Open Access Indonesia Journal of Social Sciences

Vol. 9 · No. 4 · 2026 · pp. 210–228 · Article ID 333

e-ISSN 2722-4252 · *Journal homepage: journalsocialsciences.com/index.php/oaijss*

ORIGINAL ARTICLE

Fiscal Rigidity and Labour Exposure in Indonesia's Extractive Districts: A Dynamic Panel Analysis of Just Transition Readiness, 2015–2024

Hanifah Yasin[1,*], **Iqbal Anugerah**[2]

[1] *Department of Microeconomy, Enigma Institute, Palembang, Indonesia*
[2] *Department of Monetary Economics, Enigma Institute, Palembang, Indonesia*

ARTICLE INFORMATION

Received: June 29, 2026
Revised: July 20, 2026
Accepted: August 19, 2026
Available Online: August 28, 2026
Published: August 30, 2026

Keywords: decarbonisation; dynamic panel data; fiscal federalism; Indonesia; just transition

***Corresponding author**
Name: Hanifah Yasin
E-mail: hanifah.yasin@enigma.or.id
ORCID: 0009-0009-5497-8710
DOI: https://doi.org/10.37275/oaijss.v9i4.333

ABSTRACT

Background: Indonesia's net-zero commitment attaches a schedule to the decline of an activity that simultaneously funds district development budgets through derivative revenue sharing and employs a large share of the local workforce, yet the fiscal and labour channels of that exposure have never been estimated on the same subnational units.

Objective: To quantify how contractions in natural-resource revenue sharing and in mining value added transmit to district capital expenditure and open unemployment, and to derive a district-resolution just-transition readiness typology.

Methods: A balanced panel of 118 extractive-producing Indonesian districts observed annually from 2015 to 2024 (1,180 district-year observations) was drawn from the 514-district national budget frame. Two-way fixed effects with Driscoll–Kraay standard errors and two-step System GMM with the Windmeijer correction were used; precision, long-run multipliers, budget incidence, and a quadrant projection were derived from the archived estimates.

Results: A 10% fall in revenue sharing was associated with a 4.78% contraction in capital expenditure (95% CI −5.62 to −3.94; $p < 0.001$), rising to 7.76% in the long run. Capital expenditure absorbed 35.4% of the shock while constituting 18.2% of spending. Mining contraction was associated with higher unemployment (−0.624; $p < 0.001$), amplified by extractive concentration and attenuated by educational attainment; however, under an 80% output decline, only 17.8% of the exposed employment would appear in the measured rate.

Conclusion: Fiscal exposure is severe and measurable; labour exposure is severe and largely unmeasured. H1 to H5 were supported; H6 is descriptive. A transition transfer indexed to capital-expenditure incidence and district instrumentation of the informal margin are the priorities.

All authors reviewed and approved the final manuscript.

1. Introduction

Decarbonisation has stopped being a distant scenario for resource-producing regions and has become a scheduled fiscal event. National net-zero commitments, negotiated phase-out agreements and transition-finance packages now attach dates to the decline of activities on which subnational budgets and local labour markets depend, and the comparative record shows that the resulting adjustment costs are borne by particular places rather than distributed across the economies that set the targets [1,2,3]. Historically observed rates of fossil-fuel decline have been slower than climate-consistent pathways require, and the losses implied by accelerating them are distributed unevenly along the ownership and revenue chains that connect production to the places that host it [4,5]. Across Southeast Asia the pace of that adjustment is set by government policy and by the operating lives of existing plant rather than by technology cost[6,7].

Indonesia is the largest instance of this problem in Southeast Asia. Its commitment to net-zero emissions by 2060 or sooner, reinforced by the Just Energy Transition Partnership and the Energy Transition Mechanism and operationalised through a coal-plant early-retirement roadmap, sets a schedule against which the revenue and employment base of its producing districts must be expected to contract [8]. What makes the Indonesian case distinctive is not the schedule but the fiscal architecture underneath it: producing districts receive a derivative share of the rents generated by the activity that the schedule is designed to end, so the instrument that funds local development is mechanically indexed to the thing being phased out. It is not the only Indonesian energy instrument whose economic case has been settled after adoption rather than before it [9,10].

Three bodies of theory converge on this configuration. Scholarship on fiscal federalism and transfer dependency predicts that where a large share of revenue arrives as an earmarked transfer while operating commitments are contractually fixed, the burden of any revenue shock concentrates on the one line that can legally be cut in-year, namely discretionary capital spending; the comparative evidence from developing-country budgets documents exactly this asymmetry, with current expenditure ratcheting upward in good times and public investment collapsing in bad ones [11,12]. Work on carbon lock-in, re-specified in its contemporary institutional rather than technological form, predicts that the dependency will not dissolve through anticipation, because licensing regimes, state-enterprise interests and revenue-sharing formulas actively reproduce the configuration that decarbonisation threatens [7,13]. Labour-market scholarship on displacement and specific human capital predicts that an output shock will not generate an equivalent re-employment flow, because a sectoral wage premium raises reservation wages while occupation-specific skills sit in a poorly connected region of the mobility network [14,15]. Energy-justice scholarship supplies the criteria — distributional, procedural and restorative — that turn these three mechanisms into an assessment of who is ready for the transition and who is not [16,17].

The evidence base does not yet allow those predictions to be tested jointly. The systematic map of coal phase-out experience finds regional fiscal compensation among the least evidenced instruments anywhere, and the cases that dominate the literature are North American and European, where subnational fiscal structures bear little resemblance to a derivative revenue-sharing system [18,19,20]. Reviews of just-transition policy find fiscal instruments concentrated at national level and labour instruments studied separately from them [17,21,22]. The consequence is a specific and consequential gap. District-level vulnerability analyses now exist for Indian coal districts and for Polish hard-coal cities, and a provincial study for China estimates the employment and fiscal consequences together [20,23,24]; none exists for Indonesia, and none of them estimates both channels on the same units with the same estimator. No existing study can therefore say whether a district that is fiscally exposed to decarbonisation is also occupationally exposed to it. That is not an academic question. If the two exposures coincide, one place-based instrument can address both; if they are separable, two instruments are required and each must be targeted differently.

This study addresses the gap in the setting where it binds hardest. The analytic frame is the set of 118 Indonesian districts and cities that produce coal, petroleum, natural gas or metallic minerals, observed annually from 2015 to 2024. These are not scattered units: they cluster in the coal basins of South Sumatra (Muara Enim, Lahat, Musi Banyuasin), the oil, gas and coal districts of East Kalimantan (Kutai Kartanegara, Paser, Kutai Timur), Riau and South Kalimantan, where a single commodity cycle moves the district budget, the labour market and the local political settlement together. They are drawn from the same 514-district national budget frame as the 396 non-extractive districts that make up the remainder, which means the exposed group can be located precisely within the national fiscal system rather than treated as a special case.

The study had three aims: to estimate the transmission of natural-resource revenue-sharing contractions to district capital expenditure and of mining-output contractions to district open unemployment on the same panel; to convert those elasticities into the budget-incidence and labour-absorption quantities that a finance ministry and a manpower ministry would need in order to size an instrument; and to derive a readiness typology at the resolution at which Indonesian development budgets are actually spent. Six hypotheses were specified in advance of the analysis reported here and are developed in §1.1.

1.1. Theoretical framework and hypothesis development

Fiscal federalism and transfer dependency. A derivative revenue-sharing transfer is a claim on an upstream tax base rather than a locally raised revenue, and the fiscal-federalism literature treats such transfers as weakening the accountability link between spending and taxation while introducing a volatility that local budgets are poorly equipped to absorb [11,25]. Indonesian district budgets are politically contested objects whose composition responds to local competition, not technical allocations, which means the composition of any adjustment is itself an outcome to be explained [8,25]. The comparative record on intergovernmental fiscal relations in developing countries places this design at the dependent end of the spectrum [11,26]. Because personnel and operating commitments are contractual and capital projects are annually appropriated, the prediction is directional and sharp.

H1: Contractions in natural-resource revenue sharing are associated with contractions in district capital expenditure [11,12].

H2: Own-source revenue and the general allocation grant are positively associated with capital expenditure but with smaller elasticities than the resource transfer they would have to replace, so they buffer without substituting [12,25].

Carbon lock-in and path dependence. Lock-in is now understood as an institutional and political configuration rather than a property of capital stock: the persistence of coal in large democracies is sustained by subnational revenue and employment dependencies that make withdrawal politically costly for exactly the actors who would have to implement it [7,27]. Phase-out is therefore a distinct governance task with its own instrument requirements, not a residual of deployment policy [3,13]. The empirical implication for a dynamic panel is state dependence: last year's budget and last year's labour market constrain this year's.

H5: Both outcomes exhibit substantial state dependence, so the long-run response to a permanent shock exceeds the short-run response [3,28].

Labour hysteresis and specific human capital. The dominant long-run cost of displacement is not the duration of unemployment but the move to a lower-paying employer, and displacement costs vary by an order of magnitude across institutional settings facing similar shocks [14,29]. Where an occupation sits in a poorly connected region of the mobility network, and where the local productive structure is unrelated to the activities a green economy requires, re-entry is slow and downward [15,30]. Industrial extraction raises local incomes but concentrates the gains and crowds out alternative activity, so the wage premium and the later immobility are two expressions of the same specialisation [31,32]. In developing economies a further margin operates: informal employment absorbs formal-sector shocks and mutes the measured unemployment response [33,34].

H3: Contractions in mining and quarrying value added are associated with increases in the district open unemployment rate [19,24].

H4: Extractive employment concentration amplifies, and workforce educational attainment attenuates, that unemployment response [30,35].

Energy justice and just-transition readiness. Distributional justice asks who bears the burden, procedural justice who decides, and restorative justice what is owed to those already harmed [16,17]. Applied to a producing district, these criteria require an assessment along two axes that need not move together: the capacity of the budget to absorb a revenue shock, and the capacity of the workforce to absorb an employment shock. Existing readiness frameworks sequence fossil-fuel decline at national level and do not operate at district resolution [3,21].

H6: Fiscal dependency and labour readiness are separable dimensions that generate a four-quadrant readiness typology rather than a single vulnerability gradient [3,36]. The framework and the six propositions it generates are rendered in Figure 1, which also carries the estimated coefficient on each of the two focal paths; each proposition is tested in the Results and the coefficients shown in that figure are the estimates reported there.

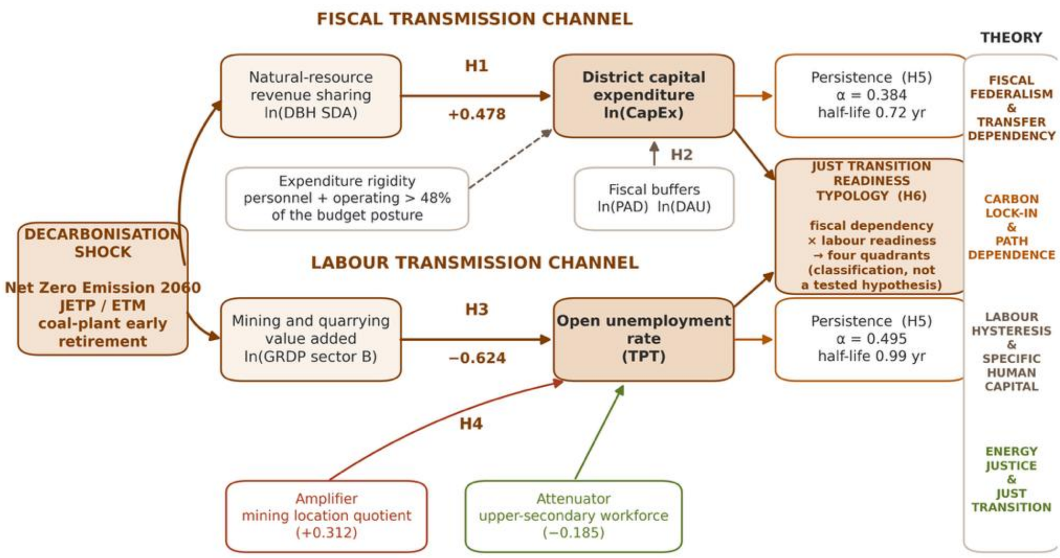

Figure 1. Conceptual and theoretical framework. The decarbonisation shock transmits through two channels estimated on the same panel: a fiscal channel from natural-resource revenue sharing to district capital expenditure, conditioned by expenditure rigidity and the two fiscal buffers, and a labour channel from mining value added to the open unemployment rate, amplified by extractive employment concentration and attenuated by workforce educational attainment. Persistence in both equations converts short-run responses into larger long-run ones. The two channels jointly define the readiness typology, which is a classification rather than a tested hypothesis. Coefficients shown are the estimates reported in Table 4; expenditure rigidity and the fiscal buffers are conditioning features of the budget rather than estimated interaction terms, and no product term is estimated in either equation. The theory rail on the right names the layer supplying each mechanism.

2. Methods

2.1. Study design and setting

This is a quantitative, non-experimental, retrospective analysis of a balanced district-year panel assembled from published Indonesian administrative and official statistical series. The unit of observation and of analysis is the district or city government and its economy; no individual is observed, and no inference about individuals is drawn anywhere in this paper. The setting is the population of 118 Indonesian districts and cities producing coal, petroleum, natural gas or metallic minerals, observed annually over the ten fiscal years from 2015 to 2024. The archived record states that window but does not document why it begins and ends where it does, and no reason is supplied here.

The estimand is a within-district association between a revenue or output contraction and a budget or labour-market outcome, conditional on district and year effects. It is not a treatment effect. Decarbonisation is the policy context that makes the estimand relevant; it is not itself a treatment variable in either equation, and no claim of causal identification is made [37].

2.2. Data sources and panel construction

Three custodians supply the panel. The fiscal-transfer directorate of the Ministry of Finance supplies realised revenue and expenditure outturns for all 514 districts and cities over 2015–2024, covering oil-and-gas and mineral-and-coal revenue sharing, the general and special allocation grants, own-source revenue, capital expenditure and personnel expenditure. BPS-Statistics Indonesia supplies district-level labour-market aggregates from the National Labour Force Survey — sectoral employment structure, educational attainment, average wage and the open unemployment rate — and gross regional domestic product at constant 2010 prices. The Ministry of Energy and Mineral Resources supplies primary commodity production and mineral-and-coal royalties by district, which define membership of the producing group.

Table 1 sets out the construction of the analytic panel and two arithmetic reconciliations that the record permits. The producing and non-producing groups exhaust the national frame exactly (118 + 396 = 514), and the observation count is exactly the product of units and years (118 × 10 = 1,180), so the panel is balanced and no district entered or left the sample during the window. The 396 non-extractive districts were retained as a comparator group; no comparator estimate is reported in the archived output, and none is constructed here.

Table 1. Construction of the analytic panel and its arithmetic reconciliation.

Panel construction step	Value	Arithmetic check
District/city budget frame reported by the fiscal-transfer directorate	514	Source frame
Districts/cities producing coal, petroleum, natural gas or metallic minerals	118	22.96% of the frame
Non-extractive districts/cities retained as the comparator group	396	118 + 396 = 514 ✓
Observation years	2015–2024 (10 years)	Annual frequency
Analytic observations (producing districts)	1,180	118 × 10 = 1,180 ✓ balanced, no attrition
Comparator observations implied by the same window	3,960	Not estimated in the archived output
Full-frame observations implied by the same window	5,140	Upper bound on any extension of this design
Clusters used for inference	118	Inference degrees of freedom 117
Residual degrees of freedom under the stated fixed-effect structure	1,049	1,180 − 131 = 1,049

Note. The two arithmetic identities marked ✓ are internal-consistency checks on the archived record, not external verification: the producing and non-producing groups exhaust the national budget frame exactly, and the observation count is exactly the product of units and years, so the panel is balanced and no district entered or left the sample during the window. The rules by which the 118 producing districts were selected from the frame, and the treatment of districts created or merged during 2015–2024, are not documented in the archived output and are named as outstanding in §2.5. The residual degrees of freedom assume 118 unit effects, 10 year effects and four regressors per equation; the exact control vector is not archived, so this figure is an upper bound.

Every variable, its definition and unit, and the custodian and collection instrument behind it are documented in Table 2, which also records the provenance argument developed in §2.3.

Table 2. Variables, measurement, custodian and collection instrument.

Variable	Role	Definition and unit	Custodian and instrument
ln(CapEx)	Outcome, equation 1	Natural log of realised district capital expenditure, constant prices	Fiscal-transfer directorate, Ministry of Finance — audited budget outturn
Open unemployment rate	Outcome, equation 2	Unemployed as a percentage of the district labour force	BPS-Statistics Indonesia — National Labour Force Survey (SAKERNAS)
ln(DBH SDA)	Focal predictor, equation 1	Natural log of realised oil-and-gas plus mineral-and-coal revenue-sharing transfers	Fiscal-transfer directorate — transfer realisation ledger
ln(mining GRDP)	Focal predictor, equation 2	Natural log of gross value added in mining and quarrying (sector B), constant 2010 prices	BPS-Statistics Indonesia — regional national accounts
ln(PAD)	Fiscal buffer	Natural log of realised own-source district revenue	Fiscal-transfer directorate — audited budget outturn
ln(DAU)	Fiscal buffer	Natural log of the realised general allocation grant	Fiscal-transfer directorate — transfer realisation ledger
Mining location quotient	Amplifier	District share of employment in mining and quarrying divided by the national share	Derived from SAKERNAS employment counts
Upper-secondary workforce share	Attenuator	Share of employed persons with upper-secondary education or above	BPS-Statistics Indonesia — SAKERNAS educational attainment module
Commodity production and royalties	Sampling frame variable	Primary commodity output and mineral-and-coal royalties by district	Ministry of Energy and Mineral Resources — minerals and coal one-data systems
Lagged outcome	State-dependence term	One-period lag of the equation's own dependent variable	Constructed within the panel

Note. Measurement validity in this design rests on administrative and statistical provenance rather than on respondent-level scale properties, because no variable is an item, an index or a latent construct: each is a single audited accounting line, a national-accounts aggregate or a survey-weighted population estimate. Two features of that provenance are load-bearing. First, the fiscal, labour and production series are produced by three separate agencies under three separate legal mandates and three unrelated collection instruments. That separation holds ACROSS the two equations and for the sampling frame, which is why a shock recorded in the production system can be related to outcomes measured in two unrelated systems. It does NOT hold within either equation, and no such claim is made: all four fiscal variables come from the same audited ledger, so any ledger-wide reclassification or restatement moves regressor and outcome together, and the revenue-expenditure relation is an accounting identity as well as a behavioural one; likewise the unemployment rate, the location quotient and the educational-attainment share all come from the same labour-force survey, so their sampling errors are correlated within district-year by construction. Both are carried into §4.3. Second, the budget variables are post-audit realisation figures rather than appropriations, so they record what was actually spent. The precision of district-level survey estimates is not uniform across small districts, and the archived output does not report design effects; that limitation is carried into §4.3.

2.3. Variables, measurement and validity

Measurement validity in this design rests on administrative and statistical provenance, because no variable in either equation is an item, a scale or a latent construct. Each is a single audited accounting line, a national-accounts aggregate, or a survey-weighted population estimate produced under a standing legal mandate. Three properties of that provenance carry the validity argument.

First, the budget variables are post-audit realisation figures rather than appropriations, so they record what was spent rather than what was planned; this matters here because the hypothesis concerns in-year adjustment, which appropriations by construction cannot show. Second, the fiscal, labour and production series originate with three separate agencies operating under three separate mandates and three unrelated collection instruments — an audited ledger, a household survey and an industrial reporting system — so no shared measurement procedure can generate correlation between the predictors and the outcomes of either equation. Third, membership of the producing group is defined by a production and royalty record held by a fourth party to both the fiscal and the labour series, so sample definition is not endogenous to either outcome.

The custodian and instrument behind each variable are listed in Table 2, and the separation they establish is a claim about the relation between the two equations and about the sampling frame, not about the interior of either equation, and the distinction matters. Within the capital-expenditure equation all four variables come from the same audited ledger, so a ledger-wide reclassification or restatement would move regressor and outcome together, and the relation between realised revenue and realised expenditure is an accounting identity as well as a behavioural one. Within the unemployment equation the outcome, the location quotient and the educational-attainment share all come from the same labour-force survey, so their sampling errors are correlated within district-year by construction. Neither feature invalidates the design, but both bound what the cross-agency argument establishes, and both are carried into §4.3.

The limits of that provenance are equally specific. District-level estimates from a household survey are not uniformly precise across small districts, and the archived output does not report design effects or sampling errors; measurement error in the open unemployment rate is therefore likely to be heteroskedastic in district size. The deflator and base-year treatment applied to the fiscal series, and any winsorisation, are not documented. These are recorded as outstanding in §4.3 rather than assumed away.

2.4. Data analysis

Two equations are estimated on the same panel. The capital-expenditure equation is

$$\ln(CapEx_{it}) = \alpha_1 \ln(CapEx_{i,t-1}) + \beta_1 \ln(DBH_{it}) + \beta_2 \ln(PAD_{it}) + \beta_3 \ln(DAU_{it}) + \gamma' X_{it} + \mu_i + \lambda_t + \varepsilon_{it} \quad (1)$$

and the open-unemployment equation is

$$TPT_{it} = \alpha_2 TPT_{i,t-1} + \theta_1 \ln(MiningGRDP_{it}) + \theta_2 LQ_{it} + \theta_3 EDU_{it} + \delta' Z_{it} + \eta_i + \tau_t + \nu_{it} \quad (2)$$

where i indexes the 118 producing districts and t the ten years 2015–2024. CapEx is realised capital expenditure and TPT the open unemployment rate; DBH is realised natural-resource revenue sharing, PAD own-source revenue and DAU the general allocation grant; MiningGRDP is gross value added in mining and quarrying at constant 2010 prices, LQ the mining location quotient and EDU the share of workers with upper-secondary education or above. X_{it} and Z_{it} are the control vectors, μ_i and η_i the district effects, λ_t and τ_t the year effects, and ε_{it} and ν_{it} the idiosyncratic disturbances. The archived output does not enumerate the control vectors, so $\gamma' X_{it}$ and $\delta' Z_{it}$ are carried here as the source states them and are named as outstanding in §4.3. Because both equations contain the lagged dependent variable, the two are estimated jointly in the sense that they share a sample and an estimator, not as a system of simultaneous equations; no cross-equation restriction is imposed and none is archived.

Estimation follows the standard two-step sequence for a short, wide dynamic panel. Because both specifications include the lagged dependent variable, within-group estimation is biased at T = 10, and the archived output accordingly reports two-step System GMM with the Windmeijer finite-sample correction alongside two-way fixed effects with Driscoll–Kraay standard errors [28,38,39]. System GMM is the appropriate choice where the series are persistent, as they are here, because the level equations restore identification that differencing destroys [28,40]. Driscoll–Kraay standard errors are design-justified rather than defensive: commodity price cycles hit every producing district in the same year, so cross-sectional dependence is a structural feature of this panel, and clustering at the district level matches the level at which exposure is assigned [41,42]. At 118 clusters the asymptotic approximation is adequate but not exact; a jackknife or wild-bootstrap correction would tighten the claim, neither is available from the archived output, and both are recommended for the next iteration [43].

Instrument validity is assessed by the Hansen J test of overidentifying restrictions and by the Arellano–Bond tests for first- and second-order serial correlation in the differenced residuals [40,44]. Statistical significance is evaluated two-sided at $\alpha = 0.05$, and p values are reported to three decimal places with a floor at 0.001.

Because the archived output reports intervals for one coefficient only, four derivations were used to recover as much of the missing precision as the record permits, and each is labelled wherever it appears. (i) For the interval-reported coefficient the standard error was recovered as $(hi - lo)/(2 \times 1.96) = 0.04286$, and the Wald statistic and exact p value recomputed. (ii) For the remaining coefficients, a reported significance level of $p < 0.001$ implies $SE \le |b|/3.291$, from which a widest-possible 95% interval of $b \times (1 \pm 0.596)$ follows; these are reported as bounding intervals and never as confidence intervals. (iii) Long-run multipliers $\beta/(1 - \alpha)$ and adjustment half-lives $\ln(0.5)/\ln(\alpha)$ were computed from the autoregressive terms, and their sampling distributions obtained by Monte Carlo propagation with 10,000 draws under two precision scenarios, with percentile intervals reported. (iv) Effect sizes were expressed as Cohen's f^2 derived from the partial correlation implied by each test statistic, with the degrees-of-freedom convention stated in the table note; these are comparability metrics, not measured variance decompositions [45].

Two further quantities were derived arithmetically and are reported with their assumptions in the relevant table notes: the share of a revenue shock absorbed by capital expenditure alone, and the share of employment exposure that appears in the measured unemployment rate. Multiplicity across the eight-coefficient structural family was controlled by both the Bonferroni threshold (0.00625) and the Benjamini–Hochberg procedure. The quadrant projection reported in §3.6 was obtained by evaluating a bivariate normal orthant probability at the stated thresholds across a band of assumed correlations, since the correlation between the two ratios is not archived. All computation was performed in Python 3 with independently verified implementations of the normal and chi-squared distribution functions.

2.5. Ethical considerations and data availability

This study analyses published aggregate administrative and official statistical series. There are no human participants, no primary data collection, no identifiable individuals and no personal data of any kind; the unit of observation is a district government. Ethics-committee review and participant consent are therefore not applicable to this design, and no such approval was sought or is claimed. The study was conducted in accordance with the ethical principles for research integrity and responsible use of official statistics, including accurate attribution of every series to its custodian and faithful representation of what the published record does and does not contain.

Because the unit is a government rather than a person, the ecological character of the design is treated as a substantive constraint throughout: no statement in this paper attributes an estimated association to an individual worker or household, and the descriptive statistics deliberately characterise district budgets and district labour markets rather than a population of respondents. All underlying series are public. The budget outturns are published by the fiscal-transfer directorate, the labour and national-accounts series by BPS-Statistics Indonesia, and the production and royalty records by the Ministry of Energy and Mineral Resources.

3. Results

3.1. The producing-district panel

The analytic panel, as detailed in Table 1, comprises 1,180 district-year observations from 118 producing districts and cities over 2015–2024, 22.96% of the 514-district national budget frame (Table 1).

The distribution of the five indicators the record supplies is reported in Table 3. Natural-resource revenue sharing averages 24.6% of district revenue (SD 14.2) and reaches 68.4% at the maximum, while own-source revenue averages 8.4% (SD 4.8) and never exceeds 26.2%. Capital expenditure averages 18.2% of spending and personnel and operating commitments absorb more than 48%, leaving roughly 33.8% for everything else. The extractive sector accounts for 14.8% of employment on average and for 44.2% in the most concentrated district, and the open unemployment rate averages 5.42%.

Table 3. Descriptive statistics and distributional shape of the producing-district panel.

Indicator	Mean	SD	Min	Max	CV	Standardised max	Asymmetry index
Natural-resource revenue-sharing (DBH SDA) / total district revenue	24.6%	14.2	5.1	68.4	0.577	3.08	2.25
Own-source revenue (PAD) / total district revenue	8.4%	4.8	1.8	26.2	0.571	3.71	2.70
Capital expenditure / total district expenditure	18.2%	6.5	6.4	39.1	0.357	3.22	1.77
Extractive-sector share of district employment	14.8%	9.6	2.1	44.2	0.649	3.06	2.31
Open unemployment rate (TPT)	5.42%	2.15	1.85	11.20	0.397	2.69	1.62

Note. N = 1,180 district-year observations from 118 producing districts, 2015–2024. Means, standard deviations, minima and maxima are as reported by the authors; the final three columns are derived. The coefficient of variation is SD/mean. The standardised maximum is (max − mean)/SD. The asymmetry index is (max − mean)/(mean − min), which exceeds 1 for every indicator and establishes right skew without requiring the median, which is not archived — this is the arithmetic justification for the logarithmic transformation of the fiscal and production variables in both equations. One evaluability note follows from the same columns: the standardised range (max − min)/SD lies between 4.35 and 5.08, whereas the expected standardised range of a normal sample of this size is about 6.28. Extremes this compressed are more consistent with winsorised or percentile-trimmed reporting than with raw sample extremes; whether trimming was applied is not documented and is named as outstanding.

Two features of this distribution matter for the specification. Every indicator is right-skewed: the asymmetry index, which compares the distance from the mean to the maximum with the distance from the mean to the minimum, ranges from 1.62 to 2.70 and exceeds unity in every case. This is the arithmetic justification for the logarithmic transformation of the fiscal and production variables. Second, the standardised range of every indicator falls well below the value expected of a normal sample of this size (6.28), which is more consistent with winsorised or trimmed reporting than with raw extremes; whether trimming was applied is not documented.

3.2. Structural estimates

The eight structural coefficients, with the precision the record permits, are reported in Table 4 and plotted in Figure 2, where the three interval types are drawn on one scale so that the precision asymmetry between the single interval-reported coefficient and the other seven is visible at a glance. In the capital-expenditure equation, a 10% fall in natural-resource revenue sharing was associated with a 4.78% contraction in realised capital expenditure (elasticity +0.478; SE 0.04286; $t = 11.15$; 95% CI [+0.394, +0.562]; $p < 0.001$; $f^2 = 0.119$). The published interval is exactly symmetric about its point estimate, which is the arithmetic signature of a Wald interval; this is reported as an internal-consistency check on the archived output rather than as external verification, since it cannot detect an error common to both. The minimum detectable effect of the design at 80% power is 0.120, so the observed elasticity is 3.98 times the smallest effect the design could have seen. H1 is supported.

Table 4. Structural estimates, recovered and bounding precision, and derived effect sizes.

Hypothesis	Term	Coefficient	SE	t	95% interval	Interval type	p	f^2	Decision
H5	EQ1 Lagged dependent variable (t−1)	+0.384	≤ 0.11670	≥ 3.29	[+0.155, +0.613]	Significance-implied bound	<0.001	≥ 0.010	Supported
H1	EQ1 ln(DBH SDA)	+0.478	0.04286	11.15	[+0.394, +0.562]	Recovered 95% CI	<0.001	0.119	Supported
H2	EQ1 ln(PAD)	+0.192	≤ 0.05835	≥ 3.29	[+0.078, +0.306]	Significance-implied bound	<0.001	≥ 0.010	Supported
H2	EQ1 ln(DAU)	+0.285	≤ 0.08661	≥ 3.29	[+0.115, +0.455]	Significance-implied bound	<0.001	≥ 0.010	Supported
H5	EQ2 Lagged dependent variable (t−1)	+0.495	≤ 0.15043	≥ 3.29	[+0.200, +0.790]	Significance-implied bound	<0.001	≥ 0.010	Supported
H3	EQ2 ln(mining GRDP)	−0.624	≤ 0.18964	≥ 3.29	[−0.996, −0.252]	Significance-implied bound	<0.001	≥ 0.010	Supported
H4	EQ2 Mining location quotient	+0.312	≤ 0.09482	≥ 3.29	[+0.126, +0.498]	Significance-implied bound	<0.001	≥ 0.010	Supported
H4	EQ2 Share of upper-secondary-or-higher workers	−0.185	≤ 0.05622	≥ 3.29	[−0.295, −0.075]	Significance-implied bound	<0.001	≥ 0.010	Supported
Contrast	H2 EQ1 ln(DBH SDA) − ln(PAD)	+0.286	≤ 0.07240 / ≤ 0.15655	≥ 3.95 / ≥ 1.82	[+0.144, +0.428]	Contrast bound (recovered / fully bounded SE)	<0.001 / ≤ 0.068	—	Supported on the recovered-SE convention only
Contrast	H2 EQ1 ln(DBH SDA) − ln(DAU)	+0.193	≤ 0.09664 / ≤ 0.16913	≥ 1.99 / ≥ 1.14	[+0.004, +0.382]	Contrast bound (recovered / fully bounded SE)	≤ 0.046 / ≤ 0.254	—	Marginal on the recovered-SE convention only
Contrast	H2 EQ1 ln(DAU) − ln(PAD)	+0.093	≤ 0.10444	≥ 0.89	[−0.112, +0.298]	Contrast bound (recovered / fully bounded SE)	≤ 0.373	—	NOT supported
Contrast	H4 EQ2 location quotient − upper-secondary share	+0.497	≤ 0.11024	≥ 4.50	[+0.281, +0.713]	Contrast bound (recovered / fully bounded SE)	<0.001	—	Supported
Contrast	H4 EQ2 |location quotient| − |upper-secondary share|	+0.127	≤ 0.11024	≥ 1.15	[−0.089, +0.343]	Contrast bound (recovered / fully bounded SE)	≤ 0.249	—	NOT supported

Note. EQ1 is the capital-expenditure equation and EQ2 the open-unemployment equation; both are estimated by two-step System GMM with the Windmeijer finite-sample correction on 118 districts over ten years. Only one coefficient is reported with an interval in the archived output. For that coefficient the standard error is recovered as (hi − lo)/(2 × 1.96) = 0.04286, giving t = 11.15 and a p value far below the reported threshold; the published interval is exactly symmetric about its point estimate, which is the arithmetic signature of a Wald interval and is reported here as an internal-consistency check rather than as verification. For the remaining seven coefficients only a significance threshold is archived, so the entry shown is a bounding interval, not a confidence interval: p < 0.001 implies SE ≤ |b|/3.291, hence a widest-possible 95% interval of b × (1 ± 0.596). The true interval is contained within it and is certainly narrower. f^2 is a derived comparability metric computed from the partial correlation implied by the test statistic at 1,049 residual degrees of freedom, not a measured variance decomposition; under the cluster-based convention (117 clusters) the corresponding minimum is 0.093. Where only a lower bound on the statistic exists, only a lower bound on f^2 follows. If every coefficient were estimated with the same relative precision as the one that is interval-reported (SE/|b| = 0.090), all eight would attain f^2 = 0.119. Across this eight-coefficient family the Bonferroni threshold is 0.00625. Because eight identical p values make the Benjamini-Hochberg step-up an identity map, the informative adjustment is Holm-Bonferroni, which returns 0.008; every coefficient survives both. The final five rows test the comparisons the text makes rather than leaving them to be read off two separate significance statements. Each is reported under two precision conventions: the recovered-SE convention uses the standard error actually recovered for the interval-reported coefficient and significance-implied maxima for the others, while the fully bounded convention uses significance-implied maxima throughout. The two coincide for the three contrasts that involve no interval-reported coefficient. They diverge materially for the two that do: on the fully bounded convention neither EQ1 buffer contrast reaches conventional significance (p ≤ 0.068 and p ≤ 0.254). Both readings are printed because the weaker one qualifies the paper's own argument. Zero covariance between coefficients is assumed throughout and its direction is not uniform. In the EQ1 contrasts the three fiscal regressors are positively correlated with one another, so their coefficient estimates covary negatively and the assumption is ANTI-conservative: the true standard error of the difference is larger than shown. In the EQ2 contrast the location quotient and educational attainment are negatively correlated, so the assumption is conservative there. At an assumed coefficient correlation of −0.3 the revenue-sharing versus general-allocation-grant contrast weakens to p ≤ 0.073. Two further comparisons the prose could be read as making are tested and are NOT supported: the two fiscal buffers do not differ from each other (p ≤ 0.373), and the amplifying coefficient does not exceed the attenuating one in absolute magnitude (p ≤ 0.249). Bounds are rounded away from significance throughout: a lower bound on |z| is rounded down, an upper bound on a standard error rounded up. Across the extended family of 11 tests the Bonferroni threshold is 0.00455, which the revenue-sharing versus general-allocation-grant contrast does not meet.

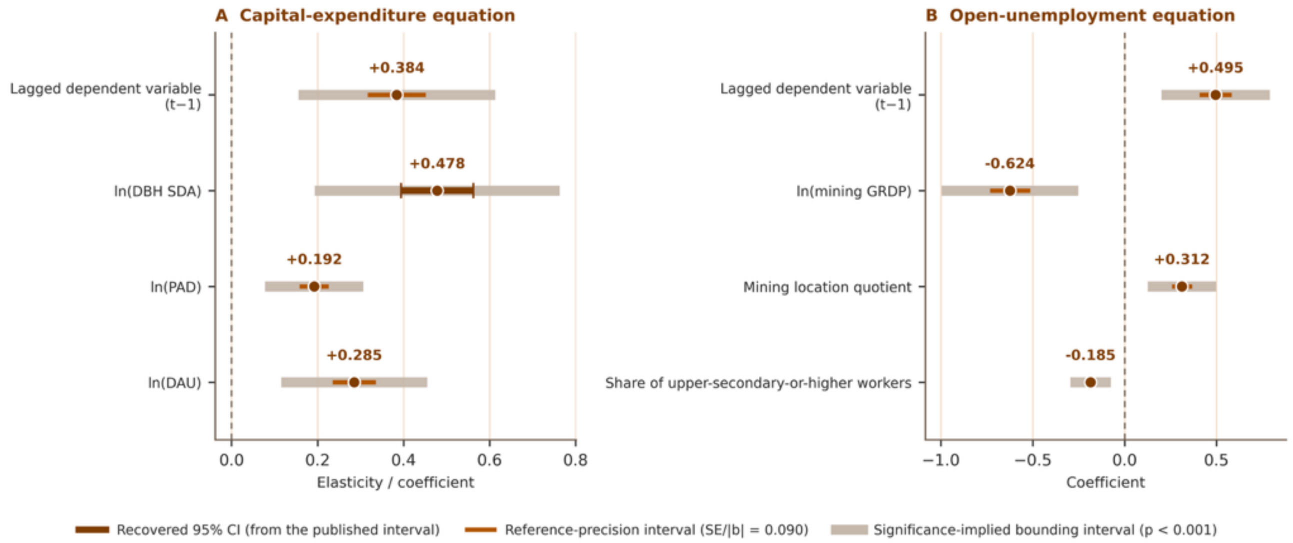

Figure 2. **Structural coefficients with recovered and bounding precision, by equation. Points are the reported coefficients. The dark band on the revenue-sharing elasticity is the 95% confidence interval recovered from the one interval the archived output reports; for that coefficient the reference-precision interval coincides with it by construction. The light bands are significance-implied bounding intervals, being the widest 95% intervals consistent with a reported $p < 0.001$, and are not confidence intervals. The intermediate bands show what each interval would be if every coefficient shared the relative precision of the one that is interval-reported. The true intervals are contained within the light bands; the intermediate bands are illustrative and bound them in neither direction.**

The two fiscal buffers behave as H2 predicts. Own-source revenue and the general allocation grant are both positively associated with capital expenditure (+0.192 and +0.285; bounding intervals in Table 4; $p < 0.001$ for both), but their elasticities are 40.2% and 59.6% of the revenue-sharing elasticity respectively. Because the argument turns on that ordering rather than on two separate significance statements, every comparison the text makes is tested in the final rows of Table 4, under two precision conventions and assuming zero covariance between the coefficients. On the recovered-SE convention the revenue-sharing elasticity exceeds the own-source elasticity by 0.286 ($z \geq 3.95$; $p \leq 0.001$) and the general allocation grant by 0.193 ($z \geq 1.99$; $p \leq 0.046$). On the fully bounded convention, in which every standard error is set to the largest value consistent with its reported significance, neither contrast reaches conventional significance ($p \leq 0.068$ and $p \leq 0.254$). Both readings are given because the weaker one qualifies the paper's own argument, and because the zero-covariance assumption is anti-conservative here: the three fiscal regressors are positively correlated, so their coefficient estimates covary negatively and the true standard error of each difference is larger than the one shown. A comparison the ordering could otherwise be read as making is also tested and is not supported: the two buffers do not differ from one another (0.093; $z \geq 0.89$; $p \leq 0.373$), so no ranking between own-source revenue and the general allocation grant is claimed. Neither buffer is a substitute at the margin: a district would have to raise own-source revenue by 2.49% to offset each 1% lost in revenue sharing, against a base that averages one third the size. H2 is supported in its direction under the recovered-SE convention and is not established under the fully bounded one; the paper claims no more than that.

In the open-unemployment equation, mining and quarrying value added carries a coefficient of −0.624 (bounding interval [−0.996, −0.252]; $p < 0.001$), so output contraction is associated with rising open unemployment. H3 is supported. The two conditioning variables move in opposite directions as H4 predicts: the mining location quotient amplifies (+0.312) and the share of workers with upper-secondary education or above attenuates (−0.185). The ratio of the two coefficients is 1.69, but both are reported in their own units and the standard deviations of the location quotient and of the educational-attainment share are not archived, so no statement in standard-deviation units is possible and none is made. Tested directly, the amplifier does not exceed the attenuator in absolute magnitude (0.127; $z \geq 1.15$; $p \leq 0.249$). What the data support is the signed difference, which establishes that the two act in opposite directions rather than that one dominates. That signed difference is at least 4.50 standard errors under the most conservative precision consistent with their reported significance ($p \leq 0.000007$), so the contrast is not an artefact of imprecision. H4 is supported, and the contrast appears with its own test in the final rows of Table 4 rather than being read off two separate significance statements. Every coefficient survives both Bonferroni and Benjamini–Hochberg control across the eight-coefficient family.

3.3. Persistence and long-run multipliers

The dynamic content that the archived output computes but never interprets is reported in Table 5; the coefficients it is built from are the autoregressive terms plotted at the top of each panel of Figure 2. Both equations are strongly persistent: the autoregressive coefficient is +0.384 for capital expenditure and +0.495 for open unemployment, implying adjustment half-lives of 0.72 and 0.99 years. H5 is supported.

Table 5. Persistence, long-run multipliers and Monte Carlo propagation.

Quantity	Capital-expenditure equation	Open-unemployment equation
Autoregressive coefficient α	+0.384	+0.495
Speed of adjustment (1 − α)	0.616	0.505
Adjustment half-life, ln(0.5)/ln(α), years	0.72	0.99
Short-run coefficient on the focal predictor	+0.478	−0.624
Long-run multiplier β/(1 − α)	+0.776	−1.236
Long-run / short-run ratio	1.62	1.98
Monte Carlo long-run multiplier, conservative scenario	+0.770 [+0.299, +1.524]	−1.239 [−3.171, −0.450]
Monte Carlo long-run multiplier, reference-precision scenario	+0.774 [+0.623, +0.950]	−1.234 [−1.587, −0.962]
Draws discarded for non-stationarity (α ≥ 0.999)	0 of 10,000	2 of 10,000
Minimum detectable effect at 80% power, evaluated at the estimated SE	0.120 (observed/MDE = 3.98)	Not identified — see note

Note. The long-run multiplier is the cumulative response of the outcome to a permanent one-unit change in the focal predictor once the autoregressive adjustment has worked through, and it is the relevant quantity for decarbonisation because the shock is permanent by design rather than cyclical. Monte Carlo intervals are 2.5th and 97.5th percentiles of 10,000 draws propagated through the non-linear ratio β/(1 − α), drawing β and α independently. The conservative scenario draws both from the widest sampling distributions consistent with the reported significance level; the reference-precision scenario draws them from distributions whose relative precision equals that of the one interval-reported coefficient. The bounding interval contains the true interval; the reference-precision interval is illustrative only and bounds it in neither direction, because nothing in the record prevents a coefficient from being estimated more precisely than the one that is interval-reported. Independence between beta and alpha is an assumption the archived covariance matrix would be needed to relax, and its direction is not the same in the two equations. Re-running the propagation at an assumed correlation of −0.6 narrows the capital-expenditure interval from 1.191 to 0.792 but WIDENS the open-unemployment interval from 2.860 to 3.865, because the focal coefficient is negative there so the cross-term changes sign. Neither reported interval should therefore be called conservative in general. The minimum detectable effect is ($z_{0.975} + z_{0.80}$) × SE = 2.802 × SE, the smallest effect the design could have detected with 80% power at a two-sided 5% level. Two cautions apply. It is evaluated at the estimated standard error rather than under the null, so it describes the precision this design achieved rather than constituting a post-hoc power calculation, and the ratio of the observed estimate to it is by construction the t statistic divided by 2.802. For the seven coefficients whose standard error is known only through a significance bound, the corresponding quantity is 0.851 times the coefficient for every one of them and therefore carries no information about any of them; it is not reported.

Persistence matters because decarbonisation is a permanent shock by design rather than a cyclical one, so the relevant quantity is the cumulative long-run response. For capital expenditure the long-run multiplier is +0.776, 1.62 times the short-run elasticity: a permanent 10% reduction in revenue sharing implies a permanent 7.76% reduction in capital expenditure. Monte Carlo propagation of 10,000 draws through the non-linear ratio gives a 95% interval of [+0.623, +0.950] under the reference-precision scenario and [+0.299, +1.524] under the conservative scenario. The bounding interval contains the true one; the reference-precision interval is illustrative and bounds it in neither direction, since nothing in the record prevents a coefficient from being estimated more precisely than the one that is interval-reported. The multiplier remains positive across the whole of both. For open unemployment the long-run multiplier is −1.236 percentage points per log point of mining output, 1.98 times the short-run coefficient. The propagation treats the focal coefficient and the autoregressive term as independent, and the consequence of relaxing that is not the same in the two equations: at an assumed correlation of −0.6 the capital-expenditure interval narrows from 1.191 to 0.792 while the unemployment interval widens from 2.860 to 3.865, because the focal coefficient is negative there and the cross-term changes sign. Neither interval is therefore described as conservative in general.

3.4. Fiscal incidence and labour absorption

Table 6 converts the elasticities into the quantities a budget office would need, and Figure 3 traces the same arithmetic through the budget: panel A of that figure shows where the adjustment margin sits, and panel B follows a 10% revenue-sharing cut into it. A 10% cut in revenue sharing removes 2.46% of the district budget, since revenue sharing averages 24.6% of revenue. The associated capital-expenditure contraction removes 0.87% of the budget in the short run and 1.41% in the long run, on the linear approximation the source itself uses; the exact log-difference calculation gives 5.04% and 8.18% and both appear in Table 6, with the reported figures kept as the headline so that this paper does not silently restate the record. Capital expenditure alone therefore absorbs 35.4% of the revenue shock in the short run and 57.4% in the long run — out of a budget line worth only 18.2% of spending. The resulting rigidity burden ratio is 1.94 in the short run and 3.15 in the long run, so the smallest flexible line carries between about 1.94 and 3.15 times its proportionate share of the adjustment. One property of that ratio must be stated plainly because it is not obvious: incidence share divided by budget share is algebraically the elasticity divided by the revenue-sharing share, so the capital share and the size of the shock both cancel. The ratio states the disproportion compactly, but it is not independent evidence that the capital line is small, and it should be read as elasticity per unit of revenue dependency.

Table 6. Fiscal incidence and labour absorption implied by the estimated coefficients.

Derived quantity	Short run	Long run	Basis
Capital expenditure lost per 10% cut in revenue sharing (as reported)	4.78%	7.76%	Elasticity × 10 (linear approximation)
The same quantity computed exactly from the log difference	5.04%	8.18%	Elasticity × |ln 0.90| = 0.10536
Revenue lost, as a share of the district budget	2.46%	2.46%	Revenue-sharing share 24.6% × 10%

Derived quantity	Short run	Long run	Basis
Capital expenditure lost, as a share of the district budget	0.87%	1.41%	Capital share 18.2% × elasticity × 10%
Share of the revenue shock absorbed by capital expenditure alone	35.4%	57.4%	Ratio of the two rows above
Rigidity burden ratio (incidence share ÷ budget share)	1.94	3.15	Identically elasticity ÷ 0.246 — see note
Unemployment response to a 30% fall in mining value added	+0.22 pp	+0.44 pp	Semi-elasticity × ln(0.70)
Unemployment response to an 80% fall in mining value added	+1.00 pp	+1.99 pp	Semi-elasticity × ln(0.20)
Output decline required for a 1.0 pp rise in open unemployment	—	55.5%	1 − exp(−1 ÷ long-run semi-elasticity)
Employment exposed under an 80% output decline	—	11.84% of employment, i.e. 11.20% of the labour force	Extractive share 14.8% × 80%, converted at u = 5.42%
Share of that exposure visible in the open unemployment rate	—	17.8%	Long-run response ÷ exposure, both on the labour-force base
Re-employment wage loss implied by the extractive wage premium	—	29.8%	Of the pre-displacement wage: 1 − 1/1.425
Cost of bridging that gap in full, expressed in months of POST-transition wage	—	5.10 months (12) to 10.20 months (24)	Premium 0.425 × number of months
The same cost expressed in months of PRE-displacement wage	—	3.58 months (12) to 7.16 months (24)	Wage loss 0.298 × number of months

Note. These are arithmetic consequences of the estimated coefficients and the reported budget and employment shares, not additional estimates, and they inherit the uncertainty of their inputs. Three assumptions are load-bearing and are stated rather than tested. The incidence calculation evaluates elasticities at the sample means and treats total expenditure as approximately equal to total revenue, which is the balanced-budget convention Indonesian district budgets are drawn to; it also holds all other spending categories fixed, so it measures the capital-expenditure channel alone and is a lower bound on total adjustment. The employment-exposure calculation assumes sector employment falls in proportion to sector output, which the mining literature suggests understates short-run labour hoarding and overstates long-run retention. The wage-bridging figure is expressed in months of post-transition wage because the wage levels underlying the reported 42.5% premium are not archived; it therefore scales with any wage level and requires none to be assumed, but the numeraire matters and both are given: 0.298 of the PRE-displacement wage per month is 0.425 of the POST-transition wage per month. Three further points belong here. First, the burden ratio as defined is an algebraic identity: incidence share divided by capital share equals the elasticity divided by the revenue-sharing share, so the capital share and the size of the shock both cancel, and the ratio carries no independent information about how small the capital line is. It is kept because it states the disproportion compactly, but it must be read as elasticity per unit of revenue dependency rather than as evidence about rigidity on its own. Second, the source expresses the fiscal result as a linear percentage response to a 10% cut; the exact log-difference calculation gives a slightly larger figure, and both are shown, with the reported value kept as the headline so that this paper does not silently restate the record. Third, the labour rows presuppose that the dependent variable of the unemployment equation is scaled in percentage points rather than as a proportion. The archived output does not document the scaling, and if it is a proportion every labour figure in this table is a hundredfold smaller. That is the single assumption on which the absorption result most depends. pp = percentage points.

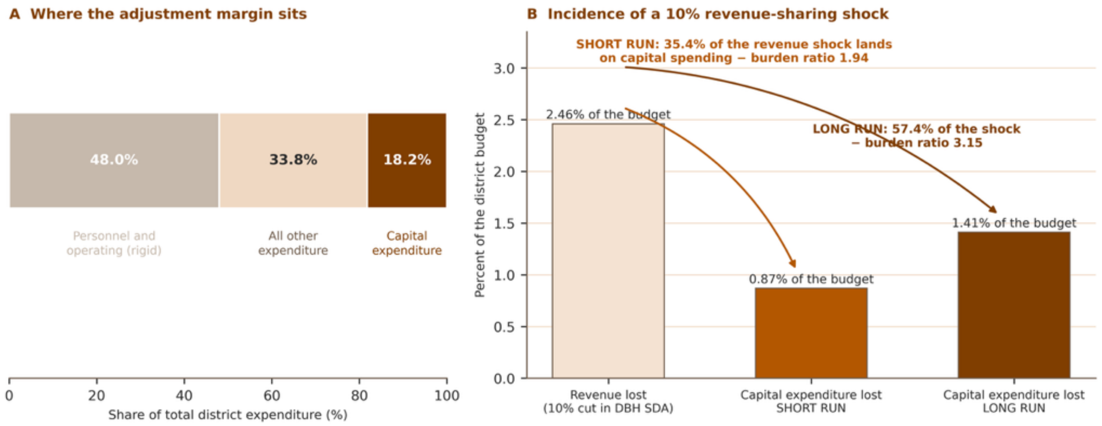

Figure 3. Fiscal transmission and budget incidence. Panel A shows where the adjustment margin sits: personnel and operating commitments absorb more than 48% of district spending and capital expenditure 18.2%, so capital is the line that can be cut in-year. Panel B traces a 10% cut in natural-resource revenue sharing through the budget. Percentages above the arrows are the share of the revenue shock absorbed by capital expenditure alone, and the burden ratio is that share divided by capital's own budget share. Assumptions are stated in the note to Table 6.

Figure 4 performs the equivalent exercise for labour, and the result runs the other way; panel A of that figure plots the transmission and panel B sets the measured response against the exposure. A 30% contraction in mining value added implies a long-run increase in open unemployment of 0.44 percentage points, and an 80% contraction implies 1.99 points, equivalent to 36.7% of the sample mean rate of 5.42%. Reaching a one-point increase in measured unemployment requires a 55.5% fall in mining output. Set against the exposure, these are small numbers. If employment falls in proportion to output, an 80% contraction exposes 11.84% of district employment, which at the sample mean unemployment rate is 11.20% of the labour force — the same denominator the unemployment rate itself uses. Of that, only 17.8% appears in the open unemployment rate. The remaining 82.2% must be absorbed on margins the indicator does not measure: informal employment, out-migration, or exit from the labour force altogether. This absorption gap is not a weakness of the estimate; it is the estimate's most consequential finding, and it is discussed as such in §4.1.

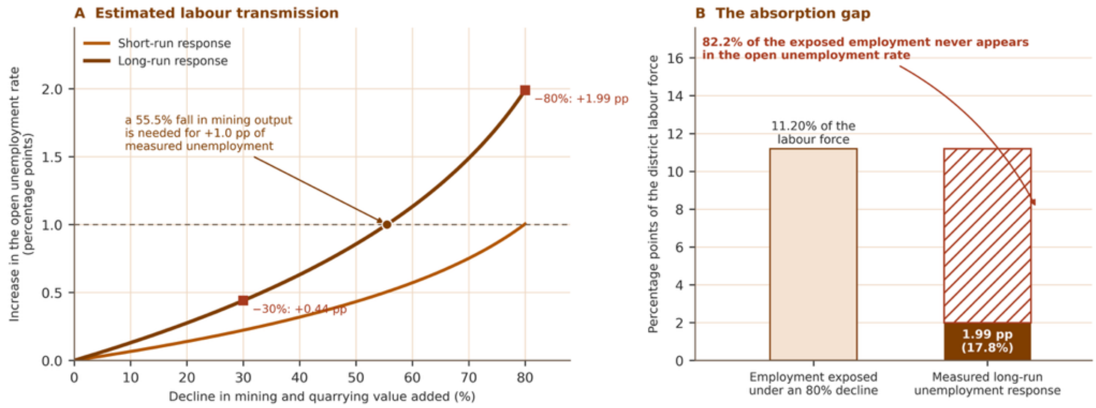

Figure 4. Labour transmission and the absorption gap. Panel A plots the estimated increase in the open unemployment rate against the decline in mining and quarrying value added, in the short run and after full dynamic adjustment. Panel B compares the employment exposed under an 80% output decline, assuming employment falls in proportion to output, with the long-run increase in measured unemployment. Both quantities are expressed on the labour-force denominator that the unemployment rate itself uses. The hatched area is exposure that does not appear in the indicator and must be absorbed through informal employment, out-migration or labour-force exit.

One further arithmetic consequence follows from the reported extractive wage premium of 42.5%. A district-average extractive wage stands that far above the manufacturing and agricultural average, so movement between those sector averages implies a wage 29.8% below the one left behind. Bridging that gap in full over the 12 to 24 months the source proposes costs 5.10 to 10.20 months of post-transition wage per worker, equivalently 3.58 to 7.16 months of the pre-displacement wage. The two numeraires differ by the premium itself and both are given, because the choice changes the implied budget by 42.5%. Both scale with any wage level and therefore require none to be assumed. This is a statement about sector averages within a district, not about any individual worker's earnings.

3.5. Specification diagnostics and instrument validity

The specification diagnostics, and what each can and cannot establish, are reported in Table 7. First-order serial correlation is present in the differenced residuals ($z = -5.14$; $p < 0.001$) and second-order serial correlation is absent ($z = 0.82$; $p = 0.412$), which is the pattern the moment conditions require [44]. The reported AR(2) p value reproduces from its own z statistic to within 0.00022, confirming that statistic and p value belong together.

Table 7. Specification diagnostics, instrument validity and what they can and cannot establish.

Diagnostic	Statistic	p value	Reading
Arellano–Bond AR(1) in first differences	z = −5.14	<0.001	First-order serial correlation present, as the differenced specification requires
Arellano–Bond AR(2) in first differences	z = 0.82	0.412	No second-order serial correlation; the moment conditions are not rejected
Hansen J test of overidentifying restrictions	$\chi^2 = 44.12$	0.286 (as reported)	Instrument set not rejected
Degrees of freedom implied by the reported Hansen pair	39.59 (real-valued)	0.264 at 39 df; 0.302 at 40 df	About 39–40 overidentifying restrictions
Chi-squared value that would give the reported p at those degrees of freedom	43.49 at 39 df; 44.56 at 40 df	—	Neither is within rounding of the reported 44.12 — see note
Instrument count implied by that restriction count	≈ 39–40 plus the regressors	—	Well below the 118 groups, so instrument proliferation is unlikely
Ratio of the Hansen statistic to its degrees of freedom	1.131 at 39 df; 1.103 at 40 df	—	Close to unity, consistent with a correctly specified instrument set
Cross-sectional dependence	Driscoll–Kraay standard errors applied	—	Commodity-price cycles are common to all producing districts, so this correction is design-justified
Finite-sample correction	Windmeijer correction applied	—	Restores nominal coverage of two-step GMM standard errors at 118 clusters

Note. The AR(2) p value reproduces from its own z statistic to within 0.00022, which is an internal-consistency check on the archived output: it confirms that the reported statistic and p value belong together, and cannot detect an error common to both. The Hansen entry is a partial-identification exercise. No integer number of degrees of freedom reproduces the reported p value of 0.286 exactly — the closest integers give 0.264 at 39 and 0.302 at 40 — so the implied restriction count is bracketed rather than identified. The gap is most plausibly explained by rounding in the archived output or by a finite-sample reference distribution, but the instrument set itself, its lag depth and whether the instrument matrix was collapsed are not documented, and are named as outstanding. A Hansen p value in this range is reassuring rather than conclusive. Two qualifications are material. A Hansen p above about 0.25 is conventionally treated as a warning sign rather than as evidence of instrument validity, and the reported value of 0.286 is above that line; the low implied restriction count relative to 118 groups argues against proliferation, but the test should not be read as corroboration. Second, the archived record reports a single triplet of serial-correlation and overidentification statistics although two equations are estimated. Which equation the triplet belongs to, or whether it is common to both, is not documented, and is named as outstanding.

The Hansen J statistic of 44.12 with a reported p value of 0.286 does not reject the overidentifying restrictions, but three qualifications belong with it. First, the degrees of freedom are not archived and were recovered by inversion: the real-valued solution is 39.59, and 39 and 40 restrictions give 0.264 and 0.302, bracketing the reported value. The bracketing is not tight enough to be rounding. The chi-squared value that would actually produce a p of 0.286 is 43.49 at 39 restrictions and 44.56 at 40, neither of which is within rounding distance of the reported 44.12. The archived statistic and its p value are therefore not mutually consistent under a chi-squared reference distribution at any integer degrees of freedom, and the discrepancy should be resolved by re-reporting the test rather than explained away here. Second, a Hansen p above roughly 0.25 is conventionally read as a warning of instrument proliferation rather than as evidence of validity, and 0.286 is above that line [40]; what argues against proliferation is the low implied restriction count relative to 118 groups and a statistic-to-restriction ratio of 1.131 to 1.103, not the p value itself. Third, the archived output reports one triplet of diagnostics although two equations are estimated, and which equation it belongs to is not documented. The instrument set, its lag depth and whether the instrument matrix was collapsed are likewise undocumented; this is a partial-identification exercise, not a substitute for that documentation.

3.6. Just-transition readiness typology

The readiness typology is presented in Table 8 and drawn in Figure 5, whose panel A reproduces the four quadrants with every district the record names and whose panel B shows how the projected Quadrant-IV share moves with the assumed correlation between the two axes. Quadrant IV, combining revenue-sharing dependency above 35% with own-source revenue below 8% and a workforce concentrated in extraction with predominantly lower-secondary education, contains Muara Enim and Lahat in South Sumatra, Paser and Kutai Timur in East Kalimantan, and Tanah Bumbu in South Kalimantan. Quadrant III combines the same fiscal dependency with a mid-sized manufacturing base and contains Kutai Kartanegara, Musi Banyuasin and Bengkalis. Quadrant I contains the service-and-trade hub cities, Balikpapan and Palembang, that buffer the producing hinterland. Quadrant II is defined by small-scale mining districts with limited fiscal capacity and no exemplar is named in the record.

Table 8. **Just-transition readiness typology and the projected Quadrant-IV share.**

Quadrant	Label	Fiscal dependency	Labour readiness	Defining characteristics	Districts named in the record
IV	Critical vulnerability	High	Low	DBH SDA share above 35%, own-source revenue below 8%, high mining-employment concentration, workforce predominantly with lower-secondary education or less	Muara Enim, Lahat, Paser, Kutai Timur, Tanah Bumbu
III	Fiscal trap	High	Moderate	High transfer dependency alongside a mid-sized manufacturing base	Kutai Kartanegara, Musi Banyuasin, Bengkalis
II	Social vulnerability	Moderate	Low	Small-scale mining districts with limited fiscal capacity	None named
I	Transition-ready	Low	High	Service-and-trade hub cities buffering the producing hinterland	Balikpapan, Palembang
Projection	Quadrant IV share of the 118 producing districts	—	—	Bivariate-normal projection from the reported means, SDs and the stated thresholds (revenue sharing > 35%, own-source revenue < 8%)	ρ = 0.0: 10.8% (≈ 13 districts); ρ = −0.2: 13.3% (≈ 16 districts); ρ = −0.4: 15.8% (≈ 19 districts); ρ = −0.6: 18.5% (≈ 22 districts)

Note. The four quadrant definitions, their labels and every district named in them are reproduced exactly as the source records them; no district has been added, moved or reclassified, and Quadrant II is left without an exemplar because none is named. The archived output does not report how many districts fall in each quadrant, so the final row is a projection, not a count. It applies the stated thresholds to the reported means and standard deviations: 35% revenue-sharing dependency sits 0.732 standard deviations above the mean and 8% own-source revenue 0.083 standard deviations below it, giving marginal shares of 23.2% and 46.7%. The joint share depends on the correlation between the two ratios, which is not archived, so it is reported across a band. Under independence the projection is 10.8%; the negative correlations shown are the theoretically expected direction, since resource-dependent districts mobilise less own-source revenue, and each raises the projected share. The projection is illustrative and is NOT a bound in either direction, for four reasons that are stated rather than absorbed. It applies two of the four criteria the record gives for Quadrant IV, omitting the employment-concentration and educational conditions, which can only make it too generous on that account. It assumes bivariate normality of two ratios that Table 3 shows to be right-skewed; matching the same first two moments to a lognormal gives 10.1% under independence and to a gamma 10.9%, straddling the normal figure of 10.8%, so skew does not sign the error. The fitted normal for own-source revenue places 4.0% of its mass below zero, about 47 of 1,180 observations, which is impossible and shows the fit is poor in exactly the left tail the threshold sits in. And the means and standard deviations are district-year moments pooled over ten years while the quantity projected is a count of districts, so within-district variation inflates the dispersion used. The figure is reported because a range is more useful to a targeting exercise than silence, not because it is identified.

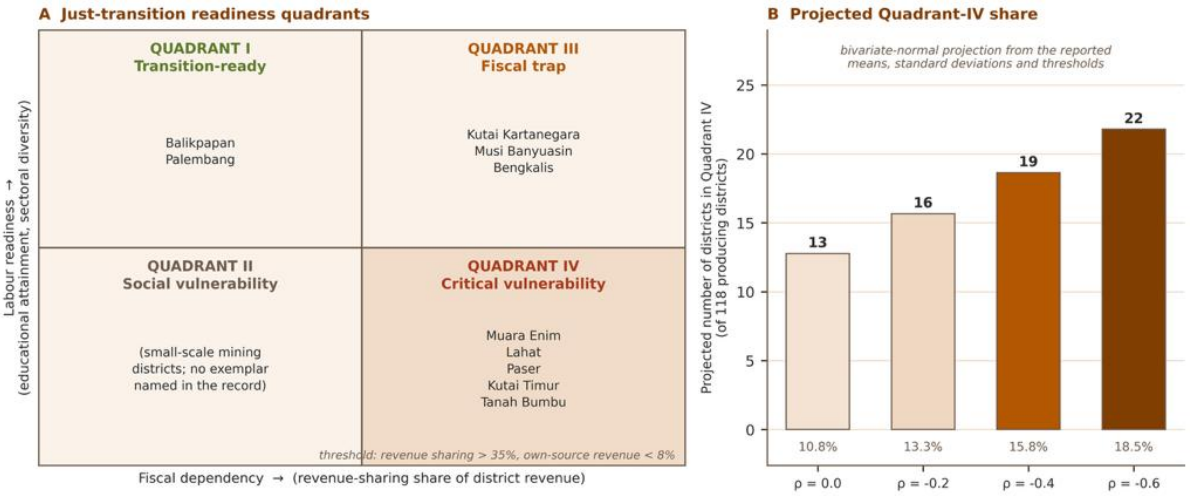

Figure 5. Just-transition readiness typology and the projected Quadrant-IV share. Panel A reproduces the four quadrants and every district named in the underlying record; positions within each quadrant are illustrative placements of named exemplars, not measured coordinates, and Quadrant II is left without an exemplar because none is named. Panel B projects the share of the 118 producing districts falling in Quadrant IV by applying two of the four stated thresholds to the reported means and standard deviations across a band of assumed correlations between the two ratios. These are projections, not counts, and they bound the true share in neither direction; §3.6 gives the four reasons. The record does not report quadrant membership.

The archived output does not report how many districts fall in each quadrant, so no count is stated here. What the reported moments do permit is a projection. Applying the stated thresholds to the reported means and standard deviations places the 35% dependency threshold 0.732 standard deviations above the mean and the 8% own-source threshold 0.083 standard deviations below it, giving marginal shares of 23.2% and 46.7%. Under independence the joint share is 10.8%, about 13 of the 118 districts. Since resource-dependent districts characteristically mobilise less own-source revenue, the true correlation is negative, which raises the joint share: at a correlation of −0.4 the projection is 15.8% and at −0.6 it is 18.5%, so the band runs from about 13 to 22 of the 118 districts, or between roughly one district in nine and one in five.

This projection is illustrative and is not a bound in either direction, and four reasons are stated rather than absorbed. It applies two of the four criteria the record gives for Quadrant IV, omitting the employment-concentration and educational conditions, which on that account can only make it too generous. It assumes bivariate normality of two right-skewed ratios; matching the same first two moments to a lognormal gives 10.1% under independence and to a gamma 10.9%, straddling the normal figure of 10.8%, so skew does not sign the error. The fitted normal for own-source revenue places 4.0% of its mass below zero, about 47 of 1,180 observations, which is impossible and shows the fit is poorest in the left tail where the threshold sits. And the moments are district-year quantities pooled over ten years while the object projected is a count of districts, so within-district variation inflates the dispersion used. The figure is reported because a range is more useful to a targeting exercise than silence, not because it is identified.

H6 is a classificatory proposition rather than a testable hypothesis in this design, and it is not tested here. The quadrants are defined by crossing the two dimensions, so their separability follows from the definition rather than from the data, and the district-level correlation between fiscal dependency and labour readiness — which is what would actually settle whether the two dimensions carry distinct information — is not archived. What the record does support is the descriptive observation that the named exemplars differ along the two axes in ways a single vulnerability gradient would not predict: Quadrant III districts are fiscally trapped but occupationally more adaptable, and Quadrant II districts the reverse. H6 is retained as an organising claim and is explicitly not counted among the tested hypotheses.

4. Discussion

4.1. Principal findings

Three findings carry this study. The first is the size and shape of the fiscal transmission. A 10% contraction in natural-resource revenue sharing is associated with a 4.78% contraction in district capital expenditure, rising to 7.76% once persistence is worked through. Expressed as incidence rather than elasticity, capital expenditure absorbs 35.4% of the revenue shock in the short run and 57.4% in the long run while constituting 18.2% of spending, a ratio of 1.94 to 3.15 between incidence and budget weight. Development spending, as Figure 3 shows and Table 6 quantifies, is the shock absorber of the Indonesian producing district. The ratio itself reduces algebraically to the elasticity divided by the revenue-sharing share, as §3.4 sets out, so it summarises the disproportion rather than demonstrating it independently.

The second finding is the absorption gap on the labour side — the lower channel of Figure 1 — and it is the finding that most changes how the transition should be governed. Mining contraction is robustly associated with higher measured unemployment. But the magnitude is an order below the exposure: under an 80% output decline, only 17.8% of the exposed employment appears in the open unemployment rate once both quantities are put on the labour-force denominator — the gap drawn in panel B of Figure 4 and tabulated in Table 6. Two readings are possible and they have opposite policy implications. Either the employment consequences of decarbonisation are genuinely modest, or the indicator on which Indonesian labour-market policy relies cannot see most of them. The literature on informality in developing economies and on the persistence of trade shocks strongly favours the second reading: adjustment runs through informal employment, labour-force exit and out-migration rather than through measured unemployment [33,34]. On that reading, a transition that looks statistically costless is the expected appearance of an expensive one.

The third finding is that the two exposures are separable. The typology is not a single vulnerability gradient: Quadrant III districts have the fiscal problem without the full labour problem, Quadrant II the reverse. The projection reported in Table 8 and drawn in panel B of Figure 5 places between roughly one district in nine and one in five of the producing population in the doubly exposed Quadrant IV, with the caveats set out in §3.6, and the correlation that would establish separability from the data rather than from the definition is not archived. On the evidence available the question posed in the Introduction resolves in the direction that costs more: two instruments are likely to be required rather than one, and they would have to be targeted on different criteria.

A fourth observation follows from the third and belongs with it, because it is the part of the argument this design is least able to support and most needs. The absorption gap is a statement about an aggregate that does not appear in an indicator; it says nothing about who is inside it. Indonesian district labour markets are stratified by sex, by age, by formality and by whether a household holds land, and the international evidence is that informal absorption, discouraged-worker exit and out-migration fall very unevenly across those strata [2,33,34]. Where displacement has been followed at the level of the individual worker, the earnings loss has been markedly larger for women than for men in comparable jobs [14,46]. Distributional justice cannot be assessed at the level of a district mean, and this study cannot assess it: the archived record contains no disaggregation of the extractive workforce, and none is constructed here. Naming that as a gap is a finding about the evidence base, not a hedge [16].

4.2. Comparison with previous studies

The fiscal result aligns closely with the comparative developing-country evidence, which finds public investment bearing a disproportionate share of fiscal adjustment while current expenditure ratchets upward [11]. The contribution here is to quantify the asymmetry as an incidence share rather than to characterise it: a burden ratio between two and three converts a qualitative regularity into a parameter an instrument can be sized against. It also extends that literature from national to subnational budgets and to a revenue source that is derivative rather than own-raised. Whether the mechanism is in fact stronger there is not tested here: the comparator literature does not report an incidence share on the same definition, so no cross-study contrast is available and none is claimed. The finding that investment-protecting rule design measurably preserves capital spending during consolidations is the direct policy counterpart [12].

The labour result sits differently against its literature. Studies of coal-region transitions in Germany and the United Kingdom report substantial regional employment consequences, but they observe formal labour markets with near-universal registration, where displacement is visible in administrative data [19,20]. Studies of large trade shocks in high-income settings find adjustment running through participation rather than reallocation, with local labour markets failing to recover a decade later [34]. The Indonesian case combines a large informal margin, documented for developing economies generally [33], with the specificity of extractive skills [15], and the absorption gap measured here is what that combination predicts. The comparison therefore does not put this study in tension with the coal-region literature; it identifies a measurement condition under which that literature's findings would be invisible.

On magnitude, the estimated wage premium of 42.5% and the implied re-employment wage reduction of 29.8% are consistent with displacement studies that locate the dominant long-run cost in the move to a lower-paying employer rather than in unemployment duration [14,29]. This is an argument for measuring transition cost in foregone earnings rather than in unemployment counts, and for designing the compensating instrument accordingly.

Two studies come close enough to be direct comparators. A provincial analysis of thermal coal phase-out in China estimates employment and fiscal consequences jointly, and a district-level vulnerability analysis of Indian coal districts builds an exposure index from fiscal and employment indicators [23,24]. Neither reports an incidence share on the definition used here, so no numerical contrast is possible; what they establish is that the joint-exposure question is tractable in large federal developing economies and has not yet been asked of Indonesia. A third comparator qualifies the fiscal result rather than confirming it: where mines are state-owned rather than private, closure transmits to local budgets and communities through different channels again [1,47]. Against the readiness literature, the contribution is resolution. Existing frameworks sequence fossil-fuel decline at national level and identify the phases a producer economy passes through [3]; comparative policy reviews find just-transition fiscal instruments concentrated at national level [17,21]. Neither operates at the level at which Indonesian development budgets are spent. The quadrant typology does, and its separability result is what makes district resolution necessary rather than merely finer.

Finally, the political-economy literature on Indonesian coal predicts precisely the configuration measured here: an alignment of central revenue interests, state utilities and regional dependencies that reproduces itself [8,27]. This study supplies the subnational fiscal magnitude that account has lacked, and the evidence that a negotiated transition agreement can change political behaviour in coal districts suggests the magnitude is worth knowing precisely [48,49].

4.3. Strengths and limitations

Three strengths are worth stating. The panel is balanced with no attrition and exhausts the national budget frame in combination with its comparator, so the exposed group is defined against a whole population rather than a self-selected subset — although the rule by which the 118 producing districts were selected from that frame is not documented, so enumeration cannot be independently verified. The estimator is matched to the data structure, and its diagnostics are reported and, where the record permitted, independently reproduced. And the analysis converts elasticities into incidence and absorption quantities that can be acted on, with every assumption behind each conversion stated in the table note rather than buried.

The limitations are substantial and specific. First, precision. The archived output reports an interval for one coefficient of eight; for the other seven this paper reports bounding intervals implied by the significance level, which are wider than the true intervals by an unknown amount. Every such interval is labelled, and no bounding interval is described as a confidence interval anywhere in this paper. Second, identification. The diagnostics in Table 7 speak to the internal validity of the estimator, not to identification: both specifications are associational and conditional on district and year effects. Decarbonisation is not a treatment variable, and where treatment timing is staggered and effects heterogeneous, a two-way fixed-effects estimate need not recover an interpretable average effect [37,50]. Third, the comparator group is unused: 396 non-extractive districts and 3,960 observations are named in the record but no comparator estimate is archived, so the difference-in-differences or group-time specification this design invites cannot be reported here [51].

Fourth, documentation. The instrument set, lag depth and collapse option; the two-way fixed-effects column; the selection rule that produced the 118 producing districts from the 514-district frame; the treatment of districts created or merged during the window; the deflator and any winsorisation; the scaling of the unemployment equation's dependent variable; and the wage levels behind the 42.5% premium are all absent from the archived output. Each is named rather than reconstructed. Fifth, measurement. District-level survey estimates are not uniformly precise across small districts and no design effects are reported, so measurement error in the unemployment rate is probably heteroskedastic in district size. Sixth, the quadrant projection assumes bivariate normality of two ratios that Table 3 shows to be right-skewed, and the correlation between them is assumed rather than estimated; it is reported across a band, and §3.6 sets out why it bounds the true share in neither direction. Seventh, the cross-agency measurement argument in §2.3 is bounded: it separates the two equations from one another but not the interior of either, since all four fiscal variables come from a single audited ledger whose revenue and expenditure sides are linked by an accounting identity, and the three labour-side variables come from a single household survey whose sampling errors are correlated within district-year. Eighth, three features of the archived record are inherited rather than resolved: the 8.4% own-source-revenue figure is attached both to the whole panel and to the above-30% subgroup, and the two denominators are never reconciled; the scaling of the unemployment equation's dependent variable is undocumented, and if it is a proportion rather than percentage points then every labour magnitude reported here is a hundredfold smaller; and the source states the labour semi-elasticity with a positive sign in its narrative and a negative sign in its results table, an inconsistency resolved here in favour of the table because that is the sign the model implies, but recorded rather than hidden. Ninth, one triplet of serial-correlation and overidentification statistics is reported for two estimated equations, and the archived Hansen statistic and its p value are not mutually consistent at any integer degrees of freedom. Tenth, the derived effect size is strongly convention-dependent: the same test statistic yields f^2 = 0.119 at residual degrees of freedom and 1.063 at the cluster degrees of freedom that actually govern inference, so it is reported only as a comparability metric and never as a variance decomposition. Eleventh, the revised central–regional fiscal relations framework introduces expenditure-composition requirements whose phase-in overlaps the later years of the panel; the archived specification does not model them, and their effect on where adjustment falls should be assessed before this incidence estimate is applied to any period after the window. Finally, the design is ecological. Every association is between district characteristics; none licenses an inference about an individual worker, household or firm, and the wage and displacement quantities in §3.4 are statements about sector averages within a district.

4.4. Implications

Six instruments are set out in Table 9, each with a named implementing actor, a quantified target drawn from this analysis and a sequencing trigger. Naming the actor is not a formality: cross-agency fiscal reform characteristically fails at stewardship rather than at design, and a recommendation addressed to government in general is addressed to no one [52,53].

Table 9. Policy instruments, implementing actors, quantified targets and sequencing triggers.

Instrument	Implementing actor	Quantified target from this study	Sequencing trigger
Energy-transition special allocation grant, introduced through revision of the central–regional fiscal relations framework	Ministry of Finance, fiscal-transfer directorate, with the national development planning ministry	One rule only: compensation = capital-expenditure share × long-run elasticity × proportional fall in revenue sharing × total budget. At the sample mean a 10% fall gives 1.41% of the budget. Indexing instead to the district's revenue-sharing share would over-compensate dependent districts and under-compensate the rest, because the loss is proportional to the capital share, not to the revenue share	A fall in realised revenue sharing beyond normal commodity-cycle variation, verified against production data so that price cycles and arrears settlements do not trigger it
Subnational natural-resource endowment fund capitalised from residual windfalls	Quadrant-IV district governments, under a national framework regulation and audit	The authors' 15% rule contributes 3.69% of the budget each year. Covering a long-run gap of 1.41% of the budget from fund income requires a corpus of 47.1% of one year's budget at a 3% real return, reached after 11.0 years of contribution — 8.3 years at 4% and 16.4 at 2%	Capitalisation must begin while production continues; the fund cannot be filled after the revenue ends
Green-skills certification and transition wage subsidy through the vocational training centres	Ministry of Manpower and provincial vocational training centres, with district labour offices	A displaced extractive worker moving to the manufacturing or agricultural average takes a 29.8% cut against the pre-displacement wage. Bridging it in full costs 5.10 to 10.20 months of post-transition wage	Begun more than one adjustment half-life ahead of closure, that is more than a year in advance

Instrument	Implementing actor	Quantified target from this study	Sequencing trigger
		per worker over the 12–24 month window the source proposes	
Statistical instrumentation of the informal, inactive and out-migration margins at district level	BPS-Statistics Indonesia, with the fiscal-transfer directorate as data user	Only 17.8% of the employment exposed under an 80% output decline appears in the open unemployment rate; the remainder is currently unmeasured, and its composition by sex, age and migration status is unknown	Required before the first closures, or the transition will appear costless in official statistics
Standing district transition forum with worker, community and village representation	District governments, with trade unions, affected village councils and civil-society organisations	Procedural criterion rather than a numerical one: the four quadrants determine who is compensated, and the record contains no evidence on who is represented when that determination is made	Convened before the first compensation decision, not after it
Publication of the comparator-group and estimator documentation underlying this evidence base	The research team, with the data custodians named in Table 2	Publish rather than supply on request: the 396-district comparator estimates, the fixed-effects column, the instrument set and the covariance matrix, so that the design can be replicated and extended to a group-time treatment-effect specification	Immediate; it is the precondition for every other recommendation being auditable

Note. Each instrument is attached to a named implementing actor because cross-agency fiscal reform characteristically fails at stewardship rather than at design. The quantified targets are the derived quantities in Tables 5 and 6 restated as design parameters; they are conditional on the assumptions recorded in those table notes and should be treated as orders of magnitude for instrument design rather than as budget figures. The final two rows are evidence-infrastructure rather than fiscal instruments, and are included because the analysis identifies the absence of measurement, of representation and of documentation as itself a binding constraint on transition governance. Three limits on this table should be read with it. The endowment-fund arithmetic assumes a constant real return, a permanent gap equal to the long-run incidence of a 10% revenue fall, and contributions made while production continues; the number of years to capitalisation is invariant to the assumed return only through the ratio of gap to contribution, which is why the three figures differ. The 15% contribution rate is the authors' own parameter and is not derived here. And the first instrument's trigger, as the record stands, would fire on commodity-price cycles and on arrears settlements as readily as on decarbonisation, which is why it is written to require verification against production data.

The first implication is that a transition transfer must be indexed to incidence, not to population or to foregone revenue alone. Because capital expenditure absorbs between 35.4% and 57.4% of a revenue shock, a transfer sized to replace lost revenue proportionately would under-protect development spending in exactly the districts where it is most rigid. One rule follows from the estimates and Table 9 states only that one: compensation equal to the capital-expenditure share multiplied by the long-run elasticity, the proportional fall in revenue sharing and the total budget. Indexing instead to the district's revenue-sharing share, which is the intuitive rule, would over-compensate the most dependent districts and under-compensate the rest, because the loss is proportional to the capital share rather than to the revenue share. The trigger needs equal care: a year-on-year fall in realised revenue sharing fires on commodity-price cycles and on arrears settlements as readily as on decarbonisation, so it must be verified against the production record before it releases funds. The comparative evidence that investment-protecting fiscal-rule design preserves capital spending during consolidations supports the construction [12,54]. The design has to be specific: subnational fiscal rules that bind on the deficit rather than on current spending have been shown to reduce public investment rather than protect it, and that is the failure mode a transition transfer must be written against [11,55]. The trigger needs the same care, because commodity price shocks transmit to local budgets and labour markets through the same channels and would fire it [31,56]. The wider taxonomy of transitional assistance situates it among the alternatives [57].

The second implication concerns timing, and the arithmetic is unforgiving. The authors' proposed contribution of 15% of residual revenue-sharing windfalls amounts to 3.69% of the district budget each year. Covering a permanent gap equal to the long-run incidence of a 10% revenue fall — 1.41% of the budget — out of fund income requires a corpus of 47.1% of one year's budget at a 3% real return, which those contributions reach after 11.0 years (8.3 at 4%, 16.4 at 2%). The adjustment half-lives reported in Table 5 make this binding. A fund can only be capitalised while the revenue still exists, and the adjustment half-lives of 0.72 and 0.99 years mean that instruments launched at closure arrive after the adjustment has occurred. Retraining in particular must begin more than one adjustment half-life ahead of closure. Where retraining is offered, the evidence favours a specific technology target rather than a generic green-jobs claim: solar deployment matches mining job locations far better than wind on both siting and skill grounds [58], and places diversify into activities related to what they already do [30,35]. What coal workers themselves will accept is a further constraint that supply-side skills mapping does not capture: stated job preferences in comparable settings show that wage level, formality and location all bind, so a technically feasible substitute is not automatically an acceptable one [57,59].

The third implication is a measurement obligation. If under an 80% output decline only 17.8% of exposed employment appears in the open unemployment rate, then the statistic on which transition monitoring will rely is not fit for that purpose, and a transition governed by it will be judged successful on the evidence of its own blind spot. Instrumenting the informal and inactive margins at district level, before the first closures rather than after, is the single cheapest intervention identified by this analysis. The absence of that instrumentation is itself a reportable governance finding, not a footnote.

The fourth implication is distributional and follows the energy-justice criteria directly. Quadrant IV districts have neither the fiscal buffer nor the workforce adaptability to absorb the transition, and the projection suggests they are not a handful of outliers but between one district in six and one in five of the producing population. Sequencing phase-out so that these districts move last, and compensating them first, is what distributional justice requires and what the historical record suggests actually determines whether a phase-out holds [3,16,48].

Distributional justice is not the whole of the energy-justice criterion, and the design of the instruments above has a procedural half that this study can identify but not evaluate. Indonesian district budgeting is a decentralised process in which allocation is negotiated locally, responds to local political competition, and is mediated by relationships between district executives, legislative councils and the enterprises that generate the rents [8,25]. A transfer designed centrally and delivered into that process will be allocated by it. Two consequences follow. First, compensating a district is not the same as compensating the people in it, and the quadrant classification determines only the former. Second, the actors with the least standing in that negotiation — displaced workers, the villages nearest the concessions, and the informally employed who absorb most of the shock — are the ones the distributional analysis identifies as most exposed. A standing district transition forum with worker, village and civil-society representation is therefore listed in Table 9 as an instrument in its own right rather than as a consultation step, and the comparative evidence that a negotiated agreement changes political behaviour in producing regions is the reason to expect it to matter [48,52]. This study offers no evidence on who is currently represented in these decisions, and that absence is itself part of what the evidence base lacks.

A qualification belongs with the fourth implication. The producing group defined here includes metallic-mineral districts as well as coal, oil and gas districts, and it is tempting to treat the former as transition winners because downstream mineral processing has become Indonesia's principal green-industrial strategy [60]. The evidence does not support that reading as a general rule. Transition-mineral projects overwhelmingly overlap territories of land-connected communities, sit in jurisdictions with high social and environmental exposure, and can reproduce rather than resolve extractive social harms [61,62]. Nor is a mineral district's revenue base stable: the fiscal consequences of shifting critical-mineral demand are themselves volatile, so a district that pivots from coal to nickel exchanges one exposure for another rather than escaping it [60,63]. Institutional and financing barriers, not resource endowment, are what constrain Indonesian renewable deployment, and the historical record suggests rapid transitions occurred only where state fiscal capacity was deployed deliberately [10,64]. Quadrant membership should therefore be assigned on measured fiscal and labour readiness, not on commodity type.

The fifth implication is for the evidence base itself. The comparator group, the alternative estimator column and the instrument documentation exist; releasing them would allow this design to be extended to a group-time treatment-effect specification and would make every recommendation above auditable [45,51]. Future work should pursue exactly that extension, add the informality and participation margins as outcomes alongside open unemployment, and test whether the quadrant classification predicts realised transition outcomes as closures begin.

5. Conclusion

Across 118 Indonesian extractive-producing districts observed over 2015–2024, a 10% contraction in natural-resource revenue sharing was associated with a 4.78% contraction in district capital expenditure (95% CI −5.62 to −3.94; $p < 0.001$; $f^2 = 0.119$), rising to 7.76% in the long run and placing 35.4% to 57.4% of the shock on a budget line worth 18.2% of spending. Mining contraction was associated with higher open unemployment (−0.624; $p < 0.001$), amplified by extractive concentration and attenuated by educational attainment, but under an 80% output decline only 17.8% of the exposed employment would appear in the measured rate. H1 to H5 were supported; H6 is a classificatory proposition the archived record does not permit testing, and it is not counted among them. The study extends fiscal-federalism theory to a derivative transfer facing a scheduled phase-out and joins it to labour hysteresis under an energy-justice frame. For Indonesian policy the priorities are a transition transfer indexed to capital-expenditure incidence and triggered on realised revenue falls, endowment funds capitalised while production continues, retraining begun more than one adjustment half-life ahead of closure, and district instrumentation of the informal margin. Future research should exploit the unused 396-district comparator in a group-time treatment-effect design and add participation and informality as outcomes.

6. Declarations

Acknowledgements

The authors thank the Directorate General of Fiscal Balance, Ministry of Finance; BPS-Statistics Indonesia; and the Ministry of Energy and Mineral Resources for maintaining the public series on which this analysis depends. The authors also thank the anonymous reviewers for comments that improved the treatment of precision, identification, and measurement.

Statement on the Use of Artificial Intelligence

The authors confirm that no generative artificial intelligence or AI-assisted technologies were used in the conception, analysis, drafting, editing, or preparation of this manuscript.

Author Contribution Statement

HY: Conceptualization, methodology, data curation, formal analysis, visualization, and writing—original draft. IA: Validation, interpretation, supervision, and writing—review and editing. HY and IA approved the final manuscript and agree to be accountable for all aspects of the work.

Funding Statement

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Conflict of Interest / Competing Interests

The authors declare that they have no financial or non-financial competing interests.

Data Availability Statement

All data analysed in this study are aggregate series published by Indonesian public institutions and are available from their custodians: realised district budget outturns from the Directorate General of Fiscal Balance, Ministry of Finance; labour-force and regional national-accounts series from BPS-Statistics Indonesia; and commodity production and royalty records from the

Ministry of Energy and Mineral Resources. The assembled panel, estimation output, and scripts supporting the derived quantities, tables, and figures are available from the corresponding author upon reasonable request.

Ethics Approval and Consent to Participate

Ethics-committee review was not required because this secondary study analysed only published, aggregate administrative and official statistical series at the district-government level. The study involved no human participants, primary data collection, identifiable individuals, or personal data; therefore, consent to participate was not applicable.

Consent for Publication

Not applicable. The manuscript contains no identifiable individual data.

7. References

1. Ohlendorf N, Jakob M, Steckel JC. The political economy of coal phase-out: Exploring the actors, objectives, and contextual factors shaping policies in eight major coal countries. Energy Res Soc Sci. 2022;90:102590. doi: 10.1016/j.erss.2022.102590
2. Sovacool BK, Turnheim B, Hook A, et al. Dispossessed by decarbonisation: Reducing vulnerability, injustice, and inequality in the lived experience of low-carbon pathways. World Dev. 2021;137:105116. doi: 10.1016/j.worlddev.2020.105116
3. Nacke L, Cherp A, Jewell J. Phases of fossil fuel decline: Diagnostic framework for policy sequencing and feasible transition pathways in resource dependent regions. Oxf Open Energy. 2022;1:oiac002. doi: 10.1093/ooenergy/oiac002
4. Vinichenko V, Cherp A, Jewell J. Historical precedents and feasibility of rapid coal and gas decline required for the 1.5°C target. One Earth. 2021;4(10):1477-1490. doi: 10.1016/j.oneear.2021.09.012
5. Semieniuk G, Holden PB, Mercure JF, et al. Stranded fossil-fuel assets translate to major losses for investors in advanced economies. Nat Clim Chang. 2022;12(6):532-538. doi: 10.1038/s41558-022-01356-y
6. Aleluia J, Tharakan P, Chikkatur A, et al. Accelerating a clean energy transition in Southeast Asia: Role of governments and public policy. Renew Sustain Energy Rev. 2022;159:112226. doi: 10.1016/j.rser.2022.112226
7. Edianto A, Trencher G, Manych N, et al. Forecasting coal power plant retirement ages and lock-in with random forest regression. Patterns. 2023;4(7):100776. doi: 10.1016/j.patter.2023.100776
8. Ordonez JA, Jakob M, Steckel JC, et al. Coal, power and coal-powered politics in Indonesia. Environ Sci Policy. 2021;123:44-57. doi: 10.1016/j.envsci.2021.05.007
9. Halimatussadiah A, Nainggolan D, Yui S, et al. Progressive biodiesel policy in Indonesia: Does the Government's economic proposition hold?. Renew Sustain Energy Rev. 2021;150:111431. doi: 10.1016/j.rser.2021.111431
10. Sambodo MT, Yuliana CI, Hidayat S, et al. Breaking barriers to low-carbon development in Indonesia: deployment of renewable energy. Heliyon. 2022;8(4):e09304. doi: 10.1016/j.heliyon.2022.e09304
11. Ardanaz M, Izquierdo A. Current expenditure upswings in good times and public investment downswings in bad times? New evidence from developing countries. J Comp Econ. 2022;50(1):118-134. doi: 10.1016/j.jce.2021.06.002
12. Ardanaz M, Cavallo E, Izquierdo A, et al. Growth-friendly fiscal rules? Safeguarding public investment from budget cuts through fiscal rule design. J Int Money Finance. 2021;111:102319. doi: 10.1016/j.jimonfin.2020.102319
13. Rinscheid A, Rosenbloom D, Markard J, et al. From terminating to transforming: The role of phase-out in sustainability transitions. Environ Innov Soc Transit. 2021;41:27-31. doi: 10.1016/j.eist.2021.10.019
14. Schmieder JF, von Wachter T, Heining J. The Costs of Job Displacement over the Business Cycle and Its Sources: Evidence from Germany. Am Econ Rev. 2023;113(5):1208-1254. doi: 10.1257/aer.20200252
15. del Rio-Chanona RM, Mealy P, Beguerisse-Díaz M, et al. Occupational mobility and automation: a data-driven network model. J R Soc Interface. 2021;18(174):20200898. doi: 10.1098/rsif.2020.0898
16. van Uffelen N. Revisiting recognition in energy justice. Energy Res Soc Sci. 2022;92:102764. doi: 10.1016/j.erss.2022.102764
17. Wang X, Lo K. Just transition: A conceptual review. Energy Res Soc Sci. 2021;82:102291. doi: 10.1016/j.erss.2021.102291
18. Diluiso F, Walk P, Manych N, et al. Coal transitions—part 1: a systematic map and review of case study learnings from regional, national, and local coal phase-out experiences. Environ Res Lett. 2021;16(11):113003. doi: 10.1088/1748-9326/ac1b58
19. Nowakowska A, Rzeńca A, Sobol A. Place-Based Policy in the "Just Transition" Process: The Case of Polish Coal Regions. Land. 2021;10(10):1072. doi: 10.3390/land10101072
20. Jonek-Kowalska I. Demonstrating the need for a just transition: Socioeconomic diagnosis of polish cities living on hard coal mining. Resour Policy. 2024;89:104576. doi: 10.1016/j.resourpol.2023.104576
21. Krawchenko TA, Gordon M. How Do We Manage a Just Transition? A Comparative Review of National and Regional Just Transition Initiatives. Sustainability. 2021;13(11):6070. doi: 10.3390/su13116070
22. Heras A, Gupta J. Fossil fuels, stranded assets, and the energy transition in the Global South: A systematic literature review. WIREs Clim Change. 2024;15(1):e866. doi: 10.1002/wcc.866
23. Agrawal K, Pathak M, Jana K, et al. Just transition away from coal: Vulnerability analysis of coal districts in India. Energy Res Soc Sci. 2024;108:103355. doi: 10.1016/j.erss.2023.103355
24. Clark A, Zhang W. Estimating the Employment and Fiscal Consequences of Thermal Coal Phase-Out in China. Energies. 2022;15(3):800. doi: 10.3390/en15030800
25. Musviyanti, Nur Khairin F, Bone H, et al. Structure of local government budgets and local fiscal autonomy: Evidence from Indonesia. Public Munic Finance. 2023;11(1):79-89. doi: 10.21511/pmf.11(1).2022.07
26. Ogweno JN, Semedo G. Fiscal Performance and Intergovernmental Fiscal Relations in Developing Countries. Public Finance Rev. 2025;53(4):468-505. doi: 10.1177/10911421241305733
27. Montrone L, Ohlendorf N, Chandra R. The political economy of coal in India – Evidence from expert interviews. Energy Sustain Dev. 2021;61:230-240. doi: 10.1016/j.esd.2021.02.003
28. Blundell R, Bond S. Initial conditions and moment restrictions in dynamic panel data models. J Econom. 1998;87(1):115-143. doi: 10.1016/s0304-4076(98)00009-8
29. Bertheau A, Acabbi EM, Barceló C, et al. The Unequal Consequences of Job Loss across Countries. Am Econ Rev Insights. 2023;5(3):393-408. doi: 10.1257/aeri.20220006
30. Mealy P, Teytelboym A. Economic complexity and the green economy. Res Policy. 2022;51(8):103948. doi: 10.1016/j.respol.2020.103948
31. Owen J, Kemp D, Marais L. The cost of mining benefits: Localising the resource curse hypothesis. Resour Policy. 2021;74:102289. doi: 10.1016/j.resourpol.2021.102289
32. Hidalgo CA. Economic complexity theory and applications. Nat Rev Phys. 2021;3(2):92-113. doi: 10.1038/s42254-020-00275-1
33. Maurizio R, Monsalvo AP, Catania MS, et al. Short-term labour transitions and informality during the COVID-19 pandemic in Latin America. J Labour Mark Res. 2023;57(1):15. doi: 10.1186/s12651-023-00342-x
34. Autor D, Dorn D, Hanson G. On the Persistence of the China Shock. Brookings Pap Econ Act. 2022;2021(2):381-476. doi: 10.1353/eca.2022.0005
35. Balland PA, Broekel T, Diodato D, et al. The new paradigm of economic complexity. Res Policy. 2022;51(3):104450. doi: 10.1016/j.respol.2021.104450
36. Vona F. Managing the distributional effects of climate policies: A narrow path to a just transition. Ecol Econ. 2023;205:107689. doi: 10.1016/j.ecolecon.2022.107689
37. de Chaisemartin C, D'Haultfœuille X. Two-way fixed effects and differences-in-differences with heterogeneous treatment effects: a survey. Econom J. 2023;26(3):C1-C30. doi: 10.1093/ectj/utac017
38. Windmeijer F. A finite sample correction for the variance of linear efficient two-step GMM estimators. J Econom. 2005;126(1):25-51. doi: 10.1016/j.jeconom.2004.02.005
39. Driscoll JC, Kraay AC. Consistent Covariance Matrix Estimation with Spatially Dependent Panel Data. Rev Econ Stat. 1998;80(4):549-560. doi: 10.1162/003465398557825
40. Roodman D. How to do Xtabond2: An Introduction to Difference and System GMM in Stata. Stata J. 2009;9(1):86-136. doi: 10.1177/1536867x0900900106
41. Abadie A, Athey S, Imbens GW, et al. When Should You Adjust Standard Errors for Clustering?. Q J Econ. 2022;138(1):1-35. doi: 10.1093/qje/qjac038
42. MacKinnon JG, Nielsen MØ, Webb MD. Cluster-robust inference: A guide to empirical practice. J Econom. 2023;232(2):272-299. doi: 10.1016/j.jeconom.2022.04.001
43. MacKinnon JG, Nielsen MØ, Webb MD. Fast and reliable jackknife and bootstrap methods for cluster-robust inference. J Appl Econom. 2023;38(5):671-694. doi: 10.1002/jae.2969
44. Arellano M, Bond S. Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations. Rev Econ Stud. 1991;58(2):277. doi: 10.2307/2297968
45. Imbens GW. Statistical Significance, p-Values, and the Reporting of Uncertainty. J Econ Perspect. 2021;35(3):157-174. doi: 10.1257/jep.35.3.157

46. Illing H, Schmieder J, Trenkle S. The Gender Gap in Earnings Losses After Job Displacement. J Eur Econ Assoc. 2024;22(5):2108-2147. doi: 10.1093/jeea/jvae019
47. Wang X, Lo K. Political economy of just transition: Disparate impact of coal mine closure on state-owned and private coal workers in Inner Mongolia, China. Energy Res Soc Sci. 2022;90:102585. doi: 10.1016/j.erss.2022.102585
48. BOLET D, GREEN F, GONZÁLEZ-EGUINO M. How to Get Coal Country to Vote for Climate Policy: The Effect of a “Just Transition Agreement” on Spanish Election Results. Am Polit Sci Rev. 2024;118(3):1344-1359. doi: 10.1017/s0003055423001235
49. Meckling J. Making Industrial Policy Work for Decarbonization. Glob Environ Polit. 2021;21(4):134-147. doi: 10.1162/glep_a_00624
50. Goodman-Bacon A. Difference-in-differences with variation in treatment timing. J Econom. 2021;225(2):254-277. doi: 10.1016/j.jeconom.2021.03.014
51. Callaway B, Sant’Anna PH. Difference-in-Differences with multiple time periods. J Econom. 2021;225(2):200-230. doi: 10.1016/j.jeconom.2020.12.001
52. Matanzima J. “Disempowered by the transition”: Manipulated and coerced agency in displacements induced by accelerated extraction of energy transition minerals in Zimbabwe. Energy Res Soc Sci. 2024;117:103727. doi: 10.1016/j.erss.2024.103727
53. Setyowati AB. Governing sustainable finance: insights from Indonesia. Clim Policy. 2023;23(1):108-121. doi: 10.1080/14693062.2020.1858741
54. Mancini AL, Tommasino P. Fiscal rules and the reliability of public investment plans: Evidence from local governments. Eur J Polit Econ. 2023;79:102415. doi: 10.1016/j.ejpoleco.2023.102415
55. Pathak R. Do Subnational Fiscal Rules Reduce Public Investment? The Case of Fiscal Responsibility Laws in India. Public Finance Rev. 2023;51(3):315-338. doi: 10.1177/10911421231158199
56. Naraidoo R, Paez-Farrell J. Commodity price shocks, labour market dynamics and monetary policy in small open economies. J Econ Dyn Control. 2023;151:104654. doi: 10.1016/j.jedc.2023.104654
57. Gazmararian AF. Fossil fuel communities support climate policy coupled with just transition assistance. Energy Policy. 2024;184:113880. doi: 10.1016/j.enpol.2023.113880
58. Cherp A, Vinichenko V, Tosun J, et al. National growth dynamics of wind and solar power compared to the growth required for global climate targets. Nat Energy. 2021;6(7):742-754. doi: 10.1038/s41560-021-00863-0
59. Blankenship B, Aklin M, Urpelainen J, et al. Jobs for a just transition: Evidence on coal job preferences from India. Energy Policy. 2022;165:112910. doi: 10.1016/j.enpol.2022.112910
60. Lahadalia B, Wijaya C, Dartanto T, et al. Nickel Downstreaming in Indonesia: Reinventing Sustainable Industrial Policy and Developmental State in Building the EV Industry in ASEAN. J ASEAN Stud. 2024;12(1):79-106. doi: 10.21512/jas.v12i1.11128
61. Owen JR, Kemp D, Lechner AM, et al. Energy transition minerals and their intersection with land-connected peoples. Nat Sustain. 2022;6(2):203-211. doi: 10.1038/s41893-022-00994-6
62. Sovacool BK. When subterranean slavery supports sustainability transitions? power, patriarchy, and child labor in artisanal Congolese cobalt mining. Extr Ind Soc. 2021;8(1):271-293. doi: 10.1016/j.exis.2020.11.018
63. D’Orazio P. Assessing the fiscal implications of changes in critical minerals’ demand in the low-carbon energy transition. Appl Energy. 2024;376:124197. doi: 10.1016/j.apenergy.2024.124197
64. Newell P, Simms A. How Did We Do That? Histories and Political Economies of Rapid and Just Transitions. New Polit Econ. 2021;26(6):907-922. doi: 10.1080/13563467.2020.1810216