A Note on Debt, Growth and Causality

Arindrajit Dube*

Preliminary (This version May 30, 2013)

Abstract

Abstract: This note documents the timing in the relationship between the debt-to-GDP ratio and real GDP growth in advanced economies during the post World War II period using the Reinhart and Rogoff dataset. I first show that the debt ratio is more clearly associated with the 5-year past average growth rate, rather than the 5-year forward average growth rate–indicating a problem of reverse causality. Indeed, there is little evidence of a lower growth rate above the 90 percent threshold when using the 5-year forward average growth rate. I use a number of simple tools to account for some of the reverse causality in the bivariate regression–such as using forward growth rate, instrumenting the current debt ratio with its lag, and controlling for lagged GDP growth rates. These simple methods of accounting for reverse causality diminish the size of the association by between 50 and 70 percent, with the linear regression estimate indistinguishable from zero. Finally non- and semi-parametric plots provide visual confirmation that the relationship between debt-to-GDP ratio and growth is essentially flat for debt ratios exceeding 30 percent when we (1) use forward growth rates, (2) control for past GDP growth, or both.

1 Introduction

Recent work by Herndon, Ash and Pollin (2013)–hereafter HAP–has demonstrated that the apparent sharp drop in average growth in observations with a debt ratio of 90 or higher was an artifact of inadvertent exclusion of certain country-years and a weighting scheme which amplified the problems of data exclusion. Once all of the available post WWII data is used, or a less sensitive weighting scheme is used, there is no indication of a sharp drop in the mean growth rate for debt-to-GDP ratio exceeding 90 percent.

In their responses, Reinhart and Rogoff argue that the negative correlation between debt-to-GDP ratio and growth in the corrected data still supports their original contention (Reinhart and Rogoff 2013a, 2013b). That contention is that in the long run, a build-up of debt slows down the economy, and this is mostly what is captured by the simple negative correlation between growth and the debt-to-GDP ratio. Indeed, as I will show using their (correct) data, a bivariate regression suggests that a 10 percent point increase is associated with a 0.18 percent point lower growth rate.

However, there are also well-known problems with interpreting this negative relationship as representing causality flowing from debt to growth. One reason is just algebraic. The independent variable debt-to-GDP is a ratio that has a numerator (debt) and denominator (GDP): any fall in GDP will mechanically boost the ratio. Even if GDP growth doesn't become negative, continuous growth in debt coupled with a GDP growth slowdown will also lead to a rise in the debt-to-GDP ratio.

*University of Massachusetts Amherst, and IZA. Contact via email: adube@econs.umass.edu.

There is also a less mechanical reason around reverse causality. A recession leads to increased spending through automatic stabilizers such as unemployment insurance, as well as discretionary stimulus spending. And governments usually finance these using greater borrowing, as undergraduate macro-economics textbooks tell us governments should do. This means negative economic shocks lead to increased debt. Casual observations suggest that this is what happened in the U.S. during the past recession, and in Japan during the 1990s.[1]

In a simple bivariate setting, the correlation between growth and debt-to-GDP ratio will reflect a mixture of both of these relationships: lower growth spurred by high debt (hypothesis A) and higher debt spurred by lower growth (hypothesis B). As long as debt does rise *at all* during recessions (and the coefficients of both the equations are consistent with stability conditions), the bivariate regression coefficient of debt on growth will *overstate* the causal effect posited in hypothesis A. Panizza and Presbitero (2013) shows the the simple algebra behind the bias of the bivariate regression coefficient. The intuition behind the bias is straightforward and shown in Figure 1. If $D = f(g) + u$ and $g = h(D) + v$ are the two loci representing the long run relationships between debt, D, and growth g, then stability requires that $f' < h'$ near the equilibrium. Both are assumed to be weakly downward sloping, as in Figure 1. As long as some of the variation in observed data comes from exogenous movements in growth (i.e., innovations in v), then the estimated regression coefficient will be "too steep" as it will be somewhere between f' and h' whereas the debt to growth relationship has a true slope of h'.

Figure 1: Debt, Growth and the Bias in the Regression Estimate

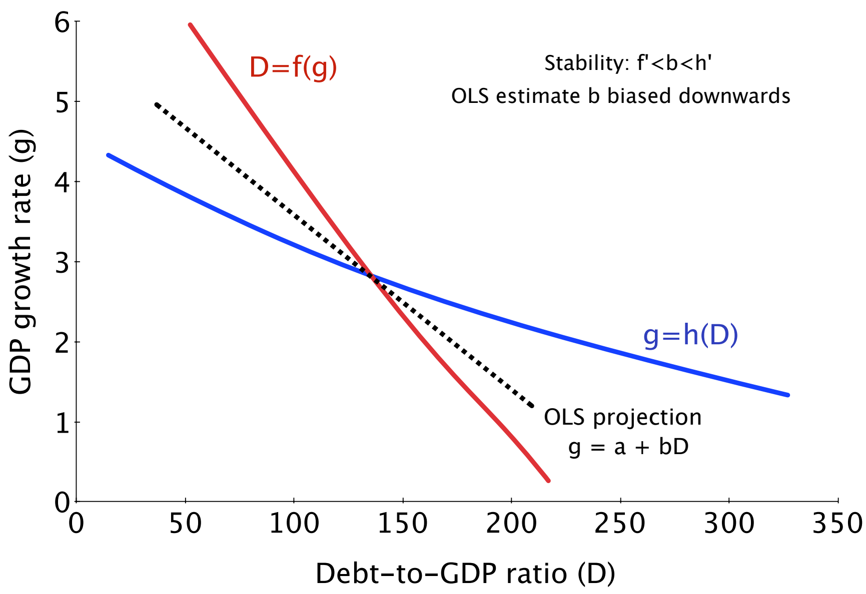

The extent of bias is, of course, an empirical question. In this short note, I will present evidence using the Reinhart Rogoff 2010 data between 1946 and 2009 that demonstrates that most of the simple correlation between debt-to-GDP and growth is not due to a causal effect of debt on growth.

First, I will show that the correlation between current debt-to-GDP and growth during the *past* 5 years is much stronger than the correlation with the *future 5* year growth. Indeed, the relationship between current

[1] For a case study of the U.S. data, see Bivens and Irons (2011).

debt ratio and 5-year forward growth is much weaker than the correlation between current debt ratio and current growth. These findings raise serious questions about the direction of causality.

Next, I will use some straightforward tools–separately and together–to account for some of the endogeneity of debt in regression:

(1) Use of the 5-year forward growth instead of just contemporaneous growth as the outcome. By using pre-determined debt-level, and considering growth over a longer period of time, it avoids some of the most obvious source of reverse causality. Other researchers have also used this outcome (e.g., Cecchetti Mohanty and Zampolli 2012).

(2) Use of lagged debt-to-GDP as an instrument for current debt-to-GDP. This is another standard tool used to address some of the reverse causality problem.

(3) Use lagged GDP growth to account for time-varying heterogeneity. I find that these simple controls for endogeneity suggests around 70% of the relationship between contemporaneous debt ratio and GDP growth is spurious. In all of the specifications with minimal control for endogeneity, the bivariate regression estimate is reduced by at least 50%.

Finally, I use a partial-linear model to graphically plot the relationship between 5 year forward growth rates and contemporaneous debt, after controlling for lagged growth rates. I find a very weak negative relationship between debt and growth, a no evidence for thresholds at higher values of the debt ratio.

As this is a short note, I will not review the sizable literature on the topic, which is summarized in a recent review (Panizza and Presbitero 2013). Where relevant, I will compare my results with specific recent papers that argue that there is a clear negative impact of debt on growth, and that find some evidence of threshold effects for higher levels of debt. Three key papers along these lines include Kumar and Woo (2010), Cecchetti Mohanty and Zampolli (2012), and Chercherita-Westphal and Rother (2012).

Since papers in this literature have certainly considered the causality issue, it is useful to clarify the contribution of this note. I highlight four such contributions:

1) I use the data from Reinhart and Rogoff which has received a lot of recent attention, and that this data covers a longer time period (1946-2009, with most countries available from 1956) than some of the other recent papers.

2) I assess nonlinearities using relatively simple visual tools: discrete bins like those used in RR, bivariate nonparametric regressions, and partial linear regressions controlling for lagged GDP growth. This reduces the black-box nature of threshold searches sometimes used in the literature.

3) I assess the direction of causality visually first using a simple bivariate plots: showing the relationship between the debt ratio and forward versus past growth. This clearly shows the extent of "the problem."

4) I show how simple alternative controls for some of the endogeneity (use of 5-year forward GDP growth rate; lagged debt ratio as an instrument; use of lagged GDP growth as a control) affect the estimates. In particular, few of the papers mentioned above have used lagged GDP *growth* as a control, as opposed to lagged GDP *level.* Yet, growth slowdown leading to increased debt is what we would expect the form of reverse causality to likely take. I show that the inclusion of lagged growth as a control is quite important.

Finally, the evidence presented in this note does not suggest that debt has no impact on growth. What it does show, however, is that most of the simple bivariate relationship between the two reflects causality in the opposite direction or other omitted variables.

2 Motivating Evidence: Correlation between Current Debt and Past versus Future Growth

As a starting point, we can contrast the predictions made by the two hypotheses: (A) high debt causes low growth, versus (B) low growth causes high debt. If hypothesis A is correct, and hypothesis B is not correct, then we should expect a correlation between debt-to-GDP at time t and lower *future* growth, as higher debt burden leads to a lower GDP growth level. To the extent that GDP growth is serially correlated, we may still see *some* correlation between current GDP and past growth, but this should be weaker than the correlation between current GDP and future growth.

In contrast, to the extent that reverse causality is a problem as stated in hypothesis B, the initial shock is in GDP. This is followed by an increase in the debt-to-GDP ratio as GDP slows down and perhaps debt levels rise from an increase in countercyclical spending. If this story is true, we would expect that higher debt-to-GDP observations are associated with low growth in *previous* years. The stronger is the reverse causal channel, the greater is the correlation we expect between debt ratio and past growth as compared to future growth.

The next figure shows non-parametrically the relationship between contemporaneous debt-to-GDP for country i in year t, $d_{i,t} = \frac{D_{i,t}}{Y_{i,t}}$ with:

(a) contemporaneous real per-capita GDP growth:

$g_{i,t} = \frac{Y_{i,t+1} - Y_{i,t}}{Y_{i,t}}$

(b) 5-year past average of real per-capita GDP growth:

$\overline{g}_{i,t-5,t-1} = \frac{1}{5} \sum_{\tau=1}^{5} (g_{i,t-\tau}) = \frac{1}{5} \sum_{\tau=1}^{5} \left(\frac{Y_{t-\tau+1} - Y_{t-\tau}}{Y_{t-\tau}} \right)$

(c) 5-year forward average of real per-capita GDP growth:

$\overline{g}_{i,t+1,t+5} = \frac{1}{5} \sum_{\tau=1}^{5} (g_{t+\tau}) = \frac{1}{5} \sum_{\tau=1}^{5} \left(\frac{Y_{t+\tau+1} - Y_{t+\tau}}{Y_{t+\tau}} \right)$

It is important to note that a 5-year forward average growth rate has advantages over the contemporaneous growth rate because a longer growth window, and a pre-determined debt ratio tend to mitigate the worst form of endogeneity induced by short run growth fluctuations. This is why many papers in the literature (e.g., Cecchetti et al.) use a 5-year forward average growth rate as the outcome. The nonparametric regressions are estimated using the *lowess* command in Stata, with the default tuning parameter of 0.8. The 95% confidence bands are constructed using bootstrapping, clustering on countries to account for arbitrary serial correlation (including those induced by the use of overlapping five year frames).

Figure 2: Relationship between Current Debt Ratio and GDP Growth: Past, Present and Future

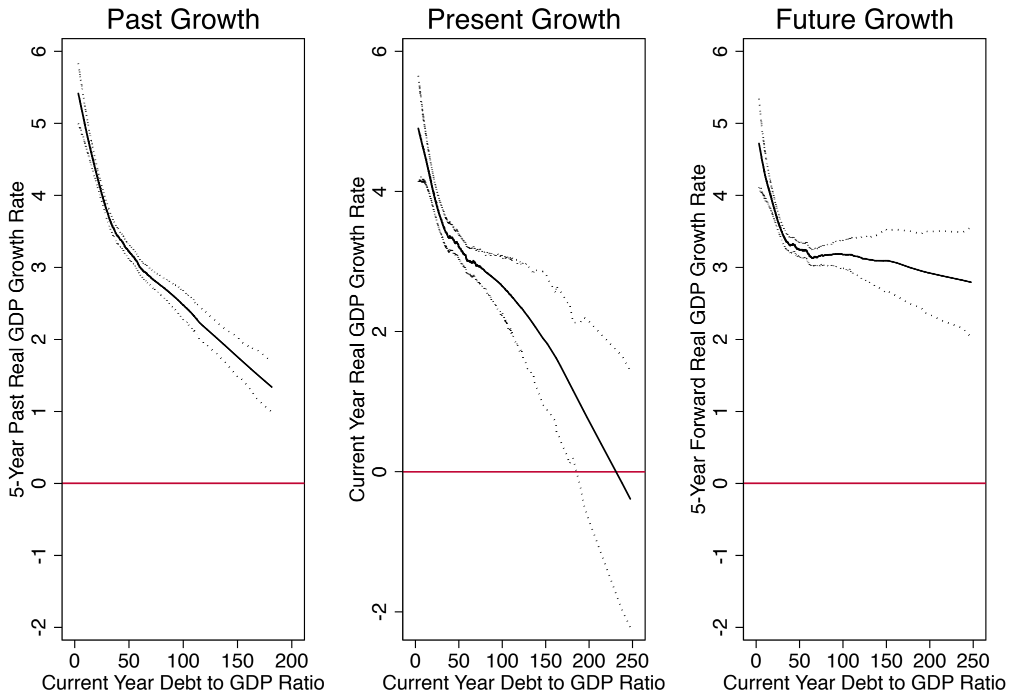

Notes. The figure reports results from lowess regressions of growth on current debt-to-GDP ratio. In the left panel, the outcome is , the outcome is the 5-year past average growth rate: average of GDP growth rates between years t-5 and t. In the middle panel, the outcome is current growth rate (between year t). In the right panel, the outcome is the 5-year forward average average growth rate (average growth rate between year t+1 and t+5). 95% confidence bands are calculated using cluster bootstrapping, and shown as dotted lines.

As the middle panel of Figure 2 indicates, there is a clear negative correlation between debt and contemporaneous growth, but as HAP shows, there is no indication of a threshold type effect at higher levels of debt. The relationship is actually close to linear, although it is less precise for higher debt ratios. When we turn to the correlation between the debt ratio and growth over the past 5 years, we find a similar negative relationship. But when we consider the 5-year forward average growth rates, the relationship is much weaker; for debt ratios greater than 50, there is barely any relationship at all between debt and growth.

Together, the evidence suggests that the contemporaneous negative relationship (middle panel of Figure 1) is driven much more by a higher debt today from weaker past growth rather than weaker future growth from higher debt today. If there were no reverse causality, it is difficult to explain why the current debt ratio is more strongly correlated with growth over the past 5 years rather than the next 5 years from now. Moreover, if there is reverse causality, then the bivariate regression is almost certainly an over-estimate of the magnitude of any true effect of increased debt on growth.

Figure 3 plots the growth rates by bins used in RR and in HAP; besides current GDP growth rates, it also shows 5-year past and forward average growth rates. Strikingly, most of decline across the bins when considering current growth rates disappear when considering 5-year forward growth rates. In contrast, past growth rates show a sharper fall-off when we consider higher debt bins. Together, this provides strong evidence that even the more moderate negative association reported in HAP in large part an artifact of using current instead of the more reliable forward growth rates–the latter having become a more standard measure in the literature.

Figure 3: GDP Growth by Debt Ratio Bins: Past, Present and Future

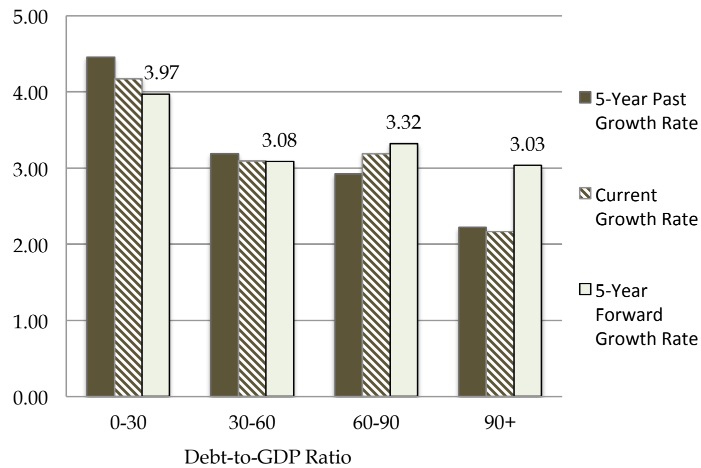

Notes. The figure reports GDP growth rates by bins of current debt-to-GDP ratio: under 30%, 30-59.9%, 60%-89.9%, 90% and above. The three outcomes are 5-year past growth rate (average between years t-5 and t); current growth rate (between years t and t+1). In the right panel, the outcome is the 5-year forward average growth rate (average growth rate between year t+1 and t+6). 95% confidence bands are calculated using cluster bootstrapping, and shown as dotted lines.

3 Regression Estimates: Debt, Growth and Endogeneity

A simple bivariate linear regression using the RR dataset produces the following relationship:

$$\begin{array}{rcccl} g_{it} & = & \underset{(0.302)}{4.270} & - & \underset{(0.005)}{0.018} \times d_{it} + e_{it} \end{array} \tag{1}$$

The standard errors in parentheses are clustered by country. [2] Since there is no visible threshold type effects in the bivariate nonparametric plot, this linear predictor closely approximates the conditional expectation of growth given debt ratio. It suggests that a 10 percentage point increase in the debt ratio, d_{jt}, is associated with a 0.18 point lower growth rate.

As we saw from the evidence in the last section, a part of this likely represents an endogeneity bias. I am now going to describe 3 ways to mitigate this bias. First, and most simply, we can replace the dependent variable g_{jt} with a five year forward average growth rate, $\overline{g}_{i,t+1,t+5}$.

$$\overline{g}_{i,t+1,t+5} = \alpha_2 + \beta_2 d_{it} + e_{it} \tag{2}$$

[2] There are only 20 countries, so this poses a problem for cluster-robust standard errors due to the relatively small number of clusters. One solution is to use the Wild cluster bootstrap, which tends to produce more accurate test size even a small number of clusters, as shown in Cameron Gelbach and Miller (2010). When I used the Wild bootstrap, the 5% cutoffs for the t-statistics were substantially larger, at 3.7 instead of the canonical 1.96, so the estimate above is significant only at the 10% level. This means that we should be skeptical of the statistical precision even when using clustered standard errors in this dataset. To keep the exposition simple, however, I report cluster-robust estimates except where specified otherwise.

Moving to a five-year forward average reduces the impact of a contemporaneous transitory shock to Y_{jt} from inducing a negative association between g_{jt} and d_{jt}. For this reason, the five-year forward rate is used as the outcome in many papers in the literature, including Cecchetti et al. (2011), to remove the contemporaneous correlation between $d_t$and g_{it} through temporary shocks in output. As shown in Table 1, column 1, the coefficient from the bivariate regression when using the 5-year forward growth rate falls by half in magnitude from -0.018 to -0.009. This is as we would expect from the visual evidence presented in Figure 1, which are just nonparametric versions of the regressions in column 1.

Second, we can address the reverse causality issue by using lagged $d_{i,t-1}$ as an instrument for $d_{i,t}$. This approach is fairly similar in spirit to using the five-year forward growth rate. Again, this tends to put some time distance between the debt and growth variables, reducing the impact from an increase in debt at time t from a negative GDP shock. Comparing columns 1 and 2, we find that instrumenting for $d_{i,t}$ in this fashion has a substantial effect when using current growth (panel A), but very little additional impact when using 5-year forward average growth rate: this is to be expected given the similarity in the variation used under the two approaches.

Table 1: Regressing Growth on Debt Ratio: Simple Controls for Endogeneity

	(1)	(2)	(3)	(4)
	Current GDP Growth Rate			
Debt-to-GDP	-0.018***	-0.012**	-0.009**	-0.005
	(0.005)	(0.006)	(0.004)	(0.004)
N	1175	1151	1151	1151
	5-year Forward Growth Rate			
Debt-to-GDP	-0.009*	-0.008	-0.006	-0.005
	(0.005)	(0.005)	(0.004)	(0.005)
N	1059	1036	1036	1036
Instrumented DEBT:		Y		Y
Lagged Growth Control:			Y	Y

Notes. The table reports results from regressions of growth on current debt-to-GDP ratio ("DEBT"). In Panel A, the outcome is the current year growth: GDP growth between year t and t+1. In Panel B, the outcome is the 5-year forward average growth rate (average growth rate between year t+1 and t+6). Column 1 reports results from simple bivariate OLS regression. Column 2 reports results from an IV specification where current DEBT is instrumented using 1-year lagged DEBT. Column 3 includes as a control growth in GDP between time t-2 and t-1. Column 4 reports the results using both instrumented DEBT and lagged GDP growth as a control. Standard errors clustered by country reported in parentheses. Significance levels indicated as follows: * p<0.1, ** p<0.05, *** p<0.01.

Finally, while using predetermined debt ratios accounts for some of the endogeneity, it is unlikely to be sufficient. One way of accounting for additional time-varying heterogeneity is by including the lagged outcome as a control. If a growth slowdown leads to an increase in the debt ratio, we can account for this using a lagged GDP growth as a control:

$$\overline{g}_{i,t+1,t+5} = \alpha_3 + \beta_3 d_{it} + \gamma_3 g_{i,t-1} + e_{it} \tag{3}$$

Now the identifying assumption is that conditional on the lagged GDP growth rate $g_{i,t-1}$, the error term $e_{2i,t}$ is uncorrelated with debt at time t, which is a weaker than the identifying assumption in the bivariate case. Interestingly, most of the literature does not include lagged growth rate as a control; it is more common to include lagged GDP level as a control. While that specification from Mankiw Romer and Weil (1992) makes sense for understanding growth behavior and catch up, it will not account for the type of reverse causality that is likely to affect the growth-debt relationship. We want to compare countries that were similar in their past growth trajectories but happen to differ in their debt ratios, and see how their future growth trajectories differ. Some exceptions to using lagged growth rates in this literature includes Baum Chercherita-Westphal and Rother (2012) and IMF analysis of growth and fiscal consolidation (IMF 2010, chapter 3).

Comparing columns 1 and 3, we find that inclusion of lagged growth also tends to reduce the coefficient in magnitude. With 5-year forward growth rates, the coefficient falls in magnitude to -0.006. Combining the instrumenting of debt ratio with inclusion of lagged growth as a control (column 4) does not have much impact once we are looking at 5-year forward growth rates; the coefficient in this case is -0.005 as shown in Column B of Panel B. Once we have used both instrumenting and lagged dependent variable, it does not matter much whether we use current year versus 5-year forward growth rates. This, too, is unsurprising as the lagged debt instrument is really a substitute for using forward growth rates.

My preferred estimate is column 3, Panel B, because it is the most transparent: it looks at the relationship between current debt ratio and 5-year forward average growth rate, and includes a single control variable–one year lagged growth. The estimate is -0.006, and is similar to the estimates using lagged debt ratio as an instrument. This suggests that around 70% of the bivariate relationship may be due to endogeneity that can be accounted for using very simple tools. Taking any of the estimates except Panel A, Column 1, we find that accounting for endogeneity using these simple methods reduces the effects by anywhere between 50 and 70%.

I did not report results using more complicated specifications that combine lagged dependent variable with country and time fixed effects, such as the Arellano-Bover/Blundell-Bond system GMM dynamic panel estimator. This because I wanted to avoid "black box" type methods here. However, I note here that dynamic panel approaches produce estimates between -0.005 and -0.009 depending on the number of lags used for the dependent variable (estimates available upon request).

4 Thresholds and Nonlinearity Once More

As a final exercise, I assess whether there is a nonlinear relationship or threshold effects visible in the relationship between 5-year growth rates and current debt ratio. I use a partial linear model, where $f(.)$ is an unknown nonlinear function:

$$\overline{g}_{i,t+1,t+5} = \alpha_4 + f(d_{i,t}) + \gamma_4 g_{i,t-1} + e_{4,i,t} \tag{4}$$

Estimation of equation (x) proceeds by first sorting by d_{it} and differencing observations, and then estimating $\hat{\gamma_4}$ using OLS. Then the residuals $\left(\overline{g}_{i,t+1,t+5} - \hat{\gamma}_4 g_{i,t-1}\right)$ are regressed non-parametrically on $d_{i,t}$ using lowess.

Figure 4: Relationship between Current Debt Ratio and 5-Year Forward GDP Growth, Controlling for Lagged Growth

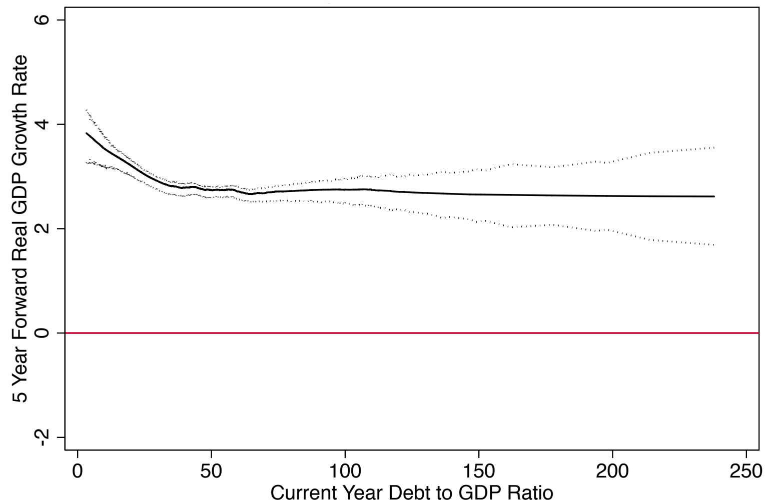

Figure 4 shows very little evidence of nonlinearities or threshold effects at higher values of the debt ratio. While there is some negative relationship between debt and growth at low levels of debt (<30), the relationship is mostly flat after 30. Once we account for past growth, there is little evidence of a remaining correlation between debt and 5-year forward average growth rates among advanced industrialized countries in the post WWII era, at least in the Reinhart Rogoff dataset.

5 Discussion

Considering the 20 countries with (unbalanced) data over the 1946-2009 period used in RR and HAP, I find that the negative association between current growth and current debt-to-GDP ratio is in large part likely driven by reverse causality. The simplest form of correction, use of 5-year forward growth rates as is used frequently in the recent literature suggests little in the way of a negative association for debt ratios beyond 30 or so. In a linear regression, moving from current to 5-year forward growth halves the magnitude of the coefficient from -0.018 to -0.009. Instrumenting the debt ratio with it's lag, while continuing to use current year GDP growth rate as the dependent variable has a similar impact. Accounting for lagged growth as a way to control for business cycle dynamics renders the estimated coefficient even smaller in magnitude (-0.005). There is no visual evidence for a turning point or structural break in the relationship between debt and growth at higher debt ratios, although the relationship becomes more imprecise at very high levels of debt. Once I account for lagged growth, the relationship between the debt ratio and the 5-year forward growth rate is essentially flat for a debt ratio of 30 or more.

There are some interesting differences in these findings when compared to some of the other prominent papers such as Cecchetti et al., and Chercherita-Westphal and Rother. Generally, both of those papers found somewhat stronger negative effects even when using 5-year forward growth rates. For example, Cecchetti et al. find a coefficient on 5-year forward growth of -0.016 for public debt using 1980-2007 sample of 18 advanced economies. Chercherita-Westphal and Rother find evidence for a turning point in the debt-growth relationship using a quadratic specification. However, they do not show whether the average partial effect in the relevant range of debt is sizeable, or even negative. Finally, Kumar and Woo find an estimate of around -0.015 for advanced economies between 1970 and 2009.

None of these 3 papers report simple non- or semi-parametric relationships of the sort presented here; and none of these control for lagged growth rate (as opposed to lagged GDP). At the same time, the samples used in these papers are also different. Assessing the source of differences between the association between debt and growth–samples and specifications–as reported here and those from the other key papers is beyond the scope of this note. However, the lack of any visible nonlinearities in the RR data when controlling for lagged growth–suggests that it would be useful to apply simple visual tools to assess nonlinearities in these other dataset as well. Future work would do well to clearly decompose any differences between results by differences in sample years, sample countries, outcome measures and econometric specifications.

References

[1] Baum, A., C. Checherita-Westphal, and P. Rother 2012. "Debt and growth: new evidence for the Euro area." *European Central Bank Working Paper Series* 1450.

[2] Cecchetti, S., M. Mohanty, and F. Zampolli 2012. "Achieving Growth Amid Fiscal Imbalances: The real effects of debt" in *Achieving maximum long-run growth - A symposium sponsored by The Federal Reserve Bank of Kansas City*. Federal Reserve Bank of Kansas City.

[3] Checherita-Westphal, C., and P. Rother 2012. "The Impact of High Government Debt on Economic Growth and its Channels: An Empirical Investigation for the Euro Area." *European Economic Review* 56(7):1392–1405.

[4] DeLong, B. J., and L. H. Summers (2012): "Fiscal Policy in a Depressed Economy," Brookings Papers on Economic Activity, Spring.

[5] Herndon, T., M. Ash, and R. Pollin 2013. "Does High Public Debt Consistently Stifle Economic Growth? A Critique of Reinhart and Rogoff." *PERI Working Paper Series* 322.

[6] International Monetary Fund 2010. "Will it Hurt? Macroeconomic Effects of Fiscal Consolidation," in *World Ecoomic Outlook*: *Recovery, Risk and Rebalancing*. International Monetary Fund.

[7] Irons, J. and J. Bivens 2010. "Government Debt and Economic Growth: Overreaching Claims of Debt 'Threshold' Suffer from Theoretical and Empirical Flaws." *Economic Policy Institute Briefing Paper* 271.

[8] Kumar, M. S., and J. Woo 2010. "Public Debt and Growth," *IMF Working Papers* 10/174.

[9] Mankiw, N. G., D. Romer, and D. N. Weil 1992. "A Contribution to the Empirics of Economic Growth." *The Quarterly Journal of Economics* 107(2): 407–37.

[10] Panizza, U. and A. F. Presbitero 2013. "Public Debt and Economic Growth in Advanced Economies: A Survey." *Money and Finance Research Group, Working Paper No.* 78.

[11] Reinhart, C. M., and K. S. Rogoff 2010. "Growth in a Time of Debt." *American Economic Review* 100(2): 573–78.