Journal of Development Economics 143 (2020) 102409

Contents lists available at ScienceDirect

Journal of Development Economics

journal homepage: www.elsevier.com/locate/devec

The Lazarus drug: the impact of antiretroviral therapy on economic growth☆

Anna Tompsett

Stockholm University, Department of Economics, Sweden

A B S T R A C T

Does better population health lead to growth in per capita income? Theory is ambiguous and empirical evidence is very limited. In 2001, a steep fall in antiretroviral (ARV) drug prices triggered rapid and massive expansion of ARV therapy coverage in lower-income countries. Exploiting the sharp resultant changes in population health, I show that ARV therapy coverage expansion led to growth in GDP per capita. The positive effects on growth most likely persist for around four years. ARV therapy coverage expansion could explain around a third of the sub-Saharan African "growth miracle".

1. Introduction

When individual health improves, so does individual productivity.[1] Many policymakers therefore expect population health improvements to lead to income growth[2] but this may not always be the case. Population health improvements often increase the size of the workforce and may lead to population growth. Larger workforces experience diminishing marginal returns to additional workers, especially when production is dependent on other inputs that are fixed in quantity, such as land, or take time to accumulate, such as physical capital. When populations grow, per capita incomes tend to decrease because wealth must be divided between more individuals. Population health improvements can therefore either increase or decrease per capita income, depending on the circumstances. Causal estimates of the effect of population health improvements on per capita income are extremely scarce. Large, sharp changes in population health are necessary to distinguish the causal effects of changes in population health from other background trends, and such changes are rare. In this study, I exploit rapid and massive expansion in access to antiretroviral (ARV) therapy for HIV/AIDS to provide new causal evidence on the effect of population health on per capita income.

HIV/AIDS is a global epidemic with its epicentre in sub-Saharan Africa. HIV/AIDS is incurable but can be treated with ARV therapy. ARV therapy is sometimes known as *the Lazarus drug*, because it converts HIV/AIDS from an imminent death sentence into a latent, albeit chronic, condition. Effective ARV therapy was developed in the late 1990s, but prices were initially prohibitively high, constraining adoption in lower-income countries. By 2001, only 30,000 people in lower-income countries were receiving ARV therapy (Reich and Bery, 2005). Drug prices faced by these countries then fell by a factor of ten in less than a year, following a price war triggered by generic competition (Médecins Sans Frontières, 2001). ARV therapy coverage took off sharply and rose rapidly, reaching 14.9 million people in lower-income countries by the end of 2014 (WHO, 2015a), as shown in Fig. 1. Over the same time period, sub-Saharan African life expectancy grew by around 17%,[3] having previously stagnated for more than a decade, and economic growth rates experienced a renaissance described as a "growth miracle" (Pinkovskiy and Sala-i-Martin, 2014; Rodrick, 2018; Young, 2012). This paper evaluates whether ARV therapy cover-

☆ Author declarations of interest: none. The author acknowledges financial support from the Tore Browaldh Foundation. The Tore Browaldh Foundation played no other role in the development of this article. Thanks for helpful comments and suggestions to: Douglas Almond, Heléne Berg, Solomon Hsiang, Mathias Iwanowsky, Andreas Madestam, Hanna Mühlrad, Pablo Ottonello, Cristian Pop-Eleches, Jakob Svensson, Alessandro Tarozzi, Jonas Vlachos, Brian Williams, Josh Graff Zivin; two anonymous referees and the editor; and seminar participants at Stockholm University, Sussex University, the University of Gothenburg, Umeå University, the Sustainable Development Research Symposium, the ASWEDE Conference on Development Economics, the EUDN Scientific Conference and the CEPR-PODER Annual Symposium on Development Economics. Particular thanks to Paul Cahu and Falilou Fall for their generosity in sharing data.

E-mail address: anna.tompsett@ne.su.se.

[1] See e.g. Bleakley (2010a), Strauss and Thomas (1998), Thomas and Frankenberg (2002).

[2] See e.g. WHO (2001).

[3] Author calculations from World Bank data, based on life expectancy at birth.

https://doi.org/10.1016/j.jdeveco.2019.102409
Received 11 July 2018; Received in revised form 26 September 2019; Accepted 26 October 2019
Available online 16 November 2019
0304-3878/

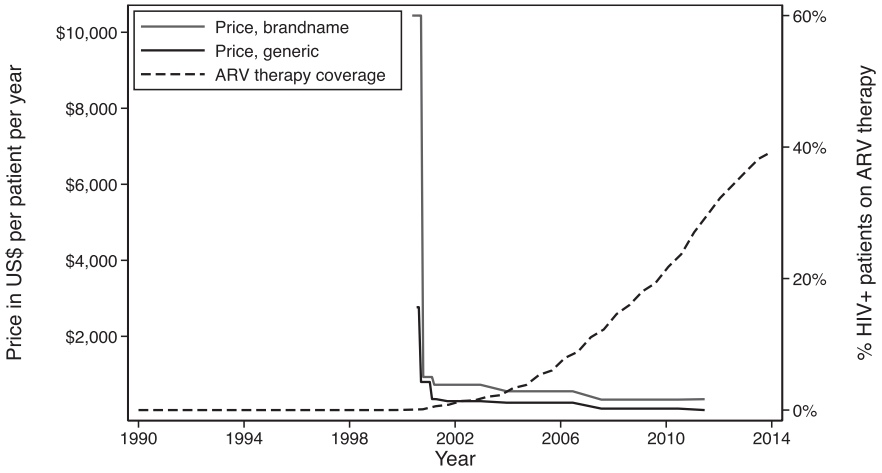

Notes Price data refers to the first-line combination of stavudine (d4T), lamivudine (3TC), and nevirapine (NVP) (Médecins Sans Frontières, 2001, n.d.). ARV therapy coverage estimates from author's calculations using UNAIDS data.

Fig. 1. ARV prices and coverage in low- and middle-income countries.

age expansion led to income growth and if so, whether the effects of ARV therapy coverage expansion can help explain the African "growth miracle".

The paper begins by describing the evolving relationship between HIV prevalence and trends in health and income growth. Sharp changes in trends that coincide with ARV therapy coverage expansion provide preliminary evidence for a positive causal effect of ARV therapy coverage expansion on income growth. Fig. 2 illustrates. Before ARV therapy coverage expansion, life expectancy and income growth are both declining in countries with high HIV prevalence, relative to other low- and middle-income countries. When ARV therapy coverage takes off, these negative relative trends reverse. While HIV prevalence is especially high in sub-Saharan Africa, the trend reversals are not explained by different regional trends. Structural tests confirm that the trend reversals coincide with the sharp fall in prices.

I then use a fixed effects-instrumental variables (IV) approach to estimate the effect of ARV therapy coverage expansion on economic growth. The central empirical challenges are the potential for reverse causality—higher incomes also lead to better health—and omitted variables bias, as other, possibly unobservable, factors may influence both health and income. My analysis accounts comprehensively for all observable and unobservable differences between countries in income levels, growth rates, and linear trends in growth rates, using yearly panel data from 90 low- and middle-income countries between 1990 and 2014. However, even after accounting for these differences, naïve estimates of the effect of ARV therapy coverage on economic growth could still be biased.

The potential for bias arises from differences in how successfully countries expand ARV therapy coverage. For example, Kenya and the Central African Republic both had HIV prevalence of 9% of the population in 2001. However, by 2014, ARV therapy coverage among the HIV positive was 55% in Kenya but only 18% in the Central African Republic. If countries that were more successful in expanding ARV therapy coverage were also experiencing other rapid changes in growth trends, this could introduce bias into estimates of the effect of ARV therapy on growth.

To eliminate this potential source of bias, I use predicted ARV therapy coverage as an instrument for observed ARV therapy coverage. To predict ARV therapy coverage, I interact the fraction of the population infected with HIV in 2001, at the time of the fall in prices, with the fraction of HIV-positive individuals receiving ARV therapy across low- and middle-income countries. Intuitively, this reconstructs a scenario in which changes in ARV therapy coverage were uniformly distributed across all HIV positive individuals in low and middle-income countries. Using predicted ARV therapy coverage as an instrument thereby exploits only the variation in coverage that is driven by *global* changes in drug price and availability.

Extending ARV therapy coverage to an additional 1% of a country's population is associated with a 6.7% increase in life expectancy (90% confidence interval: 4.6–8.7) and leads to a 1.4 percentage point increase in growth rates in GDP per capita (90% CI: 0.2–2.5). Results are similar when I limit the analysis to sub-Saharan Africa alone. Three empirical exercises suggest that the positive effects on income growth are likely to persist for at least four years. While the estimated effects are substantial when expressed in percentage terms, growth takes place relative to low baseline incomes. In sub-Saharan Africa, the effects correspond to growth in per capita income of $14 (90% CI: 0.6–27), accumulating to $56 (90% CI: 2–109) if the effects persist for four years. The effects on growth are measured relative to previous declining trends, suggesting that they should be interpreted as avoiding the negative consequences of an unchecked HIV epidemic, including the anticipated decimation of the productive adult population (U.S. Census Bureau, 2004).

The estimates are unbiased as long as countries with high and low HIV prevalence did not experience any *different, simultaneous, and unrelated changes in trends* in income growth rates that coincide with ARV therapy coverage expansion. Trend changes with these characteristics would violate the exclusion restriction in the fixed effects-instrumental variables analysis, inducing bias in the estimated effects.

I address several potential concerns about channels through which such bias-inducing trend changes could arise. First, bias-inducing trend changes could arise through reverse causality, if ARV therapy coverage expansion were itself triggered by a reversal in relative trends in health and growth in high HIV prevalence countries. However, this concern is inconsistent with the historical evidence: ARV drugs were developed primarily for use in wealthy countries; the price shock was unexpected; and the policy motivations for ARV therapy coverage expansion were, if anything, falling life expectancy and stagnant economic growth in high HIV prevalence countries. I also find no evidence for relenting negative growth trends in the decade preceding ARV therapy coverage expansion. Second, regression to mean could produce bias-inducing trend changes, if high HIV prevalence in 2001 were a consequence of poor economic growth in the preceding decade. I show that the results are robust to using alternative instruments that are exogenous to growth in the decade preceding ARV therapy coverage expansion. Finally, bias-inducing trend changes would also arise if any other factor caused trends associated with HIV prevalence to change in a way that was cor-

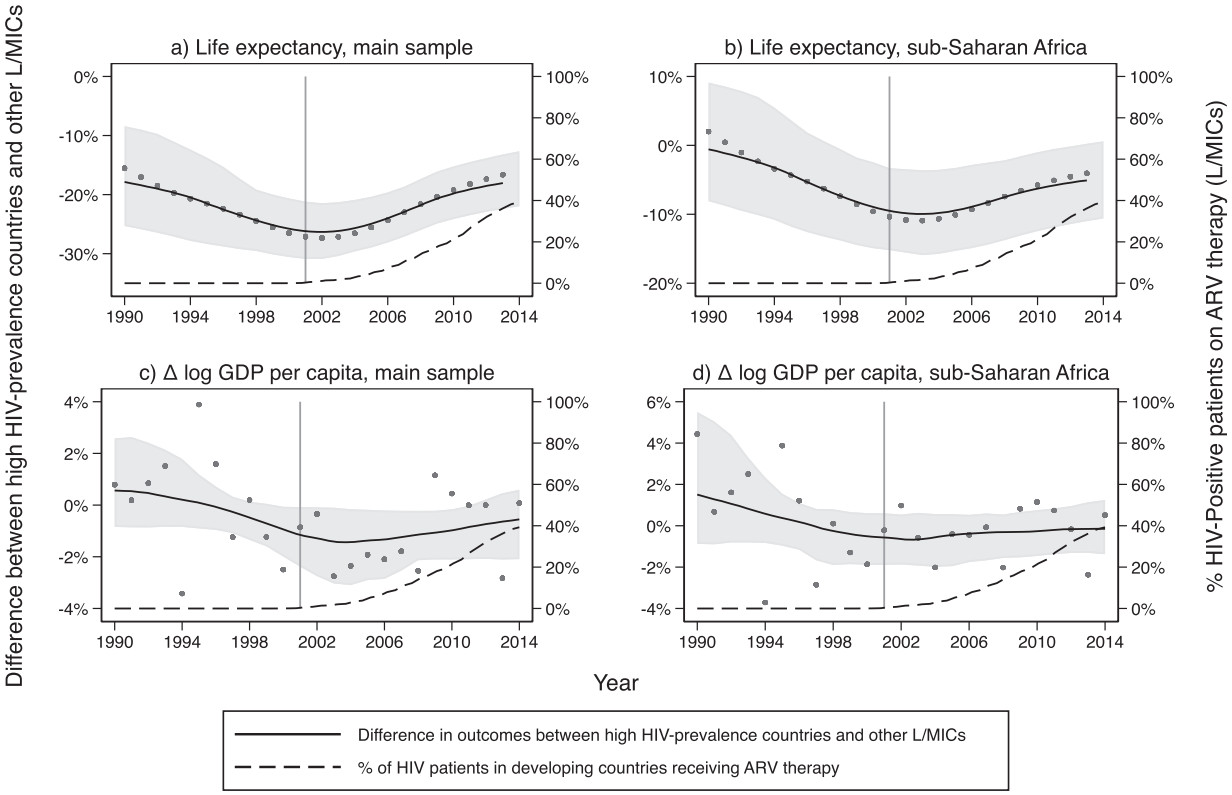

Notes Graphs show results from a local linear regression of the difference in the outcome variable between 18 high HIV prevalence countries (HIV prevalence $\geq$ 5% in 2001) and 72 (21 in sub-Saharan Africa) other low and middle income countries (HIV prevalence $<$ 5% in 2001) on time. Bootstrapped 90% confidence intervals are shown in grey. Dots show annual mean differences and the vertical line indicates the year in which the price of antiretroviral therapy fell.

Fig. 2. Differences between high and low HIV prevalence countries before and after ARV therapy coverage expansion.

related in time with ARV therapy coverage expansion. Of four potential sources of such a bias—mineral, petroleum, or export booms, and contemporaneous investment in fighting malaria—none can account for the main results.

The HIV epidemic has, thankfully, affected a relatively small number of countries, especially at very high infection levels. The relatively small number of affected countries restricts the precision with which effects can be estimated, yielding relatively wide confidence intervals. The small number of affected countries also presents an intrinsic inference problem: high HIV prevalence countries have potentially very high leverage on the estimates but are exactly the countries of interest to the analysis. I show that the results remain consistent when I systematically drop each country from the analysis. I also confirm that inference based on random permutations of HIV prevalence across countries yields similar conclusions to inference based on analytical standard errors.

The closest precedents to this study are a very small number of studies which also evaluate the causal effect of population health on mean income.[4] The best available estimate remains Acemoglu and Johnson (2007) seminal study of the major global health improvements known as the international epidemiological transition.[5] Acemoglu and Johnson (2007) pioneered the approach of using the predicted consequences of a change in health technology to instrument for its observed consequences. In contrast to the results in this paper, Acemoglu and Johnson find that the international epidemiological transition *lowered* GDP per capita relative to prevalent trends, because population growth outweighed growth in total GDP. The main contribution of this paper is therefore to demonstrate that population health improvements *can* lead to growth in mean income, at least in the short run.

Several factors could explain why this study finds positive effects on growth while Acemoglu and Johnson find negative effects. First, the unbiasedness of Acemoglu and Johnson's estimates depends on a stronger identifying assumption: that countries with high and low pre-1940 mortality from the diseases affected by the international epidemiological transition did not experience any unrelated differences in *growth* over six decades. Acemoglu and Johnson provide evidence in support of this assumption but its validity remains the subject of debate (e.g. Acemoglu and Johnson, 2014; Bloom et al., 2014). Second, Acemoglu and Johnson evaluate the long-run effects of a gradual improvement in population health that took place over several decades, over which time some countries' populations doubled. This study evaluates the immediate effects of a rapid improvement in population health, over a time period in which populations remained relatively stable. Simple economic models suggest that population health improvements

[4] Other studies show that health affects economic outcomes for households and individuals, but these studies do not recover the impact of population health on mean income (see e.g. Bleakley, 2007, 2010b; Cutler et al., 2010). An extensive literature includes health as a variable in cross-country growth comparisons (e.g. Bhargava et al., 2001; Bloom et al., 2004; Gallup and Sachs, 2001). However, the results of most cross-country studies are contentious, because exploiting only cross-sectional variation increases sensitivity to violations of the identifying assumptions. In contrast, this study joins a growing literature which uses natural experiments to establish causal relationships in macroeconomic data (see Fuchs-Schündeln and Hassan, 2016).

[5] Hansen and Prescott (2002) also show a causal effect of population health on mean income, specifically that mortality associated with the bubonic plague led to higher real wages in pre-industrial England. Hansen (2014) replicates the approach in Acemoglu and Johnson (2007) for states in the USA, finding a null effect on GDP per capita.

should increase mean income if population does not change. Finally, the diseases considered by Acemoglu and Johnson—including malaria, measles and diarrheal diseases—predominantly affect infant mortality. In contrast, HIV/AIDS primarily affects working-age adults, being the single largest cause of death in this age group in Africa at the turn of the millennium (WHO, 2004). The effect of population health on mean income is more likely to be positive when population health improvements are concentrated among the working-age population (see e.g. Ashraf et al., 2008).

This study complements earlier research on the economic consequences of ARV therapy. Most previous studies focus on individuals receiving ARV therapy or their households (e.g. Bor et al., 2012; Habyarimana et al., 2010; Thirumurthy et al., 2008). Two recent studies additionally find positive effects of ARV therapy coverage expansion on productivity (Baranov et al., 2015) and investment (Baranov and Kohler, 2018) in both the HIV positive *and* the HIV negative, attributed to changes in beliefs about mortality risks.

This paper also contributes to the wider literature on the economic impact of HIV/AIDS. Empirical estimates of the short-run effects of HIV/AIDS on income vary widely.[6] The results in this paper suggest that the short-run effects of HIV/AIDS must almost certainly be negative, since it is otherwise unlikely that ARV therapy coverage expansion would have positive effects on income growth. The results in this paper cannot, however, speak to the long-run consequences of either the HIV epidemic or ARV therapy coverage expansion.[7]

The results also shed new light on the sub-Saharan African "growth miracle". Growth rates in sub-Saharan Africa between 2002 and 2014 averaged 2.1%, compared to 0.2% between 1990 and 2001. Researchers have posited several explanations for this "growth miracle", including structural change, a commodity price boom, and a recovery in rainfall. To my knowledge, no other source considers ARV therapy coverage expansion as an explanation (see e.g. McKinsey Global Institute, 2010; McMillan and Harttgen, 2014). However, limited progress has been made in distinguishing empirically between potential explanations (McMillan and Harttgen, 2014) or in providing causal evidence for any of the posited explanations. An accounting exercise suggests that growth rates between 2002 and 2014 would have been reduced to 1.4% in the absence of ARV therapy coverage expansion. Bringing the first causal evidence to bear on this debate, my results suggest that ARV therapy coverage expansion could explain around a third of the sub-Saharan African "growth miracle".

The estimated effects on income almost certainly underestimate the true effects on welfare, as they do not capture the intrinsic benefits of improved health and reduced mortality. However, the estimated effects on income alone exceed the substantial costs of ARV therapy coverage provision. Income gains may be at risk if current levels of ARV therapy coverage are not maintained. The economic consequences of failure to maintain current coverage levels could therefore be severe.

The rest of the paper proceeds as follows. Section 2 outlines the history of the HIV epidemic and ARV therapy development, and explains why this study provides an important complement to previous evidence. Section 3 summarizes the data. Section 4 describes the relationship between HIV prevalence and trends in both health and economic growth. Section 5 sets out the fixed effects-instrumental variables estimation strategy. Section 6 reports the estimated effects of ARV therapy coverage expansion on growth, discusses how long the effects on growth can be expected to persist, and outlines some possible mechanisms. Section 7 shows a series of robustness tests and specification checks. Section 8 concludes.

2. Context

2.1. The HIV/AIDS epidemic

The Human Immunodeficiency Virus (HIV) is a retrovirus[8] that attacks and weakens the human immune system, eventually rendering it susceptible to disease. HIV is transmitted via the exchange of bodily fluids, for example through unprotected sex. HIV has a relatively long latent period—up to fifteen years—after which it develops into its full-blown form, Acquired Immunodeficiency Disease (AIDS). The transmission mechanism and long latent period mean that the vast majority of HIV infections and deaths from AIDS occur among prime working-age adults.

Identified in 1983 (Barré-Sinoussi et al., 1983; Popovic et al., 1984), HIV originates in a family of viruses endemic to primate populations in West and Central Africa (Keele et al., 2006). HIV passed to humans via cross-species transmission (Korber et al., 2000). The disease was rare until the middle of the 20th Century (Nahmias et al., 1986) but later developed into a full-blown pandemic. By the mid-1990s, HIV had spread across Africa and overseas, with an estimated 20 million people infected worldwide, including 13 million in sub-Saharan Africa (Mertens and Low-Beer, 1996). Fig. 3a shows the distribution of HIV prevalence in 2001.[9]

2.2. Antiretroviral (ARV) therapy

Antiretroviral (ARV) therapy describes a class of drugs that combat HIV. ARV therapy is a treatment, not a cure: it reduces viral load, but does not eliminate the virus from the body. The principal challenge to ARV therapy development was that HIV mutates very rapidly. This allowed drug-resistant strains to evolve quickly. The key to effective ARV therapy, established in 1996, was prescribing several drugs in combination, which reduced the likelihood of resistance developing (Gulick et al., 1997; Hammer et al., 1997). The three-drug cocktail that became standard reduced viral loads and mortality by up to 80% in high-income countries (e.g. Moore and Chaisson, 1999). ARV therapy also effectively reduces mortality in lower-income countries (e.g. Bor et al., 2013).

The price of combined ARV therapy was initially more than $10,000 per patient per year, severely limiting adoption in lower-income countries. Prices faced by lower-income countries then fell sharply by a factor of at least ten between 2000 and 2001, after Indian companies began to produce generic versions of ARV drugs. Big pharmaceutical companies initially resisted generic competition, but abruptly changed strategy when it became clear that international legal and political opinion was against them. The turning point was a Clinton administration decision to ignore violations of patent law for sub-Saharan African countries providing ARV therapy to their citizens. Pharmaceutical companies withdrew a lawsuit brought against the South African government for allowing generic versions of ARV drugs and instead reduced branded drug prices to compete. By the end of 2001, both branded and generic versions of the standard three-drug cocktail were available in low- and

[6] Early estimates probably suffered from omitted variables bias (see e.g. UNDESA, 2004). Several relatively recent, but unpublished, studies exploit instrumental variables approaches to evaluate the effect of an increase in HIV prevalence (Ahuja et al., 2009; Cahu and Fall, 2011; Mveyange et al., 2015). However, they draw very different conclusions, suggesting that not all the associated exclusion restrictions are valid.

[7] See, for example, Young (2005), who simulates growth scenarios with and without an HIV epidemic that dies out exogenously from the year 2000, concluding that very long-run GDP per capita may be higher after such an HIV epidemic dies out than if no such epidemic took place.

[8] A retrovirus is a type of virus characterized by the procedure by which it replicates inside the cells of its host. A useful and well-referenced source of information is www.avert.org, which provided much of the background to this section.

[9] Appendix Table B1 lists all countries with HIV prevalence above 1% in 2001, in rank order.

a

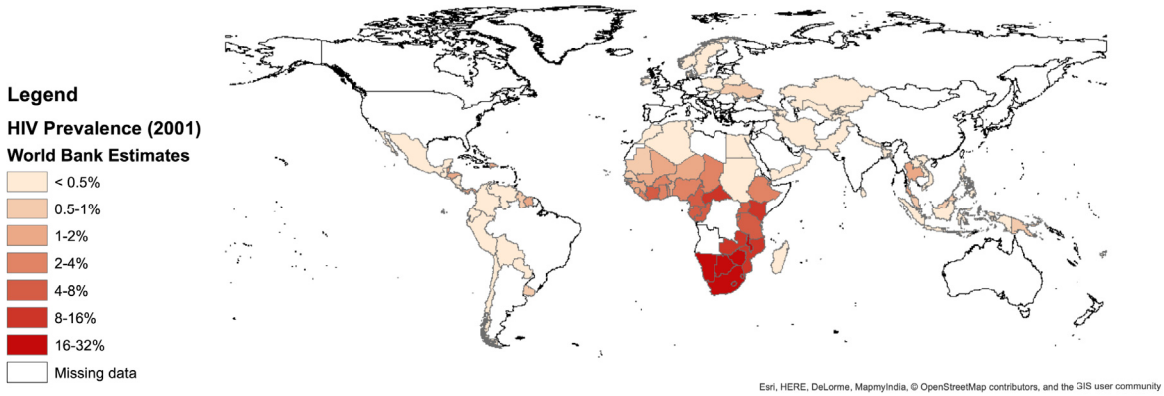

b

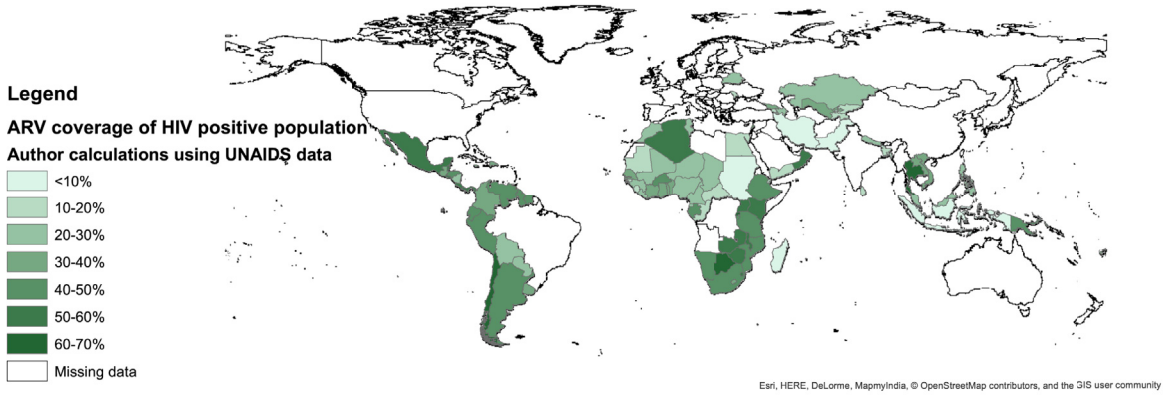

c

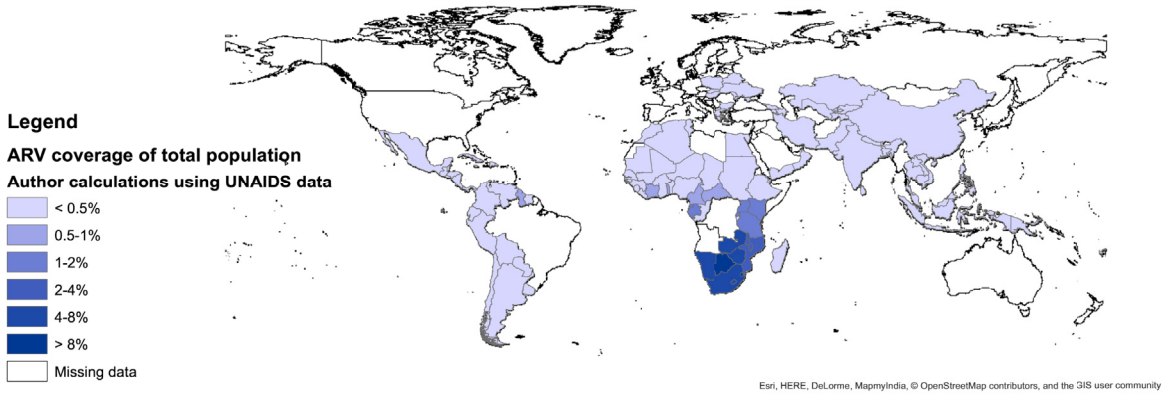

Notes World borders shape file from thematicmapping.org. Estimated HIV prevalence data among population aged 15-49 from World Bank World Development Indicators. ARV therapy coverage calculated using numbers of people living with HIV and receiving antiretroviral therapy as reported by UNAIDS and total population from World Bank data.

Fig. 3. Fig. 3a HIV Prevalence in 2001. Fig. 3b ARV coverage of people living with HIV in low- and middle-income countries (%), 2014. Fig. 3c ARV coverage of total population in low- and middle-income countries (%), 2014.

middle-income countries for less than US$1000 per patient per year (Médecins Sans Frontières, 2001). Fig. 1 illustrates.

The fall in prices precipitated a concerted global effort to increase ARV therapy coverage in low- and middle-income countries. The international effort included several major funding initiatives such as the Global Fund to Fight AIDS, Tuberculosis and Malaria, and the U.S. President's Emergency Plan for AIDS Relief (see e.g. Bendavid and Bhattacharya, 2009; Ford et al., 2011). However, price was only one among several important barriers to ARV therapy coverage expansion. ARV therapy originally required patients to take 10–15 pills per day, although this fell to as few as two with the development of fixed-dose combination therapies in 2003 (Schwartländer et al., 2006). Other barriers included a lack of trained staff, HIV testing programs, and reliable supply chains. Importantly, both the intensity of international efforts and many barriers to ARV therapy coverage expansion varied across countries. By 2014, however, 14.9 million people in low and middle income countries were receiving ARV therapy (WHO, 2015a). Fig. 3b and 3c show ARV therapy coverage as a percentage of the HIV-positive population and the total population, respectively, in 2014.[10]

Treatment protocols originally recommended ARV therapy for HIV patients only after functioning of the immune system fell below a specified threshold, or for specific vulnerable groups. However, more recent guidance recommends that all HIV positive individuals should receive treatment. These changing guidelines imply that those who first receive ARV therapy now are less likely to be symptomatic than those who first received ARV therapy when it originally became available.

2.3. When do population health improvements increase mean income?

Simple growth models show that the effect of population health improvements on mean income can be either positive or negative. Population health affects mean income through three channels, two positive—effects on productivity and human capital—and one negative—population growth. Positive effects on mean income are more likely when two conditions hold.[11] First, population health improvements are more likely to have net positive effects when effects on productivity and human capital are strong relative to effects on population growth. Second, positive effects are more likely in the short run than in the long run. This is because population growth effects take time to play out, while productivity effects typically occur much faster.

Previous studies, and in particular Acemoglu and Johnson (2007), estimate the long-run effects of improvements to population health with relatively strong effects on population growth. Evidence from ARV therapy coverage expansion provides an important counterpoint to this previous body of evidence for two reasons. First, effects on productivity and human capital are likely to be strong because of the disease profile of HIV, which primarily affects working-age adults. Second, we are able to observe the short-run effects of ARV therapy coverage expansion because it took place rapidly during a time period when we have yearly measurements of outcome variables.

3. Data

This study uses publicly-available country-level aggregate data, primarily from the World Bank and UNAIDS. Key variables include real GDP per capita (in 2005 US$), population, life expectancy, HIV prevalence, and ARV therapy coverage.[12] The main sample consists of 90 low- and middle-income countries for which HIV prevalence is reported in 2001, and for which at least 20 years of data on growth in GDP per capita is available.[13] The main unit of analysis is the country-year.

Aggregate data from lower-income countries is generally of uncertain quality. Simple mismeasurement introduces noise, and possibly attenuation bias, and reduces precision in the estimates. These are intrinsic limitations. However, a specific concern in this context is whether measurement error in the outcome variables could be correlated across space with HIV prevalence *and* in time with ARV therapy coverage expansion. If present, such an error structure could introduce systematic and potentially misleading bias into the analysis. For example, in the instrumental variables analysis, measurement error with this structure would violate the exclusion restriction.

While aggregate measures of income are certainly mismeasured (see e.g. Young, 2012), they are unlikely to be *systematically* mismeasured in such a way. To generate spurious results in this study, mismeasurement of income would need to vary with HIV prevalence within sub-Saharan Africa *and* then change sharply around the year 2001; there is no obvious reason why this should be the case.

In contrast, aggregate measures of health may be mismeasured in a systematically misleading way. Specifically, demographers calculate life expectancy using modelled mortality rates in countries which lack data on mortality rates, which includes many countries in sub-Saharan Africa. When HIV affects more than 2% of the population, the demographic models used to estimate mortality rates simulate the effects of both HIV/AIDS and ARV therapy coverage (UNDESA, 2014). Although I show results that confirm that life expectancy moves in tandem with ARV therapy coverage, I do so with the caveat that the results may be affected by these modelling adjustments.

The main outcome, GDP per capita, depends on estimates of both total income and population. Population estimates may also be systematically adjusted using demographic models. However, population estimates are less dependent on modelled inputs than health measures, because they only reflect estimated fertility, mortality, and migration rates in years since the most recent census. In years up to and including the most recent census, population estimates are based on census data. To understand what drives the effects on GDP per capita, I decompose these effects into effects on growth in total GDP and population growth.

When I analyze the main sample, I compare changes in growth trends in high HIV prevalence countries to changes in growth trends in other low- and middle-income countries. All countries with HIV prevalence above 5% are in sub-Saharan Africa. If changes in growth trends differ in sub-Saharan Africa and other low and middle-income countries for other reasons, these differences could introduce bias into the main sample estimates. Limiting the analysis to sub-Saharan Africa alone eliminates the risk of this bias, at the cost of reducing the sample size and discarding potentially important variation. Throughout the paper,

[10] Appendix Table B2 lists all countries by ARV therapy coverage as a fraction of the total population in 2014, in rank order.

[11] In Appendix A.1, I formalize this discussion in a simple model adapted from Acemoglu and Johnson (2007).

[12] I obtain ARV therapy coverage rates from UNAIDS Global AIDS Response Progress Reporting. Data on the number of people receiving ARV therapy are reported on an annual basis since 2004 for all countries in the main sample, with few missing observations. Prior to 2004, global estimates suggest that less than 1% of HIV positive patients in low- and middle-income countries received ARV therapy. Lacking better information, I treat ARV therapy coverage as zero in low- and middle-income countries before 2004. All other data are obtained via the World Bank Databank, except where otherwise specified. HIV prevalence for the population of people aged 15–49 is estimated by UNAIDS.

[13] The main sample consists of 9 countries in East Asia and the Pacific; 10 in Europe and Central Asia; 20 in Latin American and the Caribbean; 8 in the Middle East and North Africa; 4 in South Asia; and 39 in sub-Saharan Africa. Results are very similar if I include other low and middle income countries, assuming that HIV prevalence is zero if it is not reported. Results are weakened in the main sample but slightly strengthened in sub-Saharan Africa alone when I apply a stricter inclusion criteria which limits the sample to a fully balanced panel. Results in Appendix B.

I show all analyses both for the main sample and for a sub-sample of sub-Saharan African countries. In practice, both sets of results are very similar.

In all analyses, I cluster standard errors by country. Economic growth, HIV prevalence, and ARV therapy coverage may all exhibit spatial and serial correlation, which if positive tends to lead to over-rejection of null hypotheses. Clustering standard errors at the country level allows for the correct statistical inference with any type of serial correlation (Bertrand et al., 2004; Wooldridge, 2002) but does not account for spatial correlation. However, additionally accounting for spatial correlation yields slightly smaller standard errors.[14]

4. Growth trends and HIV prevalence

This section describes the relationship between growth trends and HIV prevalence before and after the fall in ARV drug prices that led to the take-off of ARV therapy coverage expansion. The first goal of the exercise is to look for patterns consistent with causal effects of ARV therapy coverage expansion on income growth. Since ARV therapy coverage changes trend sharply in 2001, we should also expect growth trends in high HIV prevalence countries to change sharply shortly afterwards, when compared to other low- and middle-income countries.[15] The second goal of this section is to establish descriptive facts about the relationship between growth and HIV prevalence that inform the instrumental variables strategy described in section 5.

Fig. 2 shows trends in life expectancy and growth in GDP per capita in countries with HIV prevalence greater or equal to 5%, relative to other low- and middle-income countries. On the *y* axis in each figure, I show the difference between high HIV prevalence countries and other low and middle-income countries. Focusing on the differences abstracts from year-to-year variation in aggregate trends.[16]

During the 1990s, both life expectancy and growth rates in GDP per capita decline in high HIV prevalence countries relative to other low and middle-income countries. High HIV prevalence countries have lower life expectancies throughout the 1990s, with the gap between high and low HIV prevalence countries steadily increasing in size. High HIV prevalence countries have similar growth rates to low HIV prevalence countries at the beginning of the 1990s, but growth rates decline in relative terms throughout the decade.

Around 2001, the negative relative trends reverse, for both life expectancy and growth in GDP per capita. Visual inspection suggests that the timing of the reversal is consistent with the timing of ARV therapy coverage expansion, which is relatively slow until 2004.

The patterns are similar within sub-Saharan Africa alone, although high HIV prevalence countries begin the 1990s with similar life expectancies to, and higher growth rates than, other low- and middle-income countries in sub-Saharan Africa. By 2001, however, high HIV prevalence countries have lower life expectancies and slightly lower growth rates, before the negative relative trends begin to reverse.

Table 1 confirms that HIV prevalence in 2001 is associated with statistically significant negative trends in life expectancy and growth rates in GDP per capita before 2001 and with statistically significant trend reversals thereafter. I estimate the following equation, shown here for growth in GDP per capita, for both life expectancy and growth:

$$\Delta \ln y_{it} = \theta_1(\varphi_{HIVi2001} \times t) + \theta_2(\varphi_{HIVi2001} \times POST_t) + \theta_3(\varphi_{HIVi2001} \times POST_t \times t) + \zeta_i + \mu_t + \epsilon_{it} \quad (1)$$

where $\Delta \ln y_{it}$ is the change in log per-capita GDP in country i in year t, defined in event time relative to 2001; $\varphi_{HIVi2001}$ is HIV prevalence in country i in the year 2001; $POST_t$ is an indicator variable which takes the value one after 2001, and zero otherwise; and ζ_i and μ_t are year and country fixed effects, respectively.

The coefficients of interest are θ_1 and θ_3.[17] The coefficient θ_1 captures relative trends in health or income growth associated with HIV prevalence prior to ARV therapy coverage expansion. Rejecting the null hypothesis that $\theta_1 = 0$ implies that HIV prevalence is associated with statistically different trends before 2001. The coefficient θ_3 captures the change in these trends after the fall in prices of ARV drugs. Rejecting the null hypothesis that $\theta_3 = 0$ implies that there are statistically different relationships between HIV prevalence and trends before and after 2001. The estimated coefficients on the changes in trend (θ_3) are larger in magnitude than the initial negative trends (θ_1), suggesting that the initial negative trends fully reverse.[18]

The point estimates on θ_1 and θ_3 are larger within the sample of sub-Saharan African countries than in the full sample (columns 2 and 4). Both the initial relative negative trends and the trend reversals are thus driven primarily by differences between high and low HIV prevalence countries within sub-Saharan Africa and not by differences between sub-Saharan Africa and other regions.[19]

The negative trends in growth in GDP per capita prior to 2001 are primarily driven by negative trends in growth in total GDP, rather than trends in population growth. The trend reversals for growth in GDP per capita decompose to larger trend reversals in growth in total GDP offset by simultaneous smaller trend reversals in population growth.[20] The positive effects on growth in GDP per capita are thus driven primarily by trends in the income data. Even if the population data is subject to modelling adjustments, as discussed in Section 3, these adjustments do not appear to drive the trends in growth in GDP per capita.

Measurements of income growth may be more reliable over the long run than in measuring annual changes in income. I observe similar relationships if I collapse the data to measures of long-run changes (Appendix A.3).

The analysis so far tests whether the data are consistent with the presence of a trend break in 2001, the year in which prices fell. In Fig. 4, I plot the results of a structural break test which instead evaluates empirically the most likely timing of the observed trend break. I focus on growth in GDP per capita. To implement the structural break test, I first estimate a series of equations similar to Equation (1), varying the trend break year between 1996 and 2006, in both the main sample and sub-Saharan Africa alone. For each regression, I redefine event time relative to the trend break year. The structural break test (or Chow test) compares the F-statistic on the coefficient corresponding to θ_3 in analyses with different trend break years. The trend break year which gives the highest F-statistic is the most likely location of the trend break.[21]

[14] Appendix A.2 compares clustered standard errors to twoway clustering by country and year, and to a combination of clustered standard errors with Conley standard errors over distances of between 1000 and 5000 km.

[15] The intuition is similar to that of a regression kink design. See e.g. Card et al. (2015).

[16] Appendix Figures B1a to B1d plot the raw data.

[17] I include the interaction term associated with the coefficient θ_2 for completeness. The coefficient θ_2 captures changes in the average relationship between outcome variables and HIV prevalence after 2001. Given that the onset of ARV therapy coverage is associated with a slope change, rather than a jump in levels, we should not expect to find important changes in relationship between HIV prevalence and outcome variables before and after 2001. As expected, the estimated coefficients θ_2 are small relative to the trend coefficients.

[18] The hypothesis that trends exactly reverse, specifically that $\theta_1 + \theta_3 = 0$, is rejected for life expectancy, suggesting that trends more than reverse. However, this hypothesis is not rejected for growth in GDP per capita.

[19] I obtain similar results from estimating Equation (1) augmented with region-year fixed effects in the full sample.

[20] See Appendix Figure B2 and Appendix Table B3.

[21] In Appendix A.4, I simulate the results of the structural test using synthetic data. I show that the test I report in the paper performs better for the growth data than two alternatives, drawing on Hansen (2001). Appendix A.4 also repeats the structural tests for log life expectancy, population growth and growth in total GDP.

Table 1
HIV prevalence and trends in life expectancy and growth before and after ARV therapy coverage expansion.

	Log life expectancy		Change in log per capita GDP	
	(1)	(2)	(3)	(4)
HIV Prevalence$_{2001}$ × Post$_{2001}$ × t	0.178***	0.172***	0.035**	0.039**
	(0.019)	(0.023)	(0.015)	(0.018)
HIV Prevalence$_{2001}$ × Post$_{2001}$	−0.187**	−0.207**	−0.092	0.021
	(0.082)	(0.097)	(0.111)	(0.162)
HIV Prevalence$_{2001}$ × t	−0.106***	−0.123***	−0.021**	−0.030**
	(0.013)	(0.014)	(0.009)	(0.015)
N	2160	936	2236	974
Sample	Main	SSA	Main	SSA

Note: Coefficients are from OLS regressions of the outcome variable on listed variables, and year and country fixed effects. Main sample consists of 90 low- and middle-income countries (39 in sub-Saharan Africa). Data are from 1990 to 2014. Standard errors are clustered by country and shown in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$.

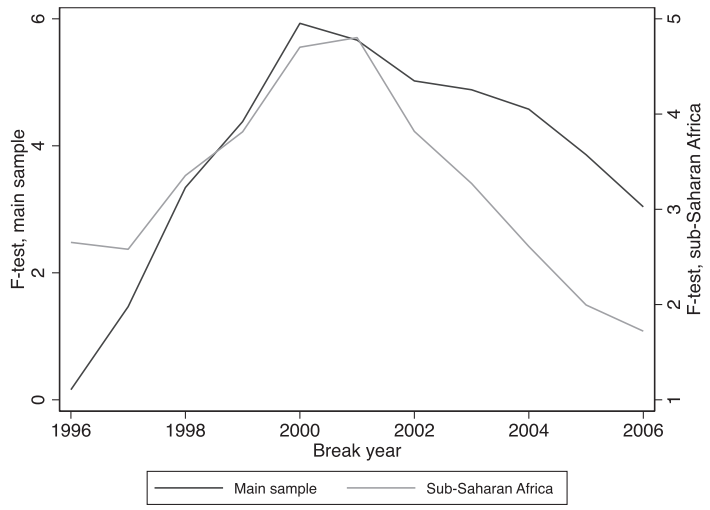

Notes Graphs show results of structural tests for timing of trend break in trends associated with HIV prevalence. Results obtained from regressions of outcome variable on year and country fixed effects; HIV prevalence in 2001 interacted with time; a post dummy which takes the value of one after the trend break year shown on the x axis; and the time and post dummy interaction. F-statistics test the null hypothesis of no change in trends associated with HIV prevalence in trend break year. Highest F-statistics identify most likely trend break year. Sample consists either of 90 low and middle-income countries (39 in sub-Saharan Africa).

Fig. 4. Structural tests for timing of trend-break in GDP per capita growth rates.

As Fig. 4 shows, this test selects the year 2000 in the main sample and 2001 in sub-Saharan Africa alone, almost exactly coinciding with the fall in prices in 2001.

The analysis in this section informs the instrumental variables strategy described in the next sections. First, the results in this section show the importance of accounting for country-specific trends as well as country fixed effects, because countries with high and low HIV prevalence clearly experience different trends prior to 2001. Second, the results suggests that the risk of introducing bias by comparing changes in trends in countries in sub-Saharan Africa to changes in trends in other low- and middle-income countries may be low. This is because the trend reversal I describe in this section is *not* driven by a generalized renaissance of growth in sub-Saharan Africa, but primarily by differences within sub-Saharan Africa. To be transparent, however, I continue to present results and robustness checks both for the main sample and for sub-Saharan Africa alone. Third, and most importantly, the evidence for a structural trend break in the early 2000s supports the identifying assumption required to give the results of the instrumental variables analysis a causal interpretation. Any posited alternative explanation for the results must explain why negative trends associated with HIV prevalence reverse at the specific moment that ARV therapy coverage expansion takes off.

5. Empirical strategy

This section describes the fixed effects-instrumental variables strategy I use to estimate the causal effect of ARV therapy coverage on per capita income. The main estimating equation is as follows:

$$\Delta \ln y_{it} = \pi \varphi_{ARVit} + \zeta_{i0} + \zeta_{i1} t + \mu_t + \epsilon_{it} \quad (2)$$

where $\Delta \ln y_{it}$ is the change in log per-capita GDP in country i in year t; φ_{ARVit} is ARV therapy coverage in country i in year t; μ_t are year fixed effects; and ζ_{i0} and ζ_{i1} are country fixed effects and country linear trends, respectively.

The parameter of interest is π, which corresponds to the effect of ARV therapy coverage on current growth. ARV therapy coverage is measured with respect to the total population in order to capture its

Table 2
First stage results.

	% of country population treated with ARVs	
	(1)	(2)
HIV Prevalence$_{2001}$ × global ARV coverage$_t$	0.71***	0.76***
	(0.06)	(0.07)
N	2236	965
Sample	Main	SSA

Note: Coefficients are from OLS regressions of outcome variable on listed variables, year and country fixed effects, and country-specific linear trends. Main sample consists of 90 low- and middle-income countries (39 in sub-Saharan Africa). Data are from 1990 to 2014. Standard errors are clustered by country and shown in parentheses. ***$p < 0.01$.

effect on the economy as a whole.[22] Including country-specific trends implies that the effect of ARV therapy coverage will be measured relative to prevailing trends.

Estimating Equation (2) by OLS may not yield unbiased estimates of π. The fixed effects and country-specific trends absorb all observable and unobservable differences in average growth rates and trends in growth rates across countries and across time. However, changes in ARV therapy coverage might still be correlated with other changes in growth trends because of other, possibly unobservable, factors. The most likely potential source of bias arises from differences in how successfully countries provide antiretroviral therapy coverage to their citizens. The direction of the bias is ambiguous. External support for ARV therapy programs might be concentrated in countries with declining growth prospects, relative to previous trends, which *ceterus paribus* would bias the estimates downwards, mirroring the classic problem with evaluating the effects of foreign aid. On the other hand, countries that are experiencing rapid growth relative to previous trends might be more successful in providing ARV therapy coverage. This would bias the estimates upwards.

To eliminate this potential source of bias, I implement an instrumental variables (IV) approach. I construct a measure of predicted ARV therapy coverage, $\varphi_{ARVit,pred}$, and use this as an instrument for observed ARV therapy coverage. I predict ARV therapy coverage by multiplying the proportion of HIV positive individuals receiving ARV therapy across all low- and middle-income countries at time t ($\overline{\psi}_{ARVt}$)[23] by HIV prevalence in country i in 2001, at the time of the sharp fall in prices ($\varphi_{HIVi2001}$). Since the vast majority of HIV positive individuals live in low and middle-income countries, I refer to ARV therapy coverage in low and middle-income countries as global ARV therapy coverage.

The first stage equation is as follows:

$$\varphi_{ARVit} = \beta\varphi_{ARVit,pred} + \widetilde{\zeta}_{i0} + \widetilde{\zeta}_{i1}t + \widetilde{\mu}_t + \nu_{it}$$
$$\text{where : } \varphi_{ARVit,pred} = \varphi_{HIVi2001} \times \overline{\psi}_{ARVt} \quad (3)$$

Predicted ARV therapy coverage essentially reconstructs what ARV therapy coverage would have looked like, if global advances in coverage were evenly distributed across all HIV positive individuals in low and middle-income countries. Using predicted ARV therapy coverage as an instrument for observed ARV therapy coverage thereby exploits only global changes in ARV therapy coverage, removing variation arising from how successfully individual countries respond to global changes in ARV therapy availability.

The IV approach will estimate local average treatment effects (LATEs) or effects on the compliers: those who are induced to expand ARV therapy coverage because of global changes in price and availability. In this case, it is reasonable to believe that most countries in the sample are compliers, given the very low take-up before the fall in prices and widespread take-up thereafter. The LATE is therefore unlikely to diverge strongly from the population average treatment effect.

Three conditions must be met for predicted ARV therapy coverage to be a valid instrument for observed ARV therapy coverage.

First, predicted ARV therapy coverage must be a strong predictor of observed ARV therapy coverage. Table 2 shows the first stage regressions. The first stage F-statistics are close to or above 100 in all regressions, in both the main sample and in the sub-Saharan Africa subsample.

Second, the monotonicity condition requires predicted ARV therapy coverage to always have a non-negative impact on ARV therapy coverage. This seems highly plausible. HIV infection is a pre-condition for ARV therapy treatment. Increases in global ARV therapy availability are also unlikely to have caused any country to decrease ARV therapy coverage.

Third, the exclusion restriction requires that conditional on the fixed effects and country-specific trends, predicted ARV therapy coverage must not be correlated with economic growth via any other channel than ARV therapy coverage or, in formal terms with respect to Equation (2), that cov $\left(\varphi_{HIVi2001} \times \overline{\psi}_{ARVt}, \epsilon_{it}\right) = 0$. The fixed effects and country-specific trends absorb, respectively, between-country differences in mean growth rates and linear time trends in growth rates. The exclusion restriction therefore corresponds to the requirement that HIV prevalence must be uncorrelated with other, unrelated *changes in trends* in economic growth that coincide with global ARV therapy coverage expansion. To illustrate, this identifying assumption could be violated if HIV prevalence and malaria prevalence were spatially correlated and if global investment in fighting malaria were contemporaneous with ARV therapy coverage expansion. In this case, I might confound the effects of improvements in malaria control with the effects of ARV therapy coverage expansion. In section 7.2, I evaluate this possibility, along with three other possible violations of the exclusion restriction, and conclude that none can account for the main results.

Other interventions may be provided alongside ARV therapy coverage, such as complementary investments in healthcare or treatment for tuberculosis, which is common among the HIV positive. If so, my estimates recover the total effect of the package of interventions. My estimates may also capture any knock-on effects of improvements to the rest of the healthcare system that result from reduced congestion. Neither of these caveats undermines the central conclusion regarding the relationship between population health and mean income.

The relationship between HIV prevalence and trends in economic growth might have changed anyway in the absence of ARV therapy cov-

[22] A country with 5% ARV therapy coverage with respect to the total population might be a country with 5% HIV prevalence and near 100% ARV therapy coverage among the HIV positive, or a country with 20% HIV prevalence and only 25% ARV therapy coverage coverage among the HIV positive.

[23] I calculate the proportion of HIV positive patients receiving ARV therapy from UNAIDS data on the number of HIV-positive individuals and the number of people receiving ARV therapy, summed across all low- and middle-income countries for which data is reported.

Table 3
Were trends associated with HIV prevalence changing before ARV therapy coverage expansion?

	Log life expectancy		Change in log per capita GDP	
	(1)	(2)	(3)	(4)
HIV Prevalence$_{2001}$ × t	−0.078***	−0.077***	−0.002	−0.063
	(0.023)	(0.027)	(0.026)	(0.046)
HIV Prevalence$_{2001}$ × t^2	−0.003	−0.004*	−0.002	0.003
	(0.002)	(0.002)	(0.002)	(0.003)
N	1080	468	1070	468
Sample	Main	SSA	Main	SSA

Note: Coefficients are from OLS regressions of the outcome variable on listed variables, and year and country fixed effects. Main sample consists of 90 low- and middle-income countries (39 in sub-Saharan Africa). Data are from 1990 to 2001. Standard errors are clustered by country and shown in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

erage expansion, for example as a consequence of the natural dynamics of the HIV epidemic. AIDS-related mortality would likely have peaked at some point in the mid-2000s (see Appendix Figure B4). However, the demographic impact of AIDS-related mortality would have been long-lasting. In the absence of ARV therapy coverage expansion, demographers projected population structures that had "never been seen before", in which the most productive adults were decimated as a proportion of the population (U.S. Census Bureau, 2004). It is therefore an open question whether negative trends in growth rates associated with HIV prevalence would have relented or accelerated. If negative trends in growth rates associated with HIV prevalence would have relented anyway, my approach might overestimate the effects of ARV therapy coverage, although this would not necessarily undermine the conclusions about the relationship between population health and mean income, because the results would still correspond to an improvement in population health leading to an increase in mean income. If, on the contrary, negative trends would have accelerated, my approach will underestimate the true effects.

It is not possible to test empirically whether or how trends associated with HIV prevalence would have changed in the absence of ARV therapy coverage. However, it is possible to test whether trends already showed signs of acceleration or deceleration before the fall in prices in 2001. Table 3 shows that negative trends associated with HIV prevalence were relatively stable in the decade preceding the fall in prices in 2001. I estimate equations with the following structure, including only the years up to and including 2001 and defining t as time since 1990:[24]

$$\Delta \ln y_{it} = \theta_4 \left(\varphi_{HIVi2001} \times t\right) + \theta_5 \left(\varphi_{HIVi2001} \times t^2\right) + \zeta_i + \mu_t + \epsilon_{it} \quad (4)$$

The term θ_5 captures whether HIV prevalence is associated with different non-linear (quadratic) trends prior to 2001. I estimate similar equations with log life expectancy as an outcome variable. Three out of four point estimates on the quadratic time trend terms are negative, suggesting that if anything, trends were worsening in the decade preceding 2001. However, the only significant quadratic trend term is an accelerating negative trend in life expectancy, within sub-Saharan Africa (column 2). Trends associated with HIV prevalence thus appear to be moderately stable during the 1990s.

6. Results

This section reports results from the instrumental variables analysis described in the previous section. Table 4 shows results for two outcome variables—life expectancy (panel A) and growth in GDP per capita (panel B)—in two samples—the main sample (columns 1 to 4) and sub-Saharan Africa alone (columns 5 to 8). For each outcome variable in each sample, I show four point estimates: the raw correlation between the outcome variable and ARV therapy coverage; OLS estimates of Equation (2), which include year and country fixed effects and country-specific trends; reduced form estimates, in which I substitute predicted ARV therapy coverage for observed ARV therapy coverage in Equation (2); and the IV estimates, which instrument for observed ARV therapy coverage using predicted ARV therapy coverage.

Comparing the raw correlations (columns 1 and 5) to the coefficients from the regressions with control variables (columns 2 and 6) shows that including the fixed effects and country-specific trends makes an important difference to the estimated relationships. In particular, the sign of the relationship between ARV therapy coverage and life expectancy reverses in the main sample.

The reduced form results (columns 3 and 7) show statistically significant relationships consistent with positive effects of ARV therapy coverage on both life expectancy and growth in GDP per capita.

The IV point estimates are shown in columns 4 and 8. The estimated effect on life expectancy is 6.7% (90% CI: 4.6–8.7) in the main sample and 6.1% (90% CI: 3.9–8.3) in sub-Saharan Africa. I report the estimated effect on life expectancy to confirm that ARV therapy coverage is correlated with contemporaneous changes in health. The estimated effects on life expectancy also allow me to scale the estimated effects on income, enabling comparison with other studies in section 6.2. However, the estimated effects on life expectancy are subject to the caveat about modelling adjustments described in section 3.

The estimated effect on growth in GDP per capita in the main sample is 1.40 percentage points (90% CI: 0.26–2.53) for a 1 percentage point increase in ARV therapy coverage. In sub-Saharan Africa, the corresponding estimate is 1.25 percentage points (90% CI: 0.05–2.45). The point estimate in sub-Saharan Africa is close in magnitude to the estimate from the main sample, although slightly smaller and slightly less precisely estimated, remaining statistically significant at the 10% level.[25]

For both outcome variables, the IV estimates are consistent in sign with but larger than the OLS estimates. Hausman tests (following Wooldridge, 2002) reject equality of the IV and OLS estimates for life expectancy, but not for GDP per capita, although the sample may simply be too small to detect differences between the OLS and IV estimates. This pattern of results suggests that observed ARV therapy coverage may be correlated with other factors that are negatively associated with changes in trends in economic growth, resulting in downward bias in the OLS estimates. For example, international support for ARV therapy coverage expansion may have been concentrated in countries

[24] Appendix Table B4 repeats this exercise for population growth and growth in total GDP, similarly finding no significant quadratic trend term. Appendix Tables B5a and B5b repeat the exercise including years up to and including 2003, showing that trends continued to be relatively stable during the initial period when ARV therapy coverage expansion was slow.

[25] The point estimates also remain positive and statistically significant in a sample that excludes sub-Saharan Africa, but the instrument is much weaker, and the confidence intervals are extremely wide.

Table 4
Impact of increasing ARV coverage on life expectancy and growth in per capita GDP.

	Main sample				Sub-Saharan Africa			
	OLS	OLS	OLS	IV	OLS	OLS	OLS	IV
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Log life expectancy								
% of population	−2.29*	5.57***		6.67***	1.39	4.96***		6.10***
treated with ARVs	(1.21)	(0.76)		(1.26)	(0.91)	(0.74)		(1.34)
HIV Prevalence$_{2001}$×			4.85***				4.72***	
global ARV coverage$_t$			(0.52)				(0.62)	
First-stage F-statistic				116				98
Hausman test p value				0.047				0.094
N	2146	2146	2160	2146	926	926	936	926
Panel B: Change in log GDP per capita								
% of population	0.24**	1.14**		1.40**	0.36***	0.97*		1.25*
treated with ARVs	(0.11)	(0.52)		(0.69)	(0.13)	(0.52)		(0.73)
HIV Prevalence$_{2001}$×			0.99**				0.94*	
global ARV coverage$_t$			(0.44)				(0.49)	
First-stage F-statistic				124				106
Hausman test p value				0.365				0.401
N	2222	2222	2236	2222	964	964	974	964
Controls	No	Yes	Yes	Yes	No	Yes	Yes	Yes

Note: Coefficients from regressions of outcome variable on listed variables. Data from 1990 to 2014. Main sample consists of 90 low- and middle-income countries (39 in sub-Saharan Africa). Data are from 1990 to 2014. Controls comprise year and country fixed effects, and country-specific linear trends. In columns 4 and 8, % of population treated by ARVs is instrumented with predicted ARV therapy coverage, equivalent to global coverage of ARVs interacted with HIV prevalence in 2001. Hausman test tests exogeneity of observed ARV therapy coverage. Standard errors clustered by country and shown in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

with relatively worse growth prospects. Alternatively, the scale of a country's ARV therapy program, conditional on HIV prevalence, may be correlated with higher than average growth in the years preceding widespread availability of ARV therapy, which would bias the OLS estimates downwards through regression to mean.

The effect on growth in GDP per capita can be decomposed into effects on growth in total GDP and on population growth (Appendix Table B6). The estimated effect on population growth is positive: 0.20 percentage points (90% CI: 0.09–0.33) in the main sample, for a 1 percentage point change in ARV therapy coverage, and 0.37 percentage points in Sub-Saharan Africa alone (90% CI: 0.19–0.55).[26] Correspondingly, the estimated effects on total GDP are larger than those on GDP per capita: 1.64 percentage points (90% CI: 0.44–2.84) in the main sample and 1.67 percentage points (90% CI: 0.35–2.99) in sub-Saharan Africa alone. The effect on growth in GDP per capita therefore occurs despite positive estimated effects on population growth. While it remains possible that the estimated effects on population growth are contaminated by modelling adjustments to aggregate data, as discussed in Section 3, the true effects on population growth would need to be several times larger in magnitude than the estimated effects to explain the growth in total GDP in the absence of growth in GDP per capita.

The marginal effect of further increasing ARV therapy coverage is likely to decrease, given that the most vulnerable individuals and those with the highest viral loads are treated first, and given diminishing marginal returns to increasing the size of the healthy workforce. Also, ceterus paribus, the marginal effect of increasing population average ARV therapy coverage should be highest in countries with highest HIV prevalence. Estimates from augmented models that allow the effect of increasing ARV therapy coverage to experience diminishing marginal returns and to vary with HIV prevalence take the expected signs (Appendix Tables B8 and B9).

[26] Appendix Figure B3 and Appendix Table B7 show that the estimated effect on population growth is primarily driven by changes in the crude death rate, with a substantial component from changes in the birth rate in sub-Saharan Africa. However, the mortality and fertility data are subject to the same overall caveats regarding adjustments, so these estimates should be interpreted with similar caution.

6.1. Persistence of effects

The results shown in Table 4 capture the average effect of ARV therapy coverage on growth. However, increasing ARV therapy coverage is unlikely to lead to persistent increases in growth rates. A more likely scenario is that a change in ARV therapy coverage affects growth only over a period of a few years. The total effect on income will be the cumulative sum of the effects on growth. To draw the correct conclusions about the impact of changes in ARV therapy coverage on income, we would like to understand how long effects on growth are likely to persist. This section summarizes the results of three exercises which together suggest that it is not unreasonable to assume that the positive effects on growth persist for around four years. If so, the cumulative effect on income of a 1 percentage point change in ARV therapy coverage is 5–5.6 percentage points. More details are provided in Appendices A.5 and A.6.

First, I estimate a distributed lag model, which generalizes Equation (2) to allow the effects of changes in ARV therapy coverage on growth to differ over time. I can only estimate the distributed lag model by OLS,[27] and the OLS estimates of the distributed lag model may not be unbiased, for the same reasons that OLS estimates of Equation (2) are not necessarily unbiased. However, the OLS estimates of the distributed lag model describe empirically how growth responds after a change in ARV therapy coverage. Estimating the distributed lag model with a varying number of lags consistently suggests positive impacts on growth over a total of four years.

Second, I show that under a set of simple conditions, assuming that the estimated growth effects persist for four years leads to conservative conclusions about the total effect on income. More specifically, I show that the total estimated 4-year effect on income approximates or underestimates the true cumulative effect on income if three conditions hold: i) a change in ARV therapy coverage leads to a constant increase in income growth rates over a fixed number of years k_{true}, ii) that a change in ARV therapy coverage does not affect income growth after

[27] Attempting to estimate the distributed lag model using an IV strategy that is conceptually similar to the IV strategy in the main analysis leads to a weak instrument problem. See Appendix A.5 for details.

Table 5
Robustness of instrumental variables analysis to specification changes.

	Change in log GDP per capita						Log GDP per capita
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: Main sample							
% of population treated with ARVs	1.40**	1.25*	2.60	1.79			17.07**
	(0.69)	(0.73)	(1.71)	(1.26)			(7.62)
First-stage F-statistic	124	103	24	13			17
N	2222	2222	2222	2172			2229
Panel B: Sub-Saharan Africa							
% of population treated with ARVs	1.25*		2.49	2.00	1.34*	1.25	13.17*
	(0.73)		(1.71)	(1.79)	(0.75)	(0.86)	(7.40)
First-stage F-statistic	106		32	7	152	15	13
N	964		964	939	987	938	964
Additional time trends	None	Region- year FE	Quad. XY-region- year FE	None	None	None	Country quad. trends
Measure of HIV prevalence	Observed 2001	Observed 2001	Observed 2001	Observed 1990	Simulated 2001	Predicted 2001	Observed 2001

Note: Coefficients from 2SLS IV regressions of outcome variable on listed variable, country and year fixed effects, country linear trends and additional time trend variables where specified. % of population treated by ARVs is instrumented with predicted ARV therapy coverage, equivalent to global coverage of ARVs interacted with a measure of HIV prevalence. Main sample consists of 90 low- and middle-income countries (39 in sub-Saharan Africa). Data are from 1990 to 2014. Simulated data for HIV prevalence in 2001 generously provided by Cahu and Fall (2011). Predicted HIV prevalence in 2001 is a linear prediction using Halperin-Bailey circumcision rates reported in Ahuja et al. (2009). Standard errors clustered by country and shown in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

k_{true} years have passed, and iii) that k_{true} is at least two years.

Third, I estimate an alternative model with log GDP per capita as the outcome variable, which includes country-specific quadratic trends (Table 5, column 7). This model directly estimates the cumulative effect of ARV therapy coverage on income but is otherwise equivalently conservative to the main model, in that it fully accounts for differences across countries in income, growth rates, and linear trends in growth rates. The estimated total cumulative effects on income vary between 13 and 17 percentage points, albeit with wide confidence intervals, suggesting that it is conservative to assume that growth effects persist for four years, with cumulative effects of 5–5.5 percentage points. These results additionally confirm that the positive effects on GDP per capita are robust to estimation using this alternative model.

6.2. Magnitude of effects

The estimated effects on growth are relatively large: increasing population coverage with ARV therapy by an additional 1% increases GDP per capita growth rates by 1.2–1.4 percentage points. If these effects persist for four years, the total change in income is 5–5.6 percentage points. However, confidence intervals are wide. While they exclude negative effects, they do not exclude much smaller positive effects: the 90% confidence interval for the 4-year effects is 1.0–10.0 percentage points in the main sample and 0.2–9.8 percentage points in sub-Saharan Africa. The effects on growth are also estimated relative to a very low baseline income.[28] Taking baseline incomes as fixed, the estimated effects on growth translate into an increase in per capita income of $23 (90% CI: 4.4–42) in the main sample or $14 (90% CI: 0.6–27) in sub-Saharan Africa. The 4-year effects suggest total gains in income of $94 (90% CI: 18–170) in the main sample or $56 (90% CI: 2–109) in sub-Saharan Africa.

The estimated effects on growth are nonetheless important relative to recent growth rates in sub-Saharan Africa. Observed growth in sub-Saharan Africa averaged 0.2% between 1990 and 2001, and 2.1% between 2002 and 2014. Extrapolating from the estimated model suggests that growth in sub-Saharan Africa would have been 1.4% (90% CI 0.1–2.0), or around a third lower, over the latter period under a counterfactual scenario with no ARV therapy coverage expansion and an unchecked HIV epidemic (Appendix A.7). Other factors have undoubtedly also been at play in sub-Saharan Africa's recent recovery (see e.g. Rodrick, 2018). However, these estimates suggest an important role for ARV therapy coverage expansion in explaining sub-Saharan Africa's "growth miracle".

The estimated effects on life expectancy are also very substantial: increasing population coverage with ARV therapy by 1% is associated with a rise in life expectancy of 6.1–6.7%, or 3.2–4.1 years.[29] These effects are consistent in magnitude with previous literature, although the caveats in Section 3 continue to apply.[30] Normalizing the effects on income (assuming four years persistence) by the effects on life expectancy yields an elasticity of per capita income to life expectancy of just over 0.8 in both the main sample and sub-Saharan Africa alone.[31] For comparison, Acemoglu and Johnson (2007), estimate the long-run elasticity of per capita income to life expectancy to lie between −1.2 and −2.7, depending on the specification. As discussed, there are several reasons why the estimated elasticities might differ in sign, but it is reassuring that the short-run elasticities estimated here are smaller

[28] Mean GDP per capita in 2001 was $1679 in the main sample and $1113 in sub-Saharan Africa (2005 US$).

[29] Baseline mean life expectancy is 61.4 years in the main sample and 52.0 years in sub-Saharan Africa.

[30] Bor et al. (2013) find a 3.0% increase in life expectancy for every 1% of the population treated with ARV therapy in a high HIV prevalence region of South Africa. However, Bor et al. (2013) note that their estimates are biased downwards by a clear prior downward trend in life expectancy, for which their analysis does not account. An approximate adjustment to their results yields a figure of 4.2%, within the confidence interval of my estimated effects on life expectancy in sub-Saharan Africa. It is possible, however, that Bor et al.'s estimates inform the demographic models used in adjusting aggregate health data, in which case I would mechanically obtain similar results. Breaking down effects on mortality by age and gender (see Appendix A.8) does show that effects are consistent with those that the health literature would predict for a direct causal effect of ARV therapy. However, the same concern about adjustment applies to all measures of adult mortality, meaning that these comparisons only add limited further information regarding the effect of ARV therapy on health.

[31] Estimating these elasticities directly by using life expectancy as the endogenous variable in Equation (2), as in Acemoglu and Johnson (2007), also yields very similar elasticities. Results available on request.

in magnitude than the long-run elasticities estimated by Acemoglu and Johnson.

The fact that the short-run effects of ARV therapy coverage on income are positive does not imply that the long-run effects on income will necessarily also be positive. If the effects on population growth persist, they might attenuate or even reverse the initial gains in income. In Appendix A.9, I compare the magnitude of the effects on income and population growth through the lens of a simple growth model. The range of predicted outcomes includes both substantial positive effects and moderate negative effects. The long-run consequences of ARV therapy coverage expansion therefore remain indeterminate.

6.3. Possible mechanisms

The previous sections show that ARV therapy coverage expansion led to growth in per capita income. We can distinguish between three potential classes of mechanism: transfer effects; direct labour force effects, which arise when individuals treated with ARV therapy recover and return to the labour force; and indirect (or externality) effects. In this section, I evaluate the feasible contribution of each of these three classes of mechanism to the overall effect on income, focusing on sub-Saharan Africa. Appendix A.10 provides more details. Transfer effects and direct labour force effects are unlikely to explain the full effect, suggesting a potentially substantial role for indirect or externality effects.

Transfer effects include externally-funded spending associated with HIV treatment programs, including but not limited to ARV therapy. A 1 percentage point increase in ARV therapy coverage constitutes a transfer of at most around $9 per capita per year. An approximate upper bound on direct labour market effects on GDP per capita is $21. Transfer effects and direct labour market effects are therefore unlikely to account for more than around $30 in per capita income gains. This is larger than the estimated one-year effects, but considerably smaller than the 4-year effects, although the 4-year effects have wide confidence intervals that include these values.

The results therefore suggest a role for indirect effects of a similar magnitude to the combined transfer and direct labour force effects. Potentially important mechanisms include indirect labour market effects, such as funeral attendance or effects on the labour supply of carers of HIV-infected individuals.[32] HIV treatment programs may also have reduced congestion in other public health services, improving health and increasing labour market participation for patients suffering from other conditions (Creese et al., 2002; Médecins Sans Frontières, 2015). Finally, indirect effects may also include changes in labour market or investment decisions that result from changes in beliefs about the risk and consequences of infection with HIV (Baranov et al., 2015; Baranov and Kohler, 2018).

7. Robustness checks and alternative explanations

This section shows that the results are robust to a series of tests and specification checks. I also examine four alternative explanations and conclude that none of them can account for the main results. In the interests of brevity, I focus on the instrumental variables analyses for the main outcome variable, growth in GDP per capita. The analyses of growth trends are similarly robust to the same series of tests and specification checks, as are the estimated effects on life expectancy, population growth, and growth in total GDP.

7.1. Robustness checks

Table 5 shows the main robustness checks for the main sample (Panel A) and for sub-Saharan Africa alone (Panel B). Column 1 replicates the main results from Table 4 for comparison.[33]

Time trends Column 2 shows that results for the main sample using region-year fixed effects are very similar to the results for sub-Saharan Africa alone, which is to be expected given that most of the variation of interest is within sub-Saharan Africa.

A natural concern, given spatial correlation in HIV prevalence, is whether the results could be driven by different trends within regions. In particular, one might worry that the results are driven by different trends in East and Southern Africa, where HIV prevalence is highest, and West Africa, where HIV prevalence is relatively low. Column 3 show that accounting flexibly for with-region trends yields larger, although less precise, estimates.[34] The estimated effect on growth becomes marginally insignificant (p values 0.127 in the main sample and 0.145 in sub-Saharan Africa). However, the increase in the point estimates suggest that the results are not driven by within-region variation in trends.

Regression to mean Section 4 shows that HIV prevalence in 2001 is associated with declining growth rates in the 1990s. I interpret this correlation as suggestive of negative effects of HIV/AIDS on growth trends. However, the same correlation could also arise if declining growth rates in the 1990s themselves led to higher HIV prevalence in 2001, for example through changes in risky sexual behaviour. If this were the case, then the estimated effects of ARV therapy coverage could be biased upwards by regression to mean in high HIV prevalence countries. To evaluate whether regression to mean drives the results, I repeat the analysis using three alternative measures of HIV prevalence to construct the instrument. All three alternative measures of HIV prevalence are, by construction, exogenous to growth in the 1990s.

The first alternative measure of HIV prevalence is HIV prevalence in 1990, the earliest year for which HIV prevalence is estimated. If the main results were biased upwards because of regression to mean, using an earlier measure of HIV prevalence in constructing the instrument should decrease this bias. In fact, using the earlier measure of HIV prevalence yields larger, albeit less precisely estimated, point estimates (column 4).[35] The first stage results are, however, considerably weakened. The second alternative measure of HIV prevalence is obtained from a simulation that aims to isolate the independent component of the epidemic's propagation dynamic in sub-Saharan Africa.[36] Using the simulated values of HIV prevalence also yields slightly larger point estimates (column 5). Finally, the third alternative measure of HIV prevalence uses variation in male circumcision rates to predict HIV prevalence. HIV prevalence is lower in countries with higher male circumcision rates.[37] The first stage is weakened and the standard errors are

[32] Note that not all such effects necessarily have the same sign: if an adult returns to the labour force, it might enable a younger family member to exit the labour force and rejoin full-time education.

[33] Appendix Tables B10a to B11d replicate the same set of robustness checks for the full set of outcome variables and the trend analyses. Appendix Tables B12a to B13d report additional robustness checks mentioned throughout the text.

[34] I allow time trends to vary with longitude and latitude within each region, specifically by including a full quadratic (second degree polynomial) function of longitude and latitude, interacted with region-year fixed effects. I obtain similar results with linear controls for longitude and latitude interacted with region-year fixed effects (Appendix Tables B12a to B13d, column 4).

[35] Data are missing for HIV prevalence in 1990 in a small subset of countries. Estimates using the original instrument, based on HIV prevalence in 2001, in the subset of countries for which HIV prevalence in 1990 is reported yield slightly smaller point estimates in the main sample and almost identical results within sub-Saharan Africa.

[36] I thank Cahu and Fall for generously sharing these data with me. See Cahu and Fall (2011) for a full description.

[37] I use data on male circumcision rates collated by Halperin and Bailey (1999) and Shelton (2002) as reported in Ahuja et al. (2009). I use Ahuja et al.'s preferred measure of male circumcision. Appendix A.11 also reports results using an alternative measure of male circumcision based on ethnic demographics.

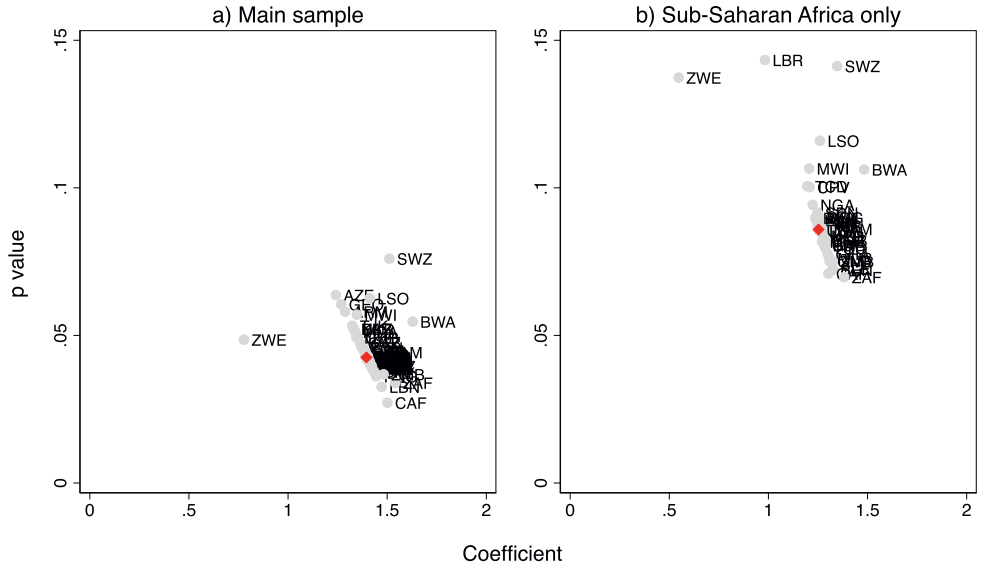

Notes Markers plot coefficients on ARV therapy coverage and associated p values from repeated 2SLS regressions in the main sample of 90 low and middle-income countries (panel a) or 39 low and middle-income countries in sub-Saharan Africa (panel b). Each marker corresponds to results from a regression excluding the labelled country, in which changes in log GDP per capita are regressed on observed ARV therapy coverage, instrumented with predicted ARV therapy coverage, year and country fixed effects, and country-specific linear time trends. Standard errors clustered by country. Main estimates shown as red diamonds.

Fig. 5. Leave-one-out estimates: Instrumental variables analysis.

larger, but the point estimate is identical to the main specification (column 6).

Since the results are insensitive to using these alternative measures of HIV prevalence to construct the instrument, it is unlikely that the results can be explained by regression to mean.

Small-sample inference The number of countries with high HIV prevalence is small, and the distribution of HIV prevalence is highly skewed, raising the concern that individual countries could be highly influential in the analysis (see e.g. Young, 2019). I run a series of regressions, identical to the main specification, but in each case excluding one country. The resulting estimates of the effect of ARV therapy on the main outcome variables (coefficients and p values) are plotted in Fig. 5, for both the main sample and sub-Saharan Africa alone.[38] Each set of estimates shows what conclusion we would have drawn without data from that particular country.

The countries with highest HIV prevalence in 2001—particularly Botswana and Zimbabwe—have the greatest influence on the estimates, although their influence is somewhat balanced: excluding Zimbabwe decreases the magnitude of the estimated impacts, while excluding Botswana increases them. In the main sample, the estimated impact nonetheless remains consistent in sign and the 90% confidence interval excludes zero when any individual country is excluded from the analysis. For the sample within sub-Saharan Africa, the estimated impact also remains consistent in sign, but p values increase to between 0.10 and 0.15 when a number of individual countries are excluded, meaning that the 90% confidence intervals no longer exclude zero. However, the IV estimates have large standard errors and these analyses discard two potentially important sources of variation, the between-region variation *and* the individual influential country. The 90% confidence intervals for the more precisely-estimated change in trends associated with HIV prevalence continues to exclude zero when any individual country is excluded, in both the main sample and the sub-Saharan Africa subsample.[39]

The small number of affected countries and skewed distribution of HIV prevalence might also raise concerns that asymptotic distributions of estimators may not apply. I confirm that the p values generated using analytical standard errors are consistent with those I obtain from randomization-based inference. I randomly permute HIV prevalence across countries, both in the main sample and within sub-Saharan Africa alone, and re-run the analysis 500 times. The p values calculated using this method are generally similar to or, for sub-Saharan Africa, smaller than those calculated using clustered standard errors. Randomization-based inference yields p values of 0.048 for the main sample and 0.040 for sub-Saharan Africa alone, while the estimates with analytical standard errors yield p values of 0.043 and 0.086, respectively.[40]

Additional robustness checks Appendix A.12 provides an extended discussion of how the empirical model accounts for Solovian convergence and why convergence cannot explain the results. Appendix A.13 replicates the analysis with six alternative measures of growth, obtaining broadly consistent results. Appendix A.14 describes the limitations of alternative approaches such as matched controls in this context. Appendix A.15 shows that the results are not sensitive to the exact time frame used in the analysis. Appendix A.16 shows that the results are insensitive to omitting countries with very low HIV prevalence,

[38] Appendix Figures B5a to B6d repeat this exercise for the trend analysis and the other outcome variables. Note that the problem is specifically the skewed distribution of disease intensity, not the use of cross-sectional variation in disease intensity in the instrument, which is common to much of the literature on health and income (e.g. Acemoglu and Johnson, 2007).

[39] See Appendix Figure B6a.

[40] For the change in HIV-specific growth trends, randomization-based inference yields larger p values in the main sample, but smaller p values for sub-Saharan Africa alone: specifically, p values of 0.072 for the main sample and 0.028 for sub-Saharan Africa alone, compared to 0.019 and 0.035 using clustered standard errors, respectively. Appendix Figures B7a to B8d show the distribution of point estimates from these exercises.

Table 6
Impact of ARVs on growth in per-capita GDP: Alternative explanations.

	Main sample			Sub-Saharan Africa		
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Mineral boom						
% of population treated with ARVs	1.40**	1.49**	1.49**	1.25*	1.29*	1.29*
	(0.69)	(0.69)	(0.69)	(0.73)	(0.75)	(0.75)
Fraction GDP mineral rents$_{2001}$ × global ARV coverage$_t$		2.14***			0.82	
		(0.81)			(0.60)	
First-stage F-statistic	124	124	123	106	104	101
N	2222	2222	2222	964	964	964
Panel B: Petroleum boom						
% of population treated with ARVs	1.39**	1.33*	1.33*	1.20	1.24*	1.24*
	(0.69)	(0.68)	(0.68)	(0.73)	(0.74)	(0.74)
Fraction GDP petroleum rents$_{2001}$ × global ARV coverage$_t$		−0.28			0.20	
		(0.43)			(0.21)	
First-stage F-statistic	124	124	123	105	105	103
N	2197	2197	2197	939	939	939
Panel C: Export boom						
% of population treated with ARVs	1.40**	1.37*	1.36*	1.25*	1.10	1.09
	(0.69)	(0.72)	(0.72)	(0.73)	(0.89)	(0.89)
Fraction GDP exports$_{2001}$ × global ARV coverage$_t$		0.05			0.13	
		(0.15)			(0.19)	
First-stage F-statistic	124	138	137	105	111	108
N	2197	2197	2197	939	939	939
Panel D: Contemporaneous investment in fighting malaria						
% of population treated with ARVs	1.40**	1.14*	1.14*	1.25*	1.31	1.31
	(0.69)	(0.65)	(0.65)	(0.73)	(0.83)	(0.83)
Malaria ecology × global ARV coverage$_t$		0.16			0.03	
		(0.10)			(0.14)	
First-stage F-statistic	124	123	122	106	82	80
N	2222	2222	2222	964	964	964
Additional trends	None	Linear	Flexible	None	Linear	Flexible

Note: Coefficients from 2SLS IV regressions of outcome variable on listed variables, year and country fixed effects, and country-specific linear trends. % of population treated by ARVs is instrumented with predicted ARV therapy coverage, equivalent to global coverage of ARVs interacted with HIV prevalence in 2001. Main sample consists of 90 low- and middle-income countries (39 in sub-Saharan Africa). Data are from 1990 to 2014. Linear controls are as specified in the table; flexible controls comprise interactions between the additional control variable and year dummies. Standard errors clustered by country and shown in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

remaining stable in samples of countries with HIV prevalence above progressively more conservative thresholds.

7.2. Alternative explanations

The results could still be biased by omitted variables if their effects were correlated across space with HIV prevalence and in time with ARV therapy coverage expansion. In this section, I evaluate four potential channels through which this type of omitted variables bias could arise and show that none can explain the main results.[41] The first three alternative explanations I evaluate relate to trade and capital flows between Africa and China, which follow a similar pattern in time to ARV therapy coverage expansion (see e.g. Chinese Ministry of Commerce, 2009). If HIV prevalence correlates with Chinese trade or capital flows, and if increased trade or capital flows in turn influenced growth, then the main results could be biased. The fourth alternative explanation concerns investment in malaria control, which increased in tandem with investment in HIV treatment (WHO, 2015b).

I test whether the main results change when I allow trends to vary with additional control variables that predict Chinese trade or capital flows, or malaria prevalence. Specifically, I include additional control variables interacted either with global ARV therapy coverage or with a full set of year fixed effects.[42] If accounting for these heterogeneous trends does not alter the estimates, or increases the estimated effect, then it is unlikely that the alternative explanation is valid.

The three additional control variables I consider which predict Chinese trade and/or capital flows are mineral resources, petroleum resources, and export dependence. Chinese FDI is correlated with the local availability of commodities (Cheung et al., 2012; Kolstad and Wiig, 2012). Countries with high export dependence are most likely to have benefited from increased trade with China. All three variables might also be correlated with HIV prevalence. Mines attract large populations of transient workers, who have high rates of HIV prevalence (Corno and de Walque, 2012). Countries with petroleum resources or high export dependency might have higher population mobility, either through lower fuel prices or the internal transport of goods, and thus higher rates of HIV transmission and infection (see e.g. Oster, 2012).

The additional control variable which predicts malaria prevalence is malaria ecology.[43] Malaria ecology captures the variation in malaria prevalence that is driven by climate and ecology and thus independent

[41] Further details are in Appendix A.17. Corresponding results for other outcome variables and for the trend analyses are in Appendix Tables B14a to B15d.

[42] This approach is preferable to controlling directly for proxies for potential time-varying omitted variables, as many of these might also constitute "bad controls" (Angrist and Pischke, 2009), in which case their inclusion could introduce bias into the estimates. If HIV prevalence were sufficiently correlated with the additional control variables, then the effects of ARV therapy might be inseparable from other trend changes associated with these additional control variables. As it happens, the correlation between HIV prevalence and the additional control variables is weaker in most cases than we might expect.

[43] The malaria ecology index was developed by Kiszewski et al. (2004) and aggregated to a population-weighted country dataset by McCord (n.d.).

of human behaviour. Perhaps contrary to expectations, while malaria ecology is positively correlated with HIV prevalence in the main sample, it is negatively correlated with HIV prevalence in sub-Saharan Africa.

Table 6 shows results both for the main sample (columns 1–3) and for sub-Saharan Africa only (columns 4–6). Columns 1 and 4 restate results from the main specification for reference, estimated on the sample with non-missing data for the additional control variable. Columns 2 and 5 show results which control for interactions between the additional control variable and ARV therapy coverage. Columns 3 and 6 show results which allow for fully flexible heterogeneous trends i.e. the interactions between the additional control and year fixed effects.

The results do not seem to be explained by differential trends in countries likely to have benefited from increased Chinese trade or capital flows. Panels A to C show that the estimated effects of ARV therapy coverage expansion are mostly either larger or very similar when I allow trends to vary as a function of mineral resources, petroleum resources, or export dependence. The results are most sensitive to allowing trends to vary with export dependence in the sub-Saharan Africa sample, where the estimated effects of ARV therapy coverage become statistically insignificant (p values 0.221 and 0.222, in columns 5 and 6 respectively), in part because the estimated coefficient is smaller, and in part because standard errors are larger. However, allowing trends to vary with export dependence has little effect on the estimates in the main sample.[44] Export dependence also has little independent explanatory power over trends.

Nor can the results be explained by differential trends in countries affected by malaria. Panel D shows that allowing trends to vary with malaria ecology decreases the estimated effects in the main sample but increases the estimated effects in sub-Saharan Africa. Allowing trends to vary with malaria ecology does render the estimate marginally insignificant within sub-Saharan Africa alone, but only because standard errors also increase (p values 0.112 and 0.114, in columns 5 and 6 respectively).[45] This pattern of results is consistent with a positive effect of increased malaria control on growth. However, any such effect explains at most a small fraction of the results in the main sample and cannot explain the results within sub-Saharan Africa.

8. Conclusion

This paper evaluates the effect of rapid, massive expansion in antiretroviral (ARV) therapy coverage on income growth in low- and middle-income countries. ARV therapy coverage expansion was triggered by a sharp, tenfold fall in ARV drug prices in 2001, after generic production of ARV drugs went internationally unopposed.

Before 2001, high HIV prevalence is associated with declining growth rates in GDP per capita, relative to other low- and middle-income countries. After 2001, these negative relative trends reverse. Structural tests pin the timing of these trend reversals down to a narrow window around the fall in prices in 2001.

Extending ARV therapy coverage to an additional 1% of a country's population is associated with a 6.7% increase in life expectancy (90% confidence interval: 4.6–8.7) and leads to a 1.4 percentage point increase in growth rates in GDP per capita (90% CI: 0.3–2.5). The estimated effect on income growth most likely persists for around four years. At mean per capita income in sub-Saharan Africa in 2001, the change in income is valued at $14 (90% CI: 0.6–27). If this effect persists for four years, the total increase in income is $56 (90% CI: 2–109). The confidence intervals on the estimated effects are, however, wide. While they exclude negative effects, they include both small and large positive effects. An accounting exercise suggests that growth in sub-Saharan Africa would have been around a third lower in the absence of ARV therapy coverage expansion, pointing to the reversal of the negative effects of the HIV/AIDS epidemic as a previously neglected explanatory factor behind sub-Saharan Africa's "growth miracle".

The positive estimated effects on income growth contrast with previous literature. There are three possible reasons. First, the unbiasedness of the estimates in this study depends on the assumption that countries with high and low HIV prevalence did not experience any *different, simultaneous, and unrelated changes in trends* in income growth rates that coincide with ARV therapy coverage expansion. Previous studies depend on stronger identifying assumptions. Second, population health improvements that improve adult health are more likely to have positive effects on income growth than those that primarily reduce infant mortality. Previous literature focuses on health improvements with large effects on infant mortality, while HIV/AIDS primarily affects working age adults. Finally, I measure short-run effects, which are more likely to be positive than long-run effects. An outstanding open question is whether the short-run positive effects on income will persist or be attenuated—or even reversed—by population growth over the long run.

The underlying mechanisms behind the effects on income growth remain an open question for future research. However, the magnitude of the estimated effects suggests a role for indirect or spillover effects, alongside transfer effects from international funding of HIV treatment programs and direct labour market effects associated with the return to fitness of those receiving ARV therapy.

Policy-makers continue to debate how to finance HIV treatment programs, including ARV therapy coverage. The current total cost of provision in sub-Saharan Africa is around $9.4 per capita per year for every 1% of the population treated with ARV therapy. The estimated benefits monetized in economic growth in a single year are at least 50% higher than these costs and are larger if the effects on growth persist over several years. At least historically, the income growth benefits of ARV therapy provision have therefore most likely exceeded the costs.

Expanding coverage further may have smaller marginal effects on income growth. ARV therapy treatment programs initially targeted individuals with high viral loads who may already have been symptomatic, whereas future coverage expansion will target individuals with lower viral loads who may still be asymptomatic, based on new treatment protocols. Additionally, we should expect diminishing marginal returns to increasing the size of the healthy labour force.

The estimated effects on income understate the true welfare effects of ARV therapy coverage, as they do not capture the direct benefits of improving health, saving lives and avoiding orphaning children, or the externality effects of avoiding new infections via interrupted transmission (Palmer et al., 2008). Gains in income may be lost if ARV therapy coverage levels are not maintained. A clear and urgent policy implication of the results is that failure to maintain present levels of ARV therapy coverage—or to address emerging problems with ARV drug resistance—could have severe economic as well as human consequences.

Appendices and Replication data

Supplementary data to this article can be found online at https://doi.org/10.1016/j.jdeveco.2019.102409.

References

Acemoglu, D., Johnson, S., 2007. Disease and development: the effect of life expectancy on economic growth. J. Political Econ. 115 (6).
Acemoglu, D., Johnson, S., 2014. Disease and development: a reply to Bloom, Canning, and Fink. J. Political Econ. 122 (6).

[44] In the analysis of trend breaks, the results also remain statistically significant in the main sample and become marginally insignificant in the sub-Saharan Africa subsample (p values 0.119 and 0.139). See Appendix Table B15a.

[45] In the analysis of trend breaks, the results remain statistically significant both in the main sample in the sub-Saharan Africa subsample. See Appendix Table B15a.

Ahuja, A., Wendell, B., Werker, E., 2009. Male Circumcision and AIDS: the Macroeconomic Impact of a Health Crisis. Harvard Business School. Working Paper 07-025.
Angrist, J.D., Pischke, J., 2009. Mostly Harmless Econometrics. Princeton University Press.
Ashraf, Q.H., Lester, A., Weil, D.N., 2008. When does improving health raise GDP? In: Acemoglu, D., Rogoff, K., Woodford, M. (Eds.), NBER Macroeconomics Annual 2008, vol. 23. University of Chicago Press. Chap. 3.
Baranov, V., Bennett, D., Kohler, H.-P., 2015. The indirect impact of antiretroviral therapy: mortality risk, mental health, and HIV-negative labor supply. J. Health Econ. 44.
Baranov, V., Kohler, H.-P., 2018. The impact of AIDS treatment on savings and human capital investment in Malawi. Am. Econ. J. Appl. Econ. 10 (1).
Barré-Sinoussi, F., Chermann, J., Rey, F., Nugeyre, M., Chamaret, S., Gruest, J., ... Montagnier, L., 1983. Isolation of a T-lymphotropic retrovirus from a patient at risk for acquired immune deficiency syndrome (AIDS). Science 220 (4599).
Bendavid, E., Bhattacharya, J., 2009. PEPFAR in Africa: an evaluation of outcomes. Ann. Intern. Med. 150 (10).
Bertrand, M., Duo, E., Mullainathan, S., 2004. How much should we trust differences-in-differences estimates? Q. J. Econ. 119 (1).
Bhargava, A., Jamison, D.T., Lau, L.J., Murray, C.J., 2001. Modeling the effects of health on economic growth. J. Health Econ. 20.
Bleakley, H., 2007. Disease and development: evidence from hookworm eradication in the American South. Q. J. Econ. 122 (1).
Bleakley, H., 2010a. Health, human capital, and development. Annual Review of Economics 2.
Bleakley, H., 2010b. Malaria eradication in the Americas: a retrospective analysis of childhood exposure. Am. Econ. J. Appl. Econ. 2.
Bloom, D.E., Canning, D., Fink, G., 2014. Disease and development revisited. J. Political Econ. 122 (6).
Bloom, D.E., Canning, D., Sevilla, J., 2004. The effect of health on economic growth: a production function approach. World Dev. 32 (1).
Bor, J., Herbst, A., Newell, M.-L., Bärnighausen, T., 2013. Increases in adult life expectancy in rural South Africa: valuing the scale-up of HIV treatment. Science 339.
Bor, J., Tanser, F., Newell, M.-L., Bärnighausen, T., 2012. In a study of a population cohort in South Africa, HIV patients on antiretrovirals had nearly full recovery of employment. Health Aff. 31 (7).
Cahu, P., Fall, F., 2011. Accounting for the effects of AIDS on growth in sub-Saharan Africa. Documents de Travail du Centre d'Economie de la Sorbonne.
Card, D., Lee, D.S., Pei, Z., Weber, A., 2015. Inference on causal effects in a generalized regression kink design. Econometrica 83 (6).
Cheung, Y.-W., de Haan, J., Qian, X., Yu, S., 2012. China's outward direct investment in Africa. Rev. Int. Econ. 20 (2).
Chinese Ministry of Commerce, 2009. Statistical Bulletin of China's Outward Foreign Direct Investment.
Corno, L., de Walque, D., 2012. Mines, migration and HIV/AIDS in southern Africa. J. Afr. Econ. 21 (3).
Creese, A., Floyd, K., Alban, A., Guinness, L., 2002. Cost-effectiveness of HIV/AIDS interventions in Africa: a systematic review of the evidence. The Lancet 359.
Cutler, D., Fung, W., Kremer, M., Singhal, M., Vogl, T., 2010. Early-life malaria exposure and adult outcomes: evidence from malaria eradication in India. Am. Econ. J. Appl. Econ. 2 (2).
Ford, N., Calmy, A., Mills, E.J., 2011. The first decade of antiretroviral therapy in Africa. Glob. Health 7 (33).
Fuchs-Schündeln, N., Hassan, T.A., 2016. Natural experiments in macroeconomics. In: Handbook of Macroeconomics, vol. 2. Elsevier.
Gallup, J.L., Sachs, J.D., 2001. The economic burden of malaria. Am. J. Trop. Med. Hyg. 64 (1,2).
Gulick, R.M., Mellors, J.W., Havlir, D., Eron, J.J., Gonzalez, C., McMahon, D., ... Condra, J.H., 1997. Treatment with indinavir, zidovudine, and lamivudine in adults with human immunodeficiency virus infection and prior antiretroviral therapy. N. Engl. J. Med. 337.
Habyarimana, J., Mbakile, B., Pop-Eleches, C., 2010. The impact of HIV/AIDS and ARV treatment on worker absenteeism: implications for African firms. J. Hum. Resour. 45 (4).
Halperin, D., Bailey, R., 1999. Male circumcision and HIV infection: 10 years and counting. Lancet 354 (9192).
Hammer, S., Squires, K., Hughes, M., Grimes, J., Demeter, L., Currier, J., ... Fischl, M., 1997. A controlled trial of two nucleoside analogues plus indinavir in persons with human immunodeficiency virus infection and CD4 cell counts of 200 per cubic millimeter or less. AIDS Clinical Trials Group 320 Study Team. N. Engl. J. Med. 337 (11).
Hansen, B.E., 2001. The new econometrics of structural change: dating breaks in U.S. labor productivity. J. Econ. Perspect. 15 (4).
Hansen, C.W., 2014. Cause of death and development in the US. J. Dev. Econ. 109.
Hansen, G.D., Prescott, E.C., 2002. Malthus to Solow. Am. Econ. Rev. 92 (4).
Keele, B.F., Van Heuverswyn, F., Li, Y., Bailes, E., Takehisa, J., Santiago, M.L., ... Hahn, B.H., 2006. Chimpanzee reservoirs of pandemic and nonpandemic HIV-1. Science 313 (5786), 523–526.
Kiszewski, A., Mellinger, A., Spielman, A., Malaney, P., Sachs, S.E., Sachs, J., 2004. A global index representing the stability of malaria transmission. Am. J. Trop. Med. Hyg. 70 (5).
Kolstad, I., Wiig, A., 2012. What determines Chinese outward FDI? J. World Bus. 47.
Korber, B., Muldoon, M., Theiler, J., Gao, F., Gupta, R., Lapedes, A., Bhattacharya, T., 2000. Timing the ancestor of the HIV-1 pandemic strains. Science 288 (1789).
McCord, G. (n.d.). Malaria ecology index dataset. Retrieved from https://sites.google.com/site/gordoncmccord//datasets.
McKinsey Global Institute, 2010. Lions on the Move: the Progress and Potential of African Economies.
McMillan, M.S., Harttgen, K., 2014. What Is Driving the 'African Growth Miracle'? NBER Working Paper No. 20077.
Médecins Sans Frontières, 2001. A Matter of Life and Death: the Role of Patents in Access to Essential Medicines.
Médecins Sans Frontières, 2015. Hivbehandling I Låaginkomstländer Inkl. Prevention Av Mödrabarnsmitta (HIV Treatment and Prevention of Maternal-Child HIV Transmission in Developing Countries). Lecture given by Kristina Bolme at Noah's Ark, Stockholm, February 18th 2015.
Mertens, T., Low-Beer, D., 1996. HIV and AIDS: where is the epidemic going? Bull. World Health Organ. 74 (2).
Moore, R., Chaisson, R., 1999. Natural history of HIV infection in the era of combination antiretroviral therapy. AIDS 13 (14).
Mveyange, A., Skovsgaard, C., Lesner, T., 2015. Does HIV/AIDS Matter for Economic Growth in Sub-saharan Africa? WIDER Working Paper 2015/086.
Nahmias, A., Weiss, J., Yao, X., Lee, F., Kodsi, R., Schanfield, M., Essex, M., 1986. Evidence for human infection with an HTLV III/LAV-like virus in Central Africa, 1959. The Lancet 327 (8492).
Oster, E., 2012. Routes of infection: exports and HIV prevalence in sub-Saharan Africa. J. Eur. Econ. Assoc. 10 (5).
Palmer, S., Maldarelli, F., Wiegand, A., Bernstein, B., Hanna, G.J., Brun, S.C., ... King, M.S., 2008. Low-level viremia persists for at least 7 years in patients on suppressive antiretroviral therapy. Proc. Natl. Acad. Sci. 105 (10).
Pinkovskiy, M., Sala-i-Martin, X., 2014. Africa is on time. J. Econ. Growth 19.
Popovic, M., Sarngadharan, M., Read, E., Gallo, R., 1984. Detection, isolation, and continuous production of cytopathic retroviruses (HTLV-III) from patients with AIDS and pre-AIDS. Science 224 (4648).
Reich, M.R., Bery, P., 2005. Expanding global access to ARVs: the challenges of prices and patents. In: Mayer, K.H., Pizer, H.F. (Eds.), The AIDS Pandemic: Impact on Science and Society. Academic Press, New York. Chap. 13.
Rodrick, D., 2018. An African growth miracle? J. Afr. Econ. 27 (1).
Schwartländer, B., Grubb, I., Perriëns, J., 2006. The 10-year struggle to provide antiretroviral treatment to people with HIV in the developing world. Lancet 368.
Shelton, J., 2002. Male Circumcision and National HIV Infection Rates. Johns Hopkins Bloomberg School of Public Health: Jim Shelton's Pearls.
Strauss, J., Thomas, D., 1998. Health, nutrition, and economic development. J. Econ. Lit. 36 (2).
Thirumurthy, H., Zivin, J.G., Goldstein, M., 2008. The economic impact of AIDS treatment: labor supply in Western Kenya. J. Hum. Resour. 43 (3).
Thomas, D., Frankenberg, E., 2002. Health, nutrition and prosperity: a microeconomic perspective. Bull. World Health Organ. 80 (2).
U.S. Census Bureau, 2004. Global Population Profile: 2002. U.S. Government Printing Office, Washington, DC.
UNDESA, 2004. Demographic impact of AIDS. In: The Impact of AIDS. United Nations Department of Economic and Social Affairs. Chap. 2.
UNDESA, 2014. World Population Prospects: the 2012 Revision: Methodology of the United Nations Population Estimates and Projections. United Nations Department of Economic and Social Affairs: Population Division.
WHO, 2001. Macroeconomics and Health: Investing in Health for Economic Development. (Report of the Commission on Macroeconomics and Health).
WHO, 2004. Global Burden of Disease: 2004 Update.
WHO, 2015a. Global Health Sector Response to HIV 2000-2015: Focus on Innovations in Africa. Progress report.
WHO, 2015b. World Malaria Report 2015. World Health Organization.
Wooldridge, J.M., 2002. Econometric Analysis of Cross Section and Panel Data. MIT Press.
Young, A., 2005. The gift of the dying: the tragedy of AIDS and the welfare of future African generations. Q. J. Econ. 120 (2).
Young, A., 2012. The African growth miracle. J. Political Econ. 120.
Young, A., 2019. Consistency without Inference: Instrumental Variables in Practical Application. (Working Paper).