The Age Distribution and Business Cycle Volatility: International Evidence

Steven Lugauer[a]* and Michael Redmond[a]

[a]University of Notre Dame, Notre Dame, IN 46556, USA

November 2011

Abstract

We estimate the age distribution's effect on business cycle fluctuations across a large number of countries. A 10 percentage point increase in the middle-aged share of the population decreases output volatility by 15 percent for the average country.

Keywords: Business Cycles, Demographics

JEL Classification: E32, J10

* Corresponding Author. University of Notre Dame, Department of Economics, 719 Flanner Hall, Notre Dame, IN 46556, USA, slugauer@nd.edu, (574) 631-1432.

1. Introduction

Over the past several decades, the age distribution for many countries has changed dramatically. However, the timing and extent of these changes have varied. The US, for example, had a large baby-boom after World War II; whereas, fertility in Japan has been declining for almost 60 years. Using standard panel data methods, we exploit the variation in the demographic changes across countries and throughout time to show that the age distribution has had a significant effect on the magnitude of gross domestic product (GDP) business cycle fluctuations.

Jaimovich and Siu (2009) [from now on JS] first connected the age distribution to the magnitude of the business cycle using data from seven developed countries. JS justify their reduced-form model by showing that the cyclical volatility of hours worked has been highest for young workers, low for the middle aged, and high again for people near retirement. The U-shape pattern for employment volatility also motivates our research question.[1] Does the age distribution affect the magnitude of the business cycle in a larger set of countries?

We estimate a model similar to JS; however, we use more countries (51) and a longer time span (1957-2000). Our findings agree with JS. When the population has relatively more prime age (30-59) workers, business cycle fluctuations tend to be low. Moreover, the age distribution effect is quantitatively large.

2. Model and Econometric Strategy

We estimate Equation (1) with Ordinary Least Squares (OLS).

$$vol_{it} = \alpha_i + \beta_t + \delta growth_{it} + \gamma share_{it} + \varepsilon_{it} \tag{1}$$

Variable *vol* equals the standard deviation of a centered 9-year window of de-trended logged annual GDP for each country and year; *vol* is our measure of cyclical output volatility. We apply the Hodrick-Prescott (HP) filter with smoothing parameter 6.25 to de-trend each country's logged GDP series. Variable *share* equals the fraction of the working age (15-64) population aged 15-29 plus those aged 60-64; *share* is our measure of the age distribution and captures the percentage of a country's working age population in their 'high volatility' years. Both *vol* and *share* have been defined as in JS to facilitate comparison. Parameter γ measures the age distribution's effect on *vol*; γ is our parameter of interest.

Variable *growth* equals the annual GDP growth rate for each country and year; *growth* accounts for the possible relationship between growth and economic stability. Ramey and Ramey (1995), for example, find a negative correlation between growth and volatility in a large panel of countries. Since GDP growth may in turn depend on the age structure, we include *growth* to eliminate the potential for spurious correlation between *share* and *vol*. Excluding *growth* does not appreciably alter our findings, although we do find evidence of the negative relationship between growth and volatility. See Section 4, Table 1, and Footnote 3 for more on the *growth* variable.

The variable α represents country fixed effect dummies to control for heterogeneity in GDP volatility levels across countries. Similarly, β represents year dummy variables to control for time varying trends common to all countries. The ε term captures other country and year specific sources of variation in *vol*.

Identification of the age distribution's effect on the magnitude of business cycle fluctuations as captured by γ comes from changes in *share* over time not common to all countries. We assume that the temporal and geographic variation in the age distribution has been caused by factors (e.g. fertility decisions made decades earlier) unrelated to the other, omitted, determinants of *vol*. Again, the specification parallels the model in JS.

3. Data

We use the United Nations World Population Prospects (2008 revision) country-by-country annual age distribution estimates from 1950 onward to compute *share*. We use annual GDP data from the Penn World Table (2009, version 6.3) to calculate *vol* and *growth*. Countries with a population less than one million in 2005 were dropped from the sample, leaving 51 countries with GDP reported from 1950 to 2007. The 9-year rolling window used to calculate *vol* eliminates eight years of observations. We drop an additional three years from the beginning and end of the sample because HP-filtered time series can have excess volatility near the endpoints. The final panel contains 51 countries and years 1957-2000, for 2,244 total observations.[2] The data is fairly representative, containing a mix of developing and developed countries from around the world.

4. Findings

Table 1 presents OLS estimates of γ based on Equation (1). The γ estimate of 3.92 reported in Column 1 captures our main result. Since JS do not include *growth* as a control variable in their model, column 2 excludes *growth*.[3] In both specifications, *share* has a statistically significant effect on *vol* at the 10 percent level or better, based on Newey-West (lag 2) standard errors (the same employed by JS). More importantly, the age distribution effect is economically large. The γ estimate of 3.92 implies that a 10 percentage point increase in *share* increases *vol* by 0.39 percentage points. Across the sample, *vol* averages 2.41 percent. Thus, a 10 percentage point increase in *share* generates more than a 15% increase in GDP volatility for the average country.

Table 1: Estimates of *share*'s effect (γ) on *vol*, 1957-2000

	Main Result (1)	No *growth* (2)	JS (3)	JS No *growth* (4)
γ	3.92	3.84	3.65	3.76
	(1.95) **	(1.99) *	(1.60) **	(1.60) **
N	2,244	2,244	264	264

Table 1 reports OLS estimates based on Equation (1) with Newey-West (lag 2) standard errors in parentheses. All regressions include a full set of country and year fixed effect dummies. Stars indicate statistical significance at the * 10% and ** 5% level.

Our results, based on different data with more countries and years, strongly agree with those reported in JS (JS estimate γ to be about 4.0). In columns 3 and 4, we restrict the sample to only countries used by JS. Our data set does not include Germany, giving us six countries compared to seven in JS. Though slightly smaller than our main result, the size of the age distribution effect remains quantitatively large and statistically significant in the restricted sample.

Lugauer (2011) finds an even larger role for the age distribution (γ estimate of 5.8) using state-by-state variation in US demographics and GDP volatility.[4] Lugauer estimates Equation (1) with a two stage approach, instrumenting *share* with lagged birth rates. Lugauer instruments with lagged birth rates primarily because differential migration across states by age could bias the OLS estimates. In fact, Lugauer finds downward bias in the OLS estimates. This bias could be one reason our OLS γ estimate is smaller than the two stage estimate in Lugauer (2011). The same migration concern applies to our cross-country analysis, and our main result should be viewed as a lower bound on the age distribution's true effect.

A second possible reason for the larger γ estimate in Lugauer (2011) is that Lugauer (2011) uses only US data, and the mechanism connecting the age distribution to business cycle volatility might be stronger in the US relative to other countries. Lugauer (2012) develops a theoretical model in which the age distribution affects aggregate output volatility through the labor market (Jaimovich, Pruitt, and Siu (2011) present an alternative theory). Accordingly, labor market differences across countries might account for the larger γ estimate in Lugauer (2011) versus the estimates (above and in JS) based on international data.

5. Conclusion

The age distribution, as measured by the 'high volatility' share of the population, has had a large effect on GDP business cycle volatility across a panel of 51 countries from 1957-2000 according to our OLS estimates. This finding agrees with the results based both on developed countries as in Jaimovich and Siu (2009) and on US states as in Lugauer (2011). Taken together these papers provide strong evidence that the age distribution affects the business cycle.

A few recent papers have also investigated the importance of demographics for macroeconomic analysis. Feyrer (2007) considers whether the age distribution affects US productivity; Shimer (2001) measures the age distribution's effect on the US unemployment rate; and Curtis, Lugauer, and Mark (2011) examine the age distribution's effect on the household saving rate in China. These related topics could be explored further using the cross-country demographic variation exploited in this paper.

Notes

1. Clark and Summers (1981) first documented the differences in cyclical employment volatility by age. The key contribution of JS is to show that the population's age distribution affects the business cycle.
2. The countries are Argentina, Australia, Austria, Belgium, Bolivia, Brazil, Canada, Colombia, Costa Rica, Democratic Republic of Congo, Denmark, Egypt, El Salvador, Ethiopia, Finland, France, Guatemala, Honduras, India, Ireland, Israel, Italy, Japan, Kenya, Mauritius, Mexico, Morocco, Netherlands, New Zealand, Nicaragua, Nigeria, Norway, Pakistan, Panama, Peru, Philippines, Portugal, Puerto Rico, South Africa, Spain, Sri Lanka, Sweden, Switzerland, Thailand, Trinidad and Tobago, Turkey, Uganda, United Kingdom, United States, Uruguay, and Venezuela.
3. The estimate of parameter δ in Equation (1) is negative and statistically different from zero at the 5% level, which is consistent with the findings reported in Ramey and Ramey (1995) and elsewhere. We also tried using 5 and 10 year growth rates, obtaining similar δ estimates and slightly larger γ estimates.
4. Lugauer (2011) uses different variable definitions. Variable *share* in Lugauer (2011) equals the fraction of the population under 35 rather than the 'volatility share' used here. Note JS use labor force 'volatility share'.

References

Curtis, Chadwick C., Steven Lugauer, and Nelson C. Mark (2011): "Demographic Patterns and Household Saving in China," *mimeo*, University of Notre Dame.

Clark, Kim B., and Lawrence H. Summers (1981): "Demographic Differences in Cyclical Employment Variation," *Journal of Human Resources*, 16, 61-79.

Feyrer, James (2007): "Demographics and Productivity," *Review of Economics and Statistics*, 89(1), 100-109.

Jaimovich, Nir, Seth Pruitt, and Henry E. Siu (2011): "The Demand for Youth: Implications for the Hours Volatility Puzzle," *mimeo*, Duke University.

Jaimovich, Nir, and Henry E. Siu (2009): "The Young, the Old, and the Restless: Demographics and Business Cycle Volatility," *American Economic Review*, 99(3), 804-826.

Lugauer, Steven (2011): "Estimating the Effect of the Age Distribution on Cyclical Output Volatility Across the United States," *Review of Economics and Statistics*, In Press.

Lugauer, Steven (2012): "Demographic Change and the Great Moderation in an Overlapping Generations Model with Matching Frictions," *Macroeconomic Dynamics*, In Press.

Ramey, Garey, and Valerie A. Ramey (1995): "Cross-Country Evidence on the Link Between Volatility and Growth," *American Economic Review*, 85(5), 1138-1151.

Shimer, Robert (2001): "The Impact of Young Workers on the Aggregate Labor Market," *Quarterly Journal of Economics*, 116(3), 969-1007.