A Data Appendix

A1 Summary Statistics

We report summary statistics for our main rural DHS sample in Table A1. For each of the 16 countries in our data, we report the share of the child population that is in DHS clusters that are classified as rural. Jordan has the smallest share of rural population in the data, with 31 percent, whereas Burkina Faso has the highest, with 80 percent. On average, countries in the data have at least half of their sample residing in DHS clusters that are classified as rural, with a mean across all countries of 71 percent. We also report the number of children in our DHS rural sample for each country. India is by far the largest country in terms of the number of observations it contributes to the data, with 175,854 children. Other countries have a sample size that is an order of magnitude smaller than India's, with values ranging from around 1,000 to about 40,000.

We report the mean and standard deviation for the two anthropometric scores we use as the primary measures of human impact: HAZ (standardized height), and the stunting dummy (when height-for-age is below -2). On average, children in our sample have HAZ values that point to malnourishment and delayed development. All countries have 10 percent or more of their children experiencing stunting, with an average across countries of 40 percent. For most of the children in the sample, experiencing exposure to a locust swarm is a rare event. On average, 1.3 percent of children are exposed to a swarm in the nine month before birth and 1.3 percent after birth. However, this low number is strongly driven by India, which accounts for 43 percent of the sample, where only about 0.55 percent of children are exposed to a swarm around the timing of their birth. In the other countries in our data, the proportion of children who are exposed to a swarm around the timing of their birth can be as high as 12 percent, in Senegal and Morocco.

A2 Locust Breeding Areas

In the main text, we describe the process we follow to choose the most appropriate clustering algorithm for outlining each season's breeding areas (optimizing clustering metrics and cross-checking with FAO maps to confirm relevance and break ties between outputs with similar clustering metrics). In practice, we obtain the best clustering output with the kmeans algorithm with four clusters for the spring season (Figure 2a), the dbscan algorithm with minimum points set at 50 and maximum radius set at 4 for the summer season, and the hdbscan algorithm with minimum points set at 25 for the winter season (Figure A1). In addition, we show that given similar parameter values, the three algorithms yield relatively similar results, which differ mainly in how they deal with outlier points (Figure A2). Among algorithms that optimize clustering metrics, choosing one that excludes outliers is more important in the summer and winter seasons for obtaining areas that are comparable with the FAO maps.

Table A1.
Summary Statistics for Main Estimation Sample

Country	Rural Pop. (%)	# Children	Standardized Height		Stunted Dummy		Swarm Exposure Around Birth			
							Observed (%)		Predicted	
			Mean	SD	Mean	SD	Before	After	Before	After
Burkina Faso	80	18,735	-1.5	1.9	.42	.49	3.9	4	4.1	4
Côte d'Ivoire	62	4,788	-1.3	1.7	.32	.47	0	0	.75	.71
Cameroon	59	6,503	-1.6	1.8	.41	.49	.17	.11	.51	.51
Chad	79	8,094	-1.4	2	.4	.49	3.2	2.9	1.6	1.6
Egypt	62	43,969	-.97	1.9	.26	.44	6.8	6.2	2.4	2.3
Ethiopia	83	25,828	-1.7	1.8	.46	.5	4.4	4.4	5.9	5.9
Guinea	74	7,555	-1.2	2	.36	.48	2.6	1.8	4	4.1
India	75	175,854	-1.5	1.8	.41	.49	.56	.54	.54	.56
Jordan	31	4,985	-.56	1.4	.13	.34	2.4	1.7	1.9	1.9
Kenya	72	20,318	-1.4	1.6	.35	.48	.045	.06	1	1.1
Mali	76	22,792	-1.6	1.8	.41	.49	6.4	6.2	5.8	5.9
Morocco	56	3,152	-.99	1.9	.27	.45	13	10	8.3	7.4
Niger	70	5,807	-1.7	1.9	.46	.5	7.4	6.6	7.2	7.1
Nigeria	69	38,240	-1.6	2.1	.45	.5	.36	.33	.92	.95
Pakistan	54	1,433	-1.9	1.5	.45	.5	3.1	3	2.1	2.2
Senegal	69	21,579	-1.1	1.3	.23	.42	12	12	9.6	10
All Countries	71	409,632	-1.5	1.8	.4	.49	1.3	1.3	1.2	1.2

Notes: For each of the countries in the data, we report the percentage of the children who are in rural DHS clusters, and the number of children in those rural clusters. For each country's rural sample, we summarize the mean and standard deviation of height-for-age and the stunting dummy. We also report the exposure to a locust swarm nine months before or after birth, using the observed FAO data, or the predicted swarm score we generate. Observations are weighted using DHS sample weights (see main text for more details).

Figure A1: Outlining Locust Summer and Winter Breeding Areas

(a) Best Clustering Output for the Summer Season

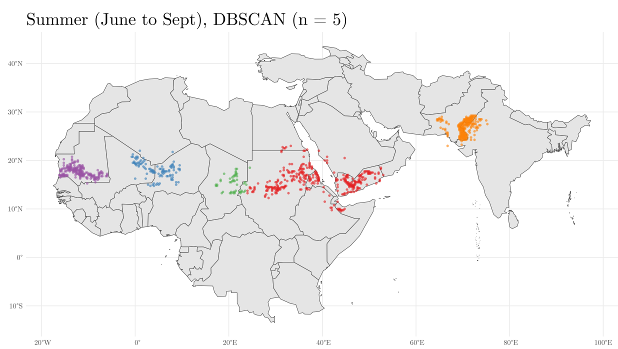

(b) Admin-1 Envelope Applied to Summer Clusters

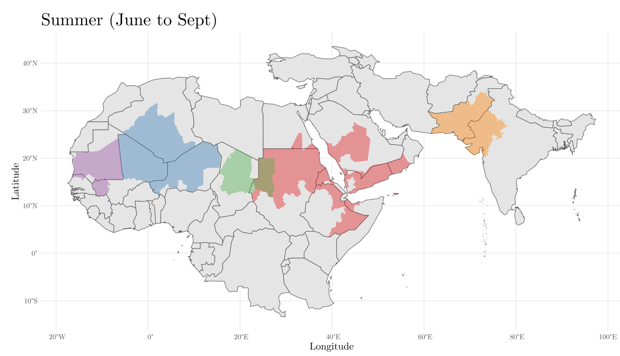

(c) Best Clustering Output for the Winter Season

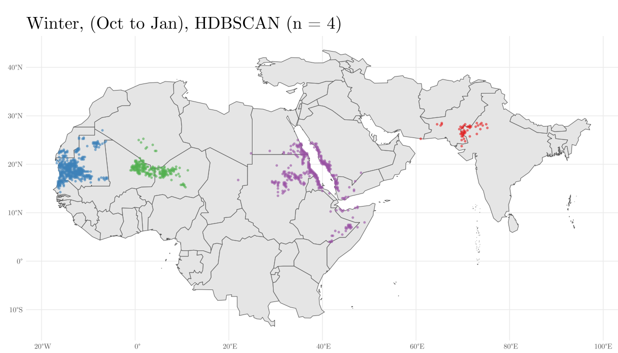

(d) Admin-1 Envelope Applied to Winter Clusters

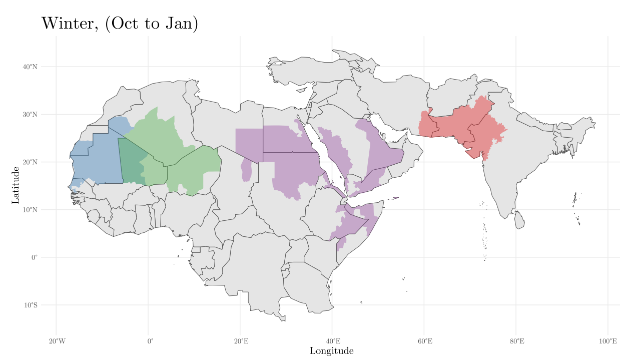

Notes: See main text for details.

Figure A2: Comparison of Breeding Area Clustering Algorithms

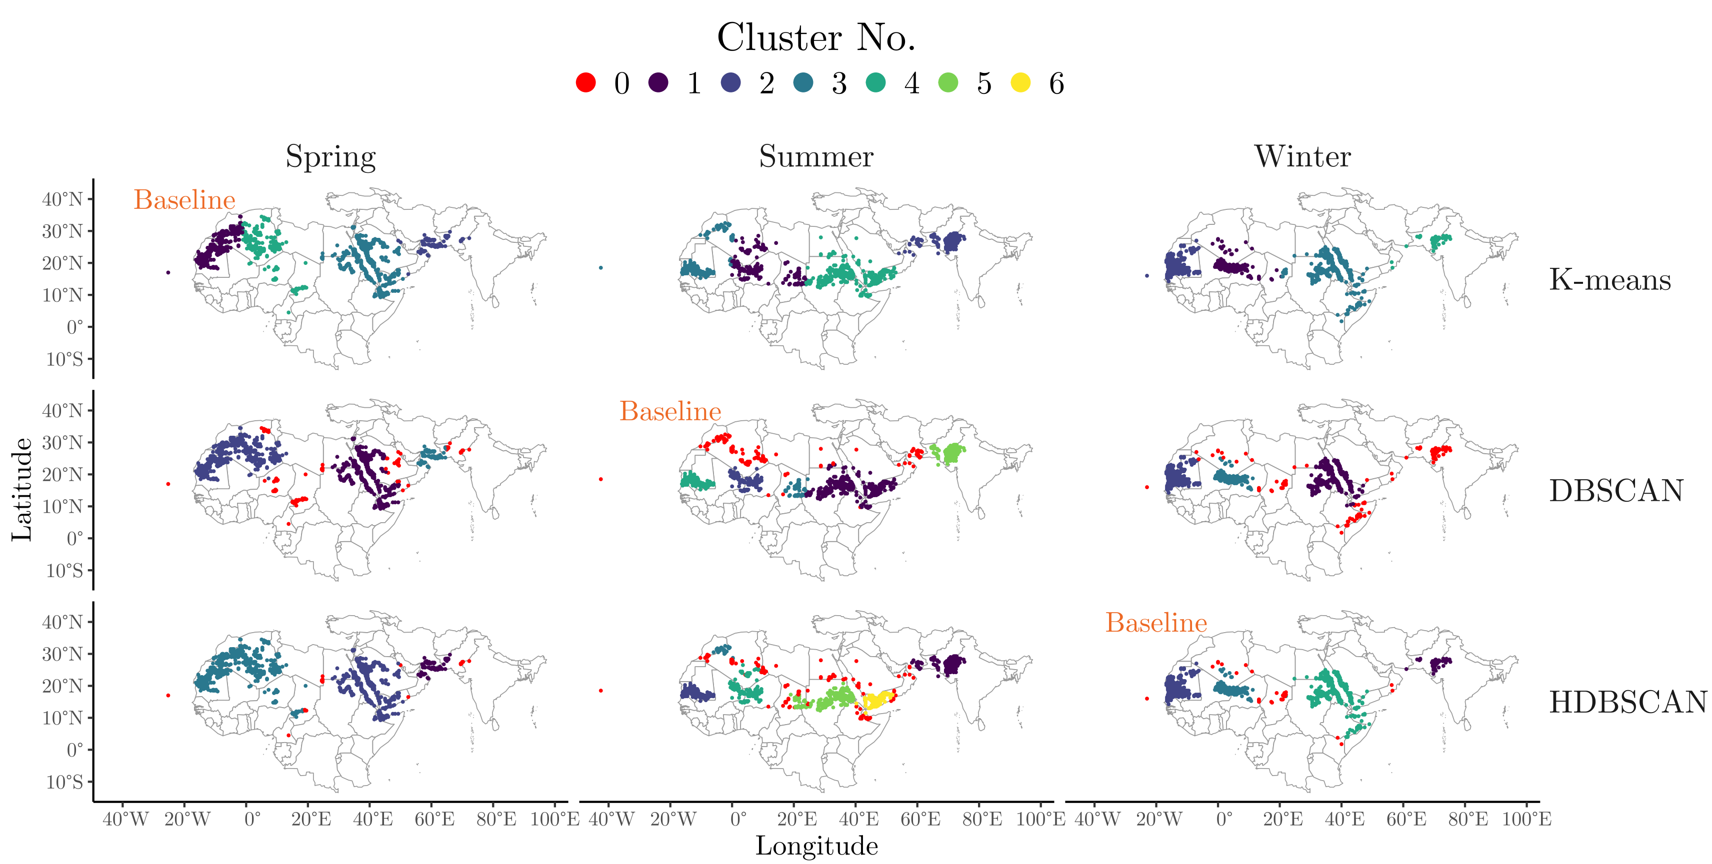

Notes: Columns correspond to different breeding seasons. Rows correspond to different clustering algorithms. Cluster “0” (red dots) corresponds to outliers that are excluded from formed clusters (note that only dbscan and hdbscan can exclude outliers). The “Baseline” annotation indicates which algorithm output is selected for each breeding season.

A3 Conflict Data

Figure A3: Spatiotemporal Variation in Armed Conflict

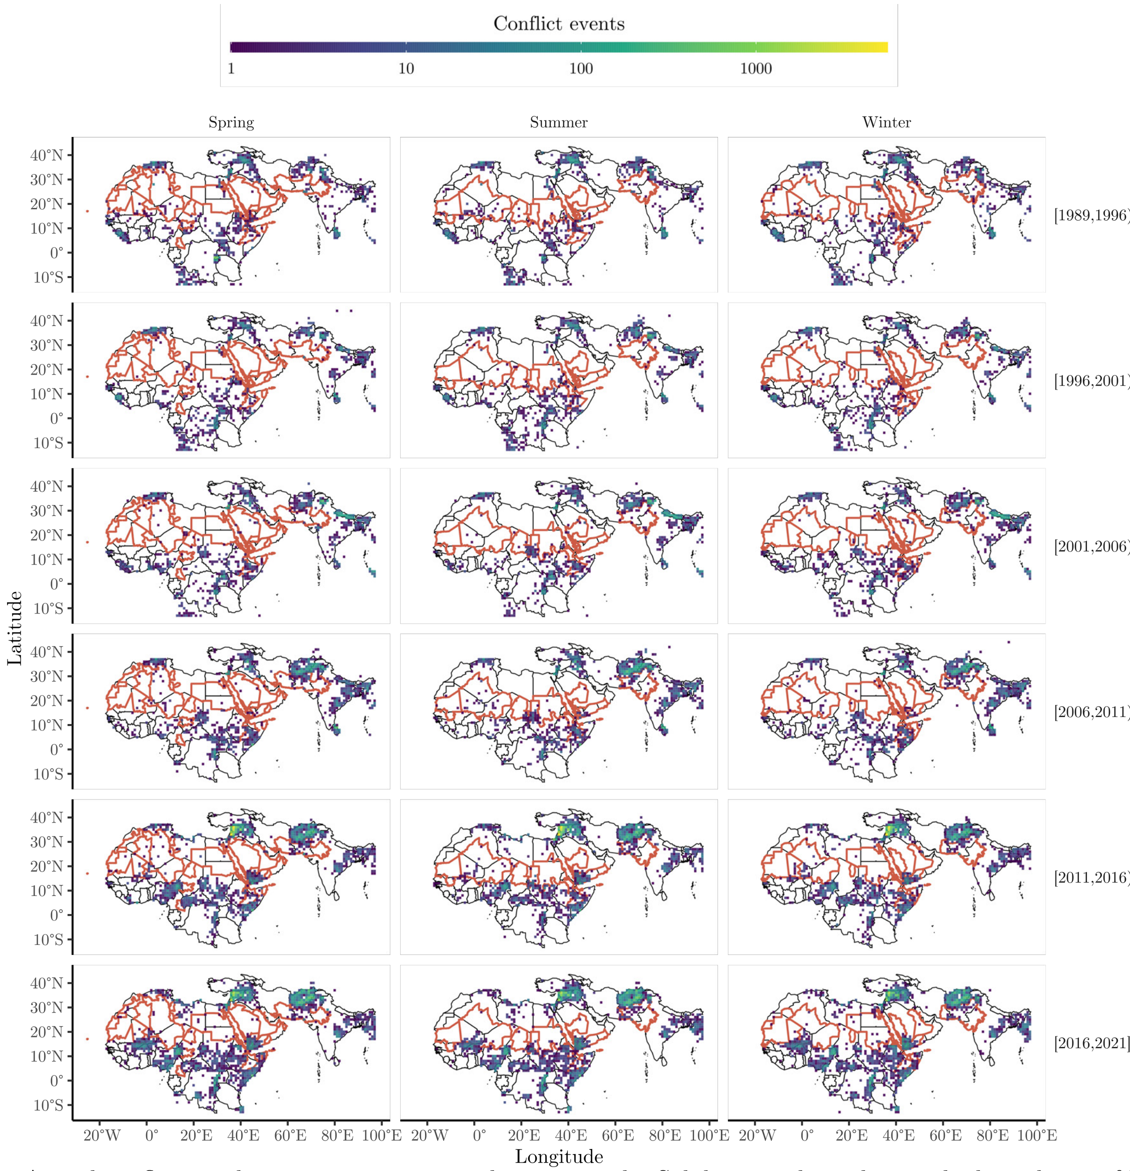

Notes: Armed conflict incidence across seasons and time periods. Solid orange lines denote the boundaries of locust breeding areas.

A4 Staple Food Prices

We use local market staple food prices comes from FEWS NET (Famine Early Warning System Network), which monitors the price of staple foods, cash crops, livestock, and labor traded at markets across Africa and Haiti. We obtain price data from 1995-2020 at 140 markets in eleven African countries.[20] We bin similar products into product groups, e.g., black beans, brown beans, and pinto beans are all mapped to the product group "beans." In 2019, Zimbabwe had multiple instances of a single product being exchanged for multiple currencies within the same market (e.g., maize grain sold in USD and ZWL at Harare, Mbare). To ensure product and market fixed effects absorb variation in currency, we run our analyses on two samples: one in which we drop Zimbabwe entirely, the other where we drop observations during and after 2019.

[20] Our set of staple foods are cassava, cassava chips, cassava flour, maize grain, maize meal, rice, roller maize meal, sorghum, sorghum flour, sweet potatoes, wheat flour, wheat grain, wheat bran, yams. All staple food quantities in our sample are reported in kilograms.

A5 Precipitation Distributed Lag Results

Figure A4: Swarm Incidence Regressed on Precipitation Shocks

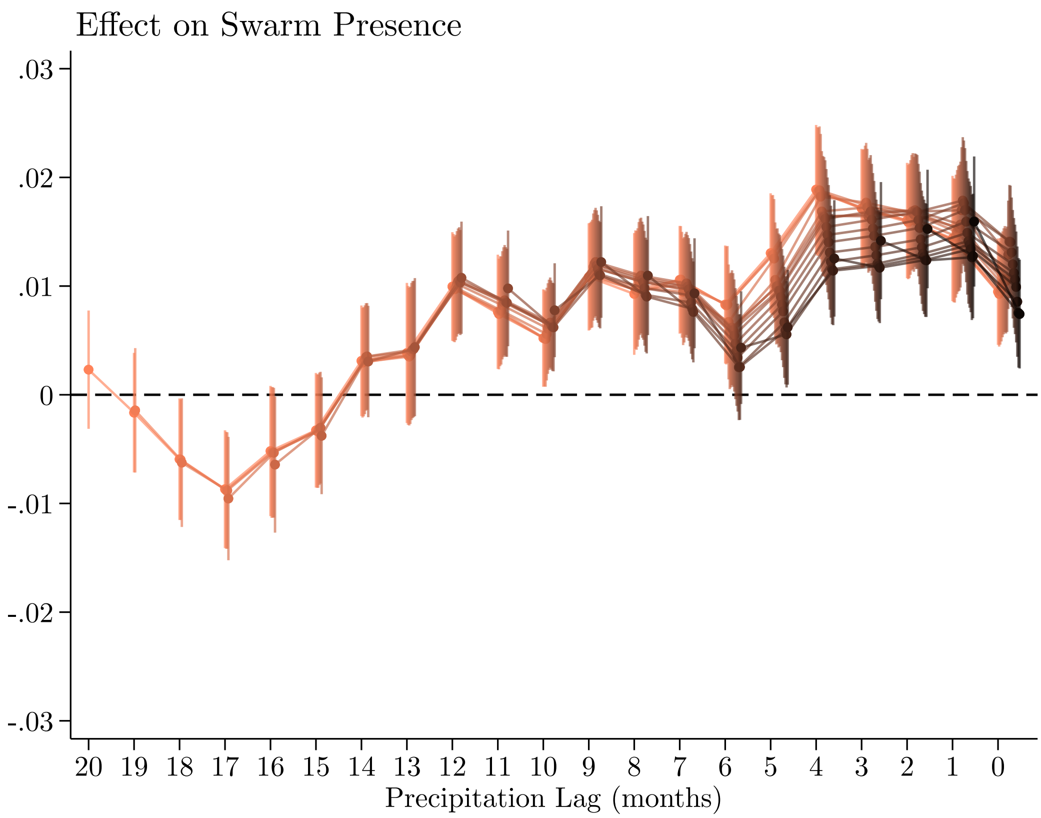

Notes: Coefficients from a set of regressions modeling $S = \sum_{l=0}^{L} p_l$, where p_l denotes the l-th precipitation shock lag and a model is run for each $L \in \{0, ..., 20\}$. Darker colors correspond to models with fewer lags. Plotted coefficients are jittered horizontally for visual clarity. Precipitation shocks are defined as a month in which any grid cell within an admin-1 experienced greater than 25mm of precipitation. Standard errors are clustered at the admin-1 level.

B Methods Appendix

B1 Construction of the Weather Data

Weather data come from ERA5, which is the fifth generation—and latest— reanalysis data product for global climate and weather produced by the European Centre for Medium-Range Weather Forecasts. In particular, this data product includes hourly data on precipitation and on minimum and maximum surface air temperature, 2 meters above the surface. In the main text, we explain that we convert this hourly weather data into precipitation totals and counts of days when average temperature falls into specific bins. In this section, we describe in more detail the temporal aggregation and spatial aggregation processes.

For temporal aggregation, we start by converting standardized time used by ERA5 into local time zones, so that daily aggregates reflect the local start and end of each day (and similarly, but less importantly, for monthly and annual aggregates). We aggregate hourly precipitation into precipitation totals by month or year, depending on the temporal unit of analysis. We aggregate hourly minimum and maximum temperature by day by first calculating hourly mean temperature, and then taking its daily average. We then further aggregate by counting, for different temperature bins, the number of days when the daily average falls into that bin, per month or year depending the temporal unit of analysis. For spatial aggregation, we take the average of precipitation totals and of counts of days spent in different temperature bins, across all ERA5 grid cells intersecting with the target spatial unit of analysis (first-level administrative boundary, locust breeding area, or 150 km buffer around clusters of Demographic and Health Surveys). When computing averages, we weight each grid cell by the extent of overlap between the grid cell and the target spatial unit.

B2 Determining the Spatial Extent of Local Locust Exposure

To determine the spatial extent of locust exposure, we evaluate the spatial clustering of the locust point data by using Ripley's K-function (Goreaud and Pélissier 1999; Haase 1995). The intuition is that we assess whether a point has a higher density of neighboring points—the K value—relative to what we would expect given the average spatial density of all points in the study area, if they were randomly distributed across space.[21] In Figure B1, we report that the difference between the observed and expected K values increases with distance, and peaks around the cutoff of 150 km. When the observed K is higher than the expected K, which is calculated under the assumption of complete spatial randomness, the data are considered to be spatially clustered. We choose the threshold of 150 km as a conservative threshold. Even though the observed K value is higher than the expected K value as far out as 350 km, there is a clear change in the deviation between observed and expected K values at 150 km. In robustness tests, we extend the 150 km threshold, run "donut" versions of the analysis, and allow treatment effects to vary by distance.

[21] A succinct review of the Ripley's K function method is available at: https://pro.arcgis.com/en/pro-app/3.0/tool-reference/spatial-statistics/h-how-multi-distance-spatial-cluster-analysis-ripl.htm. Accessed: December 12, 2025.

Figure B1: Evidence for Spatial Clustering of Locust Swarms

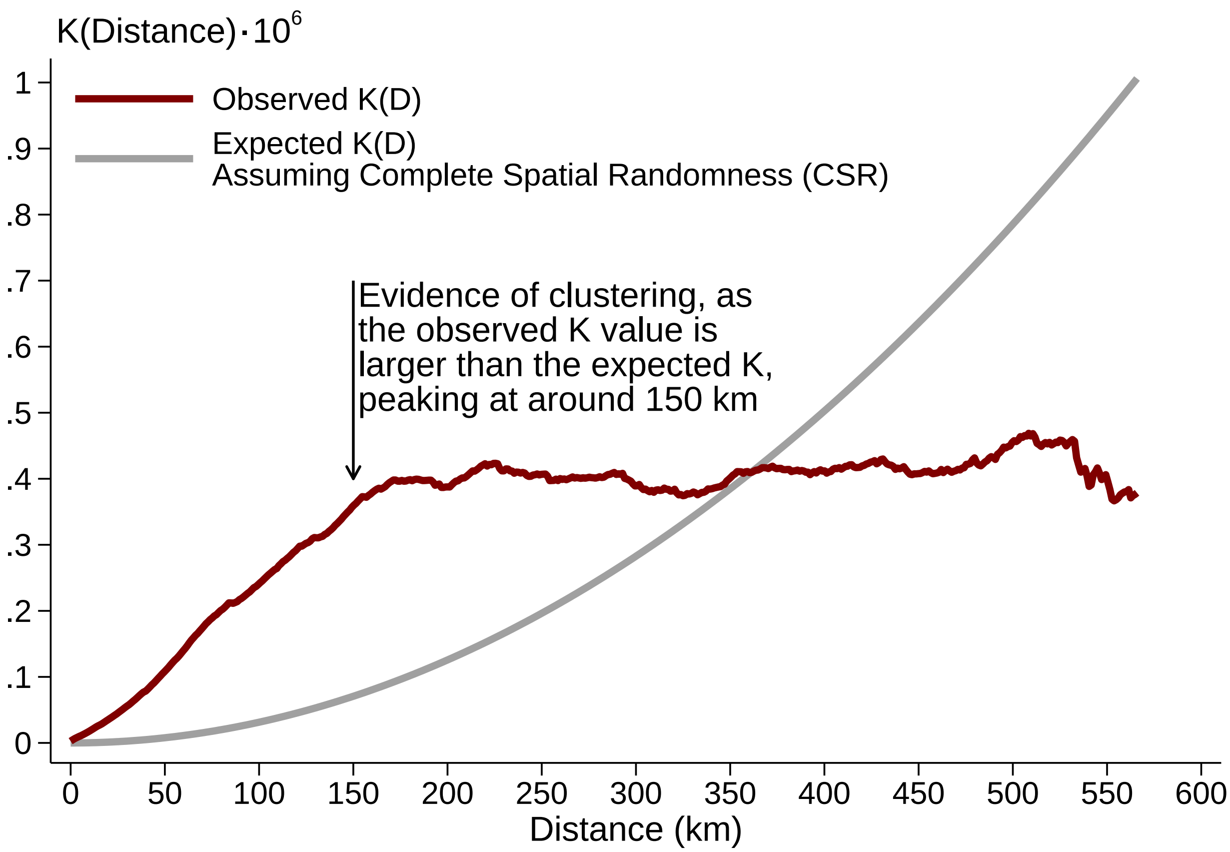

Notes: We use all years and countries in the sample to assess the degree of spatial clustering of swarm events using Ripley's K function. There is evidence of spatial clustering that extends as far as 350 km, however, we use 150 km as the threshold to define exposure because the observed K function peaks at 150 km and remains flat thereafter.

B3 Latent Exposure, Endogenous Misclassification, and an Exogenous Proxy

B3.1 Setup

Let the structural model be:

$$y_i = \alpha + \beta L_i^* + \varepsilon_i,$$

where y_i is the outcome of interest for child i, $L_i^* \in \{0, 1\}$ is true locust exposure, and ε_i is an unobserved determinant of y_i. The parameter of interest is β. To simplify notation, we suppress fixed effects and additional controls and present a single-exposure-period version of the model. The same arguments apply to the specification in the main text after residualizing all variables with respect to controls and allowing for multiple exposure periods.

In the absence of measurement error, identification of β requires:

$$\operatorname{Cov}(L_i^*, \varepsilon_i) = 0.$$

True exposure L_i^* is not observed. Instead, we observe:

1. A binary indicator L_i, constructed from FAO reports and subject to misclassification.
2. A continuous proxy score S_i, constructed from external predictors (here, locust activity in distant breeding areas) and intended to capture exposure risk.

B3.2 Misclassification of the Binary Indicator

Exogenous Misclassification We call misclassification *exogenous* if the reporting error is unrelated to unobserved determinants of the outcome, conditional on true exposure: $L_i \perp \varepsilon_i \mid L_i^*$. This condition implies $\operatorname{Cov}(L_i, \varepsilon_i) = 0$, so OLS bias arises only through the imperfect correlation between L_i and L_i^*.

To obtain a closed-form attenuation result, we additionally consider the common special case in which misclassification is *nondifferential* with respect to observables, i.e. the classification error depends only on the true state: $\Pr(L_i = 1 \mid L_i^*, X_i) = \Pr(L_i = 1 \mid L_i^*)$ for all observables X_i, so that: $\pi_{10} \equiv \Pr(L_i = 0 \mid L_i^* = 1), \pi_{01} \equiv \Pr(L_i = 1 \mid L_i^* = 0)$ are constants.

Under nondifferential misclassification, $\mathbb{E}[L_i \mid L_i^*] = (1 - \pi_{10})L_i^* + \pi_{01}(1 - L_i^*)$ and it follows that $\operatorname{Cov}(L_i, L_i^*) = (1 - \pi_{10} - \pi_{01}) \operatorname{Var}(L_i^*)$.

Consider the regression

$$y_i = \alpha + \tilde{\beta} L_i + u_i,$$

where $u_i \equiv \varepsilon_i + \beta(L_i^* - L_i)$ collects the structural error and the measurement error component.

Using $y_i = \alpha + \beta L_i^* + \varepsilon_i$, we can write

$$\operatorname{Cov}(L_i, y_i) = \beta \operatorname{Cov}(L_i, L_i^*) + \operatorname{Cov}(L_i, \varepsilon_i).$$

Since the exogeneity condition $L_i \perp \varepsilon_i \mid L_i^*$ implies $\mathrm{Cov}(L_i, \varepsilon_i) = 0$:

$$\operatorname{plim} \hat{\hat{\beta}} = \frac{\mathrm{Cov}(L_i, y_i)}{\mathrm{Var}(L_i)} = \beta \frac{\mathrm{Cov}(L_i, L_i^*)}{\mathrm{Var}(L_i)}.$$

Substituting the expression for $\mathrm{Cov}(L_i, L_i^*)$ yields

$$\operatorname{plim} \hat{\hat{\beta}} = \beta\,(1 - \pi_{10} - \pi_{01})\,\frac{\mathrm{Var}(L_i^*)}{\mathrm{Var}(L_i)}.$$

Under nondifferential misclassification with $a := 1 - \pi_{10} - \pi_{01} \in (0, 1)$, $\mathrm{Cov}(L_i, L_i^*) = a\ \mathrm{Var}(L_i^*)$ and one can show $\mathrm{Var}(L_i) \geq a\ \mathrm{Var}(L_i^*)$ (with strict inequality whenever $\pi_{10}, \pi_{01} > 0$ and $0 < \Pr(L_i^* = 1) < 1$). Hence $0 < a\ \mathrm{Var}(L_i^*)/\mathrm{Var}(L_i) < 1$, so $|\operatorname{plim} \hat{\hat{\beta}}| < |\beta|$, and the OLS coefficient is attenuated toward zero (as in Aigner et al. 1973).

Endogenous Misclassification Misclassification is *endogenous* if the probability of error depends on ε_i, implying:

$$\mathrm{Cov}(L_i, \varepsilon_i) \neq 0.$$

In this case,

$$\operatorname{plim} \hat{\hat{\beta}} = \beta \frac{\mathrm{Cov}(L_i, L_i^*)}{\mathrm{Var}(L_i)} + \frac{\mathrm{Cov}(L_i, \varepsilon_i)}{\mathrm{Var}(L_i)}.$$

The first term reflects attenuation as above. The second term has unrestricted sign. Consequently, the overall bias need not be toward zero.

B3.3 Regressions Using the Proxy Score S_i

Consider the regression:

$$y_i = a + \theta S_i + \nu_i.$$

Given the structural model for y_i and under the identifying assumption of proxy exogeneity $\mathrm{Cov}(S_i, \varepsilon_i) = 0$,

$$\theta = \frac{\mathrm{Cov}(S_i, y_i)}{\mathrm{Var}(S_i)} = \beta \frac{\mathrm{Cov}(S_i, L_i^*)}{\mathrm{Var}(S_i)}.$$

Thus θ equals β multiplied by a scaling factor reflecting how informative S_i is about true exposure. Unless S_i is a known affine transformation of L_i^*, $\theta \neq \beta$.

Because S_i is trained on a systematically undercounted label L_i, it is natural for S_i to be conservative in the sense that it understates exposure risk in high-misclassification regions. Formally, if S_i were trained on the true L_i^* rather than L_i, we would expect $\mathrm{Cov}(S_i, L_i^*)$ to be larger, and hence the scaling factor $\mathrm{Cov}(S_i, L_i^*)/\mathrm{Var}(S_i)$ to be closer to one, narrowing the gap between θ and β. Training on the mismeasured L_i instead compresses this scaling factor, so that θ understates $|\beta|$ beyond what would arise from classical measurement error alone. In addition, to the extent that the constructed score contains classical prediction error (e.g., finite-sample estimation noise), this further attenuates the regression coefficient on S_i.

B3.4 Why S_i is not a valid instrument for L_i

Suppose we attempt to recover the structural parameter β by instrumenting the misclassified exposure L_i with the proxy score S_i. The corresponding IV specification is

$$y_i = \alpha + \beta L_i + \tilde{\varepsilon}_i,$$

where, using $y_i = \alpha + \beta L_i^* + \varepsilon_i$,

$$\tilde{\varepsilon}_i = \varepsilon_i + \beta(L_i^* - L_i).$$

For S_i to be a valid instrument, two conditions must hold:

1. Relevance: $\mathrm{Cov}(S_i, L_i) \neq 0$.
2. Exclusion: $\mathrm{Cov}(S_i, \tilde{\varepsilon}_i) = 0$.

Even if $\mathrm{Cov}(S_i, \varepsilon_i) = 0$ holds under proxy exogeneity, exclusion additionally requires $\mathrm{Cov}(S_i, L_i^* - L_i) = 0$. There are two distinct reasons this condition is unlikely to hold.

Constant Misclassification Under exogenous misclassification,

$$\mathbb{E}[L_i \mid L_i^*] = (1 - \pi_{10})L_i^* + \pi_{01}(1 - L_i^*).$$

Rearranging,

$$\mathbb{E}[L_i^* - L_i \mid L_i^*] = \pi_{10}L_i^* - \pi_{01}(1 - L_i^*).$$

Taking covariance with S_i and applying the law of iterated expectations,

$$\begin{aligned}
\mathrm{Cov}(S_i, L_i^* - L_i) &= \mathbb{E}[S_i \cdot \mathbb{E}[L_i^* - L_i \mid L_i^*]] - \mathbb{E}[S_i]\,\mathbb{E}[L_i^* - L_i] \\
&= \pi_{10}\,\mathrm{Cov}(S_i, L_i^*) + \pi_{01}\,\mathrm{Cov}(S_i, 1 - L_i^*) \\
&= (\pi_{10} + \pi_{01})\,\mathrm{Cov}(S_i, L_i^*).
\end{aligned}$$

Whenever the instrument is relevant, i.e. $\mathrm{Cov}(S_i, L_i^*) \neq 0$, it is mechanically correlated with the measurement error component $L_i^* - L_i$. Exclusion fails unless misclassification is absent.

Misclassification Varies with Risk Suppose the probability of correct classification varies with S_i. For example, suppose $\Pr(L_i = 1 \mid L_i^* = 1, S_i = s)$ is increasing in s. Then $\mathbb{E}[L_i^* - L_i \mid S_i = s]$ varies systematically with s, and so

$$\begin{aligned}
\mathrm{Cov}(S_i, L_i^* - L_i) &= \mathbb{E}[S_i \cdot \mathbb{E}[L_i^* - L_i \mid S_i]] - \mathbb{E}[S_i]\,\mathbb{E}[L_i^* - L_i] \\
&= \mathbb{E}[(S_i - \mathbb{E}[S_i]) \cdot \mathbb{E}[L_i^* - L_i \mid S_i]] \neq 0,
\end{aligned}$$

since $\mathbb{E}[L_i^* - L_i \mid S_i]$ is a non-constant function of S_i. Exclusion therefore fails directly.

B3.5 Summary

OLS: y on L Regressing y_i on the misclassified indicator L_i produces attenuation under exogenous misclassification, but once misclassification is endogenous—i.e., depends on ε_i—the resulting coefficient may be biased in arbitrary directions.

OLS: y on S Regressing y_i on the proxy score S_i identifies a scaled version of β, where the scaling factor depends on how informative S_i is about L_i^*. Note that in our final cost-benefit calculation this scaling will be irrelevant, since we will multiply the coefficient from the proxy regression by predicted additional exposure ΔS_i, which is expressed in the same units as S_i. Multiplying $\hat{\theta} \cdot \Delta S_i$ therefore recovers the quantity $\beta \cdot \Delta L_i^*$—the effect of true additional exposure—without knowledge of the scaling factor.

2SLS: L instrumented by S Using S_i as an instrument for L_i is invalid under either constant or risk-dependent misclassification. Relevance implies correlation with measurement error.

B3.6 Monte Carlo Simulation

Design To illustrate the theoretical results in Sections B3.1 to B3.4, we conduct a Monte Carlo simulation. The data-generating process follows the structural model in Section B3.1:

$$y_i = \alpha + \beta L_i^* + \varepsilon_i,$$

with $\alpha = 0$, $\beta = 1$, $\varepsilon_i \sim \mathcal{N}(0, 1)$, and sample size $N = 200$. We model L^* based on an underlying risk variable $S_i \sim \text{Uniform}(0, 1)$, which we will later treat as the proxy variable.

True exposure L^* is then the binary realization:

$$L_i^* = \mathbf{1}(U_i < S_i), \quad U_i \sim \text{Uniform}(0, 1),$$

where U_i is an independent draw. Under this construction $\Pr(L_i^* = 1 \mid S_i) = S_i$, so S_i is a valid and informative proxy for true exposure by design, and $\text{Cov}(S_i, \varepsilon_i) = 0$ holds by independence.

We simulate 300 replications under each of three scenarios that differ in how the misclassified indicator L_i is constructed from L_i^*.

Scenario 1: Exogenous misclassification. The false-negative probability is constant and independent of ε_i:

$$\Pr(L_i = 0 \mid L_i^* = 1) = \pi_{10} = 0.25.$$

There are no false positives ($\pi_{01} = 0$). This scenario corresponds to the nondifferential misclassification case in Section B3.2. Setting false positives to zero is a straightforward way to model the feature of the data that the rate of false negatives exceeds the rate of false positives, but it is not critical for the results of the simulation.

Scenarios 2 and 3: Endogenous misclassification. The false-negative probability depends on the structural error ε_i through a logistic function:

$$\Pr(L_i = 0 \mid L_i^* = 1, \varepsilon_i) = \text{logit}^{-1}(c_0 + c_1\varepsilon_i),$$

with $c_0 = -0.8$ in both cases. In Scenario 2 we set $c_1 = -1.5$, so that units with higher ε_i (better unobserved outcomes) are *less* likely to be misclassified and hence more likely to appear as exposed in the data. In Scenario 3 we set $c_1 = 1.5$, so that units with higher ε_i are *more* likely to be misclassified. These scenarios correspond to the endogenous misclassification case in Section B3.2, with the sign of the induced correlation $\text{Cov}(L_i, \varepsilon_i)$ differing across the two cases to illustrate the potential for bias of uncertain sign.

For each scenario and replication, we estimate four models: OLS of y_i on L_i^* (the infeasible benchmark), OLS of y_i on L_i, OLS of y_i on S_i, and 2SLS of y_i on L_i instrumented by S_i.

Results Appendix Figure B2 displays the sampling distributions of the four estimators across the three scenarios. Each panel shows a histogram of the estimated coefficient over 300 replications; vertical red lines mark the true value $\beta = 1$ in all columns except OLS on S_i.

OLS on L^* (column 1). Across all scenarios the OLS estimator using the true exposure is unbiased and centered on $\beta = 1$, confirming that identification holds in the absence of measurement error, as long as L^* is uncorrelated with ϵ.

OLS on L (column 2). Under exogenous misclassification (Scenario 1) the distribution is centered below unity, consistent with the attenuation result in Section B3.2: with $\pi_{10} = 0.25$ and $\pi_{01} = 0$, the OLS coefficient is biased toward zero. Under endogenous misclassification the direction of bias depends on the sign of $\text{Cov}(L_i, \varepsilon_i)$. In Scenario 2 ($c_1 = -1.5$) the distribution shifts upward relative to unity, reflecting a positive bias that pushes the estimate away from zero. In Scenario 3 ($c_1 = 1.5$) the distribution shifts towards zero, illustrating how endogenous misclassification can produce bias of any sign. These patterns illustrate the result in Section B3.2 that the OLS bias under endogenous misclassification is unrestricted.

OLS on S (column 3). The distribution of the OLS coefficient on S_i is centered below unity in all scenarios, reflecting the scaling factor $\text{Cov}(S_i, L_i^*)/\text{Var}(S_i) < 1$ derived in Section B3.3. Although in the simulation S_i is the primitive from which L_i^* is drawn, the binary realisation introduces noise: $\text{Var}(S_i) > \text{Cov}(S_i, L_i^*)$ because L_i^* is a Bernoulli draw conditional on S_i rather than equal to it, so the regression coefficient on S_i understates $|\beta|$. In practice, we only observe an imperfect predictor of S_i, which would result in further attenuation. The distribution is naturally similar across all three scenarios, consistent with proxy exogeneity $\text{Cov}(S_i, \varepsilon_i) = 0$ holding by construction.

2SLS instrumenting L with S (column 4). The 2SLS estimator exhibits severe bias and high dispersion in all three scenarios. In all three cases, the distribution is substantially right-skewed and centered well above unity, consistent with the invalidity results in Section B3.4: because S_i is relevant for L_i^*, it is mechanically correlated with the measurement error $L_i^* - L_i$, violating the exclusion restriction. For brevity, we do not model the second potential exclusion restriction violation.

Summary. The simulation confirms the analytical results: exogenous misclassification produces attenuation bias in OLS on L; endogenous misclassification can produce bias of arbitrary sign; OLS on S consistently understates $|\beta|$ due to the scaling of the proxy; and 2SLS using S as an instrument for L fails to recover β in any scenario, underscoring the invalidity argument in Section B4.

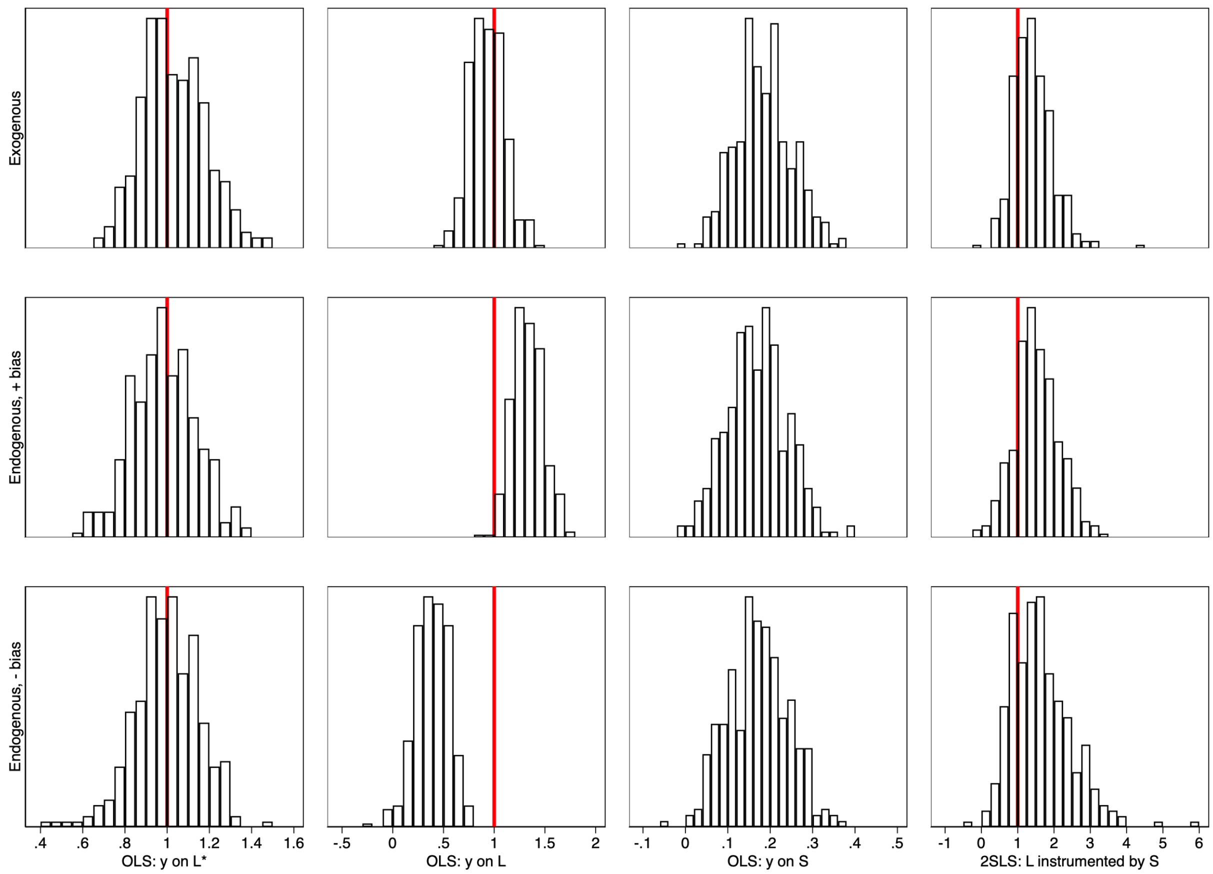

Figure B2: Sampling distributions of four estimators across three misclassification scenarios. Each histogram summarizes 300 Monte Carlo replications with $N = 200$. Columns correspond to estimators: OLS on L^* (infeasible benchmark), OLS on L, OLS on S, and 2SLS instrumenting L with S. Rows correspond to scenarios: exogenous misclassification (top), endogenous misclassification with positive OLS bias (middle), and endogenous misclassification with negative OLS bias (bottom). Red vertical lines mark the true parameter value $\beta = 1$ in columns 1, 2, and 4.

B4 Case Study: The Stunting Cost of Monitoring Interruptions in Yemen

Figure B3: Yemen breeding areas and LASSO teleconnections

(a) Breeding Areas Overlapping with Yemen

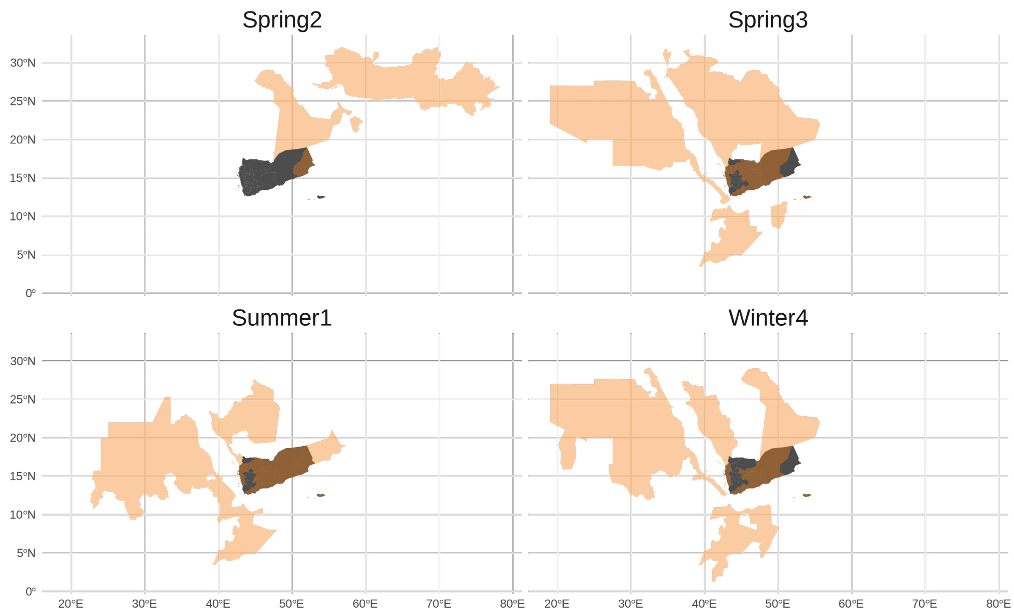

(b) Teleconnections from Yemen Breeding Areas

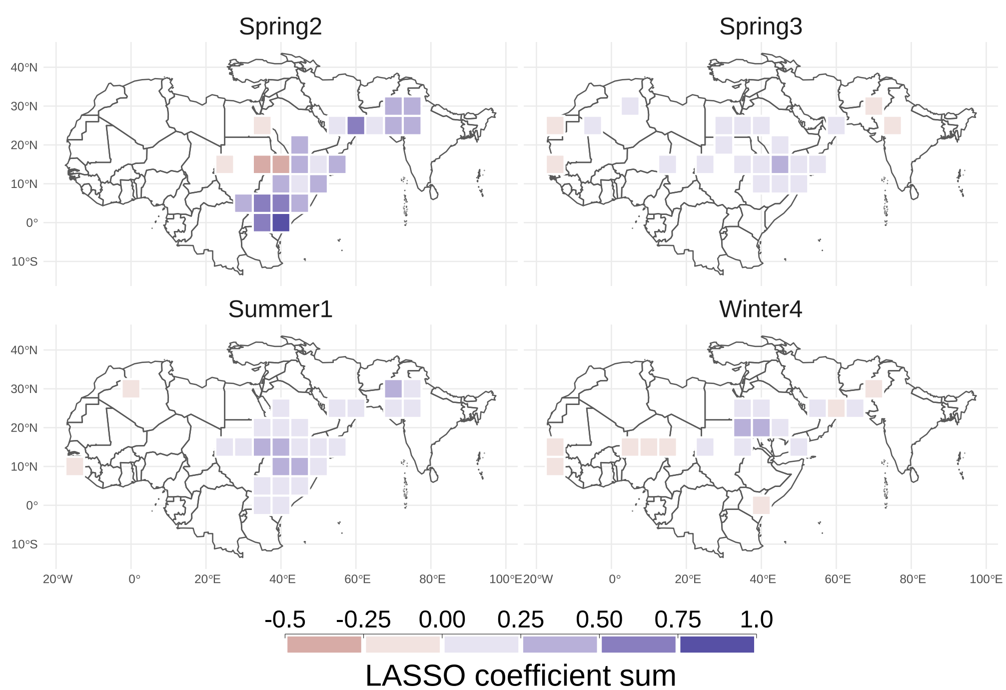

Notes: (a) The four breeding area admin-1 envelopes (orange) which overlap with Yemen (dark gray). (b) The teleconnections between Yemen breeding areas and destination grid cell are generated by summing the coefficients of a given predictor breeding area selected by a destination grid cell.

C Results Appendix

C1 Suggestive Evidence of Endogenous Monitoring

Like many non-automated and costly surveillance systems, locust monitoring is incomplete: this means that the absence of a locust observation in a particular place and time can either reflect the actual absence of locust activity or simply low surveillance or reporting effort. In this section, we provide suggestive evidence of endogenous variation in the locust monitoring data.

The nature of the locust monitoring system suggests that monitoring is likely endogenous to surveillance effort (itself likely tied to suspected locust conditions) and capacity. While the FAO provides training on locust control and assists countries with equipment, resources, and technical support during outbreaks and plagues, locust surveillance is primarily the responsibility of national governments. Much of the locust monitoring data collection process thus depends on the willingness and capability of governments to carry out monitoring activities. In particular, countries with lower state capacity likely exert lower levels of monitoring effort.

We illustrate the likely endogeneity of monitoring data to state capacity by showing strong correlations between monitoring intensity and measures of state capacity.

In Figure C1 we consider two indicators of state capacity from the Worldwide Governance Indicators set (Kaufmann and Kraay 2024)—Government Effectiveness (Figure C1a) and Regulatory Quality (Figure C1b)—and relate them to the intensity of monitoring at the country and year level. Intensity of monitoring is measured as the average share of months in a year when an admin-1 has reported any ecology observation, across all admin-1 regions of a country that ever report ecology observations. We regress this metric on each of the governance indicators and present the results in Table C1, while the corresponding regression lines are plotted in Figure C1.

Next we extract reports of monitoring interruptions made by the FAO to provide additional evidence of the incompleteness of locust monitoring. The FAO's Desert Locust Information Service issues monthly Desert Locust Bulletins that summarize and forecast global to national desert locust conditions. These bulletins include instances when field monitoring activities cannot be conducted, sometimes along with a statement on the cause of failure (e.g., "insecurity" in the case of a conflict-related interruption, or "no locust reports were received" in the case of an unspecified interruption). We scrape text from the bulletins between 1996 and 2020, and use it to create an indicator at the country and month level for whether a country experienced a monitoring interruption in a given month, whether complete (national) or partial (local).

Monitoring disruptions exhibit considerable variability over space and time (Figure C2). Figure C2a corroborates that monitoring intensity is related to need: during significant upsurges (2004-5, 2020), reports of monitoring failures are almost nil, and reports of all types (ecology and all life stages of locusts, see Figure 1a in the main text) abound. However most of the year-on-year variation seems to come from other factors. Figure C2b shows great disparities in where those disruptions happen. Low state capacity and presence of conflict cannot fully account for the patterns observed.

Figure C1: Monitoring Intensity and Quality of Governance

(a) Government Effectiveness

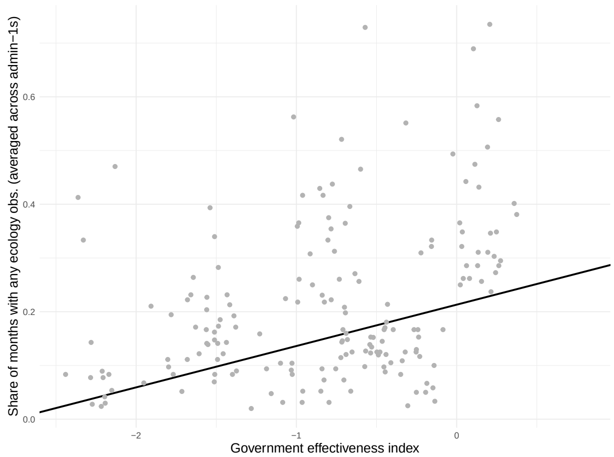

(b) Regulatory Quality

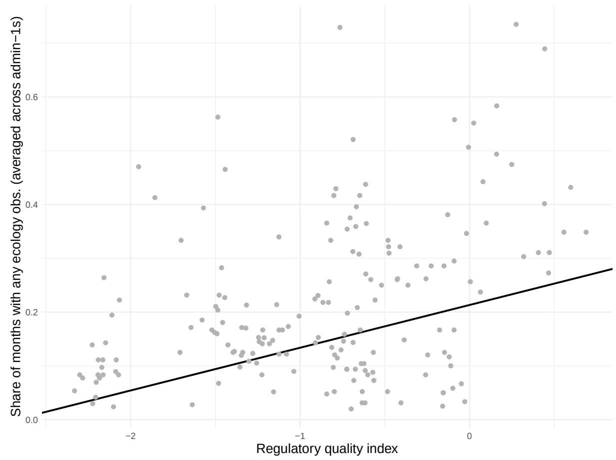

Notes: Figures plot the intensity of monitoring versus quality of governance indices (a) Government Effectiveness (b) Regulatory Quality, over 2010-20. Intensity of monitoring in both y-axes is obtained by counting the months with any ecology reporting in every administrative unit (admin-1) that ever reported ecology observations to obtain the average share of periods with reporting over a country and year.

Table C1: Ecology Observations and Quality of Governance

	Government Effectiveness	Regulatory Quality
	(1)	(2)
World Governance Indicator	0.077***	0.079***
	(0.014)	(0.014)
Observations	176	176
R^2	0.199	0.212

Notes: Table shows OLS regressions of the normalized number of ecology observations on governance indicators at the country year level, with year fixed effects (11 years, 2010-2020). The governance indicators come from the World Bank's World Governance Indicators set. The ecology observations are aggregated up at the country-year level from the raw observations: for every country, we count the months in which its admin-1s recorded at least one 'ecology' observation, obtain the share of months where each admin-1 had any ecology observation, and average over the country's admin-1s (keeping only those that ever had an ecology observation). $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

Figure C2: Monitoring Disruptions as Reported in FAO Bulletins

(a) Over Time

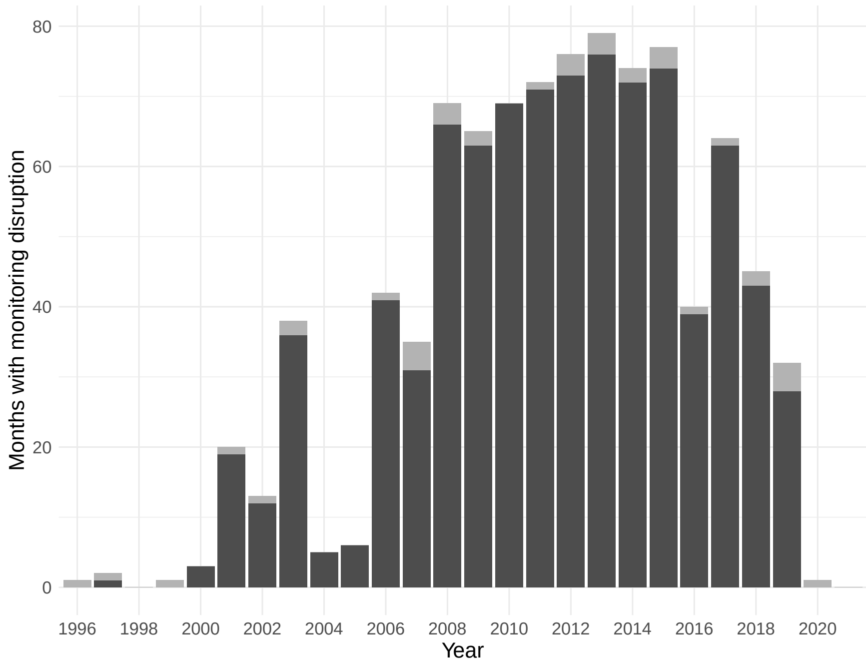

(b) By Country

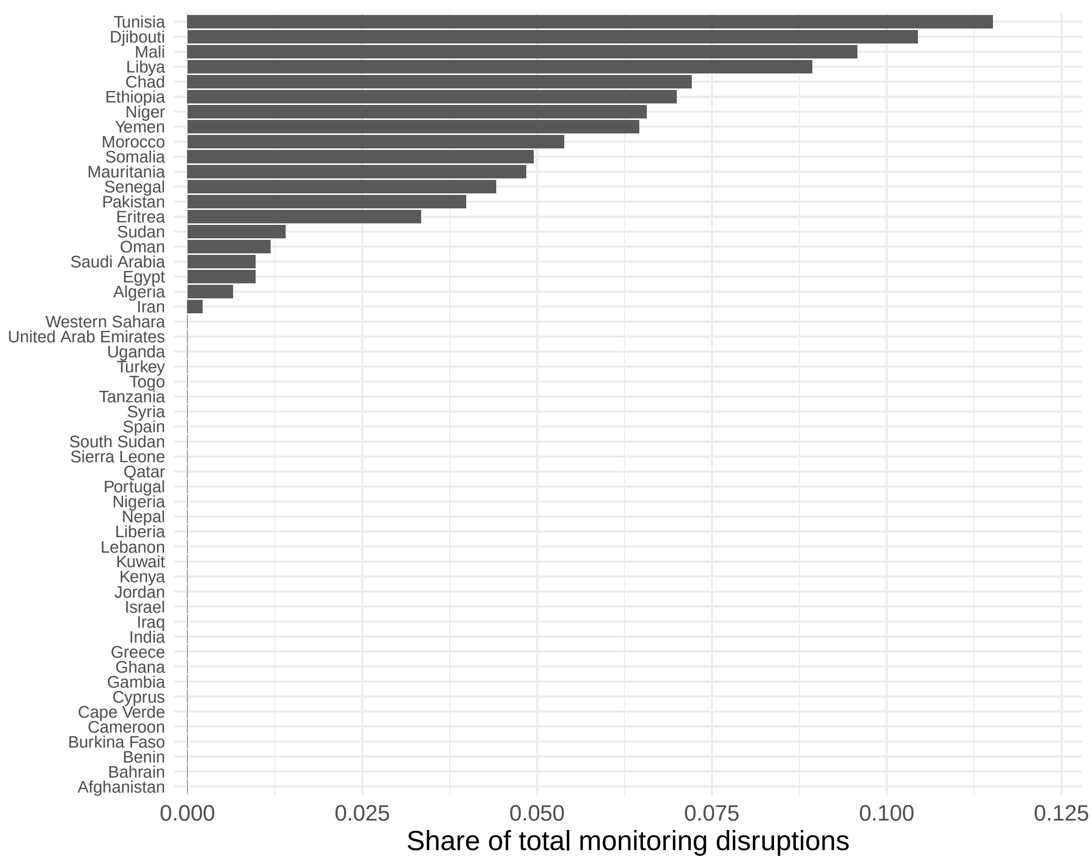

Notes: Figures show the temporal (a) and spatial (b) distributions of monitoring disruptions mentioned in the monthly FAO Desert Locust Bulletins. Panel (a) distinguishes between partial (light gray) and complete monitoring disruptions (dark gray).

C2 Sample & Specification Variations: Impact of Conflict on Monitoring and Swarms

Sensitivity to Monitoring Effort Metric In the main text, we document the impact of armed conflict on monitoring, using any locust report to capture monitoring effort. In Table C2, we provide an extended version of the results in the main paper, showing results for the number of days of recorded events as well as the dummy variable for any recorded event. In Table C2, panel A, we report coefficients from estimating Equation (1), using any monitoring or any swarm dummy variables, same as in Table 1, and extend the analysis to include any ecology or control reports. For ecology reports, which serves as an alternative measure for monitoring efforts that is not affected by the presence of any locust, the same pattern is observed: conflict suppresses ecology reports. This supports the interpretation that the sharp declines in any monitoring are predominantly driven by lower levels of monitoring and not lower levels of locusts. For control activities, we fail to detect a precisely estimated effect. As we explain in the main text, the disruption to monitoring can lower control likelihood because there is signal about a buildup of a swarm, or raise control likelihood because the disruption to conflict has resulted more swarms. In panel B, we report results for the number of days in an admin-1-month-year that have reported locust events. The number of days capture both an extensive and intensive margin. The results in panel B are similar to those in panel A, yet they capture that relative to the mean level of the number of days, conflict can reduce monitoring even in cases where monitoring is not completely shut down.

In the main text, we use a conflict dummy or variables that capture the share of recent months that have had conflict events. Here, we define three new variables that allow us to more flexibly capture the dynamics of monitoring disruptions around conflict events: Onset_{at} is a dummy variable that equals one if a conflict started at any period from $t = 0$ to $t = -12$ in admin-1 a after a period of at least 12 months in which no conflict was reported; Ongoing_{at} is a dummy variable that equals one if a conflict began before $t - 12$ and has not yet ended; and Offset_{at} is a dummy variable that equals one if a conflict ended in period t in admin-1 a and if no conflict is reported from $t = +1$ to $t = +12$. In practice, in the most saturated version of this regression, when we include three dummy variables for conflict onset, ongoing conflict, and conflict offset, the omitted category pools time periods with no conflict during the preceding or following year.

In Table C3, we report the estimation results that use the variables for onset, offset, and ongoing conflict to study how monitoring and swarm detection respond to conflict. Accounting for conflict dynamics, we estimate that a conflict event in the recent year imprecisely lowers monitoring likelihood by 3 p.p. (column 1). Saturating the regression to include offset and ongoing dummy variables, we find that conflict onset lowers monitoring likelihood by 3.8 p.p., and conflicts that have been ongoing for over a year suppress monitoring by 4.9 p.p. We find no precisely estimated persistent effect on monitoring after a conflict ends.

Table C2.

Extended Analysis for Locust Monitoring, Detection & Control on Local Conflict

Panel A. Dummies for Non-Zero Reports

	Monitoring			Swarm			Ecology			Control		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
C. Dummy	-0.031**			0.012**			-0.031***			-0.003		
	(0.012)			(0.006)			(0.012)			(0.005)		
C. Share[0,-4]		-0.072***			0.022*			-0.068***			-0.000	
		(0.025)			(0.011)			(0.024)			(0.010)	
C. Share[0,-12]			-0.092***			0.029*			-0.089***			0.005
			(0.032)			(0.015)			(0.032)			(0.015)
Observations	47,230	47,230	47,230	47,230	47,230	47,230	47,230	47,230	47,230	47,230	47,230	47,230
R^2	0.37	0.37	0.37	0.17	0.17	0.17	0.38	0.38	0.38	0.20	0.20	0.20
Mean Dep. Var.	0.28	0.28	0.28	0.04	0.04	0.04	0.25	0.25	0.25	0.05	0.05	0.05

Panel B. Counts of Days With Non-Zero Reports

	Monitoring			Swarm			Ecology			Control		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
C. Dummy	-0.425***			-0.011			-0.380***			-0.080		
	(0.127)			(0.033)			(0.145)			(0.065)		
C. Share[-1,-4]		-0.952***			-0.061			-0.926***			-0.229	
		(0.245)			(0.074)			(0.287)			(0.156)	
C. Share[-1,-12]			-1.273***			-0.043			-1.259***			-0.246
			(0.337)			(0.101)			(0.379)			(0.210)
Observations	47,230	47,230	47,230	47,230	47,230	47,230	47,230	47,230	47,230	47,230	47,230	47,230
R^2	0.40	0.40	0.40	0.15	0.15	0.15	0.39	0.39	0.39	0.18	0.18	0.18
Mean Dep. Var.	2.00	2.00	2.00	0.18	0.18	0.18	1.80	1.80	1.80	0.44	0.44	0.44

Notes: Each column reports the coefficients from one regression, capturing the effects of conflict incidence dummy (C. Dummy), or the share of previous 4 or 12 months with conflict dummy variable outcomes (C. Share) for either monitoring, locust swarm detection, ecology reporting, or locust control operations. In Panel A, outcomes a dummies for any non-zero reports in the admin-1-month level. In Panel B, outcomes are the number of days with non-zero reports in the admin-1-month level. Each regression includes admin-1 fixed effects, as well as breeding area season-by-year-by-month, and specific breeding area-by-month fixed effects. Standard errors are clustered at the admin-1 level.

* 0.10 ** 0.05 *** 0.01

Table C3.
Analysis for Locust Monitoring & Swarm Detection on Local Conflict Onset, Offset & Ongoing Variables

	Monitoring Dummy			Monitoring Days			Swarm Dummy			Swarm Days		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Onset	-0.031	-0.031	-0.038*	-0.274	-0.279	-0.386	-0.003	-0.003	-0.001	-0.054	-0.055	-0.051
	(0.020)	(0.020)	(0.020)	(0.242)	(0.243)	(0.249)	(0.008)	(0.008)	(0.007)	(0.078)	(0.079)	(0.076)
Offset		-0.002	-0.008		-0.389**	-0.484***		-0.002	-0.000		-0.130***	-0.126**
		(0.017)	(0.016)		(0.159)	(0.160)		(0.011)	(0.011)		(0.049)	(0.049)
Ongoing			-0.049***			-0.760***			0.015*			0.029
			(0.018)			(0.246)			(0.009)			(0.047)
Observations	47,230	47,230	47,230	47,230	47,230	47,230	47,230	47,230	47,230	47,230	47,230	47,230
R^2	0.37	0.37	0.37	0.40	0.40	0.40	0.17	0.17	0.17	0.15	0.15	0.15
Mean Dep. Var.	0.28	0.28	0.28	2.00	2.00	2.00	0.04	0.04	0.04	0.18	0.18	0.18

Notes: Each column reports the coefficients from one regression, capturing the effects of conflict onset, offset, or ongoing conflict on monitoring or locust swarm detection (dummies or number of days). See text for the definitions of the onset, offset, and ongoing confict dummy variables. Each regression includes admin-1 fixed effects, as well as breeding area season-by-year-by-month, and specific breeding area-by-month fixed effects. Standard errors are clustered at the admin-1 level.
* 0.10 ** 0.05 *** 0.01

Sensitivity to Precipitation Shock Definition In the main text, we use monthly rainfall above 25 mm to define precipitation shocks that lead to decreased monitoring and increased swarm outbreaks during conflict events. We pool these shocks over a period going from the contemporaneous month to previous four months. Table C4 shows the same model, but with precipitation shocks pooled over a longer time period, extending to the previous twelve months. We choose this cutoff based on regressing an indicator for reported swarms on lagged precipitation shocks and observing that significance drops beyond the twelfth month lag. We observe a similar pattern, although the recovery of monitoring during combined conflict and high precipitation conditions is more pronounced. Tables C5 and C6 show alternative results, when defining precipitation shocks locally rather than using a single absolute threshold across our large region. Instead of monthly rainfall above 25 mm, these models use rainfall events either above the median level or the 75th percentile of precipitation (labeled P_{50} and P_{75} respectively), defined using the *local* climate at the admin-1 level. The patterns observed hold under these alternative definitions, which are correlated but not equivalent (see Figure C3).

Table C4: Estimating Conflict and Weather Impacts on Swarm Outbreaks: 12 month exposure

	Any Monitoring	Control	Swarm
	(1)	(2)	(3)
Conflict × Precip Share[0,-12]	0.177***	0.033	0.069**
	(0.066)	(0.029)	(0.032)
Conflict Share[0,-12]	-0.138***	-0.013	-0.018
	(0.033)	(0.018)	(0.018)
Precip Share[0,-12]	0.278***	0.207***	0.140***
	(0.044)	(0.022)	(0.017)
Observations	47,488	47,488	47,488
R^2	0.374	0.209	0.178
Mean Dep. Var.	0.275	0.052	0.036

Notes: Each column reports the coefficients from one regression, capturing the effects of the share of recent months with both conflict and precipitation shocks (Conflict × Precip Share[0,-12]), share of recent months with conflict (Conflict Share[0,-12]), and share of recent months with precipitation shocks (Precip Share[0,-12] with Precip defined as $\mathbb{1}\{> 25\text{mm}_g\}$, where g is any 0.25° ERA5 grid cell within the admin-1). Each regression includes admin-1 fixed effects, as well as breeding area season-by-year-by-month, and specific breeding area-by-month fixed effects. Standard errors are clustered at the admin-1 level.
* 0.10 ** 0.05 *** 0.01

Table C5: Estimating Conflict and Weather Impacts on Swarm Outbreaks: 50th percentile rainfall

	Any Monitoring	Control	Swarm
	(1)	(2)	(3)
Conflict × P_{50} Share[0,-4]	-0.007	0.025	0.030*
	(0.035)	(0.020)	(0.016)
Conflict Share[0,-4]	-0.049**	-0.012	-0.003
	(0.023)	(0.011)	(0.008)
P_{50} Share[0,-4]	0.171***	0.087***	0.058***
	(0.018)	(0.010)	(0.008)
Observations	48,512	48,512	48,512
R^2	0.373	0.208	0.176
Mean Dep. Var.	0.274	0.052	0.036

Notes: Each column reports the coefficients from one regression, capturing the effects of the share of recent months with both conflict and precipitation shocks (Conflict × P_{50} Share[0,-4]), share of recent months with conflict (Conflict Share[0,-4]), and share of recent months with precipitation shocks (P_{50} Share[0,-4]). Each regression includes admin-1 fixed effects, as well as breeding area season-by-year-by-month, and specific breeding area-by-month fixed effects. Standard errors are clustered at the admin-1 level.
* 0.10 ** 0.05 *** 0.01

Table C6: Estimating Conflict and Weather Impacts on Swarm Outbreaks: 75th percentile rainfall

	Any Monitoring	Control	Swarm
	(1)	(2)	(3)
Conflict × P_{75} Share[0,-4]	-0.019	0.048	0.041*
	(0.043)	(0.032)	(0.023)
Conflict Share[0,-4]	-0.051***	-0.014	-0.000
	(0.019)	(0.010)	(0.008)
P_{75} Share[0,-4]	0.219***	0.115***	0.067***
	(0.021)	(0.015)	(0.010)
Observations	48,512	48,512	48,512
R^2	0.373	0.210	0.176
Mean Dep. Var.	0.274	0.052	0.036

Notes: Each column reports the coefficients from one regression, capturing the effects of the share of recent months with both conflict and precipitation shocks (Conflict × P_{75} Share[0,-4]), share of recent months with conflict (Conflict Share[0,-4]), and share of recent months with precipitation shocks (P_{75} Share[0,-4]). Each regression includes admin-1 fixed effects, as well as breeding area season-by-year-by-month, and specific breeding area-by-month fixed effects. Standard errors are clustered at the admin-1 level.
* 0.10 ** 0.05 *** 0.01

Figure C3: Relationship Between Absolute and Relative Precipitation Shocks

Notes: Correlation at the admin-1 level between monthly precipitation above 25 mm (in any weather cell within the admin-1) and monthly precipitation above the median level, relative to the local admin-1 climate.

C3 Predicted Swarm Time Series Using Reconstructed Teleconnections

In the main text, we describe how we use the LASSO to construct predicted swarm values that link grid cells with breeding areas. In Table C7, we document the results of regressing observed swarms on predicted swarms at the grid-year-month level, controlling for grid cell fixed effects and varying temporal trends. We obtain coefficients that are mechanically very close to 1 and precisely estimated. In Figure C4, we provide a simple descriptive visualization of the correlation between the observed swarm dummies in each grid cell, at the monthly level, and the predicted values from the LASSO procedure. Three stylized facts arise here. First, the correlation between observed and predicted values is positive, and almost linear over much of the support (captured by the local polynomial red line fit). Second, swarms are observed even for very low predicted values. Finally, as the predicted values get close to one, they always correctly predict a truly observed swarm. Even though the correlation is not perfect—some swarms get assigned with low predicted values, and vice-versa—we observe a strong pattern of positive correlation between the two, supporting the utility of predicted values as a proxy for swarm presence.

Table C7: Observed Swarms vs. Predicted Swarms

	Observed Swarm Dummy	
	(1)	(2)
Pred. Swarm Values	0.996***	0.996***
	(0.002)	(0.002)
Observations	106,260	106,260
R^2 (within)	0.338	0.338
Fixed Effects		
Grid cell	✓	✓
Year	✓	
Month	✓	
Year×Month		✓

Notes: Each column reports the coefficients from one regression, regressing observed swarm dummy values at the grid-year-month level on predicted swarm values generated by the lasso model. Fixed effects vary by column. Standard errors are clustered by grid cell.
* 0.10 ** 0.05 *** 0.01

Figure C4: Observed Swarms Versus Predicted Swarms

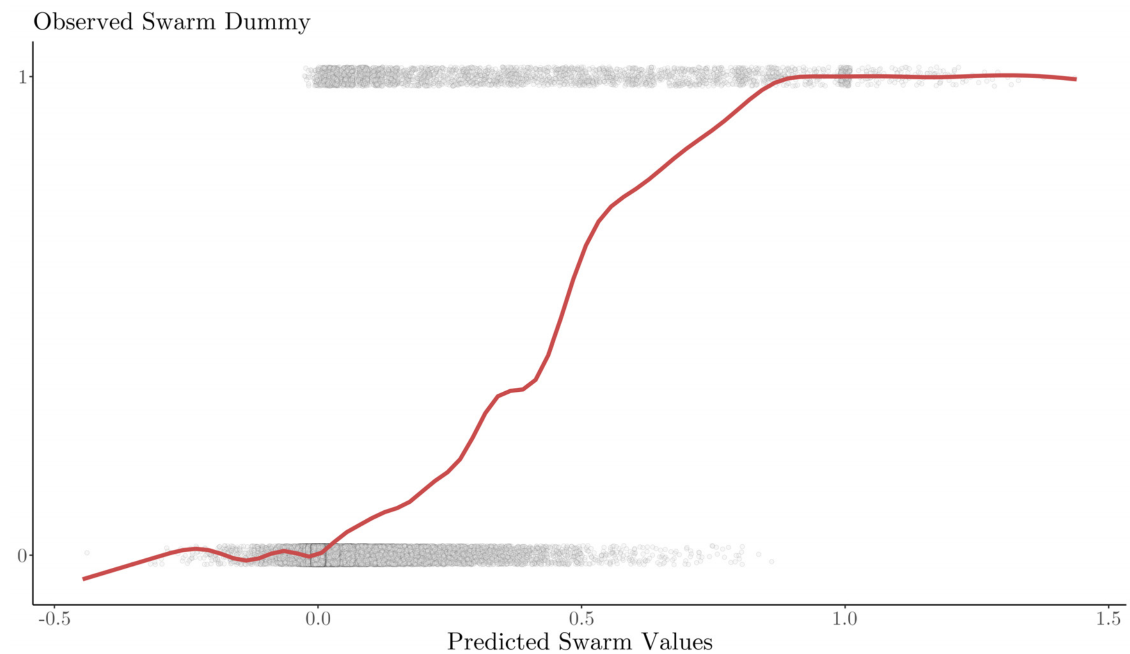

Notes: Correlation between the observed swarm dummy and the predicted swarm value using the LASSO procedure described in Section 5.

C4 Estimation Results Using Predicted Swarm Exposure to Instrument for Swarm Exposure

In Section B3, we demonstrate that the exclusion restriction for an instrumental variables approach does not hold. Simulation results shown that 2SLS results in inflated coefficients. Here, for completeness and transparency, we report results from estimating the human impacts regression when using the predicted swarm proxy to instrument for swarm exposure. The results in Table C8 validate our concern: the coefficients we obtain for standardized height are 3.9 and 6.1 times larger, and the coefficients for the stunting dummy are 3.8 and 5.3 times larger than those we report in Table 4, columns 1 and 3.

Table C8
Instrumented Locust Exposure Around Birth & Children's Height

	Standardized Height	Stunting Dummy
	(1)	(2)
$\widehat{\text{Swarm Exposure, 9m}}$ pre-birth	-1.441***	0.275**
	(0.388)	(0.107)
$\widehat{\text{Swarm Exposure, 9m}}$ post-birth	-1.702***	0.444***
	(0.525)	(0.149)
Observations	409,100	409,100
Mean Dep. Var.	-1.429	0.380

Notes: Estimation results from modifying Equation (3) so that is uses the predicted swarm exposure variables used in (5) to first predict swarm exposure, and use those prediced values in a second-stage instrumental variable regression. Each regression includes DHS cluster fixed effects, birth year, interview year, birth month, and interview month fixed effects, which are allowed to differ across WHO regions. Sample includes all rural DHS clusters. Observations are weighted using DHS sample weights (see main text for more details). Standard errors are clustered at the admin-1 level * 0.10 ** 0.05 *** 0.01

C5 Impact on Food Prices

For locust swarms to have an impact on nutritional status, they must impact access to food, either directly via local food production or indirectly via food distribution systems, which may be imperfect (Sen 1982; Meng et al. 2015). If agricultural production data were available at a sub-national level, we would have been able to test for the relationship between locust outbreaks and food production; unfortunately, such data are not available for most of the countries in our sample. However, we provide support for a food access channel by documenting how locust swarm outbreaks impact the price of staple foods, where local scarcity is typically expected to increase prices (Table C9). Local market price data on staple foods comes from the FEWS NET database (Brown 2008). In Panel A, we use all countries, truncating the sample in 2019 to avoid confounding from the COVID-19 pandemic. We find that swarm exposure drives staple food prices up by seven to nine percent. These increases are precisely estimated and robust to a series of progressively more conservative fixed effects that account for trends across product groups, markets, and time. In Panel B, we show that results are robust to excluding Zimbabwe, which experienced hyperinflation in 2019. Conte et al. (2023) find similar evidence of positive price effects in the case of the 2004 locust plague in Mali. These results suggest that local swarms generate local food scarcity that is imperfectly compensated for by food distribution networks.

Table C9
Swarm Exposure and Staple Food Prices

	Log(price)					
	(1)	(2)	(3)	(4)	(5)	(6)
A: 1992-2019						
Swarm Exposure	0.091***	0.089***	0.076***	0.095***	0.084***	0.084***
	(0.026)	(0.027)	(0.028)	(0.029)	(0.030)	(0.030)
Observations	19,096	19,096	19,095	19,038	19,038	19,036
R^2	0.978	0.986	0.979	0.982	0.989	0.989
B: 1992-2019, excluding Zimbabwe						
Swarm Exposure	0.087***	0.085***	0.071**	0.094***	0.083***	0.083***
	(0.026)	(0.028)	(0.029)	(0.030)	(0.031)	(0.031)
Observations	18,893	18,893	18,892	18,836	18,836	18,834
R^2	0.977	0.985	0.977	0.981	0.988	0.989
Fixed Effects						
Market	✓		✓			
Product Group	✓		✓			
Year	✓	✓				
Month	✓	✓				
Year×Month			✓	✓	✓	✓
Market×Product Group		✓			✓	✓
Market×Month				✓	✓	✓
Product Group×Year				✓	✓	✓
Product Group×Month						✓

Notes: Estimation results from regressing logged staple food price on swarm exposure within 150km of a market. Product group refers to a grouping of similar products (e.g., black beans and pinto beans into the product group beans). Standard errors are clustered at the admin-1 level.
* 0.10 ** 0.05 *** 0.01

C6 Unweighted Sample Results

In the main text, we use DHS sample weights when estimating the treatment effects on height-for-age, and stunting. As Solon et al. (2015) and Anttila-Hughes et al. (2021) discuss, the exact choice of weights is not obvious, and in some studies, scholars prefer to report unweighted results. To allow for easier comparison with unweighted analyses, and to examine to what degree, if any, using the DHS sample weights affects the results, we repeat the main analysis on human impacts (Table 4) in an unweighted sample (Table C10). All results maintain the same sign, statistical significance, and have similar magnitudes relative to the results in the weighted sample (Table 4).

Table C10
Locust Exposure Around Birth & Children's Height (Unweighted Sample)

	Standardized Height		Stunting Dummy	
	(1)	(2)	(3)	(4)
Swarm Exposure, 9m pre-birth	-0.240***		0.039***	
	(0.063)		(0.014)	
Swarm Exposure, 9m post-birth	-0.376***		0.101***	
	(0.060)		(0.013)	
Mean Pr. Val. 9m pre-birth		-0.838***		0.193***
		(0.221)		(0.049)
Mean Pr. Val. 9m post-birth		-1.072***		0.247***
		(0.146)		(0.044)
Observations	409,100	409,100	409,100	409,100
R^2	0.257	0.257	0.202	0.202
Mean Dep. Var.	-1.429	-1.429	0.380	0.380

Notes: Estimation results from Equations (3) and (5). Each regression includes DHS cluster fixed effects, birth year, interview year, birth month, and interview month fixed effects, which are allowed to differ across WHO regions. Sample includes all rural DHS clusters. Observations are unweighted. Standard errors are clustered at the admin-1 level
* 0.10 ** 0.05 *** 0.01

C7 Normalized Effects of Observed and Predicted Locust Exposure

To compare the magnitudes of the estimated effects of exposure to locust swarm using observed and predicted measures of exposure, we normalize each of the exposure measures by its mean and standard deviation. We then replicate the analyses shown in Table 4 using the normalized regressors. The results are shown in Table C11.

Table C11
Locust Exposure Around Birth & Children's Height

	Standardized Height		Stunting Dummy	
	(1)	(2)	(3)	(4)
Swarm Exposure, 9m pre-birth, normalized	-0.039***		0.008***	
	(0.010)		(0.002)	
Swarm Exposure, 9m post-birth, normalized	-0.032***		0.009***	
	(0.005)		(0.002)	
Mean Pr. Val. 9m pre-birth, normalized		-0.062***		0.014***
		(0.014)		(0.004)
Mean Pr. Val. 9m post-birth, normalized		-0.040***		0.011***
		(0.008)		(0.003)
Observations	409,100	409,100	409,100	409,100
R^2	0.266	0.266	0.209	0.210

Notes: Estimation results from Equations (3) and (5). Each regression includes DHS cluster fixed effects, birth year, interview year, birth month, and interview month fixed effects, which are allowed to differ across WHO regions. Sample includes all rural DHS clusters. Observations are weighted using DHS sample weights (see main text for more details). Standard errors are clustered at the admin-1 level.
* 0.10 ** 0.05 *** 0.01

C8 Locust Outbreaks and Food Security

We use data on food emergency declarations, and the reporting of famine conditions from the Emergency Events Database, EM-DAT (Delforge et al. 2025). Our goal is to test how food shortages affect our health outcomes in the main analysis. Food emergencies or famine conditions are declared at the national level, or at the admin-1 level. We use both types in our sample, which spans 1985 to 2020, and includes 0 declarations and 3 famines at the national level, and 3 declarations and 11 famines at the admin-1 level. When a declaration or famine is classified as national, we consider all admin-1s as experiencing a food emergency. The results in Table C12 support the claim that food emergency declarations and/or famine conditions degrade children's outcomes (see discussion in main text Section 6).

Table C12.
Effects of Food Insecurity on Children's Development

	Standardized Height			Stunting Dummy		
	(1)	(2)	(3)	(4)	(5)	(6)
Food Emergency Declaration	-0.33**		-0.14	0.12***		0.07**
	(0.17)		(0.14)	(0.04)		(0.03)
Famine Conditions Reported		-0.34	-0.32		0.08	0.08
		(0.22)	(0.23)		(0.05)	(0.05)
Observations	409,100	409,100	409,100	409,100	409,100	409,100
R^2	0.27	0.27	0.27	0.21	0.21	0.21

Notes: Estimation results from running a regression of height-for-age outcomes on a dummy variable for there being a food emergency declaration, or famine conditions reported, in either the admin-1 region of the DHS cluster, or in at national level if the food insecruity event was at the national level. Each regression includes DHS cluster fixed effects, birth year, interview year, birth month, and interview month fixed effects, which are allowed to differ across WHO regions. Sample includes all rural DHS clusters. Observations are weighted using DHS sample weights (see main text for more details). Standard errors are clustered at the admin-1 level.
* 0.10 ** 0.05 *** 0.01

C9 Alternative specifications

In the main text, we include DHS cluster fixed effects in our baseline specification. Here, we examine how including more granular fixed effects affects our results. In Table C13, we report three variations to the unit fixed effects. We either include admin-1, admin-2, or mother fixed effects. Each subsequent choice of fixed effects nests the one above it (admin-2 nest admin-1, DHS clusters nest admin-2, and mother fixed effects nest DHS clusters). We recover nearly identical results across all three fixed effects versions, and those are nearly identical to the results we report in the main text. In short, the choice of unit fixed effects does not appear to meaningfully affect the results. However, including mother fixed effects does reduce the sample size by 31%, from 409,110 in the main sample to 281,966, as this specification limits the sample to mothers for which we observe at least two separate births.

Table C13
Locust Exposure Around Birth & Children's Height

	Standardized Height		Stunting Dummy	
	(1)	(2)	(3)	(4)
Panel A: Admin-1 Fixed Effects				
Swarm Exposure, 9m pre-birth	-0.350***		0.067***	
	(0.097)		(0.020)	
Swarm Exposure, 9m post-birth	-0.238***		0.075***	
	(0.046)		(0.018)	
Mean Pr. Val. 9m pre-birth		-1.110***		0.235***
		(0.247)		(0.068)
Mean Pr. Val. 9m post-birth		-0.627***		0.175***
		(0.169)		(0.052)
Observations	409,618	409,618	409,618	409,618
R^2	0.102	0.102	0.069	0.069
Panel B: Admin-2 Fixed Effects				
Swarm Exposure, 9m pre-birth	-0.339***		0.065***	
	(0.090)		(0.018)	
Swarm Exposure, 9m post-birth	-0.232***		0.075***	
	(0.042)		(0.016)	
Mean Pr. Val. 9m pre-birth		-1.194***		0.260***
		(0.267)		(0.075)
Mean Pr. Val. 9m post-birth		-0.687***		0.193***
		(0.161)		(0.052)
Observations	409,501	409,501	409,501	409,501
R^2	0.125	0.125	0.087	0.087
Panel C: Mother Fixed Effects				
Swarm Exposure, 9m pre-birth	-0.317***		0.069***	
	(0.081)		(0.021)	
Swarm Exposure, 9m post-birth	-0.213***		0.076***	
	(0.057)		(0.022)	
Mean Pr. Val. 9m pre-birth		-1.350***		0.312***
		(0.289)		(0.078)
Mean Pr. Val. 9m post-birth		-0.536***		0.200***
		(0.187)		(0.054)
Observations	281,966	281,966	281,966	281,966
R^2	0.499	0.499	0.473	0.474

Notes: Estimation results from Equations (3) and (5). Each regression includes admin-1 (panel A), admin-2 (panel B), and mother (panel C) fixed effects, birth year, interview year, birth month, and interview month fixed effects, which are allowed to differ across WHO regions. Sample includes all rural DHS clusters. Observations are weighted using DHS sample weights (see main text for more details). Standard errors are clustered at the admin-1 level.
* 0.10 ** 0.05 *** 0.01

C10 Accounting for Locust Control Operations

In the main text, we discuss how control operations—most likely, insecticide spraying—could affect our results: the effect on child health could potentially be driven by exposure to insecticides, rather than degraded food security conditions. Although this would still be a consequence of locust swarm activity, it might suggest very different policy conclusions.

Here we test whether accounting for exposure to control operations affects our main estimates. Mirroring our treatment of the swarm exposure dummy, we define a control exposure dummy which is equal to one if we observe a control operation point in the FAO data within 150 km of the DHS cluster in a given year-month combination, and zero otherwise.

We do not find evidence to support the possibility that the negative health impacts we observe could be attributed to control operations. Table C14 reports the estimated health impacts of observed swarm exposure (columns 1 and 3) and predicted swarm exposure (columns 2 and 4), from specifications that include dummy variables for control exposure in the months before and after birth. The estimated coefficients on the control exposure dummies are not interpretable, and we do not report them: they potentially capture some of the direct effects of locusts as well as any effect of control, either because control operations only take place when locust infestations are sufficiently severe, or in columns 2 and 4, because they capture some of the realized variation in locust exposure that is not captured by predicted exposure. Importantly, however, the main estimates are almost entirely unchanged when we account for control operations.

One might ask how definitive this exercise is given that locust control operations could also be mismeasured. However, we note that for many locust control operations, field staff must be present to implement the control actions. We thus expect locust control records to suffer from lower levels of mismeasurement than the swarm observations.

Table C14

Locust Exposure Around Birth & Children's Height: Accounting for control operations

	Standardized Height		Stunting Dummy	
	(1)	(2)	(3)	(4)
Swarm Exposure, 9m pre-birth	-0.327***		0.071***	
	(0.107)		(0.021)	
Swarm Exposure, 9m post-birth	-0.320***		0.099***	
	(0.051)		(0.017)	
Mean Pr. Val. 9m pre-birth		-1.348***		0.301***
		(0.305)		(0.085)
Mean Pr. Val. 9m post-birth		-0.983***		0.269***
		(0.191)		(0.061)
Observations	409,100	409,100	409,100	409,100
R^2	0.266	0.266	0.209	0.210

Notes: Estimation results from Equations (3) and (5). Each regression includes DHS cluster fixed effects, birth year, interview year, birth month, and interview month fixed effects, which are allowed to differ across WHO regions. Sample includes all rural DHS clusters. Observations are weighted using DHS sample weights (see main text for more details). Standard errors are clustered at the admin-1 level.
* 0.10 ** 0.05 *** 0.01

C11 Strengthening the Identifying Assumption

The LASSO models that link locust swarm activity in recipient 5-degree grid cells with locust swarm activity in breeding areas do not impose any restriction on the breeding areas that each grid cell can draw from. While there is spatial separation between the grid cells that DHS impact exposure zones fall into and the breeding areas that they draw from, there is also some overlap (Figure 8). To rule out the possibility that results are driven by DHS impact exposure zones where the LASSO-based proxy might draw from locust activity in the overlapping breeding area, i.e., draw from local locust activity rather than from locust activity further away (which would threaten the identifying assumption), we create a proxy for local locust swarm exposure that cannot draw from local locust activity. Specifically, if a recipient 5-degree grid cell overlaps with a breeding area, then locust activity in this breeding area is dropped from the list of potential predictors that the LASSO model can choose from.

When using this more constrained proxy, we find that predictive power drops somewhat but remains very high both at the monthly level (Table C15) and when predicting exposure among children in the DHS sample (Table C16). Finally, the predicted locust swarm results are virtually unchanged for the impact of exposure in the nine months before birth, but magnitude is reduced and we lose statistical significance for the impact of exposure in the nine months after birth, suggesting that the pre-birth effects are more robust to this alternative approach to estimation (Table C17).

Excluding grid cells that receive no LASSO prediction To construct predicted exposure to swarms, we use LASSO to select, for each 5°x5° grid cell, the lagged breeding swarm dummies from breeding areas that are most predictive of locust swarm activity in the destination grid cell. For some grid cells, no predictors are selected. In that case, we set predicted exposure to swarms to zero for DHS observations that fall in that grid cell (observed swarm values are virtually always zero in those locations). In the main text, we report results that reflect this filling process. In Table C18, we report estimation results where instead of filling in missing predictions, we drop the corresponding observations from the data. The sample size for the predicted exposure regressions decreases by 17% but results are very similar to the main ones that used filled predictions.

Table C15: Observed Swarms vs. Predicted Swarms, Excluding Overlapping Breeding Areas

	Observed Swarm Dummy	
	(1)	(2)
Pred. Swarm Values	0.987***	0.986***
	(0.002)	(0.003)
Observations	106,260	106,260
R^2 (within)	0.237	0.237
Fixed Effects		
Grid cell	✓	✓
Year	✓	
Month	✓	
Year×Month		✓

Notes: Each column reports the coefficients from one regression, regressing observed swarm dummy values at the grid-year-month level on predicted swarm values generated by the lasso model. Fixed effects vary by column. Standard errors are clustered by grid cell.
* 0.10 ** 0.05 *** 0.01

Table C16.
Predicting Locust Exposure Around Birth, Excluding Overlapping Breeding Areas

	Grid Cell Fixed Effects		DHS Fixed Effects	
	9m pre-birth	9m post-birth	9m pre-birth	9m post-birth
$\overline{\text{9m }\widehat{\text{pre-birth}}}$	0.509***	0.196***	0.494***	0.186***
	(0.086)	(0.068)	(0.089)	(0.062)
$\overline{\text{9m }\widehat{\text{post-birth}}}$	-0.129***	0.661***	-0.181***	0.644***
	(0.041)	(0.117)	(0.041)	(0.117)
Observations	409,618	409,618	409,100	409,100
SW F Stat.	34.7	68.2	36.2	62.3

Notes: Estimation results from Equations (4). Columns report first stage results. Each regression includes grid cell (left) or DHS cluster (right) fixed effects, birth year, interview year, birth month, and interview month fixed effects, which are allowed to differ across WHO regions. Sample includes all rural DHS clusters. Observations are weighted using DHS sample weights (see main text for more details). Standard errors are clustered at the admin-1 level.
* 0.10 ** 0.05 *** 0.01

Table C17

Locust Exposure Around Birth & Children's Height, Excluding Overlapping Breeding Areas

	Standardized Height		Stunting Dummy	
	(1)	(2)	(3)	(4)
Swarm Exposure, 9m pre-birth	-0.369***		0.072***	
	(0.097)		(0.021)	
Swarm Exposure, 9m post-birth	-0.280***		0.084***	
	(0.041)		(0.015)	
Mean Pr. Val. 9m pre-birth		-1.260***		0.289***
		(0.364)		(0.090)
Mean Pr. Val. 9m post-birth		-0.334*		0.071
		(0.202)		(0.059)
Observations	409,100	409,100	409,100	409,100
R^2	0.266	0.265	0.209	0.209

Notes: Estimation results from Equations (3) and (5). Each regression includes DHS cluster fixed effects, birth year, interview year, birth month, and interview month fixed effects, which are allowed to differ across WHO regions. Sample includes all rural DHS clusters. Observations are weighted using DHS sample weights (see main text for more details). Standard errors are clustered at the admin-1 level.

* 0.10 ** 0.05 *** 0.01

Table C18
Locust Exposure Around Birth & Children's Height, Unfilled LASSO predictions

	Standardized Height		Stunting Dummy	
	(1)	(2)	(3)	(4)
Swarm Exposure, 9m pre-birth	-0.369***		0.072***	
	(0.097)		(0.021)	
Swarm Exposure, 9m post-birth	-0.280***		0.084***	
	(0.041)		(0.015)	
Mean Pr. Val. 9m pre-birth		-1.503***		0.310***
		(0.296)		(0.088)
Mean Pr. Val. 9m post-birth		-0.904***		0.234***
		(0.190)		(0.062)
Observations	409,100	339,402	409,100	339,402
R^2	0.266	0.256	0.209	0.202
Mean Dep. Var.	-1.512	-1.595	0.402	0.427

Notes: Estimation results from Equations (3) and (5). Each regression includes DHS cluster fixed effects, birth year, interview year, birth month, and interview month fixed effects, which are allowed to differ across WHO regions. Sample includes all rural DHS clusters. Observations are weighted using DHS sample weights (see main text for more details). Standard errors are clustered at the admin-1 level.
* 0.10 ** 0.05 *** 0.01

C12 Examining Sensitivity to Distance Used to Define Treatment

In the main text, we discuss how the data-driven choice of the 150 km threshold to define swarm exposure might result in SUTVA violations. It is plausible that DHS clusters for which we observe a swarm more than 150 km away are still experiencing some negative impacts from that swarm, especially if that remote swarm is indicative of an undetected or unreported swarm closer to the DHS cluster (a false negative). In our assignment of swarm exposure, we would treat the DHS cluster as part of the control group, which could potentially attenuate our estimates. Here we demonstrate that we detect effects even when accounting for swarm exposures that extend beyond 150 km.

We estimate a similar specification to the one in Equation (3), but including mutually exclusive dummies for swarm exposure in 100 km distance bins. Specifically, we include four dummies for swarm exposure in the nine months pre-birth, and four dummies for swarm exposure in the nine months post-birth. Each dummy is equal to one if we observe a swarm in the 0-100, 100-200, 200-300, and 300-400 km ranges in each period. We report the estimation results in Figure C5. We estimate the proximity to observed swarms results in higher impacts on height-for-age and the likelihood of stunting, but the we still recover coefficients that suggest that exposure even up to 400 km might have an negative impact.

Figure C5: Locust Exposure Around Birth & Children's Height: Exposure Distance Bins

(a) Standardized Height

(b) Stunting Dummy

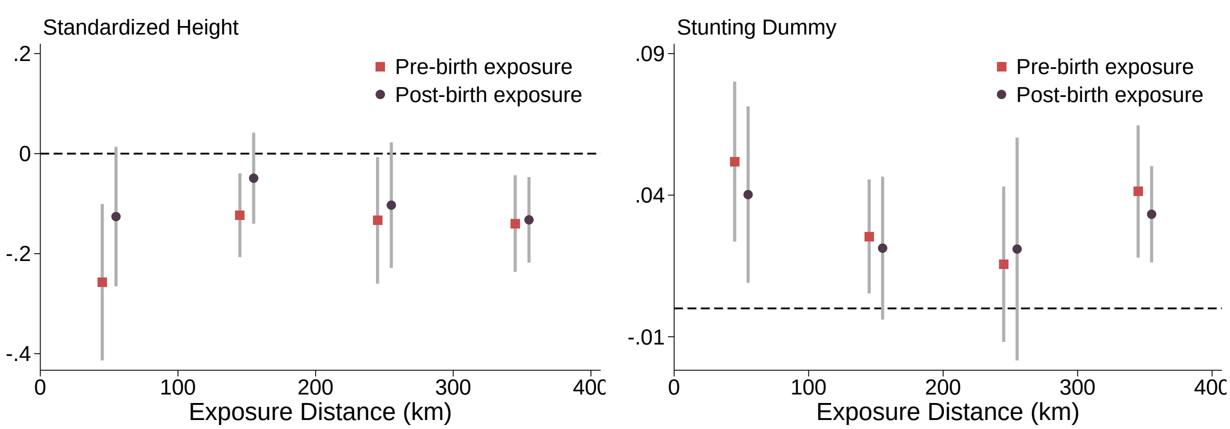

Notes: Estimation results from Equation (3). Each regression includes DHS cluster fixed effects, birth year, interview year, birth month, and interview month fixed effects, which are allowed to differ across WHO regions. Sample includes observations from rural households. Observations are weighted using DHS sample weights (see main text for details). Standard errors are clustered at the admin-1 level.

To further address how this might affect attenuation in the main sample and baseline specification, we re-estimate our model while excluding DHS clusters that are within 150 to 300 km or 150 to 400 km. In Table C19, we report results for both the swarm dummies and the predicted swarm values versions of the result. Broadly, we recover similar estimates compared to the sample where we do not exclude these "donut" DHS clusters.

Table C19
Locust Exposure Around Birth & Children's Height

	Standardized Height		Stunting Dummy	
	(1)	(2)	(3)	(4)
Panel A: Excluding 150-300km Exposures				
Swarm Exposure, 9m pre-birth	-0.398**		0.057	
	(0.173)		(0.035)	
Swarm Exposure, 9m post-birth	-0.287***		0.106***	
	(0.096)		(0.020)	
Mean Pr. Val. 9m pre-birth		-1.176***		0.202**
		(0.390)		(0.100)
Mean Pr. Val. 9m post-birth		-1.354***		0.359***
		(0.209)		(0.070)
Observations	382,057	382,057	382,057	382,057
R^2	0.269	0.269	0.213	0.213
Panel B: Excluding 150-400km Exposures				
Swarm Exposure, 9m pre-birth	-0.443**		0.041	
	(0.207)		(0.037)	
Swarm Exposure, 9m post-birth	-0.302***		0.113***	
	(0.090)		(0.025)	
Mean Pr. Val. 9m pre-birth		-0.932**		0.092
		(0.434)		(0.122)
Mean Pr. Val. 9m post-birth		-1.434***		0.372***
		(0.277)		(0.088)
Observations	370,796	370,796	370,796	370,796
R^2	0.270	0.271	0.214	0.214

Notes: Estimation results from Equations (3) and (5). Each regression includes DHS cluster fixed effects, birth year, interview year, birth month, and interview month fixed effects, which are allowed to differ across WHO regions. Sample includes all rural observations but excludes all observations exposed to a swarm between 150 and 300 km (panel A) and 150 and 400 km (panel B). Observations are weighted using DHS sample weights (see main text for more details). Standard errors are clustered at the admin-1 level.
* 0.10 ** 0.05 *** 0.01

C13 Subsample Estimates and Heterogeneity

Excluding Movers In the main text, we include the data from all the interviews in each rural DHS cluster, even for those who answer they have moved or migrated at some point. Because migration is a potential endogenous response to swarm exposure, we report results that exclude any records where the mother has answered she migrated to the surveyed village. Despite shrinking the sample size by 73%, from 409,110 to 109,297 records, we estimate similar results to those in the baseline sample, which we report in Table C20. Exposure to swarms lowers HAZ and increases the stunting likelihood. The pre-birth coefficients tend to be somewhat smaller in absolute values and the post-birth coefficients somewhat larger (with the exception of the impact of observed post-birth exposure on height), yet we cannot reject that they are the same at the 95% confidence level.

Table C20
Locust Exposure Around Birth & Children's Height: Excluding Movers

	Standardized Height		Stunting Dummy	
	(1)	(2)	(3)	(4)
Swarm Exposure, 9m pre-birth	-0.290***		0.049***	
	(0.092)		(0.018)	
Swarm Exposure, 9m post-birth	-0.232***		0.094***	
	(0.071)		(0.022)	
Mean Pr. Val. 9m pre-birth		-1.329***		0.309***
		(0.284)		(0.071)
Mean Pr. Val. 9m post-birth		-1.413***		0.399***
		(0.274)		(0.080)
Observations	109,297	109,297	109,297	109,297
R^2	0.354	0.355	0.308	0.310

Notes: Estimation results from Equations (3) and (5). Each regression includes DHS cluster fixed effects, birth year, interview year, birth month, and interview month fixed effects, which are allowed to differ across WHO regions. Sample includes observations from households in which the mother has never moved. Observations are weighted using DHS sample weights (see main text for more details). Standard errors are clustered at the admin-1 level. * 0.10 ** 0.05 *** 0.01

Excluding Nearby Conflicts One empirical concern might be that the conflicts interrupting monitoring may be themselves affecting child health directly. In fact, Wagner et al. (2018) use DHS and geolocated conflict data to show that conflicts increase infant mortality risk for up to 100 km away. In order to make sure our results are not driven by the direct effect conflicts on child health, Table C21 reports the same specification as the baseline result, but excluding observations exposed to conflict within 100 km either 5 years before the child's birth (Panel A) or before the DHS interview (Panel B). We recover coefficients that are very similar to the main impact estimates.

Table C21
Locust Exposure Around Birth & Children's Height

	Standardized Height		Stunting Dummy	
	(1)	(2)	(3)	(4)
Panel A: Excluding Observations with Nearby Conflicts at Birth				
Swarm Exposure, 9m pre-birth	-0.363***		0.072***	
	(0.085)		(0.019)	
Swarm Exposure, 9m post-birth	-0.292***		0.091***	
	(0.047)		(0.017)	
Mean Pr. Val. 9m pre-birth		-1.434***		0.305***
		(0.286)		(0.081)
Mean Pr. Val. 9m post-birth		-0.924***		0.264***
		(0.183)		(0.064)
Observations	389,321	389,321	389,321	389,321
R^2	0.267	0.267	0.210	0.211
Panel B: Excluding Observations with Nearby Conflicts at Interview				
Swarm Exposure, 9m pre-birth	-0.372***		0.073***	
	(0.090)		(0.019)	
Swarm Exposure, 9m post-birth	-0.305***		0.094***	
	(0.049)		(0.018)	
Mean Pr. Val. 9m pre-birth		-1.371***		0.300***
		(0.315)		(0.085)
Mean Pr. Val. 9m post-birth		-0.941***		0.261***
		(0.183)		(0.066)
Observations	390,855	390,855	390,855	390,855
R^2	0.265	0.265	0.209	0.209

Notes: Estimation results from Equations (3) and (5). Each regression includes DHS cluster fixed effects, birth year, interview year, birth month, and interview month fixed effects, which are allowed to differ across WHO regions. Sample excludes births with a conflict occurring within 100 km and within 5 years before birth (panel A) or within 5 years before interview (panel B). Observations are weighted using DHS sample weights (see main text for more details). Standard errors are clustered at the admin-1 level.
* 0.10 ** 0.05 *** 0.01

Africa In the main text, we pool all the DHS countries with swarm exposure across Africa and Asia. India, in particular, contributes a large number of child observations in the sample, making up 44% of the pooled sample. We test for the robustness of our results when limiting our analysis to Africa alone (Table C22). Estimates for Africa alone are very similar to those we report in the main text.

Table C22
Locust Exposure Around Birth & Children's Height: African countries

	Standardized Height		Stunting Dummy	
	(1)	(2)	(3)	(4)
Swarm Exposure, 9m pre-birth	-0.430***		0.084***	
	(0.080)		(0.016)	
Swarm Exposure, 9m post-birth	-0.303***		0.095***	
	(0.038)		(0.009)	
Mean Pr. Val. 9m pre-birth		-1.966***		0.438***
		(0.306)		(0.096)
Mean Pr. Val. 9m post-birth		-1.161***		0.342***
		(0.231)		(0.066)
Observations	180,155	180,155	180,155	180,155
R^2	0.238	0.240	0.189	0.191

Notes: Estimation results from Equations (3) and (5). Each regression includes DHS cluster fixed effects, birth year, interview year, birth month, and interview month fixed effects, which are allowed to differ across WHO regions. Sample includes observations from households in Africa. Observations are weighted using DHS sample weights (see main text for more details). Standard errors are clustered at the admin-1 level.
* 0.10 ** 0.05 *** 0.01

Urban Areas In the main text, we report results for the sample of children in DHS clusters that are classified as rural—as we expect the majority of locust swarm exposure to be concentrated in rural and not urban areas. Here we report results that also use data from the urban DHS sample. We first report results that simply extend the analysis to include children living in urban DHS clusters (Table C23). As expected, the estimated effects attenuate slightly, consistent with children in urban areas being less directly affected by exposure to locust swarms. We then estimate two models with interactions between locust exposure and a dummy for living in a rural cluster: the first (Table C24) includes the same set of fixed effects and controls as the pooled analysis reported in Table C23, the second (Table C25) allows trends to vary heterogeneously across rural and urban areas. For both specifications, effects in urban areas are consistently smaller than in rural areas, although we statistically reject the null that effects are equal in rural and urban areas for only a subset of combinations of outcomes, measures of locust exposure, and outcome variables.

Table C23
Locust Exposure Around Birth & Children's Height: Pooled Urban and Rural

	Standardized Height		Stunting Dummy	
	(1)	(2)	(3)	(4)
Swarm Exposure, 9m pre-birth	-0.329***		0.063***	
	(0.092)		(0.019)	
Swarm Exposure, 9m post-birth	-0.268***		0.084***	
	(0.054)		(0.016)	
Mean Pr. Val. 9m pre-birth		-1.367***		0.307***
		(0.298)		(0.075)
Mean Pr. Val. 9m post-birth		-0.972***		0.250***
		(0.164)		(0.056)
Observations	574,164	574,164	574,164	574,164
R^2	0.281	0.281	0.221	0.222
Mean Dep. Var.	-1.390	-1.390	0.371	0.371

Notes: Estimation results from Equations (3) and (5). Each regression includes DHS cluster fixed effects, birth year, interview year, birth month, and interview month fixed effects, which are allowed to differ across WHO regions. Sample includes all DHS clusters. Observations are weighted using DHS sample weights (see main text for more details). Standard errors are clustered at the admin-1 level.
* 0.10 ** 0.05 *** 0.01

Table C24
Locust Exposure Around Birth & Children's Height: Rural Interaction

	Standardized Height		Stunting Dummy	
	(1)	(2)	(3)	(4)
Swarm Exp., 9m pre-birth	-0.194***		0.039**	
	(0.066)		(0.016)	
Swarm Exp., 9m pre-birth × Rural	-0.170*		0.031	
	(0.102)		(0.026)	
Swarm Exp., 9m post-birth	-0.164*		0.077***	
	(0.092)		(0.021)	
Swarm Exp., 9m post-birth × Rural	-0.130		0.008	
	(0.087)		(0.023)	
Mean Pr. Val. 9m pre-birth		-0.987***		0.196***
		(0.243)		(0.061)
Mean Pr. Val. 9m pre-birth × Rural		-0.472		0.137
		(0.334)		(0.102)
Mean Pr. Val. 9m post-birth		-0.644**		0.098
		(0.252)		(0.063)
Mean Pr. Val. 9m post-birth × Rural		-0.410		0.191**
		(0.306)		(0.078)
Observations	574,164	574,164	574,164	574,164
R^2	0.281	0.281	0.221	0.222

Notes: Estimation results from Equations (3) and (5). Each regression includes DHS cluster fixed effects, birth year, interview year, birth month, and interview month fixed effects, which are allowed to differ across WHO regions. Sample includes all DHS clusters. Observations are weighted using DHS sample weights (see main text for more details). Standard errors are clustered at the admin-1 level.
* 0.10 ** 0.05 *** 0.01

Table C25
Locust Exposure Around Birth & Children's Height: Rural Interaction (including Time Trends x Rural)

	Standardized Height		Stunting Dummy	
	(1)	(2)	(3)	(4)
Swarm Exp., 9m pre-birth	-0.145**		0.023	
	(0.069)		(0.020)	
Swarm Exp., 9m pre-birth × Rural	-0.221*		0.049	
	(0.113)		(0.033)	
Swarm Exp., 9m post-birth	-0.119		0.061**	
	(0.109)		(0.027)	
Swarm Exp., 9m post-birth × Rural	-0.168*		0.025	
	(0.100)		(0.030)	
Mean Pr. Val. 9m pre-birth		-0.938***		0.235***
		(0.256)		(0.074)
Mean Pr. Val. 9m pre-birth × Rural		-0.517		0.089
		(0.360)		(0.132)
Mean Pr. Val. 9m post-birth		-0.675***		0.129**
		(0.253)		(0.062)
Mean Pr. Val. 9m post-birth × Rural		-0.337		0.149**
		(0.287)		(0.075)
Observations	574,164	574,164	574,164	574,164
R^2	0.282	0.282	0.223	0.223

Notes: Estimation results from Equations (3) and (5). Each regression includes DHS cluster fixed effects, birth year, interview year, birth month, and interview month fixed effects, which are allowed to differ across WHO regions and rural residence. Sample includes all DHS clusters. Observations are weighted using DHS sample weights (see main text for more details). Standard errors are clustered at the admin-1 level.
* 0.10 ** 0.05 *** 0.01

Sex-Specific Effects In Table C26, we report estimation results that include interactions with a female dummy variable. We find that there is meaningful heterogeneity by sex-at-birth—male births experience a larger scarring penalty relative to female births. This might suggest that female fetuses are experiencing a larger negative selection, resulting in higher rates of fetal deaths.

Table C26
Locust Exposure Around Birth & Children's Height

	Standardized Height		Stunting Dummy	
	(1)	(2)	(3)	(4)
Swarm Exp., 9m pre-birth	-0.438***		0.093***	
	(0.108)		(0.023)	
Swarm Exp., 9m pre-birth × Female	0.142***		-0.043***	
	(0.048)		(0.013)	
Swarm Exp., 9m post-birth	-0.316***		0.096***	
	(0.045)		(0.015)	
Swarm Exp., 9m post-birth × Female	0.072*		-0.024**	
	(0.042)		(0.010)	
Mean Pr. Val. 9m pre-birth		-1.723***		0.409***
		(0.383)		(0.100)
Mean Pr. Val. 9m pre-birth × Female		0.573***		-0.193***
		(0.187)		(0.040)
Mean Pr. Val. 9m post-birth		-1.026***		0.272***
		(0.203)		(0.063)
Mean Pr. Val. 9m post-birth × Female		0.139		-0.013
		(0.106)		(0.027)
Observations	409,100	409,100	409,100	409,100
R^2	0.266	0.266	0.210	0.210

Notes: Estimation results from Equations (3) and (5). Each regression includes DHS cluster fixed effects, birth year, interview year, birth month, and interview month fixed effects, which are allowed to differ across WHO regions. Sample includes all rural DHS clusters. Observations are weighted using DHS sample weights (see main text for more details). Standard errors are clustered at the admin-1 level.
* 0.10 ** 0.05 *** 0.01

C14 Selection Effects: Infant and Child Mortality

The main results measure the scarring effect of locust infestations on infants (reduced height-for-age). Here, we turn to selection effects (see Section 6.4), namely infant and child mortality.

In Table C27, we report estimates for infant mortality, the probability of dying within the first year of life (columns 3 and 4), and mortality of children under five (columns 5 and 6) as the outcome using our main regression specification Equation (3), together with our main-text results on height-for-age (columns 1 and 2) for comparison. Pre-birth exposure significantly increases both infant and under-five mortality. When using observed exposure, infant mortality increases by 0.7 percentage points for children exposed before birth, contributing to a total increase in under-five mortality of 1.1 percentage points. Post-birth exposure does not significantly affect infant or under-five mortality. When using predicted locust exposure, pre-birth impacts increase in magnitude but are no longer statistically significant for infant mortality, while post-birth impacts strengthen in both magnitude and statistical significance. In Table C28, we repeat the estimation by using standardized observed or predicted swarm variables (as we do in Table C11). Effects based on predicted locust exposure are similar in magnitude or larger than those based on observed exposure.

Table C27.
Locust Exposure Around Birth & Infant or Under-5 Mortality

	Standardized Height		Inf. Mortality		U-5 Mortality	
	(1)	(2)	(3)	(4)	(5)	(6)
Swarm Exposure, 9m pre-birth	-0.369***		0.007*		0.011**	
	(0.097)		(0.004)		(0.005)	
Swarm Exposure, 9m post-birth	-0.280***		-0.001		0.005	
	(0.041)		(0.004)		(0.004)	
Mean Pr. Val. 9m pre-birth		-1.446***		0.011		0.026*
		(0.324)		(0.010)		(0.014)
Mean Pr. Val. 9m post-birth		-0.961***		0.019*		0.042***
		(0.187)		(0.011)		(0.016)
Mean Dep. Var.	-1.512	-1.512	0.057	0.057	0.068	0.068
Observations	409,100	409,100	602,365	602,146	602,365	602,146
R^2	0.266	0.266	0.088	0.088	0.103	0.103

Notes: Estimation results from Equation (3). Columns (1) and (2) include results for height-for-age, columns (3) and (4) for infant mortality dummy, and columns (5) and (6) for under-5 children mortality dummy. Each regression includes DHS cluster fixed effects, birth year, interview year, birth month, and interview month fixed effects, which are allowed to differ across WHO regions. Sample consists of all rural DHS clusters. Observations are weighted using DHS sample weights (see main text for more details). Standard errors are clustered at the admin-1 level.
* 0.10 ** 0.05 *** 0.01

Table C28.
Standardized Locust Exposure Around Birth & Infant or Under-5 Mortality

	Standardized Height		Inf. Mortality		U-5 Mortality	
	(1)	(2)	(3)	(4)	(5)	(6)
Swarm Exposure, 9m pre-birth	-0.0660***		0.0013*		0.0019**	
	(0.0174)		(0.0007)		(0.0008)	
Swarm Exposure, 9m post-birth	-0.0511***		-0.0002		0.0008	
	(0.0075)		(0.0007)		(0.0008)	
Mean Pr. Val. 9m pre-birth		-0.1102***		0.0008		0.0020*
		(0.0247)		(0.0008)		(0.0010)
Mean Pr. Val. 9m post-birth		-0.0726***		0.0014*		0.0032***
		(0.0141)		(0.0008)		(0.0012)
Mean Dep. Var.	-1.512	-1.512	0.057	0.057	0.068	0.068
Observations	409,100	409,100	602,365	602,146	602,365	602,146
R^2	0.266	0.266	0.088	0.088	0.103	0.103

Notes: Same as in Table C27, but using z-scores for the observed or predicted swarm variables, as in Table C11.
* 0.10 ** 0.05 *** 0.01

C15 Selection Effect: Birth Rates

In the event of a fetal death—or a child who is not conceived—neither height-for-age nor child mortality are reported and that selection further creates bias (see main text Section 6.4). We explore this channel by quantifying the effect of swarm exposure on the probability that a woman of childbearing age (15-49) gave birth in any given year. Explicitly, we estimate how y_{idt}—a binary variable equal to one if woman i observed in DHS cluster d, gave birth to a living child in a given year t, zero otherwise—responds to locust swarm exposure using the following regression specification:

$$y_{it} = \beta_1 \text{Swarm}_{d,t-1} + \beta_2 \text{Swarm}_{d,t} + f(\text{weather}_{dt}) + \lambda_i + \omega_{rt} + \theta_{ry} + \varepsilon_{it} \tag{6}$$

on the women sample from the DHS, restricted to rural clusters. We reconstruct each woman's birth history (between 15 and 49) to obtain a panel at the woman-year level. In Equation (6), we consider the effect of both the contemporaneous ($\text{Swarm}_{d,t}$) and the lagged ($\text{Swarm}_{d,t-1}$) swarm exposure variables (swarm presence within 150 km of mother's DHS cluster) to account for the fact that potential pregnancies may start in the calendar year before the potential birth, as well as the contemporaneous calendar year. We account for any time-invariant cross-sectional variation in locust exposure or outcomes by including fixed effects at the woman level λ_i, and alternatively, at the DHS cluster level, η_d. We flexibly control for time trends with ω_{rt} (WHO region by calendar year) and θ_{ry} (WHO region by survey year) fixed effects. We also account for local weather conditions by including $f(\text{weather}_{dt})$, which flexibly controls for local temperature and precipitation.[22] Any

[22] Similarly to the main analyses, these include quartiles of precipitation and temperature bins. We define annual

unmodeled heterogeneity is captured by the error term, ε_{it}. We cluster our standard errors at the admin-1 level.

Overall, the analysis supports the hypothesis of a selection effect *in-utero*. In Table C29, we find that being exposed to a swarm decreases by 0.5 percentage points the probability of giving birth in the year of exposure (in column 1, with mother fixed effects, and column 3 with DHS cluster fixed effects, alike). This decline in birth probability due to swarm exposure reflects a 2.6 percent reduction relative to the mean baseline level of 19.2 percent. Lagged swarm exposure decreases the probability of giving birth by 0.9 percentage points (columns 1 and 3).

The results that use the predicted proxy for swarm exposure also recover a negative effect on the probability of giving birth, and are precisely estimated—allowing us to reject the null hypothesis of no effect on the probability of giving birth at the one percent level for both contemporaneous and lagged predicted exposure. To better compare the results between the observed and predicted swarm exposure, we standardize the variables to obtain z-scores—as we do with the analysis we report in Table C11. In Table C30, we report the results from using the standardized regressors. We find that when using the standardized predicted swarm proxy, we show that the effects are approximately twice as large as when we use the standardized observed swarm exposure (columns 2 and 4 relative to columns 1 and 3).

Table C29.
Observed and Predicted Swarm Exposure & Probability of Giving Birth

	Mother Fixed Effects		DHS Cluster Fixed Effects	
	Observed	Predicted	Observed	Predicted
	(1)	(2)	(3)	(4)
Exposure at t	-0.005*	-0.026***	-0.004*	-0.024***
	(0.002)	(0.006)	(0.002)	(0.005)
Exposure at $t-1$	-0.009***	-0.026***	-0.009***	-0.025***
	(0.003)	(0.008)	(0.003)	(0.008)
R^2	0.089	0.089	0.041	0.041
Mean Dep. Var.	0.192	0.192	0.192	0.192
Observations	9,664,857	9,664,857	9,665,755	9,665,755

Notes: Observed swarm exposure is a dummy variable, while predicted exposure is a continuous measure. Columns 1 and 3 use observed swarm exposure (binary) at t and $t-1$; Columns 2 and 4 use LASSO-predicted swarm exposure (continuous) at t and $t-1$. All regressions include WHO Region × Year and WHO Region × Interview Year fixed effects. Standard errors are clustered at the Admin-1 region level.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

precipitation quartiles globally and temperature-day bins as follows (°Celsius): <11, 11-14, 14-17, 17-20, 20-23, 23-26, 26-29, 29-32, 32-35, >35, using 17-20 as the reference bin. We include 1-period lags and leads of precipitation quartiles and temperature-day bins.

Table C30.
Standardized Observed and Predicted Swarm Exposure & Probability of Giving Birth

	Mother Fixed Effects		DHS Cluster Fixed Effects	
	Observed	Predicted	Observed	Predicted
	(1)	(2)	(3)	(4)
Exposure at t	-0.001*	-0.002***	-0.001*	-0.002***
	(0.00040)	(0.00037)	(0.00040)	(0.00036)
Exposure at $t-1$	-0.001***	-0.002***	-0.001***	-0.002***
	(0.00053)	(0.00053)	(0.00055)	(0.00053)
R^2	0.089	0.089	0.041	0.041
Mean Dep. Var.	0.192	0.192	0.192	0.192
Observations	9,664,857	9,664,857	9,665,755	9,665,755

Notes: Observed swarm exposure and predicted exposure are normalized by their means and standard deviations, such that we obtain z-scores for the two regressions. Columns 1 and 3 use standardized observed swarm exposure at t and $t-1$; Columns 2 and 4 use standardized LASSO-predicted swarm exposure at t and $t-1$. All regressions include WHO Region × Year and WHO Region × Interview Year fixed effects. Standard errors are clustered at the Admin-1 region level.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

C16 Estimating the Costs of the Desert Locust Monitoring System

Calculating the benefit-cost ratio of the current desert locust monitoring system requires estimating the monitoring costs actually incurred. In the main text, we report that we estimate the annual costs of the desert locust monitoring system to range between $37 and $77 million across our study period. We also explain how, to obtain that estimate, we separately calculate the preventive costs incurred in recession years and the emergency costs incurred in plague years. In this section, we provide more details on both calculations and on the supporting documentation.

Preventive costs are incurred by organizations at three different levels: international, regional, and national. We gather data on each organization's budget from publicly available financial documents. At the international level, the annual budget for the FAO Desert Locust Control Committee (DLCC) is $0.2 million (FAO 2023b)[23]. At the regional level, the annual budget for the FAO Commission for Controlling the Desert Locust in the Western Region (CLCPRO), the FAO Commission for Controlling the Desert Locust in the Central Region (CRC) and the FAO Commission for Controlling the Desert Locust in South-West Asia (SWAC) combined is $1.279 million (2023a); we add to that the annual budget for the Desert Locust Control Organization for Eastern Africa (DLCO-EA), estimated at $6.152 million (DLCO-EA 2023)[24]. At the national level, data on national locust control units is largely unavailable[25], and we only recover Chad's budget for 2014, evaluated at $0.89 million (Agence Nationale de Lutte Antiacridienne 2014)[26]. We rely on two distinct estimation strategies to get the total national budget: we take $3 million as our lower bound estimate, based on an FAO estimate that "the 2003-2005 upsurge [...] could have been contained at its outset with just USD 3 million in preventive control expenses."; we take $52.507 million as our upper bound estimate, obtained by scaling up Chad's annual budget, under the assumption that each country's budget is proportional to their contribution to the DLCC (FAO 2023b, Annex V). Summing annual budgets across all preventive organizations yields a lower bound of $10.631 million per year and an upper bound of $60.139 million per year in preventive costs.

To calculate emergency costs incurred during plague years, we rely on documentation of the costs actually incurred to address the plagues that occurred during our study period, as data on locust emergency funds is incomplete and challenging to annualize. This yields $533 million for 2019-2022 (2023c), which we prorate to $266.5 million for the 2019-2020 period; $570 million for 2003-2005 (2023c); and $200 million for 1987-1988 (Skaf et al. 1990). This yields an average of

[23] "The DLCC Secretariat presented the paper (DLCC 2023/25) on International Trust Fund (TF) 9161: contributions and expenditures from 2019 to 2023. The total assessed contributions by member countries is USD 207,780 per year."; "The estimated annual budget approved by the DLCC 41st Session for the biennium 2020-2021 was USD 400 000. Expenditures from 2019 to December 2022 amounted to USD 385 636 equivalent to about USD 112 073 per year.". To be conservative, we use the annualized value of the approved 2020-2021 budget (USD 200 000), which virtually amounts to the same as the annual country contributions, but exceeds the estimated annual expenditures.

[24] We use the annual forecasted budget averaged over 2023-2028. This could be an overestimate, as other documents describe DLCO-EA contributions as being modest in recent decades (FAO 2021).

[25] We contacted several national locust control units but received few responses on their budgets.

[26] Note that we remove from Chad's total budget its contribution to regional organization CLCPRO and international organization DLCC, and convert the budget from FCFA to USD.

$148 million per year in emergency costs incurred across the seven plague years in our data.

Across the 1985-2020 study period, we assign the plague cost estimate to the seven plague years (1987, 1988, 2003, 2004, 2005, 2019, 2020) and the recession cost estimate to the 29 remaining—recession—years. After averaging across years, we obtain that the annual costs of the locust monitoring that occured ranges between $37.356 million and $77.237 million, which we round for exposition purposes.