

Age-Profile Estimates of the Relationship Between Economic Growth and Child Health

Anaka Aiyar*

Joseph R. Cummins†

August 3, 2018

Abstract

Based largely on the analysis of the same underlying data, recently published papers have presented estimates that the association between economic growth and child undernutrition in developing countries is either strong and robust, or weak to nonexistent. We provide clarity on both the magnitude of the association and the underlying econometric problem. Focusing on child growth faltering as a process that unfolds over the first several years of life, we provide new evidence tracing out the relationship between macroeconomic trends and the trajectory of child growth through age 5. Using two novel regression models that each harness different kinds of within- and between-country variation, and data on over 600,000 children from 38 countries over more than 20 years, our estimates of the association are small but precise, and are consistent across both estimators. We estimate that a 10% increase in GDP around the time of a child's birth is associated with a decrease in the rate of loss of HAZ of about 0.002sd per month over the first two years of life. This generates a cumulative effect of around 0.04sd by a child's third birthday and the magnitude of the correlation largely persists through age 5. Our models are derived from both economic and biological theory and provide a new empirical framework for researchers interested in investigating the ecological-level determinants of child growth.

*Cornell University

†University of California, Riverside

1 Introduction

Economic growth is valuable insofar as it improves human wellbeing, and long-term economic development has clearly generated incredible benefits for millions of people. Yet even with steady growth in the global economy over the last several decades, child physical growth stunting induced by chronic under-nutrition and heavy disease burden still affects over 150 million children worldwide¹. Stunted growth in childhood leads to decreased wages and worsened health outcomes in later life, and contributes to the inter-generational persistence of poverty (Behrman et al., 2009; Hoddinott et al., 2008). Stunted growth is both a marker of the cumulative effects of chronic nutrient deficiency and poor health on physical development and a physical manifestation of stunted human potential.

Despite the major differences in stunting rates and mean height-for-age z-score (HAZ) between children in developed and developing countries, several recent papers have argued that there is a surprisingly weak correlation between medium-term economic growth and nutritional status within less-developed countries (Subramanyam et al., 2011; Vollmer et al., 2014). These papers stand in contrast to previous work that estimated relatively robust effects of macroeconomic conditions on child anthropometric outcomes (Smith and Haddad, 2002; Haddad et al., 2003; Klasen, 2008; Harttgen et al., 2013). A number of both critical and positive follow-up comments demonstrated that minor modeling or weighting choices or effect size re-scalings can create the appearance of either agreement or disagreement between the various estimates (Alderman et al., 2014; O’Connell and Smith, 2016; Bershteyn et al., 2015; Singh, 2014)². The overall weight of the evidence appears to indicate substantial heterogeneity in the relationship across countries and years and a relatively small average effect that requires more precise estimators to statistically differentiate from zero. Our analysis provides a step forward on precision and our models constitute a toolkit for researchers

¹UNICEF data, Date Accessed: 07/12/2016: <http://data.unicef.org/nutrition/malnutrition.html>

²Appendix Table 1 lists previous papers estimating the association between economic growth and various measures of child nutritional status

interested in investigating heterogeneity.

We argue that the conflicting empirical evidence is explained in part by failures of previous empirical strategies to properly model the underlying relationship in the data and the underlying economic-biological age-dynamics of child growth. We provide two empirical strategies derived from human capital accumulation theory that overcome these limitations. First, we employ a newly developed framework of survey-level graphical and regression analysis focused on the rate of loss of HAZ over the first two years of a child's life that allows us to directly estimate relative growth trajectories. Second, we derive a set of fixed-effects regressions from household dynamic human capital accumulation theory that specifically address the process of growth faltering. Both models identify the coefficients of interest exploiting both within- and between-country variation in GDP and HAZ, and our coefficients are interpretable in the framework of general spatio-temporal fixed effects models (e.g. linear models with year and region fixed effects).

The two models, though conceptually and statistically quite distinct, produce similar estimates. Using our survey-level outcome and regression model, we estimate that exposure to a 10% increase in GDP during early childhood is associated with a decrease in the rate of loss of HAZ relative to the World Health Organization (WHO) reference median by 0.002 sd/month. This adds up to an effect of around 0.04 - 0.05sd by the child's third birthday. Similarly, our age-profile fixed-effect model estimates a statistically, biologically and economically insignificant association before a child's first birthday, but that by a child's third birthday a 10% increase in GDP is associated with a cumulative effect of 0.03-0.04sd, which then persists through age 5.

Beyond their desirable econometric properties, our models are inspired by, and interpretable within, the framework of dynamic health capital accumulation theory. We interpret HAZ, a measure of cumulative health inputs since birth, as carrying information on the entire history of an optimally chosen stream of inputs up to the moment the child is measured. We model

changes in GDP as altering the optimal stream of investments in child health inputs, thus affecting a child’s physical growth trajectory. The age-profile models we develop are fully capable of identifying the entire set of potential age-heterogeneities across child development predicted by the abstract model.

Our goals in this work are two-fold. First, we argue that we provide more meaningful and interpretable estimates of the relationship between medium-term economic growth and child nutritional status than have previously been available. Our point estimates are small, but our confidence intervals exclude both zero and extremely large effects. Our estimates can also be understood not simply as an average treatment effect over a non-stationary population (as with previous estimates) but as the reduced-form net-effects of an early life change to the stream of child health inputs propagating over a child’s development. Second, we demonstrate how an econometric framework focused on the HAZ-age profile, instead of simply mean HAZ or stunting probability, can allow for both more precise and more nuanced estimates, regardless of the covariate of interest. Models such as those developed here may allow researchers to better trace out how inputs and investments at different ages differentially affect child development.

2 Background

2.1 Data

Any attempt to quantify a meaningful correlation between macroeconomic factors and child development outcomes requires two independent sources of information: a sample of health outcome measurements from children, and a series of measurements of the macroeconomic conditions experience by the child as they grew up. Table A1, which summarizes recent papers estimating correlations of GDP and child height, shows that two sources dominate for supplying GDP time-series data: the Penn World Tables and the World Bank Indicators.

One source, the DHS, dominates as choice of source for child anthropometric data.

2.1.1 The Demographic and Health Surveys

Our child-level dataset of outcomes and covariates was generated by appending data from 126 demographic health surveys (DHS) from 38 countries surveyed between the years 1986 and 2013 (I C F, 2011). This constitutes all of the DHS surveys meeting our basic requirement that a country had at least 2 rounds of DHS collected that included HAZ data on children from ages 0-5 years. Each individual DHS round is a large-scale, multi-stage cluster sample survey used to gather health, demographic and socioeconomic information on women (aged 15-49 years), children (aged 0-5 years) and their households. The surveys are generally conducted every few years in a given country, generating a continuity in child cohorts given the 5-year age-cutoff in the sample. The surveys can be weighted to be nationally representative. DHS data is originally designed to be aggregated up to country or region levels, and much of the previous literature takes that approach and regresses country-level indicators on country level economic indicators (Smith and Haddad, 2002; Haddad et al., 2003; Headey, 2013). Detailed information on DHS sampling design and implementation can be found on the DHS website and varies minority from country to country and year to year³.

Our estimation sample includes children between the ages of 0 to 60 months with valid HAZ scores, a measure of GDP and all included covariates. Height-for-age Z-score (HAZ) is an age- and gender-normalized measure of child height relative to the median height of a population of well-nourished and healthy children. The anthropometric standards against which children in the DHS are compared is provided by the WHO⁴. As per the WHO recommendations, anthropometry scores between -6 and 6 are considered valid and only children with valid scores have been included in the analysis (World Health Organization,

³<https://dhsprogram.com/data/data-collection.cfm>

⁴The 2005 WHO standards reflect the highest potential for physical growth and human development for children living both in developed and developing countries.

2006)⁵. Our regressions also include only those children with complete information on GDP per capita, sex and age, their mother’s education level and age and type of residence (urban or rural). After following this inclusion criteria, our estimation sample consists of a total of 685,075 children in 38 countries. Figure 1 shows the sample selection criteria and loss of sample size at each stage. Relatively few observations (1.4%) are lost to lack of GDP data or covariates.

2.1.2 World Bank Indicators

GDP per capita time-series data for individual countries is available in the World Development Indicators Series of the World Bank⁶. We use GDP per capita in purchasing power parity (PPP) terms comparable to 2005 U.S dollars for the analysis. We use World Bank data instead of the Penn World Tables data favored by Vollmer et al. (2014) for reasons explicated in O’Connell and Smith (2016) to do with comparability within countries over time.

A time-series of (ln) GDP is a canonical example of serially-correlated outcomes. Leaving a detailed discussion of the econometric ramifications of this reality to later sections, it is important to assuage two obvious concerns. First, the variation used to identify the magnitude of any particular coefficient in our regression models does not come from comparing children born in the same country one year apart. Our method compares only children of the same age. Since the DHS is only conducted every several years within one country, our (age-specific) within-country variation in GDP covers 3 or 4 semi-evenly spaced observations covering 5 to 20 years worth of economic growth. Second, our appending of multiple country’s GDP turns the individual country time-series data into a panel, allowing for non-parametric adjustments for arbitrary time-trends common across the world.

⁵The DHS multiplies standard WHO scores by 100, and we maintain this convention to make coefficients more easily interpretable, so in practice our analysis includes those that range from -600 to 600 in the DHS surveys

⁶accessed from <http://databank.worldbank.org/>; Dec 2014

2.1.3 Sample Summary

Summary information on the countries used, the survey years, characteristics of the households and children and the outcome and GDP measure can be found in Table 1, which contains the entire (small) set of covariates used in our analysis. The mean age of the children is about 29 months and the sample is evenly split by gender. 36% of the children live in urban areas. Mothers are on average 29 years old, 36% of them have no education and 35% of them have at least primary education. The average GDP per capita experienced by children in our sample is around 721 USD. The average HAZ score for all children in the sample is -144, meaning the average child in our sample is 1.4 standard deviations below the WHO reference for the median healthy and well-nourished child of that age and gender⁷.

We limited the number of covariates in order to maximize the final sample size and limit the potential for bias induced by differential omission from the analysis set. Our analysis files make it possible to easily add or remove covariates in order to test the robustness of the estimates, but in general we find that covariates do not strongly influence the estimated strength of the association conditional on the fixed-effects specification⁸.

2.2 Previous Estimates and Empirical Strategies

Previous estimates of associations between economic growth and child undernutrition differ from our analysis, and each other, in four inter-connected ways: choice of outcome variable, level of aggregation, sources of identifying variation, and timing of GDP merging relative to a child's survey or cohort. In Table A1, we provide a summary of the similarities and differences between the various approaches used in recently published papers.

Regarding timing, every previous study (though not ours) examines the relationship between an anthropometric measures of child nutritional status and values of GDP contemporaneous

⁷Following the DHS, we multiply HAZ by 100 to make units easier to display and read.

⁸Replication code can be found on the Open Science Framework: <https://osf.io/829ny/>

to the year in which the children were measured. This choice of merging strategy leads naturally to the choice of outcome measures based on child weight, since weight is capable of short-term fluctuation whereas height is not.

In the (World Health Organization, 1995) report, an expert committee on child health argued that height-for-age (HAZ), weight-for-age (WAZ), weight-for-height (WHZ) z-scores best reflected the interaction between social determinants of health and the physical development of children. They determined that indicators that use weight can accurately predict malnourishment within a population at a given time. However, since weight is highly responsive to food and nutrition availability in the short-term, weight measures cannot be used to measure the effect of past input streams on the current or future health and productivity of an individual. HAZ scores, on the other hand, can be interpreted as capturing the cumulative effects of the stream of biological inputs over the course of the child’s development. Other work has shown that child HAZ predicts lower productivity in adults (Glewwe and Miguel, 2007; Hoddinott et al., 2008) and can predict two year future mortality risk in young children (World Health Organization, 1995)⁹. Building on these arguments, in our model, we use HAZ to model the impact of changing biological inputs on a long term measure of child’s health and productivity.

2.2.1 Statistically Significant Estimates

One way to estimate the association between economic growth and undernutrition is simply to compare mean HAZ and GDP in survey year from a cross-section of countries. These cross-country studies exploit “between country” variation and implicitly compare the GDP levels of countries J and J’ with their average anthropometric outcomes of interest. These regressions tell us how countries that have grown differently in the past have experienced different health improvement trajectories, but they cannot tell us how growth in some particular country

⁹The authors of the WHO report also cautioned that any study that used HAZ as an outcome would be confounded by its relationship with age if this relationship was not accounted for properly.

affects the nutritional status of children. The model will pick up any effect of GDP on HAZ, but also any effects of anything else that is more conducive to child physical growth in richer as compared to poorer countries. In order to address this problem, the most common strategy in recent work has been to difference-out the time-invariant country-level unobservable characteristics by de-meaning the data within each country using a fixed-effects estimator.

Three previous papers have aggregated survey-level data from panel of country-years to generate within-country estimates, and these studies have estimated relatively strong associations between contemporary GDP and child wasting or undernutrition. The largest estimates are from Smith and Haddad (2002), which finds that a 10% increase in GDP per capita is associated with decreases in wasting rates of 6.3%. Using country and decade fixed effects, Haddad et al. (2003) interpret the effect size as indicating that a 10% increase in GDP per capita is associated with decreases in wasting rates of 1.5pp. Using stunting a dichotomized HAZ measure, Headey (2013) finds that a a 10% increase in GDP per capita is associated with 1.8pp reduction in child stunting.

Only one paper prior to Vollmer et al. (2014) analyzed DHS data at the individual child level to estimate correlations of GDP with HAZ, allowing for both country-level fixed-effects (and thus within-country variation) and for the inclusion of individual-level control variables in the regression equations. That work, Harttgen et al. (2013), estimates that a 10% increase in GDP is associated with a large, 1.5-1.7 log odds ratio reduction in stunting.

Instead of directly replicating these findings, here we provide a graphical representation of the empirical strategy and estimate a best-fit line that approximates these methods. Figure 2 graphs survey-round mean HAZ across survey-time log GDP/capita, with each survey round mean HAZ (now a single observation) weighted to be nationally representative. The upper panel of Figure 2 shows how mean HAZ for a country correlates with the level of GDP in the survey year. The correlation in the raw data is statistically significant and the coefficient is of

a similar magnitude to our final estimates - a 0.1 log point change in GDP is associated with a 0.03 sd increase in HAZ. The bottom panel in the graph shows changes in HAZ against changes in GDP, and the point estimates are similarly sized and also precisely estimated.

While these estimates seem relatively congruous with the previous literature, there is reason to be concerned about the comparisons made to identify the coefficients of interest. Exploiting “within country” variation helps rid regressions of certain kinds of bias arising from non-GDP related differences across countries, but comes with a different set of concerns. If we consider only country-level macro-economic conditions, then within a country the only variation in GDP comes from across birth cohorts - from children born in different years and thus exposed to the stream of GDP realizations at different points in their development. The sacrifice to be made to estimate within-country effects is that one must choose a parametric specification for the effects of secular improvement over time, and with only a few survey rounds this becomes a potentially insurmountable econometric problem. Fitting even a linear trend through 3 or 4 values of GDP per fixed-effect (the sample size of the age-country-cohort group) can lead to over-fitting and the loss of a large fraction of meaningful variation in GDP. Similarly, ignoring the age structure of the outcome variable and the non-linear effects of secular improvements on health across age-time leads to model-misspecification that can potentially generate the bias described in Cummins (2013).

2.2.2 The Vollmer et al. Estimates

Two recent papers provided important reasons to worry that some of those within-country estimates may be misleading. Subramanyam et al. (2011) finds that province level per capita GDP does not have any impact on the nutritional status of children in India. Vollmer et al. (2014) uses cross country height and weight scores from 126 DHS surveys and find for the most part statistically, economically and biologically insignificant associations. These estimates comprise two of the three studies in Table A1 that employed individual level data,

and they are the only ones that include both country and time (survey year) fixed effects.

The impact of Subramanyam et al. (2011) and Vollmer et al. (2014) on the academic debate was large, sparking a series of both positive and less positive responses in various journals (Singh, 2014; Alderman et al., 2014; Bershteyn et al., 2015; Joe et al., 2016; O’Connell and Smith, 2016; Lange and Vollmer, 2017). This was not simply because their estimates were somewhat out of line with previous research, but because these out-of-line estimates were the first to apply fixed-effect models that could identify the effect of within-country changes in GDP while still controlling non-parametrically for secular time trends. These models make assumptions about the effects of place and time that allow for group/region level effects that are persistent across time, and simultaneously allow for arbitrary secular “time” trends that are common across groups. Whereas pure within-country models compare changes in GDP to changes in HAZ, country-year fixed-effect models compare how much more (or less) change in HAZ occurred in high growth countries relative to low growth countries. Unlike in pure within-country models, differential changes in HAZ that are uncorrelated with differential changes in GDP are then ascribed to the unobserved time effect and not to GDP growth itself.

3 The Age-Profile Framework

While weighting, scaling, sample selection and a number of other modeling choices explain some of the variation in published estimates, we argue these are secondary concerns to a more foundational problem. That foundational problem is the econometric modeling of child nutritional status as an age-invariant condition, inherently ignoring the biological process of growth faltering that unfolds over the first years of a child’s life. This conceals from the analyst potential insights into the economics of health capital accumulation by aggregating over the underlying age-dynamics of interest. Our argument is not that the Vollmer et al methods, or any of the other methods employed in previous work, are wrong. Our argument

is that those estimates are both difficult to interpret in relation to the world and insufficiently precise to answer the question at hand.

3.1 Empirical Motivation

To represent the population-level association between HAZ and GDP in a manner more true to the process of growth faltering, we present the HAZ-age profiles from our sample countries in Figure 3, graphing mean HAZ across child age in months. In the top panel, we aggregate countries based on being above or below median GDP for countries in our sample¹⁰. A very clear pattern emerges. Children in both groups of countries start at similar HAZ at birth (just below 0) and then grow more slowly than the children in the well nourished, healthy WHO reference group. Mean child HAZ is then essentially constant from age 2 up to a child’s fifth birthday (Rieger and Trommlerová, 2016; Victora et al., 2010).¹¹

The second key insight is taken from the bottom panel of Figure 3, which graphs HAZ-age profiles by individual country, color coded from highest decile (magenta) to lowest (cyan) GDP/cap. This country-level graph demonstrates the stability of a key aspect of the relationship between HAZ and GDP - it is the rate of loss of HAZ, and not HAZ at birth, that primarily drives the differences in the HAZ-age profiles across GDP levels. At the left edge of the graph, poorer and richer countries are intermingled, with a slight tendency for poorer countries to have a lower intercept (interpretable as projected length-at-birth z-score). However, by the age of 4, the HAZ-age profiles have essentially sorted themselves along GDP rank. This is the result of increased severity of growth faltering in the poorer countries. Despite the fact that children in relatively richer and poorer countries were born with similar length, the children in the poorer countries grow much more slowly.

If the defining characteristic of the association between economic conditions and child height

¹⁰We use average GDP over the study period to divide countries into the two groups

¹¹There is some visual evidence of population-level “catch-up” growth, where the gap in HAZ between richer and poorer countries may close slightly before age 5,

is the process of growth faltering, then clearly the use of mean HAZ or stunting as a measure, as compared to age-specific HAZ, is averaging across the exact effect it is hoping to capture. Mean HAZ and the probability of stunting are non-linear and strongly increasing over the first two years of life, and any effect of GDP growth on HAZ or stunting by age 3 is being averaged with much smaller effects for an otherwise similar child measured at 4 months. Limiting the sample to younger children, as done in Subramanyam et al., 2011 and Vollmer et al., 2014, exacerbates the problem.

3.2 The Health Capital Accumulation Perspective

Our perspective conceives the biological process of growth faltering as the output from a dynamic health capital production function with input levels determined by a household-level utility maximization problem, drawing from health capital models such as Grossman (1972) and Becker (1962). Household decision makers optimize over time and have preferences for their own consumption (C_t) and for their children's health (H_t^A). Household's maximize $U_t(C_t, H_t^A)$ subject to a budget constraint where consumption and health inputs are positively priced. The evolution of child height is modeled by an age-specific human capital production function ($f^a()$) that takes inputs from calendar-time health purchases and public goods availability (broadly conceived).

Differences across countries in GDP, and changes within a country, are likely to generate changes in 3 elements of the optimal child health investment decision. First, increases in GDP lead to improvements in the labor market. During times of GDP growth, households are more likely to find employment, and conditional on finding employment, likely to receive more income (Topel, 1999). Second, increases in GDP are likely to increase the provision of public goods. Public goods, broadly construed, work here as a sort of in-kind transfer from the government that pays in child health. Finally, economic growth is strongly correlated with demographic change in a mutually reinforcing causal loop. Increases in economic growth

can affect both the number of children in the household and thus the distribution of the limited household resources to each child, and the availability of maternal energy during pregnancy and lactation that can be transferred to the fetus or breast-feeding child (Walker et al., 2008; Kramer et al., 2016).

The process of optimal input choice and health production iterates continuously over a child's development and throughout their adult lives. In each period the child's height from the previous period enters into the height production function and thus naturally persists to a large extent from period to period; our height last month is the basis for our height today. This insight from health capital theory provides the economic underpinning for interpreting average attained height of children in a population as resulting from the stream of biological inputs experienced over a child's life. It also serves to link the biological conception of human development (Nutrition, Disease, Cellular and Neurological Development, etc.) with the economic conception of dynamic human capital accumulation theory (inputs, shocks to health, health production functions).

The discussion so far highlights model predictions across age-at-measure, while conceptually holding age-of-exposure fixed. But the model makes predictions across a second dimension of child age as well. For a 4 year old, the (contemporary) effect of an input received at age 2 compared to the same input received at age 3 has two different components. First, there are fundamental differences in the human capital production function over child age. Secondly, re-optimization of household decision making in each period means that if unexpected external forces increase child health, households may have incentives to invest differently than they previously would have because the child is already healthier than they anticipated it being. For the purposes of this paper, we focus primarily on tracing out the correlations between GDP at birth as children age, and not on the relative impacts of GDP growth occurring at different points in a child's development. In theory, we would be able to separately, non-parametrically identify changes in GDP at each age-at-exposure by age-at-measurement permutation. However, the serial correlation in GDP (and growth rates) makes

unconstrained/non-regularized estimation of so many parameters impossible in practice. We do, though, present estimates of these parameters and discuss we can be learned and what can be done going forward.

4 Survey-Level Aggregate Model: Rate of Loss of HAZ

The defining feature of the HAZ-age profiles presented in Figure 3 is the rapid and relatively linear decrease of HAZ over the first two years of life. The age profile then becomes essentially flat (or slightly positively inclined) from ages 2 through 5. We define two parameters to characterize this empirical regularity: a) we define α as the *intercept of the HAZ-age profile on the Y-axis*, that is, the implied length-for-age Z-score (LAZ) at birth; b) we define β as the *rate of loss of HAZ from birth to age 2*, that is, how much more slowly children are growing relative to the WHO reference median child (in units of standard deviations of the reference population). As simple as they are, these summary measures provide a relatively complete characterization of the HAZ-age profile over the first two years of life.

4.1 Modeling Rate of Loss of HAZ

We estimate α & β separately for each survey round in each country as an OLS regression of HAZ for child i measured at age A .

$$HAZ_{iA}^s = \alpha^s + \beta_{age}^s * Age_{iA}^s + u_{iA}^s \quad \forall s \in S \quad (1)$$

Equation 1 allows us to estimate $\hat{\alpha}^s$, a country by survey specific estimate of the LAZ at birth and $\hat{\beta}^s$, an estimate of the rate of loss from that initial birth LAZ over the first two years of life. We then take the estimates $\hat{\alpha}^s$ and $\hat{\beta}^s$ and turn them into α_{jy} and β_{jy} , observations from country J in survey year Y , for a second stage regression on the determinants of the

shape of the HAZ-age profile over the first two years of life. We merge this constructed outcome data with a panel of (ln) per capita GDP measures from the World Bank, generating an unbalanced panel of observations at the country-year level that include the parameter estimates of the first-stage regressions and the GDP data from the World Bank panel¹².

The second stage regression takes a form involving some or all of the elements of the fully saturated regression model below, for country J in survey year Y. While each parameter in the first stage regression is estimated from a regression weighted by the probability weights provided by the DHS, the second stage regressions of the parameters on GDP are not weighted, and each survey is thus given an implicit total weight of 1.

$$P_{jy} = \delta.GDP_{jy} + \gamma_j + \lambda_y + \eta_{jy} \quad (2)$$

$\hat{\delta}$ is the estimate of the effect of GDP on the outcome P, either α or β . We interpret $\hat{\delta}$ divided by 1,000 as the effect of a 10% change in GDP on HAZ¹³. Variants of this equation, keeping or dropping different elements, allow us to estimate the effect of GDP growth on the parameters of the HAZ-age profile using fundamentally different types of identifying variation. Without γ_j and λ_y , the equation reduces to the OLS estimate of the association between economic growth and the parameters of the HAZ-age profile (α & β). This version of the regression model treats every observation as independent from the others, as though an observation from Armenia in 2000 can be naively compared to an observation from Zimbabwe in 2010. In that sense, the equation fully exploits both within- and across-country variation, but does not do so in ways that are not immediately interpretable relative to actual changes in macroeconomic conditions in the world.

¹²In this specification, we merge GDP from the survey year with the outcome data. This means that children who are infants in our regressions are being given a measure of GDP associated with their year of birth, but children who are two years old are being given GDPs that they experience at age 2. In the next section we more strictly link GDP measure to year of birth for all children, but in this section we simply note that, given both the relatively small changes in GDP across one or two years, and the high serial-correlation in GDP growth over time within a country, this should make little difference to our estimates.

¹³The calculation above is based on a change of 0.1 in the log of GDP, given that our measure of HAZ is the WHO measure multiplied by 100.

We then specify a country level fixed effects model. By including γ_j and thus de-meaning the outcome variable and the GDP time-series within a country, $\hat{\delta}$ now estimates how changes in GDP within a country over time affect LAZ at birth and the rate of loss of GDP over the first two years. Our identifying variation within a country comes not from one calendar year to the next, but from one survey year to the next, an average time difference of about 6 years. Symmetrically, we might want to focus only on the across-country differences, even if only to understand any potential differences between the OLS results and the within-country estimates. We show this by including λ_y while dropping γ_j . Our temporal fixed-effects here bin surveys into 3 year survey-time bins, leading our between country model to identify $\hat{\delta}$ by averaging correlations across countries at each point in time. ¹⁴.

The single fixed-effects estimates in the preceding paragraph are regression analogues of single-difference estimators. Exploiting both across- and within-country variation allows us to partially address the problems of omitted variables (across-country) and secular trends (within-country). With the inclusion of both λ_y and γ_j , the model now implicitly compares changes in the HAZ-age profile in a country with low growth to changes in the profile in a country with high growth.

4.1.1 Inference

We offer two strategies for estimation of standard errors for $\hat{\delta}$. First, we provide analytic standard errors clustered by country, following standard practice for spatio-temporal fixed-effects models. These standard errors are likely biased towards 0 relative to the true sampling distribution of $\hat{\delta}$, since they do not account for the uncertainty in the left-hand side variables which are themselves just estimates (Elbers et al., 2005). To account for this, we provide a second set of standard errors estimated from a 2-stage bootstrap procedure. In that

¹⁴Technically, it is possible to include individual survey year dummies into the regression. However, since surveys for different countries occur at different times, this forces the year dummies to identify off a small set of countries in each potential survey year. We thus bin time into 3 calendar-year bins, chosen so that no country appears in the same temporal bin twice. Results are generally robust to the use of individual year dummies.

procedure, we first choose (with replacement) 38 countries, and give each observation a new ID number. We then bootstrap sample, within each ID number, by the interaction of primary sampling unit (PSU) and survey round, and jointly estimate $\hat{\alpha}$ and $\hat{\beta}$ for each survey, replacing country based fixed effects with ID based fixed effects. We repeat the double bootstrap sampling 500 times and report the standard deviation of the estimates as the bootstrap standard error estimate of $\hat{\delta}$. Empirically, the large sample sizes from each survey seem to make this secondary source of variation rather small, and the two standard error estimates are similar.

4.2 Results from Rate of Loss Estimates

Table 2 presents results from estimating Equation 2 on the parameter estimates from Equation 1. The α columns show the coefficient estimates of GDP on the implied length-for-age z scores (LAZ) estimates, and the β columns show the coefficients on the estimates of the rate of loss of HAZ. The first specification provides the OLS estimate, the “within” specification (columns 3 & 4), includes γ_j , the “between” specification ((columns 5 & 6) includes λ_y , and the “DnD” specification (columns 7 & 8) present estimates when both are included.

The coefficient estimates on α are generally small, highly variable across specifications and imprecisely estimated. Taken at face value to assess potential magnitude, the point estimate on the OLS coefficient implies that a 10% increase in GDP is associated with an approximately 0.1sd increase in length-at-birth z-score. The coefficient magnitude increases to almost 0.4sd for the within estimate (with a confidence interval approximately that wide), but the preferred fixed-effects estimates (columns 7 and 8) puts the magnitude around 0.02sd, again with a relatively wide confidence interval.

The coefficient estimates on β , on the other hand, are robust across specifications and fairly precise. A 10% increase in GDP is associated with around a 0.002 sd decrease in the rate of loss of HAZ. In a country whose median child becomes exactly stunted on their second

birthday after being born 0.25sd below the reference children (reasonable given Figure 3), the rate of loss would be around 0.07 sd/month, meaning a 10% change in GDP is associated with approximately a 3% change in the base rate of loss. Cumulatively, a change in the rate of loss of HAZ of 0.002sd would generate a 0.04-0.05sd effect by the time the child reached their 3rd birthday.

5 Individual-Level Estimates

Our individual-level estimates come from a series fixed effects specifications that isolate the effect of inputs at various ages on the entire HAZ-age profile. The intuition for the fixed-effects models we estimate below can be motivated by a simple thought experiment. Suppose a researcher has a set of cross-sectional surveys with HAZ measurements from different countries covering a number of survey rounds in each country. Merging this data with country-year observations of GDP generates a country-year pseudo-panel dataset. The first insight we exploit is simply that this same procedure can apply even if we keep only the observations from each survey that are children of age A. The second insight is that there are potential gains to precision by estimating the model with all ages simultaneously, and we flesh out the corresponding fixed-effects model to allow for simultaneous estimation of all of the parameters of interest.

5.1 Single Exposure Age-Profile Fixed-Effects Models

The following equation represents the regression analogue of the initial thought-experiment, containing observations for only children aged A born into cohort year C:

$$HAZ_{icj}^A = X'_{ijc} \beta^A + \delta^A * GDP_{cj} + \mu_j^A + \lambda_c^A + \zeta_{icj}^A \quad (3)$$

This equation reduces to the standard panel fixed-effects model, but estimates the coefficients only on children of age A . There is no reason this regression cannot then be repeated for children aged $A+1$, $A-1$, $A+2$... etc. A two year old child in country J born in 2000 is compared to a two year old born in Country J in 1995. Simultaneously, the same 2 year old born in 2000 in Country J is compared to a 2 year old born in 2000 in Country J' . By merging GDP to cohort (birth year) instead of survey year, each age-group is identified off different (though correlated) variation because each cohort was born into the GDP stream experienced by their country at a different point in calendar time. If the coefficient off 2 year olds is identified off the GDP values from 1995 and 2000, the coefficient on 3 year olds is identified off GDP values from 1994 and 1999.

We then generalize the above thought experiment and regression equation into a multi-age framework where we can estimate the entire vector of δ^A simultaneously for children of all ages. We alter the preceding equation by allowing μ_j^A and λ_c^A to become μ_{ja} and λ_{ca} , that is, we make the country and cohort fixed-effects age-specific. We interpret these fixed-effects as controlling for country-specific HAZ-age profiles (μ_{ja}) and for a child's "lifespan", their growing up from birth to age A over the years T to $T+A$ (λ_{ca}).

$$HAZ_{ijca} = X'_{ijca}\beta + \sum_A \delta^A * GDP_{jc} + \mu_{ja} + \lambda_{ca} + \epsilon_{ijca} \quad (4)$$

This regression again has both within- and between-country comparison analogs, but these comparisons are now made only within a particular country's children of the same age, or across children who have lived the same "lifespan". Our "within" variation comes from comparing children of age A in country J and born in cohort C , with children of the same age A and same country J but born in cohort C' . That is, we difference out the average effect (over the whole sample period) of growing up to age A in Country J . The remaining variation in GDP and HAZ for a cohort in Country J then contains three components: exogenous noise (by assumption), any effect of GDP, and the secular improvement in HAZ that would have

occurred for children who grew up over this period even absent any within-country economic growth.

Since HAZ is not an instantaneous outcome, the calendar date on which one is measured can only affect HAZ by, in combination with cohort, fixing the calendar time through which an individual child grows; that is, by being a proxy for the time into which you were born and through which you grew up. Our “between” variation (λ_{ca}) then comes from comparing children who were born at the same time and measured at the same age (who lived the same “lifespan” over calendar time), but in different countries. All children who lived the same lifespan experienced secular changes in life-improving technologies over the exact same points in their human development. Given the remaining variation in GDP after de-meaning by country-age, the component of the secular time trend that is common across countries for children of that cohort measured at that age can be purged from the remaining variation by simultaneously estimating indicator variables for each lifespan permutation (age-cohort-period) realized in the data. Interpreted in terms of allowing for idiosyncratic time trends, controlling for child lifespan allows the non-parametric secular time trends to themselves vary non-parametrically by age.¹⁵

5.2 Results from Single Exposure Models

Figure 4 graphs, across child age, the coefficients and confidence intervals from the regressions outlined in Equation 3 and Equation 4. The estimates from age-disaggregated regressions on children at each age in years are graphed in black, the estimates from simultaneous estimation are graphed in blue, and both estimates are of similar magnitude and precision. The coefficient estimates on children under age 1 are small and statistically indistinguishable from zero. However, as the child reaches their third birthday, the magnitude of the correlation

¹⁵A graphical representation of both the within- and between-country variation in the age-profile fixed-effects models is presented in Figure A1, which plots the HAZ and GDP time series for three countries over several survey rounds.

grows and the confidence intervals remains of similar magnitude. By age 3, a 10% increase in birth year GDP is associated with a 0.04sd increase in HAZ. The point estimates for children age 4 (48 to 60 months of age) is somewhat smaller, around a 0.025sd, but is not statistically differentiable from the estimate for 36 to 48 month olds.

Table 3 provides the regression estimates from the simultaneous estimation strategy described in Equation 4 and graphed in Figure 4, along with several alternative/traditional fixed-effects estimates that also include flexible controls for child age¹⁶. The specification in column 1 has only country and survey fixed effects, adjusting for child age with a series of indicator variables for age in months. Column 2 includes country-specific age indicators, along with survey fixed effects, and column 3 includes country fixed effects with lifespan (cohort-by-survey) fixed effects. The specification in column 4 includes both the country-age and lifespan fixed effects and is our preferred, saturated model. After controlling non-parametrically for child age, all of the point estimates are similar, and standard errors for all models are of similar size.

5.3 Pseudo-Replication of Vollmer et al.: Aggregation and Stunting

In order to provide insights into why our estimates differ from Vollmer et al. (2014) and what elements of our analytic perspective drive the differences, we provide several alternative estimates more in the flavor of those provided in Vollmer et al. (2014). First, we provide a set of regression results that estimate a single coefficient across children of all ages, doing so with both HAZ and stunting (an indicator variable defined as more than 2sd below the reference population) as outcome variables in Table 4. The first and fourth columns of Table 4 ignore age-profile effects, though they control non-parametrically for child age by including a vector of indicator variables for child age in months. These replicate the statistically insignificant

¹⁶ The simultaneous and age-specific estimates are also provided side-by-side in Table A5 to facilitate comparisons.

findings in Vollmer et al. (2014) and the point estimates are generally smaller than those in our age-specific estimates for older children, as would be expected from averaging across a heterogeneous effect. The inclusion of lifespan fixed-effects and country-specific age-profiles (columns 2 and 3) improves precision (standard errors are approximately 5 – 10% smaller) and the coefficient becomes marginally statistically significant.

Columns 4, 5 & 6 of Table 4 employ the stunting variable as an outcome instead of HAZ and are the regressions most similar to those in Vollmer et al. (2014), though we continue to merge GDP with cohort and not survey year. In no specification does the effect on stunting become statistically significant, and the point-estimates can be interpreted to be rather small, a 10% increase in GDP being associated with a 0.4 – 0.6*pp* decrease in stunting rates. Since stunting is a binary measure, and changes in stunting status only occur as children cross the -2sd HAZ threshold, it is a generally less sensitive measure of health and thus harder to detect statistically. However, our age-profile estimates are powerful enough to statistically differentiate the correlation from 0. These results are provided in Table 5 and our estimates on stunting rates are statistically significant for most children over age 1 and the magnitude of the association increases roughly 50% relative to the aggregate regression. By a child’s fifth birthday, though, the relationship has weakened sufficiently such that our preferred model (column 4) does not estimate a statistically significant association.

6 Extensions and Limitations

6.1 Alternative Exposure Timing and Serial-Correlation in GDP

We have focused our analysis on estimating the correlation between GDP at the time of birth and HAZ as the child grows. However, the theoretical human capital models from which we derive our estimating equations makes predictions across a second dimension of child age as well. For a 4 year old, the effect of an input at age 2 compared to age 3 has

two different components. First, there are the fundamental differences in the human capital production function over child age ($f^a()$). Secondly, re-optimization of household decision making in each period means that households experiencing economic growth early in their child's lives may have incentives to invest differently than they previously would have absent that growth.

The serial correlation in GDP makes it difficult to disentangle when the timing of GDP growth matters relative to a child's development or if the timing matters much at all. The classic time-series serial-correlation problem is less of a concern for us since any given coefficient is estimated using GDP measures from 4-5 years apart, but there are two additional issues that arise from the serial-correlation in the country-level GDP time-series.

First, fixing GDP to birth year, the variation that identifies the coefficient on GDP for 3 year olds will be highly correlated with the variation that identifies the coefficients for 2 year olds, because the GDP values in those regressions come from shifting the (highly serially-correlated across time) panel GDP data back one period. That means there is lack of independence in our estimates across age-at-measure. Second, while we do not have strong problems with serial-correlation in calendar-time, we do have strong serial-correlation problems in "age-time", in the vector of lagged (relative to age-at-measure) GDP realizations that constitute the "stream of inputs" that were provided the the child as they aged. This makes disentangling the effects of GDP at birth from the effects of GDP growth during a child's life very difficult, at least without imposing significant structure on the underlying statistical machinery.

This is clear in Figure 5, which graphs coefficient estimates from a series of regressions on 4-year old children that shift the timing of GDP relative to the child's cohort. Each point on the graph represents the estimate of $\hat{\delta}$ from a regression where the GDP exposure is tied to a particular point in the child's development, from three years before their birth up to two or three years after they were measured. The strong serial correlation in GDP (and in

changes in GDP) means that many of the estimates are statistically significant. However, estimates from merges of the GDP stream to years between the birth of the child and the date they were measured generate the largest coefficient estimates and have lower bounds much further above the estimates those from estimates based on years before the child was born or after they were measured.

6.2 Estimation of Full Exposure Stream

In this section we present estimates of the full stream of exposures to GDP on HAZ by age. We use the following empirical strategy. We first define A as the age at which child i of cohort c and country j are measured and appear in our data. Let a refer to the age at which a child was exposed to a GDP value of interest. For example, if $a = 0$ then GDP from the year of birth (c , the child’s cohort year) will be assigned.

For children measured at age A , we define the following regression equation with $GDP_{j,c+a}$ being defined as the level of GDP experienced by a child from country J and cohort C at age a (in year $c + a$). δ_a^A , then, is defined as the parameter representing “exposure at age a on a child measured at age A ”.

$$HAZ_{ijc}^A = X'_{ijc}\beta^A + \sum_a \delta_a^A * GDP_{j,c+a} + \mu_j^A + \lambda_c^A + \epsilon_{ijc}^A \quad (5)$$

Figure 6 presents estimates, separately for each age-at-measure, of the association between GDP exposure at each period from a child’s birth until the age at which they were measured. First, we see that the estimates are incredibly noisy and they have overly large coefficients that tend to cancel one another out. This is due in large part to the serial-correlation in GDP over time, but that may be less of a problem in other empirical contexts. Second, for every age group, the contemporary, survey-time coefficient dominates the other input periods. Since the regressions are run separately by age, it could simply be an improbable “draw”

on the joint HAZ/GDP distribution in that particular set of calendar years and countries. However, we leave it to future researchers to determine whether this result is meaningful or informative, or simply the result of contemporary GDP being the only exposure-year for which each cohort experiences it in the same calendar year.

7 Conclusion

Prior studies that have estimated the relationship between economic growth and child health have often assumed away the actual process of child development and growth faltering, preferring instead static measures of population average health. In this paper we set out to develop new conceptual and methodological tools to capture the dynamic effects of health inputs on children’s growth trajectory. We parametrically model the HAZ-age profiles from 126 DHS surveys and interpret our results as evidence that the relationship between economic growth and child health is more apparent in the rate of loss of HAZ than in length at birth. In the aggregate, richer countries, and countries that grow richer, raise children that grow faster than those in poorer countries and those that did not experience recent growth.

We also make strides towards separately identifying differences across permutations of “age at measure” and “age at exposure” relevant to estimating the age-dynamic parameters of an optimal health investment model. We introduce age-profile fixed-effects models that exploit variation within and across countries and cohorts to capture the impact of medium term economic growth on child growth trajectory. Similar to the magnitudes from the aggregate regressions, we find that a 10% increase in medium term GDP is associated with an almost 0.04sd increase in HAZ by the time a child reaches the age of 3. Though the least-squares estimation we use in this paper is insufficiently structured to identify all the potential exposure-by-measurement age permutations in the presence of strong serial-correlation in GDP, we believe the econometric framework we lay out here can be greatly improved upon in terms of variance for the willingness to trade off small amounts of bias by placing prior

restrictions on the coefficient dynamics across age (at exposure and measurement).

Finally, we hope that our demonstration convinces readers of the value of such an age-profile perspective in empirical contexts such as this where cohort-panel outcome data is merged to aggregate spatio-temporal variation and an age-determined outcome of interest. Averaging effects across age can generate misleading estimates that are not as directly interpretable relative to the world as they may seem, and failure to account for the age-cohort relationship can lead to bias as seen in Cummins (2013). Such age-aggregated models also needlessly obscure nuances in human development that are revealed in our analysis without losses in statistical precision. We hope that more researchers will consider making the HAZ-age profile, and not simply mean HAZ or stunting rates, their object of investigation.

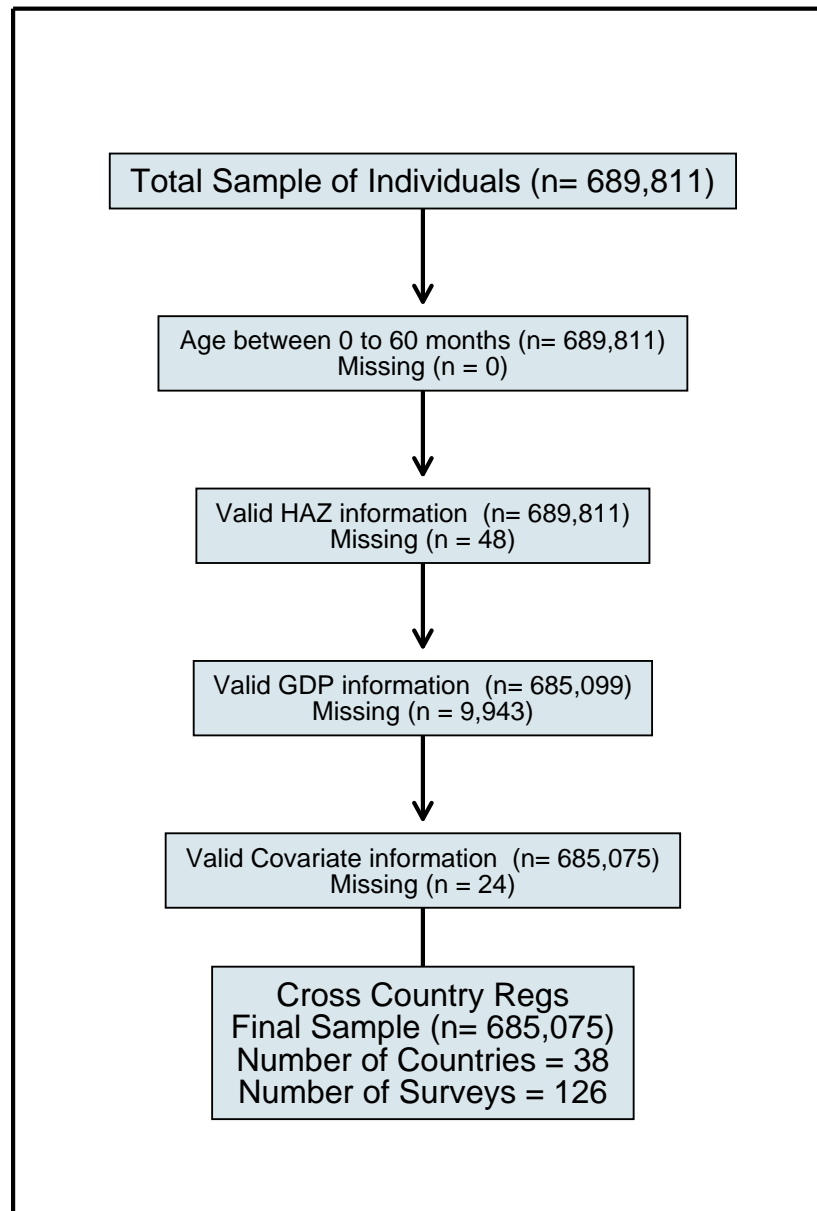
References

- Alderman, H., L. Haddad, D. D. Headey, and L. Smith (2014). Association between economic growth and early childhood nutrition. *The Lancet Global Health* 2(9), e500.
- Becker, G. S. (1962). Investment in human capital: A theoretical analysis. *Journal of political economy* 70(5, Part 2), 9–49.
- Behrman, J. R., M. C. Calderon, S. H. Preston, J. Hoddinott, R. Martorell, and A. D. Stein (2009). Nutritional supplementation in girls influences the growth of their children: prospective study in guatemala. *The American journal of clinical nutrition* 90(5), 1372–1379.
- Bershteyn, A., H. M. Lyons, D. Sivam, and N. P. Myhrvold (2015). Association between economic growth and early childhood nutrition. *The Lancet Global Health* 3(2), e79–e80.
- Cummins, J. (2013). On the use and misuse of child height-for-age z-score in the demographic and health surveys. Technical report.
- Elbers, C., J. O. Lanjouw, and P. Lanjouw (2005). Imputed welfare estimates in regression analysis. *Journal of Economic Geography* 5(1), 101–118.
- Glewwe, P. and E. A. Miguel (2007). The impact of child health and nutrition on education in less developed countries. *Handbook of development economics* 4, 3561–3606.
- Grossman, M. (1972). On the concept of health capital and the demand for health, 80 j. *Pol. Econ* 223(10.2307), 1830580223.
- Haddad, L., H. Alderman, S. Appleton, L. Song, and Y. Yohannes (2003). Reducing child malnutrition: How far does income growth take us? *The World Bank Economic Review* 17(1), 107–131.
- Harttgen, K., S. Klasen, and S. Vollmer (2013). Economic growth and child undernutrition in sub-saharan africa. *Population and Development Review* 39(3), 397–412.
- Headey, D. D. (2013). Developmental drivers of nutritional change: a cross-country analysis. *World Development* 42, 76–88.
- Hoddinott, J., J. A. Maluccio, J. R. Behrman, R. Flores, and R. Martorell (2008). Effect of a nutrition intervention during early childhood on economic productivity in guatemalan adults. *The lancet* 371(9610), 411–416.
- I C F, I. (1992-2011). Demographic and health surveys (various) [datasets].
- Joe, W., R. Rajaram, and S. Subramanian (2016). Understanding the null-to-small association between increased macroeconomic growth and reducing child undernutrition in india: role of development expenditures and poverty alleviation. *Maternal & child nutrition* 12(S1), 196–209.

- Klasen, S. (2008). Poverty, undernutrition, and child mortality: Some inter-regional puzzles and their implications for research and policy. *The Journal of Economic Inequality* 6(1), 89–115.
- Kramer, K., A. Veile, and E. Otárola-Castillo (2016). Sibling competition & growth tradeoffs: biological vs. statistical significance. *PLoS ONE* 11(3), e0150126.
- Lange, S. and S. Vollmer (2017). The effect of economic development on population health: a review of the empirical evidence. *British medical bulletin* 121(1), 47–60.
- O’Connell, S. A. and C. Smith (2016). Economic growth and child undernutrition. *The Lancet Global Health* 4(12), e901–e902.
- Rieger, M. and S. K. Trommlerová (2016). Age-specific correlates of child growth. *Demography* 53(1), 241–267.
- Singh, A. (2014). Why are economic growth and reductions in child undernutrition so weakly correlated? and what can public policy do? *The Lancet Global health* 2(4), e185–e186.
- Smith, L. C. and L. Haddad (2002). How potent is economic growth in reducing undernutrition? what are the pathways of impact? new cross-country evidence. *Economic Development and Cultural Change* 51(1), 55–76.
- Subramanyam, M. A., I. Kawachi, L. F. Berkman, and S. Subramanian (2011). Is economic growth associated with reduction in child undernutrition in india? *PLoS Med* 8(3), e1000424.
- Topel, R. (1999). Labor markets and economic growth. *Handbook of labor economics* 3, 2943–2984.
- Victora, C. G., M. de Onis, P. C. Hallal, M. Blössner, and R. Shrimpton (2010). Worldwide timing of growth faltering: revisiting implications for interventions. *Pediatrics*, peds–2009.
- Vollmer, S., K. Harttgen, M. A. Subramanyam, J. Finlay, S. Klasen, and S. Subramanian (2014). Association between economic growth and early childhood undernutrition: evidence from 121 demographic and health surveys from 36 low-income and middle-income countries. *The lancet global health* 2(4), e225–e234.
- Walker, R. S., M. Gurven, O. Burger, and M. J. Hamilton (2008). The trade-off between number and size of offspring in humans and other primates. *Proceedings of the Royal Society of London B: Biological Sciences* 275(1636), 827–834.
- World Health Organization (1995). Physical status: The use of and interpretation of anthropometry, report of a who expert committee. Technical report.
- World Health Organization (2006). Who child growth standards: Length/height-for-age, weight-for-age, weight-for-length, weight-for-height and body mass index-for-age: Methods and development. Technical report.

8 Figures & Tables

Consort Flowchart



Adapted from: www.consort-statement.org/consort-statement/flow-diagram

Figure 1: CONSORT Diagram

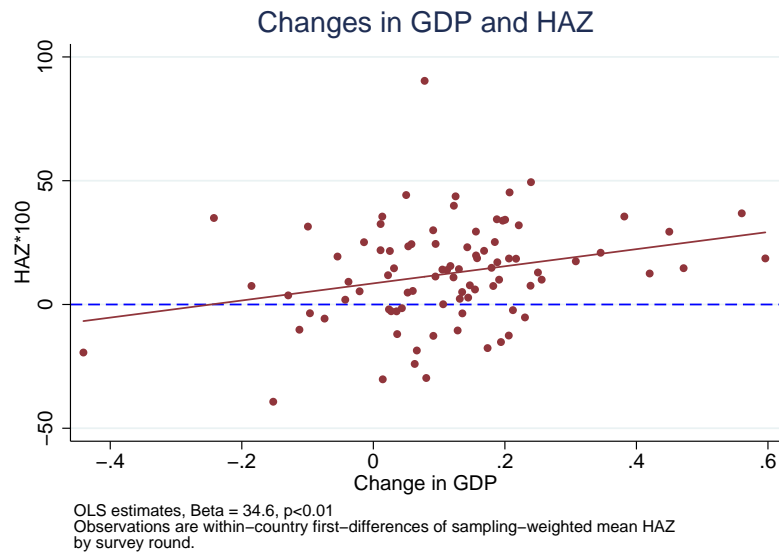



Figure 2: Birth Year GDP and Child HAZ

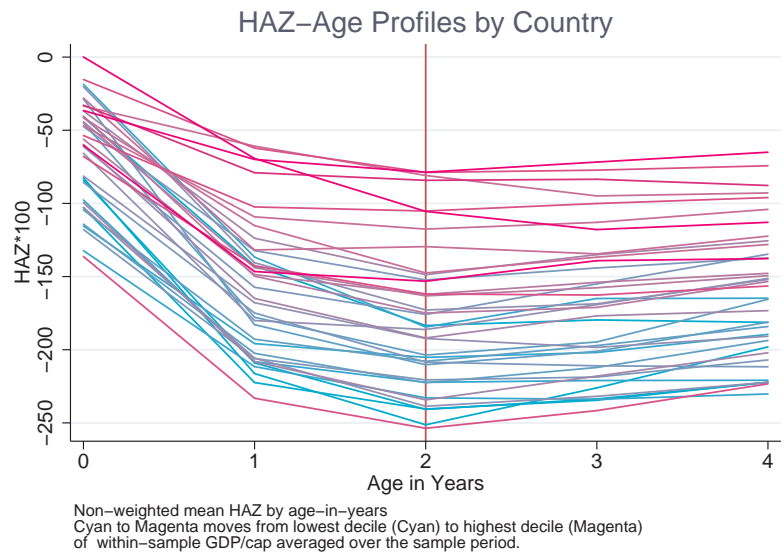



Figure 3: HAZ Age Profiles

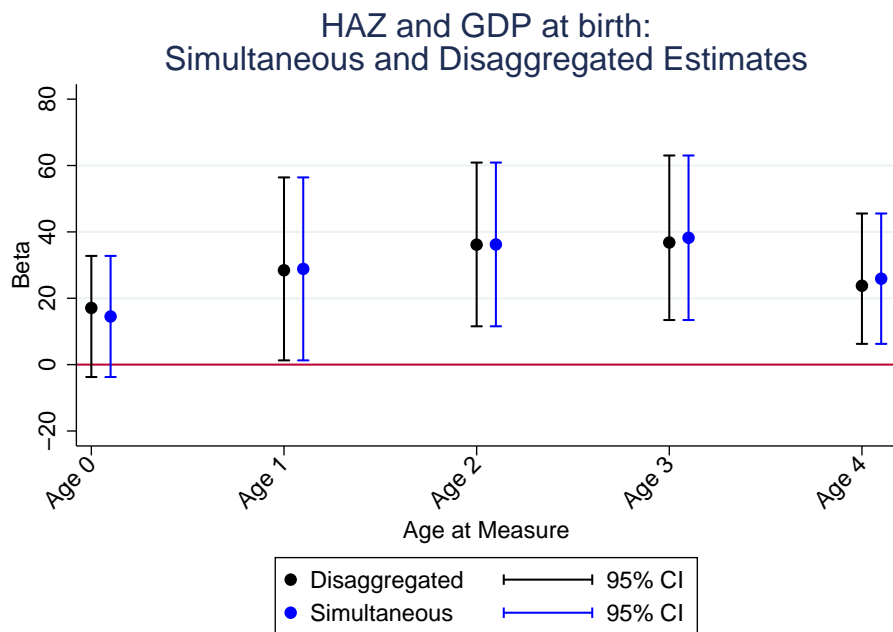


Figure 4: Results: Age-Specific Effects of GDP on HAZ

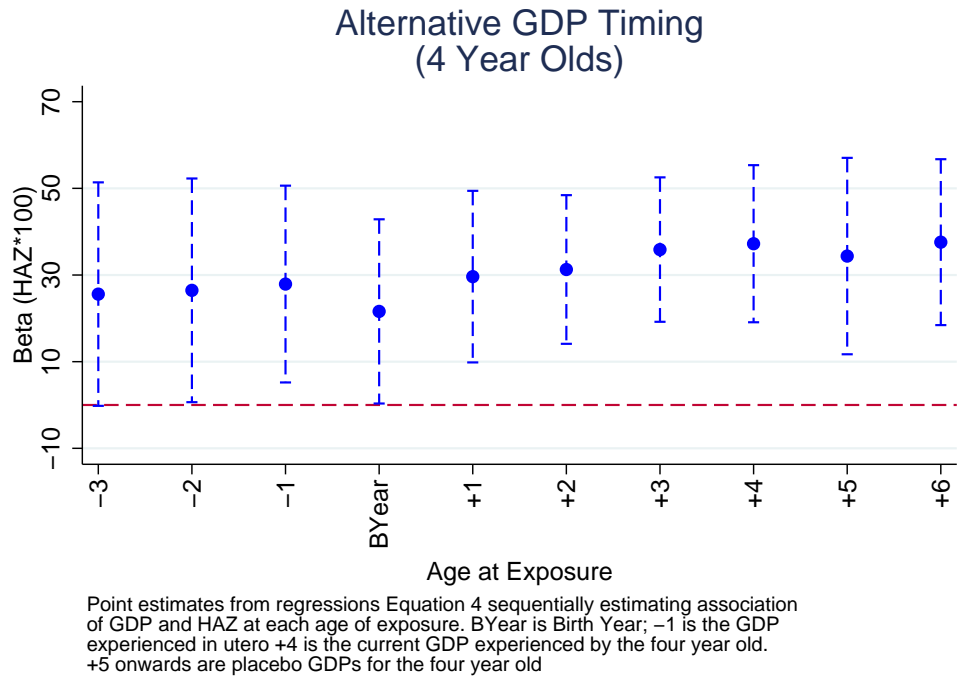


Figure 5: Results: Age-Specific Inputs of GDP on HAZ




Figure 6: Simultaneous Estimation Across Age-at-Exposure, by Age-at-Measure

Table 1: Sample Summary Statistics

Variable	Mean	SD
HAZwho	-143.89	168.16
Child age (months)	28.40	17.17
GDP per capita (log)	6.58	0.92
Survey Year Gap (years)	5.98	2.28
<i>Covariates</i>		
female (%)	0.50	0.50
Maternal Age (years)	28.80	6.83
Maternal Education		
No Education (%)	0.36	0.48
Primary Education (%)	0.35	0.48
Secondary Education (%)	0.24	0.42
urban (%)	0.36	0.48
<i>N</i>	685075	

Table 2: Results: Rate of HAZ Loss and LAZ

	OLS		Between		Within		DnD	
	Alpha b/se/bse	Beta b/se/bse	Alpha b/se/bse	Beta b/se/bse	Alpha b/se/bse	Beta b/se/bse	Alpha b/se/bse	Beta b/se/bse
GDP	9.00 (5.54) [5.60]	1.68*** (0.33) [0.35]	10.0+ (5.33) [5.40]	1.65*** (0.33) [0.36]	36.4* (17.8) [18.4]	2.21* (0.93) [0.99]	2.08 (15.7) [20.6]	2.34* (0.90) [1.15]
Mean	-19.7	-7.6						
Survey FE			✓	✓			✓	✓
Country FE					✓	✓	✓	✓
r2	0.04	0.36	0.15	0.37	0.08	0.08	0.23	0.10
N	126	126	126	126	126	126	126	126

+ 0.10, * 0.05, ** 0.01, *** 0.001; Robust standard errors clustered at the country level for 41 countries; For each specification columns represent results for children under 2, the first column presents values of the average LAZ scores (α) and the second column presents values for the rates of loss of HAZ (β) from Equation 2; Analytic cluster standard errors in (), 2-stage Bootstrap SE in [].

Table 3: Age-Profile Regressions: HAZ and GDP at Birth

	(1)	(2)	(3)	(4)
	HAZ	HAZ	HAZ	HAZ
	b/se	b/se	b/se	b/se
Age 0	8.8 (10.5)	3.2 (10.4)	12.8 (9.6)	14.5 (9.3)
Age 1	28.8***	30.3**	32.5***	28.9**
Age 2	(10.5) 33.6***	(13.4) 38.8***	(10.0) 36.7***	(14.1) 36.2***
Age 3	(10.7) 31.3***	(13.1) 38.6***	(10.0) 34.2***	(12.6) 38.2***
Age 4	(10.6) 28.8***	(11.0) 24.8**	(9.9) 31.9***	(12.6) 25.9**
Urban	(10.5) 31.9***	(9.5) 32.0***	(9.6) 31.9***	(10.0) 32.0***
Mat. Age	(2.0) 0.8***	(2.1) 0.8***	(2.0) 0.8***	(2.1) 0.8***
Female	(0.1) 14.3***	(0.1) 14.3***	(0.1) 14.3***	(0.1) 14.3***
	(1.2)	(1.2)	(1.2)	(1.2)
Sample Mean	-143.91	-143.91	-143.91	-143.91
Age Dummy (Yrs)	✓			
Survey FE	✓	✓		
Country	✓		✓	
Country-Age		✓		✓
Lifespan			✓	✓
r2	0.119	0.049	0.120	0.049
Obs	685075	685075	685075	685075

+ 0.10, * 0.05, ** 0.01, *** 0.001; Ordinary Least Squares regressions include listed covariates and fixed-effects, along with dummies for maternal education group and DHS regions. HAZ is measured in SD and multiplied by 100. Clustered robust standard errors allow for heteroskedasticity and serial-correlation at the country level.

Table 4: Results: Aggregate Correlations of GDP with HAZ and Stunting

	(1)	(2)	(3)	(4)	(5)	(6)
	HAZ	HAZ	HAZ	Stunted	Stunted	Stunted
	b/se	b/se	b/se	b/se	b/se	b/se
GDP at Birth	22.6 (11.7)	26.6* (10.9)	25.4* (11.1)	-0.04 (0.03)	-0.06 (0.03)	-0.05 (0.03)
Urban	32.6*** (2.3)	32.6*** (2.3)	32.6*** (2.3)	-0.09*** (0.008)	-0.09*** (0.008)	-0.09*** (0.008)
Mat. Age	0.8*** (0.1)	0.8*** (0.1)	0.8*** (0.1)	-0.002*** (0.0003)	-0.002*** (0.0003)	-0.002*** (0.0003)
Female	14.3*** (1.2)	14.3*** (1.2)	14.3*** (1.2)	-0.04*** (0.003)	-0.04*** (0.003)	-0.04*** (0.003)
Mean	-143.91	-143.91	-143.91	.36	.36	.36
Age	✓			✓		
Survey	✓			✓		
Country	✓	✓		✓	✓	
Country_Age			✓			✓
Lifespan		✓	✓		✓	✓
r2	0.16	0.16	0.05	0.13	0.13	0.04
N	685075	685075	685075	685075	685075	685075

+ 0.10, * 0.05, ** 0.01, *** 0.001; Ordinary Least Squares regressions include listed co-variates and fixed-effects, along with dummies for maternal education group and DHS regions. HAZ is measured in sd and multiplied by 100, and stunting is an indicator variable for $HAZ < -2$. Clustered robust standard errors allow for heteroskedasticity and serial-correlation at the country level.

Table 5: Age-Profile Regressions: Stunting and GDP at Birth

	(1)	(2)	(3)	(4)
	Stunted	Stunted	Stunted	Stunted
	b/se	b/se	b/se	b/se
Age 0	0.005 (0.03)	0.06* (0.03)	-0.008 (0.03)	-0.008 (0.02)
Age 1	-0.06* (0.03)	-0.06* (0.03)	-0.07** (0.03)	-0.06** (0.03)
Age 2	-0.08** (0.03)	-0.10*** (0.03)	-0.09*** (0.03)	-0.09** (0.04)
Age 3	-0.07** (0.03)	-0.10*** (0.03)	-0.08*** (0.03)	-0.08** (0.04)
Age 4	-0.06* (0.03)	-0.06 (0.04)	-0.07*** (0.03)	-0.05 (0.04)
Urban	-0.09*** (0.007)	-0.09*** (0.007)	-0.09*** (0.007)	-0.09*** (0.007)
Mat. Age	-0.002*** (0.0003)	-0.002*** (0.0003)	-0.002*** (0.0003)	-0.002*** (0.0003)
Female	-0.04*** (0.003)	-0.04*** (0.003)	-0.04*** (0.003)	-0.04*** (0.003)
Sample Mean	0.36	0.36	0.36	0.36
Survey FE	✓	✓		
Country FE	✓		✓	
Country-Age		✓		✓
Lifespan			✓	✓
r2	0.081	0.040	0.082	0.041
Obs	685075	685075	685075	685075

+ 0.10, * 0.05, ** 0.01, *** 0.001; Ordinary Least Squares regressions include listed covariates and fixed-effects, along with dummies for maternal education group and DHS regions. Stunting is an indicator variable for $HAZ < -2$. Cluster robust standard errors allow for heteroskedasticity and serial-correlation at the country level.

A Literature Review

Table A1: Cross Country Studies

Cross Country	(No.)	Yrs Spanned*	Data	Agg**	N	Outcomes	Child Age***	VoI	Timing ^E	Method	Effect Size ^R
Smith and Haddad (2002)	63	1970-1996	DHS	Country / Survey Year	179	WAZ	0to5	GDPpc- WDI	survey year	CountryFE	6.3% reduction in under-nutrition
Haddad et al. (2003)	61	1970-1995	WHO	Country / Survey Year	175	WAZ	0to5	GDPpc - WB	survey year	Country FE, Decade FE	32pp increase in WAZ
Headey (2013)	89	1977-2007	DHS	Country / Survey Year	160	Stunting	0to5	GDPpc- WDI	survey year	Country FE	1.8pp decrease in stunting
Harttgen et al. (2013)	28	1991-2009	DHS	Child / Survey Year	380,000	Wasting , Stunting, Under-weight	0to5	GDPpc - PWT	survey year	Country FE, Survey Year FE, Linear Age	1.5-1.7 LOR reduction in stunting
Vollmer et al. (2014)	36 (127 surveys)	1990-2011	DHS	Child/ Survey Year	460,000	Wasting , Stunting, Under-weight	0to3	GDPpc - PWT	survey year	Country FE, Survey Year FE, Linear Age	No / weak association
Subramanyam et al. (2011) on India	India	1992, 1998, 2004	DHS	Child/ Survey Year	77,000	Stunting	0to3	GDPpc (state), linear age,	survey year	State and survey year FE	No Association

No. - Number of countries, *Years Spanned, not necessarily included. **Agg - Level of Aggregation, ***Age in years (Unless specified),

VoI - Variable of Interest, ^E - Timing of exposure, ^R - Effect Size in relation to a 10% increase in GDP, DHS - Demographic Health Surveys,

WHO - World Health Organization, WDI - World Development Indicators, GDPpc - GDP per capita, FE - Fixed Effects, LOR - log odds ratio

B Summary Statistics by Country

Table A2: Summary Statistics: Asia

Country Name	Survey Years	BirthYears	N	HAZ	Birth Year GDP pc (USD 2005)
1. Armenia	3	1996-2010	4,074	-73.08	1356.94
2. Bangladesh	5	1992-2011	22,391	-190.37	391.141
3. Cambodia	3	1995-2010	10,803	-181.53	420.29
4. Jordan	5	1986-2012	27,806	-67.39	2183.46
5. Pakistan	2	1986-2012	7,114	-202.13	611.84
6. Turkey	3	1989-2003	9,943	-81.47	5605.791

Table A3: Summary Statistics: South America

Country Name	Survey Years	BirthYears	N	HAZ	Birth Year GDP pc (USD 2005)
7. Bolivia	3	1993-2007	23,032	-135.51	971.06
8. Brazil	2	1982-1995	5234	-80.06	4035.81
9. Colombia	5	1984-2009	38,430	-91.91	3321.21
10. Dominican Republic	3	1987-2013	10,109	-75.47	3081.81
11. Guatemala	2	1991-1998	12,420	-232.49	3081.80
12. Haiti	3	1997-2011	9,217	-111.84	472.59
13. Peru	4	1987-2000	34,168	-148.99	2183.584

Table A4: Summary Statistics: Africa

Country Name	Survey Years	BirthYears	N	HAZ	Birth Year GDP pc (USD 2005)
14. Benin	3	1997-2011	23,748	-167.17	535.46
15. BurkinaFaso	4	1988-2010	22,319	-149.18	358.81
16. Cameroon	3	1986-2010	10,548	-127.84	963.09
17. Congo	2	2001-2011	8,368	-109.47	1746.24
18. CoteDIvoire	2	1994-2011	4,689	-123.17	1013.27
19. Egypt	5	1988-2013	53,200	-96.21	1195.76
20. Ethiopia	3	1988-2002	22,035	-181.96	138.30
21. Ghana	3	1994-2008	8,099	-130.72	463.35
22. Guinea	3	1995-2012	8,645	-121.63	292.66
23. Kenya	3	1988-2008	14,647	-141.86	529.37
24. Liberia	2	1981-2013	7,495	-140.04	207.41
25. Madagascar	3	1988-2008	13,277	-196.03	289.29
26. Malawi	4	1988-2009	25,037	-194.54	215.22
27. Mali	3	1996-2012	24,450	-150.38	420.92
28. Morocco	3	1982-2003	14,203	-111.99	1458.99
29. Mozambique	2	1999-2011	17,372	-170.48	317.04
30. Namibia	4	1988-2012	10,675	-122.68	3317.52
31. Niger	3	1987-2011	12,579	-173.03	280.13
32. Nigeria	4	1986-2012	53,293	-148.08	855.14
33. Rwanda	4	1988-2010	18,045	-184.38	254.15
34. Senegal	3	1988-2010	10,271	-114.48	731.27
35. Tanzania	5	1987-2009	27,852	-181.88	329.15
36. Uganda	5	1984-2011	17,784	-167.55	251.265
37. Zambia	4	1987-2006	20,978	-189.17	647.78
38. Zimbabwe	4	1984-2010	13,302	-134.20	539.93

C Representation of Identifying Variation

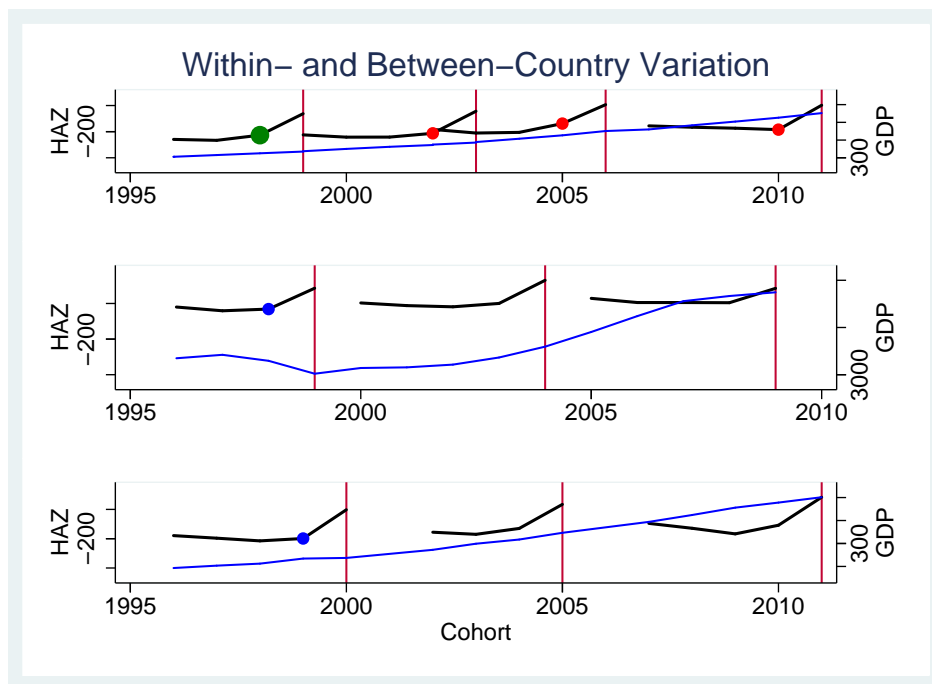


Figure A1: Identifying Variation

This figure provides a graphical representation of the identifying variation used in Equation 4. The 3 panels that plot HAZ and GDP against time (x axis) represent 3 different countries in the DHS surveys. The blue line represents the GDP trend in those countries. Each red line represents a DHS survey year and the black lines represent the HAZ age profiles of that country generated from that survey year and averaged by cohort (which, fixing survey time, is perfectly negatively correlated with age). The within country variation compares each red dot (in this case, 2 year olds) growing up in country J at different points in time. The cross country variation compares the blue observations, the effect of growing up in country J or J' at age 2 in the same time period. The green observation represents the child that is both 2 years old in Country J and grew up between the years of 1998 and 2000. Our empirical strategy differences out these two effects for any given child, and estimate the association of changes in GDP with changes in HAZ from this doubly de-meanned identifying variation.

D GDP HAZ Disaggregated vs Simultaneous Estimation

Table A5: HAZ and GDP at Birth: Simultaneous and Disaggregated Estimates

	(1)	(2)	(3)	(4)	(5)	(6)
	HAZ	HAZ	HAZ	HAZ	HAZ	HAZ
	b/se	b/se	b/se	b/se	b/se	b/se
Age 0	14.5 (9.3)	17.1* (9.9)				
Age 1	28.9** (14.1)		28.4** (13.8)			
Age 2	36.2*** (12.6)			36.1*** (12.6)		
Age 3	38.2*** (12.6)				36.8*** (12.7)	
Age 4	25.9** (10.0)					23.8** (10.1)
Urban	32.0*** (2.1)	22.8*** (2.1)	29.8*** (2.5)	36.4*** (2.4)	37.7*** (2.5)	36.4*** (2.5)
Mat. Age	0.8*** (0.1)	0.5*** (0.1)	0.6*** (0.09)	1.1*** (0.1)	1.1*** (0.1)	1.0*** (0.1)
Female	14.3*** (1.2)	21.9*** (1.4)	25.9*** (1.7)	14.9*** (1.0)	3.1** (1.4)	1.2 (1.4)
Sample Mean	-143.91	-66.79	-159.26	-178.86	-171.22	-160.86
Survey FE		✓	✓	✓	✓	✓
Country		✓	✓	✓	✓	✓
Country-Age	✓					
Lifespan	✓					
r2	0.049	0.024	0.052	0.064	0.069	0.074
Obs	685075	161657	142736	133959	132635	114088

+ 0.10, * 0.05, ** 0.01, *** 0.001; Ordinary Least Squares, Regressions clustered by country. Controls also include dummies for maternal education and regions. HAZ is in 100s, Stunting is a 0/1 dummy.