

Cross Sectional Facts for Macroeconomists^{*}

Dirk Krueger

University of Pennsylvania, CEPR and NBER
dkrueger@ssc.upenn.edu

Fabrizio Perri

University of Minnesota, Federal Reserve Bank of Minneapolis, CEPR and NBER
fperri@umn.edu

Luigi Pistaferri

Stanford University, CEPR and NBER
pista@stanford.edu

Giovanni L. Violante

New York University, CEPR and NBER
glv2@nyu.edu

Abstract

This paper provides an introduction to the special issue of the Review of Economic Dynamics on “Cross Sectional Facts for Macroeconomists”. The issue documents, for nine countries, the level and the evolution, over time and over the life cycle, of several dimensions of economic inequality, including wages, labor earnings, income, consumption, and wealth. After describing the motivation and the common methodology underlying this empirical project, we discuss selected results, with an emphasis on cross-country comparisons. Most, but not all, countries experienced substantial increases in wages and earnings inequality, over the last three decades. While the trend in the skill premium differed widely across countries, the experience premium rose and the gender premium fell virtually everywhere. At a higher frequency, earnings inequality appears to be strongly counter-cyclical. In all countries, government redistribution through taxes and transfers reduced the level, the trend and the cyclical fluctuations in income inequality. The rise in income inequality was stronger at the bottom of the distribution. Consumption inequality increased less than disposable income inequality, and tracked the latter much more closely at the top than at the bottom of the distribution. Measuring the age-profile of inequality is challenging because of the interplay of time and cohort effects.

JEL CLASSIFICATION CODES : D31, D91, E21

Author Keywords: wages; income; consumption; wealth; long-run trends in inequality; life-cycle inequality; inequality over the business cycle; government redistribution; estimation of earnings dynamics.

^{*}We thank the participants of the 2007 Philadelphia conference on “Heterogeneity in Macroeconomics” for useful comments, Jacek Rothert, Ctirad Slavik, and Chris Tonetti for excellent research assistance, and the NSF and PIER for financial support

1 Introduction

Modern macroeconomics has expanded its focus from the dynamics of aggregate variables such as GDP, consumption and wealth to the dynamics of entire equilibrium distributions across households of these variables. Theoretical models pioneered by Laitner (1979), Bewley (1986), Imrohoroglu (1989), Huggett (1993), Aiyagari (1994) and Rios-Rull (1995), among others, allow for the equilibrium determination of the joint distribution of hours worked, income, consumption and wealth. In order to confidently use these models for quantitative policy analysis it is necessary to establish that their aggregate and distributional implications are consistent with the salient features of the data. To be concrete, consider a government contemplating a reform of the income tax system. In economic models with a representative household, the distributional effects of such a policy reform cannot be analyzed, and its aggregate consequences may be misleading, since the response of households is, by construction, constrained to be identical. Heterogeneous agent macroeconomic models constitute the natural laboratory for the analysis of such a question. Suppose the model is successful in capturing the behavioral response of all households to a change in the tax code, and suppose household response depends on household characteristics such as income, wealth etc.. Unless the model (prior to the tax reform) generates an empirically realistic distribution over household characteristics, aggregating the household level behavioral responses will result in a flawed answer to the question of what the aggregate consequences of the tax reform are. Therefore, restricting heterogeneous agent macro models so that the equilibrium distributions of hours worked, income, consumption and wealth line up well with their empirical counterparts is crucial for a convincing policy analysis.

Dating back to the early work by Burns and Mitchell (1946), a wealth of research has studied how to measure aggregate economic fluctuations and has described key stylized facts for many countries. Similarly, the Penn World Tables (Heston et al., 2009) provide long run aggregate data for a large set of countries that have motivated and empirically guided the construction of models of economic growth and development. However, a systematic and comprehensive empirical analysis of the cross sectional *distributions* for wages, hours worked, income, consumption and wealth that provides stylized and comparable inequality facts for several countries is still rare. This issue of *The Review of Economic Dynamics* aims to fill this gap.¹

¹The Luxembourg Income Study (LIS) and the Luxembourg Wealth Study (LWS) provide information about the income and wealth distributions for a large set of countries. See Gottschalk and Smeeding (1997) for a cross-country comparison of income inequality using the LIS. The Cross-National Equivalent File (CNEF) at Cornell University provides household level data (mainly for income) for six countries (US, Germany, Great Britain, Australia,

Table 1: Countries and Contributors

Country	Researchers	Household Level Data Used
U.S.	Heathcote, Perri & Violante	CEX, CPS, PSID, SCF
Canada	Brzozowski, Gervais, Klein & Suzuki	FAMEX, SCF, ADSCF, LAD
U.K	Blundell & Etheridge	BHPS, FES, FRS, LFS
Germany	Fuchs-Schündeln, Krueger & Sommer	EVS, GSOEP
Italy	Jappelli & Pistaferri	SHIW
Spain	Pijoan-Mas & Sanchez Marcos	ECPF, ECHP, EFF
Sweden	Domeij & Floden	LINDA, LOUISE, HUT, HINK
Russia	Gorodnichenko, Stolyarov & Peter	RLMS
Mexico	Attanasio & Binelli	ENEU, ENIGH

More specifically, this special issue brings together the analyses of time trends in the distributions of wages, hours, earnings, income, consumption and wealth from nine countries. Table 1 summarizes the countries, the corresponding country teams and the underlying micro data used in the country studies in the order of the articles published in this volume.

Detailed documentations of the micro data sets used in each country study are contained in the nine papers of this volume. The household samples from which inequality statistics are computed, the computer codes, and the key time series are available (unless public dissemination of the data is not possible) at

<http://www.economicdynamics.org/RED-cross-sectional-facts.htm>

This website also contains a summary data set, readily available for use by researchers who want to document cross-country inequality trends in their papers, which includes time series for various inequality statistics (Gini coefficient, variance of logs, 50/10 and 90/50 percentile ratios) for wages, hours worked, labor earnings, market incomes, disposable incomes, consumption and wealth.² It also contains selected series measuring how inequality evolves over the life cycle and how wage and earnings shocks, estimated through a simple statistical model featuring permanent and transitory shocks, have evolved over time in the countries for which sufficiently rich panel data is available.

In the next section, we give a quick overview of the organizing principles for the entire empirical project. Then we briefly survey some of the main empirical findings, with special emphasis on the comparison across countries. Section 3.1 discusses how first moments, computed by aggregating

Switzerland and Canada) that are comparable across countries.

²Not all statistics for all variables are available for each of the nine countries in our study.

micro data in the different countries, compare with the corresponding NIPA data from the same countries. Sections 3.2, 3.3 and 3.4 compare and contrast the patterns of wage, earnings, disposable income, and consumption inequality across the countries studied in the volume. Section 3.5 focuses on the relations between inequality and business cycles, and Section 3.6 documents relation between inequality and the life cycle. Finally, Section 3.7 discusses the results of the estimation of a popular error-component model for wage and earnings dynamics in different countries.

2 Organizing the Data

Detailed guidelines were given to all country teams for conducting their analyses, in order to insure that the results are as comparable as possible across countries. The guidelines are also available online on the same website. In order to organize the data, our starting point is the budget constraint of a household in a typical heterogeneous agent macroeconomic model, since it is this equation that links the economic variables of interest. Broadly speaking, the literature falls into two categories, depending on whether individual or household labor supply is modeled endogenously or whether household labor earnings follow an exogenously specified stochastic process.

In the latter case, the period budget constraint reads as

$$(1) \quad c + a' = y^L + a + y^A + b + T$$

where y^L is pre-government labor earnings of all household members, a denotes the value of assets at the beginning of the period, y^A is private asset income of the household, b captures net private inter-vivos and bequest transfers from other households, T denotes transfers minus taxes from the government to the household, c denotes consumption expenditures of the household, and a' represents the value of assets that are accumulated for the next period.

In models where labor supply of a household, or individual labor supply of its members, is modeled endogenously pre-tax labor earnings of a household with two potential earners, a male and a female earner, can be written as

$$(2) \quad y^L = w_m l_m + w_f l_f$$

where (l_m, l_f) represent hours worked of the male and the female and (w_m, w_f) represent hourly wages of the male and the female member of the household, respectively. When labor supply is

modeled endogenously, it is wages (rather than earnings) that are assumed to follow an exogenously specified stochastic process.

Given that a budget constraint of the form of (1) is common to most heterogeneous agents macro models, the purpose of this project is to document trends in the distribution of the key variables appearing in this constraint. Starting from the variables typically treated as primitives (i.e. exogenous) in the model, wages (w_m, w_f) or labor earnings y^L this project documents the evolution, over time, in the distribution of hours worked, (l_m, l_f) , pre-government income $y^L + y^A + b$, disposable income $y^D = y^L + y^A + b + T$, consumption c and wealth a , for the nine countries. Distributions are high-dimensional objects; therefore the empirical analysis naturally has to be selective. In this project we focus mainly on the trends in statistics geared towards measuring the degree of dispersion or inequality of the economic variables of interest, leaving a detailed investigation of higher moments of the distributions for future work.

Since a significant share of the quantitative literature models the life cycle of a household explicitly, in addition to documenting inequality trends over time the project also provides information about how inequality in the key economic variables discussed above evolves over the life cycle. Finally, for those countries for which it is feasible we ask to estimate simple stochastic processes for wage or labor earnings, since they are crucial inputs into the quantitative macro models we would like to inform with this project.

3 Main Findings

We now describe some of the most important findings from the overall project, with particular focus on drawing cross-country comparisons.

3.1 Comparing Means from Micro Data to NIPA

Each country team was asked to first assess whether the trends in per capita income and consumption from the micro data sets employed in the analysis line up well with the corresponding trends in National Income and Product Accounts (NIPA) data. In part, the analysis of the first moments of the income and consumption distribution is motivated by the well-known fact that for the US the strong growth in real per capita consumption measured in NIPA is not reflected in micro consumption data from the US Consumer Expenditure Survey (CEX). One of the objectives of this empirical project therefore is to assess to what extent a similar divergence between micro

Table 2. Comparison of Means

Country	Levels		Growth	
	Income	Cons.	Income	Cons.
Canada	No	No	Yes	Yes
Germany	No	Yes	Yes	Yes
Italy	No	No	Yes	Yes
Mexico	No	No	Yes	Yes
Russia	No	Yes	Yes	Yes
Spain	No	No	Yes	Yes
Sweden	Yes	No	Yes	Yes
UK	No	No	Yes	No
USA	No	No	Yes	No

Notes: Qualitative classification of whether averages in micro data line up well with NIPA aggregates based on the interpretation of country teams.

and macro data exists for other countries and other variables of interest.

The comparison of per capita income and consumption between NIPA and the household level data sets in the various countries uncovers three key findings that are summarized in Table 2. First, by and large, in all countries the *time trends* in per capita income displayed in NIPA data are reproduced well by the corresponding micro data, although the absolute levels tend to be somewhat understated in most countries.

Second, in just about all countries the estimates of per capita consumption *levels* derived from micro data are significantly lower than the corresponding NIPA figures, a difference that can partially (and for some countries, almost fully) be attributed to differences in the definition of consumption between NIPA and household level data. In addition, the time series derived from the household level data sets tend to be more volatile than the NIPA data. This discrepancy could be due both to differences in definitions and to sampling variation in the micro data sets composed of at most a few thousand households.

Third, and most interestingly, in most countries the trends in per capita consumption from the household data line up rather well with the corresponding NIPA data. The two notable exceptions are the US and the UK³. Both the American CEX as well as the British Family Expenditure Survey (FES) display significantly slower growth in per capita consumption than does the corresponding

³Note that the length of the time series for per capita consumption from micro data is (significantly) longer in the US and the UK than in the other countries of this study, making it more likely to observe deviations between NIPA and household data over time.

Table 3. Wage Inequality and Wage Premia

Country	Level in year 2000				Change				
	Var. log w	College premium	Exp. premium	Gender premium	College premium	Exp. premium	Gender premium	Var. log w	Period
Canada	0.45	1.80	1.32	1.33	0.22	0.31	-0.11	0.17	1978-2006
Germany	0.27	1.38	1.27	1.28	-0.08	0.22	-0.15	0.05	1983-2003
Italy	0.17	1.51	1.34	1.03	-0.08	0.11	-0.05	0.03	1987-2006
Mexico	0.62	1.88	1.23	1.21	0.40	0.22	-0.06	0.04	1989-2002
Russia	0.77*	1.50	1.05*	1.49	-0.06	0.05*	-0.07	-0.13*	1998-2005
Spain ^(a)	0.23	1.48	1.43	1.16	-0.33	0.07	-0.21	-0.18	1985-1996
Sweden ^(b)	0.18	1.61	1.20	1.22	0.14	-0.02	-0.05	-0.09	1990-2001
UK	0.33	1.62*	1.25*	1.32	0.12*	0.20*	-0.21	0.10	1978-2005
USA	0.44*	1.80*	1.38*	1.36	0.40*	0.28*	-0.25*	0.21*	1980-2006
Average	0.38	1.62	1.27	1.27	0.11	0.17	-0.10	0.04	

A * indicates the statistic is from data on males only. Wage premia and wage dispersion for women are typically smaller.

(a) Data on changes refer to after-tax annual earnings

(b) Data on levels is for 1992

NIPA series. In the FES, the deterioration in the coverage rate starts in the mid-1990's, whereas the CEX displays no growth in real per capita nondurable consumption at all between 1980 and 2005, which is clearly at odds with US NIPA for the period following the 1982 recession. Both the US and the UK country studies in this issue discuss potential reasons for this observation in greater detail.

3.2 Wage Inequality and Wage Premia

We now briefly discuss cross-country differences and similarities in the distribution of wages and in the determinants of wage inequality. Results are summarized in Table 3.

Levels Caution should be exercised in comparing the level of wage inequality across countries, since different surveys might be more or less affected by reporting error. With this caveat in mind, one would expect the level of wage disparity to be larger in countries where institutional constraints in the labor market are less severe. This is indeed what we find. In 2000, the variance of log hourly wages in Canada, and the US was around 0.45, whereas in Spain, Italy, and Germany it was roughly half as large. The level of wage inequality in the UK is somewhere in between. Interestingly, in 2000, the countries in our sample with the largest levels of wage dispersion are the two least developed countries, Mexico and Russia.

In 2000, the college premium (i.e., the ratio between the average hourly wage of college graduates

and the average hourly wage of high-school graduates) is remarkably similar across all European countries, between 1.4 and 1.6, while in the US and Canada it reaches 1.8. In Mexico, where the fraction of the college-educated labor force is the smallest among our nine countries, the premium for those holding a college degree is close to 1.9. The experience premium (i.e., the ratio between the average hourly wage of 45-55 years old and the average hourly wage of 25-35 years old) is between 1.3 and 1.4 for most countries, except for Sweden and Mexico, where it stands at 1.2, and, for Russia, where survey data show no evidence of a positive experience premium. In 2000, the gender premium (i.e., the ratio of the average wages of men to the average wages of women) is typically between 1.2 and 1.4 in most countries. Outliers on the low end are Spain and Italy, the countries with the lowest female participation rate, and on the high end is Russia, where the gender gap reaches 1.5.

Time trends The Anglo-Saxon countries in our study, namely Canada, UK, and US, experienced a continuous and sharp increase in cross-sectional wage dispersion during the last thirty years. In these three countries, the variance of male log wages rose by roughly 40% from 1980 to 2005. This increase in inequality is only partially accounted for by observable characteristics, such as experience and education, i.e. it is largely residual in its nature. In particular, while the skill premium plays a sizeable role in the dynamics of wage inequality in the US and the UK throughout the period, in Canada it only does so in the last decade.

Among the continental European countries, the only featuring a recognizable long-run trend are Spain and Sweden, where inequality fell quite sharply. Russia also witnessed a large reduction of inequality between 1998 and 2005. All other countries display only episodic changes. For example, German data show an upward jump in the second half of the 1990s, and Italian data in the first half of the 1990s. Consistently with what emerges from US, UK and Canada, in all other countries the residual component of wage inequality plays a crucial role. The only exception is Spain, where the fall of wage inequality seems to largely coincide with that of the college premium, while residual wage dispersion is flat.

Turning to the skill premium, we note a clear dichotomy between US, UK, Canada and Mexico on the one hand, where the skill premium in 2005 is significantly larger than twenty years earlier, and all continental European countries (with the exception of Sweden), where the skill premium declined. Most of the countries in our sample saw a rise in the experience premium around 20 percent since the mid 1980s. Finally, over the same period, the gender gap shrank in every country, with the exception of Sweden.

3.3 From Inequality in Wages to Inequality in Disposable Income

As discussed above, our key organizing device is the household budget constraint which provides a natural tool for understanding how different dimensions of inequality are related via endogenous choices (e.g., labor supply and saving), financial markets, and institutions (e.g., the tax system). In the previous section, we started with individual wages, as our most primitive measure of inequality. We now take a series of steps to contrast inequality in individual wages to that in individual earnings, household earnings, pre-government income, and disposable income. Along the way, we evaluate the impact on measured inequality of individual labor supply, household income pooling, private transfers, asset income, and government redistribution.

Role of individual labor supply Inequality in earnings (measured as hourly wages times hours worked) is found to be systematically larger than inequality in wages. To see why this is the case, note that the variance of log earnings can be decomposed into the variance of log wages, the variance of log hours, and the wage-hours covariance. Hours dispersion is sizeable, especially for women, and it is the reason why earnings inequality is larger than wage inequality. Interestingly, in all the surveys considered, the correlation between log wages and log hours is found to be negative, usually between -0.1 and -0.3 . The basic neoclassical theory of labor supply suggests that when income effects dominate substitution effect, hours can co-move negatively with wages. As a partial corroboration of this view, authors often found that the correlation for women is closer to zero: women's earnings being smaller, they have a more limited income effect on household earnings. However, one should keep in mind that measurement error plays a big role in driving down this correlation: hourly wages in surveys are typically measured as earnings divided by hours worked. Therefore measurement error in reported hours artificially lowers this correlation.

With respect to time trends, one stark finding is that the US, UK and Canada show remarkably similar time paths for wage dispersion, hours dispersion and wage-hour correlation since 1975. Wage dispersion rises steadily over time. Hours dispersion is fairly constant for men, and falling for women. As more and more women work full time, dispersion in labor supply shrinks and converges to its male counterpart. Wage-hours correlation has been increasing quite steadily, especially until the mid 1980s. The combination of these forces yields a sharper rise in individual earnings inequality compared to wage inequality, especially for men.

Finally, a robust finding of this volume is that earnings inequality increases more rapidly than wage inequality during recessions because of strongly counter-cyclical movements in the dispersion of hours worked due to the surge in unemployment rates. We return to this point in Section 3.5.

Role of family labor supply The articles in this volume paint an uneven picture of the impact of family labor supply on inequality statistics. In some countries (e.g., Canada, Sweden, and UK), the level of inequality in equivalized household earnings is systematically lower than inequality in head earnings. In other countries (e.g., Germany, Italy and Spain), the data suggest the exact opposite.

It is perfectly sensible for both patterns to emerge from the data. On the one hand, incomes of the two spouses are imperfectly correlated. Moreover, often the secondary earner responds to a temporary income loss accruing to the primary earner (e.g., due to an unemployment spell) by increasing her market hours, thereby mitigating the decline in household earnings. This form of within-household income pooling and insurance reduces inequality in household earnings below inequality in head earnings. On the other hand, low-income families are more likely to receive labor earnings from one member only relative to high-income families where often both spouses actively participate in the labor market. This force raises inequality in equivalized household earnings above inequality in head earnings.

Which effect eventually dominates depends on the distribution of marital status and women's participation rates across income levels, and on the correlation between earnings of the two spouses in married families where both spouses work. In line with our finding, one would expect larger within-family insurance in countries where female participation rates are high and not very dependent on income levels, such as the Anglo-Saxon and Nordic countries.

With respect to long-run time trends, in the US and the UK, two countries with a large increase in inequality, the relevant papers uncover that the rise of inequality in equivalized household earnings has been slightly smaller than the rise in the dispersion of head earnings, mostly due to the steady growth in female labor force participation.

Role of private transfers and home production In almost all surveys, the authors found that the magnitude of private transfers across households is very small, usually insignificant. Two notable exceptions are Mexico and Italy, where measures of inequality fall sizably once private transfers are included in income.

None of our surveys, with the exception of Russia, makes a rigorous attempt to measure home production. Imputed income from home production tends to reduce cross-sectional inequality among households, because its value is only weakly correlated with income (if anything, it is negatively correlated), and hence as a fraction of total income it is more sizeable for the poor (e.g., Gottschalk and Mayer, 2002). The data on Russian households show indeed that income from

subsidiary farming, which includes both own consumption valued at market prices and sales of home grown food, has a large equalizing effect on the income distribution.

Role of capital income Turning to capital income, it is important to keep two things in mind. First, our samples are restricted to households with at least one working-age adult. Thus we mostly miss elderly households, which rely primarily on unearned income. Second, wealth, and hence capital income, is extremely concentrated at the top of the distribution and in many surveys the coverage rate of the very wealthy is extremely low.

In general, we find that adding asset income has little impact on inequality measured by the variance of log income. There are three main reasons for this finding. To begin with, median asset income is small, so for the majority of working households this source of income is dwarfed by earnings. Furthermore, capital income often suffers from severe under-reporting in survey data. An example is Spain, where from one survey (ECPF) it appears that asset income is negligible in shaping inequality, from another survey (ECHP) where it is better measured, it is found that capital income increases the level and the rise in income inequality. Finally, because asset income is so concentrated at the top, it only affects marginally the variance of log income which, by construction, is much more sensitive to the bottom than the top of the distribution. The Gini coefficient, which move more closely with inequality in the upper half of the distribution, as measured e.g. by the 90-50 ratio is more sensitive to the inclusion of capital income. For example, in the US the Gini coefficient of household earnings plus capital income is at least one point higher than the Gini for household earnings, and its rise over time is stronger.

Role of fiscal redistribution Comparing cross-sectional inequality in pre-government income to inequality in disposable income yields valuable insights into the role of the government tax and transfer system. In virtually every modern economy, the marginal tax rate rises with income, and public transfers are directed towards the poor (e.g. through cash assistance, rent and child subsidies) and towards those temporarily out of work (e.g., unemployment insurance). As a result, the public tax and social insurance system is highly progressive in most countries and one should expect it to play a key role in shaping income inequality among households. That this is indeed the case is one of the most robust findings of this project, as evident from Table 4.

First, government redistribution compresses the *level* of inequality in every country studied in this volume: cross-sectional dispersion in disposable income is significantly lower than dispersion in household earnings.⁴ The ratio between these two measures varies across countries, however.

⁴Earnings data for some countries, like Italy, Russia and Spain, are already reported net of income taxes. Hence,

**Table 4. Inequality in Pre and Post Gov. Household Income
(Variance of the Log)**

Country	Level in year 2000		Change		Period
	Pre-gov. income	Post-gov. income	Pre-gov. income	Post-gov. income	
Canada	0.50	0.25	0.16	0.05	1978-2005
Germany	0.63	0.40	0.42	0.04	1984-2004
Italy ^(a)	0.72	0.73	0.06	0.07	1987-2006
Mexico	2.10	1.70	1.15	0.75	1989-2002
Russia ^(b)	0.86	0.60	-0.11	-0.09	1994-2005
Spain ^(c)	0.73	0.56	-0.20	-0.09	1993-2000
Sweden	0.95	0.38	0.36	0.05	1978-2004
UK ^(d)	0.55	0.32	0.22	0.13	1978-2005
USA	0.67	0.41	0.11	0.11	1979-2005
Average	0.86	0.59	0.24	0.11	

(a) Data on pre-gov. income are already after tax

(b) Data on pre-gov. income are already after tax and refer to working households

(c) Data on pre-gov. income are already after tax

(d) Data refer to households with at least one worker

For example, in 2000, in terms of the variance of the log, this ratio was less than half for Sweden and Canada, and around two thirds for the US, UK, and Germany. When the data are detailed enough to disentangle the role of taxes from that of transfers, what stands out is that transfers have the largest effect in compressing inequality at the bottom of the distribution (as evident from the variance of the log), and taxes at the top (as evident from the Gini coefficient).⁵

Second, over the sample period, in several countries the tax and transfer system has reduced long-run increases in household earnings inequality. Perhaps, the most striking examples are Sweden, Canada and Germany where in the 1980s and 1990s disposable income inequality shows a much more moderate increase compared to the rise in pre-government income inequality. In contrast, in the US and the UK, while fiscal redistribution is effective in reducing the level of inequality, the increases in pre and post government income dispersion appear to be more similar in magnitude.

Third, automatic stabilizers implicit in the government transfers system (namely through unemployment benefits) seem to be very effective in smoothing *cyclical fluctuations* in household earnings, particularly at the bottom of the earnings distribution, as expected. This is evident from

one cannot properly identify the effectiveness of the government in compressing the level of inequality and in absorbing its trend.

⁵The distinction between taxes and transfers is often unclear, however. For example, in the US and the UK a primary source of government support to the poor occurs through the tax system, e.g. via the Earned Income Tax Credit.

the fact that the time series for the variance of log household income is typically much smoother when government transfers are included in the definition of income. This finding is evident, for example, from Figure 12 in the US article and Figure 4.2 in the UK article.

3.4 Disposable Income and Consumption Inequality

The previous section has showed that the connection between inequality in wages or earnings and inequality in disposable income is quite different across countries, mostly due to differences in the way the family and the government operate. In this section, we show that, conditional on the realized path of disposable income inequality, the relation between inequality in disposable income and inequality in consumption is instead very similar across the countries we study. In particular, the six following stylized facts characterize this relation in most of the countries we study:

1. The level of inequality in disposable income is larger at the bottom than at the top of the distribution
2. The level of inequality in disposable income is larger than inequality in consumption
3. The gap between the level of disposable income inequality and consumption inequality is larger at the bottom than at the top of the distribution
4. Long run changes in disposable income inequality are larger at the bottom than at the top of the distribution
5. Long run changes (positive or negative) in disposable income inequality are larger than long run changes in consumption inequality
6. The gap between long run changes in disposable income inequality and the corresponding change in consumption inequality is larger at the top than at the bottom of the distribution

Facts 2 and 5 are also pointed out by Blundell, Pistaferri and Preston (2008), and appear to be true for other countries not studied here (such as Australia and Japan). The first three stylized facts relate to the level of inequality in the two variables and are documented in Table 5. As shown in the last line of the table, on average the 50/10 ratio for income is around 3 while the 90/50 ratio is 2.33 (fact 1). The fact holds for all countries in the sample with the exception of Sweden and it probably suggests a large concentration of households with very low disposable income. Moreover, on average inequality in disposable income is higher than inequality in consumption (fact 2). Again

Table 5. Level of Inequality in Year 2000

Country	Bottom (50/10)			Top (90/50)		
	Disp Inc.	Cons.	Gap	Disp Inc.	Cons.	Gap
Canada	2.21	1.95	0.26	2.00	1.85	0.15
Germany	2.05	1.70	0.35	1.80	1.81	-0.01
Italy	2.45	1.91	0.54	1.93	1.88	0.05
Mexico	8.00	5.10	2.90	4.75	4.00	0.75
Russia	3.02	2.70	0.32	2.60	2.60	0.00
Spain*	2.04	1.82	0.22	2.00	1.90	0.10
Sweden	1.58	1.62	-0.04	1.64	1.73	-0.09
UK	2.82	NA	NA	2.08	NA	NA
USA	2.64	2.00	0.64	2.21	2.0	0.21
Average	2.98	2.35	0.65	2.33	2.22	0.15

* The level for Spain refers to year 1996

this is true for all countries with the exception of Sweden.⁶ Finally, the gap between the 50/10 ratio in disposable income and consumption is around 0.65 while the corresponding gap for the 90/50 is only 0.15 (fact 3), possibly suggesting that although there might be a high concentration of low disposable income households, there doesn't appear to be a concentration of low consumption households (note that the 50/10 and the 90/50 in consumption are fairly similar).

Table 6 documents facts 4 through 6 which relate to inequality trends. The last line shows that on average disposable income inequality has grown (the only exception is Spain) but not uniformly across the distribution, as the change in the 50/10 ratio is 0.91, much higher than the corresponding change in the 90/50 ratio which is only 0.22 (fact 4). Again this fact holds for all the countries in the sample with the exception of Sweden. The last line of the table also shows that changes in inequality in disposable income have outstripped changes in consumption inequality (fact 5), but more so for the 50/10 ratios than for the 90/50 (fact 6). This last fact is also true for all countries studied, even though there are substantial quantitative differences in the experiences of the various countries. For example, in Mexico the gap in inequality growth of income and consumption at the bottom is a very large 5.01 (i.e., at the bottom of the distribution income inequality has grown more rapidly than consumption inequality) while the same gap at the top is basically 0 (i.e. at the top income and consumption inequality have tracked each other closely). At the other extreme, in Sweden the growth gap is similar at the top and the bottom of the distribution. Overall, Table 6 suggests that the distribution of disposable income has become more unequal and that higher inequality shows

⁶One possible reason why Sweden is an exception is that in the HUT database consumption is based on bi-weekly data while income is annual.

Table 6. Long-run Changes in Inequality

Country	Bottom (50/10)			Top (90/50)			Period
	Disp. Inc.	Cons.	Gap	Disp. Inc.	Cons.	Gap	
Canada	0.38	0.20	0.18	0.10	0.07	0.03	1978-2006
Germany	0.35	0.00	0.35	0.15	0.10	0.05	1983-2003
Italy	0.22	0.09	0.13	0.05	0.01	0.04	1980-2006
Mexico	5.81	0.80	5.01	1.12	1.08	0.04	1989-2002
Russia	0.10	0.05	0.05	-0.16	-0.10	-0.06	1994-2005
Spain	-0.16	-0.13	-0.03	-0.18	0.01	-0.17	1985-1996
Sweden	0.13	0.02	0.11	0.21	0.10	0.11	1985-1998
UK	0.86	0.58	0.28	0.27	0.12	0.15	1978-2005
USA	0.55	0.25	0.30	0.40	0.15	0.25	1980-2006
Average	0.91	0.21	0.71	0.22	0.17	0.05	

up mostly with the bottom of the distribution losing ground relative to the median. In terms of the consumption distribution, the table paints a different picture: consumption has also become more unequally distributed (although much less so than disposable income) and the increase in inequality has been pretty symmetric around the median.

A tentative interpretation of our findings would suggest that either i) a large fraction of the shocks driving the level and the changes in inequality in disposable income are temporary, so that agents can keep a smoother consumption profile through self insurance and/or ii) there are formal/informal mechanisms of insurance/transfers (not reported in our disposable income data) that prevent the level or the growth of disposable income inequality to fully translate into consumption inequality. An interesting direction for future research would be to explore why these self-insurance/transfers mechanisms appear to be more effective at the bottom rather than at the top of the distribution. One possibility is that households at the bottom of the distribution are there because hit by temporary shocks (e.g., short unemployment spells), the other is that households at the bottom of the distribution rely more on transfers from informal networks of family and friends.

A final note is about changes in consumption inequality that are disconnected from changes in income inequality. An interesting case is the one of the UK over the 1995-2005 period (see Figure 5.3 in the UK paper), during which disposable income inequality was stable or falling but consumption inequality was increasing. The authors of the UK paper suggest that this might be due to appreciating housing prices that have lifted the net worth of the top and middle part of the distribution (i.e. home owners), but not that of the bottom 10% of the distribution (who do not own a home), so that consumption dynamics at the bottom is not only affected by income inequality (which is falling) but also by net worth inequality, which is growing. More generally,

when households hold different asset portfolios, a large change in prices of financial assets (like the housing prices boom and bust that most countries have recently experienced) can have an independent effect on consumption inequality.

3.5 Inequality Over the Business Cycle

One of the goals of this project is to stress the connection between inequality and the macroeconomy. Possibly the strongest evidence of this connection appears during recessions, when the overall macroeconomic activity slows down and, at the same time, inequality in many variables changes. All the studies in the volume include recessionary episodes, so in this section we'll briefly summarize the patterns of inequality during recessions and speculate on the possible causes of these patterns. Most of the countries studied in this volume witnessed sharp long-run trends, besides business-cycle fluctuations, in inequality. A formal study of the link between recessions and inequality would require a statistical decomposition of trend and cycle, which is beyond the scope of our project: in what follows, our summary will be purely descriptive.

Wages The dynamics of wage inequality during downturns are not uniform across countries nor across recession episodes. The studies in the volume document recessions in which inequality in wages fell, remained flat, and increased. This diversity suggests that the impact of recession on wage inequality is likely to depend on the specific causes of the recession and on the structure of the labor market.

Earnings and hours Differently from wages, the dynamics of earnings and hours inequality are more similar across countries and across recessions. A common pattern that we observe in all recessions and all countries is that during bad times earnings inequality at the bottom of the distribution increases sharply. This is driven, rather mechanically, by the rise in unemployment which pushes a larger number of individuals to the bottom of the hours—and hence the earnings—distribution. Some recessionary episodes are also associated with an increase in earnings inequality at the top (during the early 1990s recessions in Italy and in Sweden, the 90/50 earnings ratio rose substantially, see Figure 12 of the Italy paper and Figure 8 of the Sweden paper), whereas, in other episodes, inequality at the top is stable during downturns (e.g., in the US during the recessions of the 1970s).

Differences across business cycles and across countries also arise in the persistence of earnings inequality once the recession is over. For example, in the US, the recession of 1980-1981 coincided with an increase in earnings inequality (at the bottom) that persisted for several years after its

trough. In contrast, the rise in earnings inequality associated with the 1990 recession has been very short-lived. Since earnings inequality is closely related to unemployment, its persistence is related to the persistence of a high unemployment rate after the downturn, a feature linked to the causes of the business cycle and to the functioning of the labor markets (see Figures 7 in the Sweden and US papers that connect the persistence of earnings inequality and unemployment). In some cases though, the increase of earnings inequality after a recession outlasted the increase in unemployment, e.g., in Italy after the 1992 recession: by year 2004, unemployment had fallen to its pre-recession levels, yet earnings inequality was still well above its pre-1992 levels (see Figure 11 in the Italy paper). This disconnect is due to the fact, noted earlier, that the 1992 recession in Italy had also lead to a permanent increase in wage inequality.

Disposable income Although in all countries earnings inequality at the bottom increases during recessions, the extent to which this increase translates into a rise in disposable income inequality depends on country specific government policies (in particular, policies such as unemployment insurance). It is instructive to consider the case of the early 1990s recessions in Sweden, Canada and the US. In Sweden, the recession led to a surge in the variance of log earnings of over 50 points but the variance of log disposable income increased by less than 10 points (see Figure 10 in the Sweden paper). Similarly, in Canada, the downturn led to an increase in the variance of log earnings of over 20 points, whereas disposable income inequality rose by less than 4 points (see Figure 9 in the Canada paper). In the US instead, although the 1990 recession caused a much more modest increase in earnings inequality (only 10 log points), at least half of it translated into inequality in disposable income (see Figures 10 and 11 in the US paper). The general pattern is that, in all countries and in all recessions, inequality in disposable income during the recession rises less than inequality in earnings, reflecting the significant role played by automatic stabilizers. Quantitatively this role appears to be larger in some countries (i.e. Canada, Sweden, Germany) and smaller in others (US, Italy).

Consumption Our findings for the patterns of consumption inequality during recessions largely mirror the findings described in the previous section on the relation between long-run changes in disposable income and in consumption inequality. In particular, in most downturns, we observe an uptick in consumption inequality (especially at the bottom of the distribution) that is smaller than the corresponding increase in disposable income inequality. As argued earlier, we conjecture that the smaller increase in consumption inequality (relative to disposable income inequality) can be probably explained by the fact that a substantial fraction of the upsurge in

inequality during recessions is due to an increase in unemployment, a fairly transitory shock, which households can smooth through borrowing and saving and with the help of public transfers.

Wealth Although the experience of the Great Depression indicates that recessions can have a major impact on wealth inequality (see Wolff 1992), the studies in this volume reporting information on wealth inequality do not detect a strong connection between wealth inequality and recession episodes. For example, although the recessions of the 1990s in Sweden and Italy were major episodes in terms of the increase in earnings inequality, wealth inequality in those two countries did not significantly change during the recessions. One possible reason for this pattern is that those recession were not associated with major changes in asset prices. An interesting task for future research will be to verify whether the 2008 recession, which has coincided with a major change in prices of several assets, namely equity and housing, will bring about changes in the wealth distribution.

To summarize, the studies in this volume provide ample empirical evidence on the link between inequality and the business cycle, for several variables. This evidence should help researchers to identify the fundamental causes of this connection, since these are key to understand and measure how the welfare cost of recessions are distributed across households.

3.6 Inequality over the Life Cycle

Many macroeconomic models contain an explicit life cycle structure. Modigliani's life cycle hypothesis is an obvious and very important example. We asked the contributors of this volume to supplement evidence on time series inequality with evidence on inequality over the life cycle. In most cases, authors looked at age profiles of inequality in hours, wages, earnings, disposable income and consumption.

As argued by Deaton and Paxson (1994), the slope of the age profile for income and consumption inequality can be informative about the nature of the income process and about insurance opportunities available to households. For example, if income follows a random walk process, the cross-sectional variance of income increases linearly with age. Following a cohort over time, we would observe a fanning-out of inequality because the realization of permanent shocks makes individuals in the same cohort more and more different as they age. This linearity is lost if the variance of the permanent shock varies by age instead of being constant, in which case the only prediction is that the age profile of income inequality should be non-decreasing. If income shocks follow a persistent, but mean-reverting, autoregressive process, then the life-cycle profile would reveal some

concavity as the cohort ages. Note that, to generate a fanning out of consumption or earnings profiles over the life cycle, non-stationarity in the time series is not required; all that is needed is persistence of consumption or income innovations.

Under the strict version of the permanent income hypothesis (quadratic utility and $\beta(1+r) = 1$), if there is a unit root in disposable income, consumption and income inequality should grow at the same rate over the life cycle. However, in the presence of some insurance against permanent shocks (beyond self-insurance), consumption dispersion will grow less rapidly than income dispersion. In a more general model with both transitory and permanent income shocks and CRRA utility, consumption inequality is expected to grow less strongly than income inequality over the life cycle due to self-insurance through saving and borrowing. The greater the contribution of the transitory shock to the total income variance, the greater the detachment between income and consumption inequality over the life cycle. Partial insurance against permanent shocks makes this gap larger; binding liquidity constraints narrow the gap.

Seen from a dynamic perspective, most economic variables vary not only because of age effects but also because of time and cohort effects. Consider wages. While human capital theory predicts that wages exhibit a concave shape over the life-cycle, it is also true that wages may vary because of time effects (e.g., due to economy-wide productivity shocks) or cohort effects (e.g., due to relative supply effects in labor markets induced by large-size cohorts). Unfortunately, in general one cannot separately identify the effect of age, time and cohort because age (j), survey year (t) and birth cohort (c) are perfectly collinear in the population ($j = t - c$). While one can impose restrictions to identify certain parameters, they are somewhat arbitrary. Instead of restricting profiles in an arbitrary way, we asked contributors to either plot unrestricted life-cycle profiles, or else life-cycle profiles in which either time or cohort effects have been assumed absent. Hence, let the typical cross-sectional moment for age j at time t be $M(j, t)$ (e.g., the cross-sectional variance of log consumption for individuals aged 25 in year 2000). We asked authors to regress $M(j, t)$ on a full set of age dummies and either a full set of year dummies, or a full set of year-of-birth dummies. The age profile for the variable of interest can be obtained as the predicted age-portion of this regression (after a suitable normalization at a given age).

A number of findings emerge from the comparative analysis. First, in most countries, age profiles controlling for time effects (and omitting cohort effects) differ dramatically from age profiles controlling for cohort effects (and omitting time effects), either because the growth over the life-cycle is quantitatively different, or because the shapes of the profiles are *qualitatively* different. A

stark example is Sweden, where life-cycle wage profiles are increasing when controlling for time effects and decreasing if one controls for cohort effects (incidentally, decreasing age profiles are inconsistent with any income process exhibiting persistence).

To give a comparative summary view, in most of the countries that we study the growth in inequality in disposable income, when controlling for cohort effects, is larger than the growth when controlling for time effects. This happens in Canada (slightly so, with a 0.2 log points overall increase), Germany (also slightly so, with little or no growth at all), USA (0.25 vs. 0.1), Mexico (where the reference series is equivalized earnings, which grows by 0.3 log points as opposed to 0), and Sweden (where again the reference series is for equivalized earnings, profiles obtained controlling for time effects have an inverted U-shape while those controlling for cohort effects grow by 0.6 log points). The opposite pattern occurs in Russia (where cohort profiles are decreasing and time profiles are flat). In the US, Canada, Germany, Russia the age profiles for all variables of interest (wages, earnings, disposable income and consumption) have a qualitatively similar shape irrespective of time or cohort controls, even though the magnitude of their increase is different. In Mexico and Sweden the shapes are actually qualitatively different: age profiles when controlling for cohort effects are increasing while those obtained when controlling for time effects are decreasing (or *vice versa*).

Second, disposable income inequality grows faster than consumption inequality over the life cycle. This result, that appears to be true in all countries with the exception of Mexico, suggests the presence of insurance against income shocks. Since transitory shocks are more easily insurable than permanent shocks, a natural question, in this context, is how much of cross-sectional inequality is transitory in nature and how much is permanent in nature?

3.7 Estimation of Earnings and Wage Processes

The typical specification of income processes found in the literature, motivated by Friedman's permanent income hypothesis, is careful in distinguishing between permanent and transitory shocks to income. Of course in empirical work things are never as simple as this stark decomposition suggests: long-lived shocks may not be truly permanent and short-lived shocks may be reasonably persistent. Finally, what may pass as a permanent shock may sometimes be heterogeneity in disguise. These issues are fueling a lively debate in an active area of research that seeks to model stochastic wage or income processes, and we do not attempt to resolve these issues in this volume. Instead, we invited authors to perform a simple minimum distance estimation exercise (Chamberlain, 1984):

whenever a panel dimension on individual wages or household earnings was available, contributors were asked to estimate a statistical model for income that is the sum of a permanent (unit root) and a transitory (uncorrelated) component. Table 7.A, 7.B and 7.C present a summary of the results obtained in the various countries for wages, earnings and disposable income, respectively.

Starting with wages, two issues stand out, one methodological and one more substantial. On the methodological side, estimates obtained using covariance restrictions on the *levels* of log wages (or log earnings) in some cases differed significantly from those obtained using covariance restrictions on *first differences* (i.e. growth rates) of the data. This suggests that a simple statistical model with only purely permanent and purely transitory shocks may be mis-specified. For example, the true model could exhibit less persistence than a unit root, or an MA instead of an uncorrelated transitory component. This is, in our view, an issue that is worth exploring in future work. The main substantial finding from this exercise is that the variance of permanent wage shocks is much smaller than the variance of transitory wage shocks in all countries examined. However, the variance of transitory shocks most likely captures also variability due to measurement error, a pervasive phenomenon in micro data. As for trends, most countries exhibit a rise in the variances of both components in the 1990s, although the experience of the various countries is more nuanced.

Moving to household earnings, the discrepancy between estimates obtained using restrictions in levels and those obtained using restrictions in first differences becomes even larger. For example, in Germany the covariance restrictions in levels give implausibly high estimates of the variance of transitory shocks, while those in first-differences give implausibly high estimates of the variance of permanent shocks. In Spain there are large differences between the two surveys used for the analysis, suggesting that the measurement of earnings in ECHP may be of lower quality. For Russia, the high estimates of the variances appear to reflect genuinely higher labor market risks.

The results for disposable income largely confirm the findings for wages and household earnings. The estimates of the variances of shocks decline relative to the wage and household earnings case (with the notable exception of Russia), consistently with the role of insurance provided by taxes and transfers. In Italy the rise in disposable income inequality can be entirely explained by a rise in the variance of transitory shocks. Sweden, Canada and Russia exhibit the opposite pattern.

4 Conclusions

For many years, labor economics documented that heterogeneity in households characteristics and household behavior at the micro level is pervasive, but the standard model used in macroeconomics insisted on the representative agent abstraction. In the last decade, substantial work in macroeconomics has attempted to close this gap between theory and measurement, as heterogeneous agent, incomplete markets models have become a standard tool of quantitative macroeconomic analysis.

This volume offers, for nine countries at different levels of development, a wealth of cross-sectional facts intended for macroeconomists interested in calibrating/estimating these models and using them for descriptive analysis and normative policy evaluation. Of course, our study is just a first step, and more studies of this nature on other countries, and new data sources, are needed. After describing the objectives and the common methodology underlying the project, in this article we have tried to summarize the main findings, with an emphasis on cross-country comparisons. However, our summary is necessarily very imperfect, given how diverse the experience of each country is from every other. For this reason, we invite you to continue reading the fascinating individual articles contained in this volume.

References

- [1] Aiyagari, R. (1994), “Uninsured Idiosyncratic Risk and Aggregate Savings,” *Quarterly Journal of Economics*, 59(3), 659-84.
- [2] Bewley, T. (1986), “Stationary Monetary Equilibrium with a Continuum of Independently Fluctuating Consumers,” in Werner Hildenbrand and Andreu Mas-Colell (eds.) *Contributions to Mathematical Economics in Honor of Gerard Debreu*, North-Holland, 79-102
- [3] Burns, A.. and W. Mitchell (1946), *Measuring Business Cycles*, National Bureau of Economic Research Book Series Studies in Business Cycles, Cambridge, MA.
- [4] Chamberlain, G. (1984), “Panel Data,” in Z. Grilliches and M. Intriligator (eds.) *Handbook of Econometrics, Vol 2*, North Holland.
- [5] Deaton, A. and C. Paxson (1994), “Intertemporal Choice and Inequality” *Journal of Political Economy*, vol. 102(3), pages 437-467.
- [6] Gottschalk, P., and S. Mayer (2002), “Changes in Home Production and Trends in Economic Inequality,” in Daniel Cohen, Thomas Piketty, and Gilles Saint-Paul (eds.), *The Economics of Rising Inequalities*, Oxford University Press.
- [7] Gottschalk, P. and T. Smeeding (1997), “Cross-National Comparisons of Earnings and Income Inequality,” *Journal of Economic Literature*, 35, 633-687.
- [8] Heston, A., R. Summers and B. Aten (2009), “Penn World Table Version 6.3,” Center for International Comparisons of Production, Income and Prices at the University of Pennsylvania.
- [9] Huggett, M. (1993), “The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies,” *Journal of Economic Dynamics and Control*, 17, 953-969.
- [10] Imrohoroglu, A. (1989), “Cost of Business Cycles with Indivisibilities and Liquidity Constraints,” *Journal of Political Economy*, 97, 1364-1383.
- [11] Laitner, J. (1979), “Household Bequest Behaviour and the National Distribution of Wealth,” *Review of Economic Studies*, vol. 46(3) pages 467–483.
- [12] Rios-Rull, V. (1995), “Models with Heterogeneous Agents,” in Thomas Cooley (ed.), *Frontiers of Business Cycle Research*, Princeton, NJ: Princeton University Press.
- [13] Wolff E. (1992), “Changing Inequality of Wealth”, *American Economic Review*, 82(2), pages 552-558