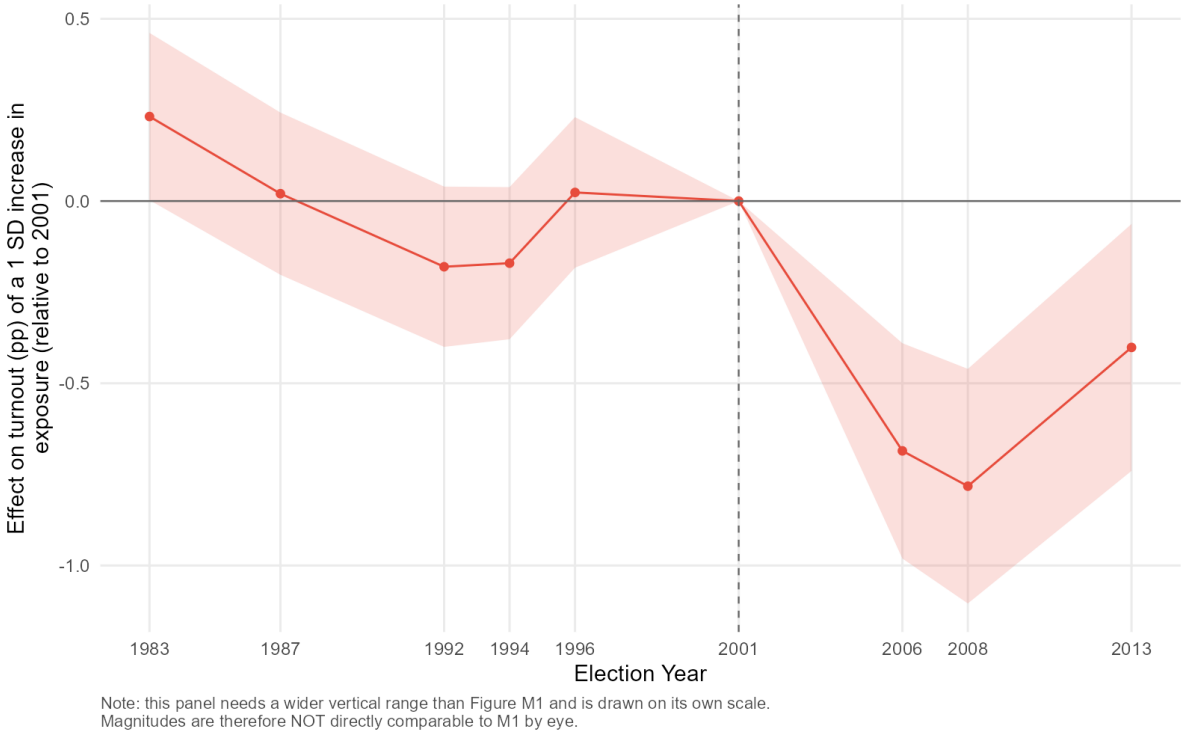
Note. Event-study coefficients from the baseline specification (Equation 2) with standard errors clustered at the **province** level (112 clusters) rather than the municipality, allowing arbitrary correlation within larger geographic units. Point estimates are identical to Figure 4 by construction; only the confidence intervals change. Shaded areas are 95% confidence intervals. $N = 69{,}504$.

Figure B.10: Spatial-Noise Diagnostics: Reference Distributions for the Cumulative Effect



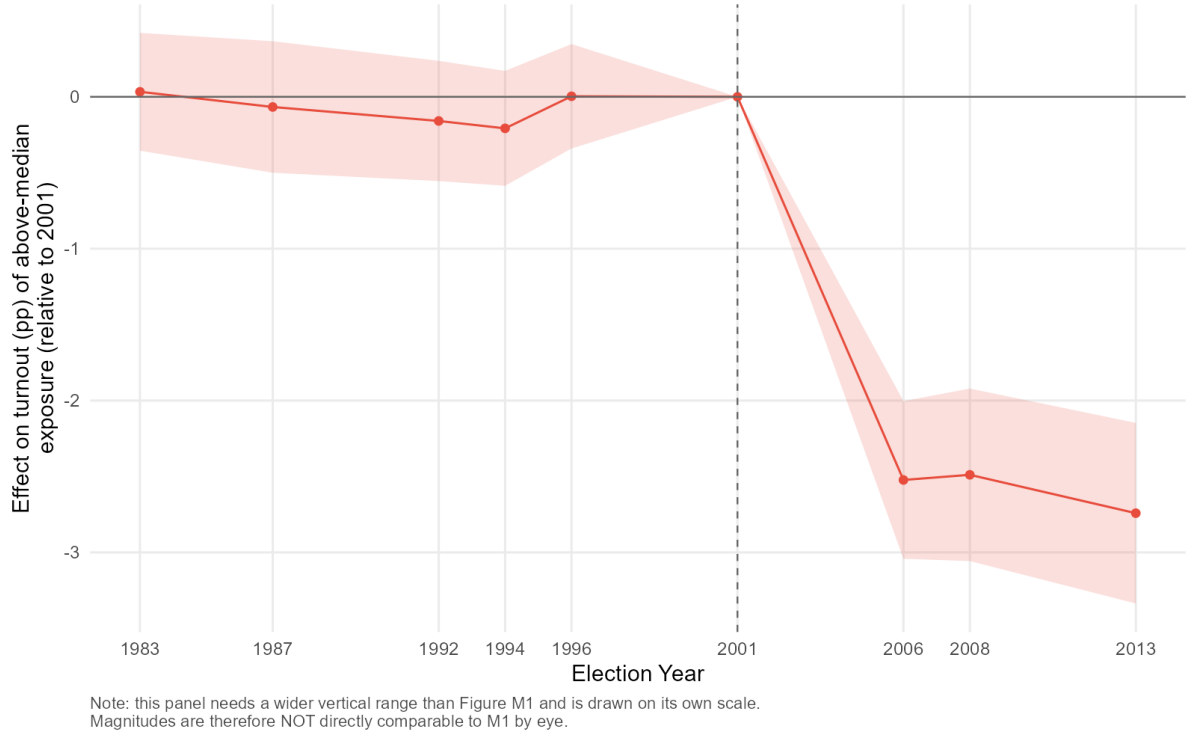
Note. The figure reports the two spatial-noise diagnostics of Conley and Kelly (2025) for the cumulative post-2001 turnout effect. Both panels plot the distribution of t-statistics on the cumulative effect across simulations, with the solid red line marking the t-statistic of the actual estimate (-4.19) and the dashed lines marking ± 1.96. *Test 1* re-estimates Equation 2 on placebo treatments drawn to have the same spatial structure as the true exposure measure, holding the outcome fixed. *Test 2* holds the treatment fixed and re-estimates on synthetic outcomes built by adding spatially correlated Gaussian-process noise to a de-meaned, trend-augmented turnout series. A red line sitting inside the mass of simulated draws would mean that spatially correlated noise alone can reproduce the estimate; here the red line lies at or beyond the left edge of both reference distributions.

Figure B.11: Robustness: Effect of Exposure on Turnout (Association Relevance Index)



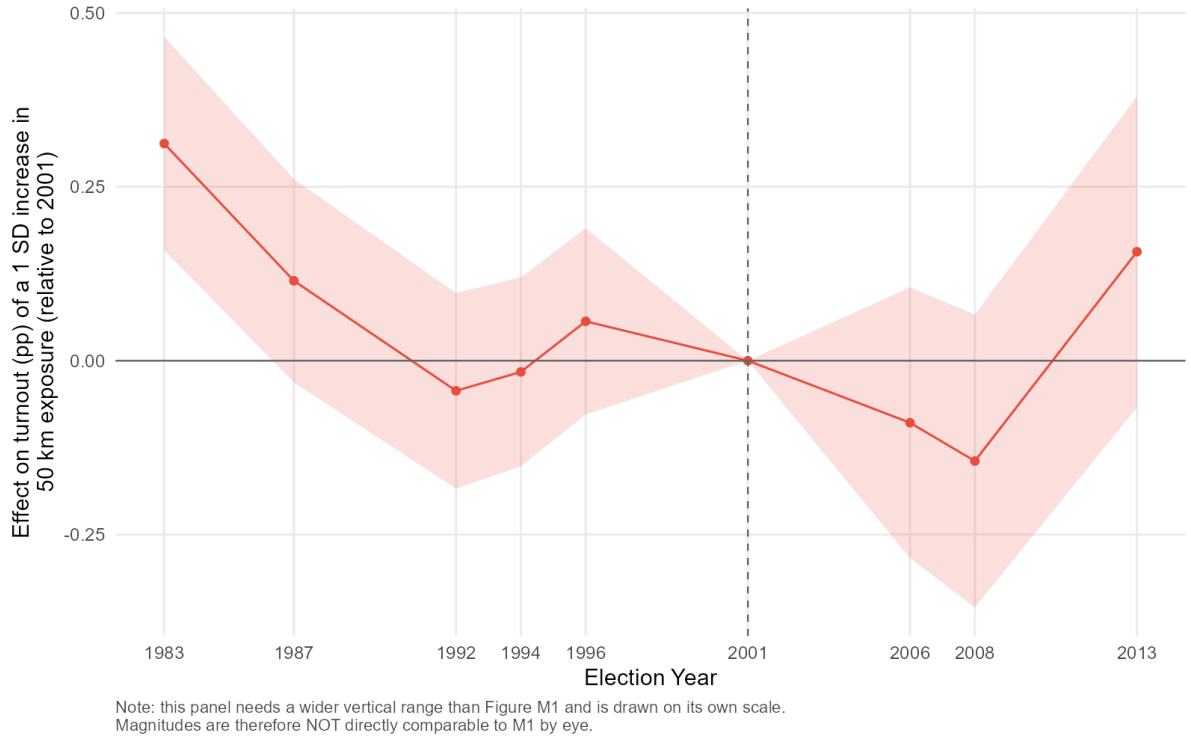
Note. Event-study coefficients from the baseline specification (Equation 2) where the treatment uses the rarity-weighted association relevance index (Section 3) in place of the simple presence dummy. Coefficients are interactions between year indicators and the standardised exposure measure, with 2001 as the reference year. Shaded areas are 95% confidence intervals with municipality-clustered standard errors. $N = 69{,}504$.

Figure B.12: Robustness: Effect of Exposure on Turnout (Binarised Treatment at the Median, 100 km)



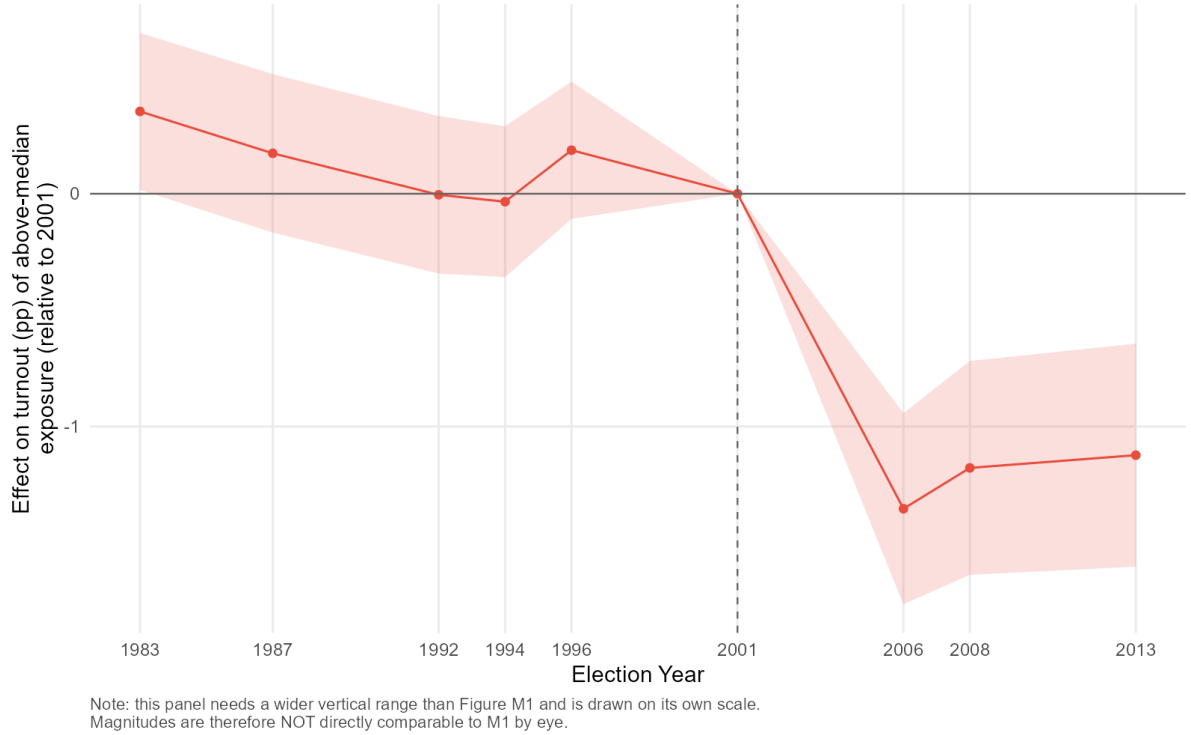
Note. Event-study coefficients where the treatment is a dummy equal to one for municipalities with above-median exposure on the 100 km index. This imposes no linear-in-dose assumption and is insensitive to the tail of the exposure distribution. The vertical axis reports the turnout effect in percentage points of above- versus below-median exposure, relative to 2001. Shaded areas are 95% confidence intervals with municipality-clustered standard errors; controls and fixed effects as in Equation 2. $N = 69{,}504$.

Figure B.13: Robustness: Effect of Exposure on Turnout (Exposure Radius Narrowed to 50 km)



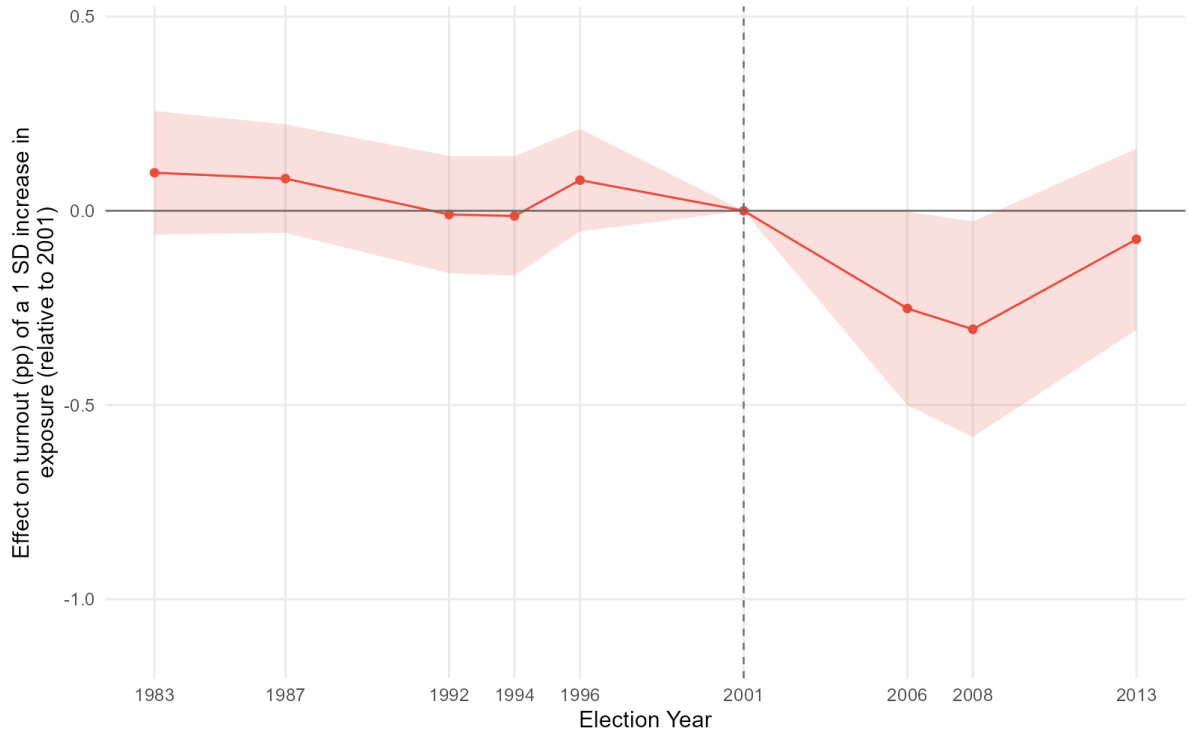
Note. Event-study coefficients from the baseline specification (Equation 2) where the market-access exposure index is computed over associations within 50 km rather than 100 km, weighted by inverse road distance and standardised on the estimation sample exactly as in the baseline. Coefficients are interactions between year indicators and the standardised 50 km exposure measure, with 2001 as the reference year. Shaded areas are 95% confidence intervals with municipality-clustered standard errors; controls and fixed effects as in Equation 2. Note the vertical scale differs from Figure 4 and magnitudes are not comparable by eye. In this continuous form the post-2001 coefficients are small and individually insignificant, and the 1983 coefficient is positive and significant, so the pre-period is also less clean than at the baseline radius. At this radius 8.2% of municipalities have no association in range, against none at 100 km, and the distribution is badly right-skewed. Figure B.14 applies a median split to this same index and does reproduce the baseline pattern; Section 6 discusses why the two forms diverge. $N = 69{,}504$.

Figure B.14: Robustness: Binarised Treatment at the Median of the 50 km Index



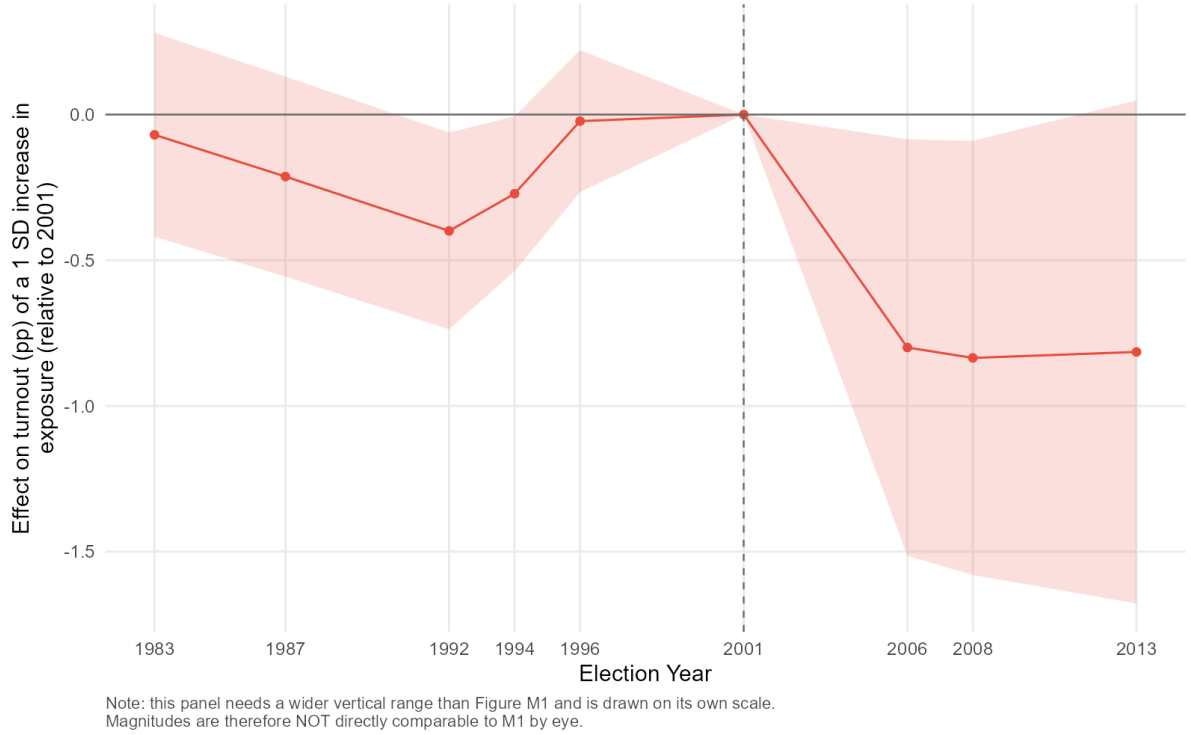
Note. Event-study coefficients where the treatment is a dummy equal to one for municipalities with above-median exposure on the *50 km* index. This is the companion to Figure B.13, which uses the same 50 km index as a continuous dose, and to Figure B.12, which applies the same median split to the 100 km index. The vertical axis reports the turnout effect in percentage points of above- versus below-median 50 km exposure, relative to 2001; the scale differs from Figure 4. Shaded areas are 95% confidence intervals with municipality-clustered standard errors; controls and fixed effects as in Equation 2. Unlike the continuous specification at the same radius, the binarised one reproduces the baseline pattern: flat pre-trends and a large, persistent post-2001 step down. Section 6 discusses why the two diverge. $N = 69{,}504$.

Figure B.15: Robustness: Effect of Exposure on Turnout (Including Municipalities Hosting Associations)



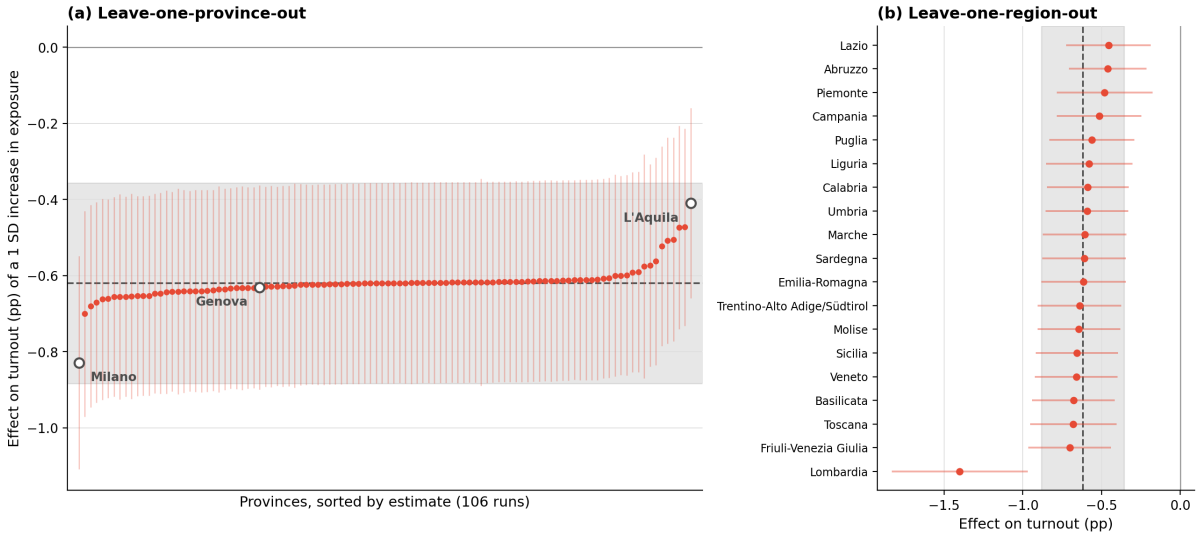
Note. Event-study coefficients from the baseline specification (Equation 2) estimated on the full sample, i.e. *including* municipalities that host at least one GSF association, which are excluded in the baseline. Attenuation is the predicted result and not a failure: re-including hosts dilutes the differential-proximity variation the design rests on. Shaded areas are 95% confidence intervals with municipality-clustered standard errors. $N = 71{,}517$.

Figure B.16: Robustness: Effect of Exposure on Turnout (Region×Year FE, Province Clusters)



Note. Event-study coefficients from a fixed effects regression of voter turnout on exposure, with **region**-by-year and municipality fixed effects and standard errors clustered at the province level. Region-by-year absorbs more aggregate time-varying shocks than province-by-year and widens the within-cell comparison set. Coefficients are interactions between year indicators and the standardised exposure measure, with 2001 as the reference year. Shaded areas are 95% confidence intervals. Controls as in Equation 2. $N = 69{,}493$.

Figure B.17: Robustness: Leave-One-Out



Note. Each point is the pooled *Exposure* × *Post* coefficient from a separate re-estimation of the baseline specification (Equation 2, pooled form) that excludes one geographic unit. Panel (a) excludes one province at a time, giving 106 estimates, sorted by size; panel (b) excludes one region at a time, giving 19. The dashed line and shaded band are the full-sample estimate, −0.619, and its 95% confidence interval. Bars are 95% confidence intervals with standard errors clustered at the municipality level. Province counts refer to provinces present in the estimation sample after association-hosting municipalities and observations with missing covariates are dropped, and are therefore fewer than the 112 clusters used for the province-clustered standard errors in Figure B.9. All 125 estimates are significant at the 5% level. Three runs are labelled in panel (a): excluding the province of Genova leaves the estimate at −0.63, so the result does not rest on the municipalities nearest the summit; excluding Milan, the densest and most heavily exposed province, *strengthens* it to −0.83, as does excluding Lombardia as a whole in panel (b), which takes the estimate to −1.40.

C Historical Background in Detail

Section 2 gives the compressed account of the summit that the empirical work requires. This appendix sets out the fuller historical record for readers who want it. It does not repeat the evidence on police planning, the trial depositions, or post-2001 protest activity, all of which are in the body of the paper.

C.1 The Anti-Globalization Movement

Through the 1990s a wave of protests emerged across Western countries criticising economic globalization and laissez-faire economic policy and demanding lower inequality between the developed and the developing world. These movements arose and developed almost independently before converging under a common label, the *Anti-Globalization Movement*.

The movement became known worldwide in November 1999, when forty thousand protesters, mainly students, unemployed people, union members and activists, forced the interruption of the WTO meeting in Seattle with an unauthorised sit-in. The initially peaceful protest escalated into urban guerrilla warfare and around 150 arrests. After Seattle, protests erupted, with varying degrees of violence, alongside every major international summit.

Italy was set to host the Group of 8 meeting in the summer of 2001, and the centre-left government in office until June 2001 chose Genova as the location. From the previous year, and with intensifying coordination, Italian and international movements began organising a protest against the summit, the *counter-summit*, to be held during the days of the G8. These associations came together to form the Genoa Social Forum.

C.2 The GSF and the Making of the Counter-Summit

The GSF was founded in December 1999 as a network of movements and associations. Its final composition, in February 2001, comprised 1,187 associations, of which 1,009 were Italy-based. It included national and international organisations: cultural associations, parties, unions, religious movements, students' associations, NGOs, and a large number of single-purpose associations.[1] Its purpose was to organise the counter-summit, a non-violent protest criticising the functioning of the summit and advancing demands consistent with the anti-globalization agenda, agreed upon by all subscribing associations.

Alongside this political preparation the GSF handled the logistics: transport, accommodation and meals for participants, and such security measures as lay within its capacity. This was done jointly with local and, in part, national authorities. Every demonstration

[1] Environmental groups, animal rights movements, LGBTQIA+ movements, health promotion associations, and so on.

and activity forming part of the counter-summit had been agreed in advance with the police and with the government's local representative. Specific facilities in Genova, public schools among them, were designated as shelters or camps where participants could sleep and eat.

C.3 Transport to Genova

Participants came to Genova from across Italy and much of Europe in the days before the summit. Security measures had been enforced in the city: the airport and port were closed and only specific stations were open. The motorway remained open but access to the city was subject to police control. Special trains were organised to carry people in. Reaching Genova therefore meant travelling by train, or by car or coach.

Most GSF associations organised buses and trains to carry their members. The remaining protesters came independently, by train or car. Our reconstruction, based on conversations with the GSF's logistical coordinators, indicates that most participants travelled on the buses and trains organised by the associations even when they were not formal members of any of them.

The trial record corroborates this from the participants' own words. Victims describe reaching Genova through their local association chapters: one testifies to travelling on coaches chartered from his province by the unions, by Rifondazione and by the Democratic Left; another to departing from Imola on one of the special trains organised for the counter-summit; another to joining an organised group that assembled at Bologna station. Others identify themselves as members of the GSF or of local union chapters. These accounts are incidental to the proceedings, which is part of why we find them credible, and they confirm that participation was routed through the geographically dispersed associational network on which our exposure index is built.

C.4 Who Was Protesting

More than 200,000 demonstrators were in Genova on 19 July, the day before the summit began. No official data on the protesters exist. Andretta et al. (2002) collected survey data from 763 participants during the protest, recording who they were, which organisation they were affiliated with, and which party they felt closest to. Participants were primarily mobilised through the GSF's own organizational networks: 41.2% declared themselves official members of a participating association; 39.5% came from the Attac/-Fiom/Rifondazione/ARCI area, 29.9% from the Lilliput network, 24.5% from the *centri sociali*, 3.8% from local organisations, with a 6% residual.

Consistent with this structure, most protesters reported learning about the GSF through their own organisation (36%), well ahead of family and friends (23%), mass media (22.7%) and the internet (4.2%). Women made up 42.8% of the sample and 96% were

Italian. 42.4% came from Northern Italy and 35.3% from the Centre, with the remaining 22.3% from the South, including Lazio, and the islands. Students were the largest occupational group (49.9%), followed by the unemployed and precariously employed (10.7%), salaried employees in the public and private sectors (12.6%), knowledge workers (6.8%), entrepreneurs (7.1%), artisans and shopkeepers (2.3%), and other statuses (10.6%). Left–right self-placement was overwhelmingly on the left and centre-left, with negligible shares identifying with the centre or the right, confirming that the GSF's street constituency was essentially left-wing.

C.5 The Three Days

The protests began peacefully. From the afternoon of the 20th there were the first police charges, as violent factions, notably the Black Bloc, engaged in vandalism and clashes with the police. The police then attacked a section of the demonstration that was authorised and peaceful, which gave rise to further clashes. The violence escalated to the most tragic moment of the day, when Carlo Giuliani, a twenty-three-year-old protester, was shot and killed by a police officer.

Giuliani's death resonated widely among protesters and in public opinion. Clashes continued the following day, 21 July. The situation worsened during the night, when police searched the Diaz school, which housed activists and journalists. The search became a raid and led to the brutal assault of the school's 93 occupants, who were then taken to the Bolzaneto barracks and there repeatedly beaten and tortured.[2]

C.6 The Trials

The counter-summit echoed loudly in Italian public debate, which polarised between those supporting the crackdown and those supporting the protesters. Media reporting was confused and often incomplete, and full responsibility would be established only through the judicial proceedings of the following years. Given the multiplicity and complexity of the events, distinct trials were opened, covering the death of Carlo Giuliani, the unjustified police charges, the devastation of the city by the Black Bloc, and the events at the Diaz school and the Bolzaneto barracks.

The trials ran for more than ten years, and no concise summary can preserve the procedural complexity and political significance of each case; for a complete account see Agnoletto and Guadagnucci (2011) and the associated judicial sources. We report the final verdicts of the most relevant proceedings.

Giuliani homicide. The officer who shot Carlo Giuliani was dismissed on the grounds of legitimate defence.

[2] AFFAIRE CESTARO c. ITALIE (Requête no 6884/11).

Unjustified police charges. The charges against the peaceful part of the march were ruled illegitimate, with the aggravating circumstance of unjustified use of violence.

Devastation by the Black Bloc. The trial of 25 protesters accused of destruction of property concluded with 15 acquitted and 10 convicted, for a cumulative sentence of 100 years.

Diaz school. The final verdict found police officers and senior officials guilty of aggravated assault, injury, slander, forgery and fabrication of false evidence. The European Court of Human Rights additionally found Italy guilty and required it to introduce the crime of torture into its legislation.

Bolzaneto barracks. The trial exposed the violence and abuse inflicted on protesters during the night of 21 July. The judges ruled that the officers' actions fell under the standard definition of torture adopted by major international institutions, and that these facts amounted to a serious suspension of the rule of law. The ECHR's handling of the Diaz case applied here too.

It emerged during the trials that officers had frequently made false statements and falsified records. In its final ruling the ECHR stressed that nobody involved in the violence at the Diaz school or the Bolzaneto barracks spent a single day in jail, because the agents involved could not be identified and the offences lapsed under the statute of limitations. Many senior police officials retained their posts while under trial and were promoted in later years.

D Construction of the Trial-Deposition Corpus

Universe. The corpus covers the depositions of the four first-instance trials arising from the Genova counter-summit: the Piazza Alimonda case (the death of Carlo Giuliani), the trial of the twenty-five protesters charged with devastation and looting, the Bolzaneto barracks case, and the Diaz school case. We restrict attention to first-instance proceedings and to the spoken depositions themselves. We exclude the court's decisions, procedural postponements (*rinvii*), and the formal presentations to the court, since these are not first-person accounts of the events and follow a fixed administrative template. Within the retained hearings, the speakers fall into five roles: witnesses, victims, civil parties (*parti civili*), defendants (*imputati*), and forensic or technical experts. The parsing and classification yield a corpus of 304 first-instance hearings spanning 2004–2008 across the four trials, comprising roughly 268,000 speaker turns and 1.83 million tokens. Of the hearings we retain for the side analysis, the large majority are victim or testimony hearings, with a smaller set of forensic hearings and of police and defendant hearings; foreign victims testifying through an interpreter account for 28 hearings. The Bolzaneto and Diaz cases dominate the victim testimony, which is consistent with the nature of the two episodes: Bolzaneto concerned the systematic mistreatment of detainees, so hearing after hearing records an individual account of custodial abuse, whereas the Diaz material combines victim testimony with the crime-scene reconstruction presented by forensic officers.

From documents to speaker turns. Each hearing is distributed as a word-processor transcript. We parse every transcript into an ordered sequence of speaker turns, where a turn is a maximal block of text attributed to a single speaker. The transcripts were produced by different court reporters over several years and follow three distinct house styles, which we detect automatically and handle separately: a colon-and-guillemet format (*SPEAKER:* <<*text*>>); a tab-indented variant of the same in which foreign witnesses speak through a court interpreter, so that their testimony appears under the interpreter's turns; and a question-and-answer format that marks each exchange with *DOMANDA* and *RISPOSTA* and identifies the deponent in a *DEPOSIZIONE DEL/DELLA* header. The parser normalises typographic variants (dashes, non-breaking spaces, stray control characters) and recovers the hearing date, the trial, and the deponent names from the document header or file name. For foreign witnesses, the interpreter's rendering is treated as the deponent's own words, since it is the only verbatim record of the testimony.

Assigning each deponent to a side. For the analysis we group deponents into three sides that narrate the same events from different standpoints: victims (including the civilians who were detained or beaten, whether speaking directly or through an interpreter), police (officers and defendants), and forensic experts (technical and medico-legal

witnesses). We assign sides in two passes. First, hearings dominated by technical or crime-scene testimony are labelled forensic on the basis of their vocabulary (photographic exhibits, measurements, site inspections). Second, for the remaining hearings we score each deponent on two self-reference lexicons: a police lexicon that captures operational, first-person institutional language (for example *comandavo*, *ricevetti l'ordine*, *ero in servizio*, *il mio reparto*) and a victim lexicon that captures the language of a person acted upon (for example *mi hanno colpito*, *ero in cella*). A deponent is labelled police when police references clearly dominate and victim when victim references dominate. The remaining ambiguous deponents are resolved with a single rule that we validated by hand: within a victim or testimony hearing, an ambiguous deponent is a victim, because foreign victims often testify through an interpreter and their accounts contain fewer explicit first-person victim markers.

Thematic coding. On the classified corpus we build three deterministic codebooks over the deponents' own words: methods of violence and aftermath effects for victims, and anticipated threat for police. Each theme is a keyword family, and a hearing is flagged when it states the theme explicitly. This rule is conservative on purpose: it records a theme only when named, so the reported shares are lower bounds on true prevalence. We read these tables as a floor, and the language-level measures (emotionality and hedging) as descriptive contrasts across sides rather than causal estimates. Our pipeline follows the text-as-data methods reviewed in Ash and Hansen (2023) and Gentzkow et al. (2019): the side and theme classifiers are dictionary methods, and the emotionality measure of Figure 2 adapts the embedding approach of Gennaro and Ash (2022).

Validation. We validated the side classification against a set of deponents whose identity is known from the case record (for example the forensic officers who presented the crime-scene survey, the journalists and detainees who testified as victims, and named police defendants). We also hand-checked the resulting roster of distinct deponents and corrected misclassifications.