

Artisanal Mining and the Hollowing of the Formal Sector: Evidence from Ghana's Galamsey Zones

Albert Duodu

Aalto University, Department of Economics

albert.duodu@aalto.fi

Michael Danquah

GIMPA/UNU-WIDER

mldaqu@gmail.com

First draft: May 2026

This version: October 22, 2026

Abstract

Informal resource booms can reshape local economies in ways that formal-sector statistics miss. We study Ghana's artisanal gold boom (galamsey), which employs roughly a million people. We measure each district's exposure using the interaction between its share of land on gold-bearing Birimian bedrock and the world gold price. Gold-suitable districts have about 19 percent fewer formal establishments than other districts in the same region, a cross-sectional gap concentrated in agriculture-linked sectors. Household surveys show a distinct change in male employment: between 2003 and 2014, as the gold price more than tripled, the male employment rate in fully gold-suitable districts rose by about 6 percentage points relative to other districts, mostly in agricultural and manual occupations. We find no differential trend in the 1990s. The establishment gap is larger in wetter districts and absent in drought-prone ones, a pattern consistent with agriculture being an important source of labour for mining. Night-lights in exposed districts fell by about 9 percent after the 2017 crackdown on illegal mining, although the estimate is imprecise. Taken together, the findings describe gold-belt districts with fewer formal establishments but a relative rise in male employment, a combination that establishment counts alone would not reveal.

Keywords: Artisanal mining; Informality; Labour reallocation; Structural transformation; Resource booms; Ghana

JEL codes: O13, Q32, R11, J22.

1 Introduction

Roughly a million Ghanaians work in artisanal and small-scale gold mining, locally called *galamsey*, and the sector produces about a third of the country's gold. Most of this activity is unlicensed and untaxed, and almost none of it appears in establishment surveys or tax records. That creates a problem for anyone asking what an informal resource boom does to a local economy. If workers leave farms and formal firms for the pits, the formal sector can shrink while total activity does not.

Most evidence on the local effects of resource booms comes from large, formal mines, where the boom reaches the local economy through wages, procurement and public revenue. Artisanal mining differs on each of these margins: it is labour intensive, decentralised and untaxed, and much of what it produces is never recorded. Studying it raises two problems. Mining sites are rarely recorded, so exposure has to be measured without observing extraction; and much of the activity the boom creates is invisible to the surveys used to track local economic development.

In this paper, we ask how Ghana's artisanal gold boom changed the formal sector and where the workers went. Our exposure measure interacts the share of each district's land on gold-bearing Birimian bedrock, a formation more than two billion years old, with the world gold price, which Ghana takes as given ([Girard et al., 2025](#); [Borusyak et al., 2025](#)). Put simply, the same rise in the gold price raises the return to digging in districts on the gold belt and leaves it unchanged in districts off it. We take this measure to three sources. A ten percent public sample of Ghana's 2024 Integrated Business Establishment Survey (IBES) covers about 184,000 establishments in all 260 districts, but at a single date. Five waves of the Demographic and Health Surveys (DHS), from 1993 to 2014, cover the same districts before and after the 2008 boom, which lets us test for differential trends. A harmonised night-lights series spans the 2017 Operation Vanguard crackdown.

The main findings are as follows. First, the formal sector is thinner in the gold belt. A fully gold-suitable district hosts about 19 percent fewer formal establishments than a non-suitable district in the same region. The gap is largest in agriculture-linked sectors, where establishment counts are about 35 percent lower; in mining-adjacent sectors it is about 11 percent and imprecise. The gap holds when we compare districts of similar size and market access, and it does not reflect establishments slipping out of the survey frame. Because geology is fixed and the survey is a single cross-section, we present it as a fact to be explained rather than as a causal estimate.

Second, the household surveys show where the workers went. Between the 2003

and 2014 survey waves, as the gold price more than tripled, the male employment rate in a fully gold-suitable district rose by about 6 percentage points relative to a non-suitable district, against a mean of 78 percent. The new work was agricultural and manual. In the 1990s, before mechanised extraction arrived, the gradient was flat. Labour did not leave the market. The pattern is consistent with workers moving into informal extraction and related work that establishment surveys do not record, although the DHS are repeated cross-sections with no mining occupation code, so they do not observe individual transitions into mining. This is the mechanism we propose for the firm-side fact, and it is also our best-identified result.

Third, the pattern sits where the mechanism predicts. The formal-sector gap is concentrated in wet, agriculturally productive districts and disappears in drought-prone districts, consistent with mining competing with farming for labour. Soils, terrain and market access could each produce a firm gradient across the gold belt, but none of them has an obvious reason to weaken in dry districts. Finally, night-lights in exposed districts fell by about 9 percent after Operation Vanguard began in 2017. The estimate is imprecise and shrinks when 2020 is dropped, so it is only suggestive evidence that informal extraction had become part of local economic activity.

Admittedly, the design has limits, which we discuss in detail below. The firm data cannot distinguish exit from entry that never happened, so they say nothing direct about whether the least productive firms left. The DHS occupation codes do not separate mining from other manual or agricultural work. And the suitability measure records where mining is possible, not where it took place.

A simple model of a district labour market organises the evidence (Section 3). Households divide their time between farming, formal wage work, artisanal mining and home production, and formal firms pay a fixed cost to operate. A higher return to mining bids up the local wage, raises participation, and pushes the least productive formal firms below the survival cutoff, most of all where a large agricultural sector supplies the labour. The model also predicts larger surviving firms. We find this within sectors, where the difference is about 2.5 percent, or less than a tenth of a worker per establishment, but not in district mean size. A decomposition of formal employment points the same way: establishment counts and total employment fall by almost the same amount, so displaced workers do not reappear in surviving firms. The firm evidence therefore does not point to productivity-enhancing selection.

The paper's main lesson concerns measurement. Studies of structural transformation and local development typically track change through formal establishments, reg-

istered employment and tax records. Where a boom runs through informal extraction, those measures can show a thinning formal sector at the same time as employment rises, so formal-sector indicators alone give an incomplete account of how a local economy responds. This complements evidence from industrial mines, which raise local incomes through backward linkages in Peru (Aragón and Rud, 2013), produce agglomeration rather than Dutch disease in US counties (Allcott and Keniston, 2018), shift women from agriculture into services across Africa (Kotsadam and Tolonen, 2016), and lower agricultural productivity through pollution in Ghana (Aragón and Rud, 2016). Those effects run through channels that formal statistics record. Bazillier and Girard (2020) compare the distributive effects of artisanal and industrial gold rushes in Burkina Faso, and Girard et al. (2025) build the continental suitability measure we use and document the environmental footprint of artisanal mining. We add the formal-sector side of an informal boom and the occupational shifts that accompany it. For policy, the paper offers a first quantitative look at the local economic footprint of anti-galamsey enforcement, which Ghana has used repeatedly since 2013 but rarely evaluated.

Section 2 describes the setting, Section 3 sets out the conceptual framework, and Sections 4 and 5 present the data and empirical strategy. Section 6 documents the formal-sector fact, Section 7 the household mechanism, Section 8 heterogeneity by agricultural conditions, and Section 9 the enforcement evidence. Section 10 concludes.

2 Background

2.1 Artisanal gold mining in Ghana

Artisanal and small-scale gold mining (ASGM) in Ghana, locally called galamsey, refers to small-group, low-capital gold extraction operating largely outside the licensing system (Hilson, 2017). Estimates put national employment at about one million workers and the share of national gold output at roughly thirty-five percent. Legal small-scale mining contributed five billion dollars to Ghana’s gold exports in 2024 according to the Ghana Gold Board; informal output is by definition unrecorded but is widely believed to be of comparable magnitude.

Galamsey activity has two structural features that matter for our analysis. First, it is geographically anchored to the gold-bearing geology of southern Ghana, a band that stretches through the Western, Western North, Ashanti, Ahafo, Bono, Bono East, Eastern, and Central regions. Second, the technology of extraction changed around 2008, when an estimated fifty thousand Chinese miners arrived with mechanised equipment

(excavators, dredges and changfa pumps) and converted what had been a low-capital activity into a much more land-intensive one (Crawford and Botchwey, 2017). Mercury contamination of rivers, sediment loading, soil degradation, and loss of cocoa-belt forest cover followed. Figure 1 plots the world gold price alongside the three Ghanaian policy events that anchor our analysis. Annual forest loss in galamsey-prone districts rose from about 106 hectares per district in 2001–07 to 353 in the 2010s, a factor of 3.3 against 2.0 in the rest of Ghana (Figure A.2 in the appendix).




Figure 1: World gold price 1990–2025 with key Ghanaian policy events. Source: World Bank Pink Sheet.

2.2 Gold geology

Gold in Ghana is hosted by the Birimian Greenstone Belt, a Precambrian formation that crystallised between 2.1 and 2.2 billion years ago (Girard et al., 2025). Mineralisation follows specific stratigraphic and lithological bands that are fixed at geological time-scales. Following Girard et al. (2025), we use the share of each district’s land that overlaps an ASGM-suitable bedrock layer as our measure of gold suitability, $Suit_d \in [0, 1]$. Roughly two thirds of Ghana’s 260 districts have non-zero suitability, and the top quartile are the southwest gold-belt districts. Because the bedrock has been in place for two billion years, $Suit_d$ predates any human economic choice. Figure 2 shows the spatial distribution of $Suit_d$ across Ghana’s 260 districts.

2.3 Enforcement

Anti-galamsey enforcement has cycled through three distinct episodes since 2010. In 2013 a Mahama administration task force deported around 4,500 Chinese miners but failed to dent domestic galamsey activity. In July 2017 the Akufo-Addo administration




Figure 2: Geological suitability for artisanal gold mining across Ghana’s 260 districts. $Suit_d$ is the share of district d ’s land area falling inside the gold-suitability polygon of [Girard et al. \(2025\)](#), which we intersect with the 2021 Ghana Statistical Service district boundaries. The classes are cut around the two mass points in the distribution: 41 districts contain no suitable ground at all and 57 are suitable throughout. Grey lines are the sixteen regions that absorb the fixed effects in every specification. Districts outlined in dark blue are the eight large-scale mining anchors, which the main analysis sample excludes so that the gradient is identified off artisanal rather than industrial extraction. The inset enlarges Greater Accra, whose municipalities are too small to read at national scale. Map lines delineate study areas and do not necessarily depict accepted national boundaries.

launched Operation Vanguard, a military-led joint task force of about four hundred soldiers operating from forward bases in the Ashanti, Eastern, and Western regions ([Hilson, 2017](#)). By February 2018 over a thousand miners had been arrested and equipment seized. In April 2021 the same administration launched Operation Halt II, a deeper crackdown extending to riverine sites. We use Operation Vanguard as our policy break: it is national in scope and has a sharp public start date. Operation Halt II is later than the end of our night-lights series and therefore falls outside our estimation window.

3 Conceptual Framework

In this section we set out a model of a district economy in which households allocate labour between agriculture, formal wage work, and artisanal mining, and formal firms enter subject to a fixed cost. Its purpose is to isolate the forces through which a mining boom acts on the formal sector, and to derive the specifications we take to the data.

Corden and Neary (1982) distinguish two routes by which a resource boom reshapes the rest of an economy. A *spending effect* operates through demand: rents accrue, expenditure rises, and non-traded prices follow. A *resource-movement effect* operates through factor markets, as the booming sector bids inputs away from other activities. Which route dominates is a question about the technology of extraction. Large-scale mining is capital intensive and its rents accrue to firms and the state, so it works largely through the first. Artisanal gold mining reverses each of those features: it is labour intensive, decentralised, and almost entirely untaxed.

Our model traces the second route. In it, a rise in gold suitability acts on formal firms through two forces. The *labour-withdrawal effect* operates through the wage: mining bids for workers, the local wage rises, and the productivity a firm needs to cover its fixed cost rises with it. The *degradation effect* operates through productivity: extraction damages land and water, and lowers the output of the firms that remain. The two act on different margins. Our evidence speaks mainly to the first; we report forest-loss evidence on the second descriptively in Section OA.7.

Allcott and Keniston (2018) study resource booms in modern America and find agglomeration: local manufacturing grew alongside extraction. Their booms raise demand in a labour market that is thick, mobile, and formal, so factor competition is weak and the spending channel dominates. Ours operate where the competing sector is smallholder agriculture and labour moves little across districts. Which channel dominates thus depends on the elasticities of labour supply and of agricultural production, and Proposition 3 converts those elasticities into a testable restriction: if factor competition is the operative channel, the formal-sector contraction must disappear where there is little agriculture to draw labour from. Bazillier and Girard (2020) study the distributional consequences of gold rushes in Burkina Faso and come closest to our setting; we differ in observing the distribution of formal establishments and in exploiting an enforcement shock.

3.1 Environment

A district d contains a unit continuum of households. Each allocates one unit of labour across formal wage employment (F), agriculture (A), artisanal mining (M), and home production (O). Sector payoffs are

$$\pi_{dt}^A = A_d \theta_{dt} (L^A)^\alpha, \quad 0 < \alpha < 1, \quad (1)$$

$$\pi_{dt}^M = S_d P_t (L^M)^\beta - \kappa(e_t) L^M, \quad 0 < \beta < 1, \quad (2)$$

$$\pi_{dt}^F = w_{dt} L^F, \quad \pi_{dt}^O = \frac{\nu}{1 + 1/\varepsilon} (L^O)^{1+1/\varepsilon}, \quad \varepsilon > 0. \quad (3)$$

A_d is land productivity, θ_{dt} a climate shock with $\mathbb{E}\theta_{dt} = 1$, and $S_d \equiv \text{Suit}_d \in [0, 1]$ the gold-suitable share of the district. Three properties of this notation carry the identification: S_d is Birimian bedrock and so predates any economic choice; P_t is a world price Ghana takes as given; and $\kappa(e_t)$, rising in enforcement intensity e_t , is what Operation Vanguard shifts in 2017.

Home production carries a finite supply elasticity ε rather than a constant reservation payoff. Under a constant reservation utility the marginal household is indifferent and L^O is indeterminate, which would leave participation an accounting residual; with $\varepsilon < \infty$ it becomes an equilibrium object that responds to the wage. Optimal allocation equates marginal returns across active sectors,

$$\alpha A_d \theta_{dt} (L^A)^{\alpha-1} = \beta S_d P_t (L^M)^{\beta-1} - \kappa(e_t) = w_{dt} = \nu (L^O)^{1/\varepsilon}, \quad (4)$$

so that, inverting,

$$L_{dt}^M = \left(\frac{w_{dt} + \kappa(e_t)}{\beta S_d P_t} \right)^{\frac{1}{\beta-1}}, \quad L_{dt}^A = \left(\frac{w_{dt}}{\alpha A_d \theta_{dt}} \right)^{\frac{1}{\alpha-1}}, \quad L_{dt}^O = \left(\frac{w_{dt}}{\nu} \right)^\varepsilon. \quad (5)$$

3.2 Firm entry and selection

Potential entrepreneurs draw productivity z from a Pareto distribution with shape $k > 1$ on $[\underline{z}, \infty)$, so $1 - G(z) = (\underline{z}/z)^k$. A firm of type z hiring n workers produces $z \cdot n$, sells at price one, and pays a fixed operating cost f . Mining degrades land and water, which we model as a productivity drag $1 - \delta(L_{dt}^M)$ common to firms in the district, with $\delta' > 0$ and $\delta(0) = 0$. A firm operates if

$$z (1 - \delta(L_{dt}^M)) \geq z_{dt}^* \equiv w_{dt} + f/n_{dt}^*, \quad (6)$$

with $n_{dt}^*(z) = (z(1 - \delta) - w_{dt})^\eta$, $\eta > 0$. The count of establishments and mean surviving size are then

$$\log N_d = \text{const} - k \log \frac{z_{dt}^*}{1 - \delta(L_{dt}^M)}, \quad \bar{n}_{dt}^* = L_d^F / N_d. \quad (7)$$

3.3 Equilibrium

The wage clears the local labour market, $L_d^F = 1 - L_d^A - L_d^M - L_d^O$, with L_d^F the labour demanded by surviving firms in (7). Everything the paper tests rests on one monotonicity, which we state as a result rather than assume.

Proposition 1 (Wage response) *Suppose mining is active ($L_{dt}^M > 0$) and formal labour demand is non-increasing in the wage. Then $\partial w_{dt} / \partial (S_d P_t) > 0$ and $\partial L_{dt}^M / \partial (S_d P_t) > 0$.*

Proof (Sketch) Totally differentiate the clearing condition. A rise in $S_d P_t$ raises the marginal return to mining in (4) at the prevailing wage, so mining labour demand shifts out. Since L^A , L^O and L^F are each non-increasing in w by (5) and by assumption, and L^M is decreasing in w for $\beta < 1$, aggregate excess demand for labour is strictly decreasing in w . The implicit function theorem then gives $dw / d(S_d P_t) = -(\partial \text{ED} / \partial S_d P_t) / (\partial \text{ED} / \partial w) > 0$, and substituting back into (5) signs L^M .

The condition that formal labour demand not slope upward is the only substantive restriction, and it fails only if the selection effect on the extensive margin is strong enough to outweigh the intensive margin. Three predictions follow.

Proposition 2 (Two margins) $\partial \log N_d / \partial S_d < 0$ and $\partial \log \bar{n}_{dt}^* / \partial S_d > 0$.

This is the labour-withdrawal effect. Both signs come from the same term in (6): the wage enters the entry cutoff directly, so a higher w_{dt} raises the productivity a firm must have to cover f and truncates the operating distribution from below. Establishment counts fall because the marginal firm exits, and mean size rises because the firms that remain are drawn from further up the productivity distribution. Two things this does not establish. The exiting firm is by construction the least productive, so *fewer but larger* is compositional and is what a misallocation accounting would score as an improvement in average formal productivity. Calling it a hollowing is a claim about externalities (mercury, deforestation, a forgone tax base, and human capital that informal extraction does not build), and the selection margin does not deliver it. And because (7) governs a stock, a single cross-section identifies a gradient in establishment counts, not a flow of exits.

Proposition 3 (Climate scope condition) $\partial^2 \log N_d / \partial S_d \partial \theta_d > 0$. Where θ_d is persistently low, (5) makes L_d^A small at baseline, so a given increase in $S_d P_t$ transfers less labour and displaces fewer firms.

Proposition 3 says where the effect should be *absent*, which is what makes it useful. Soil, terrain, settlement history, and market access all differ between the gold belt and the rest of Ghana, and each could produce a firm-count gradient on its own. None of them has a reason to weaken in dry districts. A gradient that is uniform across the climate distribution is therefore consistent with those confounds and inconsistent with the mechanism here.

Proposition 4 (Policy reversal and participation) $\kappa'(e_t) > 0$ implies $\partial L_{dt}^M / \partial e_t < 0$, hence $\partial w_{dt} / \partial e_t < 0$ and $\partial z_{dt}^* / \partial e_t < 0$: enforcement returns labour to formal employment. Separately, from (5), $\partial L_{dt}^O / \partial S_d < 0$ whenever $\partial w_{dt} / \partial S_d > 0$, so measured labour-force participation $1 - L^O$ rises with gold suitability for any finite ε .

Aggregate local activity is the sum of three of these terms and its sign depends on their relative size. Mining output $S_d P_t (L^M)^\beta$ enters positively. The labour-withdrawal effect enters negatively through the formal wage bill lost to exit, and the degradation effect enters negatively through δ , with a lag, since land degrades cumulatively. The net is positive when the marginal return to mining exceeds the average product of the exiting firm, a comparison (4) does not restrict. Existing evidence does not settle it either. Eshun (2026) finds that Ghanaian cocoa farmers in communities reached by artisanal mining lose yield but raise consumption, which is the configuration our decomposition allows when the mining term dominates. We therefore estimate the aggregate response rather than predict it, using the enforcement shift in $\kappa(e_t)$ studied in Section 9.

Proposition 1 concerns the equilibrium after the boom. Every prediction that follows therefore rests on a premise about the period before it: that gold-suitable districts were not already paying higher wages for reasons unrelated to mining. That premise is testable. Living-standards data collected in 2005/06, two years before mechanised extraction arrived, measure earnings on the same geological gradient, and we take the prediction to them in Section OA.9.

3.4 From the model to the estimating equations

Equation (7) delivers the specification directly. Taking logs and differencing across districts within a region, and noting that w_{dt} is an increasing function of $S_d P_t$ by Propo-

sition 1, gives

$$\log N_d = \beta \text{Suit}_d + \mu_{r(d)} + \varepsilon_d, \quad \beta = -k \frac{\partial \log w_d}{\partial \text{Suit}_d}, \quad (8)$$

where the region effects μ_r , for the region r containing district d , absorb the common component of z , f , and P_t , and ε_d collects district-specific variation in those objects. The regression coefficient is thus a product of two things the model separates: the Pareto shape, which is a property of the establishment-size distribution, and the wage response, which is the object of economic interest. Section 5.1 states the estimating equations in the form we take to the data, and Table 1 lists which specification tests which proposition.

Table 1: From predictions to specifications

	Object	Specification
Proposition 1	$\partial w_d / \partial S_d = 0$ before the boom	GLSS5 cross-section (Section OA.9)
Proposition 2	$\partial \log N_d / \partial S_d, \partial \log \bar{n} / \partial S_d$	(10), (11)
Proposition 3	$\partial^2 \log N_d / \partial S_d \partial \theta_d$	(13)
Proposition 4	$\partial^2 \log \text{NTL} / \partial S_d \partial e_t; \partial(1 - L^O) / \partial S_d$	(14); (12)

One prediction resists a clean test. Proposition 4 concerns formal firm counts under enforcement, and the firm data are a single 2024 cross-section, so the pre- and post-Vanguard comparison it calls for cannot be made. The available outcome is aggregate activity, which the model implies but does not uniquely sign, since a crackdown lowers mining income and raises formal employment at once. We therefore treat the enforcement estimate as a magnitude for the policy and not as a verdict on Proposition 4.

4 Data

Ghana is divided into 16 administrative regions and, on the 2021 boundaries, 260 districts. In geographic data these are the first- and second-level administrative units, labelled ADM1 and ADM2. We use “region” and “district” throughout the text; the ADM1 and ADM2 labels appear only in the titles and legends of some figures.

4.1 Establishments

Our firm-side data come from the 2024 Integrated Business Establishment Survey (IBES), administered by the Ghana Statistical Service. We use the public ten-percent sample, which contains 184,180 firms with district identifiers, ISIC Rev. 5 sector codes, and persons-engaged employment counts. The sample covers all 260 administrative districts and 82 two-digit sectors. We aggregate the firm-level data to the district level for

the establishment-count regressions and use the firm-level micro for the size-margin regressions. The 2013 IBES wave is available but uses a different sample frame (24,362 firms, 10 regions) and is not directly comparable; we therefore use the 2024 wave as a single cross-section.

4.2 Households

Household labour-force outcomes come from six waves of the Demographic and Health Surveys (DHS) in Ghana: 1993, 1998, 2003, 2008, 2014, and 2022. Each wave provides individual-level records for women aged 15–49 (the IR file) and a corresponding men’s individual recode for ages 15–59 (the MR file). We use the v717 occupation variable to classify each respondent as working, in agriculture, in manual labour, or in services. A respondent is coded as working if she or he reports work in the twelve months before the survey. Agricultural occupations are the v717 categories for agricultural self-employment and agricultural employees; manual occupations are skilled and unskilled manual work. We call the union of the two *physical* occupations. The DHS occupation codes do not separate mining from other manual or agricultural work, so these outcomes measure the occupational margin that mining draws on rather than mining employment itself. The 2022 wave (DHS-8) recoded the v717 categories in a way that breaks comparability with the earlier waves; our preferred sample is therefore 1993–2014 and we report 2022 only as a robustness check. Cluster GPS coordinates, jittered by up to 2 km in urban and 5 km in rural areas, are spatially joined to districts. Total individual-year observations on the preferred sample are 28,953 women and 16,537 men across 2,050 cluster-year observations.

4.3 Night-lights and forest loss

Aggregate economic activity at the district level is measured by the harmonised DMSP+VIIRS night-lights series of [Li et al. \(2020\)](#), which extends from 1992 to 2020 at a 30 arc-second grid resolution. We aggregate to districts by zonal sum and take the natural logarithm of 1+ the sum. The night-lights panel has 7,540 district-year observations across the 260 districts. Forest loss comes from the Hansen Global Forest Change v1.11 product ([Hansen et al., 2013](#)), extracted at district level over 2001–2023 with the standard 25% canopy-cover threshold. We use both the annual loss area in hectares and the cumulative loss share over four time windows.

4.4 Exposure to artisanal mining

The gold-suitability measure $Suit_d$ is taken directly from the AsGM map of [Girard et al. \(2025\)](#) and aggregated from PRIO-GRID half-degree cells ([Tollefsen et al., 2012](#)) to districts. The world gold price P_t is the World Bank Pink Sheet annual nominal series in U.S. dollars per troy ounce.

Our measure of artisanal mining is the interaction of the two,

$$ASGM_{dt} = Suit_d \times \log P_{t-1}, \quad (9)$$

lagged one year to allow households time to adjust. This is the proxy constructed by [Girard et al. \(2025\)](#) for the African continent, which they describe as indicating the potential value of artisanal gold mining in a given cell and year. It measures potential rather than realised activity: it says what mining is worth where the geology permits it, not where pits were dug. [Girard et al. \(2025\)](#) validate the proxy against independent records of ASGM occurrence, gold hints, and mercury emissions; we report the corresponding check for Ghana in Section [OA.4](#). Following [Borusyak et al. \(2025\)](#), the share component $Suit_d$ is exposure-weighted and predetermined by geology, while the shock $\log P_{t-1}$ is nationally exogenous, since Ghana is a price-taker on the London bullion market.

4.5 Climate

Rainfall data come from CHIRPS v2.0 at 0.05° resolution 1981–2024, aggregated to district by mean, and air temperature from ERA5 monthly reanalysis 1980–2024. We construct a district-year rainfall z-score using the 1990–2010 baseline, and binary indicators for severe drought ($z < -1$) and severe flood ($z > +1$). For cross-section interactions we use the share of years 2008–2023 in each climate state.

4.6 Sample and descriptive statistics

Our analysis sample consists of 252 districts after dropping the 8 large-scale-mine (LSM) anchor districts where industrial mining dominates: Tarkwa Nsuaem, Prestea/Huni Valley, Wassa Amenfi East, Bibiani-Anhwiaso-Bekwai, Atwima Mponua, Asutifi South, Obuasi, and Talensi. We drop these because the suitability measure conflates artisanal and industrial exposure in cells that host one of Ghana’s eight industrial gold mines. The remaining 252 districts retain the full range of $Suit_d \in [0, 0.99]$ and provide the

ASM-only identifying variation. Table 2 reports summary statistics for the analysis sample.

Table 2: Summary statistics for the analysis sample (drop-LSM)

Variable	Mean	SD	Min	p25	Median	p75	N
<i>District-level</i>							
$Suit_d$	0.69	0.39	0.00	0.30	0.93	0.99	252
Firms (IBES 2024 sample)	698.6	654.4	10	280	500	882	252
Total persons engaged	2,340	2,621	21	931	1,647	2,944	252
Mean firm size (persons)	3.21	1.04	1.6	2.6	3.0	3.6	252
Forest loss 2017–2023 (ha)	6,210	9,820	1	321	2,410	7,930	244
Drought freq. 2008–23 (share)	0.10	0.07	0.0	0.06	0.13	0.13	250
Flood freq. 2008–23 (share)	0.07	0.06	0.0	0.00	0.06	0.13	250
<i>District $\times$ year (NTL panel, 2001–2020)</i>							
$\log(1+\text{NTL sum})$	6.12	1.96	0.0	4.9	6.5	7.6	4,720
$\log(1+\text{loss area, ha})$	4.84	2.13	0.0	3.5	4.9	6.4	4,720
$ASGM_{dt}$ exposure	4.64	2.84	0.0	1.7	5.0	7.0	4,720
<i>Firm-level (IBES 2024)</i>							
$\log$ persons engaged	0.86	0.84	0.0	0.0	0.7	1.4	168k
<i>DHS cluster $\times$ year, 1993–2014</i>							
Share working	0.78	0.20	0.0	0.7	0.84	0.93	2,050
Share manual or ag (physical)	0.51	0.27	0.0	0.30	0.55	0.74	2,050

Notes: Drop-LSM analysis sample. Firm and household panels are independent. Standard errors are clustered by region in district regressions.

5 Empirical Strategy

5.1 Establishment counts and size

Our firm data are a single cross-section. IBES 2024 records establishments at one date, so nothing in it identifies entry or exit, and the gold price is national in any given year, which means the shift-share shifter $Suit_d \times \log P_t$ collapses to a constant times $Suit_d$. What the firm data can deliver is the cross-district gradient implied by (8), and we estimate it on the two margins the model separates. At the district level,

$$\log N_d = \beta Suit_d + \mu_{r(d)} + \varepsilon_d, \quad (10)$$

and at the establishment level,

$$\log Y_i = \beta Suit_{d(i)} + \mu_{r(d(i))} + \gamma_{\text{ISIC2}(i)} + \varepsilon_i, \quad (11)$$

where $\log Y_i$ is log persons engaged in establishment i , μ_r are 16 region fixed effects, and γ_{ISIC2} are 82 two-digit sector dummies. We drop the eight industrial-mine anchor

districts so the gradient reflects artisanal exposure. Standard errors are clustered by region in (10) and at districts in (11).

The principal concern with (10) and (11) is that Birimian geology determines more than gold. The belt runs through the forest zone of southern Ghana, where soils, rainfall, colonial infrastructure, and settlement density all differ from the savanna north, and each of them predicts where formal establishments locate. Region fixed effects remove the broad north–south comparison, but they leave within-region variation in terrain and market access that is correlated with the gold gradient and fixed in time. With Suit_d constant and the outcome measured once, no feature of these regressions distinguishes gold from the geology that carries it.

We therefore treat the firm cross-section as descriptive. Two features of the pattern nonetheless help discriminate between the labour-withdrawal mechanism and a confound fixed in geology. The model implies that the gap should be concentrated in the sectors that supply labour to mining, which Section 6 examines, and that it should weaken where agriculture is thin, which Section 8 examines against the climate distribution, where a geological confound has no obvious reason to switch off. Neither test turns the gradient into a causal estimate; the household evidence in Section 7 carries the causal claim about labour. Because the district-level regressions cluster on only sixteen regions, we also report wild cluster bootstrap p -values.

5.2 Household labour allocation

Household data supply the time dimension the firm data lack. DHS clusters are sampled in the same districts across waves, which allows the shift-share design that (2) motivates: gold suitability is fixed, the world price moves, and their interaction shifts the return to mining over time. At the cluster-wave level we estimate

$$Y_{ct} = \beta_1 \text{Suit}_{d(c)} + \beta (\text{Suit}_{d(c)} \times \log P_{t-1}) + \mu_{r(d(c))} + \tau_t + u_{ct}, \quad (12)$$

where Y_{ct} is the weighted share of respondents in cluster c and wave t in a given labour-market state, μ_r are region fixed effects, and τ_t are wave fixed effects. The level of Suit_d absorbs any time-invariant difference between gold-suitable and other districts within a region, so β is identified from the change in the suitability gradient as the gold price moves (Borusyak et al., 2025). Standard errors are clustered by district.

Identification rests on two conditions. The Birimian belt was laid down two billion years ago, so Suit_d is unrelated to twenty-first-century shocks (Girard et al., 2025);

and Ghana is a price-taker on the London bullion market, so $\log P_{t-1}$ is exogenous to Ghanaian district-year shocks conditional on τ_t . Neither is sufficient on its own, because $ASGM_{dt} = \text{Suit}_d \times \log P_t$ is a fixed exposure share multiplied by a series that rises almost monotonically over the sample. It is therefore collinear with any trend that differs between the gold belt and the rest of Ghana, and the fixed effects do not absorb such a trend, since it is precisely the variation $ASGM_{dt}$ uses.

This is why the early DHS waves matter. The DHS waves reach back to 1993, fifteen years before mechanised extraction arrived, so (12) can be estimated wave by wave and the pre-boom coefficients give a direct test of the identifying assumption. Section 7 reports that test. The same reasoning explains why we do not take (12) to night-lights: Ghana’s electricity network expanded unevenly and persistently over the same period, there is no comparable pre-period in which the gradient should be flat, and Table OA.8 shows that the night-lights reduced form reverses sign once gold-belt districts are allowed a trend of their own. Aggregate effects are identified instead from the 2017 enforcement break in Section 9.

6 The Formal Sector in Gold-Belt Districts

We begin with the fact the paper sets out to explain: formal establishments are fewer in gold-suitable districts. Because the establishment data are a single cross-section, this section describes where the formal sector is thinner and which sectors account for the gap; Sections 7 and 8 then ask why.

Table 3 reports estimates of (10) and (11). Columns 1 and 2 relate the number of formal establishments in a district to gold suitability, and columns 3 and 4 relate employment per establishment to suitability within two-digit sectors. In the preferred specification, which excludes the eight districts with industrial gold mines, a fully gold-suitable district has 22 log points, or about 19 percent, fewer establishments than a non-suitable district in the same region (column 2). Moving from the 25th to the 75th percentile of suitability lowers the count by 15 log points, and including the industrial-mine districts changes little (column 1).

We next check whether the gap simply reflects district size or market access, since gold-suitable districts are smaller than other districts. Controlling for district area, urban land share, travel time to the nearest city and pre-boom night-lights leaves a shortfall of 13 to 20 percent, and a Poisson model on the raw count gives 17 percent

Table 3: Formal establishments and establishment size across the gold gradient

	log N firms (district)		log persons engaged (firm)	
	Full (1)	Drop LSM (2)	Full (3)	Drop LSM (4)
$Suit_d$	-0.236** (0.105)	-0.216** (0.098)	+0.020* (0.012)	+0.025** (0.010)
Region FE	Yes	Yes	Yes	Yes
ISIC2 sector FE	–	–	Yes	Yes
Observations	260	252	168k	147k

Notes: Cross-section, IBES 2024. $Suit_d$ is the AsGM gold-suitability share (Girard et al., 2025). Drop-LSM removes 8 districts hosting industrial gold mines. Standard errors in parentheses below coefficients, clustered by region in columns 1–2 and at districts in columns 3–4. With 16 region clusters, the wild cluster bootstrap p -value for column 2 is 0.061 (Table A.3).

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

(Table A.4).¹ We also check whether establishments in gold-suitable districts have simply moved out of the survey frame as activity informalised. IBES covers unregistered as well as registered establishments, so a firm that informalises remains in the data, and the registered, micro and paid-employee shares of surviving establishments do not vary with suitability (Table A.5).² Most of the gap reflects fewer establishments within sectors that remain active rather than the disappearance of whole sectors: the number of active two-digit sectors barely changes, and about four-fifths of the decline comes from thinning within sectors (Table A.1). The sectoral mix shifts as well, with agriculture and industry accounting for about 1 and 3 percentage points less of formal activity in a fully suitable district, and services for about 4 points more (Table A.2).

Turning to establishment size, columns 3 and 4 of Table 3 show that, within two-digit sectors, establishments in gold-suitable districts employ about 2.5 percent more people, less than a tenth of a worker at the mean. Fewer but slightly larger establishments could reflect cleansing (Caballero and Hammour, 1994), in which exposure closes the least productive firms and leaves better survivors. We distinguish this from labour withdrawal by asking where the missing employment goes. Under cleansing, workers move from closing firms to surviving ones, so total employment should fall by much less than the number of establishments; under labour withdrawal, workers leave the

¹Districts with suitability above one-half average 656 km², against 1,271 km² below it, and establishments per square kilometre do not vary with suitability (Table A.4, column 5). The comparison is therefore between districts of similar size within a region. With sixteen region clusters, conventional standard errors overstate precision; under the wild cluster bootstrap, the 95 percent confidence set for the baseline estimate runs from a 35 percent shortfall to a 1 percent excess (Table A.3).

²Only 9 percent of IBES establishments are registered with the tax authority, and 91 percent employ fewer than five people.

formal sector and employment falls with the count. Table A.6 supports the second account: establishment counts fall by 21.6 log points and total employment by 20.1, leaving almost nothing for average size. The within-sector size difference reflects a shortage of one-person establishments, and survivors in exposed districts are neither more productive nor older than establishments elsewhere.³

Table 4: Sectoral incidence: the gap is concentrated in agriculture-linked sectors

Sector group	$\log N_d^{(g)}$
Agriculture-linked (crops, food manufacturing, beverages)	-0.433*** (0.081)
Other manufacturing	-0.277 (0.201)
Retail & consumer services (retail, food service, apparel)	-0.178 (0.111)
Public/social (education, health, finance)	-0.166* (0.096)
Mining-adjacent (metals, construction, transport, wholesale)	-0.121 (0.177)
Total	-0.216** (0.098)
Region FE	Yes
Districts	239

Notes: Coefficient on Suit_d in $\log N_d^{(g)} = \beta_g \text{Suit}_d + \mu_r + \varepsilon_d$, drop-LSM. Each row is a separate regression by sector group g (ISIC Rev. 5 two-digit). Standard errors in parentheses, clustered by region. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

We then ask which sectors account for the gap. If artisanal mining draws labour away from activities that rely on the same rural workers and land, the largest shortfalls should appear in those activities, whereas a general local downturn would depress formal activity across the board. We therefore re-estimate (10) separately for five sector groups, and Table 4 reports the results. Agriculture-linked establishments (crops, food manufacturing and beverages) are about 35 percent fewer in a fully suitable district, and this is the only group whose estimate remains statistically significant under the wild cluster bootstrap. For other manufacturing, retail and public services the gaps range from 15 to 24 percent and are less precisely estimated, while mining-adjacent establishments (metals, construction, transport and wholesale) show the smallest gap, about 11 percent, which we cannot distinguish from zero. At the two-digit level, crop production and food manufacturing have the largest negative coefficients, whereas wholesale

³Revenue per worker, built from banded revenue data, does not vary with suitability, although the data can rule out only productivity gains larger than about 20 percent. A cross-section records the stock of survivors rather than the flows of entry and exit, so measuring exit directly would require a linked establishment panel.

trade is slightly positive and food service is flat (Figure A.1). This is the incidence that Proposition 2 attributes to labour withdrawal.

Taken together, the establishment data show where the formal sector is thinner and which sectors account for the gap, but a single cross-section cannot show why, because it does not follow workers. Section 7 turns to the household surveys for that.

7 Labour-Market Adjustment

The labour-withdrawal effect of Section 3 gives a specific answer to why gold-belt districts have fewer formal establishments: a higher return to mining bids workers away from farms and formal firms. The household surveys can test that answer, and they are also our cleanest source of identification. Unlike the firm data, they observe gold-suitable and other districts before and after the 2008 boom, so we can check for differential trends.

The prediction is specific. Where the return to mining rises, participation should rise and workers should move into the occupations mining draws on. We estimate (12) on the five DHS waves from 1993 to 2014. The outcomes are the cluster-level share of adults who worked in the past twelve months and the share in physical (agricultural or manual) occupations.

Table 5: Household labour allocation: men in gold-suitable districts work more and shift into physical occupations

Outcome	Men (1)	Women (2)	Pooled (3)
Working at all	+0.050*** (0.018)	+0.046** (0.022)	+0.048*** (0.015)
Manual or ag (physical)	+0.057** (0.025)	+0.029 (0.027)	+0.043* (0.023)
Region FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Cluster-years	1,172	1,173	2,345
SE clustering	District	District	District

Notes: Coefficient β_2 on the shift-share exposure $\text{Suit}_d \times \log P_{t-1}$ from $Y_{dt} = \beta_1 \text{Suit}_d + \beta_2 B_{dt} + \mu_r + \tau_t + \varepsilon_{dt}$. DHS 1993–2014 (drop 2022 due to v717 recoding). Drop-LSM. Outcomes are weighted cluster-level shares. Physical labour combines manual and agricultural occupations. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Men in gold-suitable districts work more as the gold price rises, and the new work is physical. A one-unit increase in exposure, which corresponds to a fully suitable district facing a gold price one log point higher, raises the male employment rate by 5.0

percentage points and the share in physical occupations by 5.7 points (Table 5, column 1). Both estimates are statistically significant at the 5 percent level or better. Against a male employment share of 78 percent, the first is an increase of about 6 percent. Women’s participation rises by a similar amount, but their occupational mix does not shift measurably (column 2).

The timing matches the boom (Figure 3). Through the 1990s the male employment gradient on suitability is flat relative to 2003, and the pre-boom coefficients are jointly indistinguishable from zero for every outcome we examine. By 2014 the gradient had risen by 6.2 percentage points. That is close to what the pooled estimate implies. The lagged gold price rose by about 1.5 log points between the 2003 and 2014 waves, so 5.0 points per log point predicts a gap of about 7.5 points.




Figure 3: DHS event study around the 2008 boom (men, drop-LSM). Reference wave = 2003. 95% confidence intervals from standard errors clustered by district. Joint pre-trend Wald not rejected for any outcome ($p \geq 0.52$).

The two halves of the evidence fit together. The establishment data show fewer agriculture-linked firms in gold-belt districts; the household data show more men working, and working in agricultural and manual occupations, in the same districts. Labour did not leave the market. The pattern is consistent with workers moving into activity that a fixed-establishment survey is least equipped to see.⁴

Two further results point the same way. A spatial discontinuity at the edge of the Girard et al. (2025) artisanal-mining polygon, which uses no price variation at all, gives a male employment effect of 5.5 percentage points, close to the shift-share estimate

⁴The DHS are repeated cross-sections, so they do not follow individuals, and their occupation codes have no mining category. The shift into physical work therefore describes the margin mining draws on rather than a count of miners. The surveys used here end in 2014 and cover the boom but not the crackdown.

(Table A.7). The same design at the geological rock contact finds nothing, consistent with workers responding to where mining actually happens rather than to geology as such. In addition, living-standards data from 2005/06 show no earnings gradient on suitability before the boom (Table A.8), so gold-belt districts were not already higher-wage labour markets, although that test has limited power.

If labour withdrawal links the boom to the formal sector, it should bite hardest where there is the most agricultural labour to draw on. The next section tests that implication.

8 Agricultural Conditions and the Formal-Sector Gap

Proposition 3 gives labour withdrawal a testable scope condition. The formal-sector gap should be largest where agriculture supplies the most labour and should fade where farming is thin. The test is useful because the obvious confounds in the firm cross-section (soils, terrain, settlement history, market access) have no reason to weaken in dry districts. We interact suitability with standardised district climate measures,

$$\log N_d = \beta_1 \text{Suit}_d + \beta_2 C_d + \beta_3 \text{Suit}_d \times C_d + \mu_r + \varepsilon_d, \quad (13)$$

where C_d is drought frequency, flood frequency, recent rainfall, or the long-run rainfall trend.




Figure 4: Marginal Suit_d effect on $\log N$ firms by climate exposure. Drop-LSM, region fixed effects, standard errors clustered by region.

The gap is concentrated in wetter districts. The interaction with drought frequency is positive and statistically significant (Table 6, column 1).⁵ In a district one standard deviation less drought-prone than average, a fully suitable district has about 32 log points fewer formal establishments; at average drought exposure the gap is 7 log points

⁵At the 5 percent level with conventional standard errors, and at the 10 percent level under the wild cluster bootstrap.

and imprecise; in drought-prone districts it disappears (Figure 4). The other climate measures point the same way, with smaller and less precise interactions.

Table 6: Galamsey hollowing concentrates in wet, agriculturally productive districts

Climate C_d	Drought freq. 08–23 (1)	Flood freq. 08–23 (2)	Rain Z 14–23 (3)	Rain trend 1990–2023 (4)
$Suit_d$	−0.067 (0.100)	−0.227** (0.106)	−0.205** (0.096)	−0.207** (0.090)
$Suit_d \times C_d$	+0.250** (0.108)	−0.074 (0.116)	−0.077 (0.100)	−0.066 (0.094)
Region FE	Yes	Yes	Yes	Yes
Districts	229	229	229	229
SE clustering	Region	Region	Region	Region

Notes: Outcome is log firm count (IBES 2024, district level), drop-LSM. Each column is a separate regression with C_d z-scored. The interaction β_3 on $Suit_d \times C_d$ measures how the marginal Suit effect varies with climate exposure; only drought is statistically significant. Standard errors clustered by region. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Taken together, Sections 6 to 8 give a consistent picture: fewer formal establishments in the gold belt, workers moving into physical work over the same period, and a formal-sector gap that is largest where agriculture supplies the labour. The last piece of evidence asks what happened to local activity when the state suppressed mining.

9 The 2017 Crackdown

In July 2017 the government launched Operation Vanguard, a military task force against illegal mining sites. If informal extraction had become part of local economic life, disrupting it should have reduced local activity in gold-belt districts. We estimate

$$\log NTL_{dt} = \sum_{k \neq -1} \beta_k Suit_d \times \mathbf{1}\{t - 2017 = k\} + \alpha_d + \tau_t + e_{dt}. \quad (14)$$

on the district-year night-lights panel and collapse the event-time coefficients into a single post-2017 interaction. Because suitability is continuous and fixed, there is no staggered timing; identification comes from differential exposure to a common national shock.

Night-lights fell after the crackdown in gold-suitable districts. In the 2014–2020 window, where the pre-period coefficients are jointly indistinguishable from zero, a fully suitable district lost about 9 percent of its night-lights relative to a non-suitable district (Table 7, column 6). Dropping industrial-mine districts leaves the estimate unchanged.

Longer windows give larger declines, about 14 percent, but their event studies show pre-trends from the boom itself, so we do not use them.

Table 7: Operation Vanguard ATT: $\text{Suit}_d \times \text{Post-2017}$ on log nightlights

	2000–2020		2010–2020		2014–2020	
	Full	Drop LSM	Full	Drop LSM	Full	Drop LSM
$\text{Suit}_d \times \text{Post-2017}$	−0.155*** (0.054)	−0.148*** (0.055)	−0.159*** (0.048)	−0.152*** (0.049)	−0.097* (0.050)	−0.094* (0.051)
Districts	260	252	260	252	260	252
Years	21	21	11	11	7	7
Observations	5,460	5,292	2,860	2,772	1,820	1,764
R^2 within	0.016	0.014	0.024	0.022	0.010	0.010
District FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: TWFE event study collapsed to a single post-2017 dummy. Standard errors clustered by district in parentheses below coefficients. The preferred window is 2014–2020 (drop-LSM, column 6), where the joint pre-trend Wald test does not reject parallel trends ($p = 0.55$). Longer windows are reported as a robustness check; their pre-trends fail because of the 2008–2014 gold-belt boom that contaminates pre-2014 observations. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The decline is sharply timed but imprecisely estimated, and we treat it as suggestive.⁶

The direction is what Proposition 4 predicts: suppressing mining removes a source of local income. The design does not, however, isolate artisanal mining from every other change in exposed districts after 2017. The other half of the proposition, that labour returns to formal employment after a crackdown, cannot be tested with a single firm cross-section.

10 Conclusion

Ghana’s artisanal gold boom coincided with a large shift in local work. Between 2003 and 2014, as the gold price more than tripled, the male employment rate in fully gold-suitable districts rose by about 6 percentage points relative to other districts, mostly in agricultural and manual occupations. In 2024 the same districts had about 19 percent fewer formal establishments than comparable districts in their region, a gap concentrated in agriculture-linked sectors and in wetter districts.

The two margins tell complementary parts of the story. Establishment counts show

⁶The estimate is statistically significant at the 10 percent level. The pre-period is three years long, which limits the power of the pre-trend test, and dropping 2020, when COVID-19 depressed both night-lights and informal mining, reduces the decline to about 7 percent, which is not statistically distinguishable from zero (Table A.9). The eight districts with industrial gold mines are too few to separate their response from that of artisanal-prone districts (Section OA.11).

a thinner formal sector; household surveys show more men working, in occupations that establishment counts do not track. The paper does not observe individual workers moving into mining, and its contribution does not depend on doing so: it shows that the formal-establishment margin alone would give an incomplete account of how a local economy responds to an informal boom. The household evidence is the most firmly identified, the establishment evidence is descriptive, and the enforcement evidence is suggestive.

Two implications follow, each with limits. Ghana has relied on military crackdowns against galamsey since 2013, and the public case for them is environmental. Our estimates suggest that the costs of a crackdown also fall on the workers it displaces, which strengthens the case for pairing enforcement with a legal route into licensed small-scale mining or with support for farming in affected districts. We cannot evaluate those alternatives with our data, and the study does not estimate the welfare effects of enforcement. The second implication concerns measurement. In mining districts, establishment surveys and tax records miss much of the activity the boom created, so a measured decline in the formal sector overstates the loss of local economic activity.

The setting matters for external validity. Ghana's gold belt overlaps much of its forest-zone farmland, which is where our results place the largest formal-sector gap; booms in drier settings may displace fewer firms. Unfortunately, our data do not let us observe firm entry and exit or labour allocation after 2014. Linked establishment panels, such as tax records, would show whether the formal-sector gap reflects the exit of the least productive firms or entry that never occurred, and whether the crackdown sent workers back.

References

- Hunt Allcott and Daniel Keniston. Dutch disease or agglomeration? the local economic effects of natural resource booms in modern america. *The Review of Economic Studies*, 85(2):695–731, 2018.
- Fernando M. Aragón and Juan Pablo Rud. Natural resources and local communities: Evidence from a Peruvian gold mine. *American Economic Journal: Economic Policy*, 5(2):1–25, 2013.
- Fernando M. Aragón and Juan Pablo Rud. Polluting industries and agricultural productivity: Evidence from mining in Ghana. *Economic Journal*, 126(597):1980–2011, 2016. doi: 10.1111/eoj.12244.
- Abigail Barenblitt, Amanda Payton, David Lagomasino, Lola Fatoyinbo, Kofi Asare, Kenneth Aidoo, Hugo Pigott, Charles Kofi Som, Laurents Smeets, Omar Seidu, and Danielle Wood. The large footprint of small-scale artisanal gold mining in ghana. *Science of the Total Environment*, 781:146644, 2021. doi: 10.1016/j.scitotenv.2021.146644.
- Rémi Bazillier and Victoire Girard. The gold digger and the machine: Evidence on the distributive effect of the artisanal and industrial gold rushes in Burkina Faso. *Journal of Development Economics*, 143:102411, 2020.
- Kirill Borusyak, Peter Hull, and Xavier Jaravel. A practical guide to shift-share instruments. *Journal of Economic Perspectives*, 39(1):181–204, 2025. doi: 10.1257/jep.20231370.
- Ricardo J. Caballero and Mohamad L. Hammour. The cleansing effect of recessions. *American Economic Review*, 84(5):1350–1368, 1994.
- W. Max Corden and J. Peter Neary. Booming sector and de-industrialisation in a small open economy. *The Economic Journal*, 92(368):825–848, 1982.
- Gordon Crawford and Gabriel Botchwey. Conflict, collusion and corruption in small-scale gold mining in Ghana: Chinese miners and the state. *Commonwealth and Comparative Politics*, 55(4):444–470, 2017.
- Melissa Dell. The persistent effects of peru’s mining mita. *Econometrica*, 78(6):1863–1903, 2010.
- Joseph Eshun. Farm losses and welfare gains: The uneven effects of gold mining on ghana’s cocoa farmers. Working paper, University of Hawai’i at Mānoa, 2026.
- Victoire Girard, Teresa Molina-Millán, and Guillaume Vic. Artisanal mining in Africa. Green for gold? *Economic Journal*, 135:2578–2597, 2025. doi: 10.1093/ej/ueaf058.
- M. C. Hansen, P. V. Potapov, R. Moore, M. Hancher, S. A. Turubanova, A. Tyukavina, D. Thau, S. V. Stehman, S. J. Goetz, T. R. Loveland, A. Kommareddy, A. Egorov,

- L. Chini, C. O. Justice, and J. R. G. Townshend. High-resolution global maps of 21st-century forest cover change. *Science*, 342(6160):850–853, 2013.
- Gavin Hilson. Shootings and burning excavators: Some rapid reflections on the Government of Ghana’s handling of the informal galamsey mining ‘menace’. *Resources Policy*, 54:109–116, 2017.
- Andreas Kotsadam and Anja Tolonen. African mining, gender, and local employment. *World Development*, 83:325–339, 2016. doi: 10.1016/j.worlddev.2016.01.007.
- Xuecao Li, Yuyu Zhou, Min Zhao, and Xia Zhao. A harmonized global nighttime light dataset 1992–2018. *Scientific Data*, 7(168), 2020.
- Marc J. Melitz. The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica*, 71(6):1695–1725, 2003. doi: 10.1111/1468-0262.00467.
- Frits Steenhuisen and Simon J. Wilson. Development and application of an updated geospatial distribution model for gridding 2015 global mercury emissions. *Atmospheric Environment*, 211:138–150, 2019. doi: 10.1016/j.atmosenv.2019.05.003.
- Liang Tang and Tim T. Werner. Global mining footprint mapped from high-resolution satellite imagery. *Communications Earth & Environment*, 4:134, 2023. doi: 10.1038/s43247-023-00805-6.
- Andreas Forø Tollefsen, Håvard Strand, and Halvard Buhaug. PRIO-GRID: A unified spatial data structure. *Journal of Peace Research*, 49(2):363–374, 2012.

A Additional tables and figures

This appendix collects the tables and figures cited in the main text. The Online Appendix, provided as supplementary material, describes each exercise in detail and reports further results.

Tables

Table A.1: Intensive–extensive decomposition: the firm-count decline is mostly thinning within active sectors, not sector closure

Outcome		Full sample (1)	Drop LSM (2)
Total	$\log N_d$	−0.174* (0.100)	−0.216** (0.098)
Extensive	$\log N_d^{\text{active sectors}}$	−0.036 (0.037)	−0.040 (0.035)
Intensive	$\log \bar{n}_d^{\text{firms / active}}$	−0.138* (0.074)	−0.176** (0.072)
Districts		247	239
Region FE		Yes	Yes
SE clustering		Region	Region

Notes: IBES 2024 firm cross-section. Each cell reports the coefficient on Suit_d from a separate district-level regression of the indicated outcome on Suit_d with region fixed effects. The decomposition exploits the identity $\log N_d \equiv \log N_d^{\text{active}} + \log \bar{n}_d^{\text{firms / active}}$, so the extensive and intensive coefficients sum to the total up to numerical precision (drop-LSM: $-0.040 + -0.176 = -0.216$). The extensive margin counts the number of distinct ISIC2 sectors with at least one formal firm in the district; the intensive margin counts firms per active firm sector. Standard errors clustered by region. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.2: Compositional shift: agriculture and industry shares fall, services rise with $Suit_d$

Super-sector share (pp)	Full sample (1)	Drop LSM (2)
Agriculture	−1.28*** (0.45)	−1.17*** (0.43)
Industry	−3.09** (1.22)	−2.77** (1.23)
Services	+4.37*** (1.23)	+3.94*** (1.17)
Districts	247	239
Region FE	Yes	Yes

Notes: IBES 2024 firm cross-section. Each cell reports the coefficient on $Suit_d$ from a separate district-level regression of the share of formal firms in the indicated super-sector on $Suit_d$, with region fixed effects. Coefficients are scaled to percentage points of share per unit-share change in $Suit_d$ (i.e., $\beta \times 100$): a Mobile-Belt district with $Suit_d = 1$ has, relative to a non-suitable district with $Suit_d = 0$, a 1.2 pp lower agriculture share, a 2.8 pp lower industry share, and a 3.9 pp higher service share (drop-LSM). Standard errors clustered by region. The three super-sector shares sum to one by construction; the three coefficients sum to zero up to rounding. Industry includes mining-adjacent and non-mining manufacturing; service includes retail, transport, finance, and consumer services. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.3: Wild cluster bootstrap on the district-level regressions

Outcome	$\hat{\beta}$	SE	Analytic p	Bootstrap p	Bootstrap 95% CI
log N_d establishments	−0.216	(0.098)	0.043	0.061	[−0.434, +0.014]
log total persons engaged	−0.201	(0.097)	0.056	0.111	[−0.363, +0.102]
log mean establishment size	+0.015	(0.046)	0.751	0.756	[−0.080, +0.146]

Notes: Coefficient on $Suit_d$ from (10) with region fixed effects, drop-LSM, 239 districts. Analytic p -values use cluster-robust standard errors on 16 regions. Bootstrap p -values and confidence sets use the wild cluster bootstrap with 9,999 replications, Rademacher weights, imposing the null.

Table A.4: Firm cross-section with scale and market-access controls

	log N_d firms				Firms/km ²	PPML: N_d
	(1)	(2)	(3)	(4)	(5)	(6)
Suit _d	-0.216** (0.098)	-0.223** (0.093)	-0.176** (0.066)	-0.144** (0.061)	0.012 (0.187)	-0.190** (0.076)
log district area		-0.084 (0.067)	0.031 (0.091)	0.036 (0.056)		0.044 (0.048)
Urban land share			0.056* (0.028)	0.053** (0.023)	0.180* (0.089)	0.041** (0.018)
Travel time to city			-0.002 (0.001)	-0.000 (0.001)	-0.006*** (0.002)	-0.001 (0.001)
Pre-boom log NTL (2001–05)				0.174*** (0.017)	0.180*** (0.047)	0.199*** (0.024)
Districts	239	239	239	239	239	239
Within R^2	0.016	0.036	0.128	0.396	0.431	

Notes: IBES 2024 district cross-section, drop-LSM. Columns 1–4 estimate (10) by OLS with region fixed effects, adding controls cumulatively. Column 5 replaces the outcome with log firms per km². Column 6 is a Poisson (PPML) regression on the firm count with region dummies. Urban land share and travel time to the nearest city are area-weighted PRIO-GRID covariates; pre-boom log NTL is the 2001–2005 district mean and proxies baseline scale in the absence of a district population series. Standard errors clustered by region. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.5: Gold suitability does not shift the formality composition of surviving establishments

Composition outcome	Sample mean	$\hat{\beta}$ on Suit _d	SE	Bootstrap p
Share registered with the tax authority	0.088	+0.0034	(0.0084)	0.690
Share micro (fewer than 5 persons)	0.913	−0.0034	(0.0041)	0.414
Share with any paid employee	0.638	+0.0044	(0.0291)	0.888

Notes: IBES 2024 district cross-section, drop-LSM, 239 districts. Each row is a separate regression of the indicated share on Suit_d with region fixed effects. Standard errors clustered by region; bootstrap p -values from 9,999 wild cluster replications imposing the null.

Table A.6: Cleansing or hollowing? Decomposing the establishment-count effect

	Coefficient	Std. error	Bootstrap p	N
<i>Panel A. Where the establishment-count effect goes (districts)</i>				
log establishments	-0.216*	(0.098)	0.057	239
log total persons engaged	-0.201	(0.097)	0.111	239
log mean establishment size	0.015	(0.046)	0.755	239
<i>Panel B. Where the size difference sits (establishments)</i>				
Pr(persons engaged > 1)	0.032***	(0.010)	—	161,183
Pr(persons engaged > 2)	0.017**	(0.007)	—	161,183
Pr(persons engaged > 3)	0.009*	(0.005)	—	161,183
Pr(persons engaged > 5)	0.005	(0.003)	—	161,183
Pr(persons engaged > 10)	0.003	(0.002)	—	161,183
Pr(persons engaged > 20)	0.001	(0.001)	—	161,183
Pr(persons engaged > 50)	0.001	(0.001)	—	161,183
<i>Panel C. Productivity of surviving establishments</i>				
log revenue per person engaged	-0.016	(0.097)	—	159,146
revenue band index (1–12)	0.022	(0.044)	—	159,146
log revenue (band midpoint)	0.023	(0.103)	—	159,146
<i>Panel D. Vintage of the surviving stock (districts)</i>				
share founded within five years	-0.001	(0.017)	0.982	239
share founded before 2008	-0.003	(0.012)	0.804	239
median establishment age (years)	-0.116	(0.262)	0.661	239

Notes: IBES 2024, drop-LSM sample. Every coefficient is on Suit_d . Panels A and D are district regressions (239 districts) with region fixed effects and standard errors clustered by region; because there are 16 clusters, the reported p -value is a wild cluster bootstrap (9,999 Rademacher draws, null imposed). Panels B and C are establishment regressions with two-digit sector and region fixed effects and standard errors clustered at the district, for which stars are reported in the usual way. Revenue in Panel C is the geometric midpoint of the IBES band, with the open top band closed at 1.5 times its floor.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.7: DHS-cluster RDD: household labour effects track the ASGM-activity boundary, not the rock contact

Boundary	Working at all		Manual or ag (physical)	
	Men (1)	Women (2)	Men (3)	Women (4)
ASGM polygon	+0.055** (0.023)	+0.018 (0.021)	+0.049 (0.039)	+0.034 (0.038)
BGS Mobile Belt	−0.011 (0.057)	+0.031 (0.039)	+0.063 (0.084)	+0.035 (0.063)
Cluster-years (ASGM)	826	955	826	955
Cluster-years (BGS)	178	198	178	198
Wave + Region FE	Yes	Yes	Yes	Yes
Local linear	Yes	Yes	Yes	Yes
SE clustering	District	District	District	District

Notes: DHS Ghana 1993–2014 (drop 2022 wave; v717 recoding). Each cell is a separate local-linear regression: the Girard row fits the model on cluster-years within ± 10 km of the [Girard et al. \(2025\)](#) ASGM-suitability polygon edge, with treatment defined as inside that polygon and signed distance to its edge as the running variable; the BGS row fits the same model on a different sample, defined by ± 10 km of the BGS Birimian Mobile/Orogenic-Belt rock contact. The two boundaries are not mutually exclusive but they enclose different populations, hence the different sample sizes. Each regression includes treat $\times$ distance interaction, wave fixed effects, and region fixed effects; standard errors are clustered at districts level. The household labour-supply discontinuity appears at the ASGM-activity boundary but disappears at the rock contact, indicating that the household effect operates through observed ASGM activity (alluvial deposits, worked greenstone fields) rather than geological potential. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.8: No pre-boom wage gradient on gold suitability (GLSS5, 2005/06)

Outcome	$\hat{\beta}$ on Suit _d	SE	Wild-bootstrap p	Districts
log real employment income per capita	+0.048	(0.225)	0.762	98
log mean real employment income	−0.485	(0.612)	0.494	102
share of households with wage income	−0.038	(0.042)	0.389	102

Notes: District-level regressions of the indicated outcome on Suit_d with pre-2019 region fixed effects, on the GLSS5 2005/06 sample. Employment income is deflated by the survey’s Accra price index. Per-capita income is computed over households reporting positive employment income, which is why that row has four fewer districts. Standard errors clustered at region; wild cluster bootstrap p -values from 9,999 replications with Rademacher weights, imposing the null. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.9: Enforcement window and the differential-trend test on the Bartik reduced form

	Vanguard DiD		Bartik reduced form	
	(1)	(2)	(3)	(4)
$Suit_d \times \text{Post-2017}$	-0.094* (0.051)	-0.077 (0.048)		
$B_{dt} = Suit_d \times \log P_t$			-0.553*** (0.154)	1.186*** (0.214)
$Suit_d \times (\text{year} - 2000)$				-0.206*** (0.034)
District-years	1764	1512	5200	5200
Within R^2	0.010	0.008	0.017	0.058

Notes: District-year night-lights panel with district and year fixed effects throughout. Columns 1–2 collapse (14) to a single post-2017 interaction on the drop-LSM sample; column 1 reproduces Table 7 column 6 and column 2 drops 2020. Columns 3–4 estimate the Bartik reduced form on log night-lights over 2001–2020; column 4 adds a $Suit_d$ -specific linear trend. Columns 1–2 use $\log(1 + NTL_{dt}^{\text{mean}})$ and columns 3–4 $\log(1 + NTL_{dt}^{\text{sum}})$, matching the specifications in Table 7 and Table OA.8. Standard errors clustered by district.
 * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figures




Figure A.1: Sector-specific coefficients $\hat{\beta}_g$ on Suit_d in $\log N_d^{(g)} = \beta_g \text{Suit}_d + \mu_r + \varepsilon_d$, drop-LSM, with 95% confidence intervals (districts-clustered). Each row is a separate regression on the log count of firms in ISIC Rev. 5 two-digit sector g . Firm counts shown on the right. Crops, food manufacturing, and fabricated metals have the largest negative coefficients; wholesale trade flips positive, food service is flat, consistent with, though not proof of, demand from informal mining that IBES does not observe.




Figure A.2: Mean annual Hansen forest loss area per district, galamsey-prone (top-quartile $Suit_d$) versus rest of Ghana, drop-LSM. Vertical lines mark the 2008 Chinese-galamsey wave and the 2017 Operation Vanguard.

Online Appendix

Supplementary material for *Artisanal Mining and the Hollowing of the Formal Sector: Evidence from Ghana's Galamsey Zones*. Tables and figures numbered A.1, A.2, ... appear in Appendix A of the paper; those numbered OA.1, OA.2, ... appear here.

Contents

OA.1	Additional establishment results	36
OA.2	Inference with few clusters, and reclassification	36
OA.3	Selection versus labour withdrawal	37
OA.4	Validating the exposure proxy	39
OA.5	Dollar and cedi gold prices	39
OA.6	Geology and observed mining	40
OA.7	Forest loss and establishment size	42
OA.8	Forest loss, mining footprint and mercury	43
OA.9	Pre-boom earnings	44
OA.10	Spatial regression discontinuity designs	45
OA.11	Night-lights: additional results	53
OA.12	A no-galamsey counterfactual	57
OA.13	Sample construction	58

OA.1 Additional establishment results

Decomposition and composition. We decompose district-level establishment counts into the number of active two-digit sectors and the average number of establishments per active sector,

$$\log N_d \equiv \log N_d^{\text{active}} + \log \bar{n}_d^{\text{firms/active}},$$

and regress each component on Suit_d with region fixed effects. Table A.1 shows that the extensive component is small and insignificant and that roughly four-fifths of the gradient reflects fewer establishments within active sectors. Table A.2 shows that the formal sector in gold-suitable districts has smaller agriculture and industry shares and a larger services share.

Two-digit sectors. Figure A.1 re-estimates (10) separately for each two-digit sector. Crop production and food manufacturing have the largest negative coefficients; wholesale trade is slightly positive and food service is flat.

Scale, market access, and functional form. Table A.4 adds controls to the district cross-section (10). The gradient is stable in sign and significance as district area, urban land share, travel time to the nearest city, and pre-boom (2001–2005) night-lights enter: the coefficient moves from -0.216 to -0.144 , and a Poisson (PPML) specification on the firm count itself gives -0.190 . Census population is not available at this district vintage, so pre-boom night-lights stand in as a baseline-scale control.

The exception is column 5, which normalises by land area. Gold-suitable districts are considerably smaller than the rest (656 km^2 on average above $\text{Suit}_d = 0.5$ against $1,271 \text{ km}^2$ below it), and establishments per km^2 show no gradient at all (column 5). The unconditional correlation between Suit_d and $\log N_d$ is also positive ($+0.10$); the negative coefficient is a within-region comparison. Both facts limit what the firm-count result shows: gold-belt districts host fewer formal firms than other districts *in the same region and of the same size*, not fewer firms per unit of land.

OA.2 Inference with few clusters, and reclassification

Every district-level regression in the paper clusters on the sixteen regions. Cluster-robust t -statistics are known to over-reject at that number of clusters, so Table A.3 reruns the district-level outcomes with the wild cluster bootstrap.

The firm-count coefficient keeps its size but not its status. It is statistically signifi-

cant at the 5 percent level with cluster-robust standard errors, but only at the 10 percent level under the bootstrap, and the 95 percent bootstrap confidence set contains zero. For total persons engaged, the bootstrap moves the estimate from the 10 percent level to insignificance. We therefore describe the cross-sectional gradient as negative throughout and marginally significant, and we do not attach conventional stars to it in the text.

The same table shows something else. At the district level mean establishment size shows no gradient, and the coefficient on total persons engaged (-0.201) is close to the coefficient on establishment counts (-0.216). Gold-belt districts thus have proportionally fewer establishments and proportionally less employment, with average size unchanged. The larger-survivor result in Table 3 is a within-sector effect estimated on firm microdata, and it does not aggregate to a district-level size gradient.

Is this thinning or reclassification? An alternative explanation for the firm-count result is that establishments do not disappear but move out of the surveyed universe as activity informalises. Two pieces of evidence weigh against it. First, IBES 2024 does not sample on formality: only 8.8 percent of the establishments in it are registered with the tax authority, and 91 percent are micro units of fewer than five persons, so informal activity is inside the frame rather than outside it. Second, if suitability pushed establishments across a formality margin, the surviving composition should shift. Table A.5 shows it does not. The registered share, the micro share, and the share with any paid employee are all flat in Suit_d , with coefficients under half a percentage point and none statistically significant under the bootstrap. Gold-belt districts have fewer establishments that look like the establishments elsewhere, which is what thinning predicts and reclassification does not.

OA.3 Selection versus labour withdrawal

Gold-suitable districts hold fewer establishments, and those they hold are slightly larger. A competing explanation for these two facts is selection. If exposure closes the least productive firms, survivors are better firms and average size rises through composition, so what we call hollowing would be the cleansing of Caballero and Hammour (1994), with the welfare sign reversed. Truncation raises measured average size mechanically in any model with a productivity cutoff, Melitz (2003) included, so the two signs alone cannot separate the accounts.

They separate on where the missing employment goes. Total formal employment in

a district is the product of the number of establishments and their mean size,

$$\log E_d = \log N_d + \log \bar{n}_d. \quad (\text{OA.1})$$

Cleansing redeploys workers from closing firms to surviving ones, so the fall in $\log N_d$ is offset in $\log \bar{n}_d$ and $\log E_d$ moves little. Withdrawal of labour to mining removes workers from the sector, so $\log E_d$ falls with $\log N_d$ and $\log \bar{n}_d$ is flat.

Panel A of Table A.6 estimates the three terms on the same 239 districts. The establishment count falls by 21.6 log points and total persons engaged by 20.1. Mean size absorbs the remaining 1.5 log points, 7 percent of the count effect, and we cannot reject zero under the wild cluster bootstrap. Workers displaced from the missing establishments do not reappear in the ones that remain.

Panel B locates the size difference. An establishment in a fully gold-suitable district is 3.2 percentage points more likely to employ more than one person. The gradient halves at two workers, is significant only at the 10 percent level at three, and is zero from five workers upward. No establishment of any scale is larger; the lower tail is thinner. Truncation under either account produces this shape, which is why Panel A carries the argument.

Panel C asks the question cleansing answers directly. IBES records 2023 revenue in thirteen bands, which we convert to geometric midpoints. Revenue per person engaged among surviving establishments does not move with suitability, and the revenue band index is flat as well. We state this as an absence rather than a refutation: on a banded revenue measure, the estimate cannot detect productivity gains smaller than about 20 percent.

Panel D uses the year each establishment began operating. Selection that culls marginal firms should tilt the surviving stock toward older cohorts that have already passed the test. It does not. The share founded within five years, the share founded before the 2008 boom, and median establishment age are all flat in suitability, and we can rule out movements larger than 4 percentage points in the vintage shares. The deficit is spread across cohorts rather than concentrated among recent entrants.

The evidence weighs against cleansing, subject to one limit the data cannot remove. A cross-section records the stock of survivors, not the flows of entry and exit that produced it, so Panel D infers vintage composition rather than measuring turnover. A linked establishment panel would measure exit directly.

OA.4 Validating the exposure proxy

$ASGM_{dt}$ measures the potential value of artisanal mining, not observed extraction, so we check whether gold-suitable districts are in fact where gold is found. [Girard et al. \(2025\)](#) address this continent-wide by comparing their proxy with independent records of ASGM occurrence. Table OA.1 reports the Ghana counterpart, using the GOLDDATA gold-occurrence indicators distributed with PRIO-GRID ([Tollefsen et al., 2012](#)), one of the record sources they use, and one built from field geology rather than from the bedrock layer that defines $Suit_d$.

Districts above the median of $Suit_d$ have three times the recorded gold-occurrence share of districts below it, 0.268 against 0.085, and the two measures correlate at 0.275 across the 260 districts. The correlation is far from one, which is expected: PRIO-GRID cells are half-degree, so Ghana contains only 119 of them, and the records are presence indicators rather than intensities. What the comparison establishes is that the gradient runs the right way against an independent source, not that the proxy measures extraction.

Table OA.1: The exposure proxy against independent gold-occurrence records

	Value
Correlation of $Suit_d$ with the gold-record share	+0.275
Mean record share, $Suit_d \geq 0.5$	0.268
Mean record share, $Suit_d < 0.5$	0.085
Districts with any recorded gold occurrence	87
Districts	260

Notes: Gold-occurrence records are the GOLDDATA placer, vein, and surface indicators in PRIO-GRID ([Tollefsen et al., 2012](#)), aggregated to districts as the area-weighted share of overlapping cells with any recorded occurrence. $Suit_d$ is the district share of gold-suitable bedrock from [Girard et al. \(2025\)](#). The two are constructed from different sources.

OA.5 Dollar and cedi gold prices

The exposure measure in (9) uses the dollar gold price, following [Girard et al. \(2025\)](#), because that price is exogenous to Ghana. Artisanal miners are paid in cedi. The return that governs their labour allocation is the cedi price, $P_t^{GHS} = P_t^{USD} \times E_t$, and over our sample the two diverge: the dollar price rose 1.61 log points between 2001 and 2020 while the cedi price rose 3.66, as the currency fell from 0.72 to 5.60 per dollar.

The divergence is largest when it matters most. In 2014 and 2015 the dollar gold price fell while the cedi price rose, because the cedi lost a third of its value. [Barenblitt](#)

et al. (2021) date roughly half of Ghana’s recorded conversion of vegetation to artisanal mining to those years. The dollar shifter records a fall in mining incentives exactly when mining expanded fastest.

Table OA.2 takes both shifters to Hansen forest loss, the closest annual proxy for extraction we have. Each predicts loss on its own. Entered together, the cedi shifter takes a large positive coefficient while the dollar shifter turns negative, so the exchange-rate component carries the signal. Restricting to 2013–2020, after the Landsat 8 change that raised detection rates in the Hansen product, gives the same pattern.

We nonetheless keep the dollar price in the main specifications, because the exchange rate is not exogenous to Ghana. The cedi depreciates in episodes of fiscal and monetary stress that independently affect firm survival and local activity, and the gold belt overlaps the cocoa belt, whose export earnings also rise with depreciation. Substituting the cedi price would buy a better measure of the miner’s incentive at the cost of the exogeneity that identifies the design. What we take from the exercise is that the dollar price understates movements in the return to artisanal mining during depreciation episodes.

Table OA.2: Dollar and cedi gold-price shifters against forest loss

	(1) 2001–20	(2) 2001–20	(3) 2001–20	(4) 2013–20	(5) 2013–20
$ASGM_{dt}^{USD}$	0.831*** (0.115)		-0.785*** (0.176)	-0.509* (0.305)	
$ASGM_{dt}^{GHS}$		0.473*** (0.065)	0.837*** (0.131)		0.374*** (0.117)
District-years	4720	4720	4720	1888	1888

Notes: Outcome is $\log(1 + \text{annual Hansen forest loss, ha})$ on the drop-LSM district panel with district and year fixed effects. $ASGM_{dt}^{USD} = \text{Suit}_d \times \log p_{t-1}^{USD}$ and $ASGM_{dt}^{GHS} = \text{Suit}_d \times \log(p_{t-1}^{USD} E_{t-1})$, with E_t the official cedi-dollar rate (World Bank PA.NUS.FCRF). Columns 4–5 restrict to 2013 onward, within which Hansen detection is internally consistent. Standard errors clustered at district. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

OA.6 Geology and observed mining

Suit_d records where artisanal mining is geologically possible. It does not record where mining happens, and the distinction is the central interpretive question for the cross-section. Tang and Werner (2023) map the global mining footprint from high-resolution satellite imagery, which makes the two separable. Their polygons cover 1,004 km² of Ghana across 1,897 features in 39 districts.

The two measures are related as the model requires. Regressing log mapped mine area on $Suit_d$ with region fixed effects gives 0.741, statistically significant at the 10 percent level under the wild cluster bootstrap, and districts above the median of $Suit_d$ contain 4.18 km² of mapped mining on average against 0.08 km² below it. Gold-suitable districts are where mining is observed.

Table OA.3 enters them together against establishment counts. Separately, each predicts fewer establishments. Together, $Suit_d$ retains most of its magnitude while the coefficient on mapped mine area falls by nearly half and loses significance. Restricting the mining measure to polygons below 5 km², which excludes industrial-scale features, does not change this.

Two explanations fit, and the data do not separate them. Mapped mine area is measured with considerable error: only 39 of 260 districts contain any mapped polygon, and satellite delineation of small alluvial workings is incomplete, so attenuation biases its coefficient toward zero in a specification where $Suit_d$ is measured without error. Alternatively, the establishment gradient genuinely tracks the geological corridor rather than extraction, which is what Table OA.6 also indicates when the firm-side discontinuity appears at the rock contact and not at the boundary of realised mining. We cannot distinguish these with the data in hand, and we therefore do not claim that the cross-sectional gradient identifies an effect of mining as opposed to an effect of the gold-bearing corridor. The household panel and the enforcement break, where mining activity varies over time within a district, carry that claim instead.

Table OA.3: Establishment counts on geology and on observed mining extent

	(1) Geology	(2) Mining	(3) Both	(4) Both, small
$Suit_d$	-0.216** (0.098)		-0.193* (0.108)	-0.196* (0.108)
log mapped mine area		-0.055*** (0.011)	-0.031 (0.019)	
log mapped area, small polygons				-0.030 (0.022)
Districts	239	239	239	239
Within R2	0.016	0.006	0.017	0.017

Notes: IBES 2024 district cross-section, drop-LSM, 239 districts, region fixed effects, standard errors clustered by region. Mapped mine area aggregates the polygons of Tang and Werner (2023) to districts; column 4 restricts to polygons below 5 km². Both mining measures enter in logs of one plus area in km². * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

OA.7 Forest loss and establishment size

Figure A.2 plots annual forest loss in galamsey-prone districts (top quartile of $Suit_d$) against the rest of Ghana. Galamsey-prone districts start from a lower level and stay below the rest of the country in twenty of twenty-three years; what the boom changed is the growth rate. Loss in galamsey-prone districts peaks in 2017, the year Operation Vanguard began.

One possible mechanism is that artisanal mining reduces formal-sector activity by degrading local environmental conditions. To isolate this channel from the broader labour-reallocation effects captured by $Suit_d$, we augment the baseline specification with a district-level measure of cumulative forest loss:

$$\log Y_i = \beta_1 Suit_d + \beta_2 \tilde{L}_d + \mu_r + \gamma_{ISIC2} + \varepsilon_i, \quad (OA.2)$$

where $\tilde{L}_d$ denotes standardised log cumulative forest loss constructed from Hansen et al. (2013). Conditional on geological suitability, the identifying variation in forest loss reflects residual differences in deforestation intensity arising from factors such as agricultural expansion, settlement growth, or fires.

Table OA.4: Disamenity: forest loss conditional on $Suit_d$ hits firm size, not firm count

Cumulative loss window	$\log N_d$ (district)	$\log emp_d$ (district)	$\log persons_i$ (firm)
2001–2023	−0.060 (0.150)	−0.040 (0.090)	−0.022*** (0.008)
2008–2013	−0.030 (0.140)	−0.010 (0.100)	−0.021*** (0.008)
2014–2023	−0.080 (0.150)	−0.040 (0.110)	−0.024*** (0.008)
2017–2023	−0.080 (0.150)	−0.040 (0.110)	−0.025*** (0.008)
$Suit_d$ controlled?	Yes	Yes	Yes
Region FE	Yes	Yes	Yes
ISIC2 FE	–	–	Yes
Observations	223	223	147k
SE clustering	Region	Region	District

Notes: Coefficient on z-scored log cumulative Hansen forest loss $\tilde{L}_d$ from eq. OA.2, drop-LSM. Each cell is a separate regression. $Suit_d$ is included as a control throughout. Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table OA.4 reports estimates across four cumulative forest-loss windows. Forest loss has little predictive power for firm exit once gold suitability is held constant: the estimated effects on firm counts are economically small and statistically indistinguish-

able from zero across specifications. By contrast, forest loss is consistently associated with smaller surviving firms. A one-standard-deviation increase in cumulative forest loss reduces employment per surviving establishment by roughly 2.5 percent.

We report this pattern descriptively. Forest loss conditional on suitability reflects agricultural expansion and settlement as well as mining, so the coefficient is not an estimate of mining-induced degradation. Section [OA.8](#) tests what it picks up.

OA.8 Forest loss, mining footprint and mercury

Forest loss records land removed from vegetation, which is not the same as land degraded in place. Three transmissions are possible, and they differ in incidence. If the mechanism is land taken out of cultivation, the effect should travel through input supply and concentrate in agriculture-linked activity. If it is the physical footprint of extraction, it should track mapped mining rather than vegetation loss from any cause. If it is pollution, it should track emissions and reach establishments of every kind. Table [OA.5](#) separates them by adding the mapped mining footprint of [Tang and Werner \(2023\)](#) and the artisanal-mining sector of the gridded global mercury inventory of [Steenhuisen and Wilson \(2019\)](#). The outcome is log persons engaged per establishment.

The data support neither of the first two. Mapped mining enters as an indicator, because few districts contain any mapped polygon; its coefficient is close to zero and leaves the forest-loss coefficient unchanged. Splitting the sample by sector gives the same forest-loss point estimate inside and outside agriculture, about -0.03 ; only precision differs, which is not what input supply predicts.

Pollution fares better. Establishments are smaller in districts with higher modelled artisanal mercury emissions, conditional on both land measures, and the estimate is statistically significant at the 1 percent level. A one standard-deviation increase in emissions goes with establishments about 2 percent smaller. Two qualifications bound this. The mercury surface is modelled rather than measured: national emission totals are distributed spatially, and the surface correlates 0.60 with mapped mine area, so it separates from land cover but not cleanly from mining itself. And forest loss keeps its own coefficient once mercury is included, so some part of it remains a district characteristic we have not identified. Yield on land still farmed, as in [Eshun \(2026\)](#), would separate degradation from land loss directly; we hold no district-level crop-yield series.

Table OA.5: Land taken, land converted, or pollution?

	(1) All	(2) All	(3) Ag-linked	(4) Other	(5) All	(6) Non-ag
Suit _d	0.048*** (0.013)	0.047*** (0.013)	0.064 (0.044)	0.044*** (0.013)	0.050*** (0.013)	0.046*** (0.013)
Forest loss, z	-0.030*** (0.011)	-0.030*** (0.011)	-0.037* (0.021)	-0.031*** (0.011)	-0.027** (0.011)	-0.029** (0.011)
Any mapped mining		0.005 (0.014)	0.082* (0.048)	-0.001 (0.013)	0.028* (0.016)	0.014 (0.014)
ASGM mercury, z					-0.017*** (0.006)	-0.011* (0.006)
Establishments	147199	147199	13598	133601	147199	133601

Notes: Establishment-level regressions of log persons engaged on Suit_d and standardised district covariates, IBES 2024, drop-LSM, with region and two-digit sector fixed effects and standard errors clustered at district. Forest loss is standardised log cumulative Hansen loss 2001–2023; mapped mine area is the standardised log [Tang and Werner \(2023\)](#) footprint; ASGM mercury is the standardised log of the artisanal-mining sector of the [Steenhuisen and Wilson \(2019\)](#) gridded 2015 emission inventory, area-weighted to districts. Columns 3, 4 and 6 split on whether the establishment’s two-digit sector is agriculture-linked (ISIC 1, 2, 3, 10, 11, 12). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

OA.9 Pre-boom earnings

Proposition 2 turns on a wage response: gold-belt activity raises the wage formal firms must pay, which lifts the survival cutoff z_{dt}^* . The household surveys we use for the labour-allocation results record occupations but not earnings, so we test the wage margin on the Ghana Living Standards Survey instead.

The GLSS district identifier is an unlabelled integer. It follows the Ghana Statistical Service scheme region $\times 100 +$ within-region sequence, and the 2000 Population and Housing Census carries the matching label set, which decodes all 110 GLSS5 codes exactly. GLSS5 is fielded on the 110-district frame of that census, so we aggregate Suit_d from the 2021 frame back to the 2000 districts as an area-weighted mean over matched successor districts; 102 of the 110 GLSS5 districts match, covering 8,162 of 8,687 households.

GLSS5 was collected in 2005/06, before the 2008 arrival of mechanised extraction. It therefore cannot test the model’s prediction, which concerns the post-boom equilibrium. What it can do is check the identifying premise: if gold-belt districts already paid more before the boom, the post-boom labour-reallocation results would be difficult to attribute to mining. Table A.8 shows they did not. The gradient of real employment income on Suit_d is +0.048 and imprecise, and neither the level of employment income nor the share of households with wage income shows a gradient either. None of the three is statistically significant under the wild cluster bootstrap.

Two limitations bound this. Successor districts that share a name stem, such as Adansi East and Adansi West, cannot be separated in the aggregation and receive a common value of $Suit_d$, which attenuates the within-region variation. And with ten regions the design has few clusters, which is why we report bootstrap p -values. The post-boom test remains open: GLSS6 and GLSS7 carry district identifiers on the 216-district frame of the 2010 census, for which we do not currently hold the label set.

OA.10 Spatial regression discontinuity designs

The cross-sectional estimates in Section 6 exploit continuous variation in geological suitability, $Suit_d \in [0, 1]$. A remaining concern is that districts with higher suitability may also differ systematically along other dimensions correlated with economic activity, such as infrastructure, historical settlement patterns, or agro-ecological conditions. This section asks whether firm counts and firm size change discretely at the geological contact separating gold-bearing from non-gold-bearing formations, which is determined by subsurface rock created billions of years ago and predates contemporary administrative, ethnic, and infrastructural boundaries.

We report this design as a robustness exercise rather than as the paper’s identification, for two reasons. The bandwidth that isolates the contact leaves 52 districts in 16 regions, and with that few clusters the wild bootstrap does not reject a null effect at the ± 50 km bandwidth even though the cluster-robust t -statistic does. And the firm-side discontinuity appears only at the rock contact, not at the boundary of realised mining activity where the household response appears (Table OA.6). The estimates below are therefore informative about the gold-bearing corridor as a whole; they do not on their own establish that artisanal extraction is what moves firm counts.

Boundary 1: BGS Mobile Belt (rock contact). The first candidate is the Ghana 1:5M geology shapefile from the BGS Africa Groundwater Atlas (africagroundwateratlas.org). The atlas distinguishes four formations: Precambrian Mobile/Orogenic Belt, Precambrian Craton, Voltaian sedimentary basin, and coastal sedimentary basins. The Mobile/Orogenic Belt is the greenstone-belt corridor that hosts essentially all of Ghana’s gold, both in primary lode and in alluvial form: it covers 10,183 km² ($\approx 4\%$ of Ghana’s land area) and intersects 36 of the 260 districts. The Craton interior is older basement rock that carries no significant gold; we therefore code it together with the Voltaian and coastal basins as not-Birimian. Panel (a) of Figure OA.1 maps the BGS polygon and the district centroids on each side of the contact within the ± 50 km firm-side bandwidth.




Figure OA.1: Two candidate gold-belt boundaries for the sharp-jump RDD, plotted side by side at the same scale. Panel (a) shows the BGS Birimian Mobile/Orogenic Belt (yellow shading) with the rock-contact line traced in black and district centroids within the ± 50 km firm-side bandwidth coloured by treatment status (red = Mobile-Belt side, blue = Voltaian/other side, grey = outside bandwidth); this is the boundary used in the firm-side RDD. Panel (b) shows the Girard ASGM-suitability polygon (red shading) with the BGS Mobile Belt overlaid as a dashed grey outline for direct comparison; DHS clusters within ± 10 km of the Girard edge are plotted as red (inside) and blue (outside) dots. The Girard polygon is broader than the BGS rock contact because it follows observed alluvial-deposit and coastal-fringe ASGM activity in addition to the bedrock corridor. Map lines delineate study areas and do not necessarily depict accepted national boundaries.

Boundary 2: Girard ASGM-suitability polygon (activity contact). The second candidate is the [Girard et al. \(2025\)](#) ASGM-suitability polygon ($Suit_d$ in Table 5 is the area share of each district inside this polygon). The Girard polygon is broader than the BGS Mobile Belt: it follows observed gold-mining activity in Africa and includes alluvial deposits along the Volta Lake flank and the coastal fringe, not only the bedrock corridor. The polygon intersects 248 districts and is the natural boundary for asking “what happens at the edge of where mining is actually taking place?”. Panel (b) of Figure OA.1 maps the Girard polygon with the BGS Mobile Belt overlaid as a dashed outline for contrast, and shows the DHS clusters within ± 10 km of the Girard edge.

The choice of spatial boundary matters because the geological and activity-based measures capture distinct margins of exposure to mining. The [Girard et al. \(2025\)](#) ASGM-suitability polygon is constructed to trace realised artisanal mining activity, in-

cluding alluvial deposits and coastal-fringe extraction zones. By contrast, the BGS boundary follows underlying Birimian geology and therefore captures the broader mineral corridor independently of whether extraction is currently active.

The running variable is

$$d_d \equiv \text{sgn}(\text{inside}_d) \cdot \min_{x \in \partial\mathcal{B} \cap \text{int}(\text{Ghana})} \text{dist}(c_d, x),$$

the signed great-circle distance (km) from the district centroid c_d to the nearest point on the Mobile-Belt boundary $\partial\mathcal{B}$ interior to Ghana, taken positive when the centroid lies inside the Mobile-Belt polygon. Treatment is $\text{MobileBelt}_d = \mathbf{1}\{d_d > 0\}$. We restrict the sample to $|d_d| \leq 50$ km or ± 100 km of the contact and report both bandwidths. Figure OA.1 (panel a) shows the running-variable cutoff and the districts that fall inside the ± 50 km bandwidth on each side.

The Birimian Mobile/Orogenic Belt continues westward into Côte d’Ivoire and northward into Burkina Faso, so the Mobile-Belt polygon’s western and northern edges within Ghana are largely the international frontier. We explicitly exclude those national-border segments from the running-variable contact $\partial\mathcal{B}$, leaving only the internal Birimian-Voltaian and Birimian-coastal-basin contacts. The practical consequence is that the RDD comparison is asymmetric: Mobile-Belt districts in Western Region that sit deep inside the belt (close to Ghana’s frontier with Côte d’Ivoire) drop out of the ± 50 km bandwidth because they are far from the internal contact. The identifying variation comes from Mobile-Belt districts whose centroids lie within 50 km of the eastern or southeastern edge of the corridor, where they can be compared with adjacent Voltaian or coastal sedimentary districts, and not from the high-suitability Western interior. This is standard practice for spatial RDDs in single-country samples (e.g. Dell, 2010).

Estimating equation. For each firm-side outcome Y_d we estimate

$$\log Y_d = \tau \text{MobileBelt}_d + \beta d_d + \mu_{r(d)} + \varepsilon_d, \quad (\text{OA.3})$$

where $\text{MobileBelt}_d = \mathbf{1}\{d_d > 0\}$, βd_d is a local-linear control in distance, and μ_r absorbs region fixed effects. Standard errors are clustered at region. The identifying assumption is that, conditional on d_d and the controls, all determinants of Y_d other than geological gold suitability are continuous at the contact.

The primary threat is climate. Ghana’s Mobile-Belt corridor sits in the wetter, cooler southwest, while the Voltaian basin occupies the drier north. Within a ± 50 km band-

width, mean 2001–2010 temperature is 0.34°C lower and CHIRPS precipitation is 61 mm higher on the Mobile-Belt side (both differences are statistically significant at the 5 percent level), while pre-period log night-lights are balanced. Because the firm-side sample near the contact is small, the specifications reported here do not add climate covariates; the climate gradient is one reason we treat the design as robustness only. A second threat, long-run sorting onto the gold belt, is real but does not undo identification: the RDD recovers the equilibrium effect of MobileBelt_d inclusive of historical sorting, which is the parameter the cross-sectional comparison needs.

Firm-side estimates. Table OA.6 reports the firm-side estimates. At the rock contact (Panel A), crossing into the Mobile Belt is associated with 0.349 fewer log establishments at ± 50 km and 0.455 fewer at ± 100 km, and with larger mean establishment size. With 16 region clusters, the ± 50 km estimate is not statistically significant under the wild cluster bootstrap; only the ± 100 km estimate is. At the boundary of the Girard ASGM polygon (Panel B) the firm-count discontinuity is positive and imprecise at every bandwidth. The firm-side discontinuity therefore loads on the geological corridor rather than on realised mining activity, which is one reason we do not treat the cross-sectional firm gradient as an effect of mining. Table OA.7 shows that the sign is negative across polynomial orders, kernels and bandwidths, but the magnitude is not stable and the local-quadratic estimates are close to zero at narrow bandwidths: the design is underpowered. Figure OA.2 plots the firm-count and size discontinuities at the rock contact.

Table OA.6: Firm-side discontinuity at both candidate gold-belt boundaries

	$\log N_d$ firms	$\log$ mean firm size	Districts
<i>Panel A. BGS Birimian Mobile/Orogenic Belt (rock contact)</i>			
± 50 km	−0.349 (0.142) {0.188}	+0.107* (0.056) {0.062}	52
± 100 km	−0.455*** (0.055) {0.000}	+0.049 (0.020) {0.188}	66
<i>Panel B. Girard ASGM-suitability polygon (activity boundary)</i>			
± 10 km	+0.265 (0.204) {0.208}	−0.027 (0.079) {0.761}	78
± 25 km	+0.186 (0.121) {0.140}	−0.080 (0.063) {0.197}	141
± 50 km	+0.074 (0.122) {0.547}	−0.060 (0.055) {0.270}	168

Notes: Each cell is a separate local-linear RDD on the IBES 2024 cross-section with signed distance to the relevant boundary as a linear control and region fixed effects, matching Equation (OA.3). Panel A uses the BGS Birimian Mobile/Orogenic-Belt rock contact ; Panel B uses the [Girard et al. \(2025\)](#) ASGM-suitability polygon, at tighter bandwidths because that polygon is roughly ten times larger. Cluster-robust standard errors in parentheses; wild cluster bootstrap p -values in braces (9,999 replications, Rademacher weights, imposing the null). Stars are assigned on the bootstrap p -value, not the analytic one. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table OA.7: Polynomial, kernel, and bandwidth robustness for the firm-side RDD on $\log N_d$ firms

Specification	Bandwidth around $d_d = 0$ (km)			
	± 25 (1)	± 50 (2)	± 75 (3)	± 100 (4)
Linear, uniform	−0.194 (0.320)	−0.349** (0.148)	−0.476*** (0.109)	−0.455*** (0.057)
Quadratic, uniform	−0.006 (0.745)	−0.063 (0.419)	−0.175 (0.245)	−0.366* (0.206)
Linear, triangular	−0.215 (0.501)	−0.296 (0.220)	−0.386** (0.151)	−0.428*** (0.102)
Observations	37	52	59	67
Region FE	Yes	Yes	Yes	Yes
SE clustering	Region	Region	Region	Region

Notes: Each cell reports $\hat{\tau}$ on the Mobile-Belt indicator from a separate RDD on $\log N_d$ firms (IBES 2024, drop-LSM), with region fixed effects and the indicated functional form, kernel, and bandwidth. The linear-uniform ± 50 km cell corresponds to Panel A of Table OA.6; standard errors here are cluster-robust. The sign is negative in every cell, but magnitude and precision vary with bandwidth and functional form. Standard errors clustered by region. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.




Figure OA.2: Firm-side discontinuity at the Birimian-Voltaian rock contact. Binned means of (a) $\log N_d$ firms and (b) $\log$ mean firm size, after partialling out region fixed effects and recentring at the grand mean; signed distance d_d on the horizontal axis; local-linear fits with 95% pointwise confidence bands; vertical dashed line at the cutoff $d_d = 0$; black vertical bracket at $d_d = 0$ traces the RDD jump $\hat{\tau}$. IBES 2024 cross-section, drop-LSM, ± 50 km bandwidth, 5 km bins. The in-figure coefficient labels report $\hat{\tau}$, its clustered standard error, and the p -value from the cluster-robust specification. The shaded bands are pointwise prediction CIs on each side; the regression test on τ uses the regression's clustered SE, not the visual overlap of the two bands.




Figure OA.3: McCrary-style density of district centroids in signed distance to the BGS Mobile-Belt contact, restricted to the ± 100 km bandwidth. Voltaian / non-gold side in blue, Mobile-Belt side in red; vertical dashed line at the cutoff $d_d = 0$. The cutoff is a ≈ 2 -billion-year-old rock contact whose location cannot have been manipulated by modern districts; the figure is included for completeness.

Sorting check. Manipulation of the running variable is mechanically impossible in this design: d_d is the distance from a 2026 GSS district centroid to a ≈ 2 -Gya rock contact, and the rock contact was set long before any district was drawn. Figure OA.3 reports the distribution of district centroids by signed distance for completeness. Districts cluster near the boundary because Mobile-Belt districts are geographically smaller (the corridor covers 10,183 km² in 36 districts, roughly 280 km² each), but there is no spike in density at the cutoff that would indicate endogenous sorting.

Household-side discontinuity. Activity boundaries and geological boundaries identify distinct margins. Table A.7 compares two spatial discontinuity designs applied to the same household labour-supply outcomes. The first uses the ASGM-suitability polygon, which traces observed artisanal mining activity. The second uses the BGS geological contact separating gold-bearing and non-gold-bearing formations. In both cases, identification comes from local comparisons within a ± 10 km bandwidth around the respective boundary.

The two boundaries give different answers. At the ASGM activity boundary, male labour-force participation rises by 5.5 percentage points, an estimate statistically significant at the 5 percent level. The estimate is close to the 5 percentage point effect obtained from the exposure specification in Table 5, despite the fact that the two designs exploit fundamentally different variation. The exposure design combines global gold-price

movements with cross-sectional exposure, whereas the RDD relies solely on spatial discontinuities at the mining boundary. The similarity in magnitudes is consistent with both designs picking up the same labour-supply response.

No comparable discontinuity emerges at the geological boundary. Using the BGS rock contact, the estimated effect is economically small and statistically indistinguishable from zero. The contrast is informative about the mechanism. The ASGM polygon is constructed to capture realised mining suitability, including alluvial deposits and historically worked mining zones. The BGS boundary instead traces the broader Birimian geological corridor, much of which contains little active artisanal mining. Household labour supply therefore responds to the spatial footprint of realised mining activity rather than to geology alone.

Taken together, the household and firm-side results identify different margins of adjustment. Household responses are concentrated where artisanal mining operates, while the firm-side discontinuity in Table OA.6 reflects broader equilibrium effects associated with the wider gold-bearing corridor. The two designs therefore pick up different things: labour reallocation around active mining sites, and the wider pattern of the gold-bearing corridor. Figure OA.4 visualises the discontinuity in male employment at the ASGM boundary.




Figure OA.4: Household-side discontinuity at the ASGM-activity boundary. Binned means of male share-working (DHS Ghana 1993–2014) against signed distance d_c to the [Girard et al. \(2025\)](#) ASGM-suitability polygon; local-linear fits with 95% confidence bands; ± 15 km bandwidth, 2 km bins.

OA.11 Night-lights: additional results

Enforcement window and differential trends. Columns 1–2 of Table A.9 re-estimate the collapsed Vanguard DiD dropping 2020, the year in which COVID-19 depressed both night-lights and informal mining. The decline shrinks from about 9 to about 7 percent and is no longer statistically distinguishable from zero.

Columns 3–4 test the Bartik reduced form on night-lights. Because $ASGM_{dt} = Suit_d \times \log P_t$ is a fixed exposure share interacted with a near-monotone time series, it is collinear with any differential trend between gold-belt districts and the rest of Ghana; district and year fixed effects do not absorb this. Allowing gold-belt districts a linear trend of their own reverses the sign of the reduced form, from -0.553 to $+1.186$, with the trend term itself strongly negative. The negative Bartik–nightlights relationship is therefore not separately identified from secular divergence in nightlight growth. It does not identify an effect of mining exposure.

Enforcement as a control in the exposure panel. A single cross-section cannot condition on Operation Vanguard, since a national policy has no cross-sectional variation. The panel can, and doing so shows that the exposure gradient estimated across the enforcement period is not the same object as the one estimated before it.

Table OA.8 makes the comparison. Estimated on 2001–2016, before the crackdown, the exposure coefficient on log night-lights is -0.283 . Extending the window to 2020 doubles it to -0.551 . Adding $Suit_d \times \text{Post-2017}$ returns the exposure coefficient to -0.284 , its pre-enforcement value, and attributes the difference to the enforcement interaction itself, which enters at -1.034 . Roughly half of what the full-window specification attributes to mining exposure is the 2017 crackdown. Column 4 adds a gold-belt trend and reverses the sign of the exposure term, as in Table A.9.

Two implications follow. The night-lights exposure specification is not robust to either the enforcement period or a differential trend, which is why we do not use the exposure panel for night-lights. And the household panel is unaffected: the DHS waves end in 2014, three years before the crackdown, so no post-enforcement variation enters (12) there.

Table OA.8: The exposure gradient before and after conditioning on enforcement

	(1) 2001–16	(2) 2001–20	(3) + Post	(4) + Post, trend
ASGM _{dt} exposure	-0.283** (0.129)	-0.551*** (0.156)	-0.284** (0.129)	1.478*** (0.281)
Suit _d × Post-2017			-1.034*** (0.184)	0.429*** (0.142)
Suit _d × (year – 2000)				-0.253*** (0.046)
District-years	4032	5040	5040	5040

Notes: District-year night-lights panel, drop-LSM, district and year fixed effects throughout, outcome $\log(1 + \text{NTL}_{dt}^{\text{sum}})$. Column 1 restricts to the pre-enforcement window. Columns 3 and 4 add Suit_d × Post-2017 and a Suit_d-specific linear trend. Standard errors clustered by district. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Artisanal-prone versus industrial-mine districts. Panel A of Table OA.9 estimates the post-2017 change separately for artisanal-prone districts ($n = 141$) and the eight districts hosting industrial gold mines, relative to low-suitability districts. In 2014–2020 the two coefficients are almost identical (–0.091 and –0.087); the industrial-mine estimate is imprecise because the group is small. The comparison therefore cannot tell whether the decline is specific to artisanal mining, and Figure OA.5 carries the same caveat. Panel B adds a triple interaction with flood years. The coefficient is negative and statistically significant at the 5 percent level, consistent with enforcement biting hardest where rainfall-sensitive mining had been most active, but it builds on the descriptive rain pattern below and we do not lean on it.

Table OA.9: Where does Vanguard bite harder? LSM vs. ASM split and the climate triple-difference

	2010–2020 (1)	2014–2020 (2)
<i>Panel A. LSM-anchor vs. ASM-prone DiD on $\log NTL_{dt}$</i>		
LSM-anchor $\times$ Post-2017	−0.181** (0.092)	−0.087 (0.094)
ASM-prone $\times$ Post-2017	−0.122*** (0.038)	−0.091** (0.040)
Districts (ASM-prone, n)	141	141
Districts (LSM-anchor, n)	8	8
Districts (low-suit, n)	111	111
District FE	Yes	Yes
Year FE	Yes	Yes
SE clustering	District	District
<i>Panel B. Triple difference: enforcement reduces the flood-amplified gradient</i>		
$Suit_d \times Flood_{dt} \times \text{Post-Vanguard}$	—	−0.062** (0.029)
Lower-order interactions, FE	—	Yes
Districts	—	252 (drop-LSM)
SE clustering	—	District

Notes: Panel A reports the collapsed DiD coefficient on each treatment group $\times$ Post-2017, with low-suitability districts ($Suit_d < 0.5$, no LSM) as the omitted reference. The 2014–2020 column is the preferred window. Panel B reports the triple-difference on the drop-LSM 2014–2020 panel with all lower-order interactions and district + year fixed effects. Standard errors in parentheses below coefficients; clustered by district. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.




Figure OA.5: Operation Vanguard event study, 2014–2020 sample. Coefficients on $\text{LSM}_d \times \mathbf{1}\{t - 2017 = k\}$ and $\text{ASM}_d \times \mathbf{1}\{t - 2017 = k\}$, relative to low-suitability districts ($\text{Suit}_d < 0.5$, no LSM). With only eight LSM-anchor districts, the two groups cannot be separated statistically.

Rain and night-lights. Table OA.10 interacts suitability with contemporaneous climate states in the night-lights panel, $\log \text{NTL}_{dt} = \beta (\text{Suit}_d \times \text{Climate}_{dt}) + \alpha_d + \tau_t + e_{dt}$. Climate shocks do not follow the near-monotone path of the gold price, so this specification is less exposed to the differential-trend problem above. In 2014–2020, wet years and flood years amplify the suitability gradient in night-lights, while drought and heat interactions are small and insignificant. This is consistent with artisanal gold mining being water intensive (sluicing, dredging, washing), and it runs against the view that drought pushes labour into mining. We report it as descriptive.

Table OA.10: Rain, not drought: panel reduced form on log nightlights

Sample window	2000–2020 (1)	2010–2020 (2)	2014–2020 (3)
$\text{Suit}_d \times \text{Rainfall}_{dt} z_{dt}$	+0.012* (0.007)	+0.020** (0.008)	+0.053*** (0.017)
$\text{Suit}_d \times \text{Flood}_{dt} (z > +1)$	+0.024 (0.021)	+0.033 (0.032)	+0.168*** (0.056)
$\text{Suit}_d \times \text{Drought}_{dt} (z < -1)$	−0.012 (0.021)	−0.016 (0.028)	−0.032 (0.043)
$\text{Suit}_d \times \text{Tmax } z_{dt}$	+0.012 (0.021)	−0.011 (0.024)	+0.013 (0.031)
$\text{Suit}_d \times \text{Heat}_{dt}$	+0.013 (0.044)	−0.031 (0.032)	+0.016 (0.030)
District-year obs.	3,318	1,738	1,106
District FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Drop LSM districts	Yes	Yes	Yes
SE clustering	District	District	District

Notes: Each cell is a separate regression of $\log \text{NTL}_{dt}$ on the listed $\text{Suit}_d \times \text{Climate}_{dt}$ interaction, plus district and year fixed effects, on the drop-LSM district panel. The four climate states are: rainfall z (continuous), a flood indicator ($z > +1$), a drought indicator ($z < -1$), and continuous Tmax z / a binary heat indicator. The 2014–2020 column is the window highlighted in the text; it matches the Vanguard window of Table 7 and avoids the 2008–2014 gold-belt boom that contaminates pre-2014 observations. Wet years and floods amplify the suitability gradient on aggregate activity; drought and heat do not. Standard errors clustered by district in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

OA.12 A no-galamsey counterfactual

This section puts the reduced-form estimates on a common scale through the labour-allocation block of the model in Section 3. The exercise is partial equilibrium and heuristic.

Calibration. We estimate the Pareto shape parameter k with a Hill estimator on the upper quartile of the IBES persons-engaged distribution (drop-LSM sample), which gives $\hat{k} = 1.73$ above three persons. Combining the model identity in (8), $\partial \log N_d / \partial S_d = -k \partial \log w_d / \partial S_d$, with the cross-sectional coefficient of -0.216 implies a wage elasticity to gold suitability of about 0.125. If the establishment gradient were entirely a wage effect, a fully suitable district would pay about 12.5 log points more than a non-suitable one.

Counterfactual. We set $S_d = 0$ for every district and apply the estimates at the sample-mean suitability of 0.69 (Figure OA.6). Relative to a world without the gold belt, galam-

sey zones contain about 13.8 percent fewer formal establishments, each surviving establishment is about 1.7 percent larger, and male labour-force participation is about 4.3 percentage points higher. Log points are converted to percent as $100(e^x - 1)$, and participation is in percentage points because the outcome is a share. We do not report an aggregate-activity bar, because the night-lights exposure regression is not separately identified from a gold-belt trend (Table OA.8).

The exercise has clear limits. The wage elasticity comes from a single cross-sectional moment under a Pareto assumption, the propagation ignores spillovers to neighbouring districts, and the establishment gradient is itself descriptive (Section 6). What the exercise shows is that the estimates are consistent with one another at first order: a 14 percent shortfall in establishments maps through the labour-market clearing condition into a participation rise of about four percentage points.




Figure OA.6: Counterfactual percent differences at sample-mean $Suit_d = 0.69$ relative to the no-galamsey baseline ($S_d = 0$). Each bar applies the corresponding reduced-form elasticity to mean $Suit_d$. Sources: cross-section IBES (establishment count and size) and the DHS event study (participation, in percentage points). Log points are converted as $100(e^x - 1)$.

OA.13 Sample construction

Table OA.11 traces the district count through each step. Two losses matter. The main tables match IBES district names to the GSS frame with a normalisation-plus-fuzzy procedure that retains 239 districts, whereas the RDD sample merges on the exact district-name string and retains 176. The 63-district gap is a name-matching artefact rather than

a bandwidth restriction, and it applies before any bandwidth is imposed.

Table OA.11: Sample construction

Step	Districts	Used in
All GSS districts (2021 boundaries)	260	—
Drop 8 LSM-anchor districts	252	analysis sample
Matched to IBES 2024 (normalised + fuzzy names)	239	Tables A.1 to 4
with non-missing climate interaction	229	Table 6
with non-missing cumulative forest loss	223	Table OA.4
RDD firm sample (exact district-name string merge)	176	Table OA.6
within ± 50 km of the BGS rock contact	52	Table OA.6 col. 1, 3
within ± 100 km of the BGS rock contact	66	Table OA.6 col. 2, 4
within ± 50 km of the ASGM boundary	168	Table OA.6 Panel B
Nightlights panel, drop-LSM, 2014–2020	252	Table 7 col. 6

Notes: District counts at each step of the build. The IBES match normalises district names and applies fuzzy matching at a 88% similarity threshold; the RDD sample merges on the exact district-name string, which is why it starts from a smaller frame before any bandwidth is applied.