Proximity to the Plague: Measuring Locust Exposure and Child Health*

Zemar Hakim†

November 1, 2026

Abstract

I investigate early childhood health following the Kenyan 2020–2021 desert locust invasion by linking georeferenced locust observations to a nationally representative health survey. To account for differences in harvest times across the country, I build a treatment that conditions on the insect's biological capacity to damage vegetation, the growth stage of the local crop, and a subsequent decline in vegetation. Children exposed in utero to locust crop damage within 30 km are 25 percent more likely to be wasted than unexposed children. Among crop-growing households, the harm is concentrated in farming zones rather than those that also rely on herding or fishing, and among children born before the lost harvest. This pattern suggests that households cut consumption in anticipation of lost food. Relief may therefore be most effective if it reaches farming households as soon as crop damage is recorded, through fast-paying instruments such as cash transfers or expanded safety nets, rather than after the harvest fails.

Keywords: desert locust, child health, wasting, exposure measurement, Kenya

JEL Codes: I15, O12, Q12, Q54, C18

*All errors are my own.

†Department of Economics, Queen's University.

1 Introduction

Natural disasters that affect agricultural production often pose adverse nutritional effects for young children in low- and middle-income countries. The harsh conditions that these children face in early childhood compound to detrimental effects later in adulthood (Almond and Currie, 2011; Maccini and Yang, 2009; Stein and Susser, 1975). To better coordinate policy efforts, researchers have attempted to identify the effects of these agricultural disruptions on early childhood health using a variety of methods. Cross-sectional household surveys such as the Demographic and Health Surveys (DHS) and the Multiple Indicator Cluster Surveys (MICS) are widely used to identify the consequences of climate events (Michalopoulos and Papaioannou, 2018; Michler et al., 2022; Hogan and Schlenker, 2024; Josephson et al., 2026). As these surveys provide geocoded sample clusters, a researcher can link a household to the shock of interest by drawing a radius around the observed event and treating everyone inside this radial "buffer" as exposed. A small but growing literature in economics has come to employ these techniques to the incidence of locust plagues (Chatterjee, 2022; Le and Nguyen, 2022; De Vreyer et al., 2015). These papers often assign exposure from a locust's recorded presence within a radius or an administrative boundary. Conte et al. (2023) use a similar measure and rely on Mali's primarily single harvest to time specific pathways to health from the locusts. However, many countries hit by these types of agricultural shocks are bimodal; they have multiple staggered harvests that follow from two rainy seasons in a year. For example, the 2020–2021 desert locust invasion mainly hit bimodal East African countries such as Uganda, Ethiopia and Kenya. In these settings, a pest invasion at the same time of year might destroy one farmer's crops while flying over another farmer's empty field.

In this paper, I reconcile these differences in harvest timing by shifting local farming calendars according to observed rainfall. I then match georeferenced locust observations to the 2022 Kenya DHS and assign treatment that requires the insect to have arrived when a crop was susceptible to damage. A locust record counts only if it falls within a given radius of a survey cluster while the child was in utero, if it displays a biological form capable of stripping a field, if the local calendar places a crop in a susceptible growth stage on the record's date, and if satellite vegetation over the cluster then declines relative to that cluster's own pre-invasion conditions. Due to the geographical displacement of the DHS, I build the treatment at nine radii from 10 to 50 km.

Children exposed in utero using a 30 km buffer under this measure carry weight-for-height (wasting) z-scores 0.207 standard deviations below unexposed children. Adding county-specific linear time trends raises the estimate to 0.235. At the clinical threshold the probability of wasting rises 1.8 percentage points against a baseline of 7.1 percent, a relative increase of 25 percent, and 3.4 percentage points once county trends are added. The estimate is negative and significant at the 5 percent level or better at every radius from 10 to 50 km and grows as the radius tightens, reaching 0.340 standard deviations within 10 km.

Wasting records a deficit of weight relative to attained height and responds to adverse nutritional effects in the short term within months. Height-for-age (stunting) records accumulated growth over the whole of a child's life and moves slowly. The median child in this sample was 14 months old at measurement. A disruption confined to gestation and the months around birth has had little time to register in linear growth. Height-for-age z-scores are significant at 30 km when exposed only in the third trimester and among children 24 months and over. The effects on wasting are strongest and significant when children are exposed in the second and third trimesters.

The harm is also concentrated among children born after the first swarm reached their area but before the affected harvest was due. Together with a larger effect in farming zones that do not also rely on fishing or herding, this suggests that households cut consumption in anticipation of the loss. Two policy implications

follow. First, relief should be targeted by recorded crop damage: a rule based on proximity to locust reports would enroll more than twice as many children, many of them in pastoral households with no crop to lose. Second, because the harm begins before the harvest fails, support should reach crop-dependent households as soon as damage is recorded, through fast-paying instruments such as cash transfers or expanded safety nets, rather than once the shortfall arrives.

In addition to the novel design, this paper also contributes to the fetal origins literature. Barker (1990) proposed that adult disease traces to conditions in-utero, and economists have since documented effects of prenatal shocks on long-term health, schooling, and earnings (Almond and Currie, 2011; Almond et al., 2018). Rainfall and drought supply most of the identifying variation. Maccini and Yang (2009) link early-life rainfall to adult height and schooling in Indonesia, Hoddinott and Kinsey (2001) and Alderman et al. (2006) document persistent growth faltering after Zimbabwe's droughts, and Shah and Steinberg (2017) show that rainfall raises human capital in utero while lowering it at school age. Anthropometric z-scores are the standard outcome because early childhood growth predicts adult human capital (Victora et al., 2008). Rainfall shocks identify the effect of income and food availability on the fetus. But they cannot separate the agricultural channel from the waterborne-disease channel, the maternal-labor channel, or the price channel. My exposure measure isolates the agricultural channel by requiring a damaging swarm, crops on the ground, and a subsequent decline in vegetation.

The remainder of this paper is organized as follows. Section 2 describes the locust, Kenya's agricultural landscape, and the 2020–2021 invasion. Section 3 describes the data. Section 4 builds the exposure measure and the estimating equation. Section 5 reports the health results, robustness checks (Section 5.2) and potential mechanisms (Section 5.3). Section 6 discusses the results.

2 Setting

2.1 Kenya's Agricultural Landscape

Kenyan agriculture is organized around two rainy seasons. The long rains fall from roughly March to May and the short rains from October to December. Most of the country's maize, the dominant food staple, is planted with the long rains in the western and Rift Valley highlands and harvested from August onward. Eastern and coastal areas depend more heavily on the short rains and on drought-tolerant crops including sorghum, millet, and cowpeas. The timing of planting, vegetative growth, and harvest differs across the country by several months, and a locust arrival that is catastrophic in one county may be harmless in another on the same date.

County boundaries do not track how households feed themselves. Figure 1 maps the Famine Early Warning System Network's (FEWS NET) 24 livelihood zones under the household economy approach, each covering an area whose residents share a production system and a set of food-access strategies (Famine Early Warning Systems Network (FEWS NET), 2011). Each zone has a profile describing its markets, crops, sources of food and income, livestock and hazards (Figure 16 in Appendix B).

Cropping zones cover the western and central highlands, the eastern slope of the highlands, and the coastal strip. They include the Central Highlands High Potential Zone (KE19), the Western Medium and High Potential Zones (KE20 and KE21), the Southeastern mixed farming zones (KE16 and KE17), two coastal farming zones (KE12 and KE13) and two marginal mixed farming zones (KE06 and KE22). Households in these zones grow maize, beans, potatoes, and in the highlands tea and coffee, and buy the balance of their

food.

Pastoral zones cover the arid north and northeast and a belt of southern rangeland. They include the Northwestern, Northern, Northeastern, Eastern, Southeastern, and Southern Pastoral Zones (KE01, KE05, KE09, KE10, KE11, KE15). Households herd camels, cattle, and small stock, move with the rains, and obtain food from milk, livestock sales, and purchase.

Agropastoral zones sit on the transition between the two and combine rain-fed cropping with herding (KE03, KE07, KE18, KE24). Riverine zones are narrow irrigated strips along the Turkwell, the Dawa at Mandera, and the Tana (KE02, KE08, KE14). Fishing zones cover the shores of Lake Turkana and Lake Victoria (KE04, KE23).

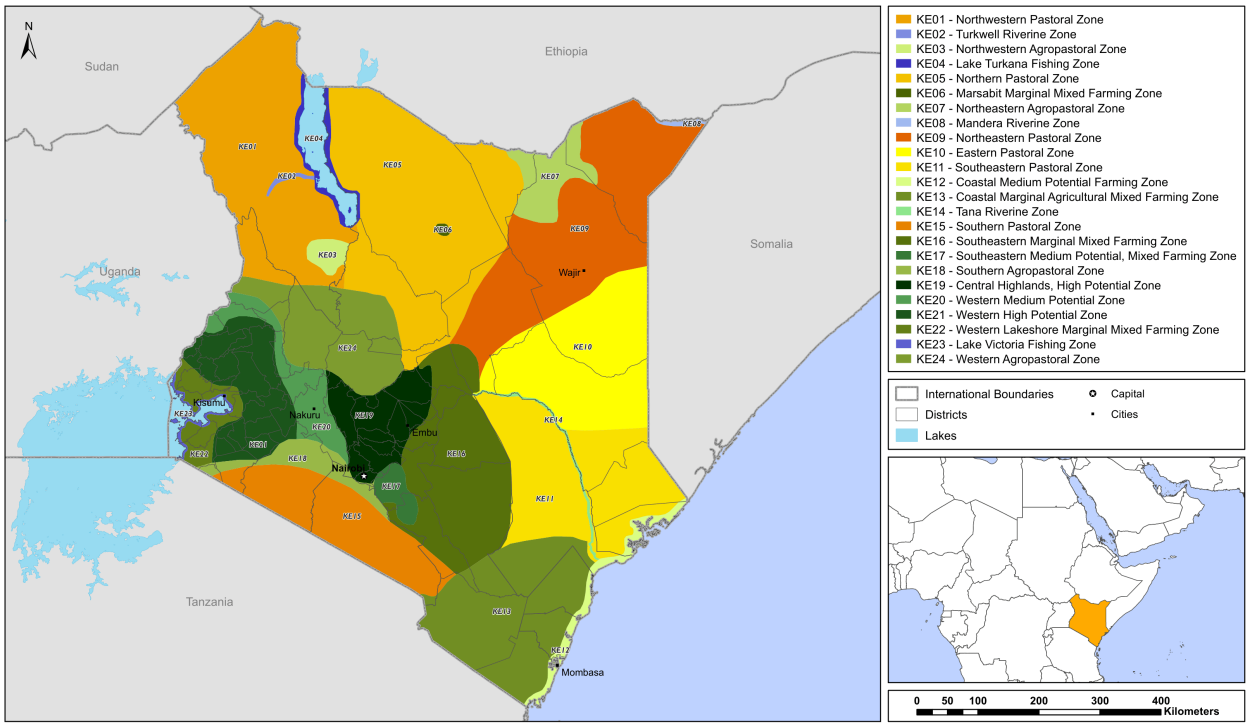


Figure 1: FEWS NET livelihood zones of Kenya. Source: (Famine Early Warning Systems Network (FEWS NET), 2011)

2.2 The Desert Locust

The desert locust, *Schistocerca gregaria*, is a grasshopper that changes phase with population density. At low densities it is solitary (solitarious). These locusts avoid one another, move relatively little, and typically remain within a restricted area. Their coloration is usually cryptic, matching the background vegetation and soil, and their bodies are slimmer with proportionally longer legs. Because they neither aggregate nor migrate in large numbers, solitary locusts cause little or no damage to crops and are not considered a plague threat. This phase persists as long as densities remain low, and it is the state in which locust populations spend most of their time.

In the gregarious phase, by contrast, locusts are crowded and highly social. When rainfall concentrates individuals onto scarce green vegetation, crowding triggers a transformation toward this phase over hours to generations (Symmons and Cressman, 2001). Gregarious individuals aggregate, march as hopper bands

before they can fly, and form adult swarms that travel with the wind over hundreds of kilometers. They develop striking black-and-yellow (or pink in immature adults) warning coloration, broader bodies, and shorter legs, and they are more active and voracious than their solitary counterparts. Swarms can contain billions of individuals and consume vast quantities of vegetation in a single day, making the gregarious phase the one responsible for the devastating agricultural damage for which desert locusts are notorious. The two phases are endpoints of a continuous, density-dependent continuum, with a transient phase occupying the middle ground.

Only gregarious swarms and hopper bands reach the densities that strip a field. Solitarious individuals occur throughout the locust's recession area in most years without agronomic consequence.

Damage also depends on timing. Locusts are polyphagous and consume cereals, legumes, and pasture, but the loss a field sustains depends on what stage the crop has reached. Standing biomass close to harvest cannot be replaced within a season. A field at planting can be resown at the cost of seed and labor. Grain already harvested and stored is not at risk, and a field out of season has nothing to lose. A swarm that passes over a harvested field still enters the record as a locust observation.

Upsurges are episodic. Global control efforts suppressed them for much of the late twentieth century (Zhang et al., 2019), and recent work links renewed activity to extreme rainfall events and their climatic drivers (Salih et al., 2020; Liu et al., 2024; Peng et al., 2020).

2.3 The 2020–2021 Kenyan Crisis

Heavy and repeated rainfall over the Arabian Peninsula in 2018 and 2019 produced several undetected breeding generations of the desert locust. Tropical cyclones delivered exceptional precipitation to usually arid terrain, creating temporary lakes in the desert and allowing locusts to breed (Salih et al., 2020). These favorable conditions sustained three generations of breeding in roughly nine months.

Swarms then crossed into the Horn of Africa through Yemen, Somalia, and Ethiopia, reaching northern Kenya in late December 2019. By early 2020 swarms had spread across much of central and eastern Kenya, prompting the national government to declare an unprecedented threat to food security and livelihoods. The invasion represented the worst locust upsurge Kenya had experienced in seventy years.

The upsurge continued into 2021. A second wave entered Kenya from Ethiopia and Somalia in mid-November 2020 and eventually reached 15 counties, reversing the gains of the preceding months (Ministry of Agriculture, Livestock, Fisheries and Cooperatives, Kenya, 2021). Kenya responded with nine surveillance and spray aircraft, 21 ground spray vehicles, and more than 500 trained personnel, operating alongside FAO-supported control operations (EnviroNews Nigeria, 2021). Ground and aerial spraying proceeded while breeding continued, and control teams worked to suppress new hopper bands before they could fledge into swarms. The campaign brought the invasion largely under control by mid 2021 with some gregarious sightings in the fall of that year. The FAO declared the invasion officially over in February 2022 (Food and Agriculture Organization of the United Nations (FAO), 2022b).

Figure 2 plots daily counts of gregarious swarm and hopper-band observations across the episode. The invasion arrives in two waves. The first begins in late December 2019, peaks in May and June 2020, and subsides by August. The second builds from December 2020, peaks in February 2021, and ends in April. The shock ran for sixteen months and crossed two long-rains seasons and one short-rains season.

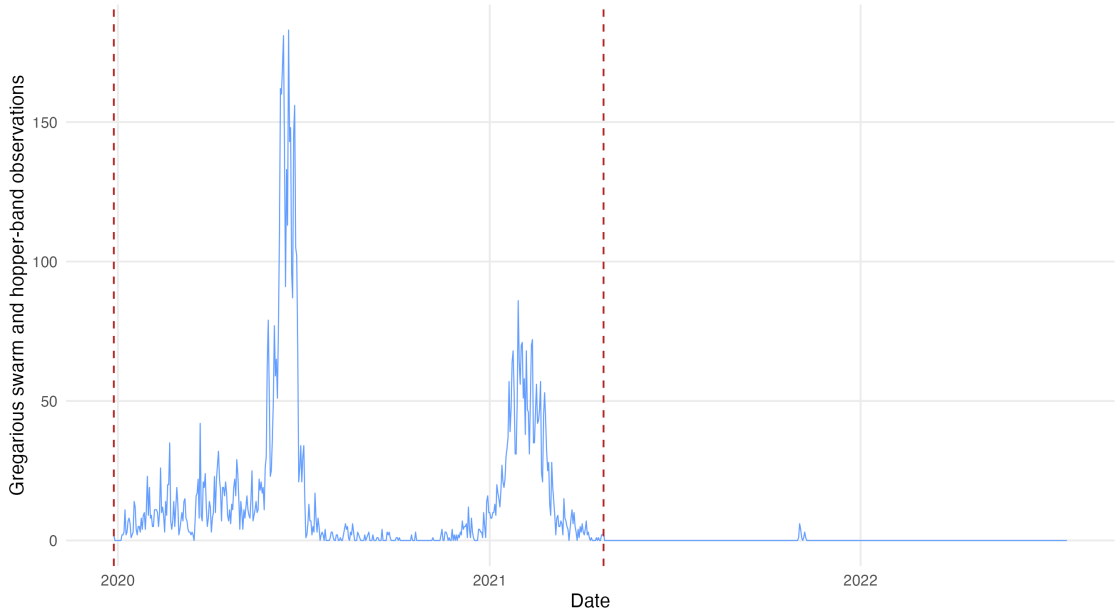


Figure 2: Daily gregarious swarm and hopper-band observations in Kenya. Dashed lines mark the derived invasion window.

3 Data

3.1 Household Survey

The child-level sample is drawn from the 2022 Kenya Demographic and Health Survey (DHS), with the 2014 round used as a pre-invasion falsification sample (Kenya National Bureau of Statistics (KNBS) and ICF, 2023; Kenya National Bureau of Statistics (KNBS) and ICF International, 2015). The DHS is a nationally representative household survey conducted under the DHS Program (ICF, formerly funded principally by USAID), fielded in Kenya roughly every five years. Each round enumerates all children born to sampled women in the preceding five years, with retrospective birth histories, anthropometric measurement (height and weight, converted to age- and sex-standardized z-scores against the WHO growth reference of WHO Multicentre Growth Reference Study Group (2006)), and household, maternal, and child covariates.

Each DHS cluster (an enumeration area, roughly a village or urban block) is assigned a single GPS coordinate, randomly displaced by up to 2 km for urban clusters and 5 km for rural clusters (up to 10 km for 1 percent of rural clusters, restricted to stay within the same second administrative region) to protect respondent confidentiality (Burgert et al., 2013). Raster and vector extraction in this paper, such as rainfall and distance to roads, uses a buffer mean or nearest-feature match over this displacement radius around the cluster coordinate.

The 7,603 children with valid anthropometry are spread across 1,644 clusters, a mean of 4.6 children per cluster and a median of 4, ranging from a single measured child in the thinnest clusters to 19 in the largest. The sample is restricted to children whose gestation overlapped the invasion window, 28 December 2019 to 23 April 2021. These children were born between December 2019 and January 2022 and were 1 to 30 months old at measurement.

3.2 Locust Records

Locust observation records come from the Reconnaissance and Management System of the Environment of Schistocerca (RAMSES), the geographic information system used by national Locust Information Officers in each affected country to record survey and control activity, and consolidated by the Food and Agriculture Organization's Desert Locust Information Service (DLIS) (Food and Agriculture Organization of the United Nations (FAO), 2022a). Records are distributed through the FAO Locust Hub. Each record carries a date, coordinates, a presence type (e.g. hopper band, adult swarm, or an explicit absence), a phase classification (gregarious, transiens, or solitary), a maturity classification, and, where control operations occurred, the control type.[1]

3.3 Climate and Vegetation

I use three remotely sensed products: CHIRPS rainfall, TerraClimate potential evapotranspiration, which I combine with CHIRPS into the SPEI drought index, and MODIS NDVI. Rainfall comes from the Climate Hazards group InfraRed Precipitation with Station data (CHIRPS), a quasi-global, 0.05°-resolution, daily-to-monthly precipitation dataset combining infrared satellite imagery with station observations, produced by the Climate Hazards Center at UC Santa Barbara and designed for drought monitoring in data-sparse regions (Funk et al., 2015). Both monthly and dekadal (10-day) CHIRPS products are used: the monthly product is the precipitation input to the drought index described next, and the dekadal product is used to detect the start of each growing season for the crop calendar.

The Standardized Precipitation Evapotranspiration Index (SPEI), at 6-month accumulation windows, is constructed from CHIRPS precipitation and potential evapotranspiration (PET) from TerraClimate, a global, ∼4 km monthly climate and water-balance dataset (Abatzoglou et al., 2018). SPEI standardizes the climatic water balance, precipitation minus PET, following the index definition in Vicente-Serrano et al. (2010); because the drought conditions overlapping the invasion period combined rainfall deficits with elevated evaporative demand, a precipitation-only index might understate its effects.

Vegetation condition comes from MODIS/Terra Vegetation Indices (MOD13Q1), a 16-day, 250 m composite Normalized Difference Vegetation Index (NDVI) product from NASA's MODIS sensor, accessed via the AppEEARS extraction service (Didan, 2015). Pixels are screened on the product's reliability flag before averaging, so cloud- or snow-contaminated observations never enter a cluster mean. NDVI levels are converted to per-cluster, per-composite z-scores against a pre-invasion climatology. Appendix B.3 describes how SPEI and the NDVI z-scores are computed.

3.4 Livelihood Zones and Crop Calendars

The exposure measure requires knowing, for a given cluster on a given date, whether a crop was standing in a field and how far it had advanced. Livelihood zone boundaries come from the FEWS NET shapefiles for the Kenya livelihood zone map described in Section 2 (Famine Early Warning Systems Network (FEWS NET), 2011).

The FEWS NET livelihood profiles also published a seasonal calendar for each of Kenya's 24 livelihood zones, showing land preparation, planting, and harvesting by crop alongside livestock, labor, and market-access cycles (Famine Early Warning Systems Network (FEWS NET), 2011). A cropping zone's calendar lists maize, beans, potatoes, and the relevant cash crops. A pastoral zone's calendar lists calving, kidding,

[1] Unfortunately, control operation data are sparse and recorded with many errors in this sample.

migration, and milk availability, and contains no crop row. Clusters in zones without a crop row are never placed in a susceptible stage and can never satisfy the stage condition.

Each seasonal calendar uses a three-color Gantt legend, land preparation, planting, and harvest, plotted month by month for every crop grown in that zone (Famine Early Warning Systems Network (FEWS NET), 2011).[2] An automated pixel-extraction routine locates the table grid and the legend's three reference colors on each of the 24 zone pages, then classifies every crop row's bar by nearest-color match at semi-month resolution. Seventeen zones return at least one crop row. The other seven, six pastoral zones and the Lake Turkana fishing zone, are detected automatically as having no crop calendar.[3] The automated output was checked against an independent manual pixel read of one page, which caught and fixed an off-by-one error in which row's bar was attributed to which crop, and all 21 automatically read zones run through the identical code path after that fix.

The three-color legend collapses to a single susceptibility ranking. Land preparation leaves bare ground and is not susceptible. Planting is moderately susceptible, since a destroyed planting can be resown at the cost of seed and labor. Harvest, when a mature crop stands in the field, is the most susceptible stage. A cluster-date outside every crop's bar on its zone's page is out of season. The calendars leave the months between the end of planting and the start of harvest blank, and a crop growing in those months counts as out of season (Appendix B.1).

Where a cluster has more than one crop group or more than one season active on a date, the governing stage is the most susceptible of those active, as Section 4 formalizes.

3.5 Food Prices

Prices come from the World Bank's Real-Time Prices (RTP) platform, which supplies a monthly price estimate for maize and beans at 233 Kenyan markets from 2007 onward, combining sparse directly-observed transactions with a model that imputes missing values from nearby markets, related commodities, and historical patterns, following the method in Andrée (2021). I take the platform's own "close" estimate, its best and most current price for a given market-month. I restrict the sample to 2015–2022. Appendix B.4 describes the imputation and how cluster-level prices are built from the market series.

3.6 Roads

Distance from each DHS cluster to the nearest major road is computed from OpenStreetMap's Kenya road network, exported through the Humanitarian OpenStreetMap Team's export tool (Humanitarian OpenStreetMap Team (HOT) and OpenStreetMap contributors, 2022). This is a crowd-sourced, continuously updated vector road network.

3.7 Administrative Boundaries

County boundaries come from the geoBoundaries Global Database of Political Administrative Boundaries, distributed via the Humanitarian Data Exchange (HDX) as the Kenya subnational administrative boundaries

[2] Figures 17 and 18 in Appendix B provide examples of one calendar for each type.

[3] Three zones, Tana Riverine, Western High Potential, and Lake Victoria Fishing, needed the corresponding page read by hand instead: the automated color-distance threshold that separates the legend colors was not stable across the environments the pipeline was run in for these three pages specifically, so those three pages are transcribed by hand and flagged as such in the underlying data.

dataset. Kenya's 47 counties correspond to the ADM1 level in this dataset, which is produced and maintained by geoBoundaries and published under an open license (geoBoundaries, 2022).

4 Empirical Strategy

4.1 Exposure construction

Let $\mathcal{L}$ denote the set of locust monitoring records, where each record $\ell \in \mathcal{L}$ carries a location, an observation date δ_ℓ, a presence indicator $P_\ell \in \{0, 1\}$, and a phase $\phi_\ell \in \{\text{gregarious}, \text{solitary/transiens}\}$.[4] A record's *form* is

$$\text{Form}(\ell) = \begin{cases} \text{damaging} & \text{if } P_\ell = 1 \text{ and } \phi_\ell = \text{gregarious} \\ \text{benign} & \text{if } P_\ell = 1 \text{ and } \phi_\ell = \text{solitary/transiens} \\ \text{absent} & \text{if } P_\ell = 0. \end{cases} \tag{1}$$

Damaging records are present records in the gregarious phase: swarms, hopper bands, and a small number of adult and hopper groups. Benign records are present records in the solitary or transiens phase, whatever their presence type, and are retained throughout as the falsification group in Section 5.2.

For child i born at date b_i to survey cluster $c(i)$, the gestational window is $W_i = [b_i - 266,\, b_i]$ under an assumed 266-day gestation.[5] For a radius r, the set of records a child could have been exposed to is

$$\mathcal{E}_i(r) = \{\ell \in \mathcal{L} : \text{dist}(c(i), \ell) \le r,\ \delta_\ell \in W_i\}, \tag{2}$$

the intersection of geographic proximity and gestational timing. Exposure is indexed by child rather than by cluster, since two children in the same cluster born months apart have different windows W_i and can differ in exposure holding cluster and radius fixed.

4.1.1 Form

The first condition is that the record describe a locust at a density capable of stripping a field. This uses the phase classification already coded in the monitoring data through equation (1). Records that are excluded are carried through the identical construction below as $\text{Form}(\ell) = \text{benign}$ and used as a falsification group.

4.1.2 Crop stage

The invasion struck at different times across the agricultural calendar, and its timing varied from place to place. The May and June 2020 peak fell while the long-rains maize crop stood under its harvest bar in the western zones and was still at planting in the Central Highlands, while the short-rains crop stood under its own harvest bar in parts of the east, and over bare ground under land preparation in counties between

[4] RAMSES leaves the phase field blank for 83 present records, about 1 percent of the records I classify as gregarious, and I assign them to the gregarious phase. All 83 are confirmed observations made between March and December 2020, during the first wave of the invasion. Of the 48 that carry an observer's comment beyond a source tag, 21 describe swarms, hopper bands, crop destruction or a heavy infestation, 12 describe spraying or a call for control, 10 describe flying, roosting or settled adults, one describes scattered adults feeding lightly, one names a tree locust, and three, one of them a test entry, give no information on phase. Classifying these records as non-damaging instead changes the treatment status of 6 of the 1,495 treated children in the weight-for-height estimation sample at the 30 km radius and moves the weight-for-height estimate by less than 0.01 standard deviations.

[5] Tables 6 to 8 in Appendix A vary each choice in the exposure construction one at a time, at the 10, 30 and 50 km radii. A 280-day gestation, counted from the last menstrual period, gives a weight-for-height estimate of −0.181 at 30 km, significant at the 1 percent level, against −0.207 under the main definition. At 50 km it falls to −0.092 and loses significance.

seasons. The February 2021 peak fell after the short-rains harvest bar had closed in most of the country. Because calendars differ across Kenya by several months, the arrival date alone does not tell whether a crop was lost. Two children conceived in the same month, in clusters 200 km apart, can face swarms on the same dates and face different agronomic consequences.

Thus, the second condition is that the arrival find a crop worth destroying. Let $\mathcal{S}(c, \delta)$ be the set of crop stages active in cluster c's crop calendar on date δ, possibly more than one under staggered planting or multiple crop groups, and possibly empty. Only staple food crops enter $\mathcal{S}(c, \delta)$: maize (including green maize and maize intercropped with beans), sorghum, wheat, teff, beans and other pulses, cassava, Irish potato, and sweet potato. Cash and horticultural crops, such as tea, coffee, cotton, pyrethrum, tobacco, sugar cane, fruit, and vegetables, are excluded, since a swarm that strips them does not destroy the household's own food.[6] Each stage carries a susceptibility rank,

$$\text{rank}(s) = \begin{cases} 3 & s = \text{harvest (standing crop)} \\ 2 & s = \text{planting} \\ 1 & s \in \{\text{land preparation}, \text{out-of-season}\}, \end{cases} \tag{3}$$

and the governing stage on a cluster-date is the highest-ranked active stage,

$$S(c, \delta) = \underset{s \in \mathcal{S}(c,\delta)}{\arg\max}\, \text{rank}(s), \qquad S(c, \delta) = \text{out-of-season if } \mathcal{S}(c, \delta) = \varnothing. \tag{4}$$

A stage is susceptible if its rank exceeds 1,

$$\text{Susceptible}(s) = \mathbb{1}[\text{rank}(s) \geq 2]. \tag{5}$$

The harvest stage, when a mature crop stands in the field, is susceptible at rank 3. Planting is susceptible at rank 2, since a destroyed planting can be resown at the cost of seed and labor. Land preparation and out-of-season are not susceptible, since a swarm arriving before a crop is planted or after it is harvested has nothing standing to destroy.[7] A cluster in a livelihood zone with no crop calendar has $\mathcal{S}(c, \delta) = \varnothing$ on every date and is never susceptible. Damaging-form arrivals excluded by this step, swarms that arrived with no crop in the field, are carried through as a second falsification group.

The transcribed FEWS NET calendar fixes a normative planting date for each zone, but the rains that actually trigger planting arrive earlier or later than that norm in any given year and cluster, so I shift the calendar cluster by cluster to track the rain. Season onset is detected separately for each cluster, year, and season from dekadal (one dekad is equal to ten days) CHIRPS rainfall, searched within a season-specific window, dekads 4 to 18 (roughly early February through late June) for the long rains and dekads 25 to 36 (September through December) for the short rains. Let ρ_d denote a cluster's rainfall total in dekad d of that window. Onset is the first dekad d satisfying two conditions:[8]

[6] A swarm that strips a cash crop destroys income but not the household's own food, so I also rebuild the treatment with cash crops (tea, coffee, pyrethrum, cotton, tobacco and sugar cane) in place of staples, keeping the form and decline conditions, to test for an income shock. The two shocks cannot be separated. Cash crops are grown in cropping zones that also grow staples, and 80 percent of the 813 children treated under the cash-crop measure at 30 km are also treated under the main measure. Entered together, the two indicators trade size from one radius to the next (Table 9 in Appendix A).

[7] Counting only the standing crop (rank 3) as susceptible raises the weight-for-height estimate to −0.225 at 30 km and −0.426 at 10 km, significant at the 5 and 1 percent levels respectively (Table 6).

[8] These are the FEWS NET start-of-season thresholds: at least 25 mm in one dekad followed by at least 20 mm over the next two dekads (Senay, 2004).

1. a trigger, $\rho_d \geq 25$ mm;

2. and a confirmation, that the two dekads immediately following together clear $\rho_{d+1} + \rho_{d+2} \geq 20$ mm, so a single wet dekad followed by drought does not count.

If no dekad in the window satisfies both, onset is not detected for that cluster-year-season, which happens in about 7 percent of cluster-season-years in 2019–2022.

A cluster's climatological onset is the median onset dekad across 1981-2018 for that cluster and season. The transcribed calendar boundaries are shifted, holding stage lengths fixed, by the gap between a given year's detected onset and that climatological median, rounded to the nearest dekad and capped at six dekads (60 days) in either direction so a single anomalous year cannot pull the calendar implausibly far from its transcribed position; a cluster-year-season where onset could not be detected is kept and left unshifted at its climatological position. A season that begins three weeks late shifts that cluster's susceptible window three weeks later.[9]

4.1.3 Verified damage

The third condition is that vegetation actually suffer. NDVI is observed in 16-day composites. Let $z_{c,t}$ be cluster c's NDVI anomaly z-score in composite t, and let t_ℓ index the composite containing event ℓ's observation date. Define the mean anomaly in the two composites preceding and the two composites following the event,

$$\bar{z}_{\text{pre}}(\ell) = \frac{1}{2}\sum_{k=1}^{2} z_{c(\ell),\, t_\ell - k}, \qquad \bar{z}_{\text{post}}(\ell) = \frac{1}{2}\sum_{k=1}^{2} z_{c(\ell),\, t_\ell + k}, \tag{6}$$

a window spanning roughly 16 to 48 days after the observation, depending on where within its own composite the event falls.[10] Composites with a missing value are dropped from the corresponding average. The change in anomaly is

$$\Delta z(\ell) = \bar{z}_{\text{post}}(\ell) - \bar{z}_{\text{pre}}(\ell), \tag{7}$$

and the event is classified as a verified decline if

$$\text{Decline}(\ell) = \mathbb{1}\big[\Delta z(\ell) \leq -\theta\big], \qquad \theta = 0.25. \tag{8}$$

A locust record counts as verified damage only if the cluster's greenness fell by at least 0.25 standard deviations. The threshold only has to separate a fall in greenness from composite-to-composite noise, and the estimate barely moves with it: across thresholds from 0 to 1.5 and the six pre- and post-event windows, the weight-for-height estimate at 30 km stays between -0.143 and -0.247.[11]

[9] The shift matters at the wider radii. With the calendar left at its transcribed FEWS NET dates, the weight-for-height estimate at 30 km is -0.064 ($p = 0.40$), and widening every stage window by 15 days on each side gives -0.110 ($p = 0.11$). Narrowing every window by 15 days gives -0.184, close to the main estimate but less precise ($p = 0.10$) on 725 treated children. Both looser calendars admit arrivals that more likely fell on bare ground. These are timing errors of the kind the stage condition exists to remove, and they water the estimate down toward the proximity null in the same way dropping the stage condition does (Figure 5). At 10 km the unshifted, widened and standing-crop variants remain significant at the 1 percent level and the narrowed variant at the 10 percent level (Table 6). The six-dekad cap binds in 19 of the 26,208 cluster-season-years from 2019 to 2022, under 0.1 percent.

[10] Using one, two or three composites on either side, in the six combinations in Table 6, gives weight-for-height estimates at 30 km between -0.143 and -0.241, all significant at the 5 percent level.

[11] For θ between 0 and 1 the weight-for-height estimate at 30 km lies between -0.194 and -0.247, significant at the 1 percent level throughout. At $\theta = 1.5$, which leaves 282 treated children, it is -0.214, significant at the 1 percent level (Table 6).

4.1.4 The exposure indicator

Treatment requires all three conditions to hold on a single record,

$$T_i^{(r)} = \mathbb{1}\Big[\exists\, \ell \in \mathcal{E}_i(r) : \text{Form}(\ell) = \text{damaging},\ \text{Susceptible}\big(S(c(i), \delta_\ell)\big) = 1,\ \text{Decline}(\ell) = 1\Big]. \tag{9}$$

All three conditions must hold on the same observation. A swarm in March and a vegetation decline in September are not evidence that the swarm caused the decline. Equation (9) admits only arrivals followed by a fall in the greenness of the same cluster within roughly six weeks. Each condition comes from how a swarm damages a crop.

Equation (8) uses a variable closer to the outcome than the other two conditions. Three safeguards apply. The test in equation (7) measures a change in the anomaly. A cluster that was already dry does not qualify on dryness alone. The window in equation (6) is short enough that a later seasonal decline unrelated to the event is not attributed to it. Equations (6) to (8) are applied identically to benign records and to damaging-form records arriving outside the susceptible window. If Decline(ℓ) captured drought, these placebo groups would satisfy it at a similar rate and predict outcomes similarly. Section 5.2 reports that they do not.

Three nested measures are defined by dropping conditions from equation (9). Proximity, $T_i^{(1,r)} = \mathbb{1}[\exists \ell \in \mathcal{E}_i(r) : P_\ell = 1]$, is gated on nothing beyond presence and is the closest analogue to the exposure variables used in existing locust research. Damaging form adds equation (1). Damaging form and stage adds equation (5). Then, by construction $T_i^{(r)} \le T_i^{(3,r)} \le T_i^{(2,r)} \le T_i^{(1,r)}$ and $T_i^{(r)} \subseteq T_i^{(3,r)} \subseteq T_i^{(2,r)} \subseteq T_i^{(1,r)}$. In the rest of the paper, a child with $T_i^{(r)} = 1$ is treated, or exposed, and a record that meets all three conditions is an event.

My measure and the unverified proximity measures treat systematically different children as exposed. The 2,523 children flagged by proximity alone are poorer, more remote and drier than the 1,547 both measures agree were exposed, and 39 percent of them live in pastoral zones, where no crop calendar applies and no child can be treated (Table 10 in Appendix A).[12]

4.2 Health Outcomes

The DHS measures every sampled child's length or height and weight under a standardized protocol: recumbent length for children under 24 months, standing height above that age, and weight to the nearest 100 grams, up to 59 months inclusive. DHS converts these raw measurements to three z-scores centrally, before the data is released, against the WHO Child Growth Standards, the international reference built from healthy, breastfed children raised without growth constraints (WHO Multicentre Growth Reference Study Group, 2006). The conversion uses the WHO's LMS method: at each value of the conditioning variable (age for HAZ and WAZ, height or length for WHZ) the reference population's distribution is summarized by three sex-specific parameters, a median M, a coefficient of variation S, and a Box-Cox skewness power L, and a child's raw measurement X is converted to a z-score

$$Z = \begin{cases} \dfrac{(X/M)^L - 1}{L\,S} & L \neq 0 \\ \dfrac{\ln(X/M)}{S} & L = 0, \end{cases} \tag{10}$$

[12] Table 11 and Figure 9 report counts at every radius. These counts, like the 1,547 here, cover all 7,603 children with valid anthropometry. Regression tables instead report treated children in the estimation sample. Of the 1,547 children treated at 30 km, 1,546 have a valid weight-for-height score, 1,544 of those also have complete controls, and 1,495 remain after children who are the only observation in their survey cluster or birth cohort drop out of the fixed-effects estimation. Across the nine radii, the 675 to 1,987 treated children in Table 11 correspond to 647 to 1,918 in the estimation sample.

so that $Z = 0$ is the reference median and $Z = -2$ is two standard deviations below it, in the reference distribution's own units, regardless of the child's age. I use the DHS's pre-computed values and set to missing any value the WHO flags as biologically implausible: height-for-age below -6 or above 6, weight-for-height below -5 or above 5, and weight-for-age below -6 or above 5.

Weight-for-height (WHZ) compares a child's weight to the reference for a child of the same sex and the same attained height or length, with age entering only through which reference curve applies. It removes the height a child has already reached and measures only whether her current body mass is proportionate to it. Since it carries no growth history of its own, WHZ is sensitive to recent nutritional deficits, and the effect of a shock on it fades over the first two years, while a height deficit accumulates. The clinical label for WHZ below -2 is *wasted*, and a child below -3 is severely wasted.

Height-for-age (HAZ) and weight-for-age (WAZ) are constructed identically, applying the same LMS transformation in equation (10) against the sex-specific WHO reference curves for length-or-height-for-age and weight-for-age respectively, in place of the weight-for-height curve, with the same WHO flags. They differ from WHZ, and from each other, only in which anthropometric ratio and which reference curve they use. HAZ compares height to the reference for a child's age. It is a cumulative record: a child's height on any given day reflects every month of growth, or growth faltering, since conception, and a single bad month barely moves it. A low HAZ score is a proxy of sustained undernutrition, repeated illness, or a poor environment operating over a long horizon, and the clinical label for HAZ below -2 is *stunted*. WAZ compares weight to the reference for a child's age; because age enters through both the numerator's growth history and the denominator's reference curve, WAZ is a composite of the other two, a child can have a low WAZ score because she is short for her age, because she is thin for her height, or both, and the score alone does not distinguish them. The clinical label for WAZ below -2 is *underweight*.

4.3 Estimation

I estimate

$$y_{ict} = \beta\, T_i^{(r)} + X_{ic}'\gamma + \alpha_c + \delta_t + \theta_{k(c)}\cdot t + \varepsilon_{ict}, \tag{11}$$

where y_{ict} is the weight-for-height z-score of child i in survey cluster c born in year-month t, $T_i^{(r)}$ is the exposure indicator at radius r defined in Section 4.1, α_c is a survey-cluster fixed effect, δ_t is a birth-cohort fixed effect defined at the year-month level, so that May 2020 and May 2021 are separate cells, and $\theta_{k(c)}\cdot t$ is a county-specific linear trend in that same year-month index.

The identifying variation is within cluster and across birth cohorts. Two children in the same enumeration area, born eight months apart, face different gestational windows. One window overlaps a verified damage event within r kilometers and the other does not. Equation (11) compares them.

I include survey cluster and birth-cohort (month×year) fixed effects. The county-specific linear trend allows counties to sit on different trajectories over the period. The main specification, column (3) of Table 2, which every other table follows, omits the trend; column (4) adds it.

The control set follows Conte et al. (2023). Child controls are sex, birth order, first-born status, twin status, and the interval since the preceding birth.

Family controls are mother's education, mother's height, household head age, household head sex, wealth quintile indicators, and the six-month SPEI over the gestational window. Mother's height enters because it captures the mother's own early-life nutrition and genetic factors that are strong predictors of a child's

growth potential.[13] Wealth quintile enters as a set of indicators because its relationship with anthropometry is not monotone in the data. Rainfall that is favorable enough to support three generations of locust breeding is also favorable for crops and pasture, which makes drought conditions correlated with both the treatment and the outcome. I include SPEI over the child's own gestational window to control for this.[14]

Standard errors are clustered on a half-degree spatial grid, which yields 166 cells across the sample. The cells coincide with those of PRIO-GRID, a standard grid for spatial analysis at this resolution (Tollefsen et al., 2012). These are coarser than the survey cluster. A swarm that arrives at one enumeration area almost always arrives at its neighbors, and the downstream effects on vegetation, local food availability, and prices propagate over a similar area (see Figures 20 and 21 in Appendix B). Residuals are therefore correlated across nearby clusters and not only within them, and clustering at the DHS cluster would overstate precision. The half-degree grid gives a wider-than-cluster but still local unit. Clustering at the county instead would leave 47 groups, too few for reliable asymptotic inference and too coarse to absorb variation.

All regressions are weighted by the DHS women's individual sampling weight. The 2022 DHS is a stratified two-stage sample: enumeration areas are drawn with probability proportional to size within each stratum and a fixed number of households is drawn in each, and some areas are sampled at higher rates so that they yield precise estimates of their own. Households enter with unequal probabilities. The weight is the inverse of the selection probability, adjusted for nonresponse in the stratum, and without it estimates lean toward the strata sampled at higher rates.[15] The weighted estimate is the effect for Kenyan children born in this window.[16]

Table 1 compares exposed and unexposed children on the covariates available before treatment. The unconditional comparison looks like a typical cross-section of Kenyan households sorted by geography: exposed children have better-educated mothers, are more likely to own land, have higher wealth, and live closer to a road, and fifteen of the nineteen variables differ at the 5 percent level. Almost none of this survives conditioning on the same cluster and birth-cohort structure the main specification uses. Mother's education, wealth, land ownership, and household size all become indistinguishable from zero, as do mother's height and household head age ($p = 0.11$ for both). Gestational NDVI remains significant after conditioning, in the opposite direction from its raw comparison, and gestational SPEI points the same way but is no longer significant ($p = 0.28$). Exposed children experience modestly less green, and somewhat drier, gestational windows than their unexposed neighbors within the same cluster and cohort. Section 5.2 addresses this.

[13] Mother's height is recorded only for women in the subsample of households where DHS measured women, so it is missing for 48 percent of children with valid anthropometry. I fill a missing value with the median among measured women. Table 12 in Appendix A re-estimates all three outcomes adding an indicator for a filled value, dropping mother's height, and restricting to children whose mothers were measured (3,513 in the weight-for-height estimation sample). At 30 km the weight-for-height estimate is −0.207, −0.206 and −0.215 under the three alternatives, all significant at the 1 percent level. The weight-for-age estimate is −0.187 and −0.177 under the first two and −0.160, significant at the 5 percent level, among measured mothers. At 30 km height-for-age stays insignificant under every alternative.

[14] Occupation and food consumption are excluded from the control set. These variables are unfortunately noisy self-reports in the DHS and are largely missing or the timing requested is irrelevant to the study; one week in most cases.

[15] See the DHS Program's Guide to DHS Statistics, which also notes that DHS no longer advises against weighting regressions.

[16] The weight distribution is slightly right-skewed. Appendix B describes it and reports the estimates under winsorized weights and without weights (Table 44). At 30 km the weight-for-height estimate is −0.186 with winsorized weights and −0.118 unweighted, both significant at the 5 percent level. The unweighted estimate is smaller, so the effect is larger among children who carry larger weights, those in strata the design sampled at lower rates.

Table 1: Covariate balance, 30 km radius

	Control	Treatment	Raw difference	FE-adjusted	(SE)
Child					
Female	0.489	0.477	−0.012	−0.029	(0.025)
Birth order	3.36	2.75	−0.609***	−0.040	(0.084)
First born	0.245	0.307	0.062***	0.017	(0.028)
Twin	0.031	0.020	−0.011**		—
Birth interval, prior (months)	46.2	52.8	6.53***	1.71	(1.79)
Birth interval, next (months)	18.5	17.2	−1.28**		—
Mother and household					
Mother's education (years)	7.20	9.46	2.27***	−0.033	(0.169)
Mother's height (cm)	160.8	160.6	−0.208*	0.346	(0.214)
Household head age	39.0	38.1	−0.883**	−1.08	(0.673)
Female household head	0.289	0.270	−0.019	−0.013	(0.022)
Household size	6.04	5.35	−0.686***	−0.051	(0.160)
Wealth quintile	2.49	3.04	0.543***	0.024	(0.032)
Assets and access					
Owns land	0.467	0.537	0.070***	−0.019	(0.019)
Owns livestock	0.687	0.655	−0.032**	0.018	(0.016)
Land owned (hectares)	1.08	0.77	−0.304*	−0.363	(0.400)
Rural	0.679	0.621	−0.058***		—
Distance to major road (km)	8.17	3.80	−4.37***		—
Gestational climate					
SPEI (6 month)	1.010	1.251	0.242***	−0.130	(0.119)
NDVI z-score	0.683	0.736	0.052**	−0.143***	(0.044)
Observations	6,056	1,547			

Raw difference is the unconditional mean difference, exposed minus unexposed, with significance from a two-sample t-test. FE-adjusted is the coefficient on the exposure indicator from regressing each covariate on exposure with the same cluster and birth-cohort fixed effects, sampling weight, and grid-clustered standard error as column (3) of Table 2. ***, **, * denote significance at 1, 5 and 10 percent. Rural and distance to a major road are fixed within a DHS cluster and have no within-cluster variation once the cluster fixed effect is included, so no FE-adjusted coefficient is reported for either. Twin status and the birth interval to the next child likewise return no FE-adjusted coefficient in the same cluster-by-cohort structure, so none is reported for either. Counts in the observation block are the 7,603 children with valid anthropometry; the estimation sample for each covariate differs where a covariate has missing values.

5 Results

5.1 Health

5.1.1 Wasting

Table 2 reports the wasting estimates. Columns (1) to (4) use the weight-for-height z-score. Column (1) includes only the fixed effects, column (2) adds the child controls, column (3) adds the family controls, and column (4) adds county-specific linear time trends. Columns (5) and (6) repeat columns (3) and (4) with an indicator for wasting, a weight-for-height score below −2, as the outcome.

In utero exposure to an event within 30 km lowers weight-for-height by 0.235 standard deviations in the most conservative specification. This corresponds to a 3.4 percentage point rise in the probability of wasting in column (6), a 48 percent increase over the baseline of 7.1 percent among unexposed children. Without the county trend the estimate is smaller on both measures, 0.207 standard deviations on weight-for-height and

1.8 percentage points on wasting, a relative increase of 25 percent. The weight-for-height z-score estimate moves by less than 0.02 standard deviations as the child and family controls enter. The standard error is stable between 0.064 and 0.071 throughout and the weight-for-height z-score estimate is significant at the 1 percent level in every column.

Table 2: In utero locust exposure and child wasting, 30 km radius

	Weight-for-height z-score				Wasted	
	(1)	(2)	(3)	(4)	(5)	(6)
In utero exposure	−0.214***	−0.218***	−0.207***	−0.235***	0.018**	0.034***
	(0.066)	(0.065)	(0.064)	(0.071)	(0.008)	(0.009)
Child controls	No	Yes	Yes	Yes	Yes	Yes
Family controls	No	No	Yes	Yes	Yes	Yes
County time trend	No	No	No	Yes	No	Yes
Observations	7,443	7,443	7,416	7,416	7,416	7,416
R^2	0.344	0.351	0.353	0.362	0.261	0.270

All regressions include birth cohort and enumeration area fixed effects. Standard errors clustered on a half-degree spatial grid in parentheses. ***, **, * denote significance at 1, 5 and 10 percent. All regressions weight by the DHS individual sampling weight. Child controls are sex, birth order, first-born status, twin status, and the preceding birth interval. Family controls are mother's education, mother's height, household head age and sex, wealth-quintile indicators, and six-month gestational SPEI. Control coefficients are estimated but not reported. County time trend denotes county-specific linear trends in calendar month of birth. The outcome in columns (5) and (6) is an indicator for a weight-for-height z-score below −2, with a baseline prevalence of 7.1 percent among unexposed children. There are 1,495 treated children in the estimation sample.

Among treated children at the main radius, the number of events during gestation predicts weight-for-height beyond the binary treatment indicator. Table 3 reports the relationship under three functional forms: the raw count of events entering linearly, the log of one plus that count, and the count capped at its 90th percentile among exposed children ($D_9 = 26$ events), to limit the influence of the small number of children exposed to an unusually large number of events. Each additional event is associated with a 0.012 standard deviation reduction in weight-for-height, significant at the 10 percent level; the log and capped forms return larger, more precisely estimated reductions of 0.214 and 0.026 standard deviations, both significant at the 1 percent level. All three specifications share the control set and fixed effects of column (3) of Table 2 and additionally control for the number of RAMSES ground visits recorded within 30 km during gestation, which separates the effect of events from the effect of monitoring intensity in the same area during the same period.

Table 3: Dose response among treated children: number of events and weight-for-height

Dependent variable: Weight-for-height z-score	Linear	Log	Capped at D_9
Events	−0.012*		−0.026***
	(0.007)		(0.008)
Log(1 + events)		−0.214***	
		(0.071)	
Observations	1,259	1,259	1,259

Standard errors clustered on a half-degree spatial grid in parentheses. ***, **, * denote significance at 1, 5 and 10 percent. D_9 is the 90th percentile of events among treated children (26). Sample restricted to children treated at the main 30 km radius; of the 1,544 treated children with complete data, those who are the only such child in their survey cluster or birth cohort drop out of the fixed-effects estimation, which leaves 1,259. All columns include the control set and the survey-cluster and birth-cohort fixed effects of column (3) of Table 2. I control for the number of RAMSES ground visits during gestation, counting every recorded visit within 30 km, of any outcome, during the child's gestation window, to separate the effect of events from the effect of monitoring intensity in that area during that period.

Weight-for-height is the main outcome in this paper because it measures acute malnutrition. It records a deficit of body mass relative to attained height and responds within months. Height-for-age records cumulative linear growth from conception onward and responds only slowly, over quarters and years. The invasion ran for sixteen months, from late December 2019 until April 2021, but the median child with valid anthropometry was 14 months old at measurement and the mean 14.7 months. For the average child, the shock fell during gestation, little more than a year before measurement (WHO Multicentre Growth Reference Study Group, 2006).

The other two standard z-scores follow this logic (Table 13, and Tables 14 and 15 for the stepped-control specifications, all in Appendix A). Height-for-age effects are negative at every radius but significant only at 15, 20 and 25 km at the 5 percent level and at 10 km at the 10 percent level. At 30 km the estimate is −0.063 with a standard error of 0.063, and the probability of stunting is unchanged. Weight-for-age falls 0.187 standard deviations, 90 percent of the weight-for-height estimate, but the probability of being underweight does not change. Weight-for-age combines body mass and linear growth against age, so an acute weight deficit registers there but is diluted by the height component that has not moved. Both outcomes also fall with the number of events (Tables 16 and 17).

Previous work on the timing of growth faltering suggests that height-for-age falls steadily until about 24 months of age and that the absolute height deficit keeps growing after that (Victora et al., 2010; Leroy et al., 2014). Indeed, splitting the sample in this study by age at measurement shows the weight-for-height effect declining with age (Table 18 in Appendix A). At 30 km the estimate is −0.529 for children measured before 12 months and −0.249 at 12 to 23 months, both significant at the 1 percent level, and 0.363 and insignificant at 24 months or older, where only 149 children are treated. Height-for-age is insignificant before 24 months and −0.673 at 24 months or older, significant at the 5 percent level.

5.1.2 Timing in Gestation

Figure 3 splits exposure by the trimester in which the qualifying event falls, at every radius. Figure 10 in Appendix A repeats it for height-for-age and weight-for-age.

First-trimester exposure is never significant at any radius. The third-trimester estimate is largest close to the damage and falls with distance. It is −0.355 at 15 km and −0.141 at 30 km, and from 35 km onward it lies between −0.022 and −0.075 and is insignificant. The second-trimester estimate does not fall with distance. It lies between −0.198 and −0.319 at every radius, is significant at the 1 percent level from 25 to 35 km, at the 5 percent level at 10 and 20 km and from 40 to 50 km, and only at the 10 percent level at 15 km. At 30 km the second-trimester estimate is −0.319 against −0.141 for the third trimester, and the pooled estimate averages over both.

The third-trimester component follows the same distance gradient as the pooled estimate (Section 5.2.1). It is large where an event was close to the household and shrinks as a wider radius adds children whose households were farther from it, which is the pattern a loss of the household's own crop during the months of fastest fetal weight gain would produce. A child exposed in the second trimester is born roughly three to six months after the event, around or after the affected crop's harvest was due, and the effect for these children reaches beyond the immediate area of a swarm. Section 5.3.1 returns to this timing by grouping children by when they were born relative to the harvest that followed the swarm.

On height-for-age, only third-trimester exposure is significant, from 10 to 30 km (−0.289 to −0.158, all at the 5 percent level or better). The third trimester is the period of fastest fetal linear growth, and Stein and Susser (1975) find a smaller but real effect of third-trimester famine exposure on length at birth alongside

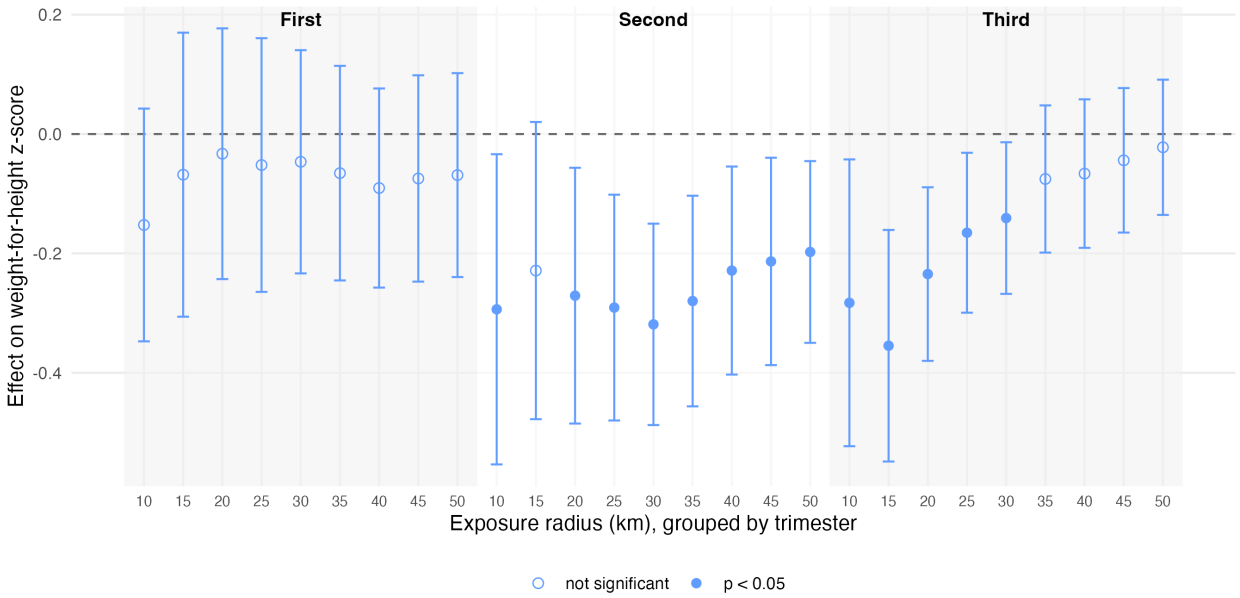


Figure 3: Weight-for-height by trimester of exposure and radius. Each point is a regression on an indicator for treatment with the event in that trimester, with the control set and fixed effects of column (3) of Table 2. Point estimates with 95 percent confidence intervals; standard errors clustered on a half-degree spatial grid.

the larger effect on birth weight.

5.1.3 Postnatal Exposure

The exposure measure is defined on the gestation window only. A child born in January 2020, for example, lived through the invasion as an infant but is included in the comparison group. If postnatal exposure also lowered weight-for-height, this would bias the estimate toward zero. I define postnatal exposure as an event within 30 km of the child's cluster in the first 12 months after birth and before the interview. Among the 7,603 children with valid anthropometry, 996 meet this definition, and 725 of them are not exposed in utero.[17] Adding it to the main specification moves the in utero estimate from −0.207 to −0.191 (standard error 0.070), significant at the 1 percent level, while postnatal exposure itself is 0.045 (standard error 0.071) and insignificant (Table 19 in Appendix A). Splitting postnatal exposure into the first and second half of the first year gives two positive estimates, neither significant, and among children not exposed in utero, postnatal exposure is 0.078 (standard error 0.085). The in utero estimate loses some precision because the two indicators overlap: 271 of the 996 children exposed after birth were also exposed in utero, and its standard error rises from 0.064 to 0.070. Because postnatal exposure is positive in every column, postnatally exposed children in the comparison group are not pushing the main estimate toward zero. Holding postnatal exposure fixed moves the in utero estimate slightly toward zero instead, from −0.207 to −0.191, so the main estimate partly picks up the difference between children exposed before birth and children exposed only after it.

[17] In the weight-for-height estimation sample of Table 19, 971 children are exposed after birth. In column (4), which drops children exposed in utero, 615 are.

5.2 Robustness

5.2.1 Exposure Conditions and Radii

I estimate the main specification at nine radii (Table 20 in Appendix A). Figure 4 plots the weight-for-height estimates with confidence bands.[18] Weight-for-height is significant at the 5 percent level or better at every radius, and at the 1 percent level at every radius up to 35 km. The estimate is largest within 10 km, at 0.340 standard deviations, and falls monotonically with the radius, to −0.155 at 40 km, −0.145 at 45 km, and −0.121 at 50 km. The number of treated children rises from 647 to 1,918 over the same range. Widening the buffer adds children who were farther from any verified damage event and dilutes the average effect, and the estimate shrinks gradually.

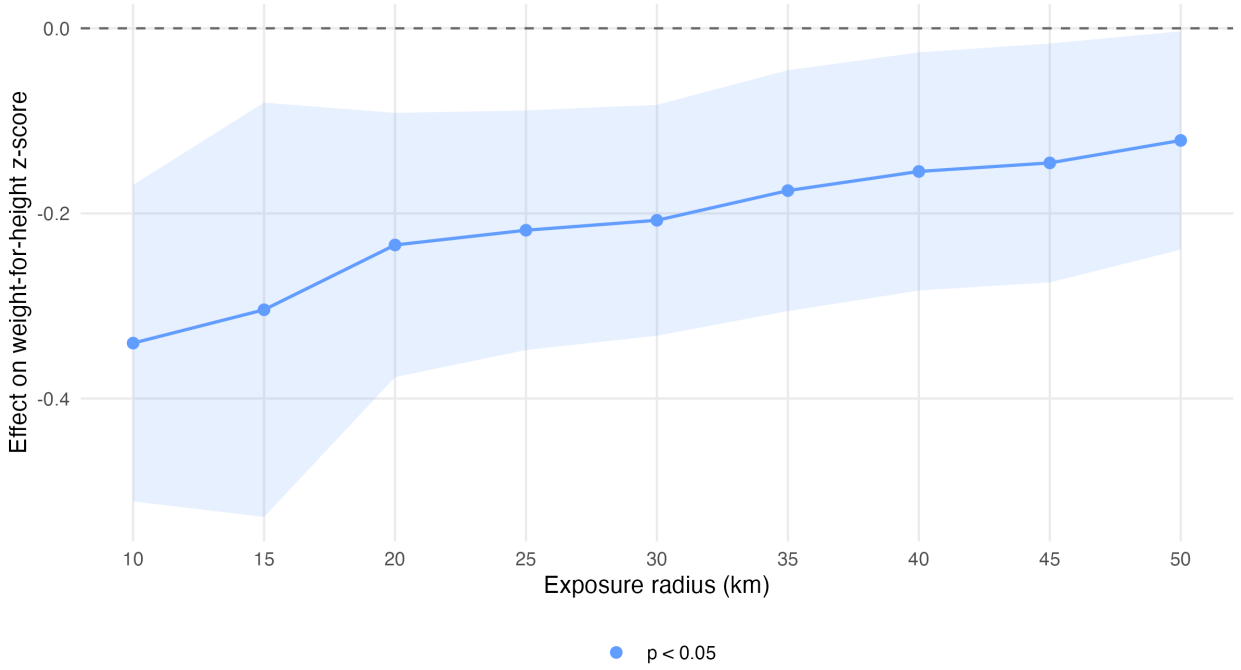


Figure 4: Weight-for-height estimate against exposure radius, with 95 percent confidence bands. Standard errors clustered on a half-degree spatial grid.

Figure 5 shows the conditions of Section 4.1 separately. Proximity alone has no detectable effect at the main radius. Its weight-for-height estimate is −0.003 with a standard error of 0.056, and it flags 4,070 of the 7,603 children with valid anthropometry. Restricting to damaging form removes 79 children and moves the estimate to −0.035. Adding the crop-stage condition removes 1,452 more and moves it to −0.077, still insignificant. Adding the verification condition removes a further 992 and moves the estimate to −0.207, significant at the 1 percent level. Figure 12 in Appendix A shows the same ladder for all three outcomes.

[18] Figure 11 adds the other two outcomes.

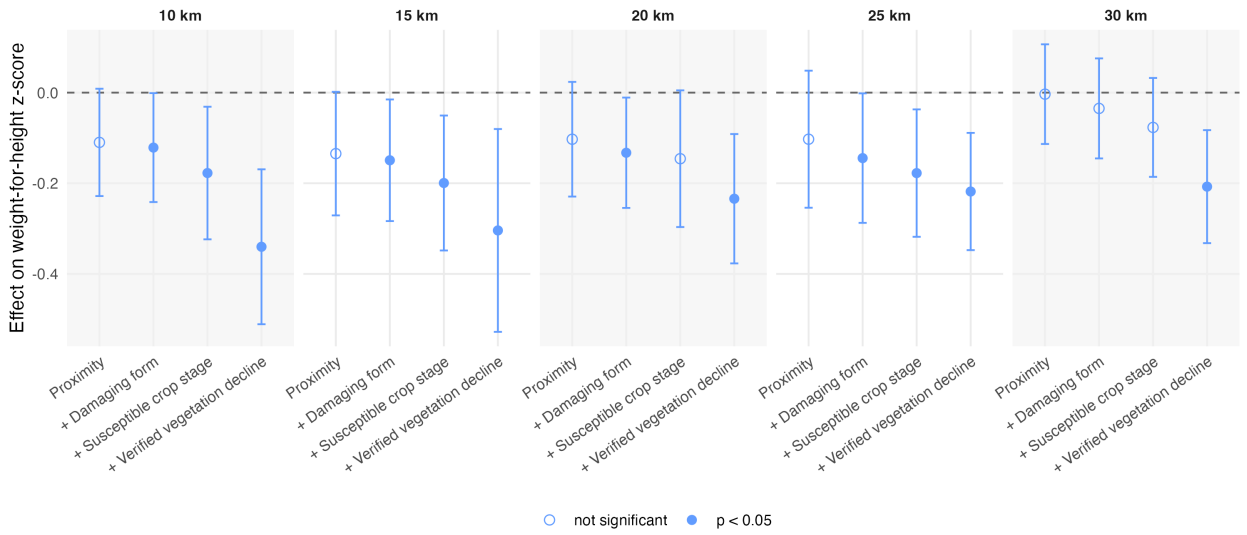


Figure 5: Weight-for-height under each exposure definition, adding one condition at a time, at radii from 10 to 30 km. Point estimates with 95 percent confidence intervals; standard errors clustered on a half-degree spatial grid.

5.2.2 Falsification

I run three falsification tests, each of which keeps part of the exposure construction and removes the part that should make a swarm harmful. If the treatment picked up something other than crop damage, these placebos should also predict lower weight-for-height.

Harmless locusts. The first test replaces damaging-form records with solitarious and transiens records, locusts at densities incapable of stripping a field, and applies the stage and NDVI conditions unchanged. Applied naively, this placebo flags children who are also flagged by the real treatment, because a cluster visited by solitarious locusts during a gestational window can be visited by swarms in the same window. At 30 km, 91 percent of the children the naive placebo flags are also treated. I purge the placebo of any overlap with the treatment before estimation, and do the same for the placebo in the second test.

Panel A of Table 21 in Appendix A reports the purged tests. With the stage and vegetation conditions applied, too few children remain below 35 km, and the solitarious placebo is never negative and significant; at 35 km it is 0.131 (standard error 0.106). Without those conditions it is 0.027 (0.099) at 30 km. One cell is positive and significant at the 5 percent level, the conditioned placebo at 40 km (0.389), which rests on a small group after purging. Purging sets the placebo indicator to zero for treated children but leaves them in the comparison group, and because treated children have lower weight-for-height, this pushes the placebo estimates upward. Panel B adds the treatment indicator at the same radius as a control, and Panel C drops treated children from the sample. The solitarious placebos remain null under both: no estimate is negative and significant at the 5 percent level on weight-for-height. Across the three outcomes and both panels, 4 of 132 placebo estimates are negative and significant at the 5 percent level, all on height-for-age or weight-for-age.

Permuting the treatment indicator within survey clusters, 1,000 times, places the 10 km estimate on weight-for-height outside the entire simulated null with a randomization p-value below 0.001. The simulated

standard deviation of the permuted estimate is 0.084, close to the clustered analytic standard error of 0.087. Running the same procedure on the solitarious placebo at 10 km, with the stage and vegetation conditions applied and before purging (−0.305), returns a randomization p-value of 0.482.

No crop in the field. The second test uses damaging-form records that arrive outside the susceptible window, when there is no crop in the field. Purged of overlap with the treatment, this placebo is 0.094 (0.052) at 30 km and is never negative and significant on weight-for-height at any radius (Panel A of Table 21). It is null at every radius under Panels B and C as well, and none of its 54 estimates in those two panels is negative and significant. Restricting the sample to cropping zones gives the same result up to 45 km (Table 22 in Appendix A): none of 27 estimates is significant at the 5 percent level. At 50 km, holding the crop-stage and treatment indicators fixed, off-season exposure is associated with lower height-for-age (−0.179) and weight-for-age (−0.235). Children in pastoral and fishing zones, who were often near locusts but had no crop to lose, show no effect either. Regressed on proximity at every radius, none of their estimates is significant and the weight-for-height estimates are positive at all nine (Figure 13 in Appendix A).

Locust-free vegetation decline. The third test investigates whether a loss of greenness during gestation predicts worse outcomes without any locust record. I treat every 16-day composite in a child's own cluster as a candidate event and apply the decline condition of equation (8) and the crop-stage condition of equation (5), without requiring a locust record. The 2014 round gives a direct version of this test, since no vegetation decline in the 2014 cohort's gestation windows can be locust damage. Declines that meet both conditions are common: 79 percent of 2014 children and 76 percent of 2022 children have at least one during gestation. I also use two other versions: a matched version, which treats the 20 percent of children in each round with the largest decline, the same share as treated children at 30 km, and a count of qualifying declines. In 2014, none of the twelve estimates across the three versions and four outcomes is significant at the 5 percent level, and one is significant at the 10 percent level: the count on death (−0.005) (Table 23 in Appendix A). The matched estimate on weight-for-height is 0.012 with a standard error of 0.048.

In 2022 the matched decline is associated with lower weight-for-height (−0.084, standard error 0.075), though not significantly, and only where locusts were recorded. Children with a matched decline and a locust record of any form within 50 km during gestation have weight-for-height 0.136 standard deviations lower, significant at the 10 percent level, while children with a matched decline and no record nearby show no difference (0.032, standard error 0.081; Table 24 in Appendix A). Adding the matched decline to the main specification leaves the treatment estimate at −0.204.[19] A cluster with no locust record is not necessarily a cluster without locusts, so the 2022 comparison is less clean than the 2014 one. In both rounds, a loss of greenness away from recorded locust activity does not predict lower weight-for-height.

5.2.3 Drought

A natural threat to identification is a shock to vegetation that coincides with the invasion but is not caused by it. The gestational SPEI control absorbs a smooth, cluster-specific measure of the climatic water balance,

[19] In 2022 the count of qualifying declines is also associated with lower weight-for-age (−0.030) and higher mortality (0.005), both significant at the 1 percent level, and lower weight-for-height (−0.019, significant at the 10 percent level) (Table 23). Split by whether a locust record fell within 50 km (Table 25), declines near locust activity predict lower weight-for-age (−0.033, significant at the 1 percent level), weight-for-height (−0.022, 5 percent level), and height-for-age (−0.027, 10 percent level). Declines away from locust activity predict lower weight-for-age (−0.025, 10 percent level), but not height-for-age (−0.024, standard error 0.026) or weight-for-height (−0.012, standard error 0.015). The count picks up some vegetation stress beyond locust damage on the slower-moving outcomes, while the weight-for-height result is confined to declines near recorded locust activity.

but it may not absorb all of it, and if treatment is systematically correlated with drought conditions that six-month index leaves unmeasured, the wasting estimate could partly reflect a drought effect misattributed to locusts.

Unconditionally, exposed children come from locations with a somewhat higher six-month SPEI on average than unexposed children, 1.251 versus 1.010 (Table 1). This reflects where events happened to occur. The gap reverses when the main specification compares children within the same cluster and birth cohort. Exposed children experienced somewhat drier gestational windows than their unexposed neighbors in the same cluster and cohort, -0.13 standard deviations on the six-month SPEI index, although the gap is not statistically significant ($p = 0.28$). This within-cluster gap runs in the direction that would make a drought confound plausible in principle. If drought lowers weight-for-height, a treated group with drier gestations would push the exposure coefficient in the negative direction this paper reports.

I remove the SPEI variable and re-estimate at every radius (see Table 26 in Appendix A). Since treated children had drier gestations within cluster, dropping SPEI should make the coefficient more negative if drought matters, and it does at every radius. The change is 0.004 to 0.007 standard deviations, about a tenth of a standard error.

Additionally, NDVI and SPEI are not interchangeable measures of the same underlying conditions. SPEI is a climatic water-balance index, precipitation minus potential evapotranspiration, standardized against a location's own history. It describes moisture stress acting on every hectare in a cluster roughly uniformly, whether or not a crop or pasture stands on it. NDVI records the greenness of the vegetation actually present on the ground, and it responds to anything that removes standing biomass abruptly, not only to a shortfall in the water available to grow it. A locust swarm defoliating a field over a matter of days produces a sharp, localized drop in NDVI on a timescale a slow-moving drought's drawdown on the water balance does not resemble. This rests on the assumption that no other large-scale consumer of vegetation, a wildlife grazing episode or an unrelated pest outbreak, stripped the same cluster in the same narrow window as a documented locust arrival. Kenya's rangelands do carry wild herbivores, but a defoliation event that happens to coincide with an event's own cluster and six-week window, systematically enough to generate the pattern in Section 5, would have to track the invasion's arrival dates and locations.

Because equation (8) tests a change in the NDVI anomaly over roughly six weeks around a specific arrival, a cluster that is dry for the whole gestational window does not qualify on that dryness alone. Both placebo groups sit in the same drought-affected places and the same drought years as the real treatment. If the decline condition picked up an ambient dry spell, they should satisfy it and predict the outcome at a similar rate. The test holds the climatic background fixed and asks whether a specific arrival was followed by a decline in vegetation. It does not depend on the direction of the seasonal SPEI comparison. [20]

5.2.4 Mortality

The main results condition on survival. A child who died before the survey has no anthropometric measurement and is removed by the sample filter that requires a valid z-score. If the shock also raised mortality among the most severely affected children, the ones who would have been measured as wasted, the survivors are positively selected and the estimates are biased toward zero.

Additionally, the weaker evidence on an effect on stunting in my analysis might be an artifact of mortality. Stunting is a long-run outcome while wasting is not. Many children might not have survived to an age where the long-run measure of health is observed.

[20] See Figure 14 in Appendix A for more detail.

I rebuild the birth panel without the anthropometry filter and re-run the identical exposure construction against it. The 2022 round has 8,213 births under this definition. Of these, 7,603 have valid measurements, 263 (3.2 percent) did not survive to interview, and 347 survived but were not measured. The mortality regressions use 8,057 of the 8,213 births: 27 are missing household head age or gestational SPEI, and a further 129 are the only child in their cluster and drop out of the cluster fixed-effects estimation. Table 27 regresses survival on exposure with the same controls, weights, fixed effects, and clustering. Nothing is significant at any radius. Estimates run from 0.003 to 0.009 against a baseline mortality rate of 3.2 percent, and the smallest p-value across the nine radii is 0.18. The three looser measures return the same null at the main radius.

Mortality selection does not explain the weight or height results, and the weight results do not extend to survival.

5.2.5 Pre-existing Conditions

If the exposed clusters were simply on a worse trajectory, a treatment constructed the same way in a pre-invasion period should reproduce the result. I test this using the 2014 DHS round, conducted six years before the invasion.

I shift the locust events back 93 months, the gap between the 2014 and 2022 KDHS fieldwork periods. This places the invasion's calendar pattern on the equivalent point in the 2014 cohort's gestation windows, and I re-evaluate the stage and NDVI conditions at the shifted dates using rainfall and NDVI from that period. Event locations and 2014 birth timing are unchanged. Whether a 2014 child's gestational window overlaps the shifted pattern then varies within cluster as real exposure does.

The four measures, nine radii and four outcomes (mortality and the three anthropometric measures) give 144 cells (Tables 28 to 31 in Appendix A), two are significant at the 5 percent level: the treatment measure on weight-for-height at 10 km (-0.277) and on mortality at 30 km (-0.021). Five (3.5 percent) are significant at the 10 percent level or better, fewer than chance alone would produce, and three of the five are negative mortality estimates under the treatment measure.

5.3 Potential Mechanisms

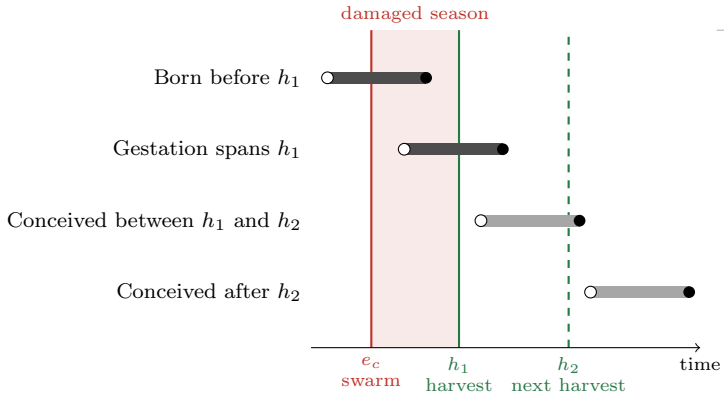
A lost harvest can reach a child in utero through three pathways. Households can cut consumption on news of the damage, before the harvest fails. Food can become unaffordable, because local prices rise or because the household's income falls. Or it can become unavailable, because the household has less of its own harvest to eat. This section tests each in turn. Section 5.3.1 uses the timing of gestation relative to the harvest that followed the swarm, together with local market prices. Section 5.3.2 compares zones by where their income comes from, and Section 5.3.3 compares zones and households by how much of their food they grow.

5.3.1 Prices

Conte et al. (2023) separate these pathways in Mali using timing. Mali grows its staples on a single calendar with one harvest a year, so a swarm during the growing season opens a common window in which every affected household knows the harvest is in trouble but the crop has not yet failed. Children in utero during that window are exposed to the news, and to any price rise it causes, but not to the lost harvest. They find that these children are harmed about as much as children who also lived through the crop failure, and that a rise in local prices during the window accounts for the effect.

Kenya has no common window. Its two rainy seasons give it two harvests a year, their timing differs across the 17 livelihood zones that grow crops, and the invasion lasted sixteen months. I therefore time each child from their own cluster's events. For each cluster c, let e_c be the date of the first event within 30 km, h_1 the next staple harvest after e_c in either season, and h_2 the next staple harvest more than 60 days after h_1.[21] Every child born on or after e_c falls in exactly one of four groups, shown in Figure 6: born before h_1, gestation spanning h_1, conceived between h_1 and h_2, and conceived after h_2. All other children, born before e_c or in clusters no event reached, form the comparison group, and the four groups enter one regression with the control set and fixed effects of column (3) of Table 2.

Only the first group separates a response to the news from the realized loss. A child born before h_1 was in utero during the swarm and the news of crop loss but was born before the harvest failed, so a negative coefficient for this group points to households cutting consumption in anticipation of the loss, or to a price rise the news causes. For the three later groups, the anticipated and realized loss arrive together.



The pregnancy overlaps	A negative coefficient points to
the swarm and the news, not the lost harvest	anticipated loss
the lost harvest	anticipated and realized loss
the months after the lost harvest	anticipated and realized loss
the next harvest onward	anticipated and realized loss

Figure 6: Timing groups and what each identifies. Each row is one example child, from conception (○) to birth (●), relative to the first event within 30 km of the child's cluster (e_c), the next staple harvest (h_1) and the harvest after that (h_2). Every child born on or after e_c falls in exactly one group. The two columns on the right give the period each group's pregnancy overlaps and what a negative coefficient for that group points to.

Because groups are timed from a cluster's first event, the invasion's later waves add two complications. Of the 907 clusters with an event within 30 km, 675 (74 percent) were hit in one season, 188 (21 percent) in two and 44 (5 percent) in three. Where a swarm returned, some children conceived after h_1 were in utero during a later swarm, which would make their harm look like lasting damage from the first one, and some children born before or around h_1 were infants when the later swarm arrived.[22] I give the first set, 225 children, their own indicator. As a check, I drop every cluster hit in more than one season, which removes both sets.

Table 4 reports the estimates. Children born before h_1 have weight-for-height 0.190 standard deviations lower (standard error 0.104), significant at the 10 percent level and similar to the pooled estimate. Children whose gestation spans the harvest are 0.135 lower, and the estimate is not significant. Children conceived

[21] Harvest dates are the ends of the harvest stage in the onset-shifted crop calendar of Section 4.1.2. The staple is maize, including maize grown with beans, where the cluster grows it, and otherwise sorghum, then beans, then any staple in the zone's calendar.

[22] Section 5.1.3 finds no effect of exposure in the first year of life (0.045, standard error 0.071), so exposure after birth is unlikely to bias these groups. Dropping clusters hit in more than one season is a precaution.

after the harvest show no harm. Their estimates are positive, 0.134 for those conceived between the two harvests and 0.177 for those conceived after h_2, and neither is significant. Children in utero during a later swarm are 0.228 lower, imprecisely estimated on 225 children. Without a separate indicator for them (Table 32 in Appendix A) the born-before estimate is −0.188, and in clusters hit in a single season (Table 33) it is −0.186, both significant at the 10 percent level. The harm appears before the harvest is lost and does not carry over to children conceived after it.

Column (2) adds the gestational maize price index: the inverse-distance-weighted maize price at nearby markets, relative to the cluster's own pre-2019 average, averaged over the child's gestation. The born-before estimate does not move (−0.194), and no other coefficient changes by more than 0.05. Local maize prices do not account for the harm to children born before the harvest.

Table 4: Timing of gestation relative to the first lost harvest, weight-for-height, 30 km radius

	(1)	(2)	Children
Born before h_1	−0.190*	−0.194*	625
	(0.104)	(0.102)	
Gestation spans h_1	−0.135	−0.150	1,183
	(0.106)	(0.106)	
Conceived between h_1 and h_2	0.134	0.130	572
	(0.098)	(0.095)	
Conceived after h_2	0.177	0.180	183
	(0.134)	(0.122)	
In utero during a later swarm	−0.228	−0.276	225
	(0.168)	(0.184)	
Maize price control	No	Yes	
Observations	7,416	7,416	

Each column is a single regression of the weight-for-height z-score on the five group indicators jointly, with the control set and fixed effects of column (3) of Table 2. Groups are defined in Figure 6 from the first event within 30 km of the child's cluster (e_c), the next staple harvest (h_1) and the next harvest more than 60 days after it (h_2). *In utero during a later swarm*: children conceived after h_1 who were exposed in utero to a later event; they are removed from the two "conceived" groups. The comparison group is children born before e_c or in clusters no event reached. Column (2) adds the gestational maize price index. Standard errors clustered on a half-degree spatial grid in parentheses. ***, **, * denote significance at 1, 5 and 10 percent.

I also test the price channel directly in the market data. The World Bank Real Time Prices series reports monthly maize and bean prices for 233 Kenyan markets from 2015 to 2022, of which 231 enter the regression. The test uses each market's own price series. A market is treated if an event ever fell within 30 km of it, and event time is measured in quarters relative to its first event. A market with no event within 30 km is control throughout its own series. Maize and beans are pooled into one regression, following Conte et al. (2023):

$$\log(\text{Price}_{m,t}) = \sum_j \theta_j \, T_m \times Q_j(t) + \mu_m + \delta_{s,t} + \xi_c \cdot t + \varepsilon_{m,t}, \tag{12}$$

where T_m is the market's treatment status, $Q_j(t)$ is the market's own relative quarter, μ_m is a market fixed effect, $\delta_{s,t}$ is a staple by year-month fixed effect, and $\xi_c \cdot t$ is a county-specific linear trend. Standard errors are clustered on market.

Figure 7 plots $\hat{\theta}_j$ for the 74 treated markets. Treated markets already run 3.6, 2.7 and 2.7 percent above control markets three, two and one quarter before their first event, all significant at the 5 percent level or better, and the gap narrows to 1.0 percent, insignificant, in the quarter of arrival. Because this premium

predates any swarm, it is a pre-existing difference between the markets that an event later reached and the rest, and it narrows once the swarms arrive. After arrival the coefficient turns negative: −1.6 percent one quarter after (not significant), and −2.9 and −4.1 percent two and three quarters after, both significant at the 1 percent level. For the median treated market the harvest follows its first event by one month, so these are also the quarters after the lost harvest. A crop failure should raise the price of the crop that failed. These coefficients run the wrong way for a supply shock. The 2020 to 2021 exposure window also overlaps the COVID-19 pandemic in Kenya, which brought its own supply disruption and a national maize price spike unrelated to any single locust event. Measuring event time from each market's own harvest boundary instead gives the same pattern (Figure 15 in Appendix A): treated markets run 2.2 to 3.2 percent above control markets before the boundary and 2.1 to 4.2 percent below them in the three quarters after it.

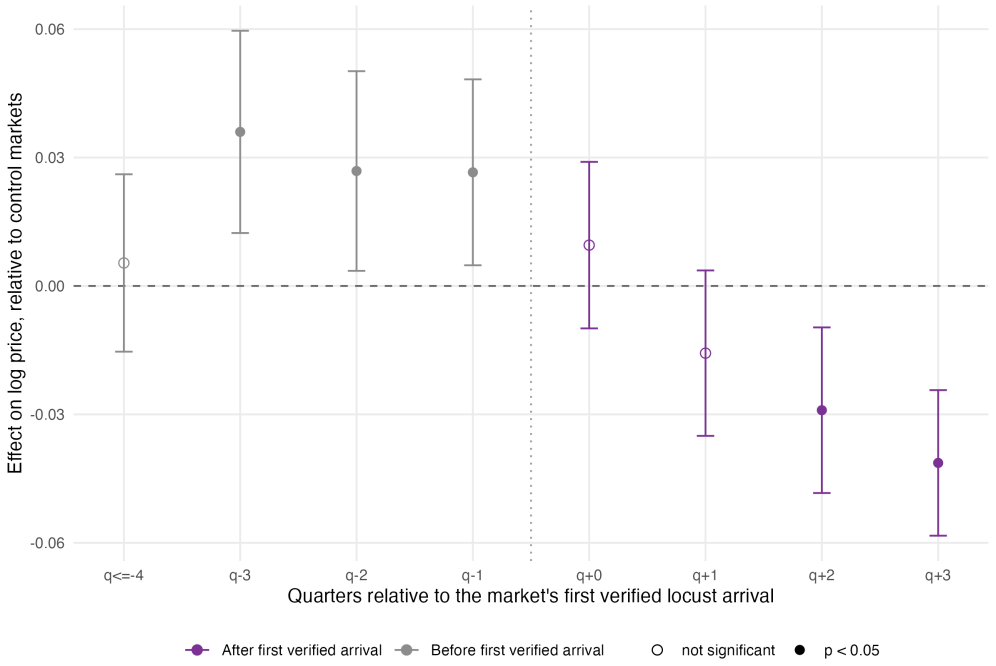


Figure 7: Market-level price event study at 30 km: log maize and bean prices in treated versus control markets, by quarter relative to each treated market's first event. Point estimates with 95 percent confidence intervals; standard errors clustered on market.

A swarm's arrival is not the first news of it. FAO issued warnings about the swarms in Ethiopia in October 2019 and in Kenya in January and November 2020. Figure 8 re-centres the same specification on each warning, comparing markets that an event later reached with the 158 markets no event reached. Prices did not rise after the October 2019 or January 2020 warnings. In the 55 markets later reached, prices fell relative to control markets over the following year, by 1.5 to 8.5 percent after the October 2019 warning and by up to 8.9 percent after the January 2020 warning. In the first two panels, prices in the markets later reached were already rising relative to control markets over the year before the warning, so those comparisons are less clean. After the November 2020 warning, prices in the 18 markets later reached rose 2.1 percent in the first quarter, significant at the 1 percent level, and the rise faded and was insignificant in the two quarters after. Most of the prices in these series are filled in by the RTP model, which adds noise to all three comparisons.

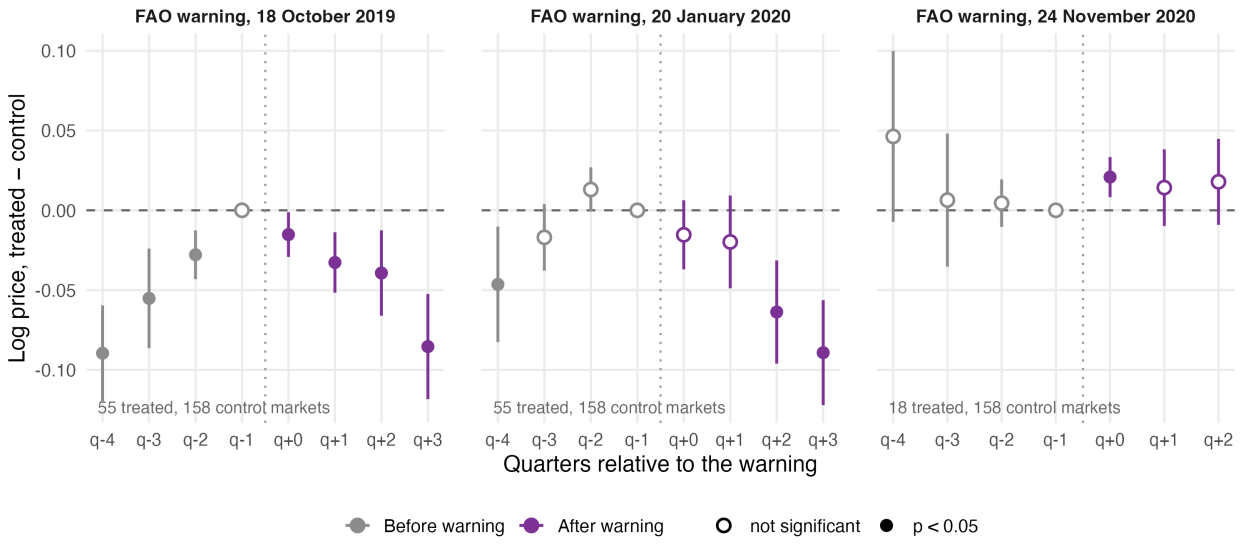


Figure 8: Market-level price event study around FAO's warnings. Log maize and bean prices in markets that an event later reached versus markets no event reached, by quarter relative to the warning, with the specification of equation (12). Point estimates with 95 percent confidence intervals; standard errors clustered on market.

Neither the swarms' arrival nor the news of them raised local prices in the markets the swarms reached, and controlling for local maize prices leaves the harm to children born before the harvest unchanged.

5.3.2 Income

If food became unaffordable because income fell, the harm should fall where earnings depend on crops. I estimate

$$y_{ict} = \beta\, T_i^{(r)} + \eta \left(T_i^{(r)} \times S_i\right) + X_{ic}'\gamma + \alpha_c + \delta_t + \varepsilon_{ict}, \tag{13}$$

where S_i is a characteristic of child i's livelihood zone. Every survey cluster lies in a single zone, so S_i is absorbed by the cluster fixed effect α_c. β is the effect of exposure in zones with $S_i = 0$ and $\beta + \eta$ the effect in zones with $S_i = 1$. I sort zones in two ways using the FEWS NET livelihood profiles (Famine Early Warning Systems Network (FEWS NET), 2011). By livelihood type, $S_i = 1$ in pastoral, agropastoral and fishing zones and $S_i = 0$ in cropping, mixed-farming and riverine zones. By main income source, $S_i = 1$ in zones where livestock or fishing is the largest stated share of household income and $S_i = 0$ in zones where crops are. If the pathway is lost crop income, β should be negative and η positive, with $\beta + \eta$ close to zero.

Table 5 reports both splits. By livelihood type, exposure lowers weight-for-height by 0.226 standard deviations in farming zones, significant at the 1 percent level, and the interaction offsets it (0.248, significant at the 5 percent level). The implied effect in herding and fishing zones is 0.021 (0.094). By main income source the pattern is the same, −0.234 in zones that earn mostly from crops and −0.043 (0.083) in zones that earn mostly from livestock or fish, though the interaction (0.191) is not significant ($p = 0.106$). The harm is concentrated where income comes from crops.

Table 5: Exposure by livelihood zone: where income comes from, weight-for-height, 30 km radius

	(1) By livelihood type	(2) By main income source
Exposure	-0.226^{***}	-0.234^{***}
	(0.068)	(0.075)
Exposure × herding or fishing	0.248^{**}	0.191
	(0.116)	(0.118)
Implied effect, herding or fishing ($\beta + \eta$)	0.021	−0.043
	(0.094)	(0.083)
Treated, farming / herding or fishing	1,263 / 232	1,185 / 310
Observations	7,416	7,416

Each column is one regression of equation (13) with the control set and fixed effects of column (3) of Table 2. Column (1): $S_i = 1$ in pastoral, agropastoral and fishing zones, 0 in cropping, mixed-farming and riverine zones. Column (2): $S_i = 1$ in zones where livestock or fishing is the main income source, 0 where crops are. No pastoral child is treated. Standard errors clustered on a half-degree spatial grid in parentheses. ***, **, * denote significance at 1, 5 and 10 percent.

Table 34 in Appendix A reports each treated zone's own exposure effect from one regression. Several estimates rest on one to three grid cells and are not reliable on their own. Interacting exposure with livelihood class relative to cropping zones gives the same picture for agropastoral zones and a larger but imprecise loss in riverine zones, which rests on 30 treated children (Tables 35 and 36 in Appendix A).

5.3.3 Food Shortage

If the harm ran through food becoming unavailable, it should be largest where households eat what they grow. I estimate equation (13) with $S_i = 1$ in zones where own production supplies half or more of the food households consume, according to the FEWS NET profiles. Exposure lowers weight-for-height by 0.216 standard deviations (0.077) in zones that buy most of their food, and the implied effect in zones that grow most of it is −0.185 (0.116), with an interaction of 0.030 (0.142) (Table 37 in Appendix A). Zones that grow most of their own food are not hit harder.

Household characteristics give the same answer. Interacting exposure, one regression at a time, with the household's wealth quintile, land ownership, rural residence and distance to a major road (Table 38 in Appendix A), no characteristic marks out a group that is hit harder. The wealth interaction is −0.045 (0.036), land ownership 0.080 (0.113) and distance to a road 0.012 (0.010), and rural households fare better, 0.215 (0.117), significant at the 10 percent level. These characteristics are correlated with one another (rural residence and wealth quintile correlate at −0.62), and when rural residence, land ownership, distance to a road and agricultural occupation are entered jointly, none is significant (Table 39). Among exposed children, those who are wasted have mothers with 3.2 fewer years of education and sit 0.74 of a wealth quintile lower than those who are not (Table 40 in Appendix A). Tables 41 and 42 repeat the comparison for underweight children and at 15 km.

The harm reaches children born before the lost harvest and is concentrated in zones where income comes from crops. Local prices did not rise, and zones that eat what they grow are not hit harder. This points to households cutting consumption in anticipation of the lost harvest, rather than to higher food prices or a realized shortage of food.

6 Discussion

In utero exposure to an event within 30 km lowered weight-for-height by 0.21 to 0.24 standard deviations and raised the probability of wasting by 1.8 to 3.4 percentage points, against a baseline of 7.1 percent. In the main specification, exposed children are 25 percent more likely to be wasted. The result depends on how exposure is measured. Locust presence alone within 30 km returns an insignificant -0.003 on the same children, outcome, controls and fixed effects. The harm is concentrated in cropping, mixed-farming and riverine zones, where exposure lowers weight-for-height by 0.23 standard deviations, and it reaches children born after the first event but before the next harvest. Both point to households cutting consumption in anticipation of the lost harvest.

The results have two implications for policy. First, relief that is targeted by proximity to locust reports reaches many households that lost nothing. Among the 7,603 children with valid anthropometry, a proximity rule at 30 km enrolls 4,070 children and the treatment measure enrolls 1,547. The 2,523 children the proximity rule adds live in drier, poorer and more remote households, and 39 percent of them live in pastoral zones where no crop was at risk. The records, crop calendars and vegetation data used here are all available during an upsurge, so relief can be targeted by damage. Second, the harm begins before the harvest fails and is concentrated where income comes from crops. Food or seed delivered after the harvest arrives too late for children who were in utero when the swarm came. Cash transfers or a temporary expansion of existing safety nets, paid to crop-dependent households as soon as damage is recorded, match the pathway these results identify.

The paper has several limitations. The mechanism tests are suggestive. The estimate for children born before the harvest is significant only at the 10 percent level and rests on 625 children, and the zone comparison rests on 232 treated children in herding and fishing zones, whose zones differ from farming zones in other ways as well. The DHS does not record household income or food spending over the exposure period, so anticipatory consumption cuts are inferred from where and when the harm appears and are not observed directly. Most prices in the market data are model estimates, treated markets were already on different paths before the FAO warnings, and the invasion overlapped with a variety of global crises that make these modeled prices noisy.

Several extensions would test and build on these results. Phone or panel surveys fielded during an upsurge could record household income, food spending, asset sales and migration month by month, and would test directly whether households cut consumption before the harvest fails. Later survey rounds and school records could show whether the health deficits from this shock persist. The 2020 to 2021 invasion also coincided with the COVID-19 pandemic and with conflict and drought in parts of the region, and separating market effects of overlapping shocks remains an open problem here, as in many other settings. The exposure construction also carries over to other regions. The 2019 to 2021 upsurge also struck Ethiopia, Somalia and Uganda, where FEWS NET publishes the same livelihood calendars, and the measure can be built there wherever geocoded survey data cover the invasion period. Other crop pests, droughts, floods and hailstorms also do damage that depends on the crop's growth stage and leaves a trace in vegetation.

Finally, spatial data, particularly on weather, have become widely accessible and inexpensive. I hope this work motivates researchers to continue developing methods that measure the consequences of climate change accurately, so that policy can mitigate them, especially for the most vulnerable.

A Supplemental Tables and Figures

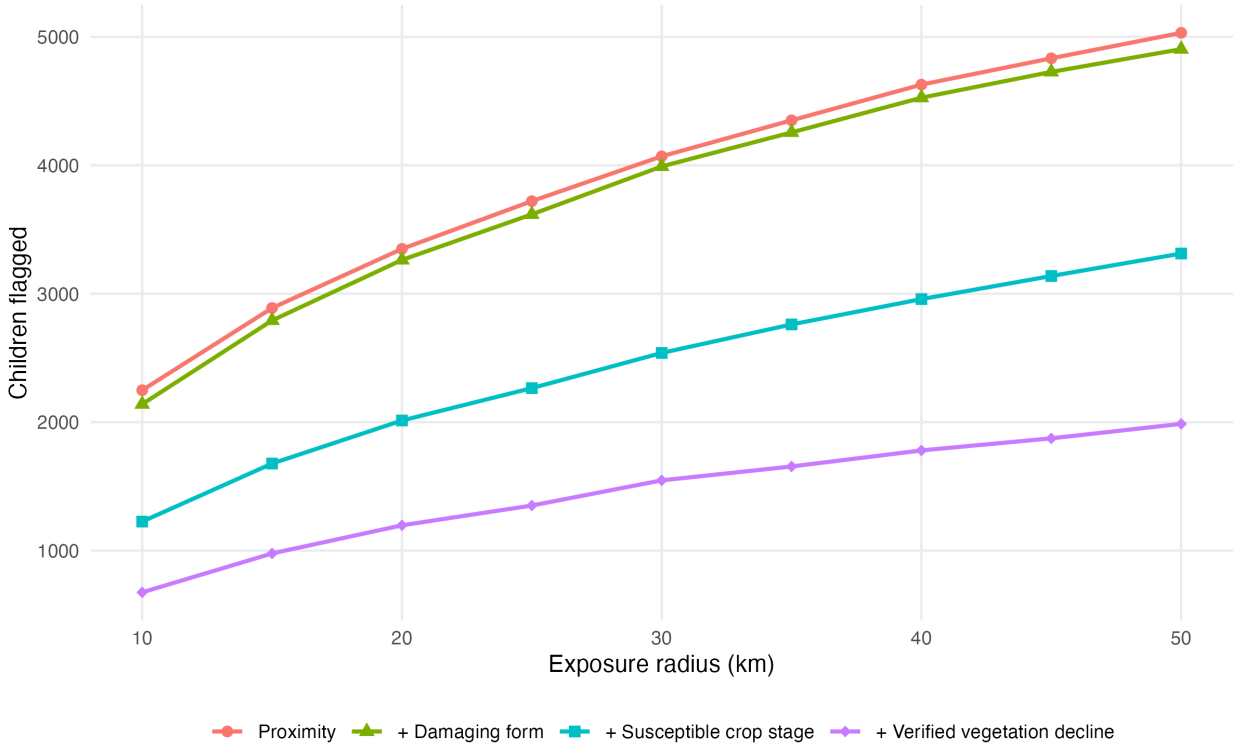


Figure 9: Children flagged as exposed, by exposure definition and radius. The four definitions nest, so every child flagged by a stricter measure is also flagged by every looser measure above it. Reproduces the counts of Table 11.

Table 6: Sensitivity of the exposure construction: weight-for-height z-score

	10 km		30 km		50 km	
Variant	Estimate	Treated	Estimate	Treated	Estimate	Treated
Main specification	−0.340***	647	−0.207***	1,495	−0.121**	1,918
	(0.087)		(0.064)		(0.060)	
NDVI decline threshold (main: $\theta = 0.25$)						
$\theta = 0$	−0.303***	760	−0.195***	1,753	−0.169**	2,247
	(0.075)		(0.066)		(0.066)	
$\theta = 0.1$	−0.317***	719	−0.204***	1,646	−0.158***	2,125
	(0.086)		(0.059)		(0.051)	
$\theta = 0.5$	−0.362***	504	−0.247***	1,207	−0.171***	1,580
	(0.086)		(0.065)		(0.054)	
$\theta = 0.75$	−0.318***	366	−0.194***	922	−0.130**	1,209
	(0.117)		(0.066)		(0.055)	
$\theta = 1$	−0.365***	254	−0.205***	657	−0.153***	893
	(0.118)		(0.054)		(0.048)	
$\theta = 1.5$	−0.311*	126	−0.214***	282	−0.091	388
	(0.166)		(0.078)		(0.085)	
NDVI composites before / after the event (main: 2 / 2)						
1 before / 1 after	−0.286***	605	−0.214***	1,559	−0.125**	2,089
	(0.078)		(0.075)		(0.063)	
3 before / 3 after	−0.292***	659	−0.168***	1,558	−0.114**	2,000
	(0.062)		(0.058)		(0.051)	
2 before / 1 after	−0.354***	521	−0.241***	1,351	−0.183***	1,813
	(0.077)		(0.065)		(0.056)	
2 before / 3 after	−0.272***	699	−0.165***	1,629	−0.094*	2,084
	(0.076)		(0.056)		(0.051)	
1 before / 2 after	−0.225***	709	−0.143**	1,694	−0.085*	2,206
	(0.078)		(0.065)		(0.051)	
Crop calendar (main: onset-shifted, planting and standing crop susceptible)						
Standing crop only	−0.426***	309	−0.225**	800	−0.196***	1,014
	(0.134)		(0.091)		(0.067)	
Unshifted FEWS NET calendar	−0.211***	618	−0.064	1,553	0.043	2,094
	(0.069)		(0.076)		(0.077)	
Stage windows widened 15 days	−0.280***	824	−0.110	1,929	−0.071	2,497
	(0.072)		(0.069)		(0.069)	
Stage windows narrowed 15 days	−0.325*	251	−0.184*	725	−0.122	1,026
	(0.167)		(0.110)		(0.084)	
Gestation length (main: 266 days)						
280 days	−0.318***	685	−0.181***	1,574	−0.092	2,011
	(0.080)		(0.054)		(0.056)	
Sampling weights (main: DHS individual weight)						
Winsorized at 1st / 99th percentiles	−0.321***	647	−0.186***	1,495	−0.102*	1,918
	(0.079)		(0.054)		(0.052)	
Unweighted	−0.270***	647	−0.118**	1,495	−0.057	1,918
	(0.071)		(0.050)		(0.048)	

Each cell is a separate regression of the weight-for-height z-score on a rebuilt verified-exposure indicator, with the control set and fixed effects of column (3) of Table 2. Each variant changes one condition and holds the others at their main values: $\theta = 0.25$, two NDVI composites before and after the event, the onset-shifted calendar with planting and the standing crop susceptible, a 266-day gestation, and the DHS individual sampling weight. The standing-crop variant counts only the harvest (standing crop) stage as susceptible. The unshifted variant uses the transcribed FEWS NET stage dates without the rainfall-onset shift. The widened and narrowed variants move both ends of every planting and standing-crop interval out or in by 15 days. The 280-day variant extends the gestation window 14 days earlier. The winsorized variant caps the sampling weight at its 1st and 99th percentiles, and the unweighted variant drops it; both keep the main treatment indicator. Standard errors clustered on a half-degree spatial grid in parentheses. ***, **, * denote significance at 1, 5 and 10 percent. The estimation sample is 7,416 children in every cell. Treated counts are children in the estimation sample.

Table 7: Sensitivity of the exposure construction: weight-for-age z-score

Variant	10 km	30 km	50 km
Main specification	−0.350***	−0.187***	−0.121**
	(0.106)	(0.064)	(0.058)
NDVI decline threshold (main: $\theta = 0.25$)			
$\theta = 0$	−0.303***	−0.180***	−0.141**
	(0.096)	(0.055)	(0.055)
$\theta = 0.1$	−0.317***	−0.178***	−0.136***
	(0.106)	(0.053)	(0.045)
$\theta = 0.5$	−0.371***	−0.209***	−0.169***
	(0.114)	(0.070)	(0.057)
$\theta = 0.75$	−0.332**	−0.170**	−0.135**
	(0.152)	(0.069)	(0.061)
$\theta = 1$	−0.364**	−0.172***	−0.146***
	(0.152)	(0.059)	(0.054)
$\theta = 1.5$	−0.359***	−0.163*	−0.051
	(0.133)	(0.088)	(0.078)
NDVI composites before / after the event (main: 2 / 2)			
1 before / 1 after	−0.296***	−0.221***	−0.139*
	(0.096)	(0.074)	(0.074)
3 before / 3 after	−0.280***	−0.151***	−0.114**
	(0.082)	(0.056)	(0.047)
2 before / 1 after	−0.368***	−0.234***	−0.192***
	(0.091)	(0.071)	(0.061)
2 before / 3 after	−0.285***	−0.148***	−0.079
	(0.101)	(0.056)	(0.052)
1 before / 2 after	−0.233**	−0.153**	−0.098
	(0.098)	(0.063)	(0.061)
Crop calendar (main: onset-shifted, planting and standing crop susceptible)			
Standing crop only	−0.430**	−0.184**	−0.159**
	(0.168)	(0.092)	(0.072)
Unshifted FEWS NET calendar	−0.189***	−0.080	0.010
	(0.063)	(0.069)	(0.069)
Stage windows widened 15 days	−0.280***	−0.133*	−0.101
	(0.099)	(0.070)	(0.067)
Stage windows narrowed 15 days	−0.382**	−0.190*	−0.167**
	(0.192)	(0.103)	(0.078)
Gestation length (main: 266 days)			
280 days	−0.317***	−0.156***	−0.092*
	(0.080)	(0.051)	(0.050)
Sampling weights (main: DHS individual weight)			
Winsorized at 1st / 99th percentiles	−0.321***	−0.166***	−0.100**
	(0.089)	(0.054)	(0.048)
Unweighted	−0.229***	−0.090*	−0.039
	(0.065)	(0.046)	(0.040)

Each cell is a separate regression of the weight-for-age z-score on a rebuilt verified-exposure indicator, with the control set and fixed effects of column (3) of Table 2. Each variant changes one condition and holds the others at their main values: $\theta = 0.25$, two NDVI composites before and after the event, the onset-shifted calendar with planting and the standing crop susceptible, a 266-day gestation, and the DHS individual sampling weight. The standing-crop variant counts only the harvest (standing crop) stage as susceptible. The unshifted variant uses the transcribed FEWS NET stage dates without the rainfall-onset shift. The widened and narrowed variants move both ends of every planting and standing-crop interval out or in by 15 days. The 280-day variant extends the gestation window 14 days earlier. The winsorized variant caps the sampling weight at its 1st and 99th percentiles, and the unweighted variant drops it; both keep the main treatment indicator. Standard errors clustered on a half-degree spatial grid in parentheses. ***, **, * denote significance at 1, 5 and 10 percent. The estimation sample is 7,421 children in every cell. Treated counts differ by at most two children from those in Table 6.

Table 8: Sensitivity of the exposure construction: height-for-age z-score

Variant	10 km	30 km	50 km
Main specification	−0.181*	−0.063	−0.057
	(0.103)	(0.063)	(0.053)
NDVI decline threshold (main: $\theta = 0.25$)			
$\theta = 0$	−0.156	−0.078	−0.038
	(0.104)	(0.052)	(0.041)
$\theta = 0.1$	−0.155	−0.054	−0.041
	(0.104)	(0.052)	(0.042)
$\theta = 0.5$	−0.194	−0.051	−0.077
	(0.129)	(0.063)	(0.055)
$\theta = 0.75$	−0.187	−0.042	−0.072
	(0.151)	(0.067)	(0.064)
$\theta = 1$	−0.182	−0.047	−0.071
	(0.151)	(0.065)	(0.064)
$\theta = 1.5$	−0.275	−0.004	0.013
	(0.183)	(0.149)	(0.129)
NDVI composites before / after the event (main: 2 / 2)			
1 before / 1 after	−0.164	−0.105**	−0.080
	(0.105)	(0.053)	(0.066)
3 before / 3 after	−0.125	−0.041	−0.043
	(0.104)	(0.055)	(0.044)
2 before / 1 after	−0.193*	−0.098	−0.106*
	(0.101)	(0.064)	(0.058)
2 before / 3 after	−0.166	−0.049	−0.014
	(0.117)	(0.055)	(0.047)
1 before / 2 after	−0.123	−0.075	−0.055
	(0.106)	(0.048)	(0.061)
Crop calendar (main: onset-shifted, planting and standing crop susceptible)			
Standing crop only	−0.197	−0.034	−0.018
	(0.156)	(0.085)	(0.071)
Unshifted FEWS NET calendar	−0.038	−0.051	−0.031
	(0.068)	(0.057)	(0.055)
Stage windows widened 15 days	−0.127	−0.093	−0.087
	(0.099)	(0.066)	(0.062)
Stage windows narrowed 15 days	−0.234	−0.083	−0.122**
	(0.155)	(0.075)	(0.059)
Gestation length (main: 266 days)			
280 days	−0.161**	−0.051	−0.050
	(0.067)	(0.055)	(0.045)
Sampling weights (main: DHS individual weight)			
Winsorized at 1st / 99th percentiles	−0.153*	−0.051	−0.044
	(0.089)	(0.059)	(0.049)
Unweighted	−0.054	−0.008	0.001
	(0.065)	(0.054)	(0.048)

Each cell is a separate regression of the height-for-age z-score on a rebuilt verified-exposure indicator, with the control set and fixed effects of column (3) of Table 2. Each variant changes one condition and holds the others at their main values: $\theta = 0.25$, two NDVI composites before and after the event, the onset-shifted calendar with planting and the standing crop susceptible, a 266-day gestation, and the DHS individual sampling weight. The standing-crop variant counts only the harvest (standing crop) stage as susceptible. The unshifted variant uses the transcribed FEWS NET stage dates without the rainfall-onset shift. The widened and narrowed variants move both ends of every planting and standing-crop interval out or in by 15 days. The 280-day variant extends the gestation window 14 days earlier. The winsorized variant caps the sampling weight at its 1st and 99th percentiles, and the unweighted variant drops it; both keep the main treatment indicator. Standard errors clustered on a half-degree spatial grid in parentheses. ***, **, * denote significance at 1, 5 and 10 percent. The estimation sample is 7,424 children in every cell. Treated counts differ by at most two children from those in Table 6.

Table 9: Treatment built from cash crops, weight-for-height

	(1) Staple crops	(2) Cash crops	(3) Both
Panel A: 10 km			
Staple-crop treatment	−0.340***		−0.326***
	(0.087)		(0.093)
Cash-crop treatment		−0.317*	−0.037
		(0.164)	(0.167)
Treated: staple / cash / both		647 / 294 / 261	
Panel B: 30 km			
Staple-crop treatment	−0.207***		−0.128**
	(0.064)		(0.052)
Cash-crop treatment		−0.283***	−0.196*
		(0.102)	(0.102)
Treated: staple / cash / both		1,495 / 813 / 649	
Panel C: 50 km			
Staple-crop treatment	−0.121**		−0.032
	(0.060)		(0.060)
Cash-crop treatment		−0.242***	−0.221**
		(0.078)	(0.092)
Treated: staple / cash / both		1,918 / 1,080 / 849	
Observations	7,416	7,416	7,416

Outcome is the weight-for-height z-score. Column (1) is the main treatment, where staple crops determine whether an arrival met a susceptible crop stage. Column (2) rebuilds the treatment with cash crops (tea, coffee, pyrethrum, cotton, tobacco and sugar cane) in place of staples, keeping the damaging-form and NDVI-decline conditions. Column (3) enters both indicators. Treated counts are children in the estimation sample; "both" counts children treated under both measures. All columns share the control set and fixed effects of column (3) of Table 2. Standard errors clustered on a half-degree spatial grid in parentheses. ***, **, * denote significance at 1, 5 and 10 percent.

Table 10: Characteristics of disputed children, 30 km radius

	Agreed	Disputed	Difference	p
Mother's education (years)	9.46	6.15	−3.31	< 0.001
Wealth quintile	3.04	2.38	−0.66	< 0.001
Distance to major road (km)	3.80	10.25	6.45	< 0.001
Owns land	0.537	0.380	−0.157	< 0.001
Owns livestock	0.655	0.675	0.020	0.201
Rural	0.621	0.658	0.037	0.017
Gestational SPEI (6 month)	1.251	0.872	−0.379	< 0.001
Gestational NDVI z-score	0.736	0.839	0.103	< 0.001
Livelihood zone shares				
Cropping	0.829	0.394	−0.435	
Pastoral	0.000	0.390	0.390	
Agropastoral	0.148	0.166	0.018	
Riverine	0.019	0.037	0.017	
Fishing	0.004	0.013	0.010	

Agreed denotes children flagged by both the proximity and verified measures (N = 1,547). Disputed denotes children flagged by proximity only (N = 2,523). Sample is the 7,603 children with valid anthropometry. p-values from two-sample t-tests.

Table 11: Children flagged as exposed, by measure and radius

Radius (km)	Proximity	+ form	+ stage	+ verified	Disputed	Jaccard
10	2,249	2,140	1,226	675	1,574	0.30
15	2,889	2,792	1,678	978	1,911	0.34
20	3,349	3,263	2,013	1,197	2,152	0.36
25	3,721	3,618	2,265	1,351	2,370	0.36
30	4,070	3,991	2,539	1,547	2,523	0.38
35	4,350	4,256	2,761	1,655	2,695	0.38
40	4,628	4,526	2,958	1,780	2,848	0.38
45	4,833	4,727	3,138	1,874	2,959	0.39
50	5,031	4,904	3,313	1,987	3,044	0.39

Sample is the 7,603 children with valid anthropometry. Measures nest, so no child is flagged by a stricter measure and not by a looser one. Disputed denotes children flagged by proximity and not by the verified measure. Jaccard is the overlap between the proximity and verified lists.

Table 12: Verified exposure and the treatment of missing mother's height

	Weight-for-height z-score	Height-for-age z-score	Weight-for-age z-score	Treated
10 km				
Main specification	−0.340***	−0.181*	−0.350***	647
	(0.087)	(0.103)	(0.106)	
Adding a missing-height indicator	−0.341***	−0.181*	−0.350***	647
	(0.088)	(0.101)	(0.105)	
Dropping mother's height	−0.339***	−0.159*	−0.338***	647
	(0.086)	(0.092)	(0.097)	
Measured mothers only	−0.310***	−0.111	−0.281***	289
	(0.110)	(0.124)	(0.107)	
30 km				
Main specification	−0.207***	−0.063	−0.187***	1,495
	(0.064)	(0.063)	(0.064)	
Adding a missing-height indicator	−0.207***	−0.065	−0.187***	1,495
	(0.064)	(0.063)	(0.065)	
Dropping mother's height	−0.206***	−0.046	−0.177***	1,495
	(0.063)	(0.057)	(0.062)	
Measured mothers only	−0.215***	0.010	−0.160**	679
	(0.071)	(0.092)	(0.077)	
50 km				
Main specification	−0.121**	−0.057	−0.121**	1,918
	(0.060)	(0.053)	(0.058)	
Adding a missing-height indicator	−0.120**	−0.058	−0.121**	1,918
	(0.061)	(0.053)	(0.058)	
Dropping mother's height	−0.120**	−0.033	−0.107*	1,918
	(0.060)	(0.048)	(0.056)	
Measured mothers only	−0.131*	0.016	−0.115	864
	(0.078)	(0.081)	(0.078)	

Each cell is a separate regression with the control set and fixed effects of column (3) of Table 2. The main specification fills a missing mother's height with the median among measured women. The second row adds an indicator for a filled value, the third drops mother's height from the controls, and the fourth restricts the sample to children whose mother was measured (3,513 for weight-for-height, 3,516 for height-for-age and 3,514 for weight-for-age, against 7,416, 7,424 and 7,421). Treated counts are for the weight-for-height estimation sample. Standard errors clustered on a half-degree spatial grid in parentheses. ***, **, * denote significance at 1, 5 and 10 percent.

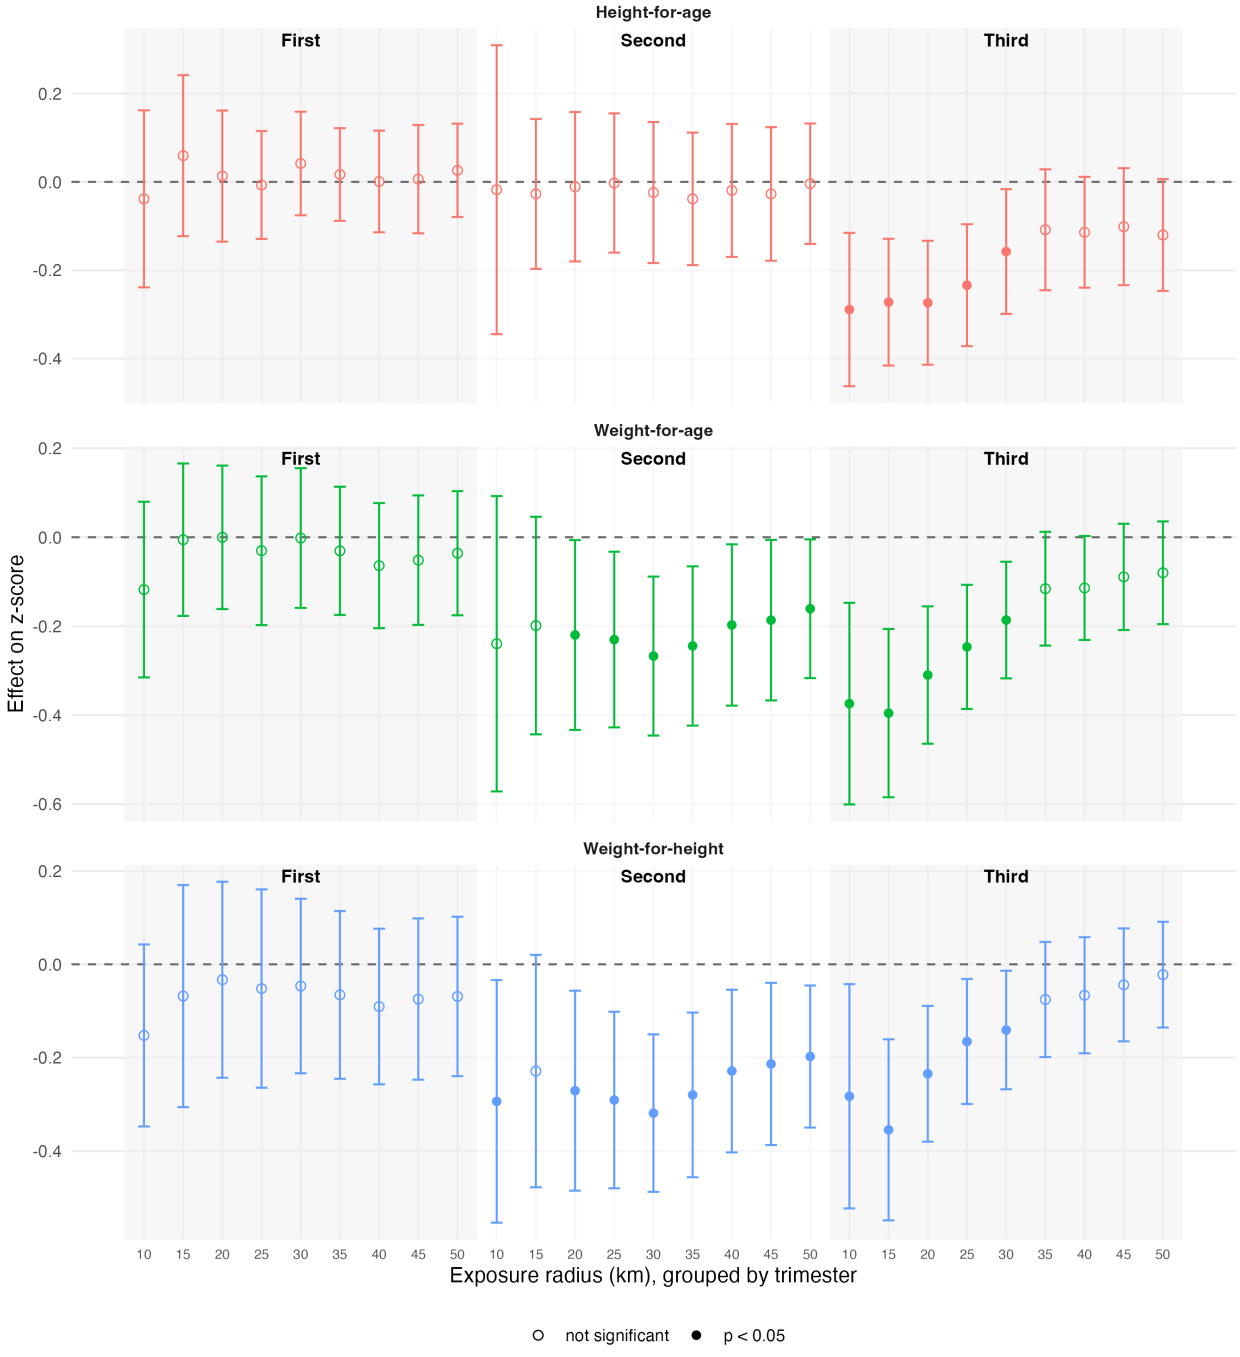


Figure 10: Trimester of exposure, all three outcomes. Regressions on an indicator for verified exposure occurring during that trimester, with the control set and fixed effects of column (3) of Table 2. Point estimates with 95 percent confidence intervals; standard errors clustered on a half-degree spatial grid.

Table 13: Alternative Anthropometric Outcomes, 30 km radius

	Height-for-age z-score	Weight-for-age z-score	Weight-for-height z-score	Stunted	Underweight	Wasted
In utero exposure	−0.063	−0.187***	−0.207***	−0.000	−0.001	0.018**
	(0.063)	(0.064)	(0.064)	(0.016)	(0.012)	(0.008)
Baseline prevalence				0.186	0.122	0.071
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Cluster and birth-month FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	7,424	7,421	7,416	7,424	7,421	7,416
R^2	0.390	0.378	0.353	0.319	0.275	0.261

Standard errors clustered on a half-degree spatial grid in parentheses. ***, **, * denote significance at 1, 5 and 10 percent. Each column is a separate regression sharing the control set and fixed effects of column (3) of Table 2. Baseline prevalence is the share of unexposed children below the relevant threshold.

Table 14: In utero locust exposure and child stunting, 30 km radius

	Height-for-age z-score				Stunted	
	(1)	(2)	(3)	(4)	(5)	(6)
In utero exposure	−0.025	−0.039	−0.063	−0.076	−0.000	0.013
	(0.054)	(0.056)	(0.063)	(0.059)	(0.016)	(0.017)
Child controls	No	Yes	Yes	Yes	Yes	Yes
Family controls	No	No	Yes	Yes	Yes	Yes
County time trend	No	No	No	Yes	No	Yes
Observations	7,451	7,451	7,424	7,424	7,424	7,424
R^2	0.335	0.355	0.390	0.396	0.319	0.327

None of the six columns is significant at conventional levels. All regressions include birth cohort and enumeration area fixed effects. Standard errors clustered on a half-degree spatial grid in parentheses. ***, **, * denote significance at 1, 5 and 10 percent. All regressions weight by the DHS individual sampling weight. Child controls are sex, birth order, first-born status, twin status, and the preceding birth interval. Family controls are mother's education, mother's height, household head age and sex, wealth-quintile indicators, and six-month gestational SPEI. County time trend denotes county-specific linear trends in calendar month of birth. The outcome in columns (5) and (6) is an indicator for a height-for-age z-score below −2, the clinical threshold for stunting.

Table 15: In utero locust exposure and child underweight status, 30 km radius

	Weight-for-age z-score				Underweight	
	(1)	(2)	(3)	(4)	(5)	(6)
In utero exposure	−0.173***	−0.185***	−0.187***	−0.216***	−0.001	0.017
	(0.060)	(0.061)	(0.064)	(0.068)	(0.012)	(0.013)
Child controls	No	Yes	Yes	Yes	Yes	Yes
Family controls	No	No	Yes	Yes	Yes	Yes
County time trend	No	No	No	Yes	No	Yes
Observations	7,448	7,448	7,421	7,421	7,421	7,421
R^2	0.343	0.361	0.378	0.385	0.275	0.291

The point estimate is negative and significant in columns (1)-(4) but the binary (underweight) columns (5)-(6) are not significant, consistent with dilution similar to Section 5's weight-for-age discussion at the main radius. All regressions include birth cohort and enumeration area fixed effects. Standard errors clustered on a half-degree spatial grid in parentheses. ***, **, * denote significance at 1, 5 and 10 percent. All regressions weight by the DHS individual sampling weight. Child controls are sex, birth order, first-born status, twin status, and the preceding birth interval. Family controls are mother's education, mother's height, household head age and sex, wealth-quintile indicators, and six-month gestational SPEI. County time trend denotes county-specific linear trends in calendar month of birth. The outcome in columns (5) and (6) is an indicator for a weight-for-age z-score below −2, the clinical threshold for underweight.

Table 16: Dose response among verified-exposed children: number of damage events and height-for-age

Dependent variable: Height-for-age z-score	Linear	Log	Capped at D_9
Verified damage events	−0.021**		−0.036***
	(0.008)		(0.008)
Log(1 + verified damage events)		−0.295***	
		(0.078)	
Observations	1,261	1,261	1,261

Standard errors clustered on a half-degree spatial grid in parentheses. ***, **, * denote significance at 1, 5 and 10 percent. D_9 is the 90th percentile of verified damage events among exposed children (26). Sample restricted to children verified-exposed at the main 30 km radius; of the 1,545 verified-exposed children with complete data, those who are the only such child in their survey cluster or birth cohort drop out of the fixed-effects estimation, which leaves 1,261. All columns include the control set and the survey-cluster and birth-cohort fixed effects of column (3) of Table 14. I control for the number of RAMSES ground visits during gestation, counting every recorded visit within 30 km, of any outcome, during the child's gestation window, to separate the effect of confirmed damage from the effect of monitoring intensity in that area during that period.

Table 17: Dose response among verified-exposed children: number of damage events and weight-for-age

Dependent variable: Weight-for-age z-score	Linear	Log	Capped at D_9
Verified damage events	−0.021**		−0.041***
	(0.009)		(0.009)
Log(1 + verified damage events)		−0.333***	
		(0.072)	
Observations	1,260	1,260	1,260

Standard errors clustered on a half-degree spatial grid in parentheses. ***, **, * denote significance at 1, 5 and 10 percent. D_9 is the 90th percentile of verified damage events among exposed children (26). Sample restricted to children verified-exposed at the main 30 km radius; of the 1,544 verified-exposed children with complete data, those who are the only such child in their survey cluster or birth cohort drop out of the fixed-effects estimation, which leaves 1,260. All columns include the control set and the survey-cluster and birth-cohort fixed effects of column (3) of Table 15. I control for the number of RAMSES ground visits during gestation, counting every recorded visit within 30 km, of any outcome, during the child's gestation window, to separate the effect of confirmed damage from the effect of monitoring intensity in that area during that period.

Table 18: In utero exposure by age at measurement, 30 km radius

Age at measurement	Weight-for-height z-score	Height-for-age z-score	Weight-for-age z-score	Treated
0 to 11 months	−0.529***	−0.249	−0.562***	483
	(0.180)	(0.228)	(0.203)	
12 to 23 months	−0.249***	−0.075	−0.214**	915
	(0.088)	(0.127)	(0.101)	
24 months or older	0.363	−0.673**	−0.085	149
	(0.520)	(0.304)	(0.394)	

Each cell is a separate regression on children in the row's age group, sharing the control set and fixed effects of column (3) of Table 2. Observations are 2,371, 3,007, and 711 for weight-for-height, 2,376, 3,012, and 712 for height-for-age, and 2,375, 3,009, and 712 for weight-for-age, by age group. Treated counts the children with valid anthropometry in each age group flagged by the verified measure at 30 km. Standard errors clustered on a half-degree spatial grid in parentheses. ***, **, * denote significance at 1, 5 and 10 percent.

Table 19: In utero and postnatal exposure, weight-for-height, 30 km radius

	(1)	(2)	(3)	(4)
In utero exposure	−0.207***	−0.191***	−0.191***	
	(0.064)	(0.070)	(0.071)	
Postnatal exposure, first 12 months		0.045		0.078
		(0.071)		(0.085)
Postnatal exposure, months 0 to 6			0.051	
			(0.069)	
Postnatal exposure, months 6 to 12			0.031	
			(0.079)	
Sample	All	All	All	Not exposed in utero
Observations	7,416	7,416	7,416	5,738

Outcome is the weight-for-height z-score. Postnatal exposure is a verified damage event within 30 km of the child's cluster after birth, within the stated months and before the interview. All columns share the control set and fixed effects of column (3) of Table 2. Column (4) restricts to the 5,921 children in the estimation sample who were not exposed in utero; 183 of them are the only such child in their survey cluster or birth cohort and drop out of the cluster fixed-effects estimation, which leaves 5,738. Standard errors clustered on a half-degree spatial grid in parentheses. ***, **, * denote significance at 1, 5 and 10 percent.

Table 20: Treatment by Radius and Health Outcome

Radius	Weight-for-height z-score	Weight-for-age z-score	Height-for-age z-score	Wasted	Treated
10 km	−0.340***	−0.350***	−0.181*	0.028**	647
	(0.087)	(0.106)	(0.103)	(0.012)	
15 km	−0.304***	−0.274***	−0.106**	0.031***	940
	(0.114)	(0.092)	(0.051)	(0.008)	
20 km	−0.234***	−0.232***	−0.125**	0.023***	1,153
	(0.073)	(0.073)	(0.061)	(0.009)	
25 km	−0.218***	−0.219***	−0.110**	0.025***	1,303
	(0.066)	(0.062)	(0.054)	(0.009)	
30 km	−0.207***	−0.187***	−0.063	0.018**	1,495
	(0.064)	(0.064)	(0.063)	(0.008)	
35 km	−0.175***	−0.159**	−0.054	0.017**	1,599
	(0.066)	(0.066)	(0.059)	(0.009)	
40 km	−0.155**	−0.153**	−0.061	0.015*	1,717
	(0.066)	(0.063)	(0.055)	(0.008)	
45 km	−0.145**	−0.149**	−0.072	0.014	1,807
	(0.066)	(0.064)	(0.057)	(0.008)	
50 km	−0.121**	−0.121**	−0.057	0.012	1,918
	(0.060)	(0.058)	(0.053)	(0.009)	

Standard errors clustered on a half-degree spatial grid in parentheses. ***, **, * denote significance at 1, 5 and 10 percent. Each cell is a separate regression sharing the control set and fixed effects of column (3) of Table 2. Sample sizes are 7,416 for weight-for-height and wasting, 7,421 for weight-for-age, and 7,424 for height-for-age. Treated counts are treated children in the weight-for-height estimation sample.

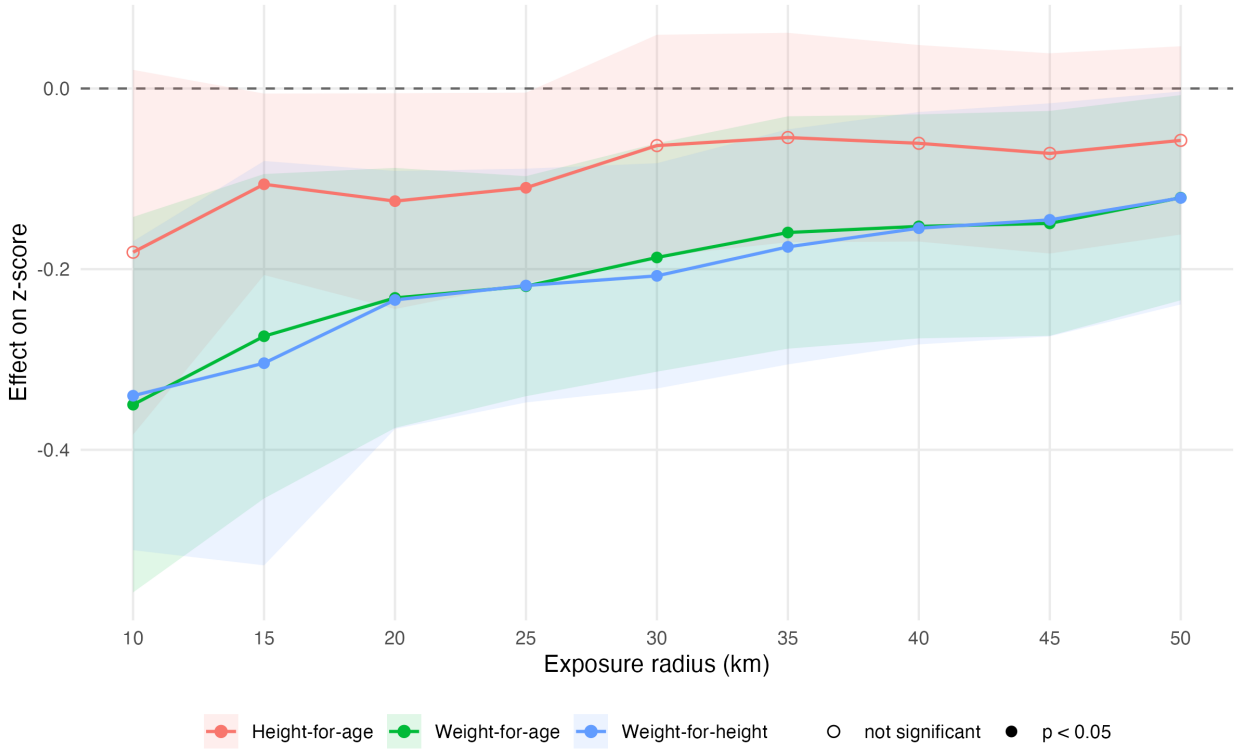


Figure 11: Estimate against exposure radius, all three outcomes, with 95 percent confidence bands.

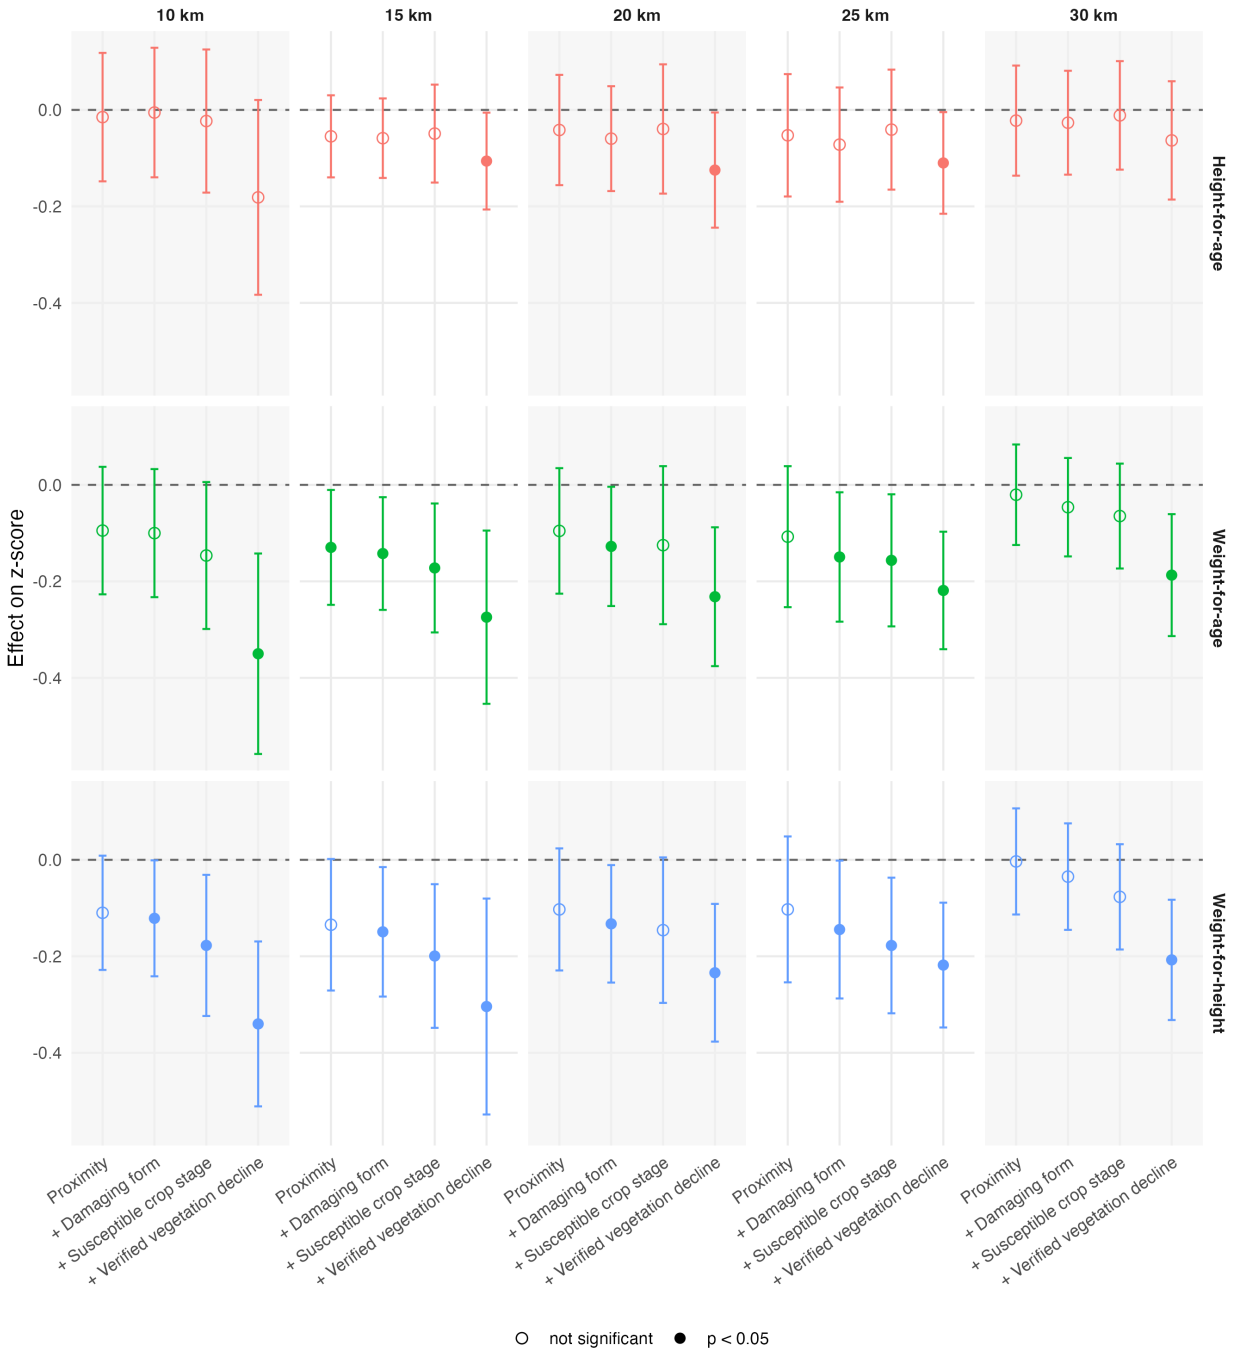


Figure 12: Exposure conditions, all three outcomes. Estimates under each exposure definition, adding one condition at a time, at radii from 10 to 30 km. Point estimates with 95 percent confidence intervals.

Table 21: Falsification tests: purged placebo estimates on weight-for-height, by radius and comparison group

Radius	Solitary/transiens, any stage	Solitary/transiens, stage and NDVI	Out-of-season damaging form
Panel A: Purged			
10 km	0.054 (0.101)	—	0.091 (0.066)
15 km	0.079 (0.088)	—	0.039 (0.059)
20 km	0.119 (0.076)	—	0.037 (0.061)
25 km	0.125* (0.075)	—	0.088 (0.061)
30 km	0.027 (0.099)	—	0.094* (0.052)
35 km	0.021 (0.087)	0.131 (0.106)	0.088 (0.063)
40 km	0.103 (0.063)	0.389*** (0.099)	0.013 (0.060)
45 km	0.062 (0.061)	0.294* (0.155)	0.063 (0.064)
50 km	0.024 (0.075)	0.054 (0.146)	0.033 (0.061)
Panel B: Purged, verified exposure controlled			
10 km	0.004 (0.097)	—	0.038 (0.067)
15 km	0.024 (0.091)	—	−0.030 (0.066)
20 km	0.080 (0.075)	—	−0.032 (0.066)
25 km	0.082 (0.074)	—	0.013 (0.070)
30 km	−0.022 (0.102)	—	0.011 (0.062)
35 km	−0.025 (0.091)	0.116 (0.111)	0.013 (0.073)
40 km	0.063 (0.069)	0.342*** (0.088)	−0.079 (0.076)
45 km	0.020 (0.065)	0.241 (0.155)	−0.021 (0.061)
50 km	−0.024 (0.087)	0.002 (0.154)	−0.050 (0.062)
Panel C: Purged, verified-exposed children dropped			
10 km	−0.029 (0.090)	—	0.087 (0.066)
15 km	0.030 (0.080)	—	0.003 (0.067)
20 km	0.057 (0.069)	—	0.007 (0.076)
25 km	0.021 (0.074)	—	0.043 (0.069)
30 km	−0.065 (0.117)	—	0.061 (0.071)
35 km	−0.089 (0.102)	0.109 (0.129)	0.057 (0.078)
40 km	0.017 (0.079)	0.385*** (0.129)	−0.038 (0.069)
45 km	−0.038 (0.077)	0.152 (0.174)	−0.019 (0.064)
50 km	−0.136 (0.093)	−0.168 (0.119)	−0.058 (0.066)

Outcome is weight-for-height (WHZ). Standard errors clustered on a half-degree spatial grid in parentheses. ***, **, * denote significance at 1, 5 and 10 percent. All three columns are purged of overlap with verified (D4) exposure at the same radius, as described in Section 5.2, and share the control set and fixed effects of column (3) of Table 2. Solitary/transiens, any stage is solitary and transiens locust records purged of D4 overlap; Solitary/transiens, stage and NDVI additionally requires the crop-stage and vegetation-decline conditions of equation (9) and is not estimated at 10 to 30 km, where purging leaves fewer than 20 children (see Figure 14); Out-of-season damaging form is damaging-form records purged of D4 overlap that arrive outside the susceptible crop stage. In Panel A, verified-exposed children sit in the comparison group. Panel B adds the verified indicator at the same radius as a control. Panel C drops verified-exposed children from the sample. Each panel has 27 nominal cells, 22 estimated after the 20-observation floor.

Table 22: Off-season damaging-form exposure in cropping zones

Radius	Weight-for-height z-score	Height-for-age z-score	Weight-for-age z-score	Treated
Panel A: Comparison group is every other child				
10 km	0.078	0.055	0.069	154
	(0.113)	(0.128)	(0.127)	
15 km	0.016	−0.005	0.002	169
	(0.112)	(0.111)	(0.124)	
20 km	−0.090	0.032	−0.038	173
	(0.131)	(0.079)	(0.109)	
25 km	0.057	−0.045	0.022	228
	(0.113)	(0.104)	(0.093)	
30 km	0.172*	0.062	0.150	268
	(0.101)	(0.084)	(0.090)	
35 km	0.085	0.076	0.093	257
	(0.116)	(0.093)	(0.114)	
40 km	0.007	−0.078	−0.044	258
	(0.135)	(0.094)	(0.129)	
45 km	0.049	−0.041	0.010	267
	(0.139)	(0.087)	(0.133)	
50 km	−0.109	−0.141	−0.169	265
	(0.127)	(0.086)	(0.114)	
Panel B: Stage and verified indicators controlled				
10 km	0.036	0.047	0.033	154
	(0.112)	(0.126)	(0.126)	
15 km	−0.030	−0.030	−0.046	169
	(0.111)	(0.109)	(0.123)	
20 km	−0.147	0.007	−0.092	173
	(0.132)	(0.081)	(0.104)	
25 km	−0.040	−0.093	−0.078	228
	(0.121)	(0.115)	(0.102)	
30 km	0.130	0.023	0.094	268
	(0.108)	(0.096)	(0.094)	
35 km	0.043	0.042	0.037	257
	(0.116)	(0.087)	(0.101)	
40 km	−0.075	−0.147*	−0.154	258
	(0.127)	(0.085)	(0.109)	
45 km	0.003	−0.079	−0.056	267
	(0.128)	(0.092)	(0.120)	
50 km	−0.155	−0.179**	−0.235**	265
	(0.126)	(0.086)	(0.101)	
Panel C: Stage- and verified-exposed children dropped				
10 km	0.173*	0.039	0.127	154
	(0.103)	(0.172)	(0.140)	
15 km	−0.008	−0.103	−0.048	169
	(0.116)	(0.135)	(0.132)	
20 km	0.012	0.035	0.042	173
	(0.109)	(0.119)	(0.111)	
25 km	0.121	−0.204	−0.035	228
	(0.125)	(0.142)	(0.095)	
30 km	0.279*	−0.120	0.141	268
	(0.151)	(0.137)	(0.119)	
35 km	0.103	−0.091	0.014	257
	(0.128)	(0.132)	(0.100)	
40 km	−0.015	−0.104	−0.086	258
	(0.140)	(0.122)	(0.112)	
45 km	0.008	−0.024	−0.022	267
	(0.119)	(0.135)	(0.113)	
50 km	−0.109	−0.076	−0.138	265
	(0.145)	(0.136)	(0.131)	

Sample is children in cropping livelihood zones (4,552 for weight-for-height, 4,557 for height-for-age and 4,555 for weight-for-age in Panels A and B). The exposure indicator equals one for a child with a damaging-form record within the radius during gestation and no such record in a susceptible crop stage, so the child is flagged by the damaging-form measure but by neither the stage nor the verified measure. In Panel A the comparison group is every other child in the sample, including children treated under the stage or verified measures. Panel B adds the stage and verified indicators at the same radius as controls. Panel C drops children flagged by the stage measure, which includes every verified-exposed child, leaving 1,818 to 3,436 children depending on the radius and outcome. Each cell is a separate regression with the control set and fixed effects of column (3) of Table 2. Standard errors clustered on a half-degree spatial grid in parentheses. ***, **, * denote significance at 1, 5 and 10 percent.

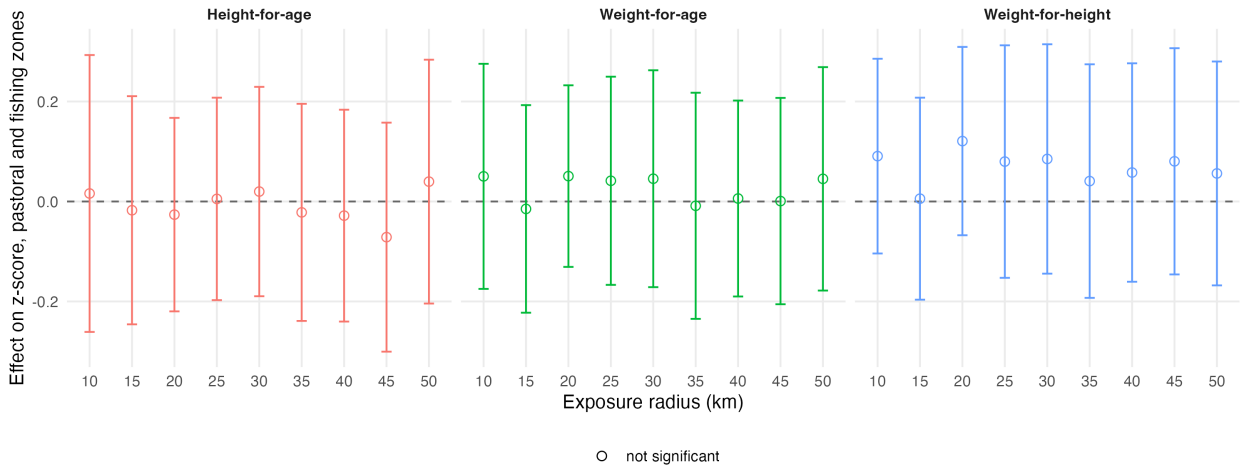


Figure 13: Proximity exposure in pastoral and fishing zones. Point estimates with 95 percent confidence intervals. No estimate is significant at the 10 percent level.

Table 23: Vegetation decline without a locust record: pseudo-event placebo

	Weight-for-height z-score	Height-for-age z-score	Weight-for-age z-score	Died	Share treated
Panel A: 2014 round					
Any decline	0.079	−0.013	0.038	−0.005	0.79
	(0.119)	(0.149)	(0.096)	(0.019)	
Matched	0.012	−0.011	0.001	0.005	0.20
	(0.048)	(0.065)	(0.047)	(0.010)	
Count	0.002	0.024	0.015	−0.005*	—
	(0.015)	(0.025)	(0.019)	(0.003)	
Panel B: 2022 round					
Any decline	−0.022	−0.051	−0.019	0.014	0.76
	(0.139)	(0.170)	(0.097)	(0.021)	
Matched	−0.084	−0.014	−0.048	0.009	0.20
	(0.075)	(0.054)	(0.069)	(0.005)	
Count	−0.019*	−0.026	−0.030***	0.005***	—
	(0.011)	(0.018)	(0.010)	(0.002)	

Every 16-day NDVI composite in a child's own cluster is treated as a candidate event and must satisfy the decline condition of equation (8) and the crop-stage condition of equation (5); no locust record is required. *Any decline* is an indicator for at least one qualifying decline during gestation. *Matched* is an indicator for the children whose largest decline during gestation is among the most severe in their round, with the cutoff (−1.56 standard deviations in 2014, −1.70 in 2022) set so that the treated share equals the verified measure's share at 30 km. *Count* is the number of qualifying declines during gestation. Each cell is a separate regression sharing the control set and fixed effects of column (3) of Table 2. Observations are 7,982 for the anthropometric outcomes and 8,698 for death in 2014, and 7,416, 7,424, 7,421, and 8,057 in 2022. Standard errors clustered on a half-degree spatial grid in parentheses. ***, **, * denote significance at 1, 5 and 10 percent.

Table 24: Verified exposure and vegetation decline with and without a nearby locust record, weight-for-height, 2022

	(1)	(2)	(3)	(4)
Verified exposure	−0.204***	−0.207***		−0.192***
	(0.062)	(0.068)		(0.065)
Matched decline	−0.075			
	(0.071)			
Matched decline, locust record within 50 km			−0.136*	−0.102
			(0.081)	(0.078)
Matched decline, no locust record within 50 km		0.005	0.032	−0.017
		(0.075)	(0.081)	(0.082)
Observations	7,416	7,416	7,416	7,416

Outcome is the weight-for-height z-score. Verified exposure is the main 30 km measure. Matched decline is defined as in Table 23. It is split by whether any locust record of any form fell within 50 km of the child's cluster during gestation. All columns share the control set and fixed effects of column (3) of Table 2. Standard errors clustered on a half-degree spatial grid in parentheses. ***, **, * denote significance at 1, 5 and 10 percent.

Table 25: Count of qualifying vegetation declines, by whether a locust record fell nearby, 2022

	Weight-for-height z-score	Height-for-age z-score	Weight-for-age z-score
Panel A: Count split by nearby locust record			
Declines, locust record within 50 km	−0.022**	−0.027*	−0.033***
	(0.011)	(0.016)	(0.009)
Declines, no locust record within 50 km	−0.012	−0.024	−0.025*
	(0.015)	(0.026)	(0.015)
Panel B: Adding verified exposure			
Verified exposure	−0.208***	−0.044	−0.177**
	(0.070)	(0.068)	(0.072)
Declines, locust record within 50 km	−0.008	−0.024	−0.021*
	(0.010)	(0.018)	(0.011)
Declines, no locust record within 50 km	−0.014	−0.024	−0.026*
	(0.015)	(0.026)	(0.014)
Observations	7,416	7,424	7,421

The count of qualifying declines of Table 23 is split by whether any locust record of any form fell within 50 km of the child's cluster during gestation. Each panel and column is a single regression sharing the control set and fixed effects of column (3) of Table 2. Verified exposure is the main 30 km measure. Standard errors clustered on a half-degree spatial grid in parentheses. ***, **, * denote significance at 1, 5 and 10 percent.

Table 26: Weight-for-height estimates with and without the drought control

Radius	With SPEI-6	Without SPEI-6	Treated
10 km	−0.340***	−0.346***	647
	(0.087)	(0.085)	
15 km	−0.304***	−0.309***	940
	(0.114)	(0.111)	
20 km	−0.234***	−0.238***	1,153
	(0.073)	(0.070)	
25 km	−0.218***	−0.223***	1,303
	(0.066)	(0.064)	
30 km	−0.207***	−0.212***	1,495
	(0.064)	(0.063)	
35 km	−0.175***	−0.181***	1,599
	(0.066)	(0.065)	
40 km	−0.155**	−0.161**	1,717
	(0.066)	(0.065)	
45 km	−0.145**	−0.152**	1,807
	(0.066)	(0.065)	
50 km	−0.121**	−0.128**	1,918
	(0.060)	(0.059)	

Outcome is the weight-for-height z-score. Standard errors clustered on a half-degree spatial grid in parentheses. ***, **, * denote significance at 1, 5 and 10 percent. *With SPEI-6* repeats the specification of column (3) of Table 2. *Without SPEI-6* is identical except for dropping the six-month gestational SPEI control, the only variable in the specification that measures drought directly. Observations are 7,416 in every cell. Treated counts are the treated children in that estimation sample, which excludes children with missing data and children who are the only child in their cluster.

Table 27: Exposure and Mortality

Radius	Estimate	Treated	p
10 km	0.005 (0.010)	736	0.628
15 km	0.004 (0.008)	1,065	0.610
20 km	0.006 (0.006)	1,303	0.311
25 km	0.004 (0.006)	1,467	0.545
30 km	0.004 (0.006)	1,679	0.489
35 km	0.003 (0.006)	1,796	0.568
40 km	0.009 (0.006)	1,928	0.184
45 km	0.008 (0.006)	2,035	0.191
50 km	0.008 (0.007)	2,154	0.256

Outcome is an indicator for death before interview. Standard errors clustered on a half-degree spatial grid in parentheses. Sample is 8,057 of the 8,213 births in the 2022 round, including children who died or were not measured (Section 5.2). Treated counts are among all 8,213 births. Baseline mortality is 3.2 percent. Each row is a separate regression sharing the control set and fixed effects of column (3) of Table 2.

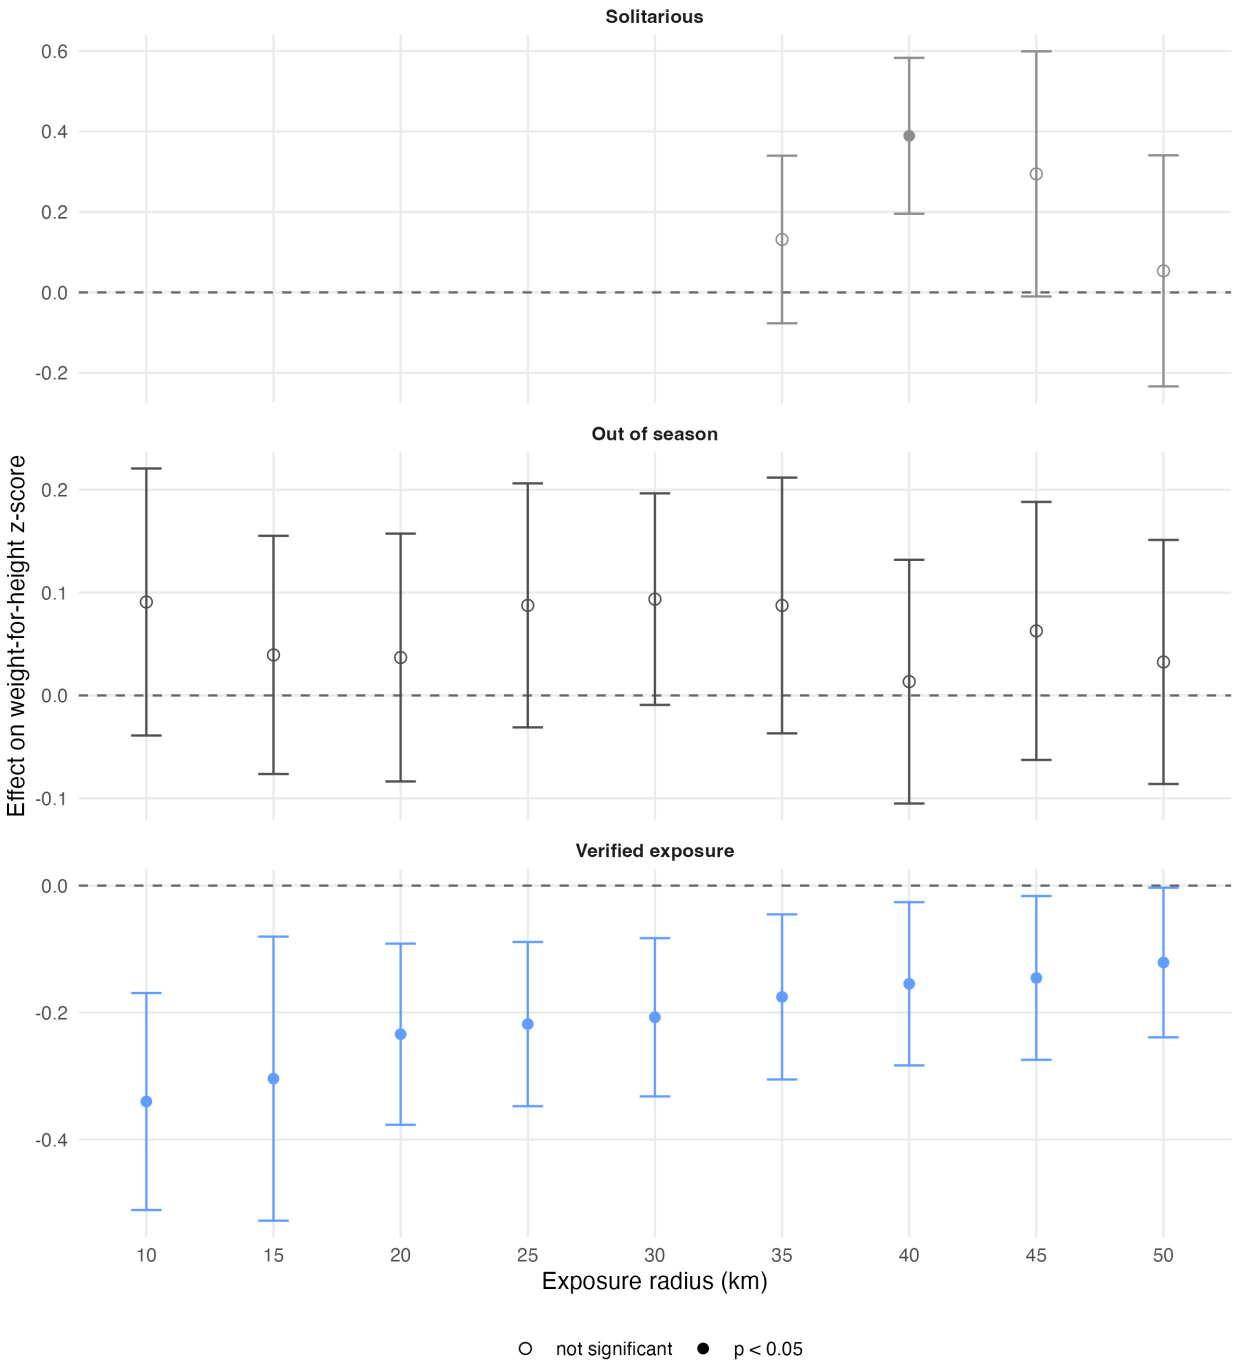


Figure 14: Falsification tests and verified exposure on weight-for-height, by radius. Solitarious and out-of-season estimates are purged of overlap with verified exposure at the same radius, as described in Section 5.2. Filled points are significant at the 5 percent level; hollow points are not. Y-axes are independent across panels and are not comparable to one another in scale. No solitarious estimate, which applies the stage and vegetation conditions, is reported at 10 to 30 km: purging leaves 9, 10, 18, 18 and 19 children respectively, all below a 20-observation floor I apply for this test.

Table 28: Shifted-calendar 2014 pre-invasion falsification: height-for-age z-score

Radius	Proximity	+ Damaging form	+ Susceptible crop stage	+ Verified vegetation decline
10 km	−0.092	−0.088	−0.068	0.071
	(0.085)	(0.086)	(0.093)	(0.126)
	[2,239]	[2,183]	[1,131]	[319]
15 km	−0.022	−0.012	−0.066	0.022
	(0.111)	(0.113)	(0.115)	(0.131)
	[2,874]	[2,798]	[1,492]	[430]
20 km	0.001	0.009	−0.073	−0.005
	(0.107)	(0.108)	(0.112)	(0.145)
	[3,292]	[3,212]	[1,751]	[522]
25 km	0.071	0.066	−0.049	0.085
	(0.084)	(0.080)	(0.100)	(0.098)
	[3,746]	[3,647]	[2,045]	[642]
30 km	0.029	0.024	−0.052	0.019
	(0.087)	(0.086)	(0.105)	(0.078)
	[4,080]	[3,978]	[2,268]	[768]
35 km	0.059	0.054	−0.022	0.111
	(0.073)	(0.069)	(0.081)	(0.067)
	[4,436]	[4,334]	[2,533]	[912]
40 km	0.047	0.053	−0.045	0.042
	(0.070)	(0.066)	(0.082)	(0.074)
	[4,768]	[4,659]	[2,753]	[1,039]
45 km	0.060	0.074	−0.019	0.138
	(0.066)	(0.060)	(0.071)	(0.085)
	[4,965]	[4,839]	[2,901]	[1,126]
50 km	0.070	0.077	−0.013	0.146
	(0.065)	(0.058)	(0.062)	(0.090)
	[5,175]	[5,043]	[3,055]	[1,225]

Each cell is a separate regression of the height-for-age z-score on the exposure measure in the column, at the radius in the row, when the real invasion's calendar pattern is shifted back 93 months onto the 2014 DHS cohort's actual gestation windows (Section 5.2). Columns add the conditions of Section 4.1 one at a time, so the last column is the verified measure. Every regression shares the control set and fixed effects of column (3) of Table 2. Standard errors clustered on a half-degree spatial grid in parentheses; treated children in brackets. ***, **, * denote significance at 1, 5 and 10 percent. Across Tables 28 to 31, 2 of the 144 cells are significant at the 5 percent level, the verified measure on weight-for-height at 10 km and on mortality at 30 km, and 5 (3.5 percent) are significant at the 10 percent level or better: the verified measure on weight-for-height at 10 km, the verified measure on mortality at 30, 35 and 50 km, and raw proximity on weight-for-age at 50 km.

Table 29: Shifted-calendar 2014 pre-invasion falsification: weight-for-age z-score

Radius	Proximity	+ Damaging form	+ Susceptible crop stage	+ Verified vegetation decline
10 km	−0.041	−0.051	−0.070	−0.143
	(0.049)	(0.049)	(0.053)	(0.107)
	[2,239]	[2,183]	[1,131]	[319]
15 km	−0.004	−0.017	−0.097	−0.098
	(0.063)	(0.062)	(0.059)	(0.094)
	[2,874]	[2,798]	[1,492]	[430]
20 km	0.009	−0.001	−0.075	−0.060
	(0.090)	(0.090)	(0.085)	(0.087)
	[3,292]	[3,212]	[1,751]	[522]
25 km	0.030	0.017	−0.055	0.021
	(0.064)	(0.064)	(0.059)	(0.077)
	[3,746]	[3,647]	[2,045]	[642]
30 km	0.007	−0.012	−0.048	0.008
	(0.069)	(0.069)	(0.067)	(0.059)
	[4,080]	[3,978]	[2,268]	[768]
35 km	0.037	0.014	−0.011	0.092
	(0.054)	(0.055)	(0.054)	(0.061)
	[4,436]	[4,334]	[2,533]	[912]
40 km	0.017	−0.003	−0.020	0.052
	(0.048)	(0.049)	(0.047)	(0.066)
	[4,768]	[4,659]	[2,753]	[1,039]
45 km	0.064	0.044	0.011	0.115
	(0.041)	(0.042)	(0.060)	(0.085)
	[4,965]	[4,839]	[2,901]	[1,126]
50 km	0.089*	0.056	0.022	0.129
	(0.046)	(0.042)	(0.057)	(0.088)
	[5,175]	[5,043]	[3,055]	[1,225]

Each cell is a separate regression of the weight-for-age z-score on the exposure measure in the column, at the radius in the row, when the real invasion's calendar pattern is shifted back 93 months onto the 2014 DHS cohort's actual gestation windows (Section 5.2). Columns add the conditions of Section 4.1 one at a time, so the last column is the verified measure. Every regression shares the control set and fixed effects of column (3) of Table 2. Standard errors clustered on a half-degree spatial grid in parentheses; treated children in brackets. ***, **, * denote significance at 1, 5 and 10 percent.

Table 30: Shifted-calendar 2014 pre-invasion falsification: weight-for-height z-score

Radius	Proximity	+ Damaging form	+ Susceptible crop stage	+ Verified vegetation decline
10 km	0.002	−0.016	−0.054	−0.277**
	(0.050)	(0.045)	(0.055)	(0.123)
	[2,239]	[2,183]	[1,131]	[319]
15 km	0.012	−0.013	−0.082	−0.169
	(0.057)	(0.047)	(0.058)	(0.106)
	[2,874]	[2,798]	[1,492]	[430]
20 km	0.010	−0.009	−0.048	−0.088
	(0.070)	(0.063)	(0.061)	(0.099)
	[3,292]	[3,212]	[1,751]	[522]
25 km	−0.009	−0.026	−0.042	−0.033
	(0.062)	(0.054)	(0.050)	(0.104)
	[3,746]	[3,647]	[2,045]	[642]
30 km	−0.011	−0.033	−0.028	−0.006
	(0.061)	(0.053)	(0.051)	(0.062)
	[4,080]	[3,978]	[2,268]	[768]
35 km	0.009	−0.019	−0.003	0.047
	(0.055)	(0.049)	(0.052)	(0.061)
	[4,436]	[4,334]	[2,533]	[912]
40 km	−0.010	−0.040	0.005	0.042
	(0.056)	(0.050)	(0.057)	(0.067)
	[4,768]	[4,659]	[2,753]	[1,039]
45 km	0.045	0.007	0.028	0.057
	(0.057)	(0.053)	(0.077)	(0.081)
	[4,965]	[4,839]	[2,901]	[1,126]
50 km	0.077	0.028	0.044	0.080
	(0.060)	(0.048)	(0.070)	(0.077)
	[5,175]	[5,043]	[3,055]	[1,225]

Each cell is a separate regression of the weight-for-height z-score on the exposure measure in the column, at the radius in the row, when the real invasion's calendar pattern is shifted back 93 months onto the 2014 DHS cohort's actual gestation windows (Section 5.2). Columns add the conditions of Section 4.1 one at a time, so the last column is the verified measure. Every regression shares the control set and fixed effects of column (3) of Table 2. Standard errors clustered on a half-degree spatial grid in parentheses; treated children in brackets. ***, **, * denote significance at 1, 5 and 10 percent.

Table 31: Shifted-calendar 2014 pre-invasion falsification: mortality

Radius	Proximity	+ Damaging form	+ Susceptible crop stage	+ Verified vegetation decline
10 km	0.005	0.003	−0.003	−0.020
	(0.009)	(0.008)	(0.010)	(0.029)
	[2,416]	[2,354]	[1,226]	[342]
15 km	0.014	0.010	0.001	−0.021
	(0.009)	(0.009)	(0.012)	(0.023)
	[3,109]	[3,025]	[1,618]	[459]
20 km	0.011	0.008	−0.001	−0.018
	(0.009)	(0.008)	(0.010)	(0.021)
	[3,562]	[3,474]	[1,897]	[560]
25 km	0.011	0.008	−0.003	−0.012
	(0.009)	(0.008)	(0.010)	(0.013)
	[4,058]	[3,950]	[2,215]	[687]
30 km	0.012	0.009	−0.002	−0.021**
	(0.008)	(0.008)	(0.010)	(0.009)
	[4,434]	[4,321]	[2,469]	[823]
35 km	0.005	0.002	−0.010	−0.027*
	(0.011)	(0.010)	(0.014)	(0.014)
	[4,810]	[4,695]	[2,750]	[975]
40 km	0.013	0.009	−0.004	−0.013
	(0.009)	(0.008)	(0.012)	(0.011)
	[5,169]	[5,047]	[2,989]	[1,118]
45 km	0.009	0.002	−0.009	−0.021
	(0.012)	(0.011)	(0.015)	(0.015)
	[5,387]	[5,243]	[3,151]	[1,208]
50 km	0.005	0.004	−0.006	−0.024*
	(0.012)	(0.010)	(0.014)	(0.013)
	[5,608]	[5,458]	[3,313]	[1,315]

Each cell is a separate regression of an indicator for death before interview, over all 2014 births alive or dead, on the exposure measure in the column, at the radius in the row, when the real invasion's calendar pattern is shifted back 93 months onto the 2014 DHS cohort's actual gestation windows (Section 5.2). Columns add the conditions of Section 4.1 one at a time, so the last column is the verified measure. Every regression shares the control set and fixed effects of column (3) of Table 2. Standard errors clustered on a half-degree spatial grid in parentheses; treated children in brackets. ***, **, * denote significance at 1, 5 and 10 percent.

Table 32: Timing test without a separate later-swarm group, weight-for-height, 30 km radius

	(1)	(2)	Children
Born before h_1	−0.188*	−0.191*	625
	(0.102)	(0.101)	
Gestation spans h_1	−0.132	−0.140	1,183
	(0.106)	(0.107)	
Conceived between h_1 and h_2	0.087	0.081	729
	(0.100)	(0.100)	
Conceived after h_2	0.078	0.075	251
	(0.119)	(0.118)	
Maize price control	No	Yes	
Observations	7,416	7,416	

As Table 4, except that the 225 children in utero during a later swarm stay in the two "conceived" groups. Standard errors clustered on a half-degree spatial grid in parentheses. ***, **, * denote significance at 1, 5 and 10 percent.

Table 33: Timing test in clusters hit in a single season, weight-for-height, 30 km radius

	(1)	(2)	Children
Born before h_1	−0.186*	−0.188*	560
	(0.104)	(0.102)	
Gestation spans h_1	−0.159	−0.171	876
	(0.109)	(0.108)	
Conceived between h_1 and h_2	0.141	0.143	469
	(0.102)	(0.097)	
Conceived after h_2	0.159	0.174	106
	(0.155)	(0.137)	
Maize price control	No	Yes	
Observations	6,549	6,549	

As Table 4, dropping the 232 clusters where verified events within 30 km preceded more than one harvest. Clusters hit in one season and clusters no verified event reached are kept. Standard errors clustered on a half-degree spatial grid in parentheses. ***, **, * denote significance at 1, 5 and 10 percent.

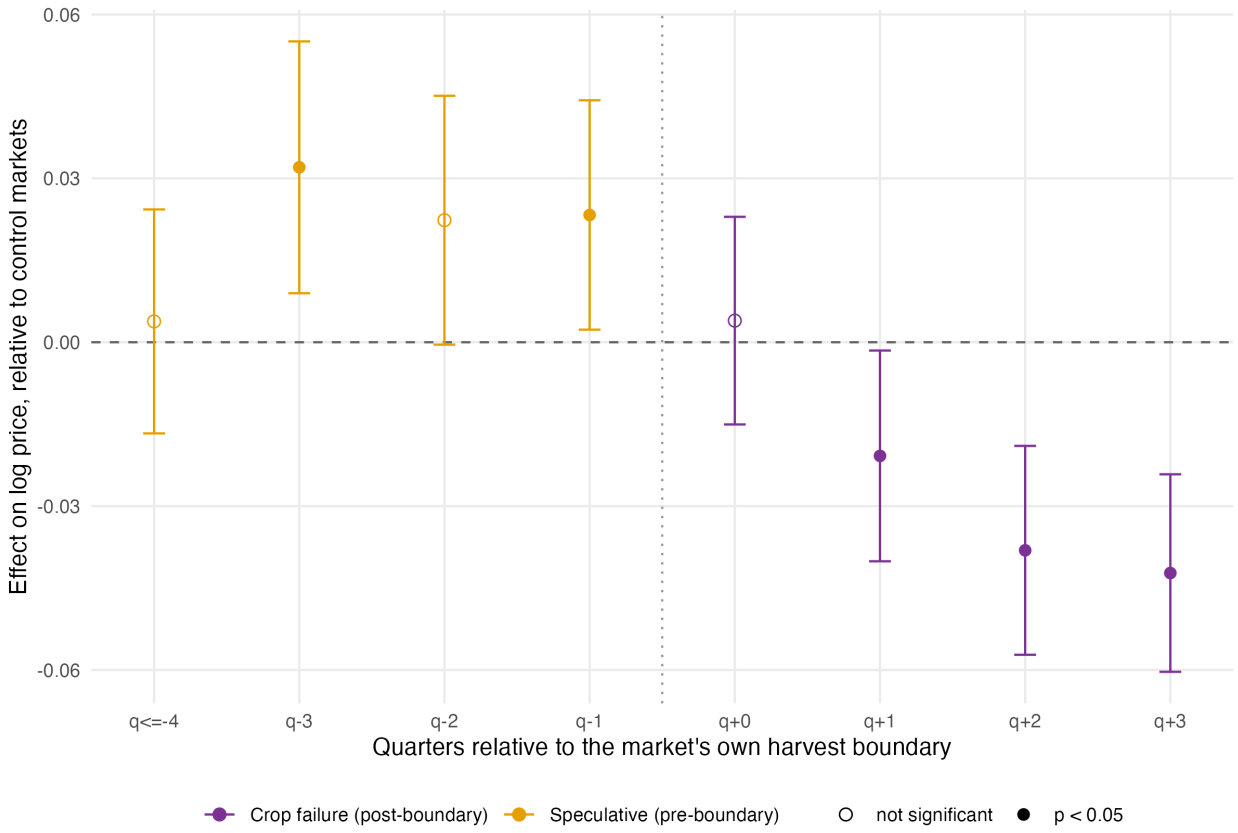


Figure 15: Market-level price event study at 30 km, by quarter relative to each treated market's own harvest boundary. Log maize and bean prices in treated versus control markets, with the specification of equation (12). Point estimates with 95 percent confidence intervals; standard errors clustered on market.

Table 34: Exposure effect by livelihood zone, weight-for-height, 30 km radius

				$S_i = 1$ under split		
Zone	Estimate	Treated	Grid cells	Type	Income	Food
Central Highlands High Potential	-0.429^{***} (0.086)	397	11			
Western High Potential	-0.107 (0.083)	354	15			
Western Medium Potential	-0.247 (0.158)	195	10			•
Western Agropastoral	0.024 (0.151)	178	12	•	•	
Southeastern Marginal Mixed	0.082 (0.211)	114	13			•
Coastal Medium Potential	-0.090 (0.119)	78	6		•	
Coastal Marginal Mixed	-0.344^{***} (0.114)	46	7			•
Western Lakeshore Marginal Mixed	0.398^{*} (0.233)	33	3			
Southern Agropastoral	0.078 (0.115)	23	4	•	•	
Tana Riverine	-0.445^{***} (0.144)	20	3			•
Northeastern Agropastoral	-0.631^{***} (0.188)	19	3	•	•	
Southeastern Medium Potential	-0.463 (0.313)	16	2			•
Northwestern Agropastoral	-0.123 (0.075)	6	1	•	•	
Lake Victoria Fishing	0.081 (0.075)	6	1	•	•	
Turkwell Riverine	0.161^{*} (0.083)	5	1			
Mandera Riverine	-0.639^{***} (0.069)	5	1			•

One regression of the weight-for-height z-score on exposure interacted with each livelihood zone, with the control set and fixed effects of column (3) of Table 2; N = 7,416. Treated zones only. Bullets mark zones with $S_i = 1$ under each split of equation (13): *Type*, pastoral, agropastoral or fishing zone (Table 5, column 1); *Income*, livestock or fishing is the main income source (column 2); *Food*, own production supplies half or more of food (Table 37). Standard errors clustered on a half-degree spatial grid in parentheses, and unreliable for zones with one to three grid cells. ***, **, * denote significance at 1, 5 and 10 percent.

Table 35: Heterogeneity: 30km exposure interacted with baseline characteristics and livelihood zone

	Weight-for-height z-score	Height-for-age z-score	Weight-for-age z-score
Panel A: Baseline characteristic			
Agricultural occupation	0.185	0.106	0.185
	(0.123)	(0.129)	(0.136)
Owns land	0.080	−0.008	0.049
	(0.113)	(0.079)	(0.114)
Rural	0.215*	0.098	0.212**
	(0.117)	(0.079)	(0.102)
Distance to major road (km)	0.012	0.016	0.017*
	(0.010)	(0.010)	(0.010)
Panel B: Livelihood zone (relative to cropping)			
Agropastoral	0.244**	0.332**	0.333***
	(0.121)	(0.139)	(0.106)
Riverine	−0.232	−0.478***	−0.406***
	(0.143)	(0.182)	(0.125)
Fishing	0.294***	0.617***	0.500***
	(0.093)	(0.104)	(0.094)

Standard errors clustered on a half-degree spatial grid in parentheses. ***, **, * denote significance at 1, 5 and 10 percent. Panel A: each cell reports the coefficient on exposure (30 km) interacted with the row characteristic from a separate regression of the column outcome on exposure, the row characteristic, their interaction, and the control set and fixed effects of column (3) of Table 2. Panel B: each cell reports exposure interacted with livelihood zone from a single regression per outcome, cropping as the reference level, with the same control set and fixed effects; each coefficient is that zone's exposure effect relative to cropping. Pastoral zones are excluded because no pastoral child is ever treated under the verified-damage measure. Among the 1,547 treated children with valid anthropometry at the main radius, 1,282 are in cropping zones, 229 in agropastoral, 30 in riverine, and 6 in fishing; the riverine and fishing coefficients rest on far fewer treated children than the cropping and agropastoral estimates.

Table 36: Riverine-versus-cropping exposure gap, by radius

Radius (km)	Weight-for-height z-score	Height-for-age z-score	Weight-for-age z-score	Treated in riverine
10	−0.175	−0.225	−0.225	23
	(0.145)	(0.144)	(0.164)	
15	−0.221	−0.294**	−0.303**	23
	(0.156)	(0.123)	(0.149)	
20	−0.269**	−0.273**	−0.331**	23
	(0.129)	(0.127)	(0.139)	
25	−0.284**	−0.280**	−0.340**	23
	(0.124)	(0.121)	(0.132)	
30	−0.232	−0.478***	−0.406***	30
	(0.143)	(0.182)	(0.125)	
35	−0.261*	−0.485***	−0.431***	30
	(0.146)	(0.182)	(0.128)	
40	0.104	−0.090	0.014	33
	(0.168)	(0.235)	(0.200)	
45	0.096	−0.080	0.011	33
	(0.167)	(0.234)	(0.199)	
50	0.157	−0.069	0.054	34
	(0.189)	(0.218)	(0.177)	

Each row re-estimates the riverine-versus-cropping interaction from the bottom panel of Table 35 at the row's radius in place of the main 30 km radius, with the same control set and fixed effects. Standard errors clustered on a half-degree spatial grid in parentheses. ***, **, * denote significance at 1, 5 and 10 percent.

Table 37: Exposure by livelihood zone: where food comes from, weight-for-height, 30 km radius

	Food source
Exposure	−0.216***
	(0.077)
Exposure × grows half or more of its food	0.030
	(0.142)
Implied effect, grows half or more ($\beta + \eta$)	−0.185
	(0.116)
Treated, buys most / grows half or more	1,099 / 396
Observations	7,416

One regression of equation (13) with the control set and fixed effects of column (3) of Table 2. $S_i = 1$ in zones where own production supplies half or more of the food households consume, according to the FEWS NET livelihood profiles, and 0 in zones that buy most of their food. Standard errors clustered on a half-degree spatial grid in parentheses. ***, **, * denote significance at 1, 5 and 10 percent.

Table 38: Exposure interacted with household characteristics, one at a time, weight-for-height, 30 km radius

	Exposure × characteristic ($\hat{\eta}$)
Wealth quintile (1–5)	−0.045
	(0.036)
Owns land	0.080
	(0.113)
Rural	0.215*
	(0.117)
Distance to major road (km)	0.012
	(0.010)

Each row is a separate regression of the weight-for-height z-score on exposure at 30 km, the row characteristic and their interaction, with the control set and fixed effects of column (3) of Table 2; $N = 7{,}416$, 1,495 treated. The four characteristics are correlated with one another; rural residence and wealth quintile correlate at −0.62. Standard errors clustered on a half-degree spatial grid in parentheses. ***, **, * denote significance at 1, 5 and 10 percent.

Table 39: Heterogeneity: baseline characteristics entered jointly

	Weight-for-height z-score	Height-for-age z-score	Weight-for-age z-score
Agricultural occupation	0.137	0.080	0.135
	(0.136)	(0.134)	(0.145)
Owns land	0.022	−0.040	−0.011
	(0.143)	(0.085)	(0.141)
Rural	0.179	0.055	0.168
	(0.174)	(0.101)	(0.156)
Distance to major road (km)	0.002	0.014	0.009
	(0.010)	(0.011)	(0.010)

Each column is a single regression of the column outcome on exposure at the main radius, all four baseline characteristics, their interactions with exposure, and the control set and fixed effects of column (3) of Table 2; each cell reports one characteristic's exposure interaction from that shared regression, so the four rows in a column are directly comparable. Compare to the one-at-a-time version in the top panel of Table 35: no characteristic is significant once the other three are controlled for directly. Standard errors clustered on a half-degree spatial grid in parentheses. ***, **, * denote significance at 1, 5 and 10 percent.

Table 40: Exposed children who are wasted versus exposed children who are not, 30 km radius

	Wasted	Not wasted	Difference	p
Mother's education (years)	6.44	9.63	−3.19	< 0.001
Wealth quintile	2.34	3.07	−0.74	< 0.001
Distance to major road (km)	5.32	3.72	1.60	0.046
Rural	0.688	0.617	0.071	0.190
Owns land	0.462	0.542	−0.079	0.173
Mother's height (cm)	160.46	160.60	−0.14	0.739
Gestational SPEI (6 month)	1.126	1.258	−0.132	0.267
Gestational NDVI z-score	0.743	0.735	0.008	0.930
Age at measurement (months)	14.93	15.69	−0.76	0.281

Sample restricted to children flagged as exposed at 30 km. Wasted denotes a weight-for-height z-score below −2 ($N = 80$). Not wasted denotes the remainder of the exposed group ($N = 1{,}466$). p-values from two-sample t-tests.

Table 41: Exposed children who are underweight versus exposed children who are not, 30 km radius

	Underweight	Not underweight	Difference	p
Mother's education (years)	6.71	9.79	−3.08	< 0.001
Wealth quintile	2.21	3.13	−0.92	< 0.001
Distance to major road (km)	5.36	3.62	1.75	0.001
Rural	0.764	0.603	0.160	< 0.001
Owns land	0.509	0.541	−0.032	0.442
Mother's height (cm)	159.49	160.73	−1.24	< 0.001
Gestational SPEI (6 month)	1.228	1.254	−0.025	0.760
Gestational NDVI z-score	0.879	0.718	0.161	0.015
Age at measurement (months)	15.75	15.64	0.11	0.834

Sample restricted to children flagged as exposed at 30 km. Underweight denotes a weight-for-age z-score below −2 ($N = 165$). Not underweight denotes the remainder of the exposed group ($N = 1{,}381$). p-values from two-sample t-tests.

Table 42: Exposed cases versus exposed non-cases, 15 km radius

	Case	Non-case	Difference	p
Panel A: Wasted versus not wasted				
Mother's education (years)	6.47	9.74	−3.27	< 0.001
Wealth quintile	2.34	3.24	−0.90	< 0.001
Distance to major road (km)	5.32	3.35	1.97	0.044
Rural	0.644	0.539	0.105	0.111
Owns land	0.424	0.514	−0.090	0.181
Mother's height (cm)	160.64	160.60	0.04	0.936
Gestational SPEI (6 month)	1.125	1.218	−0.093	0.526
Gestational NDVI z-score	0.714	0.722	−0.008	0.942
Age at measurement (months)	14.95	15.70	−0.75	0.380
Panel B: Underweight versus not underweight				
Mother's education (years)	6.77	9.90	−3.13	< 0.001
Wealth quintile	2.33	3.29	−0.96	< 0.001
Distance to major road (km)	5.50	3.21	2.29	< 0.001
Rural	0.721	0.524	0.197	< 0.001
Owns land	0.505	0.509	−0.004	0.935
Mother's height (cm)	159.52	160.75	−1.22	0.004
Gestational SPEI (6 month)	1.207	1.213	−0.006	0.952
Gestational NDVI z-score	0.845	0.706	0.140	0.094
Age at measurement (months)	15.84	15.62	0.22	0.729

Sample restricted to children flagged as exposed at 15 km. Panel A compares children with a weight-for-height z-score below −2 ($N = 59$) to the remainder of the exposed group ($N = 918$). Panel B compares children with a weight-for-age z-score below −2 ($N = 111$) to the remainder ($N = 867$). p-values from two-sample t-tests.

B Data and Descriptive Statistics

B.1 Livelihood Zones and Crop Calendars

FEWS NET publishes a profile for each of Kenya's 24 livelihood zones (Famine Early Warning Systems Network (FEWS NET), 2011), made of a description page and a seasonal calendar. Figure 16 shows the description page for Livelihood Zone 19, the Central Highlands High Potential Zone. The left column lists the zone's districts, main markets, food crops, sources of cash income, livestock and hazards. The text on the right describes its altitude, rainfall and population density, the share of income from each source, and how much of their food households grow themselves. In Zone 19, cash crops such as coffee, tea and pyrethrum provide about 30 percent of income, food crop sales about 20 percent, and milk and poultry about 30 percent, and households grow about a third of the food they eat and buy the rest. I use these pages to classify each zone by livelihood type, main income source and the share of food grown at home, the three splits in Sections 5.3.2 and 5.3.3, and the seasonal calendars to date the crop stages in Section 4.1.2.

Livelihood Zone 19: Central Highlands, High Potential Zone

Districts	Southern Meru North, Meru South, Meru Central, Kirinyaga, Muranga, Nyeri, Maragua, Kiambu, Nyandarua, southern Laikipia, Thika
Markets	• Githunguri, Limuru, Kikuyu, Kerugoya, Kutus, Embu, Sagana, Makutano, Maua, Mikinduri, Chogoria, Nyahururu, Kinamba, Nyeri, Thika
Food Crops	• Maize, Irish potato, Beans
Cash Income	• Cash crops sales • Food crop sales • Livestock product sales • Casual labor
Livestock	• Dairy cattle • Goats • Sheep • Poultry
Hazards	• HIV/AIDS • Drought • High/Variable food prices • Livestock pests and diseases • Crop diseases • Crop damage from wild animals • Malaria

This zone is located at an altitude of 1250-2300 meters above sea level and has a high population density estimated at about 280 persons/km^2. The rains which are very reliable, average 800-2000 mm per annum and are received continuously from March through November with peaks in March-May and October-November. Annual temperatures average 15-20°C and are highest between December-March.

The Kikuyu and Meru are the main ethnic communities in the zone. The majority of the inhabitants (82%) are fully settled within the zone. 10% are out migrant laborers while 7% are in-migrant laborers. The main source of income (about 30%) is cash crop sales include coffee, tea and pyrethrum. Sale of food crops (maize, beans and Irish potatoes) are also an important source of income (about 20%) along with the sale of milk and poultry (about 30%). Poorer households depend on casual labor throughout the year as a significant source of income. The peak demand for farm related labor occurs between February-November.

A typical household keeps 5-7 dairy cattle, about 15 chickens, 5-10 goats (including dairy goats) up to 3-5 sheep, a pig and a donkey. Due to the emphasis on cash crop production, only about third of food consumed by households is sourced from own production. The remaining two-thirds is purchased. A well developed road network facilitates good market integration within the zone. Nevertheless, in remote parts where roads are in poor condition, households routinely loose farm produce to spoilages during the peak wet seasons when markets become inaccessible.

With the exception of the landless and urban poor households, the zone is considered to be generally food secure. Constraints to further development of cash and food crop production include small farm sizes, endemic crop diseases and low technical skills. In remote areas of the zone, poor access to markets is also a factor limiting further crop production.

Figure 16: FEWS NET livelihood profile, description page, Livelihood Zone 19, Central Highlands High Potential Zone (Famine Early Warning Systems Network (FEWS NET), 2011).

Figures 17 and 18 reproduce the seasonal calendar for a cropping zone and a pastoral zone, and Figure 19 shows the maize calendar transcribed from these pages for every zone that grows maize.

Figure 17 shows the FEWS NET seasonal calendar for Livelihood Zone 19, the Central Highlands High Potential Zone, which covers the high-rainfall counties north of Nairobi. The zone runs a single extended long-rains season. Land preparation for the maize and bean crop occupies March and April, planting runs from late April into June, and harvesting begins in October and continues to January. Irish potato and cabbage run two cycles. Coffee and tea are harvested across most of the year, and a footnote lists permanent crops including avocado, mango, macadamia, banana, and pineapple. Food prices peak in May and June, in the gap between planting and harvest, and market access is poor during the rains. A swarm arriving in this zone in May or June meets the maize crop at planting, rank 2 in equation (3). A swarm arriving between October and January meets the crop under its harvest bar, rank 3. A swarm arriving in February meets bare ground, and one arriving between July and September, when the maize is growing but the calendar shows neither planting nor harvest, counts as out of season.

This is a limitation of the source. The FEWS NET calendars mark land preparation, planting and harvest, and leave the months between the end of planting and the start of harvest blank. A crop growing in those months stands in the field and can be eaten, but the calendars give no information on when a growing stage begins or ends for each crop and zone, and I do not interpolate one into the gap. Arrivals during the gap therefore count as out of season. They are excluded from the treatment and enter the out-of-season placebo group instead, so to the extent that they destroyed growing crops, treated children sit in the comparison group and the estimates are biased toward zero.

Figure 18 shows the calendar for Livelihood Zone 9, the Northeastern Pastoral Zone, covering Wajir and the surrounding rangeland. The calendar has no crop section. It records camel calving, kidding and lambing, disease outbreaks, milk availability, migration, livestock sales, and the lean season. Households here obtain food from milk, livestock sales, and purchase, and the FEWS assessment calendar tracks livestock condition. No date in this zone places a field in a susceptible stage, so equation (5) never holds and no child in this zone is treated. Section 5.2.2 shows that proximity to locusts does not predict lower weight-for-height in these zones.

Figure 19 collects the maize calendar transcribed from calendars like Figure 17, one row per livelihood zone. Zones without a maize row, such as the pastoral zone in Figure 18, are absent from the figure. The land preparation and planting window moves with each zone's rains, so the bimodal zones on the coast and around Lake Victoria show two separate cropping windows in a single calendar year while the single-season zones, including the Central Highlands zone above, show one.

Seasonal Calendar: Livelihood Zone 19 - Central Highlands, High Potential Zone

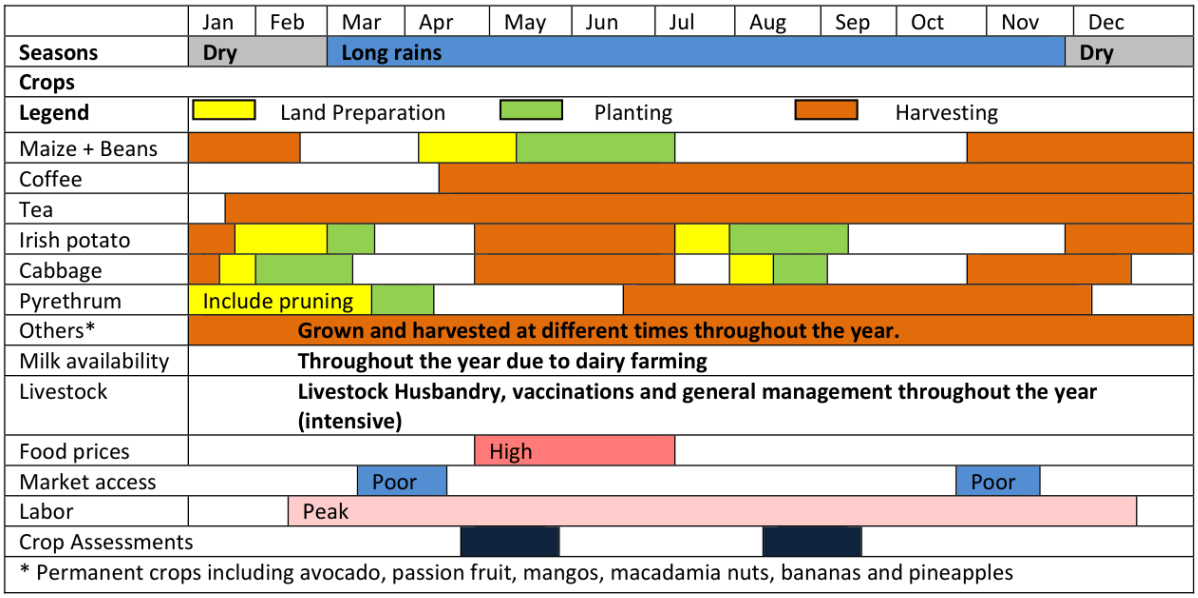


Figure 17: FEWS NET seasonal calendar, Livelihood Zone 19, Central Highlands High Potential Zone.

Seasonal Calendar: Livelihood Zone 9 - Northeastern Pastoral Zone

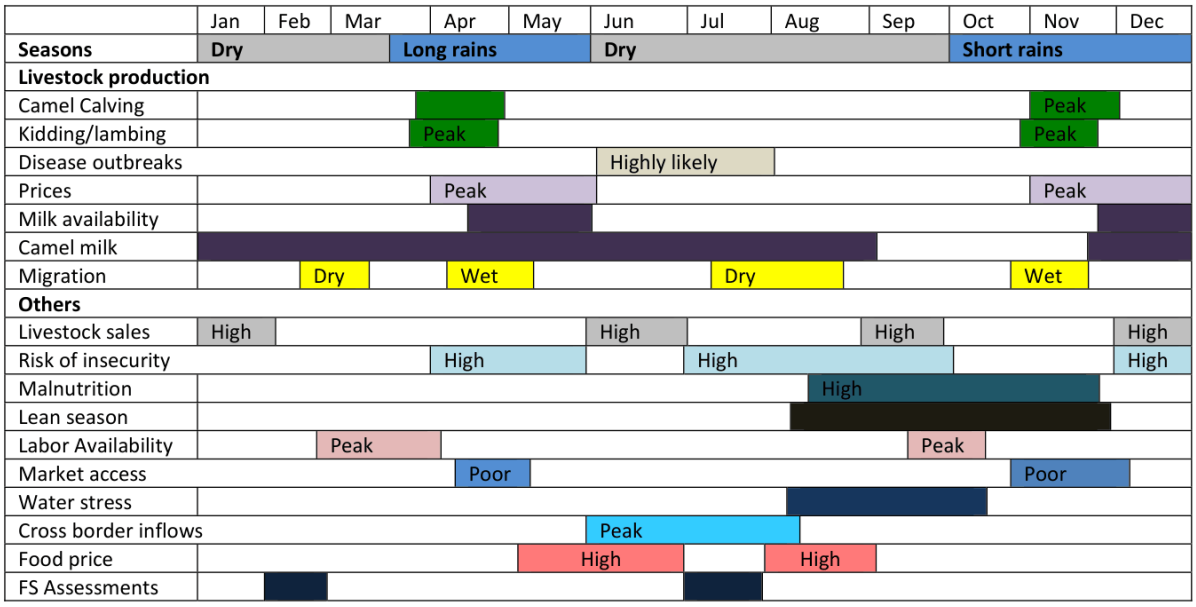


Figure 18: FEWS NET seasonal calendar, Livelihood Zone 9, Northeastern Pastoral Zone. The calendar contains no crop section.

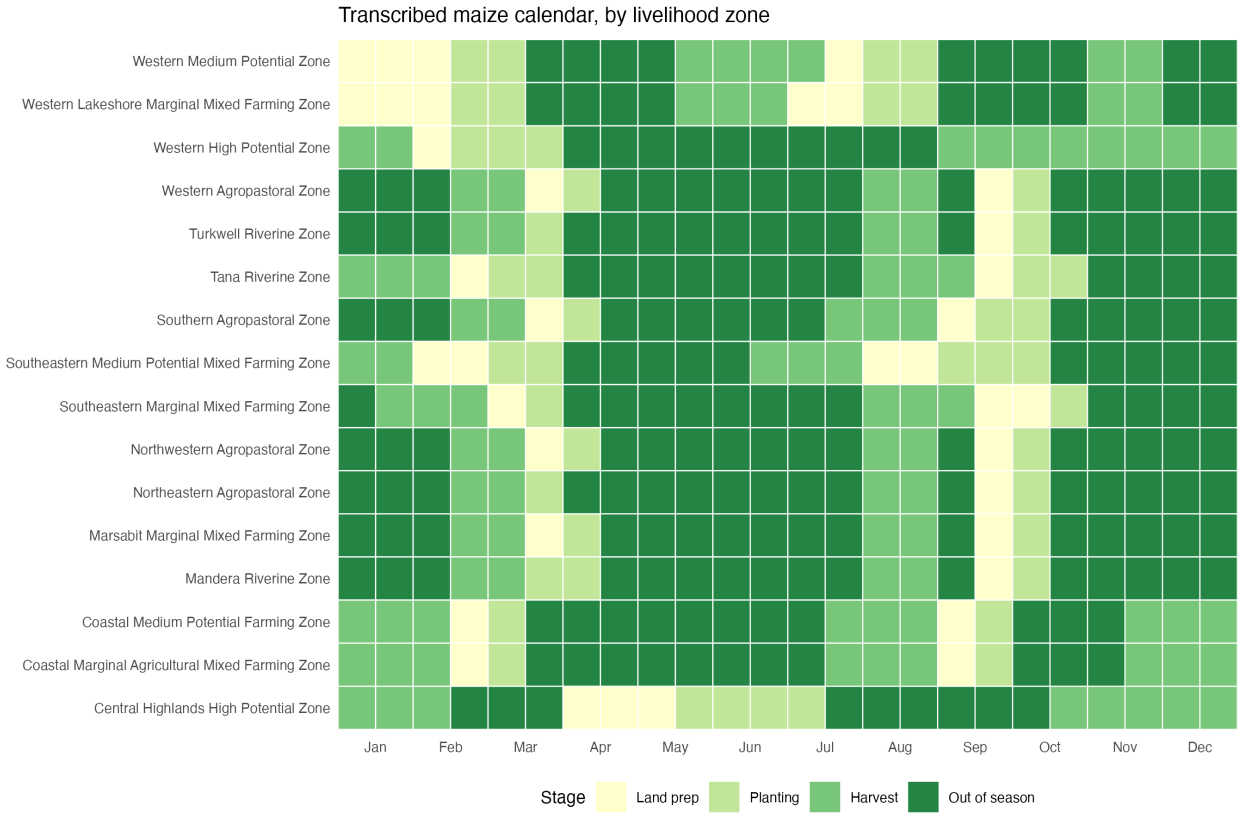


Figure 19: Transcribed maize calendar, by livelihood zone, built from calendars such as Figure 17.

B.2 Demographic and Health Survey Sample

Table 43 reports summary statistics for the 7,603 children with valid anthropometry in the 2022 DHS. The anthropometric outcomes carry Kenya's general burden. Mean height-for-age is -0.91 with a standard deviation of 1.31, and a quarter of the sample sits below -1.77. Weight-for-height is closer to the reference median at -0.21, consistent with a population where chronic stunting is more prevalent than acute wasting. Children average 14.7 months at measurement and the sample is 49 percent female. Two thirds of children are rural, 48 percent own land, and 68 percent own livestock, but only 16 percent report an agricultural primary occupation, which is why the occupation variable enters only as a heterogeneity dimension and never as a control.

Distance to a major road is heavily right-skewed. The median child lives 2.2 km from one and the 75th percentile 5.8 km, but the maximum exceeds 114 km. Gestational SPEI averages 1.06 at six months, reflecting that the invasion window overlapped a period of above-average moisture in much of the country.

Figure 20 maps the 2022 cluster locations against the major-road network and the World Bank markets used for the price analysis. Clusters follow population density, dense through the western highlands, the central highlands, and the coastal strip, and sparse across the arid north and northeast. The road and market layers follow the same pattern, which is the source of the remoteness gradient that recurs throughout the results.

Figure 21 marks the clusters with at least one child treated under the verified measure at 30 km. Treated clusters sit in the western and central highlands, the eastern slope and the coast, the cropping areas where

Table 43: Summary statistics, 2022 DHS children with valid anthropometry

	Obs.	Mean	St. Dev.	Min	Pctl(25)	Median	Pctl(75)	Max
Anthropometry								
Height-for-age (HAZ)	7,603	−0.91	1.31	−5.90	−1.77	−0.92	−0.11	5.95
Weight-for-age (WAZ)	7,600	−0.64	1.20	−5.29	−1.43	−0.63	0.13	4.85
Weight-for-height (WHZ)	7,595	−0.21	1.20	−5.00	−1.01	−0.20	0.58	4.82
Child and mother								
Age (months)	7,603	14.71	7.26	1.00	9.00	14.00	21.00	30.00
Female	7,603	0.49	0.50	0.00	0.00	0.00	1.00	1.00
First born	7,603	0.26	0.44	0.00	0.00	0.00	1.00	1.00
Twin	7,603	0.03	0.17	0.00	0.00	0.00	0.00	1.00
Birth order	7,603	3.23	2.24	1.00	1.00	3.00	4.00	14.00
Mother's education (yr)	7,603	7.66	5.08	0.00	3.50	8.00	12.00	24.00
Mother's height (cm)	7,603	160.76	4.39	139.70	160.40	160.70	161.00	180.40
Household and location								
Household size	7,603	5.90	2.57	2.00	4.00	5.00	7.00	24.00
Rural	7,603	0.67	0.47	0.00	0.00	1.00	1.00	1.00
Owns land	7,603	0.48	0.50	0.00	0.00	0.00	1.00	1.00
Owns livestock	7,603	0.68	0.47	0.00	0.00	1.00	1.00	1.00
Agricultural occupation	7,603	0.16	0.36	0.00	0.00	0.00	0.00	1.00
Distance to road (km)	7,603	7.28	15.65	0.00	0.74	2.18	5.82	114.65
Gestational window								
SPEI (6 month)	7,598	1.06	1.07	−2.11	0.25	1.20	1.98	3.06
NDVI z-score	7,603	0.69	0.80	−3.04	0.24	0.72	1.19	6.04
Maize price index	7,603	0.95	0.15	0.66	0.83	0.95	1.06	1.58

Mother's height is filled with the median among measured women for children whose mother was not measured (Section 4), which is why its interquartile range is narrow.

the stage condition can hold, while the pastoral zones of the north and northeast hold none; the few treated clusters there lie in the riverine, agropastoral and cropping zones inside that region.

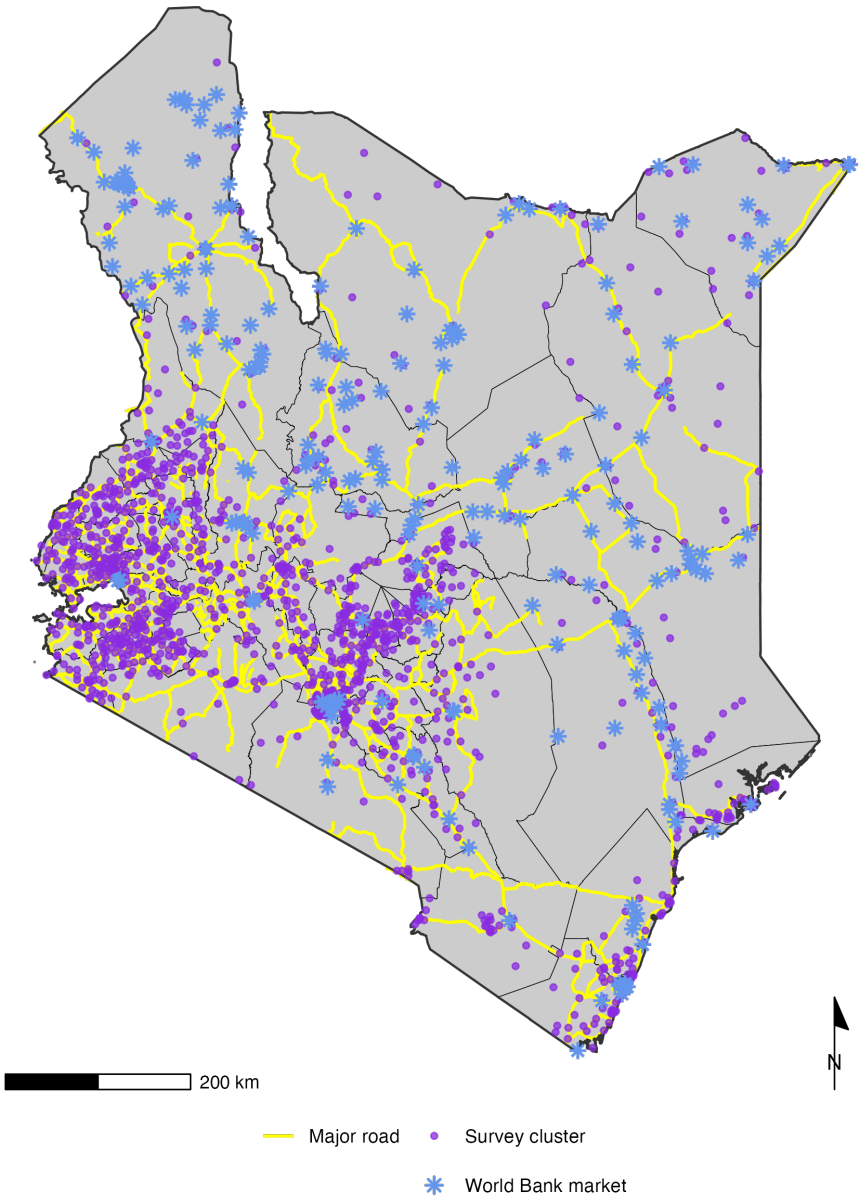


Figure 20: 2022 DHS cluster locations, major roads, and World Bank market locations.

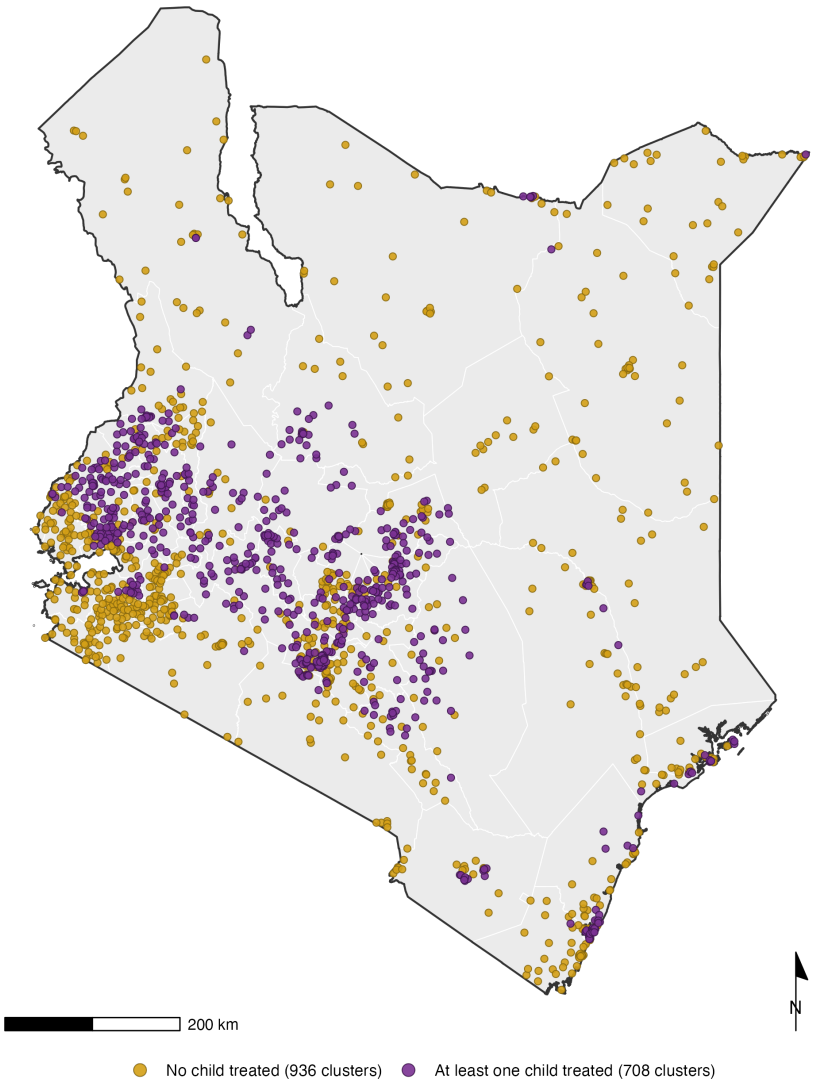


Figure 21: 2022 DHS clusters with at least one child treated under the verified measure at 30 km (purple) and with no treated child (mustard). Clusters are those with at least one child with valid anthropometry.

B.2.1 Sampling Weights

Figure 22 plots the distribution of DHS sampling weights on a log scale. The distribution is right-skewed, with the 99th percentile 6.3 times the median and the maximum 14.2 times, which motivates the winsorized robustness check in Table 44. Winsorizing weights at the first and 99th percentiles and re-running the full grid of 108 regressions leaves the median estimate at -0.076 against -0.093 untrimmed, the negative share at 93 percent against 99 percent, and the significant share at 28 percent against 33 percent. Sample size and composition are identical under winsorizing, since weights are capped rather than observations dropped. Panel B of Table 44 reports the verified estimate at 10, 30 and 50 km with DHS weights, winsorized weights and no weights. Winsorizing shrinks the weight-for-height estimate at 30 km from -0.207 to -0.186 and dropping the weights shrinks it to -0.118, both still significant at the 5 percent level. The unweighted estimates are smaller on every outcome, so the effect is larger among the children the weights scale up. I report the weighted estimates as the main results because they describe the population of Kenyan children rather than the sample as drawn.

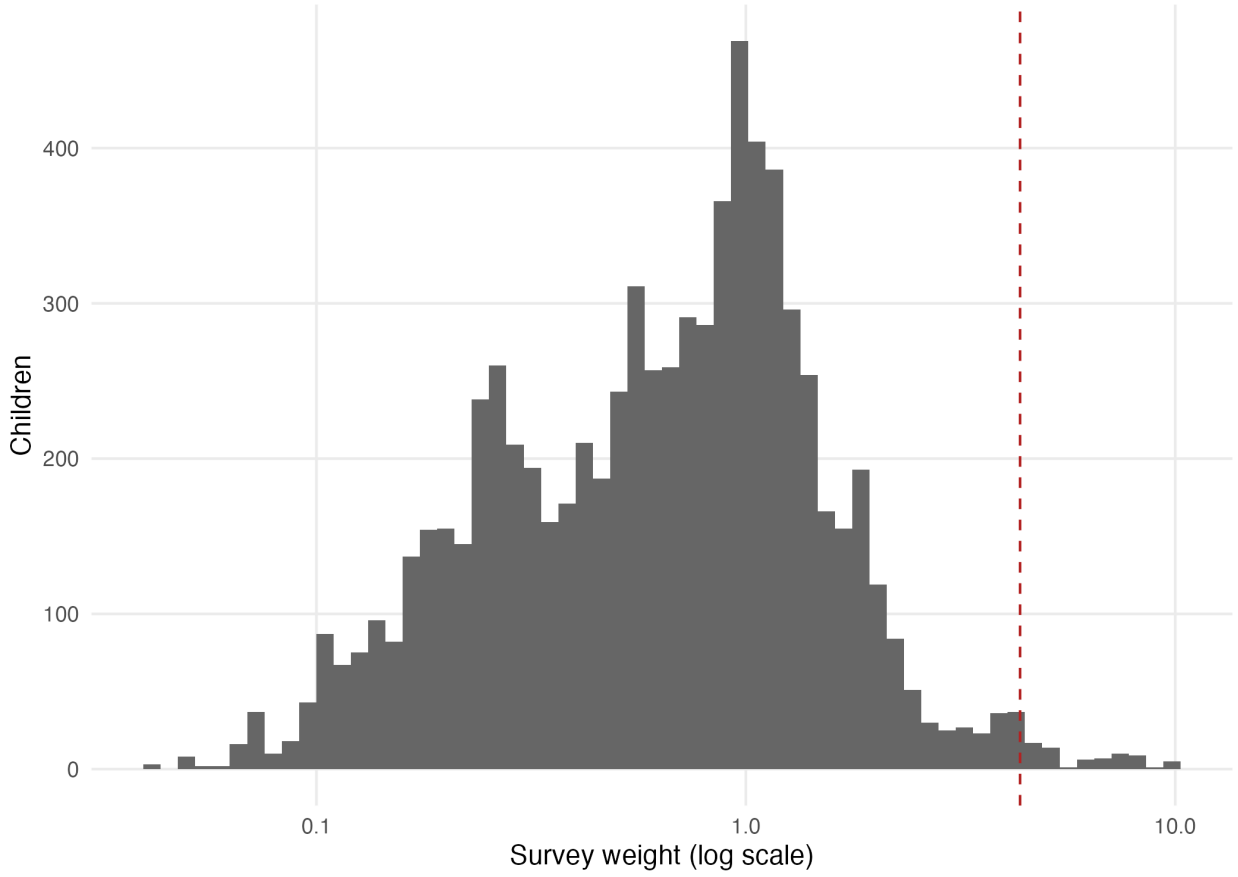


Figure 22: Distribution of DHS sampling weights. Dashed line marks the 99th percentile.

Table 44: Estimates under DHS, winsorized and no sampling weights

	DHS weights	Winsorized weights	Unweighted
Panel A: Core grid of 108 regressions			
Median estimate	−0.093	−0.076	—
Mean estimate	−0.102	−0.084	—
Share of estimates that are negative	0.99	0.93	—
Share significant at the 5 percent level	0.33	0.28	—
Panel B: Verified exposure			
Weight-for-height z-score			
10 km	−0.340***	−0.321***	−0.270***
	(0.087)	(0.079)	(0.071)
30 km	−0.207***	−0.186***	−0.118**
	(0.064)	(0.054)	(0.050)
50 km	−0.121**	−0.102*	−0.057
	(0.060)	(0.052)	(0.048)
Height-for-age z-score			
10 km	−0.181*	−0.153*	−0.054
	(0.103)	(0.089)	(0.065)
30 km	−0.063	−0.051	−0.008
	(0.063)	(0.059)	(0.054)
50 km	−0.057	−0.044	0.001
	(0.053)	(0.049)	(0.048)
Weight-for-age z-score			
10 km	−0.350***	−0.321***	−0.229***
	(0.106)	(0.089)	(0.065)
30 km	−0.187***	−0.166***	−0.090*
	(0.064)	(0.054)	(0.046)
50 km	−0.121**	−0.100**	−0.039
	(0.058)	(0.048)	(0.040)

Panel A summarizes the grid of four exposure definitions (proximity, damaging form, damaging form and crop stage, verified damage), nine radii from 10 to 50 km, and three anthropometric z-scores. Panel B reports the verified-exposure estimate at 10, 30 and 50 km. Each regression shares the control set and fixed effects of column (3) of Table 2. Winsorized weights cap the DHS individual sampling weight at its first and 99th percentiles; the unweighted column drops the weights. The sample is identical in every column. Standard errors clustered on a half-degree spatial grid in parentheses. ***, **, * denote significance at 1, 5 and 10 percent.

B.3 Climate and Vegetation Measures

Rainfall. Rainfall comes from CHIRPS at 0.05° resolution (Funk et al., 2015). For each cluster I average the CHIRPS pixels within the DHS displacement radius around the cluster coordinate, 2 km for urban and 5 km for rural clusters. The monthly product, from 1981 to 2022, is the precipitation input to SPEI. The dekadal product, at ten-day resolution, is used to detect the start of each growing season. Onset cannot be identified from monthly totals, since a season that begins in the third dekad of March and one that begins in the first dekad of April differ by roughly ten days in the field but fall in adjacent calendar months, and the crop calendar shift described in Section 4.1.2 depends on resolving that difference.

SPEI. I build the Standardized Precipitation Evapotranspiration Index of Vicente-Serrano et al. (2010) in four steps for each cluster. First, I compute the monthly climatic water balance $D_{c,m} = P_{c,m} - PET_{c,m}$, where P is CHIRPS precipitation and PET is TerraClimate reference evapotranspiration (Abatzoglou et al., 2018), averaged over the TerraClimate grid cell nearest the cluster and its eight neighbors, an area about 12 km across. Second, I sum the balance over the current and preceding months,

$$D_{c,m}^{(k)} = \sum_{j=0}^{k-1} D_{c,m-j}, \tag{14}$$

for accumulation windows of $k = 3$ and $k = 6$ months. Third, for each cluster, window and calendar month, I fit a three-parameter log-logistic distribution to $D^{(k)}$ by unbiased probability-weighted moments, using 1981 to 2018 as the reference period so the invasion years do not enter their own baseline. Fourth, I map each value's cumulative probability under the fitted distribution F to the standard normal,

$$\mathrm{SPEI}_{c,m}^{(k)} = \Phi^{-1}\Big(F\big(D_{c,m}^{(k)}\big)\Big). \tag{15}$$

Over the reference period SPEI has mean zero and standard deviation one, and a value of 1.0 means the water balance was one standard deviation wetter than normal for that cluster and time of year. Clusters whose fitted series depart from this over the reference period, with a mean beyond ± 0.15 or a standard deviation outside 0.8 to 1.25, are flagged but kept, since arid clusters with many dry months are both the hardest to fit and among the most relevant. The control in equation (11) is gestational SPEI-6, the mean of the monthly SPEI-6 values over the months of the child's gestation. The three- and six-month indices are highly correlated within cluster, so only the six-month index enters the control set, as the window more likely to cover a full growing season.

NDVI. Vegetation comes from MODIS MOD13Q1, a 16-day composite of NDVI at 250 m resolution (Didan, 2015), which I download for the Kenya extent from 2000 to 2022 through NASA's AppEEARS service. Within each 16-day period the product keeps the best available observation for each pixel, so the date of a composite is not the date of the satellite pass. Each pixel also carries a reliability flag. I keep pixels flagged good or marginal, drop those flagged as snow, ice or cloud, and take the area-weighted mean of the remaining pixels within the cluster's displacement buffer. A heavily clouded composite contributes fewer pixels, and one with no usable pixel is missing for that cluster.

I then convert levels to anomalies. Composites start on fixed days of the year, so I index each by its position in the year, $\tau = \lfloor (\text{day of year} - 1)/16 \rfloor + 1$, which takes 23 values. For each cluster c and position τ, I compute the mean $\mu_{c,\tau}$ and standard deviation $\sigma_{c,\tau}$ of NDVI over 2000 to 2018, before the invasion, and define

$$z_{c,t} = \frac{\mathrm{NDVI}_{c,t} - \mu_{c,\tau(t)}}{\sigma_{c,\tau(t)}}. \tag{16}$$

This removes both the level difference between a humid highland cluster and an arid northern one and the seasonal cycle within each, so a z-score of zero means vegetation was at its normal level for that place at that point in the year. Cluster-composite cells with fewer than ten reference years are kept and have noisier z-scores. The damage condition in equation (8) tests the change in z between the two composites before an event and the two after it, which is why the post-event window spans roughly 16 to 48 days depending on where in its composite the event falls, and why a cluster that was already dry relative to its own norm does not qualify on dryness alone. For the gestational NDVI measure in Table 1, I average the z-scores within each calendar month and then over the months of the child's gestation.

B.4 Food Prices

Prices come from the World Bank's Real-Time Prices (RTP) data for Kenya (Andrée, 2021). The platform covers 233 markets and reports monthly prices in Kenyan shillings for white maize and dry beans, each per kilogram, and for a goat. Its food price index weights maize and beans equally and gives the goat no weight. Observed transactions are sparse. From January 2007 to September 2026 the platform holds 3,055 observed

market-month prices for maize and 1,941 for beans, 4.49 percent of all food market-months. The platform fills the remaining months with Cubist, a rule-based regression model, trained on observed prices at the same and nearby markets, on related commodities and on past prices, following Andrée (2021). The reported R^2 of the imputation is 0.97 for maize and 0.98 for beans. I use the platform's final "close" estimate for each market-month, which combines observed and imputed prices. From 2015 to 2022, the 233 markets give 44,736 market-month prices across the two staples, and 6.6 percent of them rest on a transaction observed in that month.

The market-level event studies in Section 5.3.1 use each market's own series. For the gestational price control, I build a price for each DHS cluster and month from the nearest markets. I take up to the five nearest markets within 200 km of the cluster, with distances computed in UTM zone 37N, and average their prices with inverse-distance-squared weights,

$$p_{c,t} = \frac{\sum_m w_{cm}\, p_{m,t}}{\sum_m w_{cm}}, \qquad w_{cm} = \frac{1}{\max(d_{cm}, 1)^2}, \tag{17}$$

where d_{cm} is the distance in kilometres from cluster c to market m. A cluster-month blends about five markets on average, and the median distance to the nearest market is 26 km. I divide each cluster's maize price by its own average over 2015 to 2018, so the index equals 1 in a typical pre-invasion month, and the gestational maize price index in Table 4 averages this index over the months of the child's gestation.

References

Abatzoglou, J. T., Dobrowski, S. Z., Parks, S. A., and Hegewisch, K. C. (2018). Terraclimate, a high-resolution global dataset of monthly climate and climatic water balance from 1958–2015. *Scientific Data*, 5(1):170191.

Alderman, H., Hoddinott, J., and Kinsey, B. (2006). Long term consequences of early childhood malnutrition. *Oxford Economic Papers*, 58(3):450–474.

Almond, D. and Currie, J. (2011). Killing me softly: The fetal origins hypothesis. *Journal of Economic Perspectives*, 25(3):153–172.

Almond, D., Currie, J., and Duque, V. (2018). Childhood circumstances and adult outcomes: Act ii. *Journal of Economic Literature*, 56(4):1360–1446.

Andrée, B. P. J. (2021). Estimating food price inflation from partial surveys. Policy Research Working Paper 9886, World Bank, Washington, DC.

Barker, D. J. (1990). The fetal and infant origins of adult disease. *BMJ*, 301(6761):1111–1111.

Burgert, C. R., Colston, J., Roy, T., and Zachary, B. (2013). Geographic displacement procedure and georeferenced data release policy for the demographic and health surveys. Dhs spatial analysis reports no. 7, ICF International, Calverton, Maryland, USA.

Chatterjee, S. (2022). How hard did that sting? estimating the economic costs of locust attacks on agricultural production†. *Applied Economic Perspectives and Policy*, 44(1):434–459.

Conte, B., Piemontese, L., and Tapsoba, A. (2023). The power of markets: Impact of desert locust invasions on child health. *Journal of Health Economics*, 87:102712.

De Vreyer, P., Guilbert, N., and Mesplé-Somps, S. (2015). Impact of natural disasters on education outcomes: Evidence from the 1987–89 locust plague in Mali. *Journal of African Economies*, 24(1):57–100.

Didan, K. (2015). Mod13q1 modis/terra vegetation indices 16-day l3 global 250m sin grid v006.

EnviroNews Nigeria (2021). Kenya says desert locust invasion under control amid timely mitigation measures. EnviroNews Nigeria. Reporting on a press briefing by Peter Munya, Cabinet Secretary, Ministry of Agriculture, Livestock, Fisheries and Cooperatives, Kenya.

Famine Early Warning Systems Network (FEWS NET) (2011). Kenya livelihood zone descriptions. Technical report, Famine Early Warning Systems Network, Washington, DC.

Food and Agriculture Organization of the United Nations (FAO) (2022a). Desert locust information service (DLIS): RAMSES locust observation database. FAO Locust Hub, https://locust-hub-hqfao.hub.arcgis.com/. Accessed via the FAO Locust Hub public data portal. Underlying data are collected by national Locust Information Officers using the Reconnaissance and Management System of the Environment of Schistocerca (RAMSES) and transmitted to FAO DLIS for analysis and archival.

Food and Agriculture Organization of the United Nations (FAO) (2022b). FAO desert locust campaign: East africa upsurge suppressed, kenya free from desert locust. FAO Newsroom. Interview with Cyril Ferrand and Carla Mucavi.

Funk, C., Peterson, P., Landsfeld, M., Pedreros, D., Verdin, J., Shukla, S., Husak, G., Rowland, J., Harrison, L., Hoell, A., and Michaelsen, J. (2015). The climate hazards infrared precipitation with stations—a new environmental record for monitoring extremes. *Scientific Data*, 2(1):150066.

geoBoundaries (2022). Kenya — subnational administrative boundaries. geoBoundaries Global Database of Political Administrative Boundaries, distributed via the Humanitarian Data Exchange. ADM0–ADM3 boundary layers; Kenya counties correspond to ADM1.

Hoddinott, J. and Kinsey, B. (2001). Child growth in the time of drought. *Oxford Bulletin of Economics and Statistics*, 63(4):409–436.

Hogan, D. and Schlenker, W. (2024). Empirical approaches to climate change impact quantification. In Barrage, L. and Hsiang, S., editors, *Handbook of the Economics of Climate Change*, volume 1, pages 53–111. North-Holland.

Humanitarian OpenStreetMap Team (HOT) and OpenStreetMap contributors (2022). Kenya roads (openstreetmap export). Humanitarian OpenStreetMap Team Export Tool, https://export.hotosm.org/.

Josephson, A., Michler, J. D., Kilic, T., and Murray, S. (2026). The mismeasure of weather: Using earth observation data for estimation of socioeconomic outcomes. *Journal of Development Economics*, 178:103553.

Kenya National Bureau of Statistics (KNBS) and ICF (2023). Kenya demographic and health survey 2022. Technical report, KNBS and ICF, Nairobi, Kenya, and Rockville, Maryland, USA.

Kenya National Bureau of Statistics (KNBS) and ICF International (2015). Kenya demographic and health survey 2014. Technical report, KNBS and ICF International, Nairobi, Kenya, and Rockville, Maryland, USA.

Le, K. and Nguyen, M. (2022). Desert locust swarms and child health. *Economics & Human Biology*, 44:101094.

Leroy, J. L., Ruel, M., Habicht, J.-P., and Frongillo, E. A. (2014). Linear growth deficit continues to accumulate beyond the first 1000 days in low- and middle-income countries: Global evidence from 51 national surveys. *The Journal of Nutrition*, 144(9):1460–1466.

Liu, X., Zhang, D., and He, X. (2024). Unveiling the role of climate in spatially synchronized locust outbreak risks. *Science Advances*, 10(7):eadj1164.

Maccini, S. and Yang, D. (2009). Under the weather: Health, schooling, and economic consequences of early-life rainfall. *American Economic Review*, 99(3):1006–1026.

Michalopoulos, S. and Papaioannou, E. (2018). Spatial patterns of development: A meso approach. *Annual Review of Economics*, 10:383–410.

Michler, J. D., Josephson, A., Kilic, T., and Murray, S. (2022). Privacy protection, measurement error, and the integration of remote sensing and socioeconomic survey data. *Journal of Development Economics*, 158:102927.

Ministry of Agriculture, Livestock, Fisheries and Cooperatives, Kenya (2021). Government ready to combat second wave of desert locust invasion. Press release, Kilimo House, Nairobi.

Peng, W., Ma, N. L., Zhang, D., Zhou, Q., Yue, X., Khoo, S. C., Yang, H., Guan, R., Chen, H., Zhang, X., Wang, Y., Wei, Z., Suo, C., Peng, Y., Yang, Y., Lam, S. S., and Sonne, C. (2020). A review of historical and recent locust outbreaks: Links to global warming, food security and mitigation strategies. *Environmental Research*, 191:110046.

Salih, A. A. M., Baraibar, M., Mwangi, K. K., and Artan, G. A. (2020). Climate change and locust outbreak in east africa. *Nature Climate Change*, 10:584–585.

Senay, G. (2004). Crop water requirement satisfaction index (WRSI) model description. Technical report, U.S. Geological Survey, Famine Early Warning Systems Network.

Shah, M. and Steinberg, B. M. (2017). Drought of opportunities: Contemporaneous and long-term impacts of rainfall shocks on human capital. *Journal of Political Economy*, 125(2):527–561.

Stein, Z. and Susser, M. (1975). The dutch famine, 1944–1945, and the reproductive process. i. effects on six indices at birth. *Pediatric Research*, 9(2):70–76.

Symmons, P. M. and Cressman, K. (2001). Desert locust guidelines: Biology and behaviour. Technical report, Food and Agriculture Organization of the United Nations, Rome.

Tollefsen, A. F., Strand, H., and Buhaug, H. (2012). PRIO-GRID: A unified spatial data structure. *Journal of Peace Research*, 49(2):363–374.

Vicente-Serrano, S. M., Beguería, S., and López-Moreno, J. I. (2010). A multiscalar drought index sensitive to global warming: The standardized precipitation evapotranspiration index. *Journal of Climate*, 23(7):1696 – 1718.

Victora, C. G., Adair, L., Fall, C., Hallal, P. C., Martorell, R., Richter, L., and Sachdev, H. S. (2008). Maternal and child undernutrition: consequences for adult health and human capital. *The Lancet*, 371(9609):340–357.

Victora, C. G., de Onis, M., Hallal, P. C., Blössner, M., and Shrimpton, R. (2010). Worldwide timing of growth faltering: Revisiting implications for interventions. *Pediatrics*, 125(3):e473–e480.

WHO Multicentre Growth Reference Study Group (2006). WHO child growth standards based on length/height, weight and age. *Acta Paediatrica (Oslo, Norway: 1992). Supplement*, 450:76–85.

Zhang, L., Lecoq, M., Latchininsky, A., and Hunter, D. (2019). Locust and grasshopper management. *Annual Review of Entomology*, 64(Volume 64, 2019):15–34.