

WP/16/160

IMF Working Paper

Testing Piketty's Hypothesis on the Drivers of Income Inequality: Evidence from Panel VARs with Heterogeneous Dynamics

by Carlos Góes

***IMF Working Papers* describe research in progress by the author(s) and are published to elicit comments and to encourage debate.** The views expressed in IMF Working Papers are those of the author(s) and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

I N T E R N A T I O N A L M O N E T A R Y F U N D

IMF Working Paper

Western Hemisphere Department

Testing Piketty's Hypothesis on the Drivers of Income Inequality: Evidence from Panel VARs with Heterogeneous Dynamics

Prepared by Carlos Góes¹

Authorized for distribution by Alfredo Cuevas

August 2016

IMF Working Papers describe research in progress by the author(s) and are published to elicit comments and to encourage debate. The views expressed in IMF Working Papers are those of the author(s) and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

Abstract

Thomas Piketty's *Capital in the Twenty-First Century* puts forth a logically consistent explanation for changes in income and wealth inequality patterns. However, while rich in data, the book provides no formal empirical testing for its theoretical causal chain. In this paper, I build a set of Panel SVAR models to check if inequality and capital share in the national income move up as the $r-g$ gap grows. Using a sample of 19 advanced economies spanning over 30 years, I find no empirical evidence that dynamics move in the way Piketty suggests. Results are robust to several alternative estimates of $r-g$.

JEL Classification Numbers: O43, C33, C14, C12

Keywords: Income Inequality, Panel VAR, Factor Income Distribution

Author's E-Mail Address: cgoes@imf.org

¹ International Monetary Fund (WHD). I wholeheartedly thank Branko Milanovic, Peter Pedroni, Roberto Perrelli, Henrique Barbosa, Franz Loyola, Cornelius Fleischhaker, Marcelo Estevão, Troy Matheson, Iza Karpowicz, Michele Andreolli, Alexandra Martinez, Alfredo Cuevas, and participants of the WHDS1 seminar for their helpful comments. All errors and omissions are solely mine. This paper is fully reproducible: data and econometric programs are available at <https://github.com/omercadopopular/cgoes>.

Contents

	Page
1 Introduction	4
2 Piketty’s Model: What to Test	4
3 Data and Stylized Facts	6
4 Methodology	8
5 Empirical Models and Results	9
6 Robustness	21
7 Discussion: What is the Takeaway?	22
8 Conclusions	24
A Appendix: Data and Sources	26
A.1 Countries Sample	26
A.2 Data Sources	26
B Appendix: Standard Error Simulation Algorithm	27

List of Figures

1	Distribution of capital share and share of the top 1% over time	7
2	Contemporaneous correlations between $r - g$ spread and capital share or share of the top 1%	7
3	Results of Model 1: Heterogeneous composite impulse responses across sample	10
4	Results of Model 1: Median composite responses and confidence intervals	11
5	Results of Model 1: Decomposition of median composite responses	12
6	Correlation between heterogeneous responses and intergenerational elasticity of income and social expenditure, respectively	13
7	Results of Model 2: Heterogeneous composite impulse responses across sample	14
8	Results of Model 2: Median composite responses and confidence intervals	15
9	Results of Model 2: Decomposition of median composite responses	16
10	Contemporaneous correlation between GDP growth and changes in the savings rate	17
11	Results of Model 3: Heterogeneous composite impulse responses across sample	18
12	Results of Model 3: Median composite responses and confidence intervals	19
13	Results of Model 3: Decomposition of median composite responses	20
14	Robustness checks	21

1 Introduction

Capital in the Twenty-First Century, Thomas Piketty's *magnum opus*, has been widely praised for aggregating and presenting in a timely and accessible fashion the results of more than a decade of research spearheaded by Piketty and his co-authors. His work is likely to remain influential in the times to come, not only for its databases, which map the (traditionally scarce) information about income inequality but also because it is inspiring many economists to use the estimation techniques he and his co-authors popularized.

Piketty's theoretical explanations for changes in inequality patterns are interesting. In a nutshell, he argues that all other things constant, whenever the difference between the returns on capital (r) and the output growth rate (g) increases, the share of capital in national income increases. Furthermore, since capital income tends to be more unequally distributed than labor income, an increase of the capital share would likely lead to increased overall income (and, over time, wealth) inequality. Both of these are plausible relationships. However, while rich in data, *Capital* provides no formal empirical testing for these conjectures. In fact, there is little more than some apparent correlations the reader can eyeball in charts containing very aggregated multi-decennial averages.

The main contribution of this paper is to provide a rigorous empirical test of Piketty's hypotheses. For that purpose, I build Panel Structural Vector Autoregressive models to test if inequality and the capital share in the national income increase as the $r - g$ gap grows. I use Pedroni's (2013) recent Panel VAR technique, which controls for country fixed-effects and allows for full heterogeneity of dynamics across countries, to estimate these relationships in a panel of 19 advanced economies spanning over 30 years. This technique results in a distribution of impulse response functions, permitting a much more robust inference than those that rely on average estimates and assume slopes are homogeneous.

I find no empirical evidence that the dynamics move in the way Piketty suggests. In fact, for at least 75% of the countries examined, inequality responds negatively to $r - g$ shocks, which is in line with previous single-equation estimates by Acemoglu and Robinson (2015). The results also suggest that changes in the savings rate, which Piketty takes as relatively stable over time, are likely to offset most of the impact of $r - g$ shocks on the capital share of the national income. Thus, it provides empirical evidence to the model developed by Krusell and A. Smith (2015), who say Piketty relies on flawed theory of savings. The conclusions are robust to alternative estimates of $r - g$ and to the exclusion or inclusion of tax rates in the calculation of the real return on capital.

Knowing if Piketty's hypothesis is correct is of crucial importance. Such is the case because the policy solutions designed to counter increasing income inequality trends in advanced economies will need to tackle the underlying causes of inequality. The results presented here show that observed increases in income inequality in advanced economies are largely uncorrelated with changes in $r - g$, which suggests that one needs to look for the causes of inequality (and potential solutions) elsewhere.

This paper outlines Piketty's theoretical model and develops a strategy to test it. First, I present a simplified description of the basic theoretical relationships Piketty proposes. Then I describe the data and how I constructed some of the variables of interest and explain the methodology of the structural Panel VAR. I then present the results (paying special attention to how to interpret heterogeneous dynamics), perform some robustness tests, and try to relate the set of results to the bulk of the literature on income inequality, pointing to potential causes of increasing inequality which are not related to $r - g$. Finally, I conclude by summarizing the most important points of the paper and what they mean for future inequality dynamics.

2 Piketty's Model: What to Test

Building on a standard growth model, Piketty (2014) argues that the patterns of wealth and income concentration are defined by the difference in the real return on capital (r) and growth rates (g).

Here, I present a very stylized derivation of Piketty’s model and its implications, so as to develop a testing strategy.

In a closed economy where national income (Y) is a function of capital (K) and labor (L), $Y_t = K_t^\alpha L_t^{1-\alpha}$, defining the real return on capital as the marginal product of capital ($r \equiv \frac{\partial Y_t}{\partial K_t}$), means that the share of capital in the national income (α) can be represented as a function of the real rate of return on capital (r), since $r = \frac{\partial Y_t}{\partial K_t} = \alpha \frac{Y_t}{K_t}$. This basic definition is what Piketty calls the “first fundamental law of capitalism”:

$$\alpha = \frac{rK_t}{Y_t} \tag{1}$$

If the law of motion of capital is $K_{t+1} = (1 - \delta)K_t + sY_t$, where s is a constant savings rate, δ is a constant rate of depreciation, the population is constant, and $Y_{t+1} = (1 + g)Y_t$, at the steady state $\frac{d}{dt} \left[\frac{K_t}{Y_t} \right] = 0$, which implies:

$$\frac{\dot{K}}{\bar{K}} = \frac{\dot{Y}}{\bar{Y}}, \quad \frac{sY - \delta K}{K} = \frac{gY}{Y}, \quad \frac{\bar{K}}{\bar{Y}} = \frac{\bar{s}}{\bar{g} + \bar{\delta}} \tag{2}$$

where bars denote variables at their steady states. Piketty defines all of his variables in *net* terms, such that one should deduct depreciation from income, capital, and the savings rate. But, since both expressions are equivalent in their steady states (*cf.* Krusell and A. Smith 2015), and most standard textbooks use variables in gross terms, it is more expedient to explicitly to account for depreciation.

Substituting (2) into (1) yields what Piketty calls the “second fundamental law of capitalism”—an inverse relationship between the share of capital in the national income and economic growth:

$$\bar{\alpha} = \frac{\bar{r}\bar{s}}{\bar{g} + \bar{\delta}} \tag{3}$$

Dynamically, once one introduces random shocks to the steady state, one can think of the capital share (α) an moving average process centered around a steady state:

$$\alpha_t = \bar{\alpha} + \Phi(L)\xi_t \tag{4}$$

where $\Phi(L) \equiv (\sum_{j=0}^{\infty} \Phi_j L^j)$ is a polynomial of responses to stochastic innovations and ξ_t is an exogenous shock.

Taking the (net) savings rate as somewhat constant, Piketty and Zucman (2015) argue that the capital share, income inequality, and wealth inequality are rising functions of $r - g$. If Piketty is correct, then one should expect changes in the share of capital to be explained by contemporaneous and past changes in the spread between r and g (or, using the textbook model, between r and $g + \delta$). This means that all other things equal, a temporary exogenous innovation to $r - g$ should be expected to disturb the steady state and temporarily take α away from it in the same direction of the innovation. That is, if growth temporarily increases (or real return on capital decreases), the capital share is expected to temporarily fall. Piketty goes on to argue that as the returns on capital are more unequally distributed than labor income, a higher share of capital in the national income would lead to higher income and wealth inequality (Piketty 2014, ch. 7).

These relationships are empirical propositions and hence empirically testable. The testable hypothesis is that positive changes in $r - g$ lead to positive changes in inequality (z) and capital share (α). If there is no robust evidence for the baseline hypothesis (H_b), one cannot reject the alternative hypothesis (H_a) of zero (or negative) association between these variables:

$$\begin{aligned}
H_b & : \text{if } \Delta(r - g) > 0, \text{ then } \Delta z > 0, \quad \Delta\alpha > 0 \\
H_a & : \text{if } \Delta(r - g) > 0, \text{ then } \Delta z \leq 0, \quad \Delta\alpha \leq 0
\end{aligned}$$

Piketty stresses that these dynamics are long-run and asymptotic, hinting at potential difficulties in capturing them. Still, if his assumptions are correct, in a large enough sample, the observed dynamics should at least provide some evidence of such underlying relationships.

3 Data and Stylized Facts

In order to analyze the relationship between inequality, capital share and returns on capital to GDP growth differentials, I organize data from multiple sources. Data availability varies by country and the maximum range goes from 1980 through 2012. Data on the capital share and share of the top 1% are more reliable for advanced economies. For such reason, the final sample is an unbalanced long panel with annual observations for 19 advanced economies from different continents (see Appendix A for details).

Inequality is proxied by the share of national income held by the top 1%, as reported by Atkinson’s & Piketty’s World Top Incomes Database. The choice of such variable is rather obvious since this is Piketty’s measure of choice (as compared to alternative metrics, like the Gini index) to indicate rising inequality in advanced economies in the past decades. The annual capital share in the national income comes from the Penn World Tables.

To derive the second variable of interest —namely, the real return on capital net of real GDP growth—I take yearly averages of nominal long-term sovereign bond yields, calculate the post-tax nominal rates by deducting corporate income taxes and subtract from them annual percent changes in GDP deflators and real GDP growth. The r in $r - g$ is, then:

$$r_{i,t} = [(1 - \tau_{i,t})i_{i,t} - d_{i,t}] \tag{5}$$

where $r_{i,t}$ is the real return on capital, $\tau_{i,t}$ is the corporate income tax rate, $i_{i,t}$ is the nominal long-term sovereign bond yield, and $d_{i,t}$ is the annual percent change in the GDP deflator for country i at period t . GDP deflators and real growth rates are found in the IMF’s World Economic Outlook database. Tax rates come mostly from the OECD’s Tax Database—with the exception of the time series for Singapore, which was constructed independently.

The choice of sovereign bond yields as a proxy for returns on capital is not self-evident. In reality, the aggregate return on capital is a weighted average of returns across a plethora of investments. However, returns on government bonds are a good proxy for the purposes of this empirical exercise for two reasons. First, Piketty’s narrative centers on worries about a society facing the unmeritocratic rule of rentiers who profit from high returns on capital, which he illustrates by discussing how 19th century élites profited from high returns on government bonds (Piketty 2014, p. 132). Second, even if the level of return is different for different portfolios, the correlation between sovereign bonds and corporate bonds is historically very high. In fact, between 1996 and 2015, the correlation between prime AAA corporate bonds and U.S. treasuries was nearly perfect ($r = 0.93$); even when considering “junk” BB+ graded corporate bonds the co-movement between sovereigns and corporates was relatively tight ($r = 0.49$).

In the robustness section, I present results with two different proxies for r : short-term interest rates, and implied returns from the national accounts. I also present results ignoring tax rates.

Distributions of capital share and share of the top 1% in the sample show increasing trends, although with some signs of moderation in the late 2000s due to the global financial crisis. The upward trend is observed not only in the medians but also in the evolution of interquartile ranges, as shown in Figure 1.

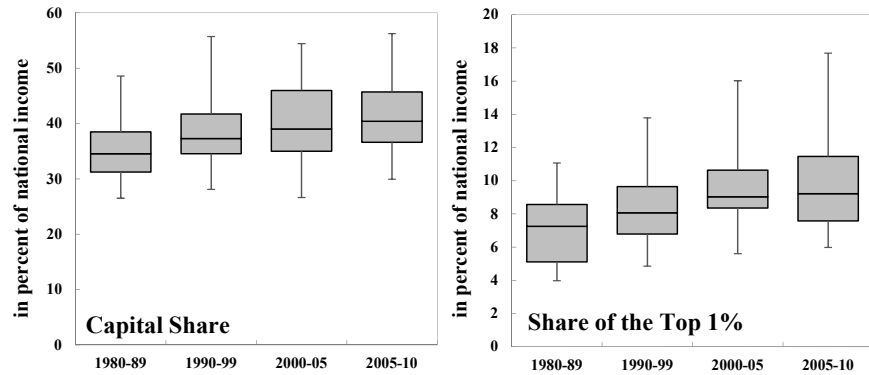


Figure 1: **Distribution of capital share and share of the top 1% over time.** Y-axis in percent, x-axis represents period averages. The sample refers to an unbalanced panel of 19 advanced economies ranging from 1981-2010. Boxplots show interquartile ranges and medians. Whiskers show minimums and maximums.

Basic stylized facts can provide some preliminary insights regarding whether or not the data support Piketty's assertions. Figure 2 plots the contemporaneous correlations between $r - g$ spreads and capital share, and share of the top 1%. Such basic correlations show no evidence of the relationship Piketty poses. Rather, the variables seem largely orthogonal.

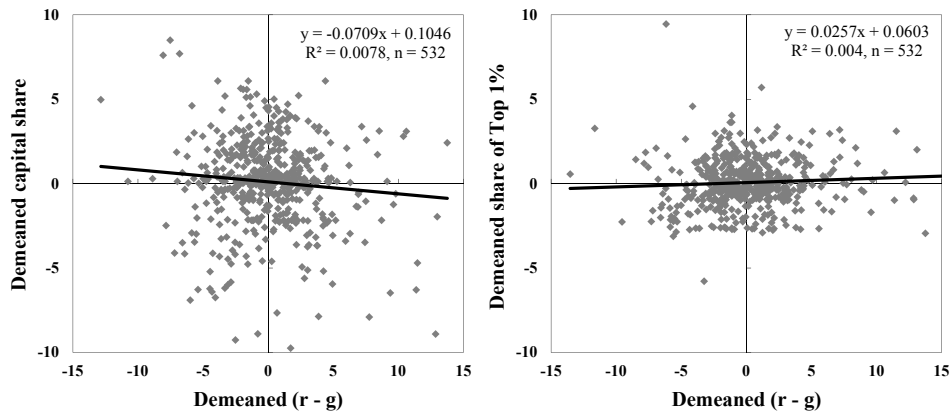


Figure 2: **Contemporaneous correlations between $r - g$ spread and capital share or share of the top 1%, respectively.** The sample refers to an unbalanced panel of 19 advanced economies ranging from 1981-2012. Variables are demeaned to account for time-invariant country-specific characteristics.

If the absence of a positive relationship persists with more refined estimation techniques, this could raise questions about Piketty's assertion that a higher $r - g$ will lead to higher inequality over the coming century. In the next section, I describe a model to test precisely that.

4 Methodology

Following Pedroni (2013), I estimate a structural Panel VAR that accommodates country-fixed effects, allows dynamics to be fully heterogeneous amongst panel members and decomposes dynamics between different responses to *idiosyncratic* and *common* shocks. In that context, for each member $i = [1, \dots, M]'$ of an unbalanced panel, let $y_{i,t}$ be a vector of n endogenous variables with country-specific time dimensions $t = [1, \dots, T_i]'$. To control for individual fixed-effects, I demean the data, resulting in $y_{i,t}^* = y_{i,t} - \bar{y}_i$, where $\bar{y}_i \equiv T_i^{-1} \sum_{t=1}^{T_i} y_{i,t} \forall i$. The baseline model is:

$$B_i y_{i,t}^* = A_i(L) y_{i,t-1}^* + e_{i,t} \quad (6)$$

where $y_{i,t}^*$ is a n -dimensional vector of demeaned stacked endogenous variables, $A_i(L)$ is a polynomial of lagged coefficients [$A_i(L) \equiv (\sum_{j=0}^{J_i} A_j^i L^j)$] with country-specific lag-lengths J_i , A_j^i is a matrix of coefficients, $e_{i,t}$ is a vector of stacked residuals, and B_i is a matrix of contemporaneous coefficients. J_i are selected based on appropriate criteria to assure that residuals approximate white noise.

To allow for heterogeneous dynamics, I first estimate and identify reduced-form VARs for each country i :

$$\begin{aligned} B_1 y_{1,t}^* &= A_1(L) y_{1,t-1}^* + e_{1,t} \\ &\vdots \\ B_M y_{M,t}^* &= A_M(L) y_{M,t-1}^* + e_{M,t} \end{aligned} \quad (7)$$

and then estimate another auxiliary VAR to recover *common* dynamics. Common dynamics are captured by averages, across individuals, for each period ($\bar{y}_t^* \equiv M^{-1} \sum_{i=1}^M y_{i,t}^*$):

$$\bar{B} \bar{y}_t^* = \bar{A}(L) \bar{y}_{t-1}^* + \bar{e}_t \quad (8)$$

After transforming the reduced-form residuals in (2) and (3) into their structural equivalents ($u_{i,t} = B_i^{-1} e_{i,t}$ and $\bar{u}_t = \bar{B}^{-1} \bar{e}_t$, respectively), I run nM OLS regressions to decompose the shocks into two terms:

$$\begin{aligned} u_{1,t} &= \Lambda_1 \bar{u}_t + \tilde{u}_{1,t} \\ &\vdots \\ u_{M,t} &= \Lambda_M \bar{u}_t + \tilde{u}_{M,t} \end{aligned} \quad (9)$$

where $u_{i,t}$ are *composite shocks*, $\bar{u}_{i,t}$ are *common shocks*, $\tilde{u}_{i,t}$ are *idiosyncratic shocks*, and Λ_i are n -by- n diagonal matrices with country specific loadings (the coefficients from the OLS regressions) denoting the relative importance of common shocks for each country. Note that $\tilde{u}_{i,t}$ vectors are truly idiosyncratic, since they are by construction orthogonal to the shocks derived from the average dynamics shared by all panel members.

Subsequently, I use standard methods described in Lütkepohl (2007) to recover the matrices of composite responses to structural shocks [$R_i(L)$] for each country, which are shown below in the vector moving average representations of M structural VARs:

$$\begin{aligned} y_{1,t}^* &= R_1(L) u_{1,t} \\ &\vdots \\ y_{M,t}^* &= R_M(L) u_{M,t} \end{aligned} \quad (10)$$

and then use the loading matrices estimated in (9) to decompose the composite responses into country-specific responses to common shocks and responses to idiosyncratic shocks:

$$\begin{aligned}
R_1(L) &= \Lambda_1 R_1(L) + (I - \Lambda_1 \Lambda_1') R_1(L) \\
&\vdots \\
R_M(L) &= \Lambda_M R_M(L) + (I - \Lambda_M \Lambda_M') R_M(L)
\end{aligned} \tag{11}$$

Equivalently, $R_i(L) = \bar{R}_i(L) + \tilde{R}_i(L)$, where $\bar{R}_i(L) \equiv \Lambda_i R_i(L)$ and $\tilde{R}_i(L) \equiv (I - \Lambda_i \Lambda_i') R_i(L)$. I then use the cross-sectional distribution of $R_i(L)$, $\bar{R}_i(L)$ and $\tilde{R}_i(L)$ to describe some properties of the collection of impulse response functions calculated, such as their medians, averages and interquartile ranges.

After recovering the point estimates of all the impulse response functions, I calculate standard errors of medians through a re-sampling simulation repeating all the steps above 500 times (see Appendix B for details).

5 Empirical Models and Results

I run three models with the methodology described above:

- In **Model 1**, $y_{i,t} \equiv [p_{i,t}, z_{i,t}]'$, where $z_{i,t}$ is the national income share of the top 1% for country i at period t and $p_{i,t} \equiv (r_{i,t} - g_{i,t})$ is the post-tax return on capital (r) net of real GDP growth (g).
- In **Model 2**, $y_{i,t} \equiv [p_{i,t}, k_{i,t}]'$, where $k_{i,t}$ is the share of capital in national income.
- In **Model 3**, $y_{i,t} \equiv [p_{i,t}, s_{i,t}, k_{i,t}]'$, incorporating the savings rate ($s_{i,t}$).

Note that, since the data are demeaned to control for country-fixed effects, if one takes the rate of depreciation as time-invariant, even if different for each country, it is not necessary to explicitly account for it in the empirical model. Depreciation is treated away once the data are demeaned as $p_{i,t}^* \equiv (r_{i,t} - g_{i,t})^* = (r_{i,t} - g_{i,t} - \delta_i)^*$.

The identification strategies of the empirical model the theoretical model proposed by Piketty: i.e., I take $r - g$ as the (most) exogenous variable. This means that I impose short-run lower triangular restrictions in the contemporaneous coefficients such that $r - g$ contemporaneously impacts the share of the top 1% and the capital share and $r - g$ responds to shocks to the share of the top 1% and the capital share only with a lag. In the robustness section, I present results with the inverse restrictions.

Pedroni's methodology allows for fully heterogeneous dynamics. This means that the results of the Panel VAR are more than *average* parameters and *average* impulse response functions —as they are in traditional panels, which impose homogeneous parameters. Rather, I have information about several moments of the distributions of impulse responses for each response horizon. With those data, I can then plot averages, medians and interquartile ranges of responses for a given horizon, as shown in Figure 3.

This is a much more informative way of reading results than in traditional panel VAR analyses. For instance, had I calculated average impulse responses from parameters estimated with traditional dynamic panels (e.g., difference or system GMM equations), I would have no way of knowing how many countries in the sample have dynamics that are similar to the average dynamics —as the underlying assumption is that parameters are equal for all countries. Effectively, knowing exactly how many countries in the sample present certain dynamics provides for much more robust inference than simply relying on average estimates. In addition to that, as shown by Pesaran & Smith (1995),

if individual dynamics are heterogeneous, aggregating or pooling slopes can lead to biased estimates, making individual regressions for each group member preferable.

Despite country-specific heterogeneity, the estimated VARs are stable and the variables are (trend) stationary. This means shocks should be interpreted as temporary and, following any shock, variables are expected to converge back to their means or deterministic trends over the long run. Assuming shocks are permanent would not affect the qualitative interpretation of results since this would simply mean accumulating temporary responses.

The results of **Model 1** show that for at least 75% of the countries an exogenous 1% positive shock to $r - g$ leads to an expected decrease in the share of the top 1% in the first five years (see Figure 3). While the median response is not statistically significant contemporaneously, it is negative and statistically significant from the first through the tenth year after the initial shock (see Figure 4). The responses to common and idiosyncratic shocks have the same (negative) sign. This is a challenge to the narrative which argues that systemic forces (as, for instance, globalization) could be exacerbating inequality as returns on capital increase (see Figure 5).

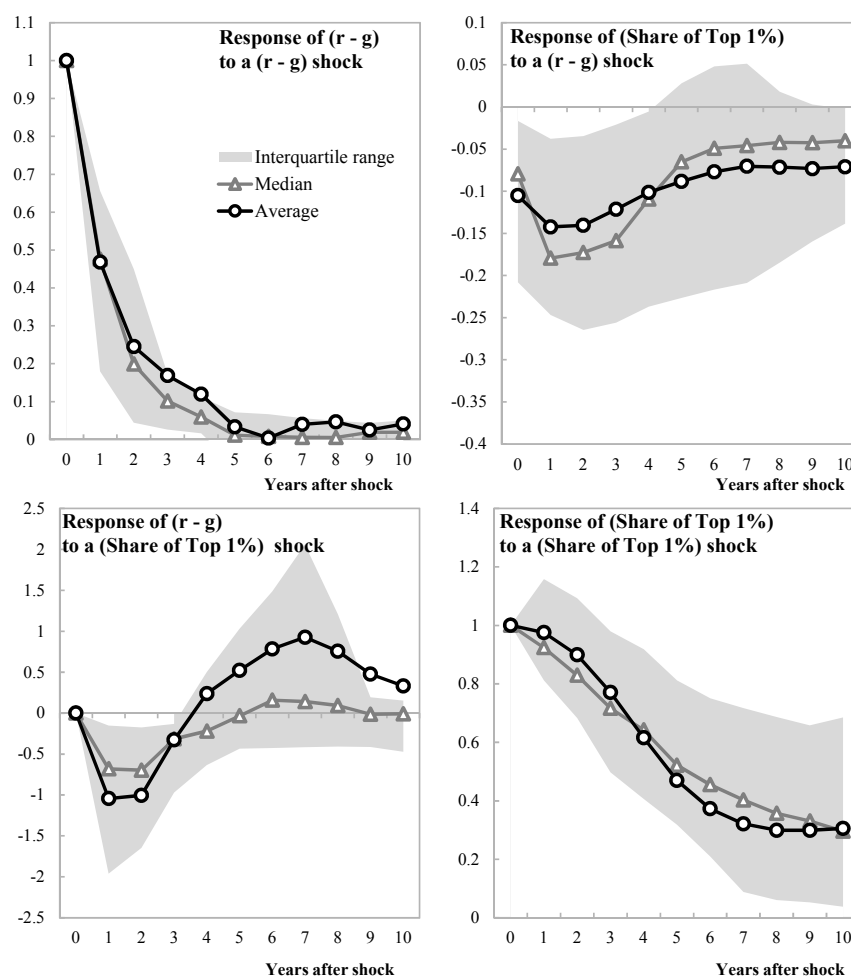


Figure 3: **Model 1: Heterogeneous composite impulse responses across sample.** The median, averages, and interquartile ranges were calculated from the distribution of IRFs of the 19 cross-sections.

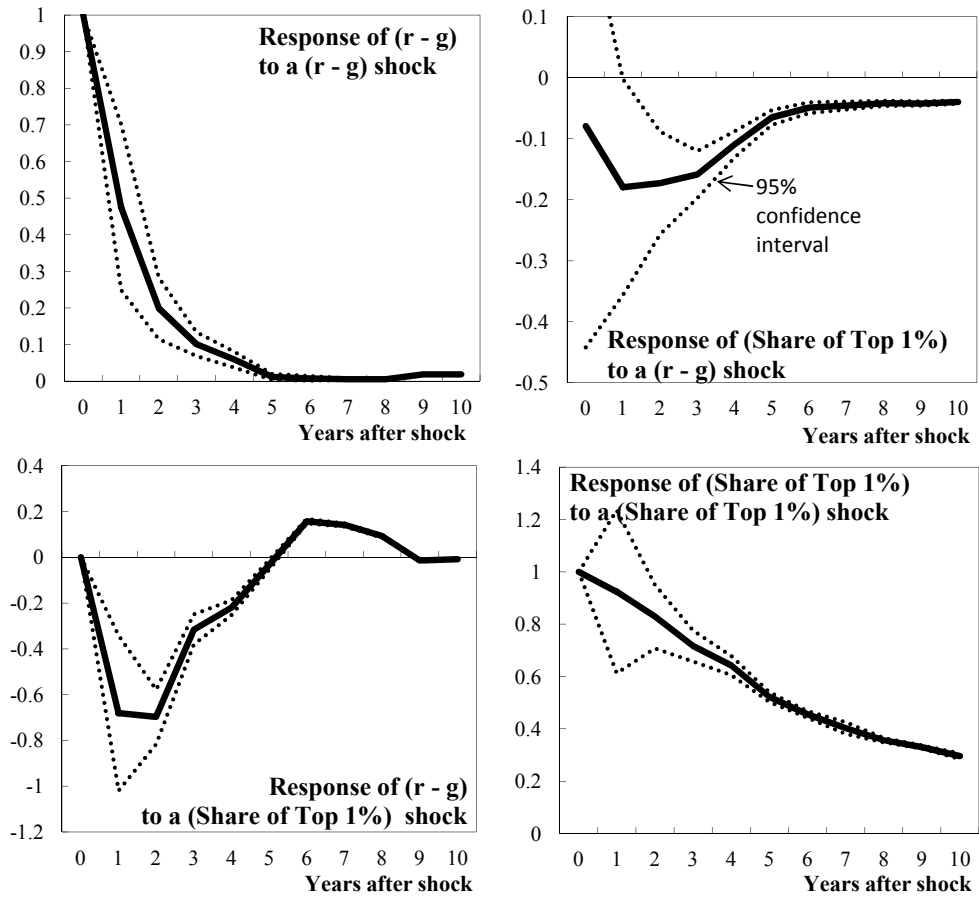


Figure 4: **Model 1: Median composite responses and confidence intervals.** Median response across a heterogeneous distribution of IRFs across 19 countries. Confidence intervals calculated from a resampling simulation with 500 repetitions. Since the distribution might change for each repetition in the simulation exercise, the confidence intervals do not represent the uncertainty around estimates for any particular country, but rather for the median estimate of the whole panel.

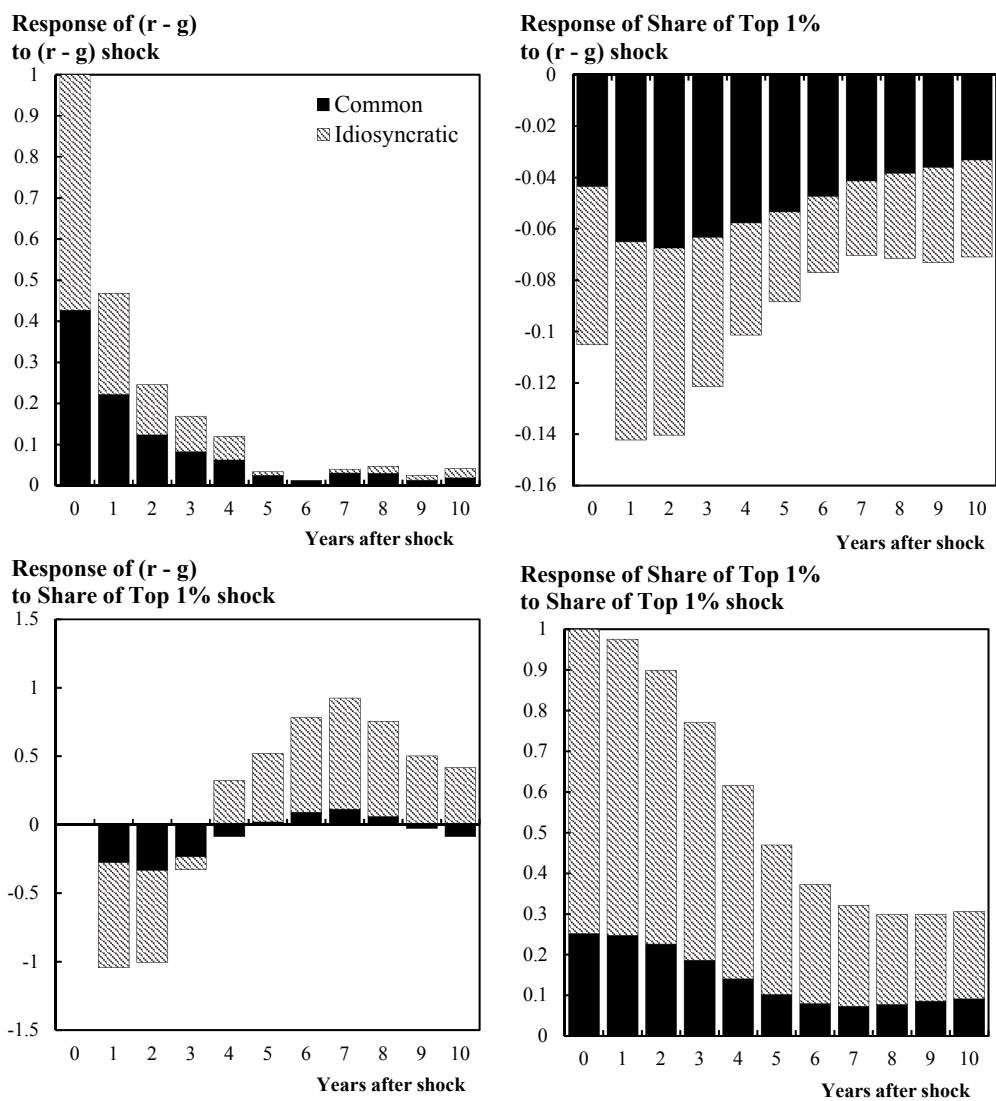


Figure 5: **Model 1: Decomposition of median composite responses.** Median responses can be decomposed into country-specific responses to common shocks and responses to idiosyncratic shocks through the use of loading factors that denote the relative importance of common shocks for each country

One of the advantages of estimating heterogeneous slopes is that one can explore the variability in parameters across the cross-sectional dimension of the panel. One can do that by correlating the response magnitudes of the different countries in the sample with potential covariates. I plot below some bivariate relationships which can provide some intuition on the causes and effects of the existent heterogeneity in responses.

There is a positive correlation between estimates of the intergenerational elasticity of income—which measures the share of a person’s income is explained by the parent’s income—and the size of the contemporaneous response of inequality to $r - g$ shocks. This hints on the idea that a higher transmission between $r - g$ and the share of the top 1% leads to a higher observed level of wage and wealth persistence.

Similarly, contemporaneous responses of inequality to $r - g$ shocks are positively correlated with social expenditure as a share of GDP for countries in the sample. This suggests that societies which are more vulnerable to larger shocks to inequality tend to be more willing to resort to government in order to provide social services. This would be a Rawlsian interpretation on how to hedge one’s exposure to inequality under a “veil of ignorance” about the future: insofar as social services are usually targeted at the lower end of income distribution, they can be seen as a kind of insurance against higher market income inequality.

Such correlations are shown as simple bivariate scatterplots in Figure 6.

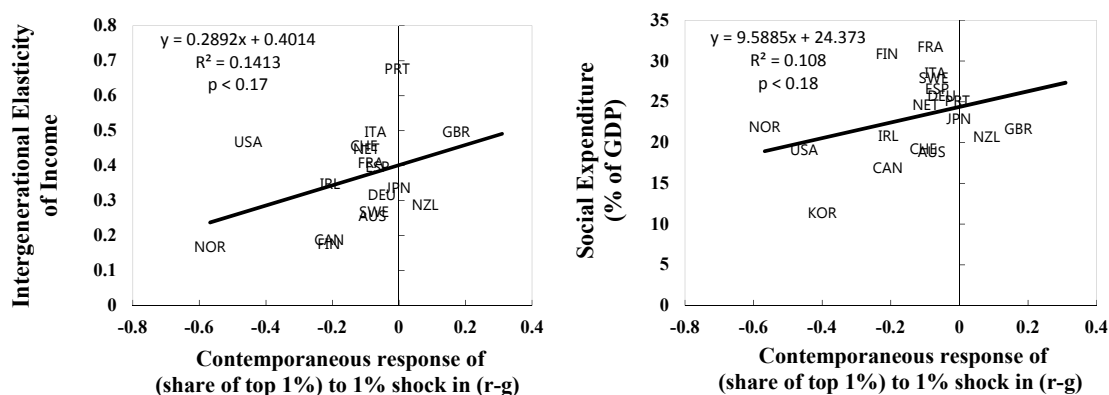


Figure 6: **Correlation between heterogeneous responses and intergenerational elasticity of income and social expenditure, respectively.** Data on intergenerational elasticity of income from Corak (2016) and Causa and Johansson (2013). Data on social spending from the OECD’s social expenditure database.

However, these results need to be taken with caution, since the regressions slopes are not statistically significant at standard critical levels. This could be either a byproduct of the small cross-sectional dimension of the sample or an indication that in a broader sample these results would not hold. Exploring further the causes and consequences of the heterogeneity in the relationship between $r - g$ dynamics and inequality is a potential avenue of future research.

I then turn to **Model 2**, which checks the relationship between $r - g$ and the capital share. The results are less clear cut than with Model 1. However, in the first two years, for at least 75% of the countries capital share responds (mildly) negatively to positive shocks to $r - g$, with the median going from negative to slightly positive then approaching zero asymptotically (see Figure 7). The median response is not statistically significant (see Figure 8).

The median response of $r - g$ to capital share shocks is negative and statistically significant. This is intuitive since it is possibly reflecting diminishing marginal returns on capital—for any fixed growth rate, increased capital shares are likely to decrease returns and hence drive $r - g$ down. When compared to Model 1, idiosyncratic shocks are more dominant (see Figure 9).

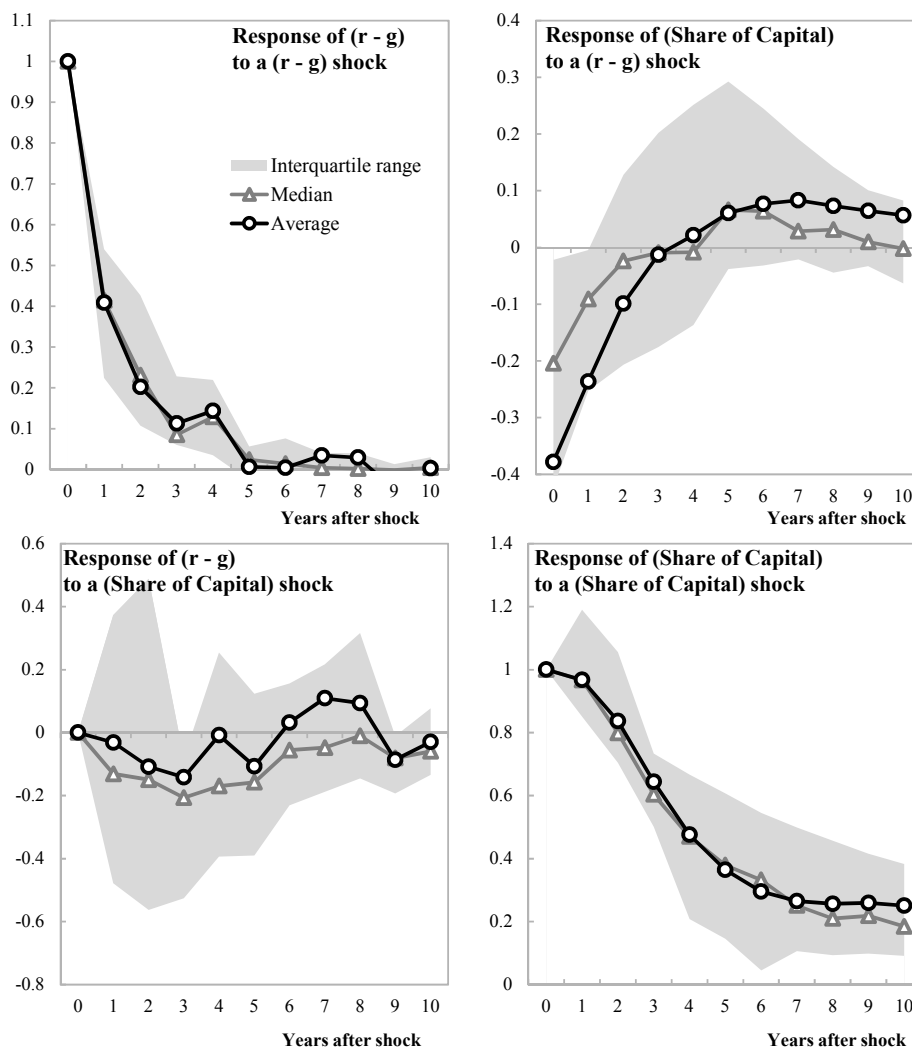


Figure 7: **Model 2: Heterogeneous composite impulse responses across sample.** The median, averages, and interquartile ranges were calculated from the distribution of IRFs of the 18 cross-sections.

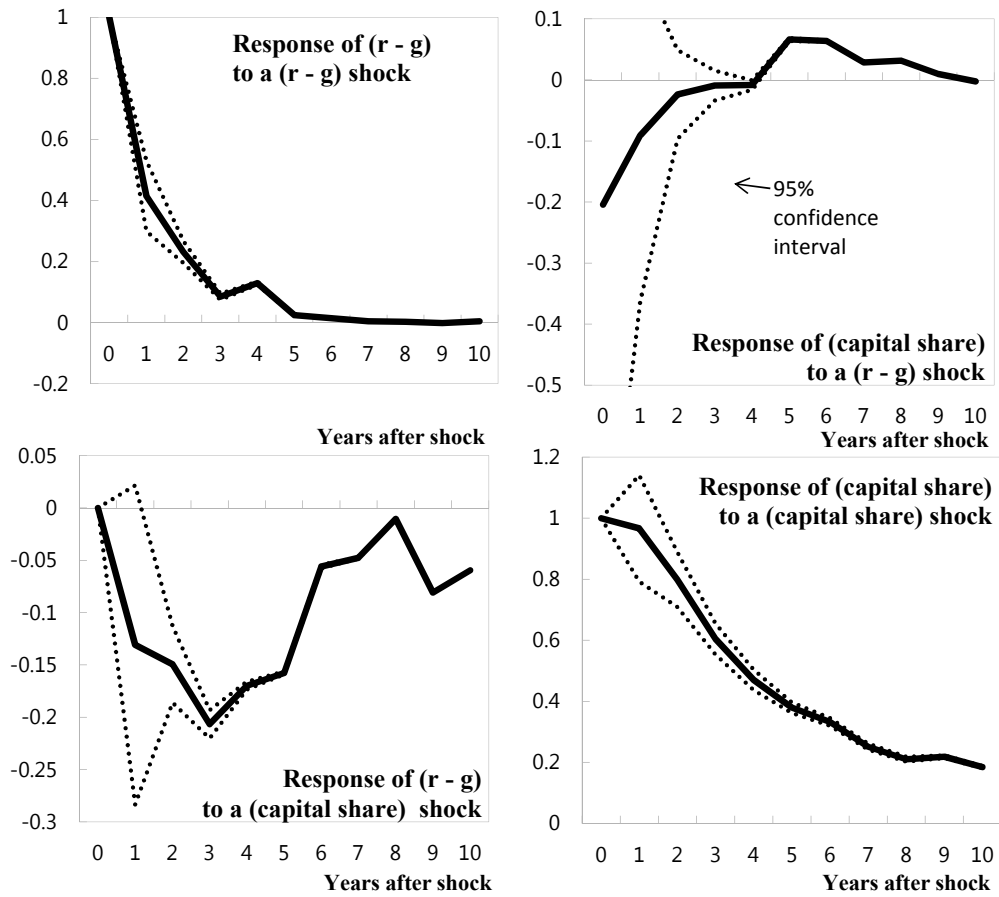


Figure 8: **Model 2: Median composite responses and confidence intervals.** Median response across a heterogeneous distribution of IRFs across 18 countries. Confidence intervals calculated from a resampling simulation with 500 repetitions. Since the distribution might change for each repetition in the simulation exercise, the confidence intervals do not represent the uncertainty around estimates for any particular country, but rather for the median estimate of the whole panel.

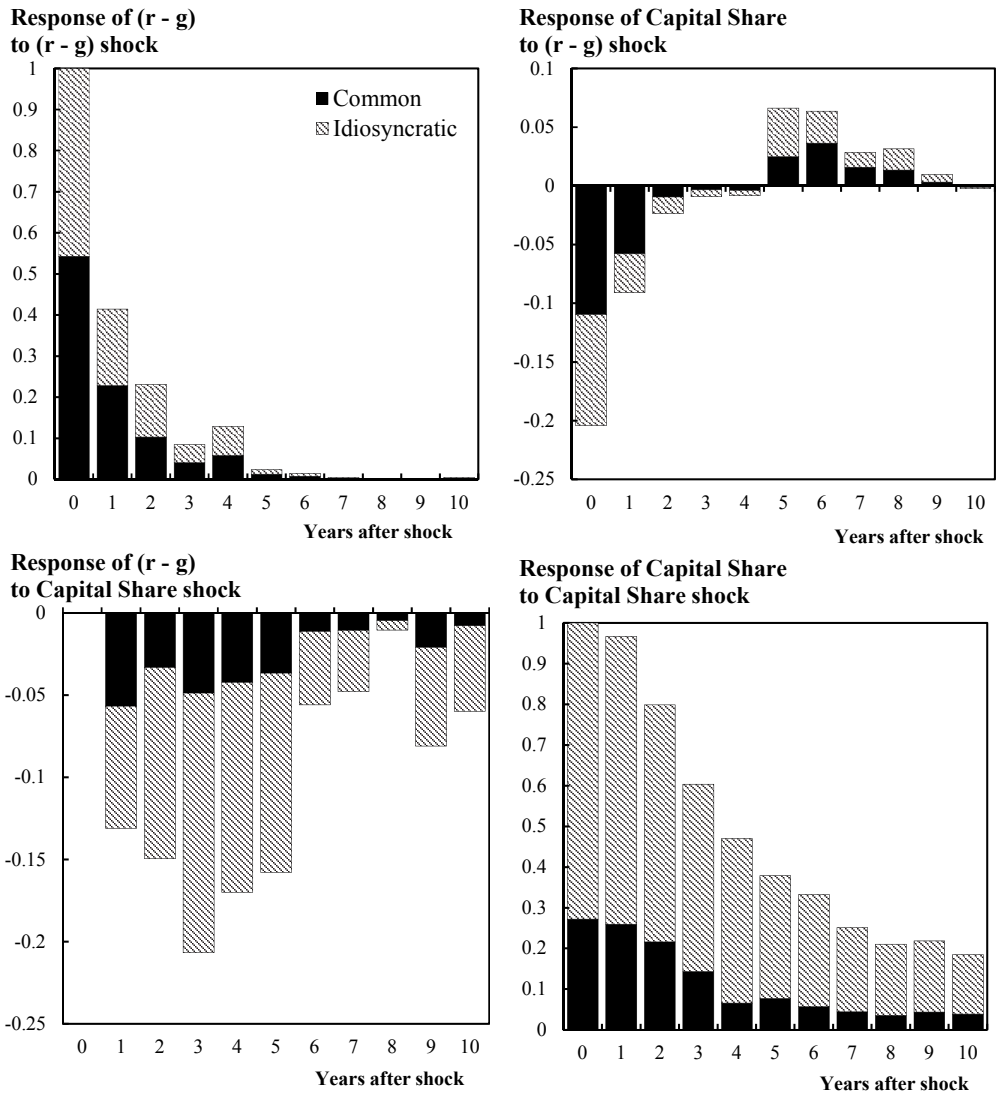


Figure 9: **Model 2: Decomposition of median composite responses.** Median responses can be decomposed into country-specific responses to common shocks and responses to idiosyncratic shocks through use of loading factors that denote the relative importance of common shocks for each country

Such result seems unusual, as even in standard models one should expect the capital share to increase following $r - g$ shocks. However, the absence of such evidence can be explained by the fact that it is unreasonable to assume that the savings rate is exogenous to $r - g$. In fact, there are strong theoretical reasons why savings rate should behave pro-cyclically. If higher growth increases expected life-time wealth, then it increases consumption in all periods of life (*cf.* Deaton 2005) —and higher savings represent higher future consumption. In line with this theoretical rationale, changes in the savings rate in this sample are positively and statistically significantly correlated with GDP growth rates (see Figure 10).

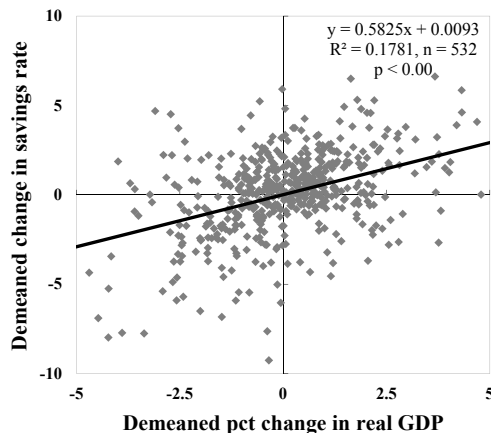


Figure 10: **Contemporaneous correlation between GDP growth and changes in the savings rate.** The sample refers to an unbalanced panel of 18 advanced economies ranging from 1981-2012. Variables are demeaned to account for time-invariant, country-specific characteristics.

Since lower growth rates are associated with a decrease in savings rates, it is hard to know *a priori* the net effect of lower growth (which, all other things equal, are equivalent to a *positive* $r - g$ shock) on the capital share. This happens because, if the decrease in the savings rate is large enough, it can offset any predicted effect of a $r - g$ shock on the capital share. To account for this, **Model 3** includes the savings rate as an endogenous variable.

The relationship observed with basic contemporaneous correlations is also reflected in the impulse response functions. For at least 75% of the countries, the savings rate responds negatively to a positive shock to $r - g$ (see Figure 11). The median response of the savings rate to $r - g$ shocks is large, negative, and statistically significant throughout the 10-year response horizon (see Figure 12). This result explains the apparent contradiction of the capital share responding negatively to such shocks.

The median response of the capital share to $r - g$ shocks is not statistically significant the first years after the initial shock and, as dynamics are incorporated, the median response of the capital share quickly net out to somewhere close to zero (see Figure 12). Common shocks contribute very little to capital share dynamics in Model 3 (see Figure 13).

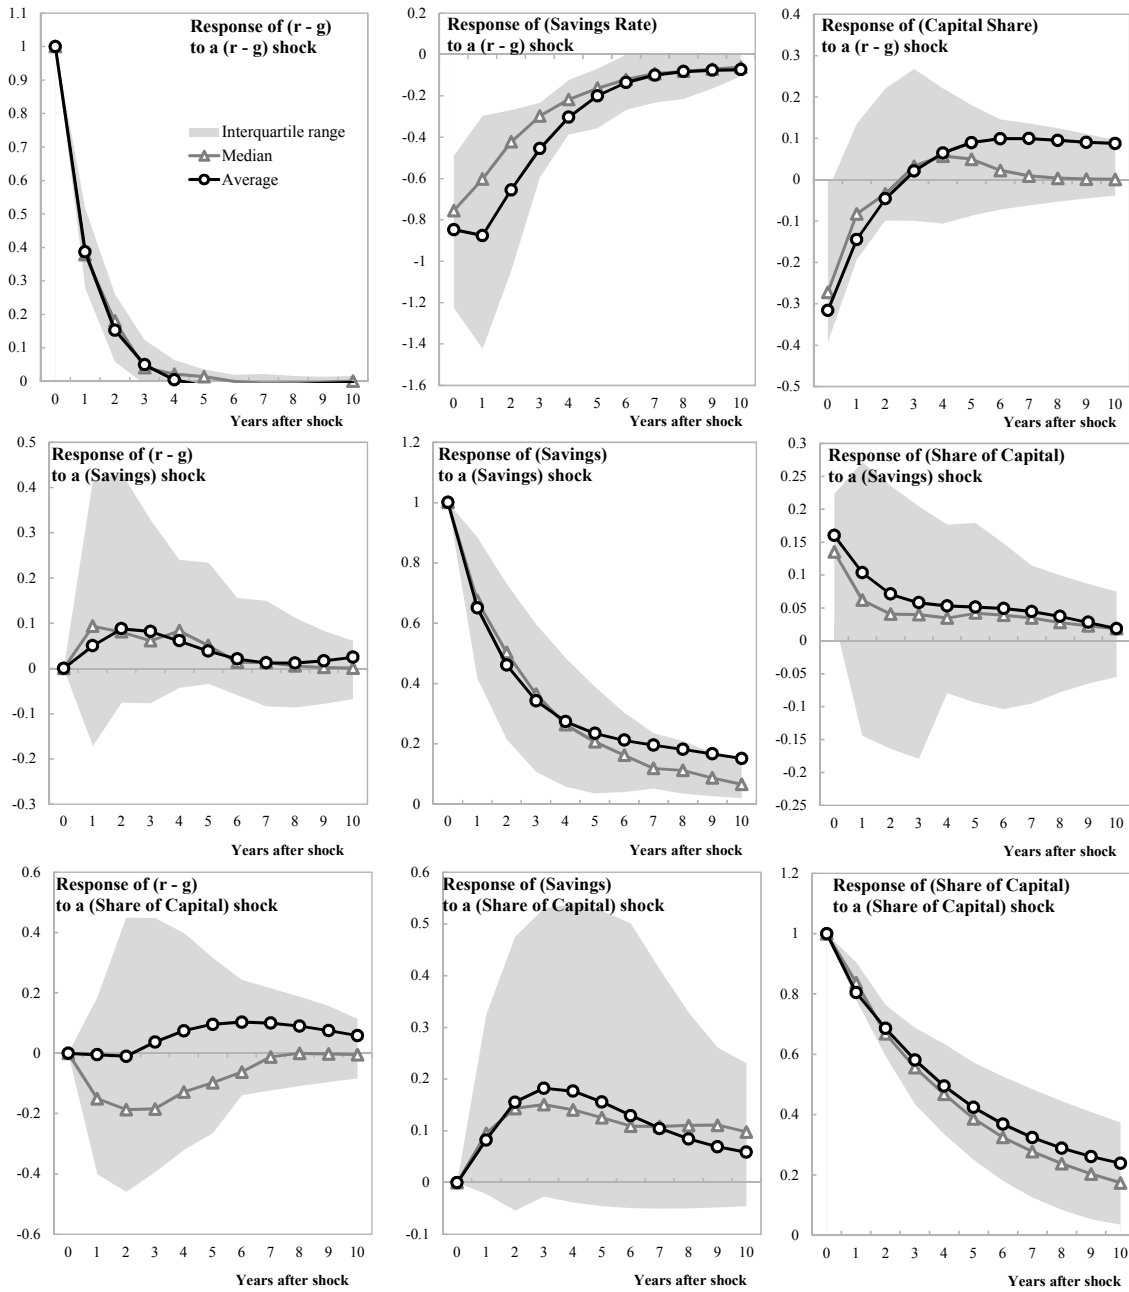


Figure 11: **Model 3: Heterogeneous composite impulse responses across sample.** The median, averages, and interquartile ranges were calculated from the distribution of IRFs of the 18 cross-sections.

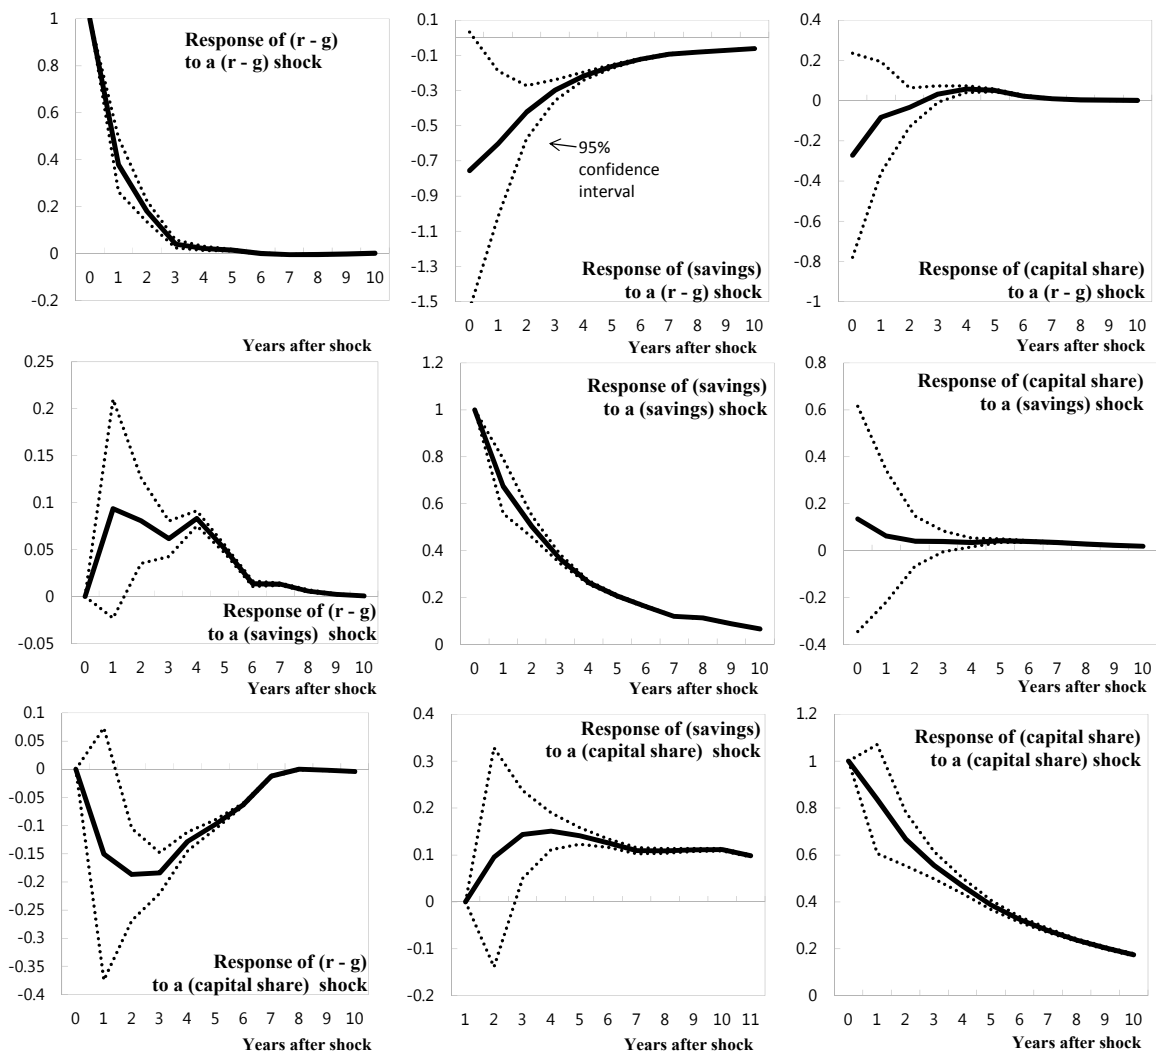


Figure 12: **Model 3: Median composite responses and confidence intervals.** Median response across a heterogeneous distribution of IRFs of 18 countries. Confidence intervals calculated from a resampling simulation with 500 repetitions. Since the distribution might change for each repetition in the simulation exercise, the confidence intervals do not represent the uncertainty around estimates for any particular country, but rather for the median estimate of the whole panel.

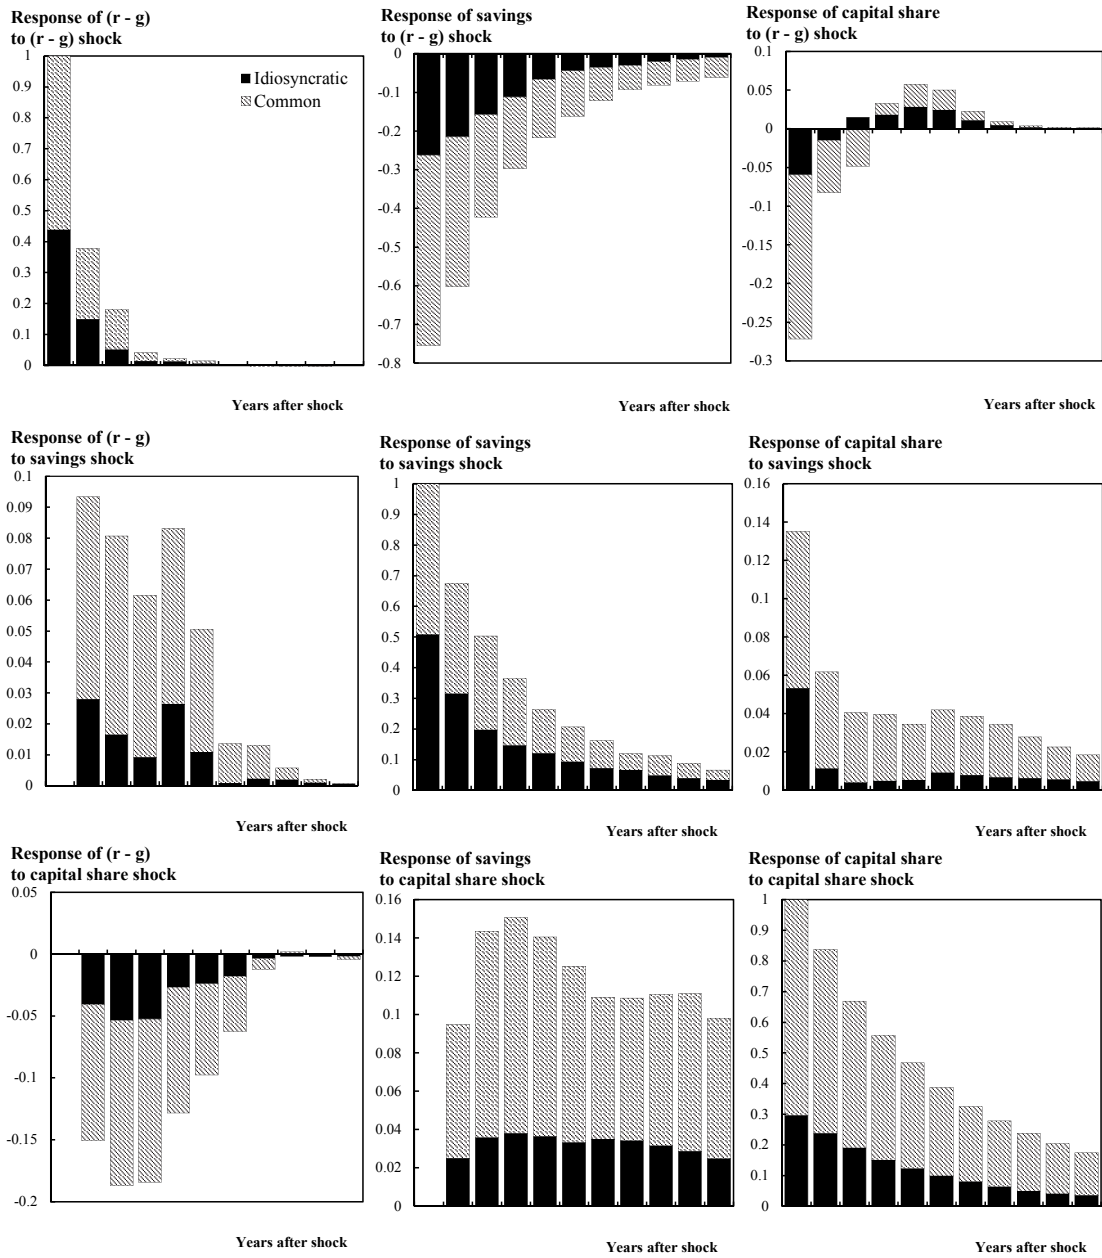


Figure 13: **Model 3: Decomposition of median composite responses.** Median responses can be decomposed into country-specific responses to common shocks and responses to idiosyncratic shocks through use of loading factors denoting the relative importance of common shocks for each country

6 Robustness

I run several alternative specifications to check the robustness of results. I do that by modifying baseline Models 1 and 3. I present all alternative results in Figure 14.

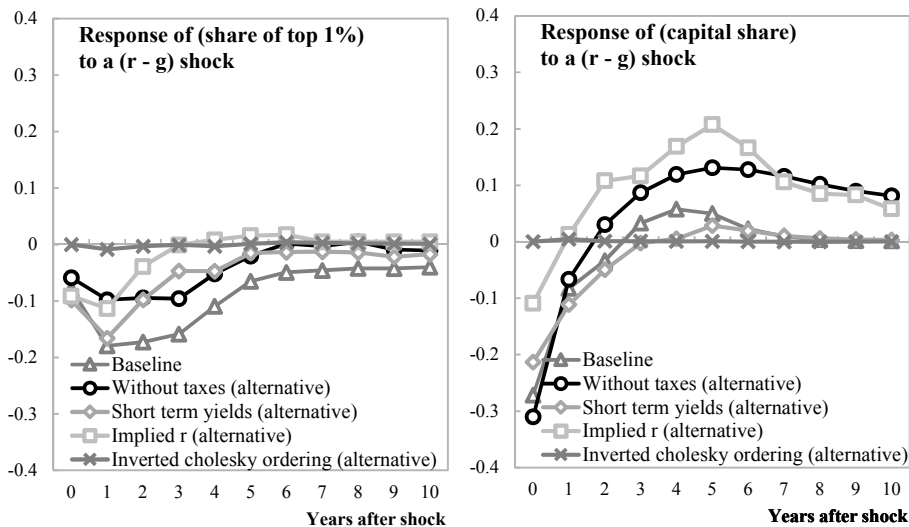


Figure 14: **Robustness checks - re-estimation of Models 1 and 3: Median composite responses.** This figure compares median responses across a distribution of IRFs of 19 (Model 1) / 18 (Model 3) countries using different specifications.

First I run an alternative model ignoring tax rates. Tax rates are important for inequality, as pointed out by Piketty and Saez (2013), but it could be that deducting tax rates from returns is adding noise to the results. However, when re-estimating the models without taking taxes into account, the results are similar to the baseline. The median responses of the share of the top 1% to $r - g$ shocks are still negative, albeit less negative than in Model 1.

When replacing long-term bond yields with short term yields (usually central bank policy rates), median responses are very similar to the baseline for the initial response horizons. In outer years, responses net out to zero faster than in the baseline.

I then re-estimate the models using an implied measure of r from the national accounts. Deriving this measure is straightforward. Since capital share is the return on capital times the capital-to-income ratio ($\alpha = rK/Y$), it follows that the return on capital is the capital share divided by the capital-to-income ratio ($r = \alpha Y/K$). After calculating implied returns on capital for all countries from the Penn World Tables, I deducted growth rates and re-estimated Model 1. Median responses still behave in a comparable fashion to all other specifications.

Overall, the results are rather robust to different measures of r . There is no evidence that the share of the top 1% responds positively to increases in $r - g$ spread.

Similarly, for the capital share, despite small differences, the results are virtually the same as in Model 3 —with the savings rate playing a counter-balancing role after $r - g$ shocks and median capital share responses going from slightly negative to slightly positive before converging towards zero.

Finally, I re-run both models with an inverted Cholesky ordering, i.e. with the share of the top 1% and the share of capital in the national income as the most exogenous variable. It could be the case that the original restrictions imposed on the baseline specifications were hiding a positive effect of $r - g$ over share of the top 1% and the share of capital. However, in the alternative specification, such

effects are very close to zero, which renders them both statistically and economically insignificant. Yet again, in this last alternative specification, there is no evidence in support of Piketty’s hypotheses of $r - g$ increasing the capital share and income inequality.

While baseline conclusions are reasonably robust to all the different specifications estimated, there are several other possible extensions of the empirical models which can provide additional robustness to the results. These include, for instance, using the Penn World Table’s estimates of time-varying depreciation rates, using measures of equity prices as a proxy for the return on capital, incorporating foreign savings to the model, or using a different measure for inequality (such as, for instance, the Gini coefficient). I intend to pursue some of these in future research.

7 Discussion: What is the Takeaway?

Piketty’s conclusion that inequality will increase in the future rests on the underlying assumption that, as growth decreases over time, driving the $r - g$ spread to increase, capital-to-income ratios will increase. However, the results from Model 2 and Model 3 fail to show robust positive responses of capital share to shocks in $r - g$ and cast doubt on Piketty’s conjecture. A possible reason for this is that Piketty could be underestimating diminishing returns of capital —thereby overestimating the elasticity of substitution between capital and labor, whose empirical estimates tend to be much lower than what he assumes (*cf.* Rognlie 2014). This relationship is illustrated in this paper by the negative median responses of $r - g$ to positive capital share shocks.

Another less emphasized but equally important problem with Piketty’s conjecture is highlighted by Krusell and Smith (2015), who argue that Piketty’s predictions are grounded on a flawed theory of savings —namely, that the savings rate net of depreciation is constant —which exacerbates the expected increase in capital-to-income ratios as growth rates tend to zero. They present an alternative model in which agents maximize inter-temporal utility and arrive in a setting in which, on an optimal growth path, the savings rate is pro-cyclical. By showing that the savings rate responds negatively to negative growth shocks (which, in turn, are translated into positive $r - g$ shocks) for at least 75% of the countries in the sample, the results of Model 3 provide empirical support to Krusell and Smith’s analysis.

Piketty (2012) says in his online notes: “with $g = 0\%$, we’re back to Marx apocalyptic conclusions,” in which capital share goes to 100% and workers take home none of the output. While this is logically consistent with the model’s assumptions, empirically there seem to be endogenous forces preventing that: non-negligible diminishing returns on capital and pro-cyclical changes in the savings rate. These are two different ways in which the transmission mechanism from $r - g$ to the capital share might get stuck: with the former, at the limit the rate of return on capital tends to zero and there is no dynamic transmission; with the latter, if growth approaches zero, the savings might ultimately become zero, offsetting any effect of lower growth on capital share. They are, however, fundamentally different: the first regards the production function and technological change, while the second has to do life-cycle behavior of capital owners.

Regarding inequality, the results from Model 1 contradict Piketty’s prediction stating that following exogenous shocks in $r - g$ inequality should increase. In fact, for at least 75% of the countries in the sample, the result is negative. These findings are in line with previous results by Acemoglu and Robinson (2015), who found negative coefficients in single-equation panel models when regressing $r - g$ on the share of the top 1%. This paper goes further, not only because the model takes all variables as endogenous, but also because it incorporates tax variability across countries. Additionally, by decomposing between common and idiosyncratic shocks, rather than using time dummies, Model 1 does not throw away potentially important information about the effect of structural forces (e.g., globalization) on these dynamics —which, as Milanovic (2014) argues, is a problem with Acemoglu and Robinson’s analysis.

The fact that a positive $r - g$ spread does not lead to higher inequality is not necessarily sur-

prising. As illustrated by Mankiw (2015) through a standard model that incorporates taxation and depreciation, even if $r > g$, one can arrive in a steady state inequality which does not evolve into an endless inegalitarian spiral. Milanovic (2017, forthcoming) explains that the transmission mechanism between $r > g$ and higher income inequality requires all of the following conditions to hold: (a) savings rates have to be sufficiently high; (b) capital income needs to be more unequally distributed than labor income; and (c) a high correlation between drawing capital income and being on the top of the income distribution. In a dynamic fashion, this paper shows that this mechanism is getting stuck because the negative responses of the savings rate to $r - g$ shocks violate the first condition, thereby preventing higher levels inequality when compared to those observed before the increases in $r - g$.

Since estimated dynamics do not confirm Piketty’s theory, observed income inequality in many advanced economies over the past decades are probably explained by factors other than the spread between r and g . In fact, there is evidence that recent inequality trends are not related to the distribution of national income between factors of production but primarily to the rising inequality of labor income (*cf.* Francese and Mulas-Granados 2015). Indeed, there are many potential explanations for the rising labor income inequality —as, for instance:

- Dabla-Norris et al. (2015), after evaluating cross-country evidence, find that past changes in inequality in advanced economies are associated the most with two labor market changes: higher skill premia and lower union membership rates. Jaumotte and Buitron (2015) also present results that correlate changes in labor market institutions, particularly lower union density, with increases in income inequality in advanced economies.
- Aghion et al. (2015) suggest innovation plays a significant role. If innovators are rewarded with higher incomes due to a temporary technological advantage (in a Schumpeterian fashion), inequality would be exacerbated. The authors show that innovation explains about a fifth of the higher inequality observed in the U.S. since 1975.
- Mare (2016) and Greenwood et al. (2012) argue that changes in mating behavior helped exacerbate income inequality. The probability that someone will marry another person with a similar socio-educational background (labeled “assortative mating”) increased *in tandem* with the rise in income inequality in U.S. in the recent decades. The interaction between higher skill premia and higher assortative mating exacerbates household income inequality, because the gap between higher and lower earners became larger and couples became more segregated.
- Chong and Gradstein (2007) use a dynamic panel to show inequality tends to decrease as institutional quality improves. The underlying logic is that if the basic rules of economic behavior are not symmetrically enforced, the rich will have a higher chance to extract economic rents, thereby increasing inequality. Acemoglu and Robinson (2015) make a similar argument. They say that economic institutions affect the distribution of skills in society, indirectly determining inequality patterns.

None of the evidence provided in the papers listed above is definite or mutually exclusive. Rather, they are most likely complementary explanations for the recent developments of a complex phenomenon such as income inequality. But what they do have in common is that none of them point to capital-to-income or $r - g$ dynamics as the driving causes of inequality patterns.

Some years after the publication of *Capital*, Piketty (2015) himself recognized the “rise in labor income inequality in recent decades has evidently little to do with $r - g$, and it is clearly a very important historical development.” He nonetheless emphasized that a higher $r - g$ spread will be important and will exacerbate future inequality changes.

However, the results in this paper show that this is likely not to be the case. The results corroborate the idea that recent inequality changes are not explained by $r - g$ but also that new shocks to $r - g$ will likely not lead to higher inequality, as there is no evidence that shocks to $r - g$

increase income inequality. Combined, the observed endogenous dynamics of $r - g$ and the share of the top 1% and the capital share, respectively, cast doubt on the reasonability of Piketty's prediction about inequality trends.

A potential caveat of this analysis is that its sample size might be too short to capture the long run relationship under scrutiny. While this is certainly possible, there are some limits to such criticism. First, since the focus of the exercise are advanced countries, one expects them to be close to their respective steady states. In such case, it is more likely that the thirty year sample is enough to capture the steady states and the effect of shocks that disturb such steady states. Second, Acemoglu and Robinson (2015) run single-equation models using 10- and 20-year averages going back all the way to the 1800s and find no statistically significant evidence that increases in $r - g$ positively impact the share of the top 1% in the national income. Unfortunately, using those multi-decennial averages in an endogenous dynamic panel framework would render the identification restrictions so strong (namely, that one variable has no impact over the other for 20 years) that it would make the entire exercise futile.

8 Conclusions

The purpose of this paper is very Popperian in nature: it tests interesting, logically consistent, and falsifiable hypotheses. Thereby, it contributes to the literature by checking the empirical veracity of a very influential theory regarding income inequality patterns.

In doing so, I found no evidence to corroborate the idea that the $r - g$ gap drives the capital share in national income. There are endogenous forces overlooked by Piketty —particularly the cyclicity of the savings rate —which balance out predicted large increases in the capital share. On inequality, the evidence against Piketty's predictions is even stronger: for at least 75% of the countries, the response of inequality to increases in $r - g$ has the opposite sign to that postulated by Piketty.

These results are robust to different calculations of $r - g$. Regardless of taking the real return on capital as long-term sovereign bond yields, short-term interest rates or implied returns from national accounting tables, the dynamics move in the same direction. Additionally, including or excluding taxes does not alter the qualitative takeaways from the results either.

Knowing if increases $r - g$ lead to inequality is very important, not only for economics as a science of human action but also for the policy repercussions of such conclusions. Without knowing the underlying causes of such trends, it is impossible to design policy actions to counter them.

Inequality is a complex phenomenon and its trends are very sluggish. It is certainly possible that the long terms relationships Piketty proposes exist and are simply not captured by the 30 years of data for the 19 advanced economies included in this sample. However, the best available data show that, if one is looking to potential solutions to increasing income inequality, one should not focus on $r - g$, but elsewhere.

References

- Acemoglu, Daron and James A. Robinson (2015). “The Rise and Decline of General Laws of Capitalism”. In: *Journal of Economic Perspectives* 29.1, pp. 3–28. DOI: 10.1257/jep.29.1.3.
- Aghion, Philippe et al. (2015). *Innovation and Top Income Inequality*. Working Paper 21247. National Bureau of Economic Research. DOI: 10.3386/w21247.
- Causa, Orsett and Åsa Johansson (2013). “Intergenerational social mobility in OECD countries”. In: *OECD Journal: Economic Studies* 2010.1, pp. 1–44. ISSN: 1468-2362.
- Chong, Alberto and Mark Gradstein (2007). “Inequality and Institutions”. In: *The Review of Economics and Statistics* 89.3, pp. 454–465. DOI: 10.1162/rest.89.3.454.
- Corak, Miles (2016). *Inequality from Generation to Generation: The United States in Comparison*. Discussion Paper 9929. Institute for the Study of Labor (IZA).

- Dabla-Norris, Era et al. (2015). *Causes and Consequences of Income Inequality: A Global Perspective*. IMF Staff Discussion Note 15/13. International Monetary Fund. DOI: 10.5089/9781513555188.006.
- Deaton, Angus S. (2005). “Franco Modigliani and the Life Cycle Theory of Consumption”. In: *Banca Nazionale del Lavoro Quarterly Review* 58.233-234, pp. 91–108.
- Francesce, Maura and Carlos Mulas-Granados (2015). *Functional Income Distribution and Its Role in Explaining Inequality*. IMF Working Paper 15/244. International Monetary Fund. DOI: 10.5089/9781513549828.001.
- Greenwood, Jeremy et al. (2012). *Technology and the Changing Family: A Unified Model of Marriage, Divorce, Educational Attainment and Married Female Labor-Force Participation*. Working Paper 17735. National Bureau of Economic Research. DOI: 10.3386/w17735.
- Jaumotte, Florence and Carolina Osorio Buitron (2015). *Inequality and Labor Market Institutions*. IMF Staff Discussion Note 15/14. International Monetary Fund. DOI: 10.5089/9781513577258.006.
- Krusell, Per and Anthony Smith (2015). “Is Piketty’s ‘Second Law of Capitalism’ Fundamental?” In: *Journal of Political Economy* 123.4, pp. 725–748. DOI: 10.1086/682574.
- Lütkepohl, Helmut (2007). *New Introduction to Multiple Time Series Analysis*. Berlin: Springer.
- Mankiw, Greg (2015). “Yes, $r > g$. So What?” In: *American Economic Review: Papers & Proceedings* 105.5, pp. 43–47. DOI: 10.1257/aer.p20151059.
- Mare, Robert D. (2016). “Educational Homogamy in Two Gilded Ages: Evidence from Inter-generational Social Mobility Data”. In: *The ANNALS of the American Academy of Political and Social Science* 663.1, pp. 117–139. DOI: 10.1177/0002716215596967.
- Milanovic, Branko (2014). *My take on the Acemoglu-Robinson critique of Piketty*. <http://glineq.blogspot.com/2014/08/my-take-on-acemoglu-robinson-critique.html>. Accessed: 11-Feb-2016.
- (2017, forthcoming). “Increasing capital income share and its effect on personal income inequality”. In: *After Piketty’s “Capital in the 21st century”: The agenda for fighting inequality*. Ed. by Heather Boushey, Bradford DeLong, and Marshall Steinbaum. Cambridge: Harvard University Press.
- Pedroni, Peter (2013). “Structural Panel VARs”. In: *Econometrics* 1, pp. 180–206. DOI: 10.3390/econometrics1020180.
- Pesaran, M Hashem and Ron Smith (1995). “Estimating long-run relationships from dynamic heterogeneous panels”. In: *Journal of Econometrics* 68.1, pp. 79–113. DOI: 10.1016/0304-4076(94)01644-F.
- Piketty, Thomas (2012). *Economics of Inequality. Course Notes: Models of Growth & Capital Accumulation. Is Balanced Growth Possible?*. <http://piketty.pse.ens.fr/fr/teaching/10/25>. Accessed: 11-Feb-2016.
- (2014). *Capital in the Twenty-First Century*. Cambridge: Harvard University Press.
- (2015). “About Capital in the Twenty-First Century”. In: *American Economic Review: Papers & Proceedings* 105.5, pp. 48–53. DOI: 10.1257/aer.p20151060.
- Piketty, Thomas and Emmanuel Saez (2013). “A Theory of Optimal Inheritance Taxation”. In: *Econometrica* 81.5, pp. 1851–1886. DOI: 10.3982/ECTA10712.
- Piketty, Thomas and Gabriel Zucman (2015). “Wealth and Inheritance in the Long Run”. In: *Handbook of Income Distribution, v. 2A*. Ed. by Jan Fagerberg, David C. Mowery, and Richard R. Nelson. Amsterdam: North Holland. Chap. 15, pp. 1303–1368. DOI: 10.1016/B978-0-444-59429-7.00016-9.
- Rognlie, Matthew (2014). *A note on Piketty and diminishing returns to capital*. http://www.mit.edu/~mrognlie/piketty_diminishing_returns.pdf. Accessed: 11-Feb-2016.

A Appendix: Data and Sources

A.1 Countries Sample

Samples are slightly different in Model 1 compared to Models 2 and 3 to account for dynamic (in)stability in different samples.

Sample for Model 1: Australia, Canada, Finland, France, Germany, Italy, Ireland, Japan, Korea, Netherlands, New Zealand, Norway, Portugal, Singapore, Spain, Sweden, Switzerland, the United Kingdom, and United States.

Sample for Models 2 and 3: Australia, Canada, Denmark, Finland, France, Italy, Japan, Korea, Netherlands, New Zealand, Norway, Portugal, Singapore, Spain, Sweden, Switzerland, the United Kingdom, and United States.

A.2 Data Sources

Variable	Scale	Source
Share of the top 1%	Percent of national income	World Top Income Database
Capital share	Percent of national income	Penn World Tables
Real GDP growth	Annual percent change	World Economic Outlook
GDP deflators	Index	World Economic Outlook
Gross national savings	Percent of national income	World Economic Outlook
Corporate income tax rates	Percent	OECD's Tax Database
10-YR sovereign bond yields	Percent	Bloomberg & Haver Analytics
Short-term interest rates	Percent	Haver Analytics & IMF
Implied return on capital	Percent	Derived from Penn World Tables

B Appendix: Standard Error Simulation Algorithm

Before starting the simulation, I run the individual VARs in equation (7) and collect predicted values ($\hat{y}_{i,t}$) and reduced form residuals ($e_{i,t}$). I then transform reduced-form residuals into their structural equivalents ($u_{i,t} = B_i^{-1}e_{i,t}$) and run the following algorithm:

1. I follow these steps $k = 500$ times:
 - (a) I re-sample the structural residuals. Let $\dot{u}_{i,t}$ denote re-sampled structural residuals $u_{i,t}$.
 - (b) I then transform re-sampled structural residuals back into reduced-form residuals by multiplying them with B : $\dot{e}_{i,t} = B_i\dot{u}_{i,t}$.
 - (c) I sum the predicted values of $y_{i,t}^*$ and re-sampled residuals to create pseudo-series for all members M . Let the pseudo-series be $\tilde{y}_{i,t} \equiv \hat{y}_{i,t} + \dot{e}_{i,t}$.
 - (d) I re-estimate the model with the pseudo-series: $\tilde{B}_i\tilde{y}_{i,t} = \tilde{A}(L)\tilde{y}_{i,t-1} + \tilde{e}_{i,t}$.
 - (e) I collect the matrices of responses, extract individual vectors for each pseudo-IRF, organize them into a matrix, and calculate the median response of this distribution.
 - (f) If this repetition $< k$, I go back to (a), otherwise, I move to step 2.
2. After I repeat this procedure k times the result will be a set of distribution matrices D of simulated medians:

$$D_m = \begin{bmatrix} \tilde{\rho}_{1m}^1 & \cdots & \tilde{\rho}_{1m}^k \\ \vdots & \ddots & \vdots \\ \tilde{\rho}_{hm}^1 & \cdots & \tilde{\rho}_{hm}^k \end{bmatrix}_{hxk} \quad (\text{B.1})$$

where m is the m^{th} response variable, h is the response horizon, and k is the number of repetitions of the simulation exercise.

3. From D_m I take square root of the second moment of each row to build a vector of standard errors:

$$\sigma_{\rho_m} = \begin{bmatrix} \text{Var}[\{\tilde{\rho}_{1m}^1, \dots, \tilde{\rho}_{1m}^k\}]^{1/2} \\ \vdots \\ \text{Var}[\{\tilde{\rho}_{hm}^1, \dots, \tilde{\rho}_{hm}^k\}]^{1/2} \end{bmatrix}_{hx1} \quad (\text{B.2})$$