Lines of Demarcation:
Causation, Design-Based Inference, and Historical Research

(Supporting Information)

Appendix A

Descriptive statistics on determinants of sabotage in wartime France

Table A.1: Descriptive statistics

Variable	Mean	Standard deviation
Vote share difference	0.032	0.343
Vote share center	0.177	0.108
Log population	12.825	0.732
Germany distance (km)	413.087	199.045
Rough terrain (% of land area)	16.999	25.025
Double track length (km)	207.886	150.375

Analysis of fighting

Tables A.2 and A.3 replicate the analysis found in Tables 4 and 5 in the main body of our article for a distinct dependent variable: armed attacks by the Resistance against German personnel and their collaborators, which FM call "*Fighting*."

Table A.2 shows the difference of means between communes with and without double-track railway lines, but with the zone of occupation held constant. The top half of the table follows FM's procedure and drops communes intersected by the LoD; the bottom half recovers those communes but does not condition on occupation zones. It should not be at all surprising that the location of important railroads does a better job of accounting for the geographical distribution of railroad sabotage than for the geographical distribution of attacks on personnel; this is, indeed, what we find. The last row of the table indicates that there were almost three times as many *Fighting* events in communes with double-track railroads as there were in communes without. Overall, even when we condition on the occupation zones and drop many observations (i.e., those communes intersected by the LoD or located outside of FM's bandwidths) the location of double-track railroads continues to be an important predictor for the location of *Fighting* events.

Table A.2: Difference of means in *Fighting* between communes with and without double-track railroads, conditional on zone of occupation

		Mean of *Fighting* | *Double track* = 0 (std. dev.)	Mean of *Fighting* | *Double track* = 1 (std. dev.)	Difference of means, *[Double track = 0] – [Double track = 1]*	T-test (unequal variances)
Communes intersecting the LoD dropped					
5km Bandwidth	Occupied (N = 49)	0.121 (0.331)	1.875 (3.202)	-1.754	t = -2.186 p < .05
	Vichy (N = 46)	0.286 (0.750)	0.364 (0.674)	-0.078	t = -0.325 p = 0.749
10km Bandwidth	Occupied (N = 148)	0.178 (0.670)	1.244 (2.508)	-1.066	t = -2.686 p < 0.05
	Vichy (N = 146)	0.328 (0.770)	0.571 (0.978)	-0.243	t = -1.085 p = 0.289
20km Bandwidth	Occupied (N = 363)	0.255 (0.886)	0.865 (1.901)	-0.610	t = -3.028 p < 0.01
	Vichy (N = 328)	0.347 (1.059)	1.088 (2.366)	-0.741	t = -1.806 p = 0.080
All Communes	Occupied (N = 832)	0.238 (0.952)	0.703 (1.676)	-0.465	t = -3.337 p < 0.01
	Vichy (N = 586)	0.347 (1.133)	0.921 (1.853)	-0.574	t = -2.628 p < 0.05
Communes intersecting the LoD included					
Communes intersecting the LoD only	(N = 130)	0.535 (1.918)	1.387 (2.376)	-0.852	t = -2.034 p < 0.05
All Communes	(N=1548)	0.304 (1.128)	0.847 (1.828)	-0.543	t = -6.293 p < .0001

Two-tailed tests.

Table A.3 shows the difference of means between the occupied and Vichy zones holding constant the presence of double-track railroad lines. As before, we perform the analysis for 5-, 10-, and 20-kilometer bandwidths around the line, and for all communes, in the departments of Charente, Cher, Saône-et-Loire, and Vienne. Only two of the eight comparisons have the sign FM anticipate; not one is statistically significant at the .05 threshold.

Table A.3: Difference of means in *Fighting* events between occupied and Vichy zones, conditional on intersection with a double-track railroad, with communes intersecting the LoD dropped

		Mean of *Fighting*, Occupied Zone (std. dev.)	Mean of *Fighting*, Vichy Zone (std. dev.)	Difference of means, Occupied - Vichy	T-test (unequal variances)
5km Bandwidth	*Double track* = 1 (N = 27)	1.875 (3.202)	0.364 (0.674)	1.511	t = 1.830 p = 0.085
	Double track = 0 (N = 68)	0.121 (0.331)	0.286 (0.750)	-0.165	t = -1.181 p = 0.244
10km Bandwidth	*Double track* = 1 (N = 62)	1.244 (2.508)	0.571 (0.978)	0.672	t = 1.508 p = 0.137
	Double track = 0 (N = 232)	0.178 (0.670)	0.328 (0.770)	-0.150	t = -1.591 p = 0.113
20km Bandwidth	*Double track* = 1 (N = 130)	0.865 (1.901)	1.088 (2.366)	-0.224	t = -0.497 p = 0.621
	Double track = 0 (N = 561)	0.255 (0.886)	0.347 (1.059)	-0.092	t = -1.123 p = 0.262
All Communes	*Double track* = 1 (N = 231)	0.703 (1.676)	0.921 (1.853)	-0.218	t = -0.866 p = 0.388
	Double track = 0 (N=1187)	0.238 (0.952)	0.347 (1.133)	-0.109	t = -1.759 p = 0.079

Two-tailed tests.

Additional tests on communes and double-track railways

Table A.4: Difference of means test, kilometers of double-track railroad, by treatment assignment

		Mean of *Double track*, Occupied Zone (std. dev.)	Mean of *Double track*, Vichy Zone (std. dev.)	Difference of means, Occupied – Vichy	T-test (unequal variances)
Communes intersecting the Line of Demarcation dropped					
(1)	All Communes (N = 1418)	689.5 (1741.5)	439.9 (1300.7)	249.7	t = 3.089 p < 0.01
(2)	20Km Bandwidth (N = 691)	1034.3 (2116.1)	371.0 (1247.8)	663.3	t = 5.075 p < 0.0001
(3)	10Km Bandwidth (N = 294)	1166.4 (2280.9)	507.8 (1474.4)	658.6	t = 2.944 p < 0.01
(4)	5Km Bandwidth (N = 95)	1495.9 (2656.5)	850.8 (1806.7)	645.1	t = 1.391 p = 0.168
Communes intersecting the Line of Demarcation split					
(5)	5Km Bandwidth (N = 353)	1039.7 (2375.5)	329.2 (1101.3)	710.5	t = 3.615 p < 0.001
(6)	3Km Bandwidth (N = 253)	750.9 (2046.1)	178.0 (691.6)	572.9	t = 3.006 p < 0.01

Two-tailed tests.

Table A.5: Difference of means in commune distance to closest double-track railroad, conditional on treatment assignment, with communes intersecting the LoD dropped

		Mean commune distance to closest double-track railroad (meters), Occupied Zone (std. dev.)	Mean commune distance to closest double-track railroad (meters), Vichy Zone (std. dev.)	Difference of means, Occupied - Vichy	T-test (unequal variances)
(1)	All Communes (N = 1418)	6891 (8394)	11307 (11144)	-4415	t = -8.107 p < 0.0001
(2)	20Km Bandwidth (N = 691)	3609 (3902)	11671 (9990)	-8062	t = -13.701 p < 0.0001
(3)	10Km Bandwidth (N = 294)	4511 (4722)	8624 (7408)	-4113	t = -5.668 p < 0.0001
(4)	5Km Bandwidth (N = 95)	4915 (5315)	7227 (6864)	-2312	t = -1.827 p = 0.071

Two-tailed tests.

Analysis of sabotage restricted to pre-D-Day period

Table A.6: Difference of means in pre-D-Day *Sabotage* between communes with and without double-track railroads, conditional on zone of occupation

		Mean of *Sabotage* | *Double track* = 0 (std. dev.)	Mean of *Sabotage*, *Double track* = 1 (std. dev.)	Difference of means, *[Double track = 0] – [Double track = 1]*	T-test (unequal variances)
Communes intersecting the LoD dropped					
5km Bandwidth	Occupied (N = 49)	0.000 (0.000)	1.5 (2.805)	-1.5	t = -2.139 p < 0.05
	Vichy (N = 46)	0.114 (0.323)	0.182 (0.603)	-0.068	t = -0.356 p = 0.728
10km Bandwidth	Occupied (N = 148)	0.037 (0.191)	0.878 (1.965)	-0.841	t = -2.735 p < 0.01
	Vichy (N = 146)	0.064 (0.277)	0.381 (0.865)	-0.317	t = -1.666 p = 0.111
20km Bandwidth	Occupied (N = 363)	0.041 (0.234)	0.938 (2.423)	-0.896	t = -3.619 p < 0.001
	Vichy (N = 328)	0.068 (0.323)	0.794 (1.647)	-0.726	t = -2.564 p < 0.05
All Communes	Occupied (N = 832)	0.035 (0.221)	0.594 (1.956)	-0.558	t = -3.547 p < 0.001
	Vichy (N = 586)	0.053 (0.272)	0.711 (1.590)	-0.658	t = -3.597 p < 0.001
Communes intersecting the LoD included					
Communes intersecting the Line only	(N = 130)	0.030 (0.172)	1.097 (2.508)	-1.066	t = -4.235 p < 0.0001
All Communes	(N=1548)	0.042 (0.240)	0.687 (1.932)	-0.645	t = -11.558 p < 0.0001

Two-tailed tests.

Table A.7: Difference of means in pre-D-Day *Sabotage* events between occupied and Vichy zones, conditional on intersection with a double-track railroad, with communes intersecting the LoD dropped

		Mean count of sabotage events, Occupied Zone (std. dev.)	Mean count of sabotage events, Vichy Zone (std. dev.)	Difference of means, Occupied - Vichy	T-test (unequal variances)
5km Bandwidth	*Double track* = 1 (N = 27)	1.5 (2.805)	0.182 (0.603)	1.318	t = 1.820 p=0.087
	Double track = 0 (N = 68)	0.000 (0.000)	0.114 (0.323)	-0.114	t = -2.095 p < .05
10km Bandwidth	*Double track* = 1 (N = 62)	0.878 (1.965)	0.381 (0.865)	0.497	t = 1.380 p = 0.173
	Double track = 0 (N = 232)	0.037 (0.191)	0.064 (0.277)	-0.027	t = -0.863 p = 0.389
20km Bandwidth	*Double track* = 1 (N = 130)	0.938 (2.423)	0.794 (1.647)	0.143	t = 0.382 p = 0.704
	Double track = 0 (N = 561)	0.041 (0.234)	0.068 (0.323)	-0.027	t = -1.133 p = 0.258
All Communes	*Double track* = 1 (N = 231)	0.594 (1.956)	0.711 (1.590)	-0.117	t = -0.486 p = 0.628
	Double track = 0 (N=1187)	0.035 (0.221)	0.053 (0.272)	-0.017	t = -1.187 p = 0.236

Two-tailed tests.

Discussion of FM's Use of Local Linear Regression as their Estimator

In section 3.3, we estimate the effect of assignment to different sides of the LoD on sabotage, conditional on the presence of double-track railroads. Formally, our estimand is:

$$\tau_1 = E[Y_i|D = 1, R = 1, X_i \leq z] - E[Y_i|D = 0, R = 1, X_i \leq z]$$
$$\tau_2 = E[Y_i|D = 1, R = 0, X_i \leq z] - E[Y_i|D = 0, R = 0, X_i \leq z]$$

τ indicates the effect of being assigned to the occupied zone rather than the Vichy zone, Y_i is the number of sabotage events observed in commune i, D is the treatment indicator ($D = 1$ when communes are assigned to the directly occupied zone; $D = 0$ when communes are assigned to Vichy), R indicates whether or not a commune's boundaries intersect with a double-track railroad, X_i is the distance from the centroid of each commune polygon to the LoD, and z is a bandwidth. We estimate τ separately for $R = 1$ (communes that intersect a double-track railroad) and $R = 0$ (communes that do not intersect a double-track railroad).

Best practices in estimating the causal effect of the forcing variable in RDDs are currently the object of a methodological dispute. Imbens and Kalyanaraman (2012: 938) prescribe the use of local linear regression with a triangular kernel, which assigns increasing weight to observations as the value of the forcing variable approaches the discontinuity. In the RDD context, Imbens and Lemieux (2008: 623-624) argue, the difference-of-means is a biased estimator of the average treatment effect if the slopes of the regression of the outcome variable on the forcing variable diverge in the vicinity of the discontinuity. In contrast, Dunning (2012: 158) argues in favor of using the difference-of-means estimator, noting that

> if the conditional expectation of the potential outcomes under treatment (or control) on either side of the regression discontinuity is much different, for units included in the study group, the natural experiment has failed—for it has not in fact generated as-if random assignment to treatment conditions. In this case, the assigned-to-control group is not a valid counterfactual for the assigned-to-treatment group.

We follow Dunning's advice, for two reasons. First, because FM cannot determine whether sabotage events occurring in communes intersected by the LoD took place in Vichy or the directly occupied zone, they drop from the analysis all such communes. Figure A.1 below shows the distribution of communes by distance to the LoD. Of 119 observations within 2.5 kilometers of the LoD, 106 or 89% are excluded from the analysis. Consequently, FM's analysis is not really an instance of estimating a regression function at a boundary point. This accounts for FM's puzzling choice of estimator: local linear regression with a *rectangular* kernel, which weighs all observations equally regardless of their distance to the discontinuity, in essence replicating the issue that leads Imbens and colleagues to criticize the use of the difference-in-means estimator and endorse local linear regression in its stead. As FM recognize, the triangular kernel recommended by Imbens and Kalyanaraman (2012), "assigns large weights to areas without substantial support" (FM 2014, 650). Under such circumstances, we contend, local linear regression presents no advantages vis-à-vis a simple difference-in-means estimator.

Second, and more importantly, we believe the FM study is a nearly perfect example of Dunning's conjecture. As FM show (651, Figure 3), the slope of the regression of *Sabotage* on distance to the LoD (the forcing variable) diverges sharply near the LoD, with a steeply positive slope on the occupied side and a near-zero slope on the Vichy side. Our maps clearly indicate the reason: in the departments of Cher and

Saône-et-Loire (which together account for over 76% of the sabotage events in their study) this is precisely where the LoD runs close to the railroad, with many high-sabotage communes lying right along it. In other words, the divergence of slopes FM observe at the discontinuity is capturing the confounding effects of double-track railroads.

Despite these arguments in favor of the difference-of-means estimator, and as a robustness check on our difference-of-means estimation, we present in Figure A.2 below local linear regression point estimates and confidence bands of the sabotage and fighting dependent variables on the occupation zone treatment, controlling for intersection with double-track railroads. Estimates were computed for 5 – 25 kilometer bandwidths. While controlling for railroads improves the efficiency of the estimates, the occupation zone treatment indicator is insignificant at the 90% confidence level for 45 out of 50 estimates.

Additional References

Imbens, Guido and Karthik Kalyanaraman. 2012. Optimal bandwidth choice for the regression discontinuity estimator. Review of Economic Studies 79(3): 933 – 959.

Imbens, Guido and Thomas Lemieux. 2008. Regression discontinuity designs: A guide to practice. Journal of Econometrics 142(2): 615 – 635.

Figure A.1: Distribution of communes within the 5km bandwidth by distance to the LoD (Y axis indicated number of communes)

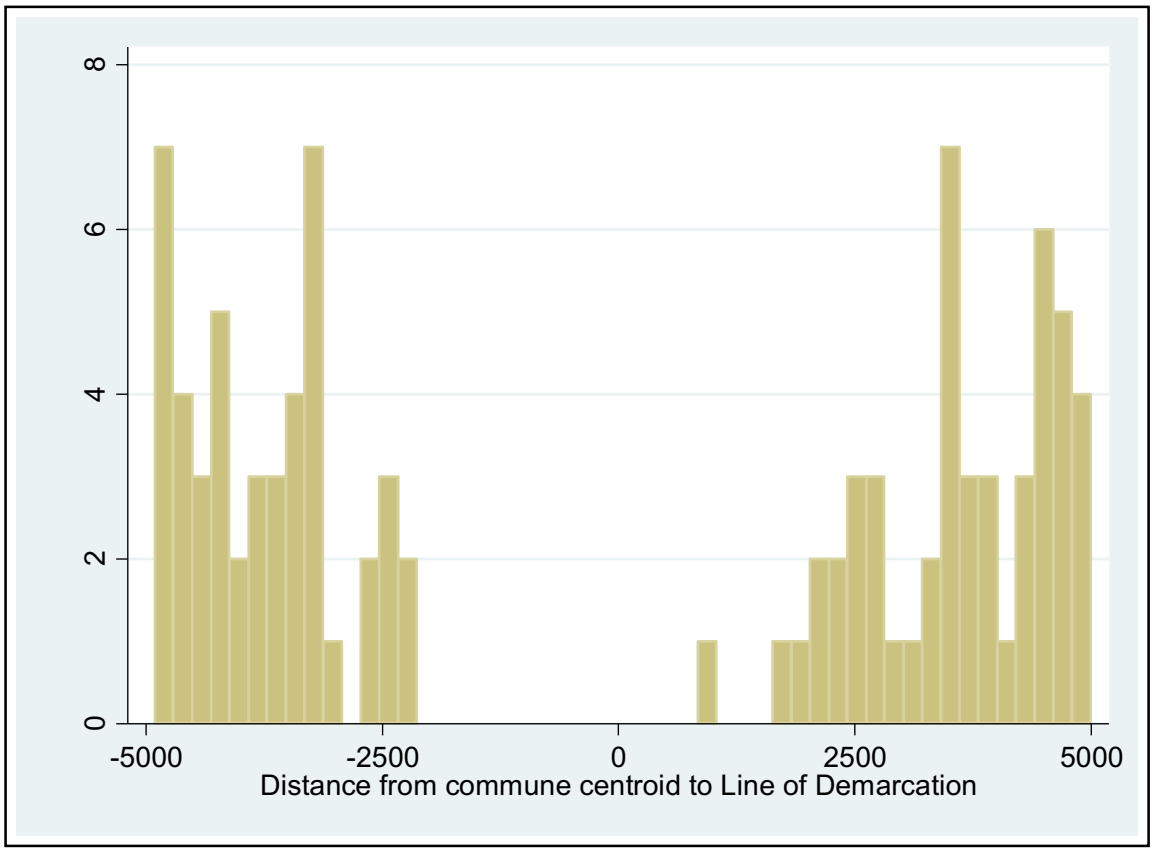

Figure A.2: Local linear regression, point estimates and confidence bands for occupation zone treatment

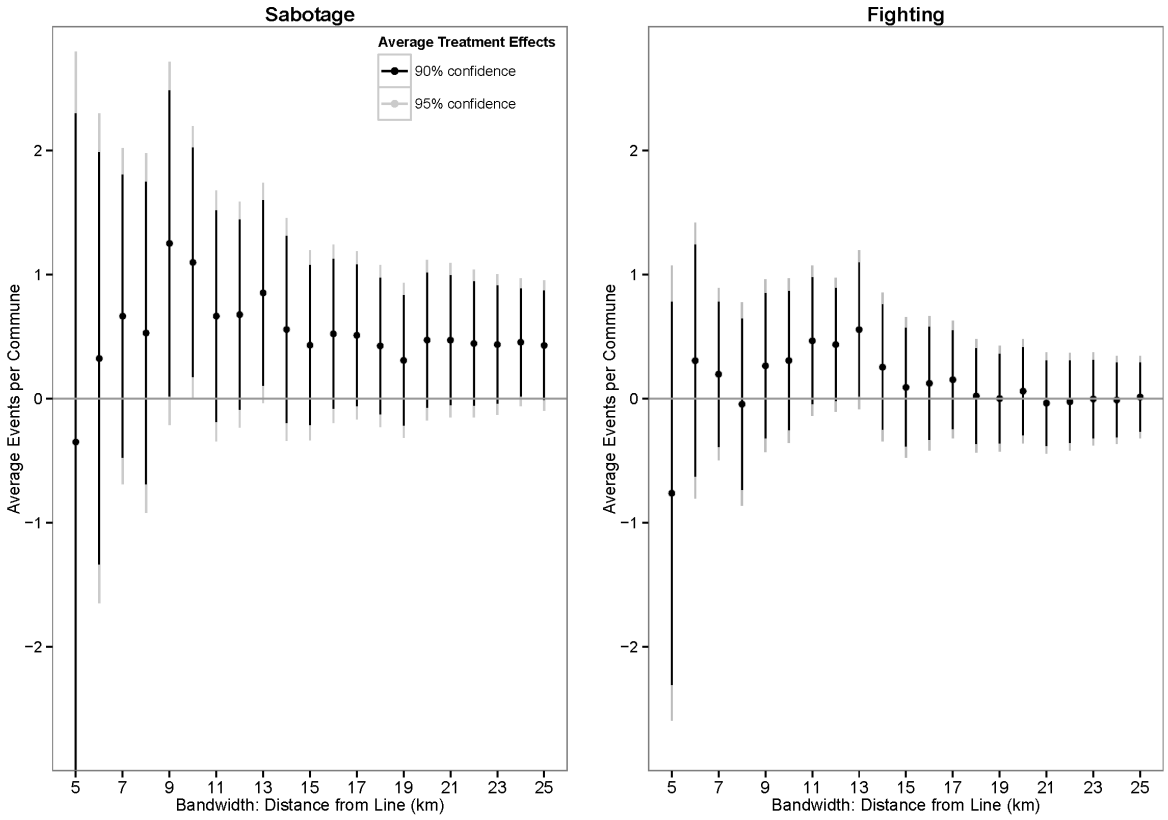

Appendix B

Discussion of differences between our data and FM's replication data

The replication data for Ferwerda and Miller (2014) became available in January 2015. As a robustness check on our results, we here discuss some discrepancies between our data and FM's data. We also replicate some of our key results using their data.

Placement of the Line of Demarcation

The first thing to note is the high degree of overlap between the two datasets. Despite the fact that FM's trace of the LoD is based on maps different from the ones we use, the coefficient of correlation between our respective indicators of the distance from commune centroids to the LoD is 0.998. Furthermore, Table B.1 below shows the correspondence between the two datasets' codings of the occupation zones. The two datasets have identical values on 1,501 out of 1,548 (97%) of the communes in the Charente, Cher, Saône-et-Loire, and Vienne departments. Not a single commune is coded by the two datasets as lying on opposite sides of the LoD. Every discrepant commune is coded by one research group as having boundaries intersected by the LoD, while the other group places the entire commune inside either the occupied zone or the Vichy zone. The relatively small number of discrepancies (47 out of 1548, or 3%) is not surprising, given that we find several cases in which the LoD intersected with only a small sliver of a particular commune's territory; consequently, differences of 100-200 meters in the trace of the LoD could result in distinct codings. A truly complete diagnosis of the differences between our datasets is not possible, given that FM's replication materials do not include their GIS data.

--- Table B.1 ---

We are also certain that FM make several coding mistakes. Map B.1 below shows a detail from our GIS of the Charente and the Vienne. We highlight the communes in these departments that FM drop from their dataset because they intersect the LoD. Given that the LoD itself was continuous, the set of communes intersected by the LoD should also have no discontinuities. But, as we can see, there are three gaps among the communes FM list as intersecting the LoD. (Two similar gaps are detectable in FM's coding of communes in the Saône-et-Loire). In other words, we can say for sure that FM include in their analysis at least five communes that, by their own coding criteria, should have been excluded. We say "at least five" because our own trace of the line intersects eight communes in these gaps.

--- Map B.1 ---

Sabotage and Fighting Events

In addition to differences in the trace of the LoD, the two datasets have somewhat different counts for sabotage and fighting events. Over 1,429 common observations, the two sabotage variables are correlated at 0.891, while the two fighting variables have a correlation coefficient of 0.795. One important difference in the coding rules for the fighting variable is that we include as instances of "fighting" all events designated in the sources as "embuscade" (ambush), "ataque" (attack), or "combat." It appears that FM tally events labeled "embuscade," "ataque," and "accrochage" (clash), while excluding events labeled "combat." The source material we both worked from does a poor job of explaining how events were sorted into these categories; FM do not explain their coding protocol in sufficient granularity for us to know which set of

events they consider to represent "fighting;" we were able to determine this only by examining their replication data. Consequently, we arrived at somewhat different operationalizations.

Beyond differences in the coding rules we apply, however, there are some differences between our event counts and the ones FM use. To some extent, such discrepancies are inevitable. The location information provided in the source material is sometimes inexact and calls for a degree of guess work. For instance, we encountered a number of events for which the location specified was not a commune. We researched those place names using a variety of sources. In some instances we were able to find villages with those names in the specified departments; in such cases, we assigned the events to the communes that contained those villages. It is easy to imagine that one group, consulting somewhat different sources, would come up with somewhat different information, resulting in slightly different counts.

Regardless of these particular issues, and given the nature of the sources, we are unlikely to generate a final consensus on coding. Thus, it makes sense to replicate our analysis using FM's data. Tables B.2 and B.3 are identical to Tables 4 and 5 from our main text, but in place of our data, we use FM's sabotage measure, their coding of the occupation zones, and their coding of the distance from commune centroids to the LoD. Since FM do not code double-track railroads, we use our own variable. There are some differences, but overall the results are quite similar to the ones we obtain using our own data. When we hold the occupation zones constant (Table B.2), communes with double-track railroads have a vastly higher rate of sabotage than communes without them. By contrast, holding double-track railroads constant (Table B.3), there is no evidence that the occupation zones explain variation in commune-level rates of sabotage.

--- Tables B.2 and B.3 ---

In short, while there are some coding differences between FM's data and ours, the differences are, for the most part, consistent with what one would expect from different research teams working from imperfect primary sources. Overall, there is no indication that our results are sensitive to coding differences. The criticisms we lay out in our paper are also supported by results based on FM's replication data.

Map B.1: Communes identified by FM as intersecting the LoD and therefore dropped from their analysis in the departments of Charente and Vienne

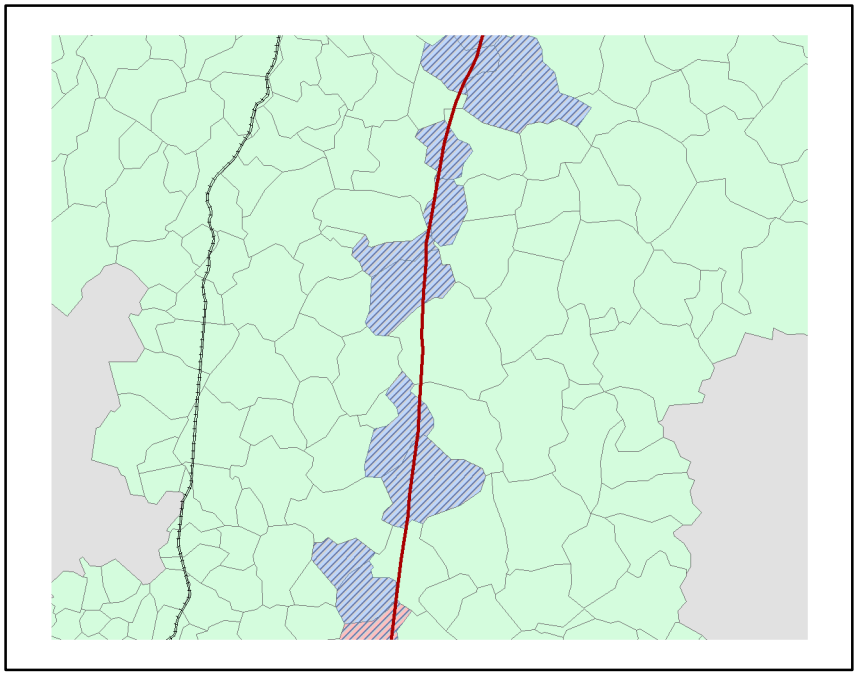

Table B.1: Correspondence of occupation zone coding between Ferwerda/Miller (FM) and Kocher/Monteiro (KM) datasets.

	Occupied (KM)	Vichy (KM)	Intersecting (KM)	Total
Occupied (FM)	819	0	20	839
Vichy (FM)	0	581	9	590
Intersecting (FM)	13	5	101	119
Total	832	586	130	1,548

Table B.2: Difference of means in *Sabotage* between communes with and without double-track railroads, conditional on zone of occupation (FM's replication data)

		Mean of *Sabotage*, *Double track* = 0 (std. dev.)	Mean of *Sabotage* | *Double track* = 1 (std. dev.)	Difference of means, *[Double track = 0] – [Double track = 1]*	T-test (unequal variances)
Communes intersecting the LoD dropped					
5km Bandwidth	Occupied (N = 55)	0.242 (0.902)	3.409 (5.612)	-3.167	t = -2.624 p < 0.05
	Vichy (N = 48)	0.026 (0.160)	2.333 (3.606)	-2.308	t = -1.920 p < 0.1
10km Bandwidth	Occupied (N = 159)	0.123 (0.547)	2.811 (5.382)	-2.689	t = -3.628 p < 0.001
	Vichy (N = 149)	0.092 (0.381)	2.611 (3.712)	-2.520	t = -2.877 p < 0.05
20km Bandwidth	Occupied (N = 372)	0.102 (0.435)	2.566 (5.020)	-2.465	t = -5.047 p < 0.0001
	Vichy (N = 338)	0.154 (0.595)	2.758 (4.381)	-2.603	t = -3.411 p < 0.01
All Communes	Occupied (N = 839)	0.132 (0.687)	1.813 (4.167)	-1.681	t = -5.180 p < 0.0001
	Vichy (N = 590)	0.128 (0.524)	2.068 (3.442)	-1.940	t = -4.840 p < 0.0001

Two-tailed tests.

Table B.3: Difference of means in *Sabotage* events between occupied and Vichy zones, conditional on intersection with a double-track railroad, with communes intersecting the LoD dropped (FM's replication data)

		Mean of *Sabotage*, Occupied Zone (std. dev.)	Mean of *Sabotage*, Vichy Zone (std. dev.)	Difference of means, Occupied – Vichy	T-test (unequal variances)
5km Bandwidth	*Double track* = 1 (N = 31)	3.409 (5.612)	2.333 (3.606)	1.076	t = 0.634 p=0.532
	Double track = 0 (N = 72)	0.242 (0.902)	0.026 (0.160)	0.217	t = 1.362 p=0.182
10km Bandwidth	*Double track* = 1 (N = 71)	2.811 (5.382)	2.611 (3.712)	0.200	t = 0.175 p = 0.862
	Double track = 0 (N = 237)	0.123 (0.547)	0.092 (0.381)	0.031	t = 0.495 p = 0.621
20km Bandwidth	*Double track* = 1 (N = 139)	2.566 (5.020)	2.758 (4.381)	-0.192	t = -0.212 p = 0.833
	Double track = 0 (N = 571)	0.102 (0.435)	0.154 (0.595)	-0.053	t = -1.215 p = 0.225
All Communes	*Double track* = 1 (N = 240)	1.813 (4.167)	2.068 (3.442)	-0.254	t = -0.494 p = 0.622
	Double track = 0 (N=1187)	0.132 (0.687)	0.128 (0.524)	0.004	t = 0.123 p = 0.902

Two-tailed tests.

Appendix C

Response to Ferwerda and Miller (2015)

In July 2015, Ferwerda and Miller (FM) released a detailed response to our criticisms, titled "Rail Lines and Demarcation Lines: A Response" (henceforth, FMR).[1] Here we evaluate the arguments they make in this reply.

The Line of Demarcation: a natural experiment?

The basic strategy of FM's research design will be familiar to anyone acquainted with the fundamentals of design-based inference. FM identify a sharp discontinuity in the data, in this case a spatial one, that corresponds to the Line of Demarcation (LoD) created by the Franco-German Armistice of 1940. They assert that the route of the LoD was locally arbitrary, such that assignment of French communes to either side was as good as random. Consequently, in close proximity to the LoD, comparing the post-treatment values of variables on one side of the line to the post-treatment value of the same variables on the other side of the line constitutes a natural experiment.

This research design would be fatally compromised if the discontinuity at the LoD happens to correspond to one or more sharp discontinuities in pre-treatment variables. We identify just such a break in the data: double-track railway lines that ran close to the LoD on one side, but not the other. Our maps (KM, Maps 2 – 4) make it clear that the LoD closely followed the contours of the double-track railway network, and our statistical tests show that assignment to the German side of the LoD is an excellent predictor of (i) whether or not a commune was intersected by a double-track railroad (KM, Table 3), (ii) how many kilometers of double-track railroad each commune contained (KM, Table A.4, and (iii) how far each commune lay from the nearest double-track railroad (KM, Table A.5).[2]

If the location of the LoD had been assigned randomly, the probability that we would observe these statistical associations with a pre-treatment variable is extremely small. Moreover, the new data FM present in their response to our critique (FMR, Table A.1) shows that multiple-track railroads were massively unbalanced, with 2.5 to 4 times as many kilometers of track on the German-occupied side as on the Vichy side of the LoD. FM do not dispute any of these results; on the contrary, they bring entirely new sources of data to bear in support of these claims.

In our critique, we argue that the local association between double-track railroads and the LoD was not an accident. We identify multiple historical sources indicating that the Germans had a keen interest in keeping specific railway connections in the part of France they would occupy directly. For strategic-level consultations, these references are remarkably specific. In the crucial meeting between Hitler, Keitel, and Jodl documented by Böhme—just three days before the German and French delegations met to discuss terms—Hitler laid out seven "guidelines" (*Richtlinien*) for an armistice. Only one of these conditions concerned the specific shape of the planned occupation zones, and it focused solely on the maintenance of "connections" (*Verbindungen*) through central France and down the French Atlantic coast. Specifically,

[1] See: http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2628508.

[2] Mentions of "KM" refer to the main body of this article and the remainder of this Supporting Information, which have been revised from the working paper versions to which FM responded. Mentions of "FM" refer to Ferwerda and Miller's original (2014) article.

Böhme (21; quoted in KM, ***) writes: "The envisioned demarcation line between occupied and unoccupied territory was drawn on a map by General Jodl. In the course of doing so, attention was paid to ensuring that the East-West connection through central France that went from Belfort through Dôle-Le Creusot-Moulins-Bourges-Tours to Nantes, and the North-South connection from Tours through Angoulême-Bordeaux to the Spanish border, would run within the territory to be occupied." Within days of this meeting, Hitler made it clear in diplomatic exchanges with Mussolini and agents of the Spanish dictator Francisco Franco that the connections he placed such a high priority on were, in fact, railroads. He specifically mentioned the Paris-Tours-Bordeaux line. Also, as we show in KM (***), the press reporting on the Armistice that appeared before the LoD was implemented is quite specific that the boundary was to be placed approximately 20 kilometers east of the Paris-Tours-Bordeaux railroad.

FM concede that the North-South "connection" was a railroad, but they insist that we don't really know if the East-West connection through central France was intended to be a railroad: "[I]t may be reasonably interpreted as a railway, but it is also possible that Boehme was referring to a line of communication or road connection" (FMR, 9, fn. 23). Why would Böhme use the same word (*Verbindung*), *in the same sentence*, to refer to a North-South railroad and some other kind of East-West connection? Since all of the towns mentioned (Belfort, Le Creusot, Moulins, Bourges, Tours, Nantes) lie on the network of double-track railway trunk lines, *and* they form an almost-direct route across central France, *and* there is no motorway that could serve this purpose, *and* the WWII German military transported their armies and material over rail networks at strategic distances, by far the most straightforward interpretation is that Böhme was talking about a railroad.

The foregoing makes it clear that keeping these two railroad connections inside their occupation zone was a priority for the German high command in shaping the LoD. FMR counters that evidence from several sources indicates that the LoD was adjusted locally by German commanders on the ground or altered over time by negotiations through the Franco-German Armistice Commission in Wiesbaden. FMR asserts that "a range of idiosyncratic factors determined the line's placement in the departments we study" (FMR, 6), and concludes that "the demarcation line was not *singularly* determined by the location of strategic railroads" (FMR, 4, emphasis added).

We agree that strategic railroads were not *the only* factor that determined the local placement of the LoD. The strategic priorities articulated by the OKW had to be implemented by lower-ranking officers. The minor modifications FM identify might have made the LoD less faithful to Hitler's "master plan," but they might also have reflected local commanders' attempts to better execute the Führer's eminently rational objectives on the ground. Here it is crucial to keep in mind that the precise contours of the LoD were not *uniquely* well-suited to maintaining these specific railroad links through France; *any* route that moved the boundary further south or east, and some routes that nudged the boundary further west, would also have kept these railroads inside the German zone. At the same time, the Germans wanted to maintain a sovereign France. To do that, they needed the French to agree to an armistice. Thus, it was in their interest to grant as large an unoccupied territory as would be consistent with their other strategic aims. Minor adjustments to the route of the LoD of a few kilometers here or there, of which there were a number we know about, would not and did not change its overall shape.[3] That *other* priorities came into play is not terribly important unless FM can show that those other rationales overruled German concerns about railroads. They cannot, which fatally undermines their identification strategy.

[3] This is confirmed by the two examples of "vagaries" of the local placement of LoD included in FMR as figure A3. In neither of these situations did the strategic railway lines change side.

Finally, on this point, FM call our attention to the fact that the Paris-Tours-Bordeaux railway—which we claimed determined the LoD's placement in the Charente and the Vienne—lay 20 kilometers away from the LoD. Since some of FM's tests use bandwidths narrower than 20 kilometers, they argue that this cannot have affected their overall results (FMR, 15). In response, we make two points. First, in their original article, FM include a number of statistical tests using "several bandwidths ranging from a distance of 10 km to 25 km from the demarcation line" (FM, 650). The reason they do this is, as they acknowledge, that "regression discontinuity designs are highly sensitive to the choice of bandwidth" (FM, 650). One of the apparent strengths of their original findings is the consistency of their results over these multiple bandwidths. If, as FMR now suggests, some of their tests at wider bandwidths should be discarded because of the imbalance we detect in the presence of strategic railways, the overall strength of their findings would be considerable weakened.[4] Second, although, as FMR claims (FMR, 15), communes within 10 kilometers of the LoD in the Charente and the Vienne did not intersect double-track railroads, it does not follow that the imbalance in these railroad does not affect FM's results. As we discuss in detail below, a location closer to major railroads is likely to be associated with many additional pre-treatment variables for which FM do not control. Therefore, proximity to the strategic railways we identify might have led to higher incidences of violent resistance activity. In short, to accept that the LoD in the Charente and the Vienne was placed with the Paris-Tours-Bordeaux railroad in mind, as FM now do, is not a trivial concession on their part.

Summing up, the evidence that the LoD was designed to keep important railroad connections inside the German zone is overwhelming. But, what if FM are right in spite of the evidence? What if the local placement of the LoD was really the result of a welter of small, "idiosyncratic," and non-systematic causes? In point of fact, their research design would still be invalidated if, by a highly improbable accident of fate, the LoD *just happened* to have ended up located right along double-track railroads that were prime targets of the Resistance. The maps and statistical tests—not only ours but also the ones included in FMR—are clear that *at least* the latter scenario must be true. This is a problem, not only because of the presence of these railroads, but also because their placement is almost certainly associated with the location of a great many other things—population centers, waterways, roads, industrial facilities, etc.—that might influence the occurrence of violent resistance. So our criticism of FM's research design is twofold. First, we can show that the LoD corresponds to at least one important pre-treatment discontinuity. Our statistical tests show that the association between the LoD and double-track railways is extremely unlikely to have occurred by chance. Second, given the historical *and* statistical evidence, it is highly improbable that the Germans had no systematic rationale for the local placement of the LoD.

Does FM's theory explain variation in resistance during WWII France?

Even though the failure of their identification strategy greatly undermines the confidence we can have in their causal claims, FM could of course still be right that the difference in political institutions between the two zones was an important determinant of the levels of violent resistance on each side of the LoD. FMR attempts to make this case by controlling statistically for single- *and* double-track railroads and showing

[4] As is clear from the maps, the LoD also bulged in two locations to within much closer than 20 km of the railroad in the Charente and the Vienne. Consequently, a number of communes located less than 20 kilometers from the LoD did intersect the Paris-Tours-Bordeaux railroad. This is an additional cause for concern regarding FM's local linear regression estimates at bandwidths between 10 and 20 kilometers.

that, when they do so, they "observe significantly elevated levels of resistance activity within the German zone" (FMR, 3).

Before we move on to discuss FM's specific analysis, let us review what we know so far about the role of railroads in shaping the patterns of violent resistance. None of these points is disputed in FMR. First, we know that a large percentage (at least 66%) of the sabotage attacks FM examine were directed against railroads. Second, we know that the Resistance (and the Allies with whom they were cooperating) had good reason to attack the strategic railroads that were the Germans' principal means of moving troops and supplies around France. We know that, in preparation for the D-Day landings, explicit plans were put in place to attack French railroads in order to cut off the Normandy beaches from reinforcement and resupply, and to prevent German troops from escaping France when the front collapsed. Third, we know that double-track railroads were massively unbalanced in the vicinity of the LoD in the departments FM studied (KM, Tables 3, A.4, and A.5; FMR, Table A.1). Finally, we know from both maps (KM, Maps 2 – 4) and statistical tests (KM, Table 4) that communes with double-track railroads running through them were *vastly* more likely to have sabotage attacks than communes without such railroads. These differences are gargantuan, dwarfing the effects FM identify for political institutions. Although only 17% of communes in FM's four departments had double-track railroads going through them, 78% of the sabotage attacks happened in those communes.

In our critique of FM, we compare the rate of sabotage on the two sides of the LoD while controlling for the presence of double-track railroads. We find little evidence of more attacks on the directly-occupied side of the LoD conditional on intersection with double-track railroads (KM, Table 5). FMR now test for a difference in the rate of sabotage per kilometer of multiple-track railroad (FMR, Table A.1) and, consistent with our criticism, find that differences between the two zones are not statistically significant at any of their bandwidths.[5]

The crux of FMR is to assert that, if one controls statistically for something *else*, something we have explicitly asserted is unlikely to matter, then sabotage attacks "remain elevated" on the German side of the line. FMR's argument involves three steps. First, rather than conditioning on whether or not each commune intersected a railroad, they measure kilometers of track running through each commune. This is a reasonable alternative measurement. Second, they examine only sabotage against railroads, rather than all sabotage attacks as they did in FM; they offer no explanation for changing the dependent variable and no tests with the original dependent variable. We are agnostic about this change; we were reluctant to base our own analysis solely on railroad sabotage because FM examined sabotage against all targets. Third, FM carry out a series of statistical tests conditional on *total* kilometers of railroad track. To account for the difference between single- and multiple-track railroads, they multiply multiple-track kilometers by two.

We should start by revisiting why we took the approach we did: conditioning on the presence of double-track but not single-track railways. In 1940-44, France was a country with a remarkably extensive railway network. Not all of the lines in this network were of equal importance. Some were intended to carry passengers and freight over long distances, while others linked smaller population centers to the national network. The most important lines in the network, its "trunk" lines, tended to be double-tracked. The crucial point here is the following: to assert that these railways were important because they had two tracks

[5] To be more specific, their t-values imply that a higher sabotage count on either side of the LoD is statistically credible. In addition, their point estimate for the 20 kilometer bandwidth indicates a higher sabotage rate on the Vichy side.

is to have it backward. Rather, these railways had two tracks because of the economic or strategic important of the places they connected. The route through central France that Hitler directed his generals to keep in the occupied zone began in Belfort (on the German frontier) and terminated in Nantes (near the Atlantic coast, at the base of the Breton peninsula). It was therefore useful in a quite general way for connecting Germany to the Atlantic and, thus, for controlling the French Atlantic coast. This line also intersected with multiple north-south rail lines that were useful for controlling southern France down to the Spanish border and the Mediterranean.[6] In contrast, the single-track railroad between, for example, St. Saviol and Le Vigeant in the Vienne was useful only for controlling those two strategically inconsequential towns.

Double-track lines were not, therefore, "twice as important" as single-track lines, as implied by FMR's method of multiplying their length by two. It would be closer to accurate to say that double-track lines mattered a lot and single-track lines hardly mattered at all.[7] Multiplying double-track length by two is thus a totally arbitrary decision with no basis in theory or the available evidence. Furthermore, since single-track railways were very common on both sides of the LoD in the territory FM examine, including their track length in the analysis conflates locales that contained important Resistance targets with locales that did not, thereby obscuring the effects of important railroads.

FMR defends this decision to pool all railroads in the departments of Charente, Cher, Saône-et-Loire, and Vienne by pointing out that single- and double-track lines are likely to be statistically associated, because they are part of a network—i.e., they must connect at some point. Thus, FM argue, controlling for double-track lines while failing to control for single-track lines will introduce bias. After all, given that our data only identify the commune in which a railway sabotage attack occurred, not its precise location within that commune, in communes where double-track and single-track railways co-occur, it is impossible to know if a Resistance attack was directed at the big railway to somewhere or the little railroad to nowhere (FMR, 14).

There are certainly places where double- and single-track railways co-occurred in these departments. However, given that most of the double-track was on the German-occupied side, most of the communes with both types of railroads were *also* on the German-occupied side. For instance, in the Charente and the Vienne, there were literally *zero* kilometers of double-track on the Vichy side of the line, from which it follows that all communes with more than zero kilometers of single-track *and* more than zero kilometers of double-track were also on the German-occupied side of the LoD. According to our data, of the 18 communes in FM's four departments in which a single-track line intersected a double-track line, 14 were located in the German-occupied zone. These communes had a very high rate of sabotage. Thus, explicitly conditioning on the coincidence of single-track and double-track lines gives additional evidence *against* FM's argument, rather than in favor of it.

---

[6] One implication of this point is that down-scaling a double-track railroad to single-track did not necessarily make it strategically unimportant. If the railroad section running from Paray-le-Monial to Montchanin was downgraded (FM say this was planned, but they do not know if it ever happened), its position in the network could still have made it vital.

[7] The deeper problem is that, of course, all double-track lines were not equally important for all strategic purposes. Likewise, some single-track lines may have been of great strategic importance in some parts of France at particular points in time. Thus, double-track lines are an imperfect quantitative proxy for the strategic problem faced by a continental power attempting to hold France against a maritime invader.

There is a more fundamental problem, however. Given the evidence we present in KM about the strategic rationale behind the location of the LoD, it is important to keep in mind that we are no longer in the Elysium of design-based inference; we've gone to messy-regression-land. Railroads might be the only confounding variable to account for, but how likely is that? The whole thrust of design-based inference is to avoid having to make such tendentious assumptions. Although in FM's view there were multiple factors determining the placement of the LoD, they now want to assume that all of those reasons are small, "idiosyncratic," and non-systematic. How likely is that?

If FM's goal is to create a statistical model of sabotage in WWII France, then there is much more work to be done, because many variables are statistically associated with double-track railroads, and a number of those variables also probably account for some portion of the variation in sabotage attacks. For example, Route Nationale 10, the most important motorway in the Charente and the Vienne, followed the path of the Paris-Bordeaux railroad, which connected the regional hubs of Châtellerault, Poitiers, and Angoulême. The Vichy zone in these departments contained no cities of comparable size. In the Saône-et-Loire, the town of Le Creusot, located near the Tours-Belfort railway, contained France's most important munitions factories, which Hitler explicitly wished to keep inside the occupied zone. As we point out in in our critique (KM, fn. ***):

> Railway lines are not built just anywhere. They often connect existing cities and towns. They influence economic and population growth as people and industries move to locations where transportation is convenient. They require rights-of-way, which tend to accrue additional infrastructure: fuel and water pipelines; electricity and telephone wires. Natural geography makes some routes more feasible and less costly than others. Civil engineers site them next to water courses and along the contours of the land to avoid building tunnels and bridges.

Thus, the plausible target set for sabotage is endogenous to railways, which are in turn endogenous to other variables that themselves plausibly influence the target set. Controlling for single-track railroad in addition to double-track railroads is not going to solve this problem. FM promise us clean identification, because the LoD was supposedly "as good as random." But clearly it was not. Because the Germans cared about key strategic railroads, they positioned the LoD along a route that created many important differences between the two zones.

In KM, we do not pretend for a moment to have provided the true causal model of sabotage in WWII France—or even a plausible identification strategy. At the same time, we have extremely good evidence that the Resistance intended to sabotage important railways, and that the vast majority of their sabotage attacks in fact happened in communes through which these important railways passed, during a time at which railway sabotage was of great military value to the Allies. FM are certainly correct that, in the vicinity of the LoD, more of these attacks occurred on the German side than on the Vichy side. But, as our maps show clearly, these attacks were also clustered in a peculiar way, in spatial "strips." These strips just happen to correspond to the routes of some of the double-tracked trunk lines of the French railway system. In the vicinity of the LoD, lines of that type just happen to be far more prevalent (2.5 to 4 times as prevalent, depending on the bandwidth) on the German-occupied side. To convincingly rebut our critique by showing that something *else* (not double-track railroads) can explain the different rates of sabotage across the LoD, FM should ideally be able to account for this odd pattern of spatial clustering. Single-track railroad lines are not the answer, because they did not follow the same strips of land. The effects of the LoD itself might be a

candidate, except that the strips of sabotage sometimes lie close to the LoD and sometimes do not. In point of fact, the strips lie close to the LoD only when a double-track railroad line *also* runs close to the LoD.

In sum, pooling single- and double-track railroads into a single variable, as FMR does, is an arbitrary modeling choice that obscures a crucial relationship we identified in the data. More importantly, once FMR concedes the necessity of controlling statistically for this particular variable, one may reasonable ask why we should stop there. FMR dismisses the usefulness of the statistical model of railroad sabotage we provide for all France, writing that "findings on overall resistance levels can only be interpreted as causal if we believe they have controlled for all relevant confounders associated with respect to occupation zone and resistance" (FMR, 4). This objection applies with equal force to the analysis they conduct in their response.

The strategic logic of violent resistance in WWII France

In KM, we show that the overwhelming majority (95%) of FM's data recorded events that occurred after the LoD was dismantled in early 1943, when Resistance groups could operate unimpeded by any physical barriers separating the two zones. We also demonstrate that more than half of the events recorded in FM's data occurred after the D-Day landings, when political rule was in flux and Resistance efforts were coordinated with and by the Allied powers. We show this by deploying a monthly chronology of the data as well as detailed local histories of two of the four departments FM study—the Vienne and the Saône-et-Loire—that together account for 83% of FM's data (Calmon 2000; Veyret 2001). These histories provide substantial detail on both (i) the role played by attacks on strategic railways and (ii) that played by Allied coordination efforts. Taken together, these factors show that FM's data is ill suited to test theories that connect the location of a Resistance attack with the geographic origins and motives of its perpetrators. Much of the violence FM observe was part of a coordinated effort to hit strategic railways in order to hinder the movement of German troops and materiel once conventional fighting in France resumed in June 1944.

Responding to this rich local evidence, FMR makes three claims. First, it argues that resistance in their departments of interest "was driven by local factors and began before D-Day" (FMR, 4). Specifically, FMR argues that "locally rooted resistance groups had emerged across France long before [the Allied landings of June 1944] and before the resistance was nationally unified" (18). In support of this claim, FMR cites two works. The first is Kedward's (1993, 163) general study of rural resistance in southern France, which emphasizes the role of "local objectives, local organization, and still more, local constraints, remained paramount" even in 1944. The second is a study of the Gironde—a department that is not included in FM's original analysis or our rebuttal—where "more than 80% [of resistance members] were local residents, with the majority organized within local units" (FMR, 19).

Let us start by noting that, in KM, we do not dispute that the Resistance had emerged locally long before D-Day and consisted mostly of local members. Most Resistance groups were in fact formed earlier in the war; and most Resistance members, unsurprisingly, joined the Resistance where they lived. These facts are hardly relevant for the matter at hand, however. What matters is that, in the four departments FM study, the Resistance did not conduct much violent activity before the LoD was abolished in early 1943. Per our count, only 32 of the 686 sabotage attacks in the data for which we have complete dates took place during this early period. The overwhelming majority of the attacks were perpetrated when the LoD was no longer a physical impediment to Resistance forces acting across the two zones. This trend in violent activity supports our view that the Resistance grew engaged in violent action in preparation for, and particularly in the aftermath of, the Allied landings of 1944, benefiting from Allied support and coordination. Since we establish that within the bandwidths FM study, the German zone was vastly richer in targets that were of

great strategic importance for the Allied effort to liberate France, FMR's counter depends on establishing that, despite this imbalance, Resistance efforts were so "local" that operatives would be unwilling or unable to travel a few kilometers from within Vichy territory to German-occupied territory in close proximity to the LoD in order to perpetrate an attack. For FMR's counter to work, the local objectives, organization, and constraints that Kedward mentions must be "local" at a scale that prevents them from operating across the small distances that separate localities within FM's bandwidths across the LoD. Is this plausible?

To answer this question we must examine FMR's second counter, their claim that "even in 1944, the border between zones still affected resistance activity" (FMR, 17). How so? Other than an orthogonal claim about how in the Gironde—again, a department not included in FM's original study or our rebuttal—"the demarcation line represented an ideological break as well as a frontier for the organization of the Resistance," all FMR has to offer about the four departments in FM's study is that in the Saône-et-Loire "the resistance was organized differently in the two zones, even after 1942" (FMR, 17, citing Veyret 2001, 16, 46).

We have little doubt that being under direct German occupation or a subject of Vichy France produced until the end of the war some ideological and organizational differences between Resistance groups depending on where they were based. But, once again, this is largely irrelevant for the matter at hand. What matters is whether the higher level of railway sabotage we observe in the German zone—where the railways were located within FM's narrow bandwidths—can be imputed to local groups originating on the same side of the LoD. It cannot. Given that the LoD no longer posed a significant physical hurdle and that, after D-Day, conditions on the ground were fluid, it is implausible that the sabotage attacks against these railways—which drive FM's findings—can be imputed to Resistance groups originating solely on the German side of the line.

This poses a serious problem for testing FM's theory. FMR retorts that "it is incorrect that our theory strictly requires that attacks by perpetrated by groups originating in the same zone" (FMR, 17). As FM wrote in their original article, because "effective insurgency requires the aid of the local population," it is no problem "if some spillover did occur *due to increased motivation to resist in the German zone*" (FM, 652, emphasis added). The problem for FM stems from the fact that, as our criticism conclusively establishes, the spillover did *not* occur due to increased motivation to resist in the German zone. Rather, the spillover was the result of a great relative abundance of strategic targets in that zone, namely, the presence of railway lines vital to the German war effort. It is impossible to tell whether these lines were being hit by Resistance operatives based in their immediate vicinity or a few kilometers farther away, in Vichy territory.[8]

[8] Attempting an intellectual jujitsu move, FM find the criticisms we lay out in KM against using the location of Resistance events to infer the political motivations of the insurgents to be "puzzling" (FMR, 18), since in our own work (Kocher et al. 2013) we do the same. FM are correct that in Kocher et al. (2013) we use the location of attacks to infer their political motivation. There is, however, a crucial difference between FM's use of location to infer motivation and ours. The difference has to do with the scale of the study. FM use a micro slice of territory—a narrow bandwidth around the LoD. We cover the entirety of French territory. At this macro level, it is implausible that most attacks deep in the directly occupied zone were perpetrated by maquis originating deep in Vichy—and vice versa. Clearly, most Resistance groups were not traveling hundreds of miles across France to conduct their business. At the micro scale FM use in their study, however, it is a highly plausible conjecture that attacks occurring on one side of the LoD (which was by the time the overwhelming majority of these attacks took place an administrative border but not a physical barrier) were perpetrated by Resistance groups originating from across the LoD *particularly in areas where strategic targets were vastly more available on one side*, as is the case with the territory FM study.

Finally, FMR introduces what it labels the "most important" objection to our claims on the role of Allied coordination of Resistance attacks in the territory FM study. FMR objects that "Plan Vert, the major coordinated campaign for destroying French railways that Kocher and Monteiro specifically reference, identified relatively few targets in the four departments we analyze" (FMR, 19). They cite a list of targets found in Durand's (1968, 432-434) history of French railroads during the war, which includes only two communes (Poitiers and Angoulême) in the sample of localities on which FM test their argument. Since both of these towns are located approximately 15 kilometers from the LoD, they are not included in FM's 5- and 10-kilometer bandwidth estimates.[9] Hence, FMR concludes that Allied coordination is "insufficient to explain the patterns of resistance" observed within FM's sample.

As we mention in KM (***), the Allies developed two plans organizing Resistance efforts against French railways in the aftermath of the D-Day landings. The first was *Plan Vert*, which targeted railway lines. The second was *Plan Grenouille*, which aimed to sabotage hoists (*appareils de levage*) and turntables (*plaques tournantes*) in selected railroad depots. Durand's target list refers to *Plan Grenouille*; he does not provide a list of the targets included in the much broader *Plan Vert*. Nevertheless, Durand does quote extensively from an assessment of the effects of *Plan Vert* produced by the French General Bondil, who says that important objectives were met in Charente, Vienne, and Saône-et-Loire (Durand 1968, 438-439): "In the Charente and the Vienne, 200 trains were immobilized on June 21 [1944] for lack of locomotives. … In the Saône-et-Loire, the important railway centers of Chalon-sur-Saône and Paray-le-Monial were rendered unusable." (Both these railroad hubs are located in the formerly occupied territory FM study, in communes intersected by the LoD, and therefore dropped from FM's analysis.) Durand's text also gives ample evidence of systematic coordination between the Allies and the Resistance in their efforts to sabotage French railways in the immediate aftermath of D-Day, as we indicated in our rebuttal.

FM are certainly correct that the Resistance had a great deal of local agency in selecting targets. As Durand (1968, 426-428) notes, the locations chosen for sabotage as part of the Allies' *Plan Vert* were "susceptible of modification" on the ground and "a great deal of initiative would be left to the operators," particularly as to where exactly to hit a railway line the *Plan* deemed a sabotage target. In this sense, *Plan Vert* was more "a rational orientation than an imperative order" (Durand, 1968, 426). In other words, the Allies provided assistance and coordination to local agents in pursuing the mutually useful objective of limiting German supply lines and troop movements. In the territory FM study, however, the railroad targets that fulfilled these common strategic objectives were much more likely to be available on the formerly occupied side of the LoD. Hence, that is where the *maquis* tended to strike. In the presence of such unequivocal evidence of the importance of a coordinated effort to target major railway lines, FMR's vague claims about the role of local agency are not particularly relevant to evaluating KM's counter to FM's study. Unless FM can show that local motivations systematically led Resistance groups to pursue agendas that were orthogonal to the Allied plans, our criticism stands.

[9] A third commune, Vierzon, is dropped from FM's analysis because their maps indicate that its boundaries intersected the LoD.

Conclusion

In the conclusion, FMR offers a general defense of the value of design-based inference with which we broadly concur. In particular, we agree that natural experiments are valuable, in part, precisely because their "underlying assumptions are explicit and transparent" (FMR, 20) enough to be susceptible to clean refutation, which is often not the case with more conventional observational research. Where we differ is more a matter of degree than kind. FM imply that an "imperfect" natural experiment is better than no natural experiment at all. All we can say is that whether this is the case depends on the degree of imperfection. If a massive discontinuity in important pre-treatment variables lies in very nearly the same place as the discontinuity that determines treatment assignment—as is the case with FM's study—then a natural experiment is worse than no natural experiment because it will produce false certainty on highly biased estimators of causal effects.

In our view, FM put the challenge of design-based inference in terms that are too stark. In their view, researchers have to choose between controlling for all possible confounding variables or exploiting an exogenous source of random variation. We think researchers have more options, and we believe they can exercise more than one of them in the same study. In KM, we use several distinctive types of evidence (documents, secondary histories, maps, and statistical tests at multiple levels of aggregation) to produce a consistent and highly credible account of the distribution of violent resistance in the section of WWII France that FM analyzed in their paper. This account demonstrates that FM's argument is not supported by the data.

It is true that our approach does not allow us to estimate precisely the causal effect of a specific treatment: the presence of double-track railways. It is not entirely clear, however, what would be the use of such an estimate. Presumably, the strategic importance of railways and other kinds of infrastructure to armed groups varies tremendously by conflict, region, technology, era, and type of warfare, among other factors. We suspect this is also true for institutions of the sort FM investigate. If this is the case, then we have even more reason to be catholic in our approach to method.