Economic Modelling 61 (2017) 376–387

Contents lists available at ScienceDirect

Economic Modelling

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

Decomposing changes in the conditional variance of GDP over time

Shahram Amini[a], Michele Battisti[b,c], Christopher F. Parmeter[d,*]

[a] Department of Economics, Virginia Polytechnic Institute and State University, Blacksburg, VA 24601, United States
[b] Dipartimento di Studi su Politica, Diritto e Societá, Palermo, Italy
[c] CeLeG LUISS, Viale Romania, 32, 00197 Rome, Italy
[d] Department of Economics, University of Miami, Coral Gables, FL 33124-6520, United States

ARTICLE INFO

JEL classification:
C14
G0
O50

Keywords:
Generalized Kernel
Nonparametric
Conditional variation

ABSTRACT

A well established fact in the growth empirics literature is the increasing (unconditional) variation in output per capita across countries. We propose a nonparametric decomposition of the conditional variation of output per capita across countries to capture different channels over which the variation might be increasing. We find that OECD countries have experienced diminishing conditional variation while other regions have experienced increasing conditional variation. Our decomposition suggests that most of these changes in the conditional variance of output are due to unobserved factors not accounted for by the traditional growth determinants. In addition to this we show that these factors played very different roles over time and across regions.

1. Introduction

Cross-country empirical growth studies commonly focus on β-convergence, in part to address such questions as "Do poor countries grow faster than richer ones?" or "How long will it take for a poor country to become rich?" Both of these questions are geared towards economies catching up with one another and highlight how relative income disparities are changing over time. However, it is well known that focusing on a coefficient in a conditional mean regression is limited (Quah, 1993a) and cannot explain concepts such as intra-distributional churning, multimodality, and expansion/contraction of the distribution over time. To more adequately study additional features of the cross-country distribution of output, growth empiricists have deployed a wide array of statistically rich modeling techniques to sharpen focus on how this distribution has changed. Within these studies a common 'distributional moment' that is of interest is the variance (see Pittau et al., 2010), leading to speculation on σ-convergence.

It its most basic form, unconditional σ-convergence is assessed by looking at differences in the variation of the logarithm of cross-country output at two periods in time. As Quah (1996a) notes, while σ-convergence may be more illuminating regarding the behavior of the cross-country distribution of output than its β-convergence counterpart, it is still only a *feature* of the distribution and as such cannot capture entirely what is happening over time to the distribution. For instance, if one were to witness σ-convergence, intra-distribution churning and/or the appearance of multiple modes could occur, either of which would not be captured concomitantly with the observance of σ-convergence.

However, one of the great appeals of studying β-convergence (even with the litany of econometric issues that impact the analysis; Durlauf, 2009), is that conditioning variables, such as quality of institutions, can be used to guide insight into how to promote growth. Consider, for instance, that if a given covariate, again using quality of institutions, has a positive effect in a cross-country growth regression the main intuition is that the speed at which a country approaches its steady state, conditional on institutions, would be higher, so the policy implication is improving institutional quality. A traditional analysis, which places very specific assumptions on the convergence equation provides limited policy insight as little in the way of heterogeneity is accounted for. If one considers parameter heterogeneity[1] then specific impacts of a given covariate, in a given country, can be made.

In this paper we investigate a counterpart of this reasoning, focusing on the conditional variation of output. When attention turns to conditional variation, questions like "If African nations had levels of human capital and population growth as in OECD countries, would we witness a diminution of income dispersion over time?" or, more generally, "Without the observed changes in human or physical capital stocks would we observe less dispersion in cross-country output?" can be addressed. The focus on conditional variation provides straightfor-

* Corresponding author.
E-mail addresses: shahram@vt.edu (S. Amini), mbattisti@unipa.it (M. Battisti), cparmeter@bus.miami.edu (C.F. Parmeter).

[1] We may think to several forms of parameter heterogeneity in growth context as a semi- or nonparametric treatment of the production structure (Liu and Stengos, 1999), mixture modelling (Pittau et al., 2010) or.

http://dx.doi.org/10.1016/j.econmod.2016.10.016
Received 8 March 2016; Received in revised form 19 September 2016; Accepted 31 October 2016
Available online 25 November 2016
0264-9993/

ward intuition on the role of a specific covariate not for promoting growth, but for influencing variation in output across countries. We see how these questions are the variational equivalents of the conditional β-regressions and are undoubtedly more interesting/insightful than their unconditional counterparts, that only focus on simple variance patterns over time. This can improve the complementary role of information of β- and σ-convergence approaches advocated by Sala-i-Martin (1996).

Using recently developed nonparametric kernel smoothing methods, we suggest a decomposition of the conditional variance of output based on covariates and time. This allows splitting overall variation in cross-country output into changes due to covariates and time. The deployment of nonparametric methods allows us to eschew both distributional and parametric functional form assumptions, which could produce misleading results due to model misspecification. To our knowledge, there currently does not exist an empirical study of cross-country variation of output in a nonparametric framework. This makes our combination of methods and application important for the growth empirics field.[2]

With our nonparametric decomposition in tow, we see two main results emerge. First, over the period 1960–2010, OECD countries observed a decreasing conditional variation, split almost equally across time and covariates. Second, all other regions of countries experienced an increasing conditional variation, with differences emerging across the role played by time and covariates. We also present a series of robustness checks over various dimensions of our empirical exercise. Qualitatively, our two main findings remain intact.

The remainder of the paper is organized as follows. Section 2 reviews the related literature, Section 3 outlines our construction of the conditional density and how we will estimate the conditional variance. Section 4 will investigate the change in the conditional variance based upon time and covariates. Section 5 summarizes our findings and offers avenues for further research.

2. Background discussion

Even with the extant shortfalls of studying different 'moments' of the distribution of output, one can still discern important information by studying the behavior of these moments over time. More directly, by focusing on the behavior of these moments in a conditional setting, empirical growth studies can glean information not available in restricted unconditional settings. Here, we use methods similar to Maasoumi et al. (2007) and Henderson et al. (2012) to estimate the conditional density of cross-country per capita output but use the focus of Pittau et al. (2010) to analyze the variance of these conditional densities over time and for different subsets of countries. The work of Maasoumi et al. (2007) focused primarily on the behavior of the conditional distribution/density of growth rates (actual and predicted) over time between OECD and non-OECD countries whereas the work of Pittau et al. (2010) decomposed the density of cross-country output into three groups and then subsequently analyzed the (unconditional) variance of these three groups over time.

Here we blend these two studies together and offer a decomposition of the variance of the conditional density. We depart from Maasoumi et al. (2007) by explicitly focusing on the variance of the distribution while we extend the work of Pittau et al. (2010) by examining conditional variances as opposed to unconditional ones. In addition to this, we decompose overall changes in the variance over time into a covariate component and a time component, logic that is close to Beaudry et al. (2005). The covariate component can be taken as measure of the impact that covariates have on overall variation in output, which cannot be discerned in an unconditional setting.

To obtain a sense for the importance of conditioning, consider the multimodality finding of Quah (1993a). While subsequently illuminating regarding the relative polarization of the distribution over time, given the unconditional framework of the analysis, only speculative evidence could be provided underlying the root for the emerging bimodal shape and increasing variance. The work of Feyrer (2008) and Henderson et al. (2008) both examined the behavior of the corresponding Solow growth determinants, along with cross-country output to see if similar patterns emerged that may provide further evidence to Quah's initial discovery of a bimodal distribution.

However, whether or not σ-convergence is an interesting phenomenon, it is useful to note that absolute β-convergence is a necessary condition for absolute σ-convergence to occur (Quah, 1993a; Furceri, 2005) and with the abundance of studies showing no tendency for absolute β-convergence across countries, it comes as no surprise that there is no supporting evidence for absolute σ-convergence. A stimulating research agenda would be to determine if a parallel necessary condition for conditional β-convergence exists for conditional σ-convergence to materialize, and moreover, if this mode of convergence is supported by the data? To begin to answer these types of questions appropriate concepts of conditional σ-convergence are needed. However, constructing a conditional counterpart has proved elusive since Barro and Sala-i-Martin (1991) formally defined this concept.[3] As Durlauf et al. (2005, p. 53) note "An economically interesting formulation of conditional σ-convergence would be a useful contribution."

Evidence in favor of absolute σ-divergence is provided in Table 1. Taking a balanced sample of 70 countries for whom we have data in Penn World Table 8.1 (Feenstra et al., 2015) and Barro and Lee (2013) from 1960 to 2010, we can look for evidence of absolute σ-convergence across decades and various measures of output per capita. Table 1 shows the variation in the logarithm of output for RGDPE, RGDPO and RGDPNA as well as for the standard growth accounting variables as investment rates (INV), population growth (POP) and years of schooling (EDUC). We immediately notice that all three of the common measures of per capita output (RGDPE, RGDPO and RGDPNA) display increasing variation over each decade, aside from a modest reduction over the last decade (where we have the financial crisis). There is not even the appearance of the variance stabilizing over time for any measure of output from 1960 to 2000. While this yields conclusions regarding the lack of absolute σ-convergence, identifying the underlying causes for this increase remain elusive in an unconditional setting.[4]

It is interesting to note that the variation in investment rates seems to be declining over time, in line with the research of Caselli and Feyrer (2006) while levels of education do not display a clear pattern of absolute σ-convergence. The apparent σ-divergence is in accord with the development accounting findings of Stamatakis and Petrakis (2005) and Henderson et al. (2008) and the underlying reasons for this divergence represents an interesting research agenda not explored here.

Taking into account the reduction in variance of population, we see how in a traditional augmented Solow model a lá Mankiw et al. (1992) the diverging pattern of output is not mimicked by the diverging pattern of the determinant variables. Notwithstanding that, we could still have that the returns to these variables could be changing over time, and as such, may explain this increasing variation. In order to provide more intuition consider for instance the situation where a country that is in the extreme right of the distribution (a high income country) in the previous decade, grows much more than the average, becoming even richer. Now, *ceteris paribus*, we could have for instance

[2] See Boccanfuso et al. (2013) and Li et al. (2016) for additional applications of nonparametric kernel smoothing methods in alternative growth and income contexts.

[3] See Phillips and Sul (2007), who develop a panel time varying idiosyncratic convergence test.

[4] We caution that the appearance of σ-divergence can be attributed to the measure of dispersion used and not the actual phenomena of output diverging over time (Dalgaard and Vastrup, 2001).

Table 1
Variance of logarithm of output and growth accounting variables over time.

Year	RGDPE	RGDPO	RGDPNA	INV	POP*1000	EDUC
1960	0.813	0.825	1.048	0127	0.009	6.08
1970	0.920	0.911	1.151	0.156	0.010	6.94
1980	1.137	1.113	1.260	0.118	0.011	7.68
1990	1.290	1.289	1.402	0.112	0.007	7.01
2000	1.510	1.539	1.544	0.063	0.003	6.52
2010	1.335	1.412	1.413	0.064	0.005	6.70

that, from the implications of the augmented Solow model, this country invested much more than other countries in physical and human capital. On the other it could be the case that this country had average investment rates but the returns to these investments are much higher than in other countries. This second possibility implies a change in the coefficients of the Solow model over time. Our decomposition in Section 3 may shed light on which of these two effects is taking hold.

An interesting aspect of considering variation of conditional moments over time is how they correspond to the common notions of conditional convergence prevailing in the growth empirics literature currently. In the β-convergence literature the conditioning was designed to account for differences in the steady state levels of output across countries. However, when one migrates from a first moment setting to the investigation of higher moments, such as the conditional variance, it is not clear what conditional convergence of a ‘moment’ implies. While researchers have offered regression based tests of both unconditional and conditional σ-convergence (Cannon and Duck, 2000; Egger and Pfaffermayr, 2007; Huber and Pfaffermayr, 2010), as noted by Bliss (2000), the fact that the variance of the left hand side variable is changing over time renders asymptotic analysis of the test difficult to understand. Broadly, this criticisms passes to any type of regression (conditional mean) setting which analyzes convergence of higher order moments (conditional or not). More importantly, Durlauf et al. (2005) have raised several issues with constructing suitable notions for conditional σ-convergence. Rather, to crystallize our focus we simply point out that instead of focusing on convergence or divergence, we instead seek to understand how conditional variation changes over time, irrespective of whether this may be termed to conditionally diverge or converge.

It is useful to discuss what a conditional measure of variance should look like. As noted by Quah (1993a,b), regressions that are done over time can fail to capture the underlying distribution dynamics, resulting in researchers misinterpreting their results, or, more starkly, for the results to be meaningless entirely. In fact, the notion of unconditional σ-convergence can be seen as the variance from the unconditional density of the logarithm of output at two points in time. Similarly, one may then suggest that the appropriate concept of conditional σ-convergence be defined by looking at the variance of the conditional density at two points in time. The key distinction here though, however, is that now that conditioning variables are present one needs to define exactly how the two variances are compared.

Taking this notion further, in a β-convergence study (either absolute or conditional) a regression of growth rates over time is run on some initial value of output or income and ‘conditioning’ variables. A negatively signed coefficient estimate on the initial income variable is taken as evidence for conditional β-convergence. The definition of absolute σ-convergence however, is a relationship between the variance of the logarithm of output at *two* points in time. Thus, one should not be looking to run a regression where a negative coefficient on some level of income at a point in time signifies σ-convergence.[5] While it is appealing to focus on aspects of ‘convergence’ in growth empirics, it also draws interest away from the appropriate settings to study the behavior of conditional moments of the distribution of cross-country output. Without formal links to the variables which influence cross-country output, it is hard to attach any meaning to why we see increased dispersion of incomes over time, let alone lay claim that these increases are attributed to a specific variable.

If we define unconditional σ-convergence as

$$\sigma^2_{Y,t} > \sigma^2_{Y,T}, \quad \text{for } T > t, \tag{1}$$

then an appropriate way to look at conditional variation over time would be

$$\sigma^2_{Y_i|X_i,t} \gtreqless \sigma^2_{Y_i|X_i,T}, \quad \text{for } T > t, \tag{2}$$

where the conditioning takes place at the same level of the covariates across time.[6] Thus, what we are looking at *within* the conditional density of the logarithm of output is, for a given level of the covariates over time, has the dispersion of the log of output diminished. Thus, it is possible that the conditional variation is increasing within certain parts of the conditional distribution but not at others.

3. Empirical methods

While many growth studies begin with a regression framework, here we focus on the insights of Quah (1993a,b) and use a distributional approach. Previous work that has studied various features of the growth process through a density (distribution) framework are Quah (1996a, 1996b, 1996c), Bianchi (1997), Jones (1997), Pritchett (1997), Paap and van Dijk (1998), Desdoights (1999), Johnson (2000), Gisbert (2003), Anderson (2004), Azariadis and Starchurski (2004), Canova (2004), Fiaschi and Lavezzi (2004), Johnson (2005), Pittau (2005), Pittau and Zelli (2006), Maasoumi et al. (2007), Henderson et al. (2008), and Pittau et al. (2010). Some of the papers have focused on unconditional density features while others have focused on conditional densities. Regardless, this list shows the growing body of literature that attempts to learn about the growth process via distributional methods. We introduce a generalized product kernel approach to construct the conditional density over time that has not previously been deployed in the growth empirics literature.[7]

3.1. Estimating a conditional density

Let $f(\cdot)$ denote the joint density of (X,Y) and $g(\cdot)$ denote the marginal density of X. Here Y is the logarithm of output and X are the set of conditioning variables that will form the basis of the conditioning set to study conditional σ-convergence. Define the generalized product kernel as

$$K_\gamma(x, X_i) = \prod_{s=1}^{q} h_s^{-1} l^c\left(\frac{x_s^c - X_{si}^c}{h_s}\right) \prod_{s=1}^{r} l^u(x_s^u, X_{si}^u, \lambda_s^u) \prod_{s=1}^{p} l^o(x_s^o, X_{si}^o, \lambda_s^o). \tag{3}$$

Here, X is partitioned into three components, continuous (X^c), ordered discrete (X^o), and unordered discrete (X^u). The distinction between data types is important in nonparametric applications as it lessens the curse of dimensionality. We have used q, r and p to denote the number of continuous, unordered, and ordered conditioning variables, respectively. $h_s^{-1} l^c(\cdot)$ is the standard normal kernel function with window width h_s, associated with the s^{th} component of x^c that is commonly used in unconditional density studies (Bianchi, 1997; Jones, 1997;

[5] In addition to this, if the shape of the error distribution at two points of time is different, for example normally distributed at time t but a bimodal density at time T, regression-based tests for conditional σ-convergence will not be informative and most likely will use incorrect asymptotic distributions as the basis for constructing p-values.

[6] This is inline with conditional β-convergence studies that incorporate interactions and nonlinearities that involve the initial value of income (Durlauf and Johnson, 1995; Stengos and Li, 1998; Durlauf et al., 2001; Kourtellos, 2003).

[7] Two exceptions are Maasoumi et al. (2007) and Henderson et al. (2012), but both papers use a generalized product kernel in a regression setting.

Pritchett, 1997). l^u is a variation of Aitchison and Aitken's (1976) kernel function which equals one if $x_{si}^u = x_{sj}^u$ and $\lambda_s{}^u$ otherwise, and l^o is the Wang and Van Ryzin (1981) kernel function which equals one if $x_{si}^o = x_{sj}^o$ and $(\lambda_s{}^o)^{|x_{si}^o - x_{sj}^o|}$ otherwise. h^c, λ^u and λ^o are the bandwidths associated with each kernel. See Li and Racine (2003, 2007) for further details.

We denote our estimators of the joint and marginal densities, $\hat{f}(x, y)$ and $\hat{g}(x)$ as

$$\hat{f}(x, y) = n^{-1} \sum_{i=1}^{n} K_\gamma(x, X_i) k_{h_y}(y, Y_i), \tag{4}$$

$$\hat{g}(x) = n^{-1} \sum_{i=1}^{n} K_\gamma(x, X_i), \tag{5}$$

where h_y is the smoothing parameter associated with Y and $k_{h_y}(\cdot)$ is the simple univariate, continuous kernel which smooths our outcome variable ($k_{h_y}(u) = h_y^{-1}\ell^c(u)$). Noting that the conditional density of Y is defined as $m(y|x) = f(x, y)/g(x)$, we estimate the conditional density as

$$\hat{m}(y|x) = \hat{f}(x, y)/\hat{g}(x). \tag{6}$$

Once we have obtained an estimate of the conditional density we can construct the conditional variance of the density for any level of the covariates ($x = \bar{x}$, say). Thus, we can estimate the conditional density at two points in time and hold the covariates fixed at their initial time period levels to determine if the conditional variance of y has changed. All that remains is to discuss how the bandwidths, which are an integral part of the conditional density estimator, are arrived at and how we choose to estimate the variance at any particular level of covariates.

3.2. Bandwidth selection

Bandwidth selection is commonly perceived as the most important aspect of any kernel based nonparametric modelling (Henderson and Parmeter, 2015). While a variety of methods exists concerning construction of optimal bandwidths, we advocate a data-driven approach which has recently been hailed for its desirable asymptotic properties (Hall et al., 2004) and ability to detect the inclusion of irrelevant variables. The approach is termed Least Squares Cross-Validation (LSCV). The bandwidths are selected by minimizing a sample analog of integrated squared error:

$$ISE = \int \{\hat{m}(y|x) - m(y|x)\}^2 g(x) dx dy. \tag{7}$$

ISE can be written as the sum of three components ($ISE_1 + ISE_2 + ISE_3$), only two of which (ISE_1 and ISE_2) depend on the unknown bandwidths.

For a given sample, the analog estimators of ISE_1 and ISE_2 are

$$\widehat{ISE}_1 = n^{-1} \sum_{i=1}^{n} \hat{G}_{-i}(X_i)/\hat{g}_{-i}(X_i)^2, \quad \widehat{ISE}_2 = n^{-1} \sum_{i=1}^{n} \hat{m}_{-i}(X_i, Y_i)/\hat{g}_{-i}(X_i)^2,$$

where

$$\hat{G}(X_i) = n^{-2} \sum_{i_1=1}^{n} \sum_{i_2=1}^{n} K_\gamma(X_i, X_{i1}) K_\gamma(X_i, X_{i2}) \int k_{h_y}(y, Y_{i1}) k_{h_y}(y, Y_{i2}) dy, \tag{8}$$

and a subscript $-i$ denotes a leave-one-out estimator, for example,

$$\hat{g}_{-i}(X_i) = (n-1)^{-1} \sum_{j=1, j\neq i}^{n} K_\gamma(X_i, X_j).$$

The integral which appears in (8) can be simplified if one uses a Gaussian kernel to smooth over y, i.e., $\ell^c(u) = \phi(u)$. In this case $\int k_{h_y}(y, Y_{i1}) k_{h_y}(y, Y_{i2}) dy$ produces the convolution kernel, $\bar{k}_{h_y}(\cdot)$ which is itself a Gaussian kernel (albeit with different variance).[8]

The LSCV objective function is given as

$$LSCV(h_y, \gamma) = \widehat{ISE}_1(h_y, \gamma) - 2\widehat{ISE}_2(h_y, \gamma), \tag{9}$$

where $\gamma = \{h_1, \ldots, h_q, \lambda_1^o, \ldots, \lambda_p^o, \lambda_1^u, \ldots, \lambda_r{}^u\}$ is the vector of bandwidths associated with the covariates. Once the bandwidths have been determined the conditional density can be estimated at any given level of the covariates. Given the importance of the bandwidths to the empirical performance of the estimator, and the known criticisms that LSCV can produce bandwidths which are 'too small' relative to the optimal bandwidths (Henderson and Parmeter, 2015), alternative bandwidth selection mechanisms are important to deploy as robustness checks. We do this in Section 4.4 using maximum likelihood cross-validation (MLCV).

3.3. Estimating the conditional variance

Constructing moments from a conditional density estimator is common. For example, when one estimates a regression model, this is the mean from the conditional density of the independent and dependent variables, i.e.,

$$\hat{E}[Y|X = x] = \int y\hat{m}(y|X = x) dy = \sum_{i=1}^{n} w_i(x) Y_i,$$

where $w_i(x) = K_\gamma(x, X_i)/\sum_{i=1}^{n} K_\gamma(x, X_i)$, which is the typical local constant kernel regression estimator of Nadaraya (1965) and Watson (1964). In our setting we are interested in the conditional variance. A little bit of algebra reveals that the estimator of the conditional variance (when Y is continuous), $\widehat{Var}[Y|X = x]$, is[9]

$$\begin{aligned} \hat{v}(x) &= \int [y - \hat{E}[Y|X = x]]^2 \hat{m}(y|X = x) dy \\ &= h_y \sigma_K^2 + \sum_{i=1}^{n} w_i(x)[Y_i - \hat{E}[Y|X = x]]^2. \end{aligned} \tag{10}$$

Here, σ_K^2 is the variance of the kernel used to smooth Y. In our empirical example we use a second order Gaussian kernel, $K(x) = h^{-1}(x/h)$, which has unit variance. We eschew integration by noting that our conditional variance can be calculated in three steps. First, conduct least-squares cross-validation[10] to determine the optimal bandwidths for the conditional density. Second, estimate the conditional mean of Y using the bandwidths corresponding to the Xs. Note that h_y is not used in this second stage. Third, using the estimates of the conditional mean obtained in the second stage, regress the squared residuals on X using the same set of bandwidths as those used for the conditional mean.

3.4. Decomposing changes in variation over time

To investigate changes in conditional variation over time we can look at the difference between our conditional variance estimates at different points in time. That is,

$$\begin{aligned} \hat{v}_t(x_t) - \hat{v}_T(x_T) &= \underbrace{\hat{v}_t(x_t) - \hat{v}_T(x_t)}_{\text{Time Effect}} + \underbrace{\hat{v}_T(x_t) - \hat{v}_T(x_T)}_{\text{Covariate Effect}} \\ &= \underbrace{\hat{v}_t(x_t) - \hat{v}_t(x_T)}_{\text{Covariate Effect}} + \underbrace{\hat{v}_t(x_T) - \hat{v}_T(x_T)}_{\text{Time Effect}}. \end{aligned} \tag{11}$$

A nice feature of this strategy is that it allows the decomposition of the change in the variance into factors attributable to movements within the density against factors attributable across time. For example, one

[8] Precisely, the integral takes the form $\bar{k}_{h_y}\left(\frac{Y_{i1} - Y_{i2}}{h_y}\right)$, where $\bar{k}_h(u) = h^{-1}e^{-u^2/4}/\sqrt{4\pi} = h^{-1}\phi(u/2)/\sqrt{2}$.

[9] See Hyndman et al. (1996) for more on this derivation in the simple, univariate kernel setting.

[10] Alternative data-driven bandwidth selection mechanisms could be deployed as well, for instance, maximum likelihood cross-validation. We present results in Section 4.4 using alternative bandwidths and kernels to assess robustness to the smoothing parameters.

can check if an increase/decrease in human capital led to a widening of the conditional variance over time.

Our primary concern with the decomposition is with the covariate effect. It tells us the contribution to the overall change in the conditional variance based on how the covariates have changed over time, holding the shape of the distribution fixed from one time period to another. Thus, in order to say that the variance went down, the change in the covariates must have moved the country in such a way that the conditional variance decreased, holding time fixed. The time effect is how the conditional variance function itself has changed over time, holding the level of the covariates fixed. It is linked to features of the output process not captured in our covariate set. An obvious candidate is TFP. In addition to this, if time dependent heteroscedasticity were present the time effect could trigger either the appearance of an increasing or decreasing conditional variance if it were large enough in magnitude. This point is further discussed in Durlauf et al. (2005).

By focusing on specific levels of covariates we can determine if their shift has been promoting a narrowing of the variance of the conditional density, since we can hold time fixed. If the covariate effect is zero then we can conclude that changes in the conditional variance are due to underlying time varying factors. This could be as simple as time varying heteroscedasticity or something more complicated such as comparable productivity stocks producing asymmetric changes in the cross-country production function. When the covariate effect is not zero this provides insight into how the conditioning set influences dispersion of cross-country output in more directly interpretable ways.

Due to the fully nonparametric setup of our problem, the choice of covariate levels in both time periods will dictate the appearance of changes of the conditional variance. However, it is not obvious which set(s) of points to look over to construct our variance decomposition. We elect to look over the actual data points across time. While other strategies exist this feels the most natural to us.

An alternative viewpoint for (11) is that of a distributional generalization of the decomposition approach of Beaudry et al. (2005) who build counterfactual income distributions by growth rate accumulation according to Barro type regressions. That approach implies a homogeneity assumption of the production function across countries, given by the fact that they apply the same coefficients to each part of the income distribution (so that a shift in the mean is representative for all distributional shifts). In our case, the decomposition in (11) eschews assumptions on the production structure across countries. This is also consistent with the findings of Massoumi et al. (2007) and Henderson et al. (2012) who show that traditional Barro regressions are statistically misspecified.

4. Empirics

4.1. Data

The data used for this study come from Penn World Table 8.1 (Feenstra et al., 2015). We use RGDPE as our primary measure of output per capita, and we use population growth rates and physical capital formation from the corresponding variables in the PWT. Our measure of human capital stocks are the average years of schooling for the population from Barro and Lee (2013). RGDPE is real GDP measured from the expenditure-side, which is constructed using prices for final goods that are constant across countries and over time. An alternative measure, RGDPO, is output-side real GDP, using prices for final goods exports and imports that are, similar to RGDPE, constant across countries and over time could also be used (we present a robustness check using this measure in Section 4.4). Both of these measures are useful for making comparisons across countries (Feenstra et al., 2015; Table 1).

Before determining the change of the conditional variance over time we look at the behavior of each of the variables in order to determine if any interesting patterns emerge within geographical groups or over time. The reported changes are relative to each variable so that they may be compared. Table 2 shows the results for the sample of 70 countries for the period 1960–2010. We immediately see that there is substantial heterogeneity across region and time.[11] There is evidence of global σ-divergence for GDP, while only the OECD shows sustained absolute σ-convergence amongst regions.[12]

Table 2
Relative changes in variances of logarithm of output and Solow growth determinants over time by geographical regions 1960–2010. LAC stands for Latin American and Caribbean countries while MENA is Middle East and North African Countries.

	1970 vs 1960	1980 vs 1970	1990 vs 1980	2000 vs 1990	2010 vs 2000
RGDPE					
All	0.132	0.236	0.134	0.170	−0.116
OECD	−0.395	−0.069	−0.149	0.080	−0.085
LAC	0.049	0.177	−0.161	−0.011	0.206
Asian	0.223	1.888	0.481	0.319	−0.080
African	0.108	0.323	0.402	0.424	−0.149
MENA	0.053	0.143	0.130	−0.108	−0.258
POPULATION GROWTH					
All	0.220	−0.240	−0.057	−0.438	0.024
OECD	−0.260	−0.272	−0.324	−0.130	0.443
LAC	1.594	−0.619	0.038	−0.164	−0.108
Asian	−0.855	1.139	0.262	−0.061	−0.373
African	1.017	1.101	0.256	0.109	0.655
MENA	−0.324	−0.487	0.601	−0.788	1.303
YEARS OF EDUCATION					
All	0.140	0.107	−0.087	−0.070	0.027
OECD	0.084	−0.007	−0.268	−0.292	0.071
LAC	0.133	0.060	0.176	−0.028	−0.031
Asian	1.210	0.461	0.519	0.220	−0.040
African	0.241	0.099	0.430	−0.006	0.520
MENA	0.179	0.239	−0.003	−0.131	−0.071
INVESTMENT RATES					
All	0.165	0.110	−0.397	−0.582	0.731
OECD	0.455	−0.229	−0.293	−0.536	1.071
LAC	−0.061	0.457	−0.782	−0.178	0.448
Asian	−0.276	−0.512	0.380	−0.641	1.926
African	0.266	−0.524	−0.329	−0.237	0.284
MENA	0.737	0.430	−0.637	−0.825	3.054

It is also worth noting that the traditional Solow growth determinants behave very different across time and regions, even though there is a general reduction in the variation of these three variables over the later part of our study. This finding aligns well with the work of Battisti and Parmeter (2013), who, in a multivariate clustering framework found asymmetries among the relationship between GDP and physical and human capital stocks across richer and poorer countries across time. However, it is difficult to ascribe the behavior of any one of the Solow growth determinants' behavior to the general pattern of output variation over time since all of these variances are calculated unconditionally. To more aptly characterize the impact that one of the Solow determinants has on output variation, we now turn to our main variance decomposition.

4.2. Main variance decomposition

Using the methods described in Section 3, we estimate the

[11] A decreasing pattern inside each group could be still consistent with an increasing variance if the distribution becomes more multimodal, so if there is a divergent(across)-convergence(within) phenomenon.

[12] Our regions are defined according to a geographical criterion, outside of the OCED, but are very similar to clusters determined according to the location inside the distribution of output, given the income differences among continents as in Pittau et al. (2010) or in Battisti and Parmeter (2013).

Table 3
Overall, time and covariate effects for conditional variance for OECD countries.

Country	1960–2010			1960–1980			1980–2010		
	Overall	Covariate	Time	Overall	Covariate	Time	Overall	Covariate	Time
Australia	0.035	0.018	0.017	0.044	0.015	0.029	−0.009	0.004	−0.013
Austria	0.059	0.018	0.041	0.018	0.003	0.015	0.041	0.032	0.009
Belgium	0.065	0.03	0.036	0.018	0.006	0.012	0.047	0.045	0.002
Canada	0.051	0.026	0.025	0.035	0.01	0.025	0.016	0.021	−0.005
Switzerland	0.052	0.024	0.028	0.045	0.009	0.035	0.007	0.005	0.003
Germany	0.072	0.036	0.036	0.01	−0.003	0.013	0.062	0.054	0.008
Denmark	0.073	0.039	0.034	0.013	0.003	0.009	0.061	0.06	0
Spain	0.09	0.06	0.03	−0.004	0.039	−0.043	0.094	0.09	0.004
Finland	0.088	0.055	0.033	0.007	0.022	−0.015	0.082	0.081	0.001
France	0.083	0.047	0.035	−0.007	0.022	−0.028	0.089	0.087	0.002
United Kingdom	0.07	0.04	0.03	0.031	0.004	0.027	0.039	0.04	0
Greece	0.065	0.024	0.041	−0.056	0.002	−0.059	0.121	0.116	0.005
Ireland	0.064	0.042	0.022	−0.007	0.015	−0.022	0.071	0.079	−0.008
Iceland	0.045	0.021	0.023	−0.036	−0.007	−0.029	0.081	0.086	−0.006
Italy	0.063	0.023	0.04	−0.019	0.006	−0.025	0.082	0.07	0.012
Japan	0.072	0.033	0.039	0.015	0.013	0.002	0.057	0.053	0.004
Korea Republic	0.289	0.257	0.032	0.153	0.228	−0.075	0.136	0.136	0
Luxembourg	0.064	0.051	0.012	0.043	0.022	0.021	0.021	0.036	−0.015
Netherlands	0.065	0.03	0.035	0.036	0.012	0.023	0.029	0.028	0.002
Norway	0.058	0.032	0.026	0.026	0.009	0.018	0.032	0.037	−0.006
New Zealand	0.036	0.009	0.027	0.043	0.009	0.035	−0.008	−0.003	−0.005
Portugal	−0.03	−0.067	0.038	−0.033	−0.099	0.066	0.003	−0.037	0.041
Sweden	0.065	0.036	0.029	0.051	0.018	0.033	0.014	0.017	−0.002
Turkey	−0.013	−0.047	0.034	−0.008	−0.217	0.208	−0.005	−0.097	0.092
United States	0.052	0.027	0.025	0.052	0.02	0.031	0	0.001	0
Average	0.065	0.035	0.031	0.019	0.006	0.012	0.048	0.042	0.006

conditional distribution of output based on 10 years averages of population growth, human capital stocks and investment rates. We do this for the decades 1960, 1970, 1980, 1990, 2000 and 2010 and then construct our measure of conditional variation, and the subsequent decomposition, accordingly. Since we are subtracting year 2010 numbers from year 1960 numbers, when $\hat{v}_{1960}(x_{1960}) - \hat{v}_{2000}(x_{2000})$ is positive, this implies a reduction in conditional variation. A negative value thus signifies an increase in the conditional variance.

For ease of discussion we first discuss changes in the conditional variance for each country, grouping them regionally. To provide further insight we also split the full period into two subperiods: 1960–1980 and 1980–2010, in order to see what happens before and after the oil shocks. We begin with our regional analysis. The OECD country results are presented in Table 3. We notice several distinct features. First, almost all of these countries have experienced a decrease in the conditional variance (aside from Portugal and Turkey). Second, every country has an estimated positive time effect for the variance decomposition. This time effect captures, amongst other things, time dependent heteroscedasticity and captures almost half of the change in conditional variation that has occurred over time for OECD countries. Third, the covariate effect is larger in the 1980–2010 period suggesting convergence in the accumulation of productive factors after 1980. Lastly, the great majority of the decrease in conditional variation occurs in the last 30 years, after the productivity slowdown.

All of these results suggest that the unconditional decrease in variation across OECD countries seems to have occurred recently (0.048 vs. 0.019), and is due primarily to changes in the structural determinants of growth. We note that the sub-period total effects do not have to equal the full period total effect because of rounding. A broad implication from Table 3 is that investing in capital and human capital can lessen fluctuations in output over time relative to similar countries with lower levels of capital investment.

Given that recent studies have highlighted the difference between African countries and the rest of the world, we isolate the Sub-Saharan countries within our dataset in Table 4. The results are almost identical to those for the OECD countries except in one facet, almost all African nations experience increases in conditional variation over time. In this case the covariate effects are less negative than the OECD countries, while for the great majority of the countries (8 out of 10), the time effects are negative. On average, the measure of divergence is quite large. These results are suggestive that even if Africa had factor levels in 2010 starting in 1960, this would not have been enough of an impetus to witness a reduction in the conditional variation as the time effects over the 1960 to 2010 period appear to dominate the overall conditional variance.

However, if we look at our results for the sub-period analysis we see that the increasing variation in output for sub-Saharan countries as dramatically shifted, with a strong time effect existing in over 1960–1980 and a large covariate effect from 1980–2010. This could be evidence that global productivity shocks starting in 1960 did not have a large enough effect to decrease the variability in output for African countries (perhaps given low levels of factors of accumulation) and then when these countries were in a position to benefit from these shocks, they no longer existed. Regardless, we see that our focus on the sub-periods highlights that the overall effects for our analysis need not be constant throughout time.

Tables 5–7 present the results for the Latin American countries (LAC), the Asian countries and the group of countries comprising the Middle East and North Africa (MENA). We see that each area has experienced a different overall situation, the majority of these countries have witnessed an increase in the conditional variance. The time effects is negative for LAC countries over the 1960–2010 period, being larger (on average) in the period 1980–2010. Perhaps what is most interesting from the sub-period analysis for the LAC is that while in both sub-periods there was a decrease in the conditional variance, with a similar role of time and covariate effects, even if proportionally bigger after the oil shocks, with covariates led to the reduction in variance and time effect playing the opposite role.

The experience of the Asian countries, as a bloc, is different than the LAC. While there was also conditional variation divergence for the while 1960–2010 period as well as the 1960–1980 sub-period, there was a slight conditional variation convergence in the second sub-

Table 4
Overall, time and covariate effects for conditional variance for sub-Saharan African countries.

Country	1960–2010			1960–1980			1980–2010		
	Overall	Covariate	Time	Overall	Covariate	Time	Overall	Covariate	Time
Benin	−0.356	−0.034	−0.322	0.029	−0.156	0.185	−0.386	−0.086	−0.3
Congo Democratic Rep.	−0.334	−0.023	−0.311	0.021	−0.034	0.055	−0.355	−0.089	−0.266
Ghana	0.133	0.126	0.007	0.15	0.148	0.002	−0.017	−0.121	0.104
Kenya	−0.046	−0.041	−0.005	0.008	−0.052	0.06	−0.054	−0.133	0.08
Mauritius	−0.113	−0.082	−0.03	−0.089	−0.037	−0.053	−0.023	−0.005	−0.018
Malawi	−0.009	0.123	−0.132	0.145	0.067	0.077	−0.154	−0.11	−0.044
Uganda	−0.037	0.002	−0.039	0.014	−0.017	0.031	−0.051	−0.109	0.057
South Africa	−0.081	−0.052	−0.03	−0.093	−0.009	−0.084	0.011	0.057	−0.046
Zambia	−0.001	−0.011	0.011	−0.001	0.017	−0.018	0	−0.118	0.118
Zimbabwe	−0.093	−0.063	−0.029	−0.019	−0.032	0.013	−0.074	−0.047	−0.027
Average	−0.094	−0.006	−0.088	0.017	−0.011	0.027	−0.123	−0.076	−0.034

period. The change, within this sub-period, is entirely due to the time effect that reverted its impact on the overall variation. The Asian bloc also saw a balance between covariate effects in the sub-periods. What is noteworthy by comparing LAC and the Asian bloc is that for the 1960–2010 period, an additional driver of conditional variation expansion was through covariates, while in LAC they acted to reduce dispersion. This might not be surprising given the levels of capital deepening and population expansion that took place over these five decades.

The MENA bloc appears to have the most stability (on average) of the five blocs. Looking at single countries, very diversified experiences arise given the high heterogeneity of countries within this group. For instance, Israel is a developed country, but only entered the OECD in 2010; similar concerns exist for Cyprus and Malta as both countries are developed but are not members of the OECD. The result of almost no change in the conditional variance over the 1960–2010 period is an average of this country changes as well as the two sub-periods that show opposite patterns: lower dispersion either in covariates or in time effects in the first period and the opposite in the second.

Of the four blocs of lesser developed countries, Asia has the highest average increase in conditional variation while MENA has the smallest increase, on average. The LAC and Africa blocs are roughly similar in terms of average measures though the sub-period results suggest different patterns for the overall global effect. Thus, while both blocs are exhibiting similar behavior, the underlying sources for this increase in output variation differs. Finally we may note as in all these four blocs the impact of time effect in absolute terms (not looking to the sign) is bigger than the covariate one, suggesting that most of these changes in the conditional variance of output is due to unobserved factors not accounted for by the traditional growth determinants.

One will notice that if we sum the overall effects for the sub-periods 1960–1980 and 1980–2010 we obtain the same overall effect as the global period 1960–2010 (out of small difference of 0.01 for Africa). However, this is not the case with either of our two distinct effects, time and covariate. See for instance the results for African countries in Table 4. Notice that the global covariate effect is −0.006, but if we are to add together the covariate effects for the two sub-periods we have the covariate effect as −0.087. Moreover, the change in the covariate effect is dramatic as we switch from one sub-period to another. To obtain a greater sense for what is going on with our decomposition we further decompose each of our effects across the sub-periods to see what more can be learned from what we see in the tables. If we look the framework of the decomposition we see that the overall effect is (again) given by:

$$\hat{v}_{60}(x_{60}) - \hat{v}_{10}(x_{10}) \tag{12}$$

and this impact is preserved within the subsamples:

$$\hat{v}_{60}(x_{60}) - \hat{v}_{10}(x_{10}) = \hat{v}_{60}(x_{60}) - \hat{v}_{80}(x_{80}) + \hat{v}_{80}(x_{80}) - \hat{v}_{10}(x_{10}). \tag{13}$$

Table 5
Overall, time and covariate effects for conditional variance for Latin American countries.

Country	1960–2010			1960–1980			1980–2010		
	Overall	Covariate	Time	Overall	Covariate	Time	Overall	Covariate	Time
Argentina	−0.092	0.069	−0.161	−0.004	−0.005	0.001	−0.088	0.017	−0.105
Bolivia	−0.147	0.05	−0.198	−0.064	0.049	−0.113	−0.083	0.099	−0.182
Brazil	−0.155	0.06	−0.215	−0.112	0.068	−0.18	−0.043	0.161	−0.204
Chile	−0.107	0.046	−0.153	−0.007	−0.029	0.022	−0.1	−0.004	−0.096
Colombia	−0.126	0.035	−0.161	−0.115	0.018	−0.133	−0.011	0.09	−0.101
Costa Rica	−0.189	−0.016	−0.174	−0.062	−0.031	−0.031	−0.127	0.016	−0.144
Dominican Republic	−0.149	0.061	−0.21	−0.076	0.072	−0.148	−0.073	0.131	−0.204
Ecuador	−0.172	−0.024	−0.148	−0.061	−0.033	−0.028	−0.111	0.01	−0.12
Guatemala	0.132	0.117	0.015	0.006	0.027	−0.021	0.126	−0.026	0.152
Honduras	0.19	0.303	−0.113	0.178	0.321	−0.144	0.012	0.117	−0.105
Jamaica	−0.175	0.019	−0.194	−0.062	−0.047	−0.014	−0.113	0.06	−0.173
Mexico	−0.041	0.134	−0.175	−0.038	0.134	−0.172	−0.003	0.124	−0.126
Panama	−0.119	0.03	−0.149	−0.029	−0.032	0.003	−0.09	−0.015	−0.075
Peru	−0.12	0.028	−0.148	−0.042	−0.009	−0.033	−0.077	−0.011	−0.067
Paraguay	−0.178	−0.02	−0.159	−0.097	−0.009	−0.088	−0.081	0.069	−0.151
El Salvador	0.042	0.323	−0.281	0.016	0.323	−0.308	0.026	0.313	−0.287
Trinidad and Tobago	−0.135	0.07	−0.206	−0.202	−0.014	−0.188	0.067	0.282	−0.215
Uruguay	−0.213	0.019	−0.232	−0.02	−0.034	0.013	−0.192	0.019	−0.212
Venezuela	0.203	0.34	−0.137	0.217	0.335	−0.118	−0.014	0.053	−0.067
Average	−0.082	0.087	−0.168	−0.030	0.058	−0.088	−0.051	0.079	−0.131

Table 6
Overall, time and covariate effects for conditional variance for Asian countries.

Country	1960–2010			1960–1980			1980–2010		
	Overall	Covariate	Time	Overall	Covariate	Time	Overall	Covariate	Time
Bangladesh	−0.087	−0.071	−0.016	−0.019	−0.155	0.136	−0.069	−0.302	0.234
China	−0.13	−0.104	−0.026	−0.32	−0.041	−0.279	0.19	−0.155	0.345
India	−0.077	−0.089	0.012	−0.022	−0.167	0.145	−0.055	−0.362	0.306
Sri Lanka	−0.271	−0.057	−0.214	−0.398	−0.079	−0.319	0.127	0.29	−0.163
Malaysia	−0.168	0.048	−0.216	−0.157	0.086	−0.244	−0.011	0.147	−0.158
Pakistan	−0.082	−0.041	−0.041	−0.021	−0.19	0.169	−0.062	−0.253	0.192
Philippines	−0.237	−0.049	−0.188	−0.292	−0.036	−0.256	0.055	0.044	0.011
Thailand	−0.057	0.058	−0.115	0.08	0.086	−0.006	−0.137	−0.267	0.13
Taiwan	−0.273	−0.052	−0.22	−0.316	−0.058	−0.258	0.043	0.284	−0.241
Average	−0.154	−0.040	−0.114	−0.163	−0.062	−0.101	0.009	−0.064	0.073

Table 7
Overall, Time and Covariate Effects for conditional variance for MENA countries.

Country	1960–2010			1960–1980			1980–2010		
	Overall	Covariate	Time	Overall	Covariate	Time	Overall	Covariate	Time
Cyprus	−0.059	0.005	−0.064	−0.013	0	−0.013	−0.045	0.053	−0.098
Egypt	0.167	0.241	−0.074	0.266	0.096	0.17	−0.098	−0.036	−0.063
Iran	0.173	0.253	−0.08	0.295	0.176	0.119	−0.123	−0.066	−0.057
Israel	−0.062	−0.022	−0.04	0.007	−0.013	0.02	−0.069	0.011	−0.08
Jordan	−0.172	0.006	−0.178	−0.001	−0.037	0.037	−0.172	0.009	−0.181
Morocco	−0.03	−0.051	0.021	0	−0.051	0.051	−0.03	−0.152	0.122
Malta	−0.089	−0.035	−0.054	−0.04	−0.034	−0.006	−0.049	0.046	−0.095
Average	−0.010	0.057	−0.067	0.086	0.023	0.054	−0.084	−0.019	−0.065

In essence, the overall change in the conditional variance that we find is identical to looking over small time frames and then adding up the individual changes.

This stability does not hold when we focus on either the time or covariate effects. To cement this, focus on the time effect component of the conditional variance change:[13]

$$\text{Time Effect}_{60-00} = \hat{v}_{60}(x_{10}) - \hat{v}_{10}(x_{10}). \tag{14}$$

Now, the time effect estimated from the subsample time effects is

$$\text{TimeEffect}_{60-80} + \text{Time Effect}_{80-10} = \hat{v}_{60}(x_{80}) - \hat{v}_{80}(x_{80}) + \hat{v}_{80}(x_{01}) - \hat{v}_{10}(x_{10}) \tag{15}$$

so that the difference between the full time effect and the subsample components is:

$$\hat{v}_{60}(x_{10}) + \hat{v}_{80}(x_{80}) - (\hat{v}_{60}(x_{80}) + \hat{v}_{80}(x_{10})). \tag{16}$$

This further decomposition shows that time effects (and covariate effects) in the subsamples are equal only if these four terms sum to zero.

Our results suggest that the reason we observe a big time effect for the African continent is that we have that the relative impact of covariates and time switched importance as we switched sub-periods. If the distribution was time constant, then all changes in the conditional variance would lie with movements in the covariates as in (11), however, when the conditional density changes over time, the sub-period results may not give supporting evidence as we now account for intermediate changes between the sub-periods as in (15). Given that three of the four components in (16) are small while the fourth component is quite large for the African bloc, this suggests that in the 1980–2010 sub-period, the main driver behind the increasing variance are changes in the covariates. This is compounded by a change in the conditional density which appears to by different from the remaining blocs of countries, leading to a larger increase in the conditional variational; this is reinforced from the findings in Table 4.

4.3. Single variable changes

The previous results show how changes in the conditional variance are impacted both over time and as the entire set of covariates changed. This turned out to be important for Latin American and sub-Saharan African as both the time and covariate effects vary greatly amongst individual countries within these regions. A further refinement of our approach is to single out specific changes in a given covariate to see which variables within our conditioning set have the greatest impact on changes of the conditional variance. This approach is undertaken by implementing the decomposition holding all variables fixed except one as we move across time.[14] This strategy allow to answer questions such as "What would happen if we only observed changes in human capital over time?" This approach will provide insights about the leading determinant(s) of the observed changes in conditional variance for each country and time period.

We perform the decomposition for each country holding two of the three covariates fixed at their 1960 levels allowing the third to change. We can think of our decomposition in this case as:

$$\hat{v}_t(x_{jt}, x_{-jt}) - \hat{v}_T(x_{jT}, x_{-jt}) = \underbrace{\hat{v}_t(x_{jt}, x_{-jt}) - \hat{v}_T(x_{jt}, x_{-jt})}_{\text{Time Effect}_j} + \underbrace{\hat{v}_T(x_{jt}, x_{-jt}) - \hat{v}_T(x_{jT}, x_{-jt})}_{\text{Covariate Effect}_j}, \tag{17}$$

where x_{jt} refers to the variable of interest in time period t and x_{-jt} refers to the remaining variables in time period t. Our primary results

[13] The covariate effect component will be identical, but with an opposite sign.

[14] This is in the spirit of the quantile decompositions of Machado and Mata (2005).

Table 8
Conditional variance effects by variable and region.

Bloc	Education			Population			Investment		
	Overall	Covariate	Time	Overall	Covariate	Time	Overall	Covariate	Time
1960–2010									
Overall	−0.035	0.042	−0.077	−0.048	0.002	−0.051	−0.039	−0.002	−0.037
OECD	0.066	0.035	0.031	0.061	0.006	0.055	0.063	0.003	0.060
SSA	−0.108	0.001	−0.109	−0.149	−0.001	−0.148	−0.136	−0.011	−0.125
LAC	−0.080	0.098	−0.178	−0.111	0.004	−0.115	−0.088	−0.001	−0.087
ASIA	−0.162	−0.028	−0.134	−0.137	−0.001	−0.136	−0.134	−0.001	−0.133
MENA	−0.010	0.059	−0.070	−0.008	−0.005	−0.003	−0.005	−0.005	0.000
1960–1980									
Overall	−0.010	0.014	−0.024	−0.008	0.000	−0.008	−0.006	−0.001	−0.005
OECD	0.014	0.006	0.008	0.006	0.003	0.004	0.003	0.004	−0.000
SSA	0.014	0.002	0.012	0.038	0.000	0.037	0.034	−0.007	0.041
LAC	−0.014	0.064	−0.078	−0.040	0.001	−0.041	−0.030	−0.003	−0.028
Asia	−0.147	−0.053	−0.094	−0.087	−0.004	−0.083	−0.085	−0.002	−0.083
MENA	0.055	0.007	0.047	0.061	−0.006	0.067	0.072	−0.003	0.075
1980–2010									
Overall	−0.022	0.000	−0.023	−0.035	−0.005	−0.031	−0.026	0.003	−0.029
OECD	0.048	0.039	0.008	0.042	−0.002	0.044	0.044	0.003	0.041
SSA	−0.125	−0.083	−0.041	−0.166	0.001	−0.167	−0.152	−0.001	−0.151
LAC	−0.049	0.049	−0.098	−0.081	−0.009	−0.072	−0.057	0.009	−0.067
Asia	0.000	−0.082	0.082	0.026	−0.011	0.037	0.029	0.006	0.023
MENA	−0.084	−0.045	−0.039	−0.081	−0.000	−0.081	−0.078	−0.010	−0.068

from this exercise are shown in Table 8.

Generally, we see that human capital had the largest impact on changes in the conditional variance, both across groups as well as for the first sub-period. Outside of the positive overall effect of all three covariates for the OECD, the three growth determinants lead to increases in the conditional variance over the 1960–2010 period for the other four regions. The impact of population growth and investment exhibit smaller impacts on the conditional variance and fluctuate more frequently across regions and time periods.

A further insight from our single variable decompositions lies in the 'covariate' and 'time' effects. In this setting the covariate effect is solely due to changes in the variable of interest while the time effect still captures changes in the conditional distribution over time. Looking at the covariate effects for the 1960–2010 time period, education appears to have the largest effect in general. The population covariate effect is quite small, as is the effect of investment, viewed purely from the covariate effect. These results appear in line with previous research that has investigated the impact of factors of accumulation on the bimodality of the unconditional distribution of output (Feyrer, 2008; Henderson et al., 2008). The larger time effects are consistent with unobserved factors over time, such as total factor productivity.

These results for our subgroup of African countries are in line with the evidence of Houssa et al. (2014). Using only a sample of African countries, they deploy a counterfactual decomposition framework based on a production frontier approach (which differs from the methods here as they use an unconditional density and do not deal with variance) and find that efficiency played a key role in changing the shape of the density of output per worker. In our setting efficiency would fall into the time effect, outside of the breadth of physical and human capital accumulation. Given the large time effects across all three components, for each time period for the sub-Saharan bloc, we also have evidence of unexplained factors contributing to our individual effects.

Another interesting finding is that the overall effects for both population and investment are largest for the SSA bloc for both 1960–2010 as well as for the later sub-period, 1980–2010. In fact the overall effect for education is the largest for the SSA bloc in the 1980–2010 sub-period as well, while it is the largest for the Asian bloc for the entire 1960–2010 period. It appears the individual, overall effects had the smallest impact, in general, for the MENA bloc, whereas the SSA and Asian blocs have the largest overall effects in both the full period results as well as in the individual sub-periods.

A few additional insights from the individual results. We see that for the LAC, SSA and MENA blocs, that the impact of education has the same direction of effects both for the 1960–2010 and 1960–1980 periods, while the LAC, Asia and MENA blocs have the same direction of effects for both population and investment. Looking over time, we see that several of the overall effects change in direction as we progress from the 1960–1980 to the 1980–2010 period. This is suggestive of some form of a change in the conditional density in the way in which the conditioning set impacts cross-country output.

Another way to understand the single variable results is through the common perception of the influence of each of the individual effects on output/growth from conditional mean models. Finding a positive influence on output/growth is common for both investment and population growth. However, Delgado et al. (2014) have provided extensive evidence that education does not influence the conditional mean of cross-country growth. Thus, one can interpret our findings here that education does matter, albeit through channels beyond the first moment, which is the most commonly studied in the growth empirics literature.

4.4. Robustness checks

It is natural to question how robust our findings our to various modeling decisions made at various points of the analysis. For example, do the results hold up if we were to use an Epanechnikov kernel rather than a Gaussian kernel to smooth the data, or if we use a different measure of cross-country output, such as RGDPO instead of RGDPE. In this section we explore the robustness of our findings. For brevity we only focus on the impacts on the regional averages over the full 1960–2010 period. Results for a specific country or different time periods are available upon request.[15]

First, we assess if the choice of kernel has an impact, at the regional level, on our findings. Table 9 presents the regional averages of our variance decomposition. It is clear that the main findings hold across regions. This is not surprising since it is well known in the nonpara-

[15] We thank three anonymous referees for suggesting the various robustness checks which appear here.

Table 9
Variance decomposition by region using Epanechnikov kernel. Bandwidths selected via least-squares cross-validation.

Region	Overall effect	Covariate effect	Time effect
OECD	0.09	0.04	0.05
SSA	−0.09	0.02	−0.11
LAC	−0.07	0.11	−0.18
ASIA	−0.17	−0.04	−0.14
MENA	−0.03	0.07	−0.10
Overall	−0.03	0.05	−0.07

Table 10
Variance decomposition by region using bandwidths selected via maximum likelihood cross-validation. Gaussian product kernel used to smooth the data.

Region	Overall effect	Covariate effect	Time effect
OECD	0.06	0.03	0.03
SSA	−0.08	−0.01	−0.08
LAC	−0.10	0.06	−0.15
ASIA	−0.14	−0.05	−0.10
MENA	0.02	0.04	−0.02
Overall	−0.04	0.02	−0.06

metric community that the choice of kernel has little impact on one's analysis (Henderson and Parmeter, 2015). We see that the conditional variation in the OECD decreased (though the magnitude has changed), and that the balance between time and covariates remains. The overall effect for the sub-Saharan countries is of equal magnitude thought now the covariate effect is positive (in Table 4 it was positive). The results for LAC are virtually identical with Table 5 as are the findings for MENA and Asian countries. It is clear that the choice of kernel has little impact on the overall insights of our variance decomposition.

While the impact of the kernel is known to be minimal in applied settings, the choice of bandwidth plays an integral role. To study the impact that different bandwidths have on our findings, we deploy MLCV rather than LSCV to estimate our bandwidths, but still use the Gaussian kernel to smooth the data. The new conditional variance findings appear in Table 10. The findings for OECD members is almost identical to our estimates calculated with LSCV bandwidths appearing in Table 3. Further, we see the same general patterns for the other groups of countries, sub-Saharan and Asian countries, on average, have an increase in conditional variation, arising from both covariates and time, while Latin American countries have an increase in conditional variation, with opposite impacts of time and covariates. One difference is with MENA countries, where now we have a decreasing variation. We note here that this could be due to the fact that there are only seven total countries for which to calculate the MENA results.

Paying less attention to the tools used to smooth the data (the kernel and the bandwidth), we instead focus on how we measure output, and if this might lead to different insights in our findings. Recall that our main analysis used RGDPE in PWT 8.1. An alternative measure for comparison across countries is RGDPO, measuring GDP on the output side. Keeping our estimation method identical to our baseline results (Gaussian kernel with bandwidths selected via LSCV), we present regional averages in Table 11 replacing RGDPE with RGDPO in our decomposition. For OECD members and sub-Saharan African and Asian nations we have the same findings. However, the change in variation for

Table 11
Variance decomposition by region using alternative measure of output per capita (rgdp_o). Gaussian product kernel used to smooth the data with bandwidths selected via LSCV.

Region	Overall effect	Covariate effect	Time effect
OECD	0.04	0.02	0.02
SSA	−0.13	−0.01	−0.12
LAC	−0.33	0.11	−0.44
ASIA	−0.16	−0.04	−0.12
MENA	0.06	0.13	−0.07
Overall	−0.11	0.04	−0.15

Latin American Countries, while qualitatively the same, is an order of magnitude larger than previously found with RGDPO while the overall effect for MENA countries is positive, whereas our earlier findings were an increase in conditional variation of output over time.

5. Implications and future research

This paper has taken a step towards the systematic study of changes in the conditional variance of cross-country output. We suggested construction of the conditional density to parse the second moment and investigate changes over time. This measure of conditional variation was further decomposed into two distinct pieces, a 'covariate' effect and a 'time' effect. The covariate effect uncovers the impact of the covariates on the conditional variance while the time effect examines unexplained changes in the conditional variation, such as total factor productivity and pure time dependent heteroscedasticity.

We found that changes in covariate accumulation plays a relevant role, explaining almost 50% of the overall change in the conditional variance over time. This usually presents itself in the form of divergence, leading to an expanding conditional variance over time, especially for the sub-period 1960–1980, while for the 1980–2010 sub-period this effect suggest conditional convergence. We compared results across blocs of countries and found that the sub-Saharan bloc experienced the greatest conditional divergence while the OECD bloc was the only one to experience conditional convergence. The time effect also represents a substantial predictor for changes in the conditional variation of the conditional density of cross-country output. The time effect is the largest for the sub-Saharan bloc.

We also furthered the insights from our decomposition results by considering further decompositions of each of the individual effects, time and covariate, as well as looking at the impact of changes for single covariates. Our findings suggest that the sub-Saharan bloc of countries experienced the largest change in the covariate and time effects moving from the 1960–1980 and 1980–2010 period. The individual covariate effects suggest that education contributed the greatest to changes in the conditional variation over time. Additionally, our exercise shows the usual catching up of the Asian bloc after 1980, mainly driven by factors of accumulation and a modest time effect (implying a role for other variables such as institutions that drives this accumulation process).

While conditional β-convergence is likely to continue to dominate empirical discussions of growth, we hope that the approach presented here will make the study of distributional variation over time more appealing to growth empiricists. These results can be used to determine if factor accumulation will indeed diminish the apparent increase in variation of output levels over time that has been at the crux of many growth debates. Additionally, further inclusion of common determinants of growth would be a useful extension.

Appendix A. Alternative estimation technique of conditional variance

While our focus on estimating the conditional variance was based on estimation of the conditional density, an alternative approach would be to estimate the skedastic function directly, ignoring modeling of the conditional density. To understand the implications this might have on our

Table 12
Variance decomposition by region using local-linear least-squares to estimate the conditional variance function. Gaussian product kernel used to smooth the data with bandwidths selected via LSCV and output measured as RGDPE.

Region	Overall effect	Covariate effect	Time effect
OECD	0.11	−0.02	0.14
SSA	−0.11	−0.02	−0.08
LAC	−0.07	−0.10	0.03
ASIA	−0.15	−0.10	−0.05
MENA	0.01	−0.04	0.05
Overall	−0.01	−0.05	0.04

Table 13
Variance decomposition by region using local-linear least-squares to estimate the conditional variance function. Gaussian product kernel used to smooth the data with bandwidths selected via AIC_c and output measured as RGDPE.

Region	Overall effect	Covariate effect	Time effect
OECD	0.17	0.06	0.10
SSA	0.01	0.14	−0.14
LAC	−0.04	0.13	−0.17
ASIA	−0.17	−0.04	−0.13
MENA	0.09	0.00	0.09
Overall	0.04	0.08	−0.04

findings we present estimates of our decomposition of conditional variance using traditional regression methods. Tables 12 and 13 present estimates using local-linear least-squares (Li and Racine, 2007) with bandwidths selected using either LSCV or AIC_c (Hurvich et al., 1998) bandwidth selection.

Using LSCV to select the bandwidths, Table 12 confirms many of our original findings, a decrease in conditional variation for OECD members, with increasing conditional variation for sub-Saharan African, Latin American and Asian regional blocs, on average, and a small change in the conditional variation over the 1960–2010 period for MENA countries. Some differences do arise pertaining the direction and magnitude of the covariate and time effects, but in general our main conclusions hold.

With bandwidths selected using AIC_c, Table 13 confirms some of our original findings but there are some differences as well. While we see a decrease in conditional variation for OECD members the change is much larger using the AIC_c bandwidths than using bandwidths selected with LSCV. Further, there is almost no change in the conditional variation for the sub-Saharan African bloc of countries over the 1960–2010 period, with almost equal, but opposite time and covariate effects. The overall change for Asian countries is quite similar, both in sign and magnitude, but while the Latin American countries have an increase in conditional variation, the signs of the covariate and time effects are reversed compared to the decomposition results using LSCV bandwidths appearing in Table 12.

References

Aitchison, J., Aitken, C.G.G., 1976. Multivariate binary discrimination by the kernel method. Biometrika 63 (3), 413–420.
Anderson, G., 2004. Making inferences about the polarization, welfare and poverty of nations: a study of 101 countries 1970–1995. J. Appl. Econom. 19, 537–550.
Azariadis, C., Starchurski, J., 2004. A Forward Projection of the Cross-Country Income Distribution, Working Paper 570, Kyoto University.
Barro, R., Lee, J.-W., 2013. A new data set of educational attainment in the World, 1950–2010. J. Dev. Econ. 104, 184–198.
Barro, R., Sala-i-Martin, X., 1991. Convergence across states and regions. Brook. Pap. Econ. Act. 1, 107–158.
Battisti, M., Parmeter, C.F., 2013. Clustering and polarization in the distribution of output: a multivariate perspective. J. Macroecon. 35 (1), 144–162.
Beaudry, P., Collard, F., Green, D.A., 2005. Changes in the world distribution of output per worker: how a standard decomposition tells and unorthodox story. Rev. Econ. Stat. 87 (4), 741–753.
Bianchi, M., 1997. Testing for convergence: evidence from nonparametric multimodality tests. J. Appl. Econom. 12 (4), 393–409.
Bliss, C., 2000. Galton's fallacy and economic convergence. Oxf. Econ. Pap. 51, 420–422.
Boccanfuso, D., Richard, P., Savard, L., 2013. Parametric and nonparametric income distribution estimators in CGE micro-simulation modeling. Econ. Model. 35, 892–899.
Cannon, E., Duck, N., 2000. Galton's fallacy and economic convergence. Oxf. Econ. Pap. 53, 415–419.
Canova, F., 2004. Testing for convergence clubs in income per capita: a predictive density approach. Int. Econ. Rev. 45 (1), 49–77.
Caselli, F., Feyrer, J., 2006. The marginal product of capital. Q. J. Econ. 122 (2), 535–568.
Dalgaard, C.-J., Vastrup, J., 2001. On the measurement of σ-convergence. Econ. Lett. 70, 283–287.
Delgado, M., Henderson, D.J., Parmeter, C.F., 2014. Does education matter for economic growth? Oxf. Bull. Econ. Stat. 76 (3), 334–359.
Desdoights, A., 1999. Patterns of economic development and the formation of clubs. J. Econ. Growth 4 (3), 305–330.
Durlauf, S., Johnson, P., 1995. Multiple regimes and cross country growth behaviour. J. Appl. Econom. 10 (4), 365–384.
Durlauf, S., Johnson, P., Temple, J., 2005. Growth econometrics. In: Aghion, P., Durlauf, S. (Eds.), Handbook of Economic Growth vol. 1. Elsevier, Amsterdam.
Durlauf, S., Kourtellos, A., Minkin, A., 2001. The local Solow growth model. Eur. Econ. Rev. 46 (4–6), 928–940.
S.N., Durlauf, 2009. The Rise and Fall of Cross-Country Growth Regressions. Hist. Polit. Econ. 41, 315–333.
Egger, P., Pfaffermayr, M., 2007. On testing conditional sigma-convergence. Oxf. Bull. Econ. Stat., 1–21.
Feenstra, R.C., Inklaar, R., Timmer, M.P., 2015. The next generation of the penn world table. Am. Econ. Rev. 105 (10), 3150–3182.
J.D., Feyrer, 2008. Convergence by Parts. B.E. J. Macroecon. 8, 1–35.
Fiaschi, D., Lavezzi, M., 2004. Distribution dynamics and nonlinear growth. J. Econ. Growth 8, 379–401.
Furceri, D., 2005. β and σ-convergence: a mathematical relation of causality. Econ. Lett. 89, 212–215.
Gisbert, F.J.G., 2003. Weighted samples, kernel density estimators and convergence. Empir. Econ. 28, 335–351.
Hall, P., Racine, J.S., Li, Q., 2004. Cross-validation and the estimation of conditional probability densities. J. Am. Stat. Assoc. 99 (468), 1015–1026.
Henderson, D.J., Papageorgiou, C., Parmeter, C.F., 2012. Growth empirics without parameters. Econ. J. 122 (559), 125–154.
Henderson, D.J., Parmeter, C.F., 2015. Applied Nonparametric Econometrics. Cambridge University Press, Cambridge, United Kingdom.
Henderson, D.J., Parmeter, C.F., Russell, R.R., 2008. Modes, weighted modes, and calibrated modes: evidence of clustering using modality tests. J. Appl. Econom. 23

(5), 607–638.
Houssa, R., Badunenko, O., Henderson, D.J., 2014. Significant drivers of growth in Africa. J. Prod. Anal. 42, 339–354.
Hyndman, R.J., Bashtannyk, D.M., Grunwald, G.K., 1996. Estimating and visualizing conditional densities. J. Comput. Graph. Stat. 5 (4), 315–336.
Huber, P., Pfaffermayr, M., 2010. Testing for conditional convergence in variance and skewness: the firm size distribution revisited. Oxf. Bull. Econ. Stat. 72 (5), 648–668.
C.M., Hurvich, J.S., Simono, C.-L., Tsai, 1998. Smoothing Parameter Selection in Nonparametric Regression Using an Improved Akaike Information Criterion. J. R. Stat. Soc. Series B 60 (2), 271–293.
Johnson, P., 2000. A nonparametric analysis of income convergence across the US states. Econ. Lett. 69, 219–223.
Johnson, P., 2005. A continuous state space approach to convergence by parts. Econ. Lett. 86, 317–321.
Jones, C.I., 1997. Convergence Revisited. J. Econ. Growth 2 (2), 131–153.
Kourtellos, A., 2003. Modeling Parameter Heterogeneity in Cross-Country Growth Regression Models, Mimeo, University of Cyprus.
Li, Q., Racine, J.S., 2003. Nonparametric estimation of distributions with categorical and continuous data. J. Multivar. Anal. 86, 266–292.
Li, Q., Racine, J.S., 2007. Nonparametric Econometrics: Theory and Practice. Princeton University Press, New Jersey.
Li, K.-W., Zhou, X., Pan, Z., 2016. Cross-country output convergence and growth: evidence from varying coefficient nonparametric method. Econ. Model. 55, 32–41.
Liu, Z., Stengos, T., 1999. Non-linearities in cross-country growth regressions: a semiparametric approach. J. Appl. Econom. 14 (5), 527–538.
Maasoumi, E., Racine, J.S., Stengos, T., 2007. Growth and convergence: a profile of distribution dynamics and mobility. J. Econom. 136 (2), 483–508.
Machado, J.A.F., Mata, J., 2005. Counterfactual decomposition of changes in wage distributions using quantile regression. J. Appl. Econom. 20 (4), 445–465.
Mankiw, N.G., Romer, D., Weil, D., 1992. A contribution to the empirics of economic growth. Q. J. Econ. 107 (2), 407–437.
Nadaraya, E.A., 1965. On nonparametric estimates of density functions and regression curves. Theory Appl. Probab. 10, 186–190.
Paap, R., van Dijk, H., 1998. Distribution and mobility of wealth of nations. Eur. Econ. Rev. 42 (7), 1269–1293.
Pittau, M.G., 2005. Fitting regional income distributions in the European Union. Oxf. Bull. Econ. Stat. 67 (2), 135–161.
Pittau, M.G., Zelli, R., 2006. Empirical evidence of income dynamics across EU regions. J. Appl. Econ. 21, 605–628.
Pittau, M.G., Zelli, R., Johnson, P., 2010. Mixture models and convergence clubs. Rev. Income Wealth 56 (1), 102–122.
Phillips, P.C.B., Sul, D., 2007. Transition modeling and econometric convergence tests. Econometrica 75 (6), 1771–1855.
Pritchett, L., 1997. Divergence, big time. J. Econ. Perspect. 11 (3), 3–17.
Quah, D.T., 1993a. Galton's fallacy and tests of the convergence hypothesis. Scand. J. Econ. 95, 427–443.
Quah, D.T., 1993a. Empirical cross-section dynamics in economic growth. Eur. Econ. Rev. 37, 426–434.
Quah, D.T., 1996a. Twin peaks: growth convergence in models of distribution dynamics. Econ. J. 106, 1045–1055.
Quah, D.T., 1996b. Empirics for economic growth and convergence. Eur. Econ. Rev. 40, 1353–1375.
Quah, D.T., 1996c. Convergence empirics across economies with (some) capital mobility. J. Econo. Growth 2 (1), 27–59.
Sala-i-Martin, X., 1996. The classical approach to convergence analysis. Econ. J. 106 (437), 1019–1036.
Stamatakis, D., Petrakis, P.E., 2005. Human Capital, Growth and Convergence Traps: Implications From a Cross-Country Analysis, Mimeo, Athens National and Kapodistrian University.
Wang, M.C., Van Ryzin, J., 1981. A class of smooth estimators for discrete distributions. Biometrika 68, 301–309.
Watson, G.S., 1964. Smooth regression analysis. Sankhya 26 (15), 359–372.