

A Report on “Persecution Perpetuated:
The Medieval Origins of Anti-Semitic
Violence in Nazi Germany” by
Voigtländer and Voth (2012)

Reviewer 2

February 07, 2026

v1



isitcredible.com

Disclaimer

This report was generated by large language models, overseen by a human editor. It represents the honest opinion of The Catalogue of Errors Ltd, but its accuracy should be verified by a qualified expert. Comments can be made [here](#). Any errors in the report will be corrected in future revisions.

I am wiser than this person; for it is likely that neither of us knows anything fine and good, but he thinks he knows something when he does not know it, whereas I, just as I do not know, do not think I know, either. I seem, then, to be wiser than him in this small way, at least: that what I do not know, I do not think I know, either.

Plato, *The Apology of Socrates*, 21d

To err is human. All human knowledge is fallible and therefore uncertain. It follows that we must distinguish sharply between truth and certainty. That to err is human means not only that we must constantly struggle against error, but also that, even when we have taken the greatest care, we cannot be completely certain that we have not made a mistake.

Karl Popper, 'Knowledge and the Shaping of Reality'

Overview

Citation: Voigtländer, N., and H.-J. Voth. (2012). Persecution Perpetuated: The Medieval Origins of Anti-Semitic Violence in Nazi Germany. *Quarterly Journal of Economics*. Vol. 127, No. 3, pp. 1339-1392.

URL: <https://www.jstor.org/stable/23251987>

Abstract Summary: This paper finds that anti-Semitic violence in Nazi Germany had local continuity over 600 years, with plague-era pogroms reliably predicting later anti-Semitic acts.

Key Methodology: Quantitative empirical analysis using standard regression techniques (*OLS*, *Poisson ML*), propensity score matching, and geographical matching on a new dataset of almost 400 German towns.

Research Question: How persistent are cultural traits, specifically anti-Semitism in Germany, and under what conditions does this persistence weaken?

Summary

Is It Credible?

Voigtländer and Voth present a striking thesis: that local anti-Semitic attitudes in Germany persisted for more than half a millennium, linking medieval violence directly to the atrocities of the Nazi era. The authors argue that towns and cities that witnessed pogroms during the Black Death (1348–50) were significantly more likely to engage in anti-Semitic acts in the 1920s and 1930s. They support this claim with a wide array of 20th-century indicators, showing that localities with a history of medieval Jew-burning saw higher rates of 1920s pogroms, higher Nazi Party vote shares in 1928, more letters to the virulently anti-Semitic newspaper *Der Stürmer*, more attacks on synagogues during *Reichskristallnacht*, and a higher proportion of Jews deported to extermination camps (pp. 1340, 1363, 1365). The authors posit that this continuity reflects the long-term transmission of cultural norms, which was only broken in cities with high levels of trade openness or migration (p. 1341).

The study establishes a robust statistical correlation between medieval and modern events. The association holds across various specifications and estimation techniques, including matching by geography, which compares nearby towns with different histories. For instance, the analysis suggests that medieval pogroms are associated with a probability of 1920s pogroms that is over 6 percentage points higher, a substantial increase given the baseline frequency (p. 1366). The authors also provide evidence of intermediate persistence, showing that these same towns were more likely to host anti-Semitic *Judensau* sculptures and witness the Hep-Hep riots of 1819 (p. 1377).

However, the interpretation of this correlation as the isolation of a specific, latent cultural trait requires caution. A primary challenge is the nature of the explanatory variable itself. The pogroms of 1349 were not random exogenous shocks; the

authors' own analysis shows they were more likely in cities with a prior history of violence and in commercially developed centers like Free Imperial cities (pp. 1374, 1377). While the authors control for observable medieval characteristics, the indicator for 1349 pogroms inevitably captures a bundle of historical, economic, and political legacies. It is difficult to definitively disentangle the persistence of cultural attitudes from the persistence of the unobserved institutional or economic conflicts that generated the violence in the first place.

Furthermore, the specificity of the finding—that it represents *anti-Semitism* rather than general xenophobia or political extremism—is open to debate. The authors acknowledge they cannot distinguish between anti-Semitism and a general hatred of minorities, as Jews were the primary minority group in both periods (p. 1344). Consequently, the results might reflect a persistent local intolerance toward any “out-group” rather than a unique animus toward Jews. The evidence provided to distinguish anti-Semitism from general right-wing radicalism is also somewhat fragile. The authors rely on a “political experiment” involving the 1924 election to separate anti-Semitic voters (DVFP) from nationalists (DNVP). Yet, in the baseline regression, the correlation between medieval pogroms and the anti-Semitic DVFP vote is not statistically significant at conventional levels (p. 1365).

The explanatory power of the findings also warrants context. While the historical legacy is statistically significant, the models typically explain only a small fraction of the variation in 20th-century outcomes. For example, the adjusted R^2 for the 1928 Nazi vote share model is 0.043, implying that over 95% of the variation in local support for the Nazis is driven by factors other than the medieval history identified in the study (p. 1365). Additionally, the narrative occasionally simplifies complex results; for instance, while the text highlights that Protestantism correlates with higher Nazi voting, it omits discussion of the regression results showing that Protestant areas had significantly *fewer* deportations and *fewer* anti-Semitic letters (pp. 1365, 1367).

Despite these limitations, the article makes a genuine contribution by identifying

conditions under which cultural persistence fails. The finding that the link between medieval and modern hatred vanishes in “open” trading cities and those with high migration provides valuable insight into how economic incentives and population mixing can erode deep-seated prejudices (p. 1378).

The Bottom Line

Voigtländer and Voth provide compelling evidence that German towns with a history of medieval anti-Jewish violence were statistically more likely to exhibit anti-Semitic behaviors in the Nazi era. However, the claim that this reflects the isolation of a specific, culturally transmitted anti-Semitic trait is difficult to distinguish from persistent general xenophobia or unobserved historical confounders. While the historical signal is detectable and robust, it explains only a very small portion of the variation in 20th-century atrocities, suggesting that while history matters, it was not the primary determinant of local participation in the Holocaust.

Potential Issues

Debatable endogeneity and confounding of the main explanatory variable: The article's central claim rests on the variable for Black Death pogroms (POG1349) being a valid indicator of a latent, persistent cultural trait of anti-Semitism. However, the variable's interpretation is complex. The authors' own analysis demonstrates that pogroms in 1349 were significantly more likely in cities with a prior history of anti-Jewish violence, as well as in more commercially developed centers like Free Imperial cities (pp. 1374, 1377). This suggests that the 1349 pogroms were not an exogenous shock that revealed a latent trait, but rather an outcome of a longer historical trajectory where pre-existing violence and economic motives were intertwined. While the authors use the correlation with pre-1349 violence as evidence supporting their persistence thesis and explicitly control for observable medieval economic and political factors in their regressions, the POG1349 variable remains a composite of cultural attitudes, economic incentives, and political structures (p. 1375). This makes it difficult to isolate the persistence of culture alone, as the variable may also be capturing the persistence of economic or political factors that generate conflict.

Inability to distinguish anti-Semitism from general xenophobia: The article's conclusions are specific to the persistence of anti-Semitism, but the research design may not be able to rule out a plausible alternative explanation: the persistence of general xenophobia or intolerance toward minorities. The authors acknowledge this limitation, stating, "We are not able to distinguish between anti-Semitism and a hatred of minorities in general" (p. 1344, fn 10). They argue that because Jews were the single largest minority in both the medieval and interwar periods, the two concepts are "observationally equivalent." While this is a practical constraint of the historical setting, it remains a significant interpretative issue. The findings could reflect a persistent local tendency toward intolerance of any "out-group," with Jews being the most salient target in both historical eras, rather than a continuous, specific animus

toward Jews themselves.

Fragile statistical evidence for distinguishing anti-Semitism from general radicalism: A key argument in the article is that it measures specific anti-Semitism rather than general right-wing or extremist political attitudes. This claim is supported by a “political experiment” using the 1924 election, where the right-wing DNVP party split into a more moderate version and the radically anti-Semitic DVFP. The article argues that medieval pogroms predict more votes for the DVFP and fewer for the DNVP (pp. 1367, 1384). However, the statistical evidence for the crucial link to the DVFP is fragile. In the article’s main baseline regression, the coefficient for medieval pogroms on the DVFP vote share is not statistically significant ($p > 0.10$) (p. 1365). The result only gains significance when the authors employ matching estimators. While this is a valid robustness strategy, the lack of a significant finding in the standard OLS specification weakens the claim that the data provides a clean separation between specific anti-Semitism and general political extremism.

Low explanatory power of the main findings: While the article identifies statistically significant relationships between medieval pogroms and 20th-century outcomes, the overall explanatory power of its models is low. The adjusted R^2 values reported for the main regressions are consistently small, typically explaining only 4–8% of the variation in the dependent variables. For instance, the baseline model for the 1928 Nazi vote share has an adjusted R^2 of 0.043, and the model for the main composite measure of anti-Semitism has an adjusted R^2 of 0.052 (pp. 1365, 1372). This indicates that over 90% of the variation in why some German towns were more anti-Semitic than others in the 1920s and 1930s is left unexplained by the historical pogrom indicator and the associated controls. The article does not discuss this aspect of the results, and while low R^2 values are not uncommon in cross-sectional historical studies, this context is important for assessing the practical significance of the findings.

Selective interpretation of the role of Protestantism: The article’s narrative regard-

ing the role of Protestantism appears incomplete, as it does not fully reflect all of its own empirical results. The text highlights the finding that a higher share of Protestants is associated with a higher Nazi vote share in 1928, noting this confirms the existing literature (p. 1367). However, it omits discussion of contradictory findings from the same table (p. 1365). These results show that a higher percentage of Protestants is significantly associated with *fewer* Jewish deportations and *fewer* letters to the virulently anti-Semitic newspaper *Der Stürmer*. By focusing only on the voting result, the article presents an incomplete picture of the complex relationship between Protestantism and different forms of anti-Semitic behavior found in the data.

Reliance on indirect and potentially noisy proxies for 20th-century outcomes: The study's measures of 20th-century anti-Semitism—such as the number of deportees, letters published in *Der Stürmer*, and synagogue attacks—are indirect proxies for local popular sentiment and are subject to significant measurement error. The authors acknowledge these limitations. They note that deportations were part of a centrally directed policy, that letters were subject to editorial selection, and that the absence of synagogue attacks could be due to “technical” constraints like fire hazards (pp. 1354–1356). While they provide arguments and cite historical sources to support the view that these variables still capture meaningful local variation, the inherent noise in these proxies is an important caveat to the interpretation of the results.

Ambiguous construction of the main composite outcome variable: The article's primary summary measure of 20th-century anti-Semitism is the first principal component derived from six different outcome variables. However, the procedure for handling missing data in its construction is not clearly specified. A footnote states that when one of the six variables is missing, “we use the five remaining variables to construct the principal component measure,” but then immediately notes that “Replacing the missing values... yields 311 observations” (p. 1371, fn 47). These statements could describe two different procedures—recalculating the principal component analysis for each observation based on its available data, or some form of data

imputation. The lack of clarity on this methodological step makes it difficult for the reader to fully assess the construction and reliability of this key dependent variable.

Potential bias in the main deportation analysis due to pre-war emigration: The main regression analysis for the number of Jewish deportees uses the 1933 Jewish population as a control variable (p. 1365). The authors acknowledge that this is a potential source of bias, as over half of Germany's Jews emigrated after 1933 but before mass deportations began (p. 1368, fn 44). If emigration was higher in more anti-Semitic towns, this would reduce the number of Jews available for deportation, potentially causing the analysis to understate the true effect of historical anti-Semitism. The article addresses this in an online appendix by showing the results are robust to using the 1939 Jewish population instead (p. 62). However, the potentially biased specification is featured in the main results table, while the more robust specification is relegated to the online appendix, a presentational decision that could have been made with greater transparency.

Presentation of effect sizes and minor clerical errors: The article contains several minor inaccuracies and presentational choices that could affect a reader's interpretation of the magnitude of the effects. First, the text contains inconsistencies regarding the magnitude of some of its own descriptive findings. It claims 1920s pogroms were "6 times more likely" in towns with a medieval history, when its own data show the factor is 7.45 (8.2% vs 1.1%) (pp. 1340, 1363). Similarly, it states letters to *Der Stürmer* were "about 30% more frequent," when the data show an increase of approximately 46% (0.86 vs 0.59 per 10,000 inhabitants) (p. 1363). Second, the sample sizes for some analyses vary between tables without explicit explanation. For example, the number of observations for deportations is 301 in one table but 278 in another (pp. 1358, 1365). While likely due to missing data for control variables, this lack of clarity can create confusion.

Lack of presented robustness checks for matching parameters: The article relies on propensity score and geographic matching to bolster its claims, using four and

two nearest neighbors, respectively. The authors assert in footnotes that the “results are much the same” or “almost identical” when the number of matches is changed (p. 1366, fn 40–41). While these checks are important for establishing the robustness of the matching estimates, the results of these sensitivity analyses are not presented in the article or its appendix. Providing this evidence would lend greater credibility to the chosen specifications.

Future Research

Disentangling xenophobia from anti-Semitism: Future work could test the specificity of the cultural trait by examining attitudes toward other minority groups. By analyzing historical data on the treatment of non-Jewish minorities (e.g., Poles in eastern Germany, or Catholics in predominantly Protestant areas during the *Kulturkampf*), researchers could determine if the “persistence” identified here is specific to Jews or indicative of a broader “culture of intolerance.”

Mechanisms of transmission: To better understand *how* such attitudes persist over centuries, research could focus on the micro-foundations of transmission. This might involve analyzing the content of local school curricula, parish records, or oral histories in towns with high vs. low persistence to identify specific vehicles of cultural memory, such as local festivals, songs, or family narratives that kept the memory of medieval conflicts alive.

Testing the “openness” hypothesis: The finding that trade and migration erode prejudice could be tested in other contexts. Researchers could examine other historical instances of long-term ethnic conflict (e.g., in the Balkans or South Asia) to see if commercial hubs and cities with high rates of in-migration consistently show lower rates of inter-generational transmission of hatred compared to isolated hinterlands.

© 2026 The Catalogue of Errors Ltd

This work is licensed under a

Creative Commons Attribution 4.0 International License

(CC BY 4.0)

You are free to share and adapt this material for any purpose,
provided you give appropriate attribution.

isitcredible.com