JOURNAL OF Development ECONOMICS

Journal of Development Economics
Vol. 61 (2000) 417–446

www.elsevier.com/locate/econbase

Does income inequality raise aggregate saving?

Klaus Schmidt-Hebbel [a,*], Luis Servén [b,1]

[a] *Central Bank of Chile, Agustinas 1180, Santiago, Chile*
[b] *The World Bank, 1818 H Street, NW, Washington, DC 20433, USA*

Abstract

This paper reviews analytically and empirically the links between income distribution and aggregate saving. Consumption theory brings out a number of direct channels through which income inequality can affect overall household saving — positively in most cases. However, recent political-economy theory points toward indirect, negative effects of inequality — through firm investment and public saving — on aggregate saving. On theoretical grounds, the sign of the saving–inequality link is therefore ambiguous. This paper presents new empirical evidence on the relationship between income distribution and aggregate saving based on a new and improved income distribution database for both industrial and developing countries. The empirical results, using alternative inequality and saving measures and various econometric specifications on both cross-section and panel data, provide no support for the notion that income inequality has any systematic effect on aggregate saving. These findings are consistent with the theoretical ambiguity.

PACS: E21; D31
Keywords: Saving; Inequality; Income distribution

1. Introduction

The theoretical relation between saving and income distribution has received attention both in the historical growth literature and more recent neoclassical consumption theory. In fact, the link between the functional distribution of income

* Corresponding author. Tel.: +56-2-670-2586; fax: +56-2-670-2853.
E-mail addresses: kschmidt@condor.bcentral.cl (K. Schmidt-Hebbel), lserven@worldbank.org (L. Servén).
1 Tel.: +1-202-473-7451; fax: +1-202-522-2119.

PII: S0304-3878(00)00063-8

and saving (and growth) is at the heart of the neoclassical growth model (Solow, 1956) as well as neo-Keynesian growth models (Lewis, 1954; Kaldor, 1957; Pasinetti 1962). In more recent neoclassical consumption theory, the focus is on the links between personal distribution of income and aggregate saving as a result of consumer heterogeneity in regard to endowments or institutional constraints (see Deaton, 1992 for a discussion). Most (but not all) of the mechanisms pointed out by the above mentioned theories suggest positive direct effects of income inequality on overall personal saving. However, recent political-economy research brings out negative indirect links from inequality — through investment, growth, or public saving — to aggregate saving (e.g., Dornbusch and Edwards, 1991; Persson and Tabellini, 1994; Alesina and Perotti, 1996). Taken together, these two strands of the literature imply that the overall impact of inequality on aggregate saving is theoretically ambiguous and can only be assessed empirically.

Most of the empirical literature on the links between personal income inequality and personal saving based on cross-section micro data suggest a positive relation between them (e.g., Bunting, 1991; Dynan et al., 1996). In turn, the evidence from aggregate (typically cross-country) data is more mixed. Some studies also find positive and significant effects of personal income inequality on aggregate saving (Sahota, 1993; Cook, 1995; Hong, 1995), but other studies do not (Della Valle and Oguchi, 1976; Musgrove, 1980; Edwards, 1996). Reconciling these conflicting results is difficult because empirical studies based on macro data use different samples and specifications, different measures of saving and inequality and, in most cases, income distribution information of questionable quality.

This paper reexamines the empirical evidence from macro data on the links between the distribution of personal income and aggregate saving, controlling for relevant saving determinants. It provides an encompassing framework and a robustness check for previous empirical studies, and extends them in five dimensions: (i) testing alternative saving specifications; (ii) using alternative inequality and saving measures; (iii) making use of newer, better and larger databases; (iv) conducting estimations jointly and separately for industrialized and developing countries; and (v) applying various estimation techniques on both cross-country and panel data. On the whole, we do not find any consistent effect of income inequality on aggregate saving — a result that is consistent with the theoretical ambiguity.[2]

The paper is organized as follows. Section 2 provides a brief literature survey on the links between income distribution and saving determination. Section 3 reviews previous empirical studies of the impact of income distribution on saving.

[2] This paper thus extends significantly, along several dimensions, preliminary findings reported in Schmidt-Hebbel and Servén (1999) which also showed little effects of income distribution on saving. Such provisional conclusion is considerably strengthened in this article based on an expanded database, presenting results that exploit both the medium and long time frequencies of the data, and performing systematic robustness tests over alternative specifications and estimation techniques.

A description of the large database on inequality and saving, used in this paper, follows. Section 5 presents new econometric evidence and Section 6 concludes.

2. Income distribution and saving: a brief survey

Aggregate saving is the outcome of individual saving efforts by heterogeneous members of different classes of savers. Heterogeneity may be due to the fact that different individuals determine their consumption/saving plans according to different objectives (i.e., their preferences are not identical). Even if all individuals possess identical preferences, their behavior may differ because they face different institutional constraints (e.g., in their access to borrowing). Heterogeneity is of course important because when agents are dissimilar, the aggregate levels of those variables that are relevant for individual saving decisions are typically not sufficient to determine aggregate saving — the latter also depends on the distribution of such variables across individual savers.[3] However, even if all agents share the same preferences and face identical constraints, distribution still matters as long as agents' (common) decision rule for saving is not linear in the relevant variables.[4]

Next, we review briefly the literature on saving and income distribution, analyzing the channels through which different forms of income distribution affect aggregate saving. We examine four topics: (i) links between saving and the functional distribution of income; (ii) links between saving and the personal distribution of income; (iii) implications of precautionary saving and borrowing constraints for distribution and saving; and (iv) indirect effects of distribution on saving.

The link between the functional distribution of income and saving (and growth) is at the heart of the neoclassical growth model (Solow, 1956), as well as the neo-Keynesian growth models of Lewis (1954), Kaldor (1957) and Pasinetti (1962). These models are general equilibrium by nature, with both saving and income distributions as endogenous variables.

In the neoclassical framework, workers and capitalists do not necessarily differ in their saving patterns. By contrast, the neo-Keynesian growth models of Lewis and Kaldor assume from the outset that workers and capitalists have different

[3] Only under very particular (and restrictive) forms of heterogeneity do aggregate consumption and saving depend exclusively on aggregate quantities. For a formal discussion of the circumstances under which the economy can be summarized by a ''representative consumer'', see Mas-Colell et al. (1995). See also Kirman (1992) for a recent sharp criticism of the representative–agent paradigm.

[4] Carroll and Kimball (1996) show that saving is non-linear for almost all commonly used utility functions as long as there is uncertainty about both labor and capital income. They also show that linearity is a special case that arises in only three particular combinations of assumptions. An example of non-linearity between saving and income is provided by Ogaki et al.'s (1996) consumption preferences with a subsistence consumption term.

saving behavior. Lewis (1954) argues that most saving comes from the profits of the entrepreneurs in the modern and industrial sector of the economy, who save a high fraction of their incomes, while other groups in the economy save less. Similarly, in the simplest form of Kaldor's (1957) model, workers spend what they earn (their propensity to save is zero) and the share of profits in national income depends positively on the investment–output ratio and inversely on the propensity to save of the capitalists. Pasinetti (1962) assumes that saving propensities differ among classes of individuals rather than among classes of income. Workers' saving is not zero; indeed, they are assumed to own shares on the capital stock and receive part of the profits. Nevertheless, the implications for the share of profits in income are the same as those obtained by Kaldor.

With consumer heterogeneity, recent neoclassical consumption theories also generate links between personal income distribution and aggregate saving that, unlike the classical theories just referred to, do not depend on the exogenous distinction between savers and non-savers. These links result from a non-linear relationship between individual saving and income.

A starting point is the life-cycle hypothesis (LCH), amended to include bequests (e.g., Kotlikoff and Summers 1981, 1988).[5] If bequests are a luxury, saving rates should be higher among wealthier consumers and richer countries, which empirically seems to be the case (Dynan et al., 1996; Carroll, 1998). In this regard, the fact that saving is concentrated among relatively few richer households, who may be accumulating mostly for dynastic motives, is in agreement with a central role of bequests in driving saving.

An alternative route through which income distribution may matter for aggregate saving was suggested by Becker (1975). If there are decreasing returns to human capital, the poor will invest relatively more in human capital than will the rich. Since human capital expenditures are considered as consumption in standard national accounting, the measured saving rates of the poor will appear lower than those of the rich, even if their ''overall'' saving rates (including human capital accumulation) are identical.

In turn, precautionary saving also implies a link between distribution and saving. Consumers with low assets tend to compress consumption to avoid running down their precautionary balances, so that their marginal propensity to consume out of wealth is higher than that of those consumers holding large asset stocks — they would devote most of any extra wealth or income to consumption (see Carroll and Kimball, 1996 for a proof). Thus, redistribution toward the poor would

[5] The view that bequests as a saving motive are perhaps more important than life-cycle considerations, and that the elasticity of bequests to lifetime resources exceeds unity, helps explain some empirical puzzles on the simple LCH model (see Deaton, 1992, 1999). One of them is that there is little evidence that the old dissave, as implied by the simple LCH; on the contrary, their saving rates appear to be as high or even higher than those of young households.

depress aggregate saving. The opposite could happen, however, if the poor face greater uncertainty, are more risk-averse, or have more limited access to risk diversification than the rich; in such circumstances, a transfer from the latter to the former would lead to higher aggregate saving.

Binding borrowing constraints — the inability of some consumers to borrow — forge another link between income distribution and saving. Given the inability to borrow, consumers use assets to buffer consumption, accumulating when times are good and running them down to protect consumption when earnings are low (Deaton, 1991).[6] If borrowing constraints affect mostly poorer households, and not the rich who hold large asset stocks, redistribution from the latter to the former makes borrowing constraints less likely to bind, thus lowering aggregate saving rates. However, while this is true in the short run, it is not likely to be the case in the long run. Buffer-stock savers have a target wealth-to-income ratio to which they will eventually return after a temporary income shock, so there should be no long-run effect on their saving level.

Finally, the recent political-economy literature brings out some indirect links between distribution and saving operating through third variables that affect saving. One line of argument is that a highly unequal distribution of income and wealth causes social tension and political instability; the result is lower investment in response to increased uncertainty, along with adverse consequences for productivity and thus growth (Alesina and Rodrik, 1994; Alesina and Perotti, 1996; Persson and Tabellini, 1994; Perotti, 1996). In addition, income distribution may affect growth also through taxation and government expenditure: in a more unequal society, there is greater demand for redistribution and therefore higher taxation, lower returns to investments in physical and human capital, and less investment and growth. From the point of view of saving, the implication is that if saving is positively dependent on growth, then higher inequality will, through the above channels, depress aggregate saving — in contrast with the positive impact of inequality on saving implied by most of the theories examined so far.

The inverse relationship between inequality and investment, suggested by the above literature, could also imply a negative association between inequality and saving through firms' earnings retention. The latter is typically the primary source of financing for private investment, so that if higher inequality lowers investment it should also reduce firm saving. What happens with aggregate saving, however, depends on whether firm owners (i.e., households) can pierce the ''corporate veil'' that separates household and firm decisions. If this was not the case, lower firm saving would not be fully offset by higher household saving, implying lower aggregate saving.

[6] However, Carroll (1997) shows that essentially the same kind of buffer-stock saving that emerges in Deaton's model can be obtained in a model without borrowing constraints.

Finally, distributive inequality may also tend to lower public saving, as governments engage more actively in redistributive expenditures — as in the populist experiences examined by Dornbusch and Edwards (1991). In the absence of strict Ricardian equivalence, this would in turn reduce aggregate saving.

3. Empirical studies

Empirical tests of the impact of income distribution on saving are rather scarce. Some early studies followed the Kaldor–Lewis approach and focused on the functional distribution of income. Along these lines, Houthakker (1961), Kelley and Williamson (1968), Williamson (1968) and Gupta (1970) found some evidence that the propensity to save from non-labor income exceeds that from labor income.

More recent empirical studies have shifted their focus from functional to personal income inequality. Blinder (1975), using US time-series data for 1949–1970, finds that higher inequality appears to raise aggregate consumption (and thus lower saving), although the estimated effect is in general statistically insignificant. He proposes as a preferable empirical test the estimation of separate consumption equations by income class, a suggestion taken up by Menchik and David (1983), who use disaggregated US data to test directly whether the elasticity of bequests to lifetime resources is larger or smaller for the rich than for other income groups. They find that the marginal propensity to bequeath is unambiguously higher for the wealthy, so that higher inequality leads to higher lifetime aggregate saving.

Data from household surveys typically show that high-income households save on average more than low-income households. Bunting (1991), who uses consumer expenditure survey data for the US, finds strong evidence that households' marginal propensities to save uniformly increases as their quintile share of income rises. Dynan et al. (1996) find some evidence that high-income US households save a larger fraction of their permanent income than poorer households do.

Several cross-country studies have used aggregate saving data from national-accounts sources. Early contributions by Della Valle and Oguchi (1976) and Musgrove (1980), using cross-country data on both industrial and developing countries, find no statistically significant effect of income distribution on saving. The exception are the OECD countries in the study by Della Valle and Oguchi, for which they find some evidence that increased inequality may increase saving. In turn, Lim (1980) finds that inequality tends to raise aggregate saving rates in a cross-section sample of developing countries, but his coefficient estimates are significant at conventional levels only in some subsamples.

Venieris and Gupta (1986) use aggregate data for 49 countries to draw inferences about the average saving propensities of different income. Their results suggest that poorer households have the lowest saving propensities, while — somewhat surprisingly — the highest one corresponds to the middle-income group. Hence, redistribution against the rich may raise or lower the aggregate saving ratio depending on whether the favored group is the middle class or the poor, respectively. However, the interpretation of their results is somewhat unclear due to their use of constant-price saving as the dependent variable, which has no clear analytical justification.

Sahota (1993) uses data on 65 industrial and developing countries for the year 1975. He regresses the ratio of gross domestic saving (GDS) to gross domestic product (GDP), GDS/GDP, on the Gini coefficient and a quadratic polynomial in per capita income, also including regional-dummy variables as a crude attempt to control for cultural and habit effects. The parameter estimate on the Gini coefficient is found to be positive, but the estimate is somewhat imprecise and significantly different from zero only at the 10% level.

More recently, Cook (1995) presents estimates of the impact of various inequality measures on the GDS/GDP ratio in 49 LDCs, using a conventional saving equation including also the level and growth rate of real income, dependency ratios, and a measure of capital inflows; the exogeneity of the latter variable is clearly questionable. A dummy for Latin American countries is also added to the regressions, although its justification is unclear since no other regional dummies are included. Using decade averages for the 1970s, he finds a positive and significant effect of inequality on saving, which appears robust to some changes in specification and to the choice of alternative indicators of income inequality.

Hong (1995) reports econometric results on the effect of the share of the top 20% income group on GDS/GDP ratios in cross-country samples of 56–64 developing and industrial countries, using 1960–1985 averages for each country. He finds that the income share of the top 20% of the population has a positive effect on saving rates, controlling for population dependency, the level and growth of income, and education level.

Lastly, Edwards (1996) estimates private saving equations using panel data for developing and OECD countries for the years 1970–1992. While the main focus of the study is not the relation between income distribution and saving, he reports two regressions, mostly for OECD countries, that are controls for income inequality. He finds that the latter has a significant positive effect on private saving if combined with one set of regressors, but negative and insignificant when combined with a different set.

In summary, most empirical studies based on micro household data find evidence of a positive effect of income concentration on household saving. Regarding the studies based on cross-country aggregate data, the results are more mixed, although some do find a positive impact of income inequality on total saving. Independently of their results, however, most cross-country studies utilize a questionable saving measure (GDS) and use income distribution data of highly heterogeneous quality, mixing both income- and expenditure-based measures; we return to these issues below. Their robustness to alternative specifications and data samples is also unclear. These drawbacks justify a more systematic empirical search of the effect of income inequality on aggregate saving across different specifications, saving and income distribution measures, data samples, and estimation techniques. This is our next task.

4. Data

Our basic information (described in more detail in Appendix A) are annual cross-country time-series data for the 1965–1994 period from the World Bank macroeconomic and social databases, and income distribution data from the database recently assembled by Deininger and Squire (1996). The latter represents a major improvement in rigor, quality, and coverage over preceding data sets — including in particular those used in the previous empirical literature on saving and income distribution.

Unless otherwise noted, here and in the rest of the paper we use the terms ‘‘income inequality’’ and ‘‘income distribution’’ for all samples and statistical results, even when they refer to Gini coefficients obtained from both income- and expenditure-based information. In order to make the Gini coefficients from income- and expenditure-based household data comparable (income is typically more concentrated than expenditure), we follow the simple procedure suggested by Deininger and Squire by adding a constant equal to 6.7 to the expenditure-based coefficients. The latter figure is the average difference between income- and expenditure-based Gini coefficients reported by Deininger and Squire, for those country-year observations for which both data are available. However, it is methodologically much less clear what type of correction should be applied to the expenditure-based income shares by quintiles. Hence, we opt for dropping them and restrict our database of income shares only to income-based data.

Likewise, we use the term ‘‘saving’’ or ‘‘saving ratio’’ to refer to gross national saving (GNS) or its ratio to gross national product (GNP). We choose national saving and national product data as the relevant variables, differing from most previous empirical studies that are based on the less-adequate GDS and GDP measures. The latter are not good saving and income measures for open economies because they exclude net income from abroad. GNS and GNP are preferable measures because they involve a broader definition of income, closer to the one whose distribution across households is captured by the income distribution data, and closer also to the income concept relevant for agents’ consumption and saving decisions. The empirical implications of this choice are examined below.

Table 1 presents descriptive statistics of our sample of income distribution indicators, which is a subsample of Deininger and Squire’s database, as discussed in Appendix A. We work with two different samples. Our cross-section sample is comprised by 52 (40) country observations on expenditure- and income-based distribution indicators (income-based indicators alone), constructed as 1965–1994 country averages. For some countries, some of the variables of interest (notably

Table 1
Income distribution indicators: descriptive statistics

	Number of observations[a]		Gini coefficient[b]		Income share ratio of top 20% to bottom 40%[b]		Income share of middle 60%[b]	
	Cross-section	Panel of 5-year averages	Mean	Standard deviation	Mean	Standard deviation	Mean	Standard deviation
World	52 (40)	270 (178)	0.417	0.085	3.33	1.56	0.48	0.07
OECD countries	19 (17)	89 (78)	0.336	0.042	2.12	0.37	0.53	0.02
Developing countries	33 (23)	181 (100)	0.464	0.066	4.22	1.50	0.43	0.05

[a] Income-based observations in parentheses.
[b] Summary statistics computed on cross-section averages. For the income shares, they have been computed on the income-based data only.

the income distribution indicators) are not available every year of the 1965–1994 period; in such cases, country averages were computed using the available observations. In addition, we use a panel data set based on 5-year averages to retain the time-series variation in the data, while reducing to the extent possible the cyclical component of the annual information — as well as to mitigate potential measurement errors present in the latter. The resulting unbalanced panel data set comprises 270 (178) observations of 5-year averages from 82 (54) countries on expenditure- and income-based distribution indicators (income-based indicators alone). Roughly one-third of the observations belong to industrial (OECD) countries and two-thirds to developing countries.

Table 1 reports means and standard deviations of three conventional indicators of inequality: the Gini coefficient, the ratio between the income shares of the richest 20% and poorest 40% of the population, and the income share of the ''middle class'', defined as the middle 60% of the population (often used as an indicator of equality).[7] The statistics reflect that developing countries are more unequal than industrial countries by any of the three indicators presented, and show also a larger dispersion in their levels of inequality.

Table 2 presents the cross-correlation patterns of saving rates and their major conventional determinants, to which we add the inequality indicators. Eight revealing features emerge, which should be kept in mind for the discussion of the multivariate regression results below.[8]

First, on this paper's core relationship — between the saving ratio and income distribution — the full-sample correlation between the saving ratio and the Gini coefficient of income distribution is in fact negative. This correlation is not sensitive to the choice of the Gini coefficient as the relevant inequality statistic — Table 2 reports very similar correlations between saving and the ratio of the top 20% to bottom 40% income shares, or the income share of the middle 60% of the population. However, the correlation pattern between saving and income concentration is rather different in industrial countries (where it is positive) from that observed in developing countries (where it is negative), although in neither case is it significantly different from zero.[9]

The similar informational content of all three income distribution indicators (noted, e.g., by Clarke, 1992) is confirmed by our second stylized fact: the

[7] For a discussion of the properties of these and other indices of income inequality, see Cowell (1977).

[8] In the subsequent discussion, we focus on the facts that are reflected by both the cross-section (lower half of the table) and the panel (upper half) correlations — they are qualitatively very similar. The levels of statistical significance can be inferred from the standard errors reported at the bottom of Table 2.

[9] These and other subsample correlations are not reported in the tables to save space.

Table 2
Correlation matrix of basic saving determinants
Notes: Cross-section correlations appear below the main diagonal; 5-year panel correlations appear above. Those involving income shares are computed on the income-based data only. Standard errors are as follows: 0.064 for the panel correlations (except for those involving income shares, whose standard error is 0.078), and 0.139 for the cross-section correlations (except for those involving income shares, whose standard error is 0.158).

	GNS/ GNP	GNP per capita	GNP per capita squared	Growth rate GNP per capita	Gini coefficient	Income share top 20%/ bottom 40%	Income share middle 60%	Old-age dependency ratio	Young-age dependency ratio
GNS/GNP	–	0.289	0.222	0.401	−0.253	−0.234	0.105	0.230	−0.429
GNP per capita	0.360	–	0.959	−0.057	−0.608	−0.514	0.693	0.891	−0.858
GNP per capita squared	0.287	0.972	–	−0.085	−0.546	−0.456	0.622	0.813	−0.759
Growth rate of per capita GNP	0.496	0.016	−0.023	–	−0.094	−0.034	−0.112	−0.037	−0.119
Gini coefficient	−0.316	−0.698	−0.653	−0.289	–	0.909	−0.867	−0.618	0.633
Income share top 20%/bottom 40%	−0.348	−0.650	−0.597	−0.235	0.936	–	−0.774	−0.535	0.582
Income share of middle 60%	0.291	0.800	0.726	0.057	−0.903	−0.850	–	0.676	−0.687
Old-age dependency ratio	0.229	0.862	0.813	0.031	−0.715	−0.553	0.734	–	−0.902
Young-age dependency ratio	−0.420	−0.885	−0.814	−0.238	0.775	0.592	−0.760	−0.940	–

correlation between the Gini coefficient and the ratio of the income share of the richest 20% of the population to that of the poorest 40% is very high, and so is that between the Gini and the middle 60% of the population.

Third, saving rates tend to rise with per capita income, an association that has been found in a number of empirical studies of saving (e.g., Collins, 1991; Schmidt-Hebbel et al., 1992; Carroll and Weil, 1994; Masson et al., 1995; Edwards, 1996). Fourth, a strong positive association is observed between saving ratios and real per capita growth, a fact also amply documented in cross-country empirical studies (e.g., Modigliani, 1970; Maddison, 1992; Bosworth, 1993; Carroll and Weil, 1994). However, its structural interpretation remains controversial, as it has been viewed both as proof that growth drives saving (e.g., Modigliani, 1970) and that saving drives growth through the saving–investment link (e.g., Levine and Renelt, 1992; Mankiw et al., 1992).[10] Fifth, demographic dependency ratios and saving rates are not unambiguously correlated. While old-age dependency and saving are positively correlated, the correlation between young-age dependency and saving is negative.

Sixth, the negative correlation between young-age and old-age dependency ratios is also very large. Seventh, dependency ratios are closely correlated with real per capita income. Finally, the correlation between income distribution measures and dependency ratios are also high.

Summing up, the world sample shows a negative association between aggregate saving rates and standard measures of income inequality, although the relationship is not robust across industrial and developing country subsamples. However, this refers only to the simple correlation between saving and income distribution. The more substantive question is whether the negative association between both variables continues to hold — or is reversed in sign — once other standard saving determinants are taken into consideration. This task is undertaken next.

5. Econometric results

We now turn to the empirical assessment of the relationship between aggregate saving and income distribution. We proceed in three stages. First, we investigate if the results of some recent studies that find a positive impact of income concentration on aggregate saving still hold when using our improved data set and alternative saving definitions. Next, we introduce a simple saving specification, including measures of income distribution along with other standard saving determinants; we estimate it on cross-country data, and test the robustness of our

[10] On the saving–growth causality, see the recent overviews by Carroll and Weil (1994), Schmidt-Hebbel et al. (1996) and Deaton (1999).

empirical results to alternative specifications, income distribution measures, country samples, and econometric techniques. Finally, we perform similar empirical tests using panel data.

5.1. Replicating previous results

For our replication, we select three recent studies focused on the effect of income inequality on aggregate saving, namely those by Sahota (1993), Cook (1995) and Hong (1995) cited above. In each case, we maintain the authors' original specification, which in the first two studies involves the use of regional-dummy variables. Our data samples are somewhat smaller in country observations than the corresponding samples in the original studies because we limit ourselves to the high-quality income and expenditure distribution data subset of Deininger and Squire's database, as discussed above. However, our data set is much larger in the time dimension and includes more recent years.

Columns 1–3 in Table 3 present our attempt at replicating these authors' results using the saving rate definition (GDS/GDP) adopted by them, and computing the individual country observations (time averages) over a time period as close as possible to that in each original study. Our results in column 1, on a sample of 45 OECD and developing countries, are very similar to Sahota's: the parameter estimate on the Gini coefficient is positive and close to that reported by the author (0.19), but well below conventional levels of significance. Regarding Cook's specification, applied to a sample of 28 LDCs, the effect of the Gini is also positive but barely reaches a 10% confidence level (column 2). Finally, in Hong's specification, applied to 50 OECD and developing countries, the relevant distributive variable is the income share of the top 20%, for which our results do replicate a positive effect significant at a 5% level.

As noted earlier, however, GDS and GDP are questionable measures of saving and income, and their use can be viewed as introducing measurement error in the saving rate (as well as the income level). In the sample, such measurement error, defined as the difference between the GDS and GNS ratios, is significantly positively correlated with the inequality measures, which strongly suggests that their estimated coefficients in columns 1–3 of Table 3 are biased upward.[11] Therefore, in columns 4–6 of Table 3, we redo the above estimations using GNS and GNP measures. The common finding across all three specifications is indeed a decline in the magnitude of the coefficient estimates on the inequality indicators, particularly in the case of Hong's specification. Although their estimated standard

[11] The correlation between this measurement error and the Gini coefficient equals 0.26 in the full sample, and 0.45 for LDCs. Likewise, the correlation between the error and the top 20/bottom 40 income share ratio equals 0.30. The other regressors are not significantly correlated with the difference between both saving ratios, with the exception of the growth rate, which shows a significant negative correlation — hence its coefficient estimates in columns 1–3 are likely biased downward.

Table 3
Replication of previous results (t-statistics in parentheses)
Note: t-statistics computed using heteroskedasticity-consistent standard errors.

Sample	Dependent variable: GDS/GDP		
	(1) Sahota (1972–1978) full	(2) Cook (1970s) LDCs	(3) Hong (1965–1994) full
Constant	−0.037 (−0.417)	0.208 (1.035)	−0.153 (−0.735)
Real GDP or GNP per capita (1987 constant dollar)	4.47E−05 (5.206)	4.24E−05 (4.757)	
Real GDP or GNP per capita squared	−2.05E−09 (−4.537)		
Real GDP or GNP growth rate		0.077 (0.136)	0.818 (1.584)
ln(GDP or GNP per capita)			0.030 (2.728)
Old-age dependency ratio		−3.365 (−2.594)	
Young-age dependency ratio		−0.211 (−0.470)	
Total dependency ratio			0.020 (0.060)
Gini coefficient	0.182 (0.938)	0.384 (1.736)	
Income share of top 20%			0.221 (2.199)
Current account balance/(GNP or GDP)		0.601 (2.838)	
Latin America regional dummy	0.088 (2.044)	−0.035 (−1.516)	
Africa regional dummy	0.153 (3.391)		
Asia regional dummy	0.131 (4.771)		
Adjusted R^2	0.369	0.556	0.401
Standard error	0.059	0.057	0.047
Number of observations	45	28	50

Dependent variable: GNS/GNP				
(4) Sahota (1972–1978) full	(5) Cook (1970s) LDCs	(6) Hong (1965–1994) full	(7) Sahota (1965–1994) full	(8) Cook (1965–1994) LDCs
0.025 (0.249)	0.020 (0.120)	0.157 (0.767)	0.048 (0.681)	0.197 (1.128)
3.68E − 05 (3.035)	4.45E − 05 (6.058)		2.27E − 05 (3.122)	8.91E − 06 (1.207)
− 1.79E − 09 (− 2.947)			− 9.28E − 10 (− 2.578)	
	0.589 (1.592)	1.421 (4.187)		0.626 (1.749)
		0.011 (0.978)		
	− 1.789 (− 1.966)			− 0.624 (− 0.689)
	0.213 (0.589)			− 0.318 (− 0.806)
		0.088 (0.895)		
0.158 (0.815)	0.244 (1.378)		0.161 (1.149)	0.363 (1.772)
		0.088 (0.895)		
	0.695 (4.995)			0.667 (2.866)
0.031 (0.577)	− 0.045 (− 2.876)		0.007 (0.181)	− 0.050 (− 3.110)
0.081 (1.873)			0.021 (0.559)	
0.088 (2.363)			0.081 (2.509)	
0.233	0.638	0.511	0.355	0.374
0.062	0.046	0.040	0.046	0.046
44	28	50	52	31

errors also decline slightly, the significance of the coefficient estimates falls well below conventional levels.[12]

Next we expand the time dimension of the country averages by applying Sahota's and Cook's specifications to our longer 1965–1994 sample period. This allows us to add more countries to the sample, while still using GNS and GNP measures. The results show that under Sahota's specification (column 7) the Gini coefficient estimate remains virtually unaffected and statistically insignificant; under Cook's specification, it rises back to reach 10% significance (column 8). However, other robustness checks (not reported to save space) on Cook's specification show that dropping the Latin America dummy, whose inclusion seems arbitrary, or the current account surplus — which, being the difference between saving and investment ratios, is clearly an endogenous variable — would make the parameter estimate on the Gini coefficient considerably smaller and statistically insignificant at any reasonable level.

We conclude that the basic finding of these three empirical studies — a positive effect of income concentration on aggregate saving — is not robust. In two cases (Hong's and Sahota's specifications), the result vanishes altogether when using improved saving and income distribution measures, while in the third one (Cook's equation) it falls short of conventional significance levels and is critically dependent on a questionable empirical specification. The natural question is whether firmer evidence on the effect of income inequality on saving can be found using more standard specifications and testing for different samples and estimation techniques. This is our next topic.

5.2. Testing alternative specifications using cross-country data

We first examine the evidence from cross-country data, using country averages for the period 1965–1994. We start from a simple specification similar to those found in comparable cross-country studies of saving (see, e.g., Schmidt-Hebbel et al., 1992; Masson et al., 1995; Edwards, 1996). It also encompasses the income, demographic and inequality variables included by Sahota (1993) and Cook (1995); but for the reasons already noted, exclude their more controversial variables (the

[12] One might argue that this discrepancy between the results based on GNS and GNP, and the earlier ones obtained with GDS and GDP, could be due instead to errors of measurement of net income from abroad, which would add random noise to the national saving rates. However, there are no clear reasons why net foreign transfers and factor payments, which are usually reflected in fairly reliable balance-of-payments information, should be measured with less precision than domestic output. Further, as noted in the text, the national income-based measures yield more precise (albeit smaller) parameter estimates than the domestic income-based ones, which seems to contradict such ''random-noise'' interpretation. Thus, we view as more plausible the alternative interpretation in the text that GDS and GDP are noisy measures of the theoretically preferable GNS and GNP.

current account balance and regional dummies). The basic equation to be estimated is the following:

$$\mathrm{GNS/GNP} = \alpha_0 + \alpha_1 \mathrm{gnp} + \alpha_2 (\mathrm{gnp})^2 + \alpha_3 \mathrm{growth} + \alpha_4 \mathrm{old} + \alpha_5 \mathrm{young} + \alpha_6 \mathrm{distrib}.$$

To recapitulate, ‘GNS/GNP’ is the ratio of current-price GNS to current-price GNP, ‘gnp’ is real per capita GNP, ‘growth’ is the (geometric) average annual rate of growth of real per capita GNP, ‘old’ is the old-age dependency ratio, ‘young’ is the young-age dependency ratio, and ‘distrib’ is an income distribution variable. This basic specification embeds both a linear and a quadratic term in real per capita income to encompass the non-linear relation between the saving rate and income commonly found in the literature; accordingly, we expect $\alpha_1 > 0$, $\alpha_2 < 0$. The majority of empirical studies suggest that the coefficient on ‘growth’ should be positive, while those on the dependency ratios should be negative, according to standard life-cycle arguments.[13] Finally, as income distribution indicators, we broaden our empirical search by using both the Gini coefficient and the ratio of the income share of the richest 20% of households to that of the poorest 40%. The latter variable, however, is available only for the smaller income-based sample.[14]

Table 4 shows estimation results using the basic specification for a variety of samples and estimation methods. As a benchmark, the first column reports OLS estimates for the full sample of OECD and developing countries excluding income distribution indicators; this simple specification accounts for nearly 50% of the observed cross-country variation in national saving rates. As expected, the second and third rows show that saving ratios rise with income levels (a result also found by Carroll and Weil, 1994 and Edwards, 1996) but taper off at high income, as indicated by the negative coefficient on squared GNP. In turn, the fourth row in the table indicates that saving ratios are positively associated across countries with per capita GNP growth rates. A 1% increase in real growth raises the national saving ratio by about 0.9 percentage points. Finally, both old- and young-age dependency ratios have the expected negative effect on national saving rates, although only the former reaches conventional levels of significance.[15]

[13] See Leff (1969) and Modigliani (1970). Gersovitz (1988) includes an analytical discussion of the effects of these and other demographic variables on saving.

[14] The empirical results obtained on the smaller subsample of income-based Gini coefficients were very similar to those reported below, which use the larger sample comprising both income-based and adjusted expenditure-based Ginis. To save space, we only report the latter.

[15] The estimated coefficients on per capita income and its square, as well as on the young-age dependency ratio, exhibit large standard errors. The reason for this lack of precision is the strong cross-correlation between age-dependency ratios, real income and its square shown in Table 2. Indeed, the null hypothesis that real income and its square have no impact on saving rates can be rejected with a p-value of 0.04; adding the restriction that young-age dependency also has no effect further reduces the p-value to 0.009, overwhelmingly rejecting the null. We attempted to increase precision by imposing equality of the coefficients on the old- and young-dependency ratios; somewhat surprisingly, however, this parameter restriction was rejected in most specifications.

Table 4
Cross-section estimates of saving equations. Dependent variable: GNS/GNP (t-statistics in parentheses)

Sample	1	2	3	4	5
	OLS[a]				
	Full	Full	Full	OECD	LDC
Constant	0.397 (2.636)	0.367 (2.375)	0.488 (2.620)	0.394 (2.214)	0.392 (2.410)
Real GNP per capita (1987 constant dollar)	1.19E − 05 (1.356)	1.20E − 05 (1.356)	6.70E − 06 (0.791)	1.01E − 05 (1.306)	6.05E − 05 (4.527)
Real GNP per capita squared	− 4.51E − 10 (− 1.134)	− 4.41E − 10 (− 1.112)	− 2.33E − 10 (− 0.656)	− 3.34E − 10 (− 0.947)	− 6.46E − 09 (− 4.778)
Real GNP growth rate	0.878 (2.132)	0.947 (2.195)	1.106 (2.061)	3.085 (3.251)	0.575 (1.247)
Old-age dependency ratio	− 1.264 (− 2.469)	− 1.244 (− 2.487)	− 1.382 (− 2.539)	− 1.181 (− 2.547)	− 2.167 (− 3.510)
Young-age dependency ratio	− 0.468 (− 1.542)	− 0.511 (− 1.666)	− 0.753 (− 1.883)	− 0.819 (− 1.919)	− 0.387 (− 1.200)
Gini coefficient		0.095 (0.864)		0.066 (0.708)	− 0.067 (− 0.471)
Income share ratio of top 20%/bottom 40%			0.005 (0.654)		
Adjusted R^2	0.423	0.419	0.491	0.555	0.569
Standard error	0.043	0.043	0.042	0.027	0.041
Sargan test (p-value)	–	–	–	–	–
Hausman test (p-value)	–	–	–	–	–
Number of observations	52	52	40	19	33

6	7	8	9
GMM[a]		Trimean	
Full	Full	Full	Full
−0.008 (−0.029)	−0.089 (−0.244)	0.332 (2.094)	0.631 (3.930)
2.79E−05 (2.813)	3.40E−05 (2.236)	8.70E−06 (1.397)	−5.00E−06 (−0.079)
−1.16E−09 (−2.596)	−1.22E−09 (−2.097)	−2.65E−10 (−0.874)	1.09E−10 (0.369)
2.463 (2.687)	2.537 (2.966)	1.307 (3.225)	1.162 (2.953)
−0.233 (−0.315)	−0.554 (−0.757)	−1.066 (−2.091)	−1.736 (−3.492)
0.402 (0.713)	0.337 (0.426)	−0.492 (−1.577)	−1.089 (−3.241)
−0.113 (−0.499)		0.127 (1.184)	
	0.011 (0.934)		0.008 (1.383)
–	–	–	–
0.048	0.048	–	–
0.527	0.759	–	–
0.596	0.988	–	–
47	39	52	40

[a] *t*-Statistics computed using heteroskedasticity-consistent standard errors. The equations using the top 20% to bottom 40% ratio (columns 3, 7 and 9) as the relevant distribution measure were estimated using only observations with income-based distribution data.

Columns 2–3 in Table 4 add the income distribution indicators. The sign pattern of the parameter estimates of the conventional variables remains unchanged, and their magnitude is rather similar to that in column 1. However, neither the Gini coefficient (column 2) nor the ratio of income shares of the top 20% and bottom 40% of the population (column 3) carry a significant coefficient. Columns 4–5 split the sample between industrial and developing countries, using the Gini coefficient as inequality indicator. The resulting estimates suggest that the saving–income and saving–growth relationships are not robust across country groups. The saving–income link is weak among industrial countries (column 4) but very strong among developing countries (column 5), while the opposite applies to the saving–growth relationship which appears strong among OECD countries but weak among LDCs — the same cross-country pattern found by Carroll and Weil (1994). The influence of demographic dependency on saving is negative in both groups, although it varies in size and statistical significance. Finally, the parameter estimate of the Gini coefficient is positive for the OECD subsample and negative for LDCs — the same sign pattern of the simple correlation mentioned earlier — but in both cases insignificantly different from zero. Interestingly, estimates (not reported) on the two subsamples using the ratio of income shares instead of the Gini as inequality indicator yield the same sign pattern and lack of significance.

One potential source of bias in these regressions is the possible endogeneity of right-hand side variables. The level, growth rate, and distribution of income are all determined jointly with the saving rate, and reverse causality from the latter variable to the rest is quite possible — for example, through the conventional saving–investment–output link mentioned earlier. In order to address this possible simultaneous equation bias, columns 6 and 7 of Table 4 report generalized method of moments (GMM) estimates on the full country sample, using instruments for the level, square, growth rate, and distribution of income.[16]

As the table shows, the sign pattern of the GMM estimates is similar to that of the OLS estimates. However, the coefficients of the instrumented GNP-related variables are more than double in size relative to their OLS counterparts, and they all become significant. This contrasts with the loss in precision of the coefficients of the dependency ratios (assumed exogenous). Finally, and most important, the two income distribution measures remain insignificant. The Sargan statistics presented in the table reveal no evidence against the validity of the instruments. As a more formal comparison between the GMM and the OLS estimates, we

[16] The instruments used were the following: initial per-capita GNP and its square; population growth; years of schooling; secondary enrollment rate; black market premium; terms of trade shocks; initial life expectancy; political instability; and civil liberties. All of these variables (except for the first two from World Bank data) were taken from Barro and Lee (1994). Unavailability of data on some of the instruments led to the loss of five observations in the full sample (one in the income-based sample).

computed Hausman statistics testing the endogeneity of the income level and its square, growth rate, and distribution indicator. As can be seen from the table, the statistics show very little evidence against the OLS estimates — although this might be partly due to the low power of the Hausman test.

Heteroskedasticity and extreme observations are common problems in cross-country empirical studies, and while the point estimates and standard errors just discussed remain consistent in the presence of heteroskedasticity (and, furthermore, the GMM estimates remain asymptotically efficient), their accuracy in small samples can be severely distorted in the presence of outlying observations. Thus, as a further check on our results, we computed robust estimators based on regression quintiles, still using our basic empirical specification. In columns 8–9 of Table 4, we present estimation results using the trimean — a linear combination of the regression median (or LAD estimator) and the 25th and 75th quintiles, with weights 0.5, 0.25 and 0.25, respectively — which outperforms the OLS estimator in the presence of even modest outlier contamination (see Koenker and Bassett, 1978 for evidence).[17] Most parameter estimates are in fact very similar to those obtained from OLS, and the income inequality indicators carry positive signs but remain insignificant. However, the specification using the ratio of income shares as inequality indicator (column 9) shows a reversed sign pattern in the per-capita GNP polynomial.

Overall, the above results indicate that neither simultaneity nor outlying observations are the cause of our failure to find significant effects of income inequality on aggregate saving. In view of this fact, we experimented with a number of other OLS specifications, using alternative income inequality measures, adding other saving regressors found in previous studies, and allowing for nonlinear effects of the income distribution indicators. The regression results (Table 5) did not yield significant coefficients for the income inequality measures in any of these alternative specifications.

To summarize, our extensive empirical tests on cross-section data find little evidence of income inequality affecting aggregate saving. It might be argued, however, that these results are based on inefficient estimators, because we have ignored the time-series dimension of the data. Further, the fact that in previous studies ad-hoc regional dummies seem to affect the significance of the inequality indicators suggests that like in other cross-country empirical work, country heterogeneity could be a potential problem in the regressions. The best way to address these concerns is by exploiting the full panel dimension of our data.

[17] The LAD estimator itself, which in general is superior to the trimean under extreme outlier contamination, but inferior in less-extreme cases, yielded virtually identical results. In both cases, the covariance matrix of the estimator was computed along the lines described by Koenker and Bassett (1982).

Table 5
Alternative OLS specifications. Dependent variable: GNS/GNP (t-statistics in parentheses)
Note: The above t-statistics were computed using heteroskedasticity-corrected standard errors.

Sample	1	2	3	4	5	6
	Full	Full	Full	Full	Full	Full
Constant	0.6037 (2.933)	0.394 (2.009)	0.361 (2.236)	0.189 (0.954)	0.266 (1.784)	0.302 (1.874)
Real GNP per capita (1987 constant dollar)	9.59E − 06 (1.161)	8.90E − 06 (1.075)	1.39E − 05 (1.301)	1.26E − 05 (1.364)	1.10E − 05 (1.379)	1.48E − 05 (1.838)
Real GNP per capita squared	− 3.39E − 10 (− 0.996)	− 3.22E − 10 (− 0.935)	− 4.64E − 10 (− 1.252)	− 4.61E − 10 (− 1.118)	− 4.50E − 10 (− 1.216)	− 5.66E − 10 (− 1.539)
Real GNP growth rate	1.154 (2.154)	1.171 (2.182)	0.958 (2.109)	0.963 (2.283)	0.941 (2.409)	0.572 (1.224)
Old-age dependency ratio	− 1.328 (− 2.463)	− 1.316 (− 2.430)	− 1.258 (− 2.593)	− 1.154 (− 2.340)	− 0.787 (− 1.476)	− 1.107 (− 2.070)
Young-age dependency ratio	− 0.741 (− 1.968)	− 0.770 (− 2.025)	− 0.512 (− 1.662)	− 0.474 (− 1.530)	− 0.315 (− 0.999)	− 0.551 (− 1.771)
Gini coefficient			0.112 (0.667)	0.863 (1.610)	0.154 (1.213)	0.249 (1.792)
Income share of middle 60%	− 0.250 (− 1.257)					
Income share of top 20%		0.2193 (1.312)				
GNP*Gini coefficient			− 4.37E − 06 (− 0.193)			
Gini coefficient squared				− 0.902 (− 1.400)		
Current account/balance GNP					0.561 (2.443)	
Latin America regional dummy						− 0.016 (− 0.428)
Africa regional dummy						0.005 (0.143)
Asia regional dummy						0.038 (1.090)
Adjusted R^2	0.508	0.512	0.406	0.417	0.411	0.463
Standard error	0.041	0.041	0.044	0.043	0.042	0.042
Number of observations	40	40	52	52	48	52

5.3. Panel data results

Table 6 reports unbalanced-panel regressions using 5-year averages for OECD and developing countries, jointly and separately. In all cases, the regressions include a full set of time dummies that were jointly significant in most cases and hence retained for ease of comparison across estimations. The first column reports simple pooled OLS estimates on the full sample, using our basic specification. The results are qualitatively similar to the cross-section OLS estimates, although the magnitude of the coefficients of the income level and growth variables is considerably smaller. In any case, the estimated impact of the Gini index remains small and insignificant. Analogous results were obtained using instead the ratio of income shares.

Column 2 adds regional dummies to the regression, thus reproducing the crude attempt of previous studies to capture (regional) heterogeneity. While the dummies themselves are insignificant individually and jointly (a test of their joint significance yields a *p*-value of 0.130, implying that they could be safely dropped from the regression), their addition does raise considerably the estimated parameter on the Gini coefficient, which becomes significant at the 10% level.

This result confirms the need to control more satisfactorily for country heterogeneity likely present in the data. To do so, in columns 3–6 in Table 6, we compute fixed-effect panel estimates; this, however, removes from the sample 27 developing countries possessing only one 5-year observation. Column 3 reports full-sample estimates using the Gini coefficient as distribution indicator. Comparison with the preceding column reveals that the addition of country effects, with themselves as highly significant (a Wald test of their joint significance yields a *p*-value below 0.001), greatly improves the overall precision of the estimates. The clear exception is the Gini coefficient, whose point estimate is now close to zero and insignificant.

These estimates also allow us to test whether the regional dummies capture satisfactorily the country heterogeneity present in the sample. This amounts to testing a set of linear restrictions (specifically, 51 of them) on the estimated country dummies from column 3. The computed Wald statistic yielded a marginal significance level below 0.001, rejecting overwhelmingly the regional-dummy in favor of the fixed-effect specification.

Column 4 is analogous to column 3 but uses the ratio of the income share of the top 20% to the bottom 40% of the population as the relevant income inequality measure — which results in the loss of another 14 countries from the sample. While for the conventional regressors the results are fairly similar to the earlier ones; the income inequality variable now carries a negative coefficient, statistically significant at the 5% level.

Columns 5–6 report fixed effect estimates using the industrial and developing country subsamples, respectively, with the Gini coefficient as inequality indicator. For the former sample, the results are rather poor, as only the income growth rate

Table 6
Panel estimates of saving equations. Dependent variable: GNS/GNP (t-statistics in parentheses)

Estimation procedure	1	2	3	4	5	6	7	8
Sample	Pooled OLS[a]		Fixed effects[a]				Panel GMM[a]	
	Full	Full	Full	Full	OECD	LDCs	Full	Full
Constant	0.545	0.467	NA	NA	NA	NA	0.408	0.324
	(6.968)	(5.033)					(1.540)	(2.672)
Real GNP per capita	4.23E − 06	6.76E − 5	2.51E − 05	1.64E − 05	4.14E − 06	6.56E − 05	2.85E − 05	1.33E − 05
	(1.005)	(1.244)	(2.357)	(1.742)	(0.549)	(2.704)	(2.040)	(0.989)
Real GNP per capita squared	−5.27E − 11	−1.40E − 10	−5.94E − 10	−3.33E − 10	6.00E − 09	−3.29E − 09	−8.75E − 10	−2.98E − 10
	(−0.276)	(−0.595)	(−2.087)	(−1.342)	(0.286)	(−2.004)	(−1.840)	(−0.570)
Real GNP growth rate	0.509	0.511	0.491	0.474	0.327	0.500	0.696	1.838
	(2.221)	(2.379)	(4.227)	(5.695)	(2.495)	(3.196)	(1.842)	(2.777)
Old-age dependency ratio	−1.744	−1.627	−2.918	−2.801	0.267	−3.58	−2.486	−1.341
	(−6.507)	(−4.801)	(−3.164)	(−3.236)	(0.472)	(−1.804)	(−4.027)	(−1.775)
Young-age dependency ratio	−0.91	−0.840	−0.573	−1.107	−0.257	−0.284	−0.586	−0.461
	(−5.986)	(−5.503)	(−1.737)	(−5.073)	(−1.312)	(−0.666)	(−1.686)	(−1.881)
Gini coefficient	0.054	0.154	0.042		0.119	0.070	0.146	
	(0.739)	(1.897)	(0.348)		(1.642)	(0.475)	(0.452)	
Income share ratio of top 20%/bottom 40%				−0.009				0.012
				(−2.013)				(1.038)
Latin America regional dummy		−0.021						
		(−0.509)						
Africa regional dummy		−0.006						
		(−0.147)						
Asia regional dummy		0.017						
		(0.482)						
Wald test of joint significance (p-value)	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Time effects (p-value)	0.015	0.027	0.305	0.031	0.000	0.399	0.451	0.009
Standard error	0.071	0.070	0.037	0.030	0.060	0.041	0.056	0.046
Sargan test (p-value)	–	–	–	–	–	–	0.097	0.591
2nd order autocorrelation (p-value)	–	–	–	–	–	–	0.301	0.116
Number of observations (countries)	248 (82)	248 (82)	221 (55)	157 (40)	86 (19)	135 (36)	124 (36)	107 (32)

[a] t-Statistics computed using heteroskedasticity-consistent standard errors.

is found significant, while the old-age dependency ratio carries a positive (albeit insignificant) coefficient. The developing-country results are much more precise, with the GNP-related variables carrying significant coefficients of the expected sign. In both subsamples, and especially in the OECD, the parameter estimates on the Gini coefficient increase somewhat relative to the full-sample estimate, but remain below conventional significance levels. Similar results (not reported) were obtained using the ratio of income shares as inequality indicator.

These panel data experiments clearly reject the regional-dummy specification as a satisfactory device to control for heterogeneity, but yield no evidence of any systematic impact of inequality on saving. Thus, as a final check, we return to the issue of simultaneity in the panel context, to re-assess if endogeneity of the regressors might be the cause of the latter result. Under adequate assumptions, panel data allow the use of ''internal'' instruments to correct for simultaneity (Arellano and Bond, 1991). Following a GMM procedure recently proposed by Blundell and Bond (1997), we reestimate our basic equation using a system framework that combines the original specification in levels with its first-differenced version.[18]

Columns 7 and 8 of Table 6 report the resulting full-sample GMM estimates. The estimation procedure summarized above requires at least three observations per country on each variable, and this unfortunately leads to a considerable decline in sample size. Regarding the conventional regressors, the GMM estimates follow the already-familiar pattern. In the specification using the Gini coefficient (column 7), parameter estimates are fairly precise, with the clear exception of the Gini itself, whose coefficient rises somewhat in magnitude but remains thoroughly insignificant. Interestingly, the time effects are insignificant as well, and other experiments reveal that dropping them would reverse the sign of the Gini coefficient estimate, although the parameter would remain insignificant. In turn, the regression including the ratio of income shares (column 8) is somewhat less precise, and the inequality variable itself is also insignificant. The Sargan statistics testing the validity of the instruments show some moderate evidence against the null for the specification using the Gini, which suggests some caution regarding

[18] We assume that the per-capita income level and growth variables, as well as the relevant income distribution indicator, are all endogenous; the demographic indicators are assumed exogenous. We use once-lagged differences of the endogenous variables as additional instruments in the level specification, and twice-lagged levels of the endogenous variables as additional instruments in the first-differenced specification. Validity of the latter instruments requires only that the residual of the first-differenced equation displays no autocorrelation of higher than first order, while validity of the former instruments requires in addition that the correlation (if any) between the country-specific effect and the regressors be time-invariant (see Blundell and Bond, 1997 for further details). These assumptions can be assessed empirically using Sargan-type and residual autocorrelation tests, as done here.

the GMM results, although the residual autocorrelation tests reveal no evidence of misspecification.

This concludes our comprehensive empirical search for the influence of income inequality on aggregate saving, controlling for other saving determinants. On the whole, we find little evidence that income concentration has any systematic impact on aggregate saving. Only exceptionally have we found significant effects — although of opposing signs: a positive one (barely significant) when using ad-hoc regional dummies (alone or in combination with another highly suspect variable, the current account deficit), and a negative one in some panel data subsamples when controlling for country-specific effects and using the ratio of income shares as the relevant inequality measure.

6. Concluding remarks

The historical growth literature and more recent neoclassical consumption theory point out various channels through which income inequality affects personal saving. Most of these mechanisms (but not all) suggest positive direct effects of income inequality on overall personal saving. However, recent political-economy research brings out negative indirect links from inequality (through investment, growth, and public saving) to aggregate saving. Taken together, these two strands of the theoretical literature imply that the overall impact of inequality on aggregate saving is ambiguous and can be assessed only empirically.

The empirical literature based on household data typically finds a positive relation between personal income inequality and overall personal saving. In turn, some empirical studies based on macro (national accounts) saving data, typically conducted on cross-country samples, also report positive effects of personal income inequality on aggregate saving. Other studies, however, find the opposite result or no effect whatsoever. Reconciling these conflicting results is difficult because macro-based empirical studies use widely different samples and specifications, different measures of saving and inequality and in most cases income distribution information of questionable quality.

This paper has reexamined the empirical evidence from macro data on the links between the distribution of personal income and aggregate saving, controlling for relevant saving determinants, providing an encompassing framework and a robustness check for previous empirical studies, and extending them in five dimensions: (i) testing alternative saving specifications; (ii) using alternative inequality and saving measures; (iii) making use of newer, better, and larger databases; (iv) conducting estimations jointly and separately for industrialized and developing countries; and (v) applying various estimation techniques on both cross-country and panel data. On the whole, we do not find any consistent effect of income inequality on aggregate saving — a result that agrees with the theoretical ambiguity.

Acknowledgements

We are grateful to Klaus Deininger and Lyn Squire for kindly making available to us their database on income distribution. We thank Aart Kraay, Steve Marglin, Branko Milanovic, Vito Tanzi, and two anonymous referees for useful comments and suggestions on an earlier version. Excellent research assistance by Wanhong Hu and Faruk Khan is gratefully acknowledged.

Appendix A

Deininger and Squire's (1996) cross-country time-series database on annual income inequality measures embodies major improvements over existing databases. A clear distinction is made between income- and expenditure-based inequality measures, as well as between household- and individual-based, and the underlying primary data are checked for important quality criteria: they have to be based on household or individual surveys (not on national accounts), their coverage has to be comprehensive (i.e., based on nation-wide samples), and measurement of income (or expenditure) has to be comprehensive as well (including all income or expenditure categories).

While the total number of country-year observations in Deininger–Squire is 2621, applying the three latter quality criteria reduces the number to 682 high-quality country-year observations, corresponding to 108 countries and years within the period 1890–1995. For these observations, both Gini coefficients and income shares by population quintiles are available. Of the latter 682 observations, we include in our subsample only those 468 country-year observations (corresponding to 82 countries) that fall into the 1965–1994 period (the one for which we have complete macroeconomic data). However, this set includes observations based on income data along with others based on expenditure data. We made the Gini coefficients from income- and expenditure-based data comparable by following the procedure described in Section 4.

For the cross-country sample, averages over 1960–1994 were used for the following variables: GDS/GDP ratio, real GDP per capita, and growth rate of real GDP per capita. For all other variables, averages over 1965–1994 were used, except where indicated otherwise. An additional requirement is imposed on the cross-country sample to achieve a minimum of time representation: countries are included only if they have at least one observation in each of the following two 15-year periods: 1965–1979 and 1980–1994. This leaves us with 52 country observations (19 industrial and 33 developing economies) for the income- and expenditure-based subsample and 40 country observations (17 industrial and 23 developing economies) for the income-based subsample.

The panel data sample comprises 5-year country averages. Unlike with the cross-country data, we expand our data set to include all countries regardless of the number of 5-year periods for which they possess observations. The number of 5-year country observations for income distribution variables (Table 1) is 270 (178) for income- and expenditure-based data (income-based data). The number of observations used in the panel regressions (Table 6) is lower because of missing data on some of the regression variables. Eighty-two countries are included in the full panel sample, of which, 20 are industrial and 62 are developing countries.

Variable definitions and sources are summarized next.

Variable	Definition and Source
GNS ratio	GNS/GNP ratio, each series in current prices and local currency. Source: World Bank National Accounts (WBNA)
GDS ratio	GDS/GDP ratio, each series in current prices and local currency. Source: WBNA
Real GNP per capita	In constant 1987 US dollars. Source: WBNA
Real GDP per capita	In constant 1987 US dollars. Source: WBNA
Growth rate of real GNP per capita	Annual growth rate computed from real GNP per capita
Growth rate of real GDP per capita	Annual growth rate computed from real GDP per capita
Gini coefficient	Source: Deininger and Squire (1996)
Ratio of income share of top 20% to bottom 40% of population and income share of middle 60%	Source: Deininger and Squire (1996)
Old age (and young age) dependency ratio	Population aged 65 and over (14 and below) as share of total population. Source: World Bank Social Indicators

References

Alesina, A., Perotti, R., 1996. Income distribution, political instability, and investment. European Economic Review 40, 1203–1228.

Alesina, A., Rodrik, D., 1994. Distributive politics and economic growth. Quarterly Journal of Economics 109, 465–490.

Arellano, M., Bond, S., 1991. Some tests of specification for panel data: Monte Carlo evidence and an application to employment equations. Review of Economic Studies 58, 277–297.

Barro, R., Lee, L., 1994. Data set for a panel of 138 countries, http://www.nber.org./barro.lee, January.

Becker, G., 1975. Human Capital. NBER, Cambridge.

Blinder, A., 1975. Distribution effects and the aggregate consumption function. Journal of Political Economy 87, 608–626.

Blundell, R., Bond, S., 1997. Initial conditions and moment restrictions in panel data models. University College London Discussion Paper 97-07.

Bosworth, B.P., 1993. Saving and Investment in a Global Economy. Brookings, Washington, DC.
Bunting, D., 1991. Savings and the distribution of income. Journal of Post Keynesian Economics 14, 3–22.
Carroll, C., 1997. Buffer-stock saving and the life-cycle/permanent income hypothesis. Quarterly Journal of Economics CXII (1), 1–56.
Carroll, C., 1998. Why do the rich save so much? In: Slemrod, J. (Ed.), Does Atlas Shrug? Economic Consequences of Taxing the Rich. Cambridge Univ. Press, London.
Carroll, C., Kimball, M., 1996. On the concavity of the consumption function. Econometrica 64 (4), 981–992.
Carroll, C., Weil, D., 1994. Saving and growth: a reinterpretation. Carnegie–Rochester Conference Series on Public Policy 40, 133–192.
Clarke, G., 1992. More evidence on income distribution and growth. Policy Research Working Paper 1064, The World Bank.
Collins, S., 1991. Saving behavior in ten developing countries. In: Bernheim, B.D., Shoven, J.B. (Eds.), National Saving and Economic Performance. University of Chicago Press, Chicago, pp. 349–372.
Cook, C., 1995. Saving rates and income distribution: further evidence from LDCs. Applied Economics 27, 71–82.
Cowell, F., 1977. Measuring Inequality: Techniques for the Social Sciences. Wiley, New York.
Deaton, A., 1991. Saving and liquidity constraints. Econometrica 59, 1121–1142.
Deaton, A., 1992. Understanding Consumption. Clarendon Press, Oxford.
Deaton, A., 1999. Saving and growth. In: Schmidt-Hebbel, K., Servén, L. (Eds.), Understanding Consumption. Clarendon Press, Oxford.
Deininger, K., Squire, L., 1996. A new data set measuring income inequality. World Bank Economic Review 10, 565–591.
Della Valle, P., Oguchi, N., 1976. Distribution, the aggregate consumption function, and the level of economic development: some cross-country results. Journal of Political Economy 84, 1325–1334.
Dornbusch, R., Edwards, S., 1991. The macroeconomics of populism. In: Dornbusch, R., Edwards, S. (Eds.), The Macroeconomics of Populism in Latin America. University of Chicago Press, Chicago.
Dynan, K., Skinner, J., Zeldes, S., 1996. Do the rich save more? Unpublished manuscript.
Edwards, S., 1996. Why are Latin America's saving rates so low? An international comparative analysis. Journal of Development Economics 51 (1), 5–44.
Gersovitz, M., 1988. Saving and development. In: Chenery, H., Srinivasan, T.N. (Eds.), Handbook of Development Economics. North-Holland, Amsterdam.
Gupta, K., 1970. Personal saving in developing nations: further evidence. Economic Record 46, 243–249.
Hong, K., 1995. Income distribution and aggregate saving, unpublished manuscript. Harvard University, Cambridge, MA, November.
Houthakker, H., 1961. An international comparison of personal saving. Bulletin of the International Statistical Institute 38, 55–69.
Kaldor, N., 1957. A model of economic growth. Economic Journal 57.
Kelley, A.C., Williamson, J.G., 1968. Household savings behavior in developing countries: the Indonesian case. Economic Development and Cultural Change 16 (3), 385–403.
Kirman, A., 1992. Whom or what does the representative individual represent? Journal of Economic Perspectives 2, 117–136.
Koenker, R., Bassett, G., 1978. Regression quantiles. Econometrica 46, 33–50.
Koenker, R., Bassett, G., 1982. Robust tests for heteroskedasticity based on regression quantiles. Econometrica 50, 43–61.
Kotlikoff, L., Summers, L., 1981. The role of intergenerational transfers in aggregate capital accumulation. Journal of Political Economy 90, 706–732.
Kotlikoff, L., Summers, L., 1988. The contribution of intergenerational transfers to total wealth: a

reply. In: Kessler, D., Masson, A. (Eds.), Modeling the Accumulation and Distribution of Wealth. Clarendon Press, Oxford.

Leff, N.H., 1969. Dependency rates and savings rates. American Economic Review 59, 886–896.

Levine, R., Renelt, D., 1992. A sensitivity analysis of cross-country growth regressions. American Economic Review 82, 942–963.

Lewis, W.A., 1954. Economic development with unlimited supplies of labor. The Manchester School 22, 139–191.

Lim, D., 1980. Income distribution, export instability and savings behavior. Economic Development and Cultural Change 26, 359–364.

Maddison, A., 1992. A long-run perspective on saving. Scandinavian Journal of Economics 94, 181–196.

Mankiw, N., Romer, D., Weil, D., 1992. A contribution to the empirics of economic growth. Quarterly Journal of Economics 107, 407–438.

Mas-Colell, A., Whinston, M., Green, J., 1995. Microeconomic Theory. Oxford Univ. Press, Oxford.

Masson, P., Bayoumi, T., Samiel, H., 1995. Saving behavior in industrial and developing countries, International Monetary Fund, Staff Studies for the World Economic Outlook. International Monetary Fund, September.

Menchik, P., David, M., 1983. Income distribution, lifetime savings, and bequests. American Economic Review 73, 672–690.

Modigliani, F., 1970. The life-cycle hypothesis of savings and intercountry differences in the savings ratio. In: Eltis, W.A., Scott, M.F.G., Wolfe, J.N. (Eds.), Induction, Growth and Trade. Oxford Univ. Press, Oxford.

Musgrove, P., 1980. Income distribution and the aggregate consumption function. Journal of Political Economy 88, 504–525.

Ogaki, M., Ostry, J., Reinhart, C.M., 1996. Saving behavior in low- and middle-income developing countries: a comparison. IMF Staff Papers 43, pp. 38–71.

Pasinetti, L., 1962. Rate of profit and income distribution in relation to the rate of economic growth. Review of Economic Studies 29, 267–279.

Perotti, R., 1996. Growth, income distribution, and democracy: what the data say. Journal of Economic Growth 1, 149–187.

Persson, T., Tabellini, G., 1994. Is inequality harmful for growth? Theory and evidence. American Economic Review 84, 600–621.

Sahota, G., 1993. Saving and distribution. In: Gapinski, J.H. (Ed.), The Economics of Saving. Kluwer Academic Publishers, Boston.

Schmidt-Hebbel, K., Servén, L. (Eds.), The Economics of Saving and Growth. Cambridge Univ. Press, London.

Schmidt-Hebbel, K., Servén, L., Solimano, A., 1996. Saving and investment: paradigms, puzzles, policies. World Bank Research Observer 11, 87–117.

Schmidt-Hebbel, K., Webb, S., Corsetti, G., 1992. Household saving in developing countries: first cross-country evidence. World Bank Economic Review 6, 529–547.

Solow, R., 1956. A contribution to the theory of economic growth. Quarterly Journal of Economics 70, 65–94.

Venieris, Y., Gupta, D., 1986. Income distribution and sociopolitical instability as determinants of savings: a cross-sectional model. Journal of Political Economy 94, 873–883.

Williamson, J., 1968. Personal saving in developing nations: an intertemporal cross-section from Asia. Economic Record 44, 194–202.