Modern Economy, 2026, 17(3), 430-452
https://www.scirp.org/journal/me
ISSN Online: 2152-7261
ISSN Print: 2152-7245

Chinese Development Finance and MDG Gap Dynamics: Evidence from Child Mortality and Poverty

Prabuddha Sanyal

Straight Deal Mortgage, Huntington Beach, CA, USA
Email: prabuddha@straightdealmortgage.com

How to cite this paper: Sanyal, P. (2026). Chinese Development Finance and MDG Gap Dynamics: Evidence from Child Mortality and Poverty. *Modern Economy, 17,* 430-452.
https://doi.org/10.4236/me.2026.173023

Received: November 27, 2025
Accepted: March 3, 2026
Published: March 6, 2026

Abstract

This paper examines whether Chinese development finance is associated with faster progress toward Millennium Development Goal style targets in low- and middle-income countries. We combine AidData's Chinese Official Finance to the Global South data with World Bank indicators and study progress through the lens of MDG gaps differences between observed outcomes and a target level. We model these gaps as dynamic processes, allowing Chinese finance to shift the trajectory while accounting for strong persistence over time. The results show that under five mortality gaps follow highly persistent AR(1) dynamics consistent with a constant proportional closing of the remaining distance to the target an empirical analogue to Zeno's paradox in which progress continues but reaching zero is slow and asymptotic. At the same time, improvements in average gaps do not translate into broad cross-country convergence. The distribution of gaps retains high Shannon entropy, indicating that dispersion remains substantial even as mean outcomes improve. Overall, the findings suggest that MDG progress is governed by a sticky law of motion in which external finance may affect levels or rates at the margin, but global development outcomes can improve without the system collapsing into a low entropy configuration where most countries cluster near the target.

Keywords

Chinese Development Finance, Millennium Development Goals, Zeno Paradox, Shannon Entropy, Aid Effectiveness, Under-Five Mortality

1. Introduction

The Millennium Development Goals were designed as time bound and numerically explicit commitments on poverty, education, and health. Implicitly, they framed development as a problem of closing gaps by a deadline: by 2015, countries would move measurably toward benchmark outcomes such as lower child mortality and reduced extreme poverty. Over the same period, China emerged as a major source of development finance across the Global South, raising a second set of questions about whether a new and large financier could accelerate progress on these internationally defined goals.

This paper brings these two themes together by focusing on the dynamics of MDG related gaps and by asking how Chinese official finance fits into that dynamic. Rather than evaluating progress only in levels or average trends, we treat the gap to a target as an evolving state variable. This shifts the emphasis from whether outcomes improved to how they improve over time and whether the global distribution of progress is converging or remaining dispersed.

The first question is whether the typical pattern of progress is proportional: do countries close a roughly constant share of the remaining gap each year? If so, improvements will be largest when the gap is large and will mechanically slow as the target approaches. This produces an asymptotic approach to the goal a trajectory that looks like steady progress but may rarely hit zero in finite sample data. In that sense, the MDG gap behaves like an empirical counterpart to Zeno's paradox, where the remaining distance keeps shrinking but full closure becomes increasingly difficult.

The second question concerns the system level distribution of gaps. Even if the global mean gap falls, the world may or may not converge to a configuration in which most countries cluster near the target. To capture this, we track the cross-country distribution using Shannon entropy[1]. A decline in entropy would indicate convergence toward a low dispersion world; persistently high entropy would indicate that progress is uneven and that substantial dispersion remains amid broad improvements.

The third question is how Chinese development finance interacts with these dynamics. Chinese funding could meaningfully change the law of motion for gaps by accelerating closure, or it could act primarily as a level shifter within a process that is otherwise dominated by persistence and structural constraints. Distinguishing between these possibilities matters for how we interpret aid effectiveness claims and for whether large scale finance can realistically deliver convergence rather than improvement with continued dispersion.

Empirically, we link AidData's geocoded panel of Chinese official finance to World Bank outcome indicators, focusing on under-five mortality and the poverty gap index in our event-study analysis. For each outcome we construct an MDG style gap relative to a target and estimate dynamic panel models in which the gap

[1]Shannon entropy is a summary measure of how concentrated or spread out the cross-country distribution of gaps is across categories (e.g., bins of gap sizes). It is higher when countries are distributed relatively evenly across bins, indicating a more dispersed pattern, and lower when countries are concentrated in a small number of bins, indicating clustering. Entropy therefore complements the mean gap by describing the shape of the distribution; a low-entropy outcome can reflect convergence near the target, but it can also reflect clustering away from the target, so it should be interpreted alongside the location of the mass.

follows an autoregressive process and Chinese finance enters as a covariate. This framework allows us to quantify persistence, evaluate whether gap closure is approximately proportional, and assess whether external finance is associated with shifts in trajectories.

The results highlight a central tension in global development progress. Under five mortality gaps exhibit strong persistence and a proportional closing pattern consistent with an asymptotic approach to targets. Meanwhile, the entropy of the cross country distribution remains high, implying that average improvements do not translate into broad convergence. Taken together, the paper's contribution is to link a simple but powerful dynamic interpretation of MDG progress with a distributional measure of global dispersion, and to situate Chinese development finance within that combined Zeno and entropy perspective.

The contribution of this paper is threefold. First, we provide an explicit, data-driven formulation of MDG progress as a Zeno-type process in which shortfalls are reduced proportionally over time rather than eliminated in finite time. Second, we show that cross-country dispersion in MDG shortfalls remains high even as the average shortfall declines, consistent with persistently high entropy and limited global convergence toward the targets. Third, we evaluate the role of Chinese development finance within this structure, finding that while it is associated with meaningful improvements in MDG indicators, it does not by itself generate convergence toward a low-entropy distribution in which most countries cluster near an "all-targets-met" state.

2. Literature Review

This paper speaks to three strands of literature: empirical aid effectiveness, the political economy of Chinese development finance, and theoretical work on convergence and entropy in economic systems.

2.1. Aid Effectiveness and MDG-Type Outcomes

The aid effectiveness literature has long debated whether external assistance improves growth and social outcomes in developing countries, with evolving methods and conclusions.

Early cross-country studies, most famously Burnside and Dollar (2000), argued that aid's impact on growth is conditional on "good policies" (macroeconomic stability, openness, institutional quality). This line of work used panel regressions with aid to GDP ratio and policy indices, typically finding that aid is growth-enhancing only in favorable policy environments. Later re-examinations raised doubts about the robustness of these findings to sample selection, model specification, and instrument choice, underscoring the sensitivity of aggregate aid-growth regressions (Easterly, 2009).

In parallel, a second generation of studies shifted from macro growth to sectoral and social outcomes more directly aligned with the MDGs: health, education, and poverty. Cross-country work has examined whether higher aid flows correlate with lower infant and under-five mortality, increased school enrollment, or reduced poverty headcounts. Results are generally mixed: some studies find statistically significant associations, but effect sizes are modest and often sensitive to model choice, data definitions, and endogeneity corrections (Boone, 1996; Dreher et al., 2008; Mishra & Newhouse, 2009; Rajan & Subramanian, 2008).

More recently, micro-level and project-level analyses have provided a more granular picture. Clemens et al., (2011) and others distinguish between "early-impact" aid (e.g. infrastructure, agriculture) and slower-moving categories[2], finding that some types of aid have clearer growth or social effects over medium horizons. Randomized controlled trials and quasi-experimental evaluations of specific health and education interventions (e.g. deworming, conditional cash transfers, school construction) tend to show local, often substantial improvements, but these do not always aggregate cleanly into national-level MDG progress.

Many authors move beyond aggregate "aid totals" to emphasize which types of aid arrive when, arguing that only certain components are likely to affect measured growth over short horizons (Clemens et al., 2011) and that macro-level findings are highly sensitive to identification choices and proposed transmission mechanisms (Rajan & Subramanian, 2008; Easterly et al., 2004). Meanwhile, project-level RCT and quasi-experimental evidence shows sizable local gains in health, schooling, and behavior (Miguel & Kremer, 2004; Cohen & Dupas, 2010; Gertler, 2004), but scaling those impacts into national accounts or MDG aggregates is not straightforward.

Across these waves, three themes stand out for this paper. First, the literature increasingly recognizes heterogeneous and context-dependent impacts rather than a single average treatment effect of "aid." Second, most macro work treats aid as a regressor in growth or level equations rather than as part of the dynamic process by which gaps to international targets evolve. Third, the focus is overwhelmingly on aggregate official development assistance from OECD-DAC donors, with far less attention to non-traditional donors such as China (Burnside & Dollar, 2000; Roodman, 2007; Rajan & Subramanian, 2008; Bräutigam, 2011).

2.2. Chinese Development Finance and Its Effects

A second, rapidly growing body of work focuses on China's role as a development financier in the Global South. This literature is anchored by AidData and related efforts that systematically track Chinese official finance (including both concessional "aid-like" flows and more commercial-like loans) to Africa, Asia, and Latin America (AidData. 2017; Dreher et al., 2016; Strange et al., 2017).

Descriptive work documents the scale, sectoral composition, and modalities of Chinese official finance, emphasizing large infrastructure (roads, power, ports),

[2]Clemens et al. (2011) the idea is that some aid categories can plausibly affect growth (and some social outcomes) within a few years, while other categories either operate mainly through very long lags or are not primarily aimed at raising measured GDP. So they split aid into buckets based on expected time-to-impact and mechanism.

resource-related projects, and industrial parks/overseas economic zones, with more limited emphasis on social sectors (Bräutigam, 2009; Dreher et al., 2018; Bräutigam & Gallagher, 2014). This literature also highlights differences from OECD-DAC donors in financing terms, conditionality narratives, and implementation practices (Bräutigam, 2011), and situates Chinese development finance within a broader "going out" strategy and geopolitical objectives (Taylor, 2006; Alden, 2007; Downs, 2011).

Relative to OECD-DAC aid, Chinese official finance is often characterized as less policy-conditional and more project-centered, with implementation frequently bundled through Chinese firms, which may shorten project timelines. These features align with China's Go Global strategy, where overseas finance supports commercial expansion and diplomatic objectives alongside development goals (Bräutigam, 2011; Buckley et al., 2007; Wellner et al., 2022).

Empirical work on Chinese official finance and related overseas engagement increasingly evaluates local economic impacts using geocoded project data and subnational outcomes (often including remotely sensed measures such as nighttime lights), building on both the emerging China-finance micro-evidence and established spatial approaches to measuring local development from satellites (Isaksson & Kotsadam, 2018; Dreher et al., 2019; Henderson et al., 2012). Second, a large literature looks at debt sustainability and financial risk. It studies when Chinese lending is linked to a higher risk of debt distress, and it also examines how specific contract terms shape the lender's leverage and the dynamics of debt crises. Those contract terms can include collateral-like arrangements, repayment and escrow mechanisms, and clauses that matter for debt restructuring (Horn et al., 2021; Gelpern et al., 2022; Brautigam, 2020).

Third, work in governance and political economy examines how Chinese projects are allocated and how they shape domestic politics, including whether project placement tracks political alignment or strategic regions, whether it strengthens incumbents, and how it interacts with institutional quality (Isaksson & Kotsadam, 2018; Dreher et al., 2019; Brazys & Vadlamannati, 2021; Sanyal & Babu, 2012).

By contrast, evidence on social-sector outcomes is comparatively thinner. Within health and education, the literature contains numerous case studies of Chinese-funded hospitals, training programs, medical teams, and scholarships. However, systematic cross-country or panel-based analyses that connect Chinese finance to MDG-style outcomes such as under-five mortality, school enrollment, or poverty headcounts remain limited. Most quantitative research instead emphasizes macroeconomic and financial indicators—growth, trade, investment, and debt—rather than population-level welfare outcomes tracked by the MDGs (King, 2013; Dreher et al., 2018; Horn et al., 2021).

Two gaps in this literature are directly addressed here. First, relatively few studies link Chinese finance to MDG-type social indicators in a dynamic framework that recognizes persistence and noise in the evolution of gaps. Second, almost none embed Chinese finance within a formal convergence/entropy perspective, asking not only whether outcomes improve on average, but how the distribution of gaps across countries evolves and how likely it is that many countries will simultaneously reach target levels (Dreher et al., 2019; Isaksson & Kotsadam, 2018; Quah, 1993; Quah, 1996).

2.3. Convergence, Complexity, and Entropy in Economic Systems

A third, smaller but conceptually important literature deals with convergence, complexity, and entropy in economic systems.

The growth and inequality literatures distinguish between beta-convergence (poor units growing faster than rich ones) and sigma-convergence (declining dispersion over time). Empirical work often finds some evidence of beta-convergence conditional on covariates, but weaker or absent sigma-convergence: the cross-sectional variance of income or productivity frequently remains high or even increases. Similar patterns show up in cross-country health and education indicators: substantial average improvements can coexist with persistent cross-sectional dispersion (Buckley et al., 2007; Horn et al., 2021; Barro & Sala-i-Martin, 1992; Pritchett, 1997; Becker et al., 2005).

Alongside this, a more theoretical and methodological strand uses information-theoretic measures, particularly Shannon entropy, to describe economic distributions and dynamics. Entropy has been used to analyze the diversification of export structures, the distribution of income or wealth, portfolio allocation under uncertainty, and even the "disorder" of macroeconomic regimes. The core insight is that low-entropy, highly ordered states occupy a small region of the state space: they are possible but statistically rare under many plausible stochastic processes (Shannon, 1949; Theil, 1967; Maasoumi, 1993; Cover & Thomas, 2006; Sims, 2003).

Complexity and econophysics approaches build on this intuition by treating the economy as a high-dimensional system made up of many interacting agents, sectors, and institutions. Because these components influence one another, aggregate outcomes can exhibit path dependence (history matters), nonlinear feedback (small changes can be amplified), and large responses to shocks, generating rich and sometimes unpredictable dynamics over time. In this view, uniform "success" outcomes such as everyone rich or everyone at target represent a relatively narrow set of underlying micro-level arrangements; they are possible, but not typical. More commonly, even under broadly similar conditions, economies evolve toward heterogeneous distributions across places and groups, with persistent dispersion and shifting patterns as interactions and shocks propagate through the system (Arthur, 1994; Arthur, 1999; Beinhocker, 2006; Farmer & Foley, 2009; Bouchaud, 2013; Yakovenko & Rosser, 2009).

The connection to development goals is straightforward, but it has not been explored much. Take a specific MDG target, for example getting under five mortality below 25 deaths per 1000 in every low- and middle-income country. One can think of that goal as a "low-entropy" outcome, meaning most countries end up tightly clustered near the target. From that perspective, progress is not only about whether the average country improves. It also depends on how the entire distribution of outcomes across countries changes, including whether the spread narrows enough that lagging countries catch up.

If these indicators follow an AR(1) process with random shocks, countries may move closer to the target over time, but unpredictable shocks make it unlikely that every country will be below the threshold by any specific deadline. Thus, gradual progress is possible, but reaching the target everywhere at a set time remains improbable (Doob, 1953).

3. Methods and Equations

3.1. MDG Gaps

For each country i and year t, we define MDG gaps as deviations from simple numerical targets.

Under-five mortality gap: $g_{i,t}^{\mathrm{u5}} = \mathrm{U5MR}_{i,t} - \tau^{\mathrm{u5}}$ where $\mathrm{U5MR}_{i,t}$ is the under-five mortality rate (deaths per 1000 live births), and τ^{u5} is an MDG-style target (e.g., 25)[3]. Positive gaps indicate underperformance relative to the target, whereas negative gaps indicate overachievement.

3.2. Zeno-Style Dynamic Specification (AR(1) with Chinese Aid)

To capture the temporal evolution of MDG gaps, we estimate a first-order autoregressive process with Chinese development finance as a covariate. For a generic gap $g_{i,t}$ (which is $g_{i,t}^{\mathrm{u5}}$ in our case), the baseline empirical specification is:

$$g_{i,t+1} = \alpha + \beta g_{i,t} + \gamma \log\left(1 + \mathrm{Aid}_{i,t}\right) + \varepsilon_{i,t+1}, \tag{1}$$

where: α is a constant term, β is the persistence (or convergence) coefficient, γ captures the effect of Chinese aid, $\mathrm{Aid}_{i,t}$ is the volume of Chinese official finance to country i in year t, and $\varepsilon_{i,t+1}$ is an idiosyncratic shock with mean zero and variance σ^2.

Conditional on the current gap and aid, the one-step-ahead expectation is:

$$\mathbb{E}\left[g_{i,t+1} \mid g_{i,t}, \mathrm{Aid}_{i,t}\right] = \alpha + \beta g_{i,t} + \gamma \log\left(1 + \mathrm{Aid}_{i,t}\right). \tag{2}$$

If we focus on the pure Zeno mechanism and abstract from aid and the intercept for intuition, we have:

$$g_{t+1} = \beta g_t.$$

Iterating forward from an initial gap g_0 yields:

[3]We also attempted primary school enrolment gaps and poverty gap measures as our outcome measures. because these series are too sparsely and irregularly observed in our borrower-country panel once merged to annual Chinese-finance flows. The resulting estimation samples are small and unbalanced, with few country clusters and limited within-country time variation after country and year fixed effects, making standard cluster-robust inference unreliable and estimates highly sensitive to sample composition. Where feasible, we report these outcomes only as descriptive patterns/robustness checks.

$$g_t = \beta^t g_0.$$

For $0 < \beta < 1$,

$$\lim_{t\to\infty} g_t = 0, \text{ yet } g_t > 0 \text{ for all finite } t \text{ if } g_0 > 0.$$

This is the discrete-time Zeno property: the gap converges to zero asymptotically but is never exactly zero at any finite time under a continuous process[4].

The change in the gap can be decomposed as:

$$\Delta g_{i,t+1} \equiv g_{i,t+1} - g_{i,t} = (\beta - 1) g_{i,t} + \alpha + \gamma \log(1 + \text{Aid}_{i,t}) + \varepsilon_{i,t+1}, \tag{3}$$

where the term $(\beta - 1) g_{i,t}$ captures the proportional reduction in the remaining gap.

3.3. Continuous-Time Zeno Limit

For conceptual clarity, the Zeno idea can also be expressed in continuous time. Let $g_i(t)$ denote the gap at continuous time t. A standard exponential decay process is:

$$\frac{\mathrm{d}g_i(t)}{\mathrm{d}t} = -\lambda g_i(t) \tag{4}$$

with solution:

$$g_i(t) = g_i(0)\mathrm{e}^{-\lambda t}, \lambda > 0.$$

Again, the remaining gap as a fraction of the initial gap is given by:

$$\frac{g(t)}{g(0)} = \mathrm{e}^{-\lambda t} \tag{5}$$

The discrete AR(1) coefficient ρ and the continuous decay parameter λ are related by the following:

$$\rho \approx \mathrm{e}^{-\lambda \Delta t} \tag{6}$$

where Δt is the length of one period (e.g., one year) or equivalently,

$$\lambda \simeq -\frac{\ln \rho}{\Delta t} \tag{6'}$$

3.4. Shannon Entropy of Cross-Country MDG Gaps

To characterize the dispersion of MDG gaps across countries at a given time, we

[4]One can think of the gap as how far something is from its goal. Each step, the gap is updated using some rule. What happens in the long run mostly depends on how this rule behaves when the gap is already small, meaning the system is close to the goal. If the goal is stable, then once the gap is small, each step tends to shrink it by roughly the same fraction. The gap then falls quickly at first and more slowly later (this is geometric decay). If the rule instead multiplies small gaps by a number bigger than one, the gaps grow and the system moves away from the goal (geometric growth). There is also a borderline case where the rule does not clearly shrink the gap in percentage terms. In that case, the simple "same percentage each step" story does not work. Any convergence that happens is slow and drawn out, with progress becoming especially slow as the gap gets small.

use Shannon entropy. Let $g_{i,t}$ denote the gap for country i in year t, and consider its empirical cross-country distribution at fixed t.

Because the gaps are continuous, we approximate the distribution with a histogram. Partition the support of $g_{i,t}$ into K bins, and let $p_{k,t}$ be the fraction of countries whose gaps fall into bin k in year t. The discrete Shannon entropy of the cross-country gap distribution at time t is:

$$H_t = -\sum_{k=1}^{K} p_{k,t} \log p_{k,t} \tag{7}$$

with the convention that terms with $p_{k,t} = 0$ are omitted.

High values of H_t indicate a spread-out, high-dispersion (high-entropy) distribution of MDG gaps across countries. Low values of H_t correspond to many countries clustered in a narrow range, indicative of a more ordered, low-entropy configuration.

In a continuous idealization, if $f_t(g)$ is the probability density function of $g_{i,t}$ across countries in year t, the differential entropy is:

$$h_t = -\int f_t(g) \log f_t(g) \mathrm{d}g. \tag{8}$$

In practice, the empirical analysis uses the discrete, binned form H_t, computed from the observed gaps.

3.5. Near-Zero Events and Practical MDG "Success"

Because the probability of exactly hitting $g_{i,t} = 0$ is essentially zero under a continuous AR(1) noise process, we also define a "near-zero" set for a tolerance $\varepsilon > 0$:

$$\mathcal{N}_\varepsilon = \{ g : |g| < \varepsilon \}.$$

Given a panel or a set of simulated trajectories, the empirical frequency of near-zero events is:

$$\hat{P}_\varepsilon = \frac{1}{NT} \sum_{i=1}^{N} \sum_{t=1}^{T} \mathbf{1}\{ g_{it} \mid \varepsilon \} \tag{9}$$

where $\mathbf{1}(\cdot)$ is the indicator function. This quantity summarizes how often countries come close to their MDG targets, even if they almost never hit them exactly.

4. Data Description

Our analysis combines project-level data on Chinese official finance with country-year indicators of progress toward the Millennium Development Goals (MDGs).

The core financing data come from AidData's Chinese Loans to Low- and Middle-Income Countries (CLG-LMIC) Dataset, Version 2.0, distributed as AidDatas_CLG_LMIC_Dataset_v1.0.xlsx. We use the CLG-LMIC 2.0_Records sheet, which reports one observation per AidData record; records can correspond to loans, grants, technical assistance, scholarships/training, and debt actions, and multiple records may map to the same underlying financing event via shared Loan_Event_ID (e.g., when events are reported in tranches)[5].

Each record reports the recipient country (Country_of_Activity_ISO3) and commitment year (Commitment Year), along with detailed financing attributes such as Flow Type and Flow Class (ODA vs OOF). Concessionality is primarily captured through classification fields, while granular loan-term variables needed to compute a grant element are missing for most records. Amounts are provided in the original currency and in several U.S. dollar conversions; throughout, we use Amount_Constant_USD_2023, which measures commitments in inflation-adjusted 2023 U.S. dollars.

4.1. Aggregating the Country-Year Panel Data with MDG Metrics

To construct a country-year panel of Chinese finance, we aggregate the project-level CLG-LMIC data by recipient ISO3 code and commitment year. For each country-year pair, we sum Amount_Constant_USD_2023 across all records to obtain an annual series of total Chinese official finance, which we denote china_official_finance_total_usd_2023. Years with no recorded Chinese finance for a given country are coded as zero commitments in the main specification, although in some robustness checks we treat them as missing to assess sensitivity to the extensive margin. The resulting panel covers a broad set of low-income and middle-income countries over the period in which AidData's CLG-LMIC collection is most complete; for the empirical exercises below, we focus on the 2000-2023 window, where overlap with MDG-related outcome data is strongest.

The MDG outcome data are taken from the World Bank's World Development Indicators (WDI)[6]. As a concrete MDG-relevant metric, we use the under-five mortality rate (SH.DYN.MORT, deaths of children under age five per 1,000 live births). For each country and year, we extract the corresponding under-five mortality rate, retaining all observations with non-missing values. These WDI series are provided at annual frequency, with coverage that typically extends from 1990 to the most recent year available. We harmonize country identifiers by mapping the WDI country ISO3 codes to the AidData Country_of_Activity_ISO3 codes and then merge the two datasets on country-year pairs.

To cast under-five mortality in MDG gap form, we set a target of 25 deaths per 1000 live births. For each country i and year t, with observed under-five mortality u5mr_{it}, we define the MDG gap as $gap_{it}^{\text{u5}} = \max\{0, \text{u5mr}_{it} - 25\}$, so the gap measures the remaining excess mortality relative to the target. Positive values of $g_{i,t}$ indicate that a country remains above (worse than) the target; non-positive values indicate that the target has been met or surpassed. In the empirical analysis,

[5]The following data source was used in the study: https://www.aiddata.org/data/aiddatas-global-chinese-development-finance-dataset-version-2-0.

[6]The exact data source is: https://databank.worldbank.org/source/world-development-indicators.

we work directly with this continuous u5_gap variable. The merged country-year panel therefore contains, for each country with available data, under-five mortality (u5mr), the corresponding MDG gap (u5_gap), and the annual volume of Chinese finance in constant 2023 USD (china_aid_usd).

4.2. Data Limitations

Several limitations of the data should be noted. First, the CLG-LMIC dataset covers officially financed projects that can be identified and geocoded by AidData; some Chinese flows, especially more opaque or short-term instruments, may be missing or measured with error. Second, the timing of commitments (Commitment_Year) may not align perfectly with disbursements, so annual commitment totals are only an approximation to the effective yearly financial impulse. Third, WDI under-five mortality estimates for low-income countries are partly model-based and subject to revision; measurement error in u5mr will translate into noise in the constructed MDG gaps.

5. Main Results

This section provides the main results of the study. This section is organized as follows: (a) Descriptive Statistics; (b) Zeno Dynamics of the MDG Gap; and (c) Shannon entropy of the cross-country gap distribution.

5.1. Descriptive Statistics

We begin by describing the joint distribution of under-five mortality, MDG gaps, and Chinese official finance in the merged country-year panel. The unit of observation is a country-year between 2000 and 2023. Table 1 reports summary statistics for under-five mortality (u5mr), the MDG gap (u5_gap), and Chinese official finance commitments (china_aid_usd, constant 2023 USD).

Table 1. Under-5 Mortality, MDG Gaps. And Chinese official finance commitments.

Statistic	u5mr	$g_{i,t}^{u5}$	China Aid (in USD billions)
count	5760.0	5760.0	
mean	39.194	14.194	0.8898
std	40.723	40.723	3.4209
min	1.4	−23.6	0.0
10%	4.4	−20.6	0.0001
25%	9.6	−15.4	0.0076
50%	22.9	−2.1	0.0539
75%	57.169	32.169	0.4463
90%	99.11	74.11	1.9325
max	478.9	453.9	89.7245

Source: Computed by the Authors' using CLG-LMIC dataset.

5.2. Zeno Dynamics of the MDG Gap

To characterize the temporal dynamics of MDG gaps, we estimate a simple autoregressive model of order one AR(1) at the country-year level. Let g_{it} denote the under-five mortality gap in country i and year t. Our baseline specification is a pooled ordinary least squares (OLS) regression of the form:

$$g_{it} = \alpha + \phi g_{i,t-1} + \varepsilon_{it} \tag{10}$$

where α is a constant, ϕ is the persistence parameter, and ε_{it} an error term. The regression pools all available country-year observations, with the lagged gap constructed within countries. Table 2 reports coefficient estimates for the baseline AR (1) model and an aid-augmented variant that includes lagged Chinese finance as an additional regressor.

Table 2. Pooled AR(1) gap dynamics with and without Chinese aid.

	Baseline AR(1)	Aid-augmented
α (intercept)	−0.49**	−0.723***
	(0.194)	(0.107)
ϕ (lagged gap)	0.938***	0.936***
	(0.02)	(0.02)
γ (log lagged aid)		0.033
		(0.022)
Adjusted R^2	0.943	0.943
F-statistic	90953.22	45515.1
Observations	5520	5520

Source: Computed by the Authors' using CLG-LMIC dataset. ***denote significance at 1% level, **denote significance at 5% level, *denote significance at the 10% level.

The estimated persistence parameter ϕ is approximately 0.94 in both specifications, implying that only about six percent of the remaining gap is closed each year. This high persistence is the empirical analogue of Zeno dynamics: gaps decay geometrically but converge to the target only asymptotically. The coefficient on lagged Chinese aid is small in magnitude, and the Adjusted R^2 of the aid-augmented model is almost identical to the baseline, indicating that adding Chinese aid does not materially alter the strongly persistent nature of gap dynamics in this reduced-form specification.

5.3. Shannon Entropy of the Cross-Country Gap Distribution

While the AR(1) results speak to within-country dynamics, we also ask how the cross-country configuration of MDG gaps evolves over time. For each year, we compute the Shannon entropy of the binned distribution of gaps across countries. Let $p_{k,t}$ denote the share of countries whose gap falls in the bin k in year t. The entropy measure (in bits) is defined as in Equation (7).

Higher values of H_t correspond to a more dispersed, high-entropy configuration of country gaps. Table 3 reports entropy and the mean MDG gap for the first few years of the sample as an illustration.

Table 3. Evolution of gap entropy and mean under-five mortality gap, 2000-2004.

Year	Entropy	Mean u5_gap
2000	2.464	33.533
2001	2.48	31.308
2002	2.47	29.024
2003	2.466	26.575
2004	2.459	24.679

Notes: Entropy is computed from the cross-country distribution of MDG gaps using base-2 logarithms. Mean u5 gap is the cross-country average of the under-five mortality gap in each year.

Entropy declines only modestly over time and remains at a relatively high level even as the mean MDG gap shrinks, indicating that the global distribution of gaps remains dispersed rather than collapsing into a low-entropy, near-target state. Taken together with the AR(1) estimates, this supports a Zeno-plus-entropy interpretation of MDG progress: countries individually close a constant fraction of the remaining gap each period, but the cross-country macrostate remains high-entropy and far from a uniform near-target configuration.

5.4. Robustness Checks

This section summarizes a set of robustness checks on the Zeno-style MDG gap dynamics and the associated entropy patterns. We examine alternative dynamic specifications, additional fixed effects, different gap definitions, and alternative entropy constructions. Across these variations, the core findings of slow geometric convergence and a persistently high-entropy cross-country configuration of gaps remain remarkably stable.

Robustness checks here mainly serve as power and identification diagnostics. Baseline dynamic models show small, statistically indistinguishable finance effects for under-five mortality gaps. Some exposed-subsample and aggregated-finance variants produce marginal significance, but they hinge on few countries. Poverty analyses are underpowered due to limited coverage and clustering constraints. Taken together, these exercises reinforce the main takeaway from the entropy analysis persistence and heterogeneity dominate, and short-run finance-lag effects are not precisely estimated.

This perspective also clarifies why the development-finance coefficients are difficult to estimate precisely in fixed-effects regressions. If most countries receive little or no Chinese finance in most years, within-country variation in finance is sparse. When outcome dynamics are strong, the AR terms absorb much of the predictable change, leaving limited residual variation for finance lags to explain. In a zero-entropy framing, this is the empirical analogue of the distribution being "sticky": the system's state (the gap) evolves slowly and unevenly, so identifying external financing effects requires unusually large and sustained finance shocks or a sample concentrated in exposed countries.

The gap dynamics and entropy results can be interpreted through a zero-entropy lens. In the benchmark of convergence, the cross-country distribution of gaps becomes increasingly concentrated around the target (low dispersion and low uncertainty). In that limit, the Shannon entropy of the distribution falls because mass accumulates in a narrow range of outcomes. Conversely, when progress is uneven and the distribution remains spread out or bifurcated, entropy stays elevated. High persistence slows convergence within countries, and heterogeneous shocks and policies sustain dispersion across countries.

Table 4 reports estimate from dynamic panel models of the under-five mortality gap that include both country and year fixed effects. In all specifications, the dependent variable is the gap $g_{i,t}$ and the key regressor is its own lag, capturing the persistence of deviations from the MDG target. The "FE baseline" model includes only the lagged gap and fixed effects, while the "FE aid-aug" model additionally incorporates lagged Chinese official finance in logarithmic form.

Table 4. Dynamic models with country and year fixed effects.

Model	φ (lagged gap)	γ (log lagged aid)	R-squared	Observations
FE baseline	0.655		0.952	5520
FE aid-aug	0.655	−0.002	0.952	5520

Notes: Dependent variable: under-five mortality gap. The following equations are estimated: $g_{it} = \alpha g_{i,t-1} + \mu_i + \lambda_t + \epsilon_{it}$ and the third column is estimated as:
$g_{it} = \alpha g_{i,t-1} + \beta \text{ChinaAid}_{it} + \mu_i + \lambda_t + \epsilon_{it}$.

The estimated persistence parameter ϕ is 0.655 in both models, indicating a substantial degree of inertia in the under-five mortality gap even after controlling for time-invariant country characteristics and common shocks. This value is lower than in the pooled specifications without fixed effects, consistent with some of the apparent persistence in the raw data being driven by between-country heterogeneity and global trends rather than purely within-country dynamics. Nonetheless, a coefficient of about 0.65 still implies that gaps close only gradually over time. The coefficient on lagged Chinese finance in the aid-augmented fixed-effects model is approximately −0.002, which is numerically tiny and statistically negligible by conventional standards. The inclusion of this term has virtually no effect on the estimated persistence and does not improve model fit: the R-squared remains at 0.952 and the number of observations is unchanged at 5,520 in both specifications. Overall, Table 4 indicates that, once country and year fixed effects are introduced, Chinese official finance does not materially influence the dynamic evolution of MDG gaps.

Table 5 examines the robustness of the entropy-based characterization of the global distribution of MDG gaps to alternative binning schemes. For each year from 2000 to 2004, Shannon entropy is computed for the cross-country distribution of gaps using three different numbers of bins: 5, 10, and 20[7]. As expected, entropy levels are higher when more bins are used, reflecting the larger maximum attainable entropy under finer partitions of the support. However, the qualitative time pattern is very similar across all three binning schemes. In each case, entropy declines only modestly over the sample period and remains at comparatively high levels, even as average gaps fall.

Table 5. Entropy robustness to alternative binning schemes.

Year	5 bins	10 bins	20 bins
2000	1.179	1.825	2.464
2001	1.179	1.831	2.48
2002	1.172	1.814	2.47
2003	1.154	1.802	2.466
2004	1.152	1.807	2.459

Notes: Shannon entropy of the cross-country gap distribution computed using alternative numbers of bins. For each year, the table reports entropy under three binning schemes: 5, 10, and 20 bins.

This pattern implies that the cross-country distribution of MDG gaps does not rapidly collapse toward the target; instead, mass remains spread across a wide range of gap values. The global "macrostate" of gaps therefore stays high-entropy and dispersed over time, rather than converging to a low-entropy configuration concentrated near the MDG goal.

Taken together with the dynamic panel results in Table 4, these findings reinforce the interpretation that MDG gaps exhibit slow convergence and that their cross-country dispersion remains substantial, with Chinese official finance playing at most a minimal role in altering these aggregate dynamics.

5.5. Integrating the Zero-Entropy Framing

Across specifications, MDG-related gaps in low- and middle-income countries behave as highly persistent processes. Accordingly, any relationship between Chinese development finance and under-five mortality rates is best interpreted as shifting the trajectory of gap closure rather than generating discrete jumps.

Our "Zeno" terminology is intentionally interpretive: it is empirically equivalent to a standard AR(1) partial-adjustment process with geometric decay. When gaps follow an AR(1) law of motion with $0 < \rho < 1$, expected progress each period is proportional to the remaining shortfall, implying mechanically slower improvement as targets are approached. This is the same adjustment structure emphasized in the conditional convergence literature (e.g., Barro and Sala-i-Martin-

[7]We compute entropy from binned gap distributions. Because entropy depends on how the bins are defined, we check robustness to alternative binning schemes, including different numbers of equal-width bins and standard data-driven rules for bin width. The main patterns are unchanged.

style convergence regressions); "Zeno" is a re-framing of this established dynamic, not a new econometric model.

Consistent with this perspective, event-study estimates around the first observed Chinese finance "onset" show gradual rather than discontinuous changes. Using a staggered-adoption-robust design (Sun-Abraham), poverty exhibits the clearest evidence of post-onset improvement, with event-time coefficients tending to become more negative after onset, though effects vary across cohorts and horizons and uncertainty grows at longer lags[8].

Finally, the entropy lens helps explain why countries can improve on average without becoming more similar. Differences across countries remain large, so the gaps can stay widely spread even while the average gap gets smaller. This fits our "Zeno-plus-entropy" view, where progress is slow and gradual and the overall distribution does not tighten much toward a common level of target attainment.

6. Conclusion and Policy Implications

This paper finds that Millennium Development Goal (MDG) gaps in child mortality, and extreme poverty in low-income and middle-income countries decrease asymptotically rather than reaching the targets within a fixed time. Each year, countries reduce a consistent portion of the remaining gap, but full achievement is never realized. Despite declining average gaps, high Shannon entropy shows persistent variation between countries, indicating that progress remains uneven.

The combined effects of Zeno dynamics and high-entropy outcomes mean that fully synchronized global achievement of the MDGs is highly unlikely. Chinese development finance may help reduce poverty and improve social indicators, but it works within this system rather than changing its fundamental dynamics or entropy; it can accelerate progress slightly without altering overall patterns.

The core message is simple. The standard MDG story assumes that, with enough effort, the world can hit a fixed set of time-bound targets. The evidence here suggests otherwise. MDG gaps behave like Zeno's paradox: they keep shrinking, but only by fractions, not in a final jump to zero. At the same time, high Shannon entropy shows that country outcomes stay messy and dispersed rather than converging to a clean end-state where almost everyone meets the targets.

Policy should stop treating "zero gaps by year T" as a realistic organizing principle. In a Zeno-style, high-entropy world, eradication deadlines are rhetoric, not empirics. A more honest framework focuses on continuous improvement: faster fractional reductions in gaps and lower dispersion across countries, rather than pass/fail judgments against fixed thresholds.

Because the global system is structurally heterogeneous, insisting on universal, synchronized target attainment by a common date wastes effort and creates artificial "off-track" narratives. Policy should instead embrace diversity in country trajectories and support steady, resilient progress along those paths.

[8]We present these results in **Appendix**.

Finally, theory and empirics in development economics need to reflect these Zeno and entropy features. Models that build in linear, finite-time convergence to externally set goals are inconsistent with the data. A more realistic paradigm treats progress as continuous, uneven, and asymptotic, morally urgent and measurable, but not well described by claims that the MDGs, or their successors, can literally be "achieved" everywhere by a specific deadline.

Declaration

The opinions expressed in this study are solely those of the authors. Any remaining errors or omissions remain their sole responsibility. The authors did not receive any financial assistance for completion of this study.

Conflicts of Interest

The author declares no conflicts of interest regarding the publication of this paper.

References

AidData (2017). *AidData's Global Chinese Official Finance Dataset, Version 1.0 (Dataset).* AidData at William & Mary. https://www.aiddata.org/data/chinese-global-official-finance-dataset-version-1-0

Alden, C. (2007). *China in Africa.* Zed Books Ltd. https://doi.org/10.5040/9781350219106

Arthur, W. (1994). *Increasing Returns and Path Dependence in the Economy.* University of Michigan Press. https://doi.org/10.3998/mpub.10029

Arthur, W. B. (1999). Complexity and the Economy. *Science, 284,* 107-109. https://doi.org/10.1126/science.284.5411.107

Barro, R. J., & Sala-i-Martin, X. (1992). Regional Growth and Migration: A Japan-U.S. Comparison. *Journal of the Japanese and International Economies, 6,* 312-346. https://www.sciencedirect.com/science/article/abs/pii/088915839290002L

Becker, G. S., Philipson, T. J., & Soares, R. R. (2005). The Quantity and Quality of Life and the Evolution of World Inequality. *American Economic Review, 95,* 277-291. https://doi.org/10.1257/0002828053828563

Beinhocker, E. D. (2006). *The Origin of Wealth: Evolution, Complexity, and the Radical Remaking of Economics.* Harvard Business School Press.

Boone, P. (1996). Politics and the Effectiveness of Foreign Aid. *European Economic Review, 40,* 289-329. https://doi.org/10.1016/0014-2921(95)00127-1

Bouchaud, J. (2013). Crises and Collective Socio-Economic Phenomena: Simple Models and Challenges. *Journal of Statistical Physics, 151,* 567-606. https://doi.org/10.1007/s10955-012-0687-3

Bräutigam, D. (2009). *The Dragon's Gift: The Real Story of China in Africa.* Oxford University Press.

Bräutigam, D. (2011). Aid 'with Chinese Characteristics': Chinese Foreign Aid and Development Finance Meet the OECD-DAC Aid Regime. *Journal of International Development, 23,* 752-764. https://doi.org/10.1002/jid.1798

Brautigam, D. (2020). A Critical Look at Chinese 'Debt-Trap Diplomacy': The Rise of a Meme. *Area Development and Policy, 5,* 1-14.

https://doi.org/10.1080/23792949.2019.1689828

Bräutigam, D., & Gallagher, K. P. (2014). Bartering Globalization: China's Commodity-Backed Finance in Africa and Latin America. *Global Policy, 5,* 346-352. https://doi.org/10.1111/1758-5899.12138

Brazys, S., & Vadlamannati, K. C. (2021). Aid Curse with Chinese Characteristics? Chinese Development Flows and Economic Reforms. *Public Choice, 188,* 407-430. https://www.jstor.org/stable/48750085

Buckley, P. J., Clegg, L. J., Cross, A. R., Liu, X., Voss, H., & Zheng, P. (2007). The Determinants of Chinese Outward Foreign Direct Investment. *Journal of International Business Studies, 38,* 499-518. https://doi.org/10.1057/palgrave.jibs.8400277

Burnside, C., & Dollar, D. (2000). Aid, Policies, and Growth. *American Economic Review, 90,* 847-868. https://doi.org/10.1257/aer.90.4.847

Clemens, M. A., Radelet, S., Bhavnani, R. R., & Bazzi, S. (2011). Counting Chickens When They Hatch: Timing and the Effects of Aid on Growth. *The Economic Journal, 122,* 590-617. https://doi.org/10.1111/j.1468-0297.2011.02482.x

Cohen, J., & Dupas, P. (2010). Free Distribution or Cost-Sharing? Evidence from a Randomized Malaria Prevention Experiment. *Quarterly Journal of Economics, 125,* 1-45. https://doi.org/10.1162/qjec.2010.125.1.1

Cover, T. M., & Thomas, J. A. (2006). *Elements of Information Theory* (2nd ed.). Wiley-Interscience.

Doob, J. L. (1953). *Stochastic Processes.* Wiley.

Downs, E. S. (2011). *Inside China Inc China Development Bank's Cross-Border Energy Deals.* Brookings Institution. https://www.brookings.edu/wp-content/uploads/2016/06/0321_china_energy_downs.pdf

Dreher, A., Fuchs, A., Hodler, R., Parks, B. C., Raschky, P. A., & Tierney, M. J. (2019). African Leaders and the Geography of China's Foreign Assistance. *Journal of Development Economics, 140,* 44-71. https://doi.org/10.1016/j.jdeveco.2019.04.003

Dreher, A., Fuchs, A., Hodler, R., Parks, B., Raschky, P. A., & Tierney, M. J. (2016). Aid on Demand: African Leaders and the Geography of China's Foreign Assistance. *SSRN Electronic Journal, 118,* 50-65. https://doi.org/10.2139/ssrn.2900351

Dreher, A., Fuchs, A., Parks, B., Strange, A. M., & Tierney, M. J. (2018). Apples and Dragon Fruits: The Determinants of Aid and Other Forms of State Financing from China to Africa. *International Studies Quarterly, 62,* 182-194. https://doi.org/10.1093/isq/sqx052

Dreher, A., Nunnenkamp, P., & Thiele, R. (2008). Does Aid for Education Educate Children? Evidence from Panel Data. *The World Bank Economic Review, 22,* 291-314. https://doi.org/10.1093/wber/lhn003

Easterly, W. (2009). How the Millennium Development Goals Are Unfair to Africa. *World Development, 37,* 26-35. https://doi.org/10.1016/j.worlddev.2008.02.009

Easterly, W., Levine, R., & Roodman, D. (2004). Aid, Policies, and Growth: Comment. *American Economic Review, 94,* 774-780. https://doi.org/10.1257/0002828041464560

Farmer, J. D., & Foley, D. (2009). The Economy Needs Agent-Based Modelling. *Nature, 460,* 685-686. https://doi.org/10.1038/460685a

Gelpern, A., Horn, S., Morris, S., Parks, B., & Trebesch, C. (2022). *How China Lends: A Rare Look into 100 Debt Contracts with Foreign Governments.* https://www.piie.com/publications/working-papers/how-china-lends-rare-look-100-debt-contracts-foreign-governments

Gertler, P. (2004). Do Conditional Cash Transfers Improve Child Health? Evidence from Progresa's Control Randomized Experiment. *American Economic Review, 94,* 336-341. https://doi.org/10.1257/0002828041302109

Henderson, J. V., Storeygard, A., & Weil, D. N. (2012). Measuring Economic Growth from Outer Space. *American Economic Review, 102,* 994-1028. https://doi.org/10.1257/aer.102.2.994

Horn, S., Reinhart, C. M., & Trebesch, C. (2021). China's Overseas Lending. *Journal of International Economics, 133,* Article 103539. https://doi.org/10.1016/j.jinteco.2021.103539

Isaksson, A., & Kotsadam, A. (2018). Chinese Aid and Local Corruption. *Journal of Public Economics, 159,* 146-159. https://doi.org/10.1016/j.jpubeco.2018.01.002

King, K. (2013). *China's Aid and Soft Power in Africa.* Boydell & Brewer.

Maasoumi, E. (1993). A Compendium to Information Theory in Economics and Econometrics. *Econometric Reviews, 12,* 137-181. https://doi.org/10.1080/07474939308800260

Miguel, E., & Kremer, M. (2004). Worms: Identifying Impacts on Education and Health in the Presence of Treatment Externalities. *Econometrica, 72,* 159-217. https://doi.org/10.1111/j.1468-0262.2004.00481.x

Mishra, P., & Newhouse, D. (2009). Does Health Aid Matter? *Journal of Health Economics, 28,* 855-872. https://doi.org/10.1016/j.jhealeco.2009.05.004

Pritchett, L. (1997). Divergence, Big Time. *Journal of Economic Perspectives, 11,* 3-17. https://doi.org/10.1257/jep.11.3.3

Quah, D. (1993). Empirical Cross-Section Dynamics in Economic Growth. *European Economic Review, 37,* 426-434. https://doi.org/10.1016/0014-2921(93)90031-5

Quah, D. T. (1996). Twin Peaks: Growth and Convergence in Models of Distribution Dynamics. *The Economic Journal, 106,* 1045-1055. https://doi.org/10.2307/2235377

Rajan, R. G., & Subramanian, A. (2008). Aid and Growth: What Does the Cross-Country Evidence Really Show? *Review of Economics and Statistics, 90,* 643-665. https://doi.org/10.1162/rest.90.4.643

Roodman, D. (2007). The Anarchy of Numbers: Aid, Development, and Cross-Country Empirics. *The World Bank Economic Review, 21,* 255-277. https://doi.org/10.1093/wber/lhm004

Sanyal, P., & Babu, S. (2012). Aid Effectiveness and Capacity Development: Implications for Economic Growth in Developing Countries. *Modern Economy, 3,* 567-577. https://www.scirp.org/journal/paperinforcitation?paperid=23326

Shannon, C. E. (1949). A Mathematical Theory of Communication. In C. E. Shannon, & W. Weaver (Eds.), *The Mathematical Theory of Communication* (pp. 1-55). University of Illinois Press.

Sims, C. A. (2003). Implications of Rational Inattention. *Journal of Monetary Economics, 50,* 665-690. https://doi.org/10.1016/s0304-3932(03)00029-1

Strange, A. M., Parks, B., Tierney, M. J., Fuchs, A., Dreher, A., & Ramachandran, V. (2017). *Tracking Underreported Financial Flows: China's Development Finance and the AidData Methodology.* AidData, William & Mary.

Taylor, I. (2006). *China and Africa Engagement and Compromise.* Routledge.

Theil, H. (1967) *Economics and Information Theory.* North-Holland Publishing Company.

Wellner, L., Dreher, A., Fuchs, A., Parks, B., & Strange, A. (2022). Can Aid Buy Foreign

Public Support? Evidence from Chinese Development Finance. *SSRN Electronic Journal,* 78 p. https://doi.org/10.2139/ssrn.4070236

Yakovenko, V. M., & Rosser, J. B. (2009). *Colloquium*: Statistical Mechanics of Money, Wealth, and Income. *Reviews of Modern Physics, 81,* 1703-1725. https://doi.org/10.1103/revmodphys.81.1703

Appendix: Robustness Checks and Mathematical Details of Sun-Abraham Estimation for the Poverty Gap at $1.90/Day

A.1. Outcome, Treatment, and Event Time

We have also estimated the poverty gap index measured at the $1.90 per day international poverty line. Let, i index countries and t denote index years.

Let us also define the following:

Outcome: Y_{it} is the poverty gap index at $1.90/day for country i in year t.

Treatment timing: Let G_i be the first year in which country i is treated (the event year). For never-treated countries, define $G_i = \infty$.

Event time: Define event time

$$k \equiv t - G_i \tag{A1}$$

We normalize the coefficient at $k = -1$ to zero (the omitted/reference period).

Treatment Indicator: Let us define

$$D_{it} \equiv 1[t \geq G_i] \tag{A2}$$

A.2. Why Sun-Abraham Estimation Is Done Rather than Two-Way Fixed Effect (TWFE)

As Chinese finance is associated with progress in some contexts, but where the deeper structure is still "Zeno-plus-entropy": persistent dynamics that make rapid catch-up difficult, and a global distribution that does not mechanically collapse toward uniform target attainment, we use a staggered adoption (different G_i across countries) and potentially heterogeneous treatment effects. The Sun-Abraham addresses this by estimating cohort-specific dynamics using appropriate not-yet treated (and/or never treated) comparisons and then aggregate this across countries.

A.3. Estimation Equations

Conceptual Event-Study Specifications

A conceptual event-study regression model is of the following form:

$$Y_{it} = \alpha_i + \lambda_t + \sum_{k \neq -1} \beta_k 1[t - G_i = k] + X'_{it}\gamma + \varepsilon_{it} \tag{A3}$$

where, α_i are country fixed-effects, λ_t are year fixed-effects, X'_{it} is a vector of time-varying controls, and ε_{it} is the error term.

A.4. Sun-Abraham Cohort-Specific Event-Time Interactions

Sun-Abraham estimates cohort-by-event-time effects. The following equation was estimated:

$$Y_{it} = \alpha_i + \lambda_t + \sum_{g \in G} \sum_{k \neq -1} \theta_{gk} 1[G_i = g] \cdot 1[t - g = k] + X'_{it}\gamma + \varepsilon_{it} \tag{A4}$$

where, G is the set of treated cohorts (event years). The reported $\widehat{ATT(k)}$ is an aggregation of $\widehat{\theta_{gk}}$ across cohorts g using valid comparisons (not-yet-treated and/or never treated units)[9].

In Figure A1 and Table A1, we present Sun-Abraham event-study estimates of the effect of aid per capita on the \$1.90/day poverty gap. Event time k is defined relative to the first year of treatment, with $k=-1$ normalized to zero; negative k values are pre-treatment leads and nonnegative k values are post-treatment lags.

Table A1. Event-study point estimates for poverty-gap at \$1.90/day.

Event time k	Baseline	Trim support	Alt event (p75)	Placebo (+5y)
−5	3.706	3.85	1.4	−7.598
−4	6.721	1.183	2.285	−0.572
−3	6.575	0.456	−1.267	−1.979
−2	3.952	0.506	3.356	−1.283
−1	0.0	0.0	0.0	0.0
0	−7.152	−2.147	1.149	−8.157
1	−4.586	−2.258	0.015	−5.596
2	−2.351	−2.34	−0.896	−7.999
3	−6.788	−3.323	−3.772	−10.453
4	−4.185	−3.12	2.33	−6.321
5	−10.905	−3.98	−1.824	−8.529
6	−6.906	−4.367	−1.358	−7.241
7	−13.159	−4.533	−6.732	−6.99
8	−13.177	−5.96	−7.251	−13.38
9	−9.13	−6.374	−4.803	−11.508
10	−10.897	−6.732	−4.585	−10.869

Notes: Coefficients are normalized to 0 at $k=-1$. These are point estimates. To report the confidence intervals and test pre-trends, one needs to re-estimate the model while saving the variance-covariance matrix for event-time coefficients (typically, these are standard errors at the country level).

The baseline profile turns negative immediately after treatment and remains negative for most post-treatment horizons, consistent with a sustained reduction in the poverty gap following treatment.

[9] $\widehat{\theta_{gk}}$ is the estimated *event-time effect* k periods relative to treatment for the cohort g that is first treated in year g. Intuitively, it compares the outcome for cohort-g units at calendar time $t=g+k$ to their outcome just before treatment, and then subtracts the corresponding change for a set of *valid controls* at the same calendar time. "Valid" means units that have not yet been treated by time t (and, if available, never-treated units), so the comparison is not contaminated by already-treated units. Thus, $\widehat{\theta_{gk}}$ is a difference-in-differences estimate constructed separately for each treated cohort and each relative time k; the reported $\widehat{\theta_k}$ these cohort-specific estimates across g using only those valid comparisons.

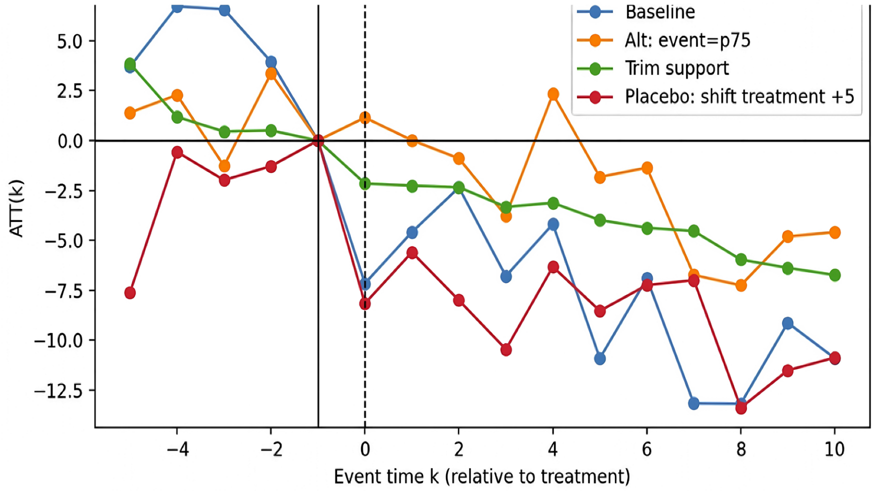

Figure A1. Event-study coefficients on log(1 + aid) on poverty-gap at \$1.90/day. Source: Computed by the Authors using Aiddata.

However, the baseline estimates exhibit sizable pre-treatment deviations from zero, and a +5-year placebo shift also generates pronounced movements in the estimated effects. This raises concerns about the credibility of the parallel-trends assumption and indicates sensitivity to timing and support restrictions.

We therefore interpret the post-treatment declines as suggestive evidence of improved poverty outcomes, while emphasizing robustness checks (alternative event definitions, support trimming, and placebo timing) as essential for assessing credibility.