

# Income Inequality and Saving\*

FRANCISCO ALVAREZ-CUADRADO†, ‡ and MAYSSUN EL-ATTAR VILALTA†

†Department of Economics, McGill University, 855 Sherbrooke St. West Montreal, QC Canada H3A 2T7, (e-mail: francisco.alvarez-cuadrado@mcgill.ca; mayssun.el-attarvilalta@mcgill.ca)

‡CIREQ, Centre interuniversitaire de recherche en économie quantitative, Pavillon Lionel-Groulx, Université de Montréal, 3150, rue Jean-Brillant Montréal, QC, H3T 1N8

## Abstract

Over the last three decades, the average income for the bottom half of the US distribution increased by 8% while their average saving rate decreased by 8 percentage points. Over the same period, the US experienced a substantial increase in inequality and a continuous decrease in the aggregate saving rate. We propose an explanation based on interpersonal comparisons consistent with these trends. When households care about their consumption relative to others, individual saving rates decrease with reference income while aggregate saving decreases with income inequality. We provide evidence from the PSID and a panel of OECD countries consistent with these predictions.

## I. Introduction

In his seminal work *Income, Saving and the Theory of Consumer Behaviour* (1949), James Duesenberry introduces the relative income hypothesis in an attempt to rationalize the well-established differences between cross-sectional and time-series properties of saving rates. On the one hand, a wealth of studies based on 1935–36 and 1941–42 cross-sectional budget surveys present a saving ratio that increases with income. On the other hand, the time-series data on savings and income collected by Kuznets (1942) present a trend-less saving ratio. Duesenberry (1949) proposes an individual consumption function that depends not only on current individual income but also on the level of income of the reference group. As a result ‘for any given relative income distribution, the percentage of income saved by a family will tend to be a unique, invariant, and increasing function of its percentile position in the income distribution. The percentage saved will be independent of the absolute level of income. It follows that the aggregate saving ratio will be independent of the absolute level of income’ (Duesenberry, 1949, p. 3).

Despite its empirical success, see for instance Mayer (1966, 1972), the relative income hypothesis was quickly replaced by the well-known permanent income hypothesis (Modigliani and Brumberg, 1954; Friedman, 1957) as the economists’ workhorse to un-

JEL Classification numbers: D91, E21, C23.

\*We would like to thank Mutlu Yuksel for pointing us to some relevant data.

derstand saving behaviour.<sup>1</sup> According to this view the cross-sectional correlation between saving rates and income is driven by transitory deviations from permanent income, while in the aggregate, most transitory components cancel out, leading to the close relation between consumption and income observed in time series data. In an influential paper, Dynan, Skinner and Zeldes (2004) address this old empirical question. Applying new methods to a variety of saving measures derived from the Panel Study of Income Dynamics (PSID), the Survey of Consumer Finances (SCF), and the Consumer Expenditure Survey (CEX) they find a strong positive relationship between saving rates and lifetime income.

The goal of this paper is twofold. First, we identify interpersonal comparisons as one of the channels behind the documented gradient between saving rates and lifetime income. Second, we explore the impact of income inequality on this gradient. For this purpose, we propose a simple overlapping generations model of interpersonal comparisons along the lines of Alvarez-Cuadrado and Long (2011, 2012). Analyzing the saving choices of one generation, we find several testable implications. First, individual saving rates increase with lifetime income. Second, saving rates decrease with the level of income of the reference group and they do so more for poor households. Third, an increase in income inequality leads to an increase in inequality of saving rates. And fourth, once we assume that interpersonal comparisons are only upward-looking, a mean preserving spread in the distribution of labour income leads to a decrease in the aggregate saving rate.

Using several measures of saving derived from the PSID, we test the first three predictions. We find, in line with Dynan *et al.* (2004), a strong positive gradient between saving rates and income conditional on life-cycle and demographic characteristics. On average, an increase in income of 10% is associated with an increase in the saving rate that ranges between 0.4 and 0.5 percentage points depending on the measure of saving used. Approximating income of the reference group by average income in the state of residence, we find that, after controlling for household income and other individual characteristics, saving rates decrease in comparison income and do so more for poor households. For instance, our results suggest that a 10% increase in income of the reference group leads to a decrease in the total saving rate for a household in the 10th income percentile of one fifth of a percentage point. This is a large decrease, particularly if one considers that the average saving rate in the bottom decile is close to zero. When we include a measure of income inequality in our regressions, we find that those regions or periods with higher income inequality are characterized by a more unequal distribution of saving rates controlling for income. Furthermore, all these results remain robust when we instrument for permanent income using average labour earnings, different measures of consumption, and education. Most importantly, we conduct several exercises to rule out the possibility of reverse causality and provide additional evidence that the impact of inequality on saving rates that we identify is indeed due to interpersonal comparisons and not to other factors.

Finally, we test the fourth proposition constructing a panel that includes personal saving rates for six developed countries compiled from OECD statistics and measures of inequality from Alvarado *et al.* (2013). Our sample, although limited in terms of country coverage,

<sup>1</sup> One should point out that Friedman's view is more nuanced than the profession seems to believe: '... and finally, the evidence that we have cited seems to fit it (the permanent income hypothesis) somewhat better... however, this evidence is by no means sufficient to justify a firm rejection of the relative income hypothesis' (Friedman, 1957, p. 169).

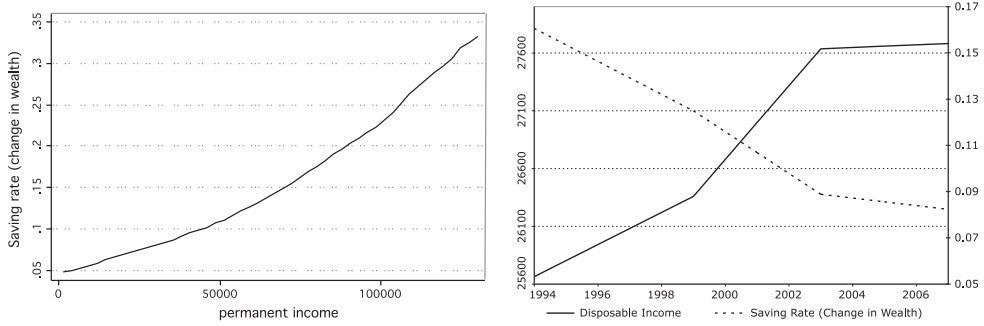
spans over 50 years and accounts for roughly 40% of world GDP. Our results consistently suggest a negative correlation between saving rates and inequality. This correlation persists when we include country fixed effects and other determinants of aggregate saving such as age dependency ratios (Modigliani, 1970), GDP growth (Carroll and Weil, 1994), or real interest rates (Boskin, 1978).

Our empirical method is closely related to Dynan *et al.* (2004). While they focus on documenting and exploring the robustness of the gradient between saving rates and permanent income,<sup>2</sup> our paper investigates one of the channels through which changes in permanent income lead to changes in saving rates, namely interpersonal comparisons, and studies the impact of income inequality on this gradient. Our theoretical model is closely related to the extensive literature on relative consumption. Abel (1990), Gali (1994), Corneo and Jeanne (1998), Carroll, Overland and Weil (2000), Liu and Turnovsky (2005) and Wendner (2014) are just a few examples. In a recent paper, De Giorgi, Frederiksen and Pistaferri (2016) use a very comprehensive Danish dataset that includes information that allows constructing networks of co-workers. They exploit this network structure to explore the impact of relative consumption. Their results suggest an elasticity of own consumption to peer's consumption of about one third. Furthermore, their data allows to explore whether these effects are driven by conspicuous (visible) consumption – status –, average consumption – 'keeping up with the Joneses' –, or risk-sharing agreements among peers. In line with the modelling strategy we will follow, their evidence suggests that peer effects are driven by average consumption. Additional evidence on the empirical relevance of relative consumption is provided by Levine, Frank and Dijk (2014) and Bertrand and Morse (2016).<sup>3</sup>

Finally, Dynan *et al.* (2004) find both a strong positive relationship between saving rates and lifetime income and evidence suggesting that the marginal propensity to save out of lifetime income is higher for rich households than for poor ones. Figure 1a illustrates these results. The rich not only save more as a fraction of their lifetime income but they also save a higher fraction of each additional dollar of income. Other things equal, an increasing marginal propensity to save has at least two direct implications. First, as households become richer their saving rates should increase. So at the individual level, we should find an

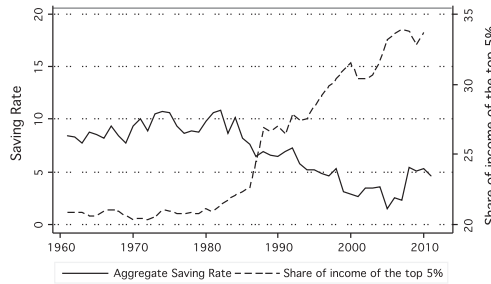
<sup>2</sup> Specifically these authors find that this gradient is not caused by life-cycle differences or variation in Social Security replacement rates (Huggett and Ventura, 2000), or by a precautionary motive together with uncertain medical expenses (De Nardi, French and Jones 2010; Kopecky and Koreschkova, 2010), or by the presence of dynastic smoothing (Becker and Tomes, 1986) or by the inclusion of high-income entrepreneurs in their sample (Quadrini, 1999; Hurst and Lusardi, 2004).

<sup>3</sup> Coibion, Gorodnichenko, Kudlyak and Mondragon (2014) perform a 'difference-in-difference' exercise across income groups and regional inequality levels finding that low-income households in high-inequality regions accumulated less debt relative to income than their counterparts in lower-inequality regions. At first sight, one may think that their results contradict ours. In our view, the results of Coibion *et al.* (2014) only suggest that the relaxation of borrowing limits is not homogeneous across regions leading to larger increases in borrowing in more equal regions. In fact, they document that credit card limits, aside from balances, increase more in low inequality areas suggesting this geographically asymmetric change in credit supply. This asymmetric change in credit supply would only attenuate the negative effect of inequality on saving rates that we document. Finally, our work also connects with the recent literature on self-reported well-being (see, for instance, Oswald, 1997; Luttmer, 2005; Clark, Frijters and Shields, 2008; Easterlin, 1995, 2010) formalizing some of their insights and exploring its implications for savings. Finally, our results are also informative for the literatures on income and consumption inequality (see, for instance, Katz and Autor, 1999; Autor, Katz and Kearney, 2008; Attanasio, Hurst and Pistaferri, 2012) and saving and wealth inequality (De Nardi and Fella, 2017).



(a) Savings across the income distribution.

(b) Evolution of income and savings for households with low income.



(c) Evolution of the household saving rate and inequality in the US.

Figure 1. Motivation figures

Notes: (a) Saving is the change in wealth over a period. Permanent income is the average labour income of the household over the period in which the household is in the sample (in 2007 dollars). (b) Saving rate is the average change in wealth over a period for households below the median of permanent income divided by their average disposable income. Disposable income is the average after tax household income in 2007 dollars for households below the median of permanent income. The figure reports five-year moving averages except 1999, which is an average of 1994, 1999 and 2003. (c) Inequality is measured as the share of income of the top 5% of the distribution from the update of Piketty and Saez (2003). The personal saving rate is obtained from the Bureau of Economic Analysis.

increase in saving rates as permanent income increases. Second, if we redistribute income from households with a low marginal propensity to save to those with a high marginal propensity to save the aggregate saving rate should increase. So at the aggregate level, we should find that increases in income inequality should be associated with increases in the average saving rate. Using PSID data from 1994 to 2007, Figure 1b depicts the evolution of average income and the average saving rate for households below the median of the US (lifetime) income distribution. Over this period, their average income increased by 8% while their average saving rate fell by almost 8 percentage points.<sup>4</sup> Figure 1c combines the evolution of the share of income of the top 5% of the distribution from the update of Piketty and Saez (2003) with the personal saving rate from the Bureau of Economic Analysis. In the early 80s, personal saving was roughly 10% of disposable income falling to barely 2%

<sup>4</sup> Asimilar pattern has been reported by Bosworth and Anders (2008). These authors show that, once some extreme values are trimmed, average saving by individuals in the lower third of the income distribution has fallen over the period 1984–2005.

by 2007. Over the same period the share of income of the top 5% increased by more than 50% from 21% to 33%.

Our model provides a simple explanation combining all three trends. In the presence of interpersonal comparisons, individual saving depends not only on individual lifetime income but also on the lifetime income of the reference group. As a result and despite the 8% increase in lifetime income of households below the median, their saving rates have fallen as a consequence of the much faster increases in income for those in the upper tail of the distribution. For instance, the average income for the top decile increased by roughly 38% over the same period. At the aggregate level, and despite the fact that the model's marginal propensity to save is independent of the level of income, an increase in the share of income received by high-income households induces a reduction in the levels of saving of low-income households that goes beyond this propensity. The reason lies in the increase in consumption of high-income individuals that stimulates envious consumption further depressing saving for low-income individuals. As a result, increases in income inequality are associated with decreases in the aggregate saving rate.

The paper is organized as follows. In section II, we set out the basic model and derive its implications for the interaction among saving, income and inequality. In section III, we use data from the PSID to test these implications. In section IV, we extend the model to allow for a upward-looking comparisons and explore its implications for aggregate saving. These implications are tested in section V. In section VI we offer some concluding remarks, while the Appendices provide details on the derivations, the data, and some additional empirical results.

## II. Individual saving and income inequality

Consider a simplified version of the OLG economy presented in Alvarez-Cuadrado and Long (2011, 2012) that abstracts from bequests and prices – the wage per unit of effective labour,  $w_t$ , and the gross return to saving,  $R_t$  – are exogenously determined. Individuals live for two periods, 'youth' and 'old-age'. At the end of their youth, each individual gives birth to a single offspring and therefore the labour force is constant. Each generation consists of a set of agents, indexed by  $i$ , distributed uniformly along the unit interval with a total population mass of 1.

Individuals within a given generation differ in their endowment of productive labour,  $l_t^i$ . Specifically, we assume that their labour productivity is the realization of a stationary i.i.d. random variable with mean  $\bar{l} = 1$  and standard deviation  $\sigma^l$ . The resulting distribution of wages for period  $t$  has mean  $\bar{w}_t = w_t$  and standard deviation  $\sigma_t^w = w_t \sigma^l$ .

Let us focus on the  $i$ -th individual born in period  $t$ . When young, he inelastically supplies his labour endowment, earning an income  $w_t^i = l_t^i w_t$  that is divided between first-period consumption,  $c_t^i$ , and saving,  $s_t^i$ . His first-period budget constraint is given by,

$$c_t^i + s_t^i = w_t^i \equiv y_t^i, \quad (1)$$

where we denote by  $y_t^i$  his lifetime resources, which coincide with labour income.

In the second period of his life, the individual is retired. His only source of income is the return on the savings he made when young,  $R_{t+1} s_t^i$ , which he allocates to second-period consumption,  $d_{t+1}^i$ .

$$R_{t+1}s_t^i = d_{t+1}^i \tag{2}$$

The preferences of an individual born in period  $t$  are given by the following life-cycle utility function,

$$U(\hat{c}_t^i, \hat{d}_{t+1}^i) = \ln(c_t^i - \gamma \bar{g}_t^i) + \beta[\ln(d_{t+1}^i - \eta \bar{o}_{t+1}^i)] \tag{3}$$

where  $0 < \beta < 1$  is the subjective discount factor.

Our key behavioural assumption is that the satisfaction derived from consumption does not depend on the absolute level of consumption itself but rather on how it compares to the consumption of some reference group. For the moment,  $\bar{g}_t^i = \bar{c}_t = \int_0^1 c_t^j dj$  and  $\bar{o}_{t+1}^i = \bar{d}_{t+1} = \int_0^1 d_{t+1}^j dj$  are the reference consumption levels when young and old respectively for any household born at time  $t$ . The importance of positional concerns when young and old is given  $0 < \gamma < 1$  and  $0 < \eta < 1$  respectively.<sup>5</sup> Concerning the relative magnitude of  $\gamma$  to  $\eta$ , one may consider the work of development psychologists and sociologists (Coleman, 1961; Simmons and Blyth, 1987; Corsaro and Eder, 1990) which suggests that interpersonal comparisons and peer effects are more pronounced early in life. In the ‘first-period’ of life, people work, find partners, and raise children, being exposed to and therefore influenced by a wide variety of social networks. More direct evidence comes from Charles, Hurst and Roussanov (2009) who find that differences in visible consumption disappear when they restrict their sample to older households, suggesting that the relative importance of interpersonal comparisons decreases with age. In line with this evidence, we assume the following relationship between the degrees of envy in both periods of life,  $\eta = \xi\gamma$ , where  $0 < \xi \leq 1$ . Finally, we place restrictions on the stationary distribution of productive endowments to guarantee that everyone’s relative consumption is positive.

**Saving and inequality**

The  $i$ -th individual of the generation born in period  $t$  takes prices and the choices of the other members of his generation as given and chooses consumption in both periods of his life to maximize (3) subject to equations (1) and (2). The necessary conditions for this problem imply

$$\frac{(d_{t+1}^i - \xi\gamma\bar{d}_{t+1}^i)}{(c_t^i - \gamma\bar{c}_t)} = \beta R_{t+1} \tag{4}$$

with the standard interpretation that along an optimal solution the  $i$ -th agent equates the marginal rate of substitution between first and second period choices to its relative price, the interest rate adjusted for the rate of time preference. We begin by characterizing the behaviour of the average household, i.e. the household with average income,  $\bar{y}_t = \bar{w}_t$ , then we turn to the  $i$ -th individual of the same generation whose choices are given by<sup>6</sup>

<sup>5</sup> As Frank (1985, p. 111) points out, ‘the sociological literature on reference group theory stresses that an individual’s personal reference group tends to consist of others who are similar in terms of age’. Along these lines our specification restricts interpersonal comparisons to individuals within the same generation. See Abel (2005) and Alonso-Carrera, Caballe and Raurich (2008) for an alternative specification of the reference group.

<sup>6</sup> See Appendix S1 for a detailed derivation.

$$s_t^i = \frac{d_{t+1}^i}{R_{t+1}} = \frac{\beta}{1 + \beta} (y_t^i - \gamma(1 - \xi)\chi\bar{y}_t), \quad (5)$$

$$c_t^i = \frac{1}{1 + \beta} [y_t^i + \gamma\beta(1 - \xi)\chi\bar{y}_t], \quad (6)$$

where  $\chi \equiv 1/(1 + \beta - \gamma(\xi + \beta))$ .

First-period consumption of the  $i$ -th household, equation (6), is made out of two components. The first increases in the household's lifetime income. The second reflects the influence of interpersonal comparisons and increases in the lifetime income of his reference group. As a result saving depends on relative, rather than absolute, income. When individual  $i$ 's satisfaction depends on consumption comparisons across households the relevant variable driving saving choices is the comparison between individual  $i$ 's lifetime income and the lifetime income of his reference group. The agents populating our economy are not only 'disposed, as a rule and on the average, to be forward-looking animals' as those in Modigliani and Brumberg (1954, p. 430) or Friedman (1957), but are also *outward-looking animals* as their choices are partially driven by the choices of other members of the community they live in.

The following propositions summarize the implications of the model for saving rates and their interaction with income inequality.

*Proposition 1.* Under our assumptions, the (individual) saving rate,  $s_{\text{rate}}^i$ , increases with lifetime income.

$$\frac{\partial s_{\text{rate}}^i}{\partial y^i} = \frac{\beta\gamma(1 - \xi)\chi}{1 + \beta} \frac{\bar{y}}{(y^i)^2} > 0$$

*Proof.* Equation (5) implies the following saving rate out lifetime income,

$$s_{\text{rate}}^i \equiv \frac{s^i}{y^i} = \frac{\beta}{1 + \beta} \left( 1 - \gamma(1 - \xi)\chi \frac{\bar{y}}{y^i} \right) \quad (7)$$

The result follows from differentiation of equation (7).

In the absence of interpersonal comparisons,  $\gamma = 0$ , saving is proportional to lifetime income and therefore, as in the standard version of the permanent income hypothesis, the saving rate is independent of lifetime resources. Once we allow for consumption externalities,  $\gamma > 0$ , the saving rate varies across the income distribution with poor households saving a smaller fraction of their resources than rich households. The introduction of interpersonal comparisons diverts resources from less positional uses – savings – to more positional ones – consumption when young. Furthermore, this diversion varies across the income distribution being stronger for poor households than for rich ones as the following proposition illustrates.<sup>7</sup>

<sup>7</sup> There are two intuitive interpretations of these results. First, along the lines suggested by Alvarez-Cuadrado and Long (2012), in the presence of interpersonal comparisons the income elasticities of demand for consumption and saving vary across the income distribution. See their proposition 1 for details. Second, in line with the work of Atkeson and Ogaki (1996, 1997), the introduction of interpersonal comparisons induces an intertemporal elasticity of substitution (IES) that increases with income. To illustrate this point assume that  $\xi = 0$ . The Euler equation for household  $i$  is given by  $\frac{d_{t+1}^i}{c_t^i} = \beta \left( 1 - \gamma \frac{\bar{c}_t}{c_t^i} \right) R$  where the term in brackets is proportional to the IES.

*Proposition 2.* Under our assumptions, individual saving rates decrease with the level of income of the reference group and they do so more (less) for poor (rich) households.

*Proof.* Differentiate equation (7) to reach  $\frac{\partial s_{\text{rate}}^i}{\partial \bar{y}} = -\frac{\beta\gamma(1-\xi)\zeta}{(1+\beta)y^i} < 0$ .

Finally, we can use this simple model to explore the impact of income inequality on individual saving. A convenient measure of the dispersion of the distribution of lifetime income is the coefficient of variation, i.e. the ratio between the standard deviation and the mean of the wage distribution,  $CV(w^i) = \frac{\sigma^w}{\bar{w}}$ . We combine the distribution of wages with equation (5) to derive the coefficient of variation of saving,

$$CV(s^i) = \Theta \frac{\sigma^w}{\bar{w}}, \quad (8)$$

where  $\Theta \equiv (1 + \beta - \gamma(\xi + \beta))/(1 + \beta - \gamma(1 + \beta)) > 1$ . So, an increase in income inequality leads to a more than proportional increase in saving inequality. Nonetheless, we are interested in the relationship between income inequality and saving rates. The following proposition illustrates this relationship.

*Proposition 3.* An increase in income inequality leads to an increase in inequality of saving rates.

*Proof.* Applying the Delta method to equation (7), we approximate the standard deviation for the saving rate by  $\sigma^{\text{rate}} \approx \frac{\beta\gamma(1-\xi)\zeta}{1+\beta} \frac{\sigma^w}{\bar{w}}$  and therefore the coefficient of variation of the saving rates is given by,

$$CV(s_{\text{rate}}^i) \approx \frac{\gamma(1-\xi)}{(1+\beta)(1-\gamma)} \frac{\sigma^w}{\bar{w}}. \quad (9)$$

As opposed to the standard version of the permanent income hypothesis,  $\gamma = 0$ , where saving rates are independent of the level of income and therefore unaffected by the degree of income inequality, in the presence of interpersonal comparisons,  $\gamma > 0$ , changes in the distribution of income lead to changes in the distribution of saving rates. Specifically an increase in income inequality is associated with an increase in the coefficient of variation of saving rates.

*Remark to Proposition 3.* The aggregate (average) saving rate is independent of the income distribution.

Since preferences are quasi-homothetic and capital markets are perfect saving is an affine function of income, a property that ensures that the distribution of income does not affect aggregate outcomes.<sup>8</sup> Along this dimension, our modelling approach is similar to that of Chatterjee (1994) and Caselli and Ventura (2000).

In the next section we will use PSID data to test these propositions.

<sup>8</sup>Notice that the timing of income is equivalent to perfect capital markets.



### III. Empirical Analysis: Individual Saving and Inequality

#### Empirical method

There are several issues that we need to address before testing the empirical implications of the model presented in the previous section.

The first issue concerns the appropriate definition of saving. We use three different measures that capture different aspects of savings behaviour. First, we use a broad measure of saving, the change in wealth between two consecutive waves of the PSID. This measure incorporates information on housing, other real estate, financial assets and owned businesses, including unrealized capital gains while the asset remains in the household's portfolio and realized capital gains once it is sold. Second, we add estimates of saving through Social Security and private pensions to the change in wealth to obtain our most comprehensive measure of saving. Finally, we compute a measure of active saving as the difference between income, including realized capital gains but not unrealized ones, and consumption.

Second, we would like to recover the effect of permanent income on savings, but we do not observe permanent income. This is the key problem tackled by Dynan *et al.* (2004) and we will follow their approach. The PSID reports current income which is a noisy measure of permanent income since it includes transitory components of income and possibly some measurement error. We deal with this problem using consumption, food consumption, years of education of the head of the household and average labour income as instruments for permanent income. As pointed out by Dynan *et al.* (2004), different instruments may suffer from different problems. For instance, to be able to use average labour income we need to assume that the transitory component in income is uncorrelated over time. One could also think that education could have a direct effect on saving decisions (Lusardi, 1999), and this is why we include years of education of the head of the household as an explanatory variable in our regressions and only exclude it when we use education as an instrument. Furthermore instrumental variables provide an estimate for a specific group, those individuals whose behaviour is affected by the instrument (Imbens and Angrist, 1995). Since not every individual is affected by the instruments in the same way, the estimated coefficients will vary according to the instruments used. Because of the different potential problems associated with each instrument and the heterogeneous treatment effects, we perform our empirical analysis using one instrument at a time. Our goal is to show that the predictions of the model are robust to changes in specification.

Since current income, rather than permanent income, is also used in the construction of the saving rate, similarly, although less acute, problems arise. We reduce the impact of transitory income on the measured saving rate by computing savings over multi-year intervals and using averages of current income over these same periods. As a result, the remaining transitory income component or the presence of classical measurement error in our dependent variable is unlikely to cause a bias in our estimates.<sup>9</sup>

<sup>9</sup> Notice that the effect of the transitory component of income in the measures of the saving rate will tend to induce a negative correlation between saving rates and income biasing the coefficient of income downward. The concerns discussed so far were also faced by Dynan *et al.* (2004). We refer the interested reader to their paper for a more detailed discussion.

Two other key variables for testing the propositions derived from our model are reference income and inequality. We measure income of the reference group as average income in the state of residence and we use the coefficient of variation of disposable income in the state of residence as our measure of inequality. Both variables vary across states and over time.

In contrast to Dynan *et al.* (2004), whose main focus is on the identification of the causal effect of income on saving rates, our focus is on the impact of reference income and inequality on saving rates and on the saving-income gradient. A causal interpretation of our results requires ruling out the possibility of reverse causality. One concern about reverse causality arises from the fact that saving rates vary systematically with income. Since income includes both labour and capital income, a portion of the current changes in income, and therefore in income inequality, is likely the result of past variation in saving rates. To deal with this concern, we show that results are similar when we use a measure of inequality calculated using labour income, which excludes capital income and therefore is independent of past saving choices.<sup>10</sup> The existence of factors that simultaneously affect both the dependent and independent variables could also affect the causal interpretation of our results. To deal with this, we include state dummies in the regression to capture potential unobserved drivers of differences in saving across states (e.g. cultural differences, say between Minnesota and California). To rule out a potential relationship between saving and inequality driven by business fluctuations we also include time dummies.<sup>11</sup>

Once the causal nature of our relationships of interest is established, we will provide additional evidence that they are indeed due to interpersonal comparisons and not to other factors. Even if the relationship between inequality and saving is causal, it could still be driven by factors other than interpersonal comparisons. For instance, differences in inequality across states may be associated with differences in redistributive policies that also affect saving. In section IV, we address this concern by including political variables in the regression. Moreover, we verify that our results are not driven by interstate variation in housing prices that may be caused by differences in inequality across states (Matlack and Vigdor, 2008). In addition to this, results on reference income and saving also provide evidence on the role of interpersonal comparisons.

Finally, all our regressions include as controls several factors that may affect saving (or consumption) decisions. We include age and age squared of the head of the household to control for life-cycle variation in saving rates (Huggett and Ventura, 2000), years of education (Lusardi, 1999), and a measure of household size (Attanasio and Browning, 1995).

For most of our results we estimate a median regression with the savings rate as our dependent variable, and a measure of income, reference group income or regional inequality, and the other controls as explanatory variables.<sup>12</sup> When we implement the IV estimation

<sup>10</sup> Education is a particular type of saving and labour income reflects the returns to educational investment. Since education is not included in our measure of savings, reverse causality through this channel does not seem to be an issue. Nonetheless, in all our regressions we control for years of education.

<sup>11</sup> Furthermore, Runkle (1991), using PSID data, finds that aggregate shocks are not very important for individual consumption (and therefore saving) choices.

<sup>12</sup> See the distribution of saving is very skewed, see Table 1, an OLS regression, which estimates the effect at the mean, would estimate the effect of the independent variables in the upper part of the distribution of the saving rate.

to recover the effect of permanent income, we use a two-stage procedure. Standard errors are clustered at the household and state level and bootstrapped using 500 replications.

## Data

We obtain all our information from the PSID, which is unique in providing a measure of assets for a panel of households. We use eight wealth supplements covering the years 1984, 1989, 1994, 1999, 2001, 2003, 2005 and 2007.<sup>13</sup> Using the panel dimension of the data, we estimate each household's change in wealth and active saving. As noted by Bosworth and Anders (2008), very wealthy households are underrepresented in the PSID, but for the bottom 95% of the wealth distribution, the PSID wealth data yields very similar results to data from the Survey of Consumer Finances (SCF) where this sample bias is not an issue.

We select our sample using the same criteria as Dynan *et al.* (2004). We restrict the sample to households with a head between 30 and 60 years of age, measured at the middle of the preceding five-year period. As a result we avoid dealing with issues relevant to very young households, such as liquidity constraints, and to very old ones, such as retirement or acute health problems. We select households with an absolute value of active saving not larger than \$1,050,000 and with disposable income of at least \$1,400.<sup>14</sup> From the remaining sample, we drop households in the extreme 0.5% of the distribution of saving. We also drop observations for households that had a change in head. All the variables are adjusted for inflation using the price deflator for personal consumption expenditures and are reported using 2007 dollars. We end up with a sample of 24,036 observations about 8,063 households, with a minimum of 2,815 households in any given year.

For ease of comparison, we construct saving rates following the procedure outlined by Dynan *et al.* (2004). For our first measure of saving,  $\Delta\text{Wealth}$ , we use information on the change of eight asset categories net of their associated liabilities. These assets are home equity, other real estate, private business/farm, vehicles, transaction accounts, corporate equities, annuities/IRAs, and other savings. We also include information about non-collateralized debt. The resulting measure of saving is very broad.

Our second measure extends the previous one to include contributions to Social Security and company pension schemes,  $\Delta\text{Wealth} + \text{pensions}$ . We include this as a separate measure since not all of these contributions result from individual choices. To calculate these contributions, we add PSID data on employee's contributions to private pension plans and expected Social Security benefits.<sup>15</sup> We compute these benefits subtracting net marginal Social Security tax rates reported by Feldstein and Samwick (1992) and Cushing (2005) from gross Social Security tax rates. While the gross Social Security tax rate is easy to compute – 11.2% (12.5% after 1995) of labour income up to \$50,000 – the net marginal tax rate (and the expected benefits) varies with income, age and marital status of the tax

The median regression allows us to estimate the effect at the median of this distribution, where most of the mass is concentrated.

<sup>13</sup> After 2007 the wealth supplement of the PSID was discontinued.

<sup>14</sup> These figures, expressed in 2007 dollars, are equivalent to those in Dynan *et al.* (2004), expressed in 1994 dollars.

<sup>15</sup> Unfortunately the PSID does not include employer's contributions to private pension plans. Since these contributions are mostly concentrated in high-income households, this omission will tend to reduce the gradient between saving rates and income.

payer. For instance, in 1994, a household in Tennessee with four members whose head is 35 years old and earns the median income, around \$46,000, faced a net marginal Social Security tax rate of 5.59%. As a result, his Social Security savings (or expected benefits) were 5.61% of his labour income.

The third measure, *Active Saving*, corresponds to the difference between income, including realized capital gains but not unrealized ones, and consumption. This measure includes net investments in private pensions and annuities, home equity, businesses, vehicles, stocks, cash assets and other financial assets plus redemption of other debt. It is also adjusted for transfers of assets in and out of the household.<sup>16</sup>

For the denominator of the saving rate, we use either five-year or two-year averages of household income, depending on the periodicity of the PSID wealth supplements. Household income is total income after taxes and transfers including income from labour earnings, asset flows, private transfers, private pensions, public transfers and Social Security pensions minus total taxes. When the numerator of the saving rate includes unrealized capital gains or pensions, its denominator, the measure of income, is adjusted accordingly. For each household, the saving rate is calculated as average saving over a multi-year period divided by average income over the same period. Finally, we refer to the three saving rates associated to these measures of saving as *total saving rate*, *total saving rate plus pensions* and *active saving rate*, respectively.

Table 1 presents summary statistics for each of these variables. To reduce the influence of observations with very low income, the reported mean is computed dividing average annualized saving by average income over the period. Since the income measures used to compute the different saving rates differ, median income differs across measures. The reported median and mean saving rates suggest that the measures of saving differ in breadth and that the distribution of saving rates is skewed to the right. Comparing rows 1 and 2 one can see that the inclusion of pensions increases the saving rate for the median saver by 4 percentage points. The summary statistics for the distribution of saving rates are similar to those reported by Dynan *et al.* (2004) and Bosworth and Anders (2008).

Finally, this table also presents summary statistics for households' disposable income and the other variables included in our regressions. We report different measures of income that will be used at different stages of our analysis. Average labour income is computed using all observations on labour income for the time period that a household remains in the sample (1968–2007). On average, this measure includes fourteen observations of yearly income.

### **Saving rates, current income and inequality**

All the regressions include age, age squared and years of education of the head of the household, size of the household, and year and state dummies as controls. Although not reported,

<sup>16</sup> Cronqvist and Siegel (2015) use a similar measure of active savings but theirs excludes realized capital gains in the household's primary home, implicitly treating this home as consumption. In a similar spirit, Browning and Leth-Petersen (2003) use administrative data on income and wealth to impute consumption expenditures. See the Data Appendix for a more detailed description of the construction of our saving measures.

TABLE 1  
Summary statistics

<i>Saving</i>	<i>Median</i>	<i>Mean</i>	<i>Median income</i>
Total saving rate	0.07	0.21	47,115
Total saving rate + pension	0.11	0.24	49,323
Active saving rate	0.04	0.12	46,199
<i>Income</i>			
Disposable income (after tax)	46,526	56,497	54,137
Disposable income (pre tax)	56,332	72,093	83,346
Labour income (pre tax)	53,510	69,304	82,824
Avg. labour income	52,288	61,570	56,512
<i>Demographics</i>			
Head's years education	12	13.09	2.44
Head's age	45	45.53	8.05
Household size	3	3.02	1.52

*Notes:* All saving and disposable income measures are expressed in 2007 dollars. Measures of income used to compute the saving rates:  $\Delta$ Wealth: disposable income + capital gains from housing, other real estate, business and stocks,  $\Delta$ Wealth + pensions: disposable income + capital gains + employer contributions to Social Security and pensions. Median saving rate equals median of the ratio of saving to income. Average saving rate equals average saving divided by average income.

their associated coefficients are statistically significant throughout all our specifications.<sup>17</sup> In Table 2, we test the relationship between income and the saving rate using our three different measures of saving. The first two columns show that the total saving rate increases with income. When we introduce a quadratic income term the relationship is convex. This is consistent with the findings of Dynan *et al.* (2004) who capture this nonlinearity by allowing for quantile specific slopes. Column 1 indicates that an increase in income of 10% translates into an increase in the saving rate of half of a percentage point. Hence, increasing income of a household roughly at the median of the income distribution (\$46,000) by \$20,000 increases its saving rate by 2.1 percentage points from about 10% to about 12%.<sup>18</sup> Using the estimates of Dynan *et al.* (2004), this same household would increase its saving rate by 2.2 percentage points.<sup>19</sup> When we extend our saving measure to include Social Security benefits and private pensions, saving rates become slightly less responsive to income, but the relation remains positive and significant. Combined with a larger constant (not reported), this indicates that the inclusion of Social Security raises the saving

<sup>17</sup> Furthermore, the coefficients on the socio-demographic controls have the expected signs. The effect of age is hump-shaped. The effect of education is positive, while that of household size is negative.

<sup>18</sup> The saving rate is computed for a 45-year old head of a household with 4 members who has 12 years of education and lives in Oregon.

<sup>19</sup> When using the PSID, Dynan *et al.* (2004) find that increasing a household's income by 10,000 1982 dollars (about 20,000 2007 dollars) implies a change in the saving rate of 2.2 percentage points.

TABLE 2  
Saving rates and income

	$\Delta\text{Wealth}$		$\Delta\text{Wealth} + \text{pensions}$		Active	
	(1)	(2)	(3)	(4)	(5)	(6)
ln(income)	0.049*** (0.002)	-0.418*** (0.032)	0.038*** (0.002)	-0.369*** (0.044)	0.039*** (0.002)	-0.299*** (0.028)
ln(income) <sup>2</sup>		0.023*** (0.002)		0.020*** (0.002)		0.017*** (0.001)
Pseudo- <i>R</i> <sup>2</sup>	0.010	0.012	0.010	0.010	0.001	0.001
<i>N</i>	24,036	24,036	24,036	24,036	24,036	24,036

Notes: The column heading indicates the saving measure used.  $\Delta\text{Wealth}$  refers to the change in wealth or total saving. Active saving is defined as total saving minus capital gains. Income is current after tax household income in 2007 dollars. All regressions also include controls for age, age squared, household size, head's education and year and state dummies which are not reported. Clustered standard errors are reported in parentheses. \*\*\*Significant at less than 1%; \*\*significant at 5%; \*significant at 10%.

rate of the poor, weakening the saving-income relationship.<sup>20</sup> The relationship between saving and income remains positive and significant, although marginally weaker, when we consider active savings, which excludes unrealized capital gains.

Proposition 2 states that individual saving rates decrease with the level of income of the reference group, and more so for poor households. To test this proposition, we regress the saving rate on own income, our measure of income of the reference group – average state income – and an interaction between the two variables. This interaction aims at capturing the differential impact of reference income across the income distribution. For ease of interpretation, we demean log average state income by the population mean of this variable. Our theory predicts that the coefficient on income of the reference group should be negative and the coefficient on the interaction term positive. Results from this regression are reported in the first three columns of Table 3. The effect of own income is positive. Considering an individual with income and regional income equal to the median of our sample, we find that the estimated marginal effects of income are similar to those reported in Table 2. The sign of the marginal effect of average state income varies across the income distribution. It is negative for poor households and positive for richer ones. This positive effect on the saving rates of richer households, particularly those with incomes above the mean, may be driven by the fact that average state income is an inadequate measure of their reference income. In the next section, once we instrument for permanent income, we will consider different reference groups for different individuals, all involving upward-looking comparisons, and show that Proposition 2 holds for a much larger fraction of the income

<sup>20</sup> Dynan *et al.* (2004) find that adding estimates of saving through Social Security and private pensions to the total saving rate steepens the saving-income gradient. This difference might arise from the fact that their estimates of private pensions are more accurate than ours since they include employer's contribution to private plans.

TABLE 3  
Saving rates, income, reference group and inequality

	Reference group			Inequality		
	$\Delta$ Wealth	$\Delta$ Wealth + pensions	Active	$\Delta$ Wealth	$\Delta$ Wealth + pensions	Active
	(1)	(2)	(3)	(4)	(5)	(6)
ln(income)	0.049*** (0.002)	0.039*** (0.002)	0.042*** (0.002)	0.037*** (0.005)	0.028*** (0.005)	0.030*** (0.005)
ln(state income)	-0.969*** (0.099)	-0.902*** (0.107)	-0.712*** (0.103)			
ln(state income) × ln(income)	0.095*** (0.010)	0.087*** (0.008)	0.073*** (0.010)			
Inequality				-0.124** (0.053)	-0.117** (0.051)	-0.104** (0.052)
Inequality × ln(income)				0.011** (0.005)	0.010** (0.005)	0.010* (0.005)
Pseudo- $R^2$	0.011	0.010	0.001	0.011	0.010	0.001
$N$	24,036	24,036	24,036	24,036	24,036	24,036

Notes: The column heading indicates the saving measure used.  $\Delta$ Wealth refers to the change in wealth or total saving. Active saving is defined as total saving minus capital gains. Income is current after tax household income in 2007 dollars. State income is the state level average of this measure. The regression includes demeaned log state income. Inequality is measured as the coefficient of variation of disposable income per region an year. All regressions also include controls for age, age squared, household size, head's education and year and state dummies which are not reported. Clustered standard errors are reported in parentheses. \*\*\*Significant at less than 1%; \*\*significant at 5%; \*significant at 10%.

distribution. Finally, the signs, size, and significance of the coefficients are similar for different measures of the saving rate.<sup>21,22</sup>

Proposition 3 states that an increase in income inequality leads to an increase in inequality of saving rates. In a way, our results in Table 2, that simply reproduce those of Dynan *et al.* (2004), already suggest this through a compositional effect. If the rich save more, increasing income inequality will increase saving inequality. In our empirical exercise, we can go one step further and test whether income inequality has a direct effect, beyond its impact through income, on saving rates and the saving-income relationship. To test this, we compute the coefficient of variation of disposable income for each year and state.<sup>23</sup> We then regress the household saving rate on income, the coefficient of variation, and the interaction between the two. Our theory predicts that the coefficient on this interaction

<sup>21</sup> We have also run the regression including state-year fixed effects instead of state reference income and state and year fixed effects. The coefficient on the interaction of state reference income and individual income remains significantly positive, in line with Proposition 2.

<sup>22</sup> While the correlation between log income and the interaction between log income and log state income is very high due to the low variation in state income relative to individual income, this is not driving results. To show this, we also estimated the regressions in Table 3 and the following tables excluding log income. This yielded very similar marginal effects of state income, with similar variation across the income distribution. Because it is theoretically strongly preferable to include log income in the regression, all results are reported including this variable.

<sup>23</sup> Since households with very low-income realizations, which we drop from our sample, may be important for the adequate measurement of income inequality we include them in our measures of inequality. Nonetheless, our results are not sensitive to this choice.

term should be positive, i.e. as inequality increases the gradient between saving rates and income becomes steeper. As always, we also include our set of controls which includes state dummies. We conduct this exercise for each one of the measures of saving. The last three columns of Table 3 show the results of these regressions. The coefficient of interest is positive and significant as predicted by the theory. The negative coefficient associated with inequality together with the positive interaction term implies that income inequality has a negative effect on saving rates for households with low levels of income. This negative effect becomes weaker, and after a certain point positive, as income increases. Since our results control for the level of income they suggest that income inequality has more than just a compositional effect on saving inequality.<sup>24,25</sup>

### Saving rates, permanent income and inequality

In this section we reproduce the previous exercises using data on consumption, food consumption, average labour income, and education as instruments for permanent income. We follow a two-stage estimation procedure. In the first stage, we regress current income on one of the instruments, the set of control variables, and measures of reference income or inequality. The fitted values of this regression are the measures of permanent income that we use in the second stage, where we estimate a median regression with the saving rate as our dependent variable and predicted permanent income, reference income or inequality, and the other controls as explanatory variables.

Table 4 summarizes these results. While point estimates vary with the choice of instrument, the magnitude and direction of the effects remain similar across specifications. The high value of the *F*-test of the excluded instrument implies that all the instruments are powerful. In the table, we report results for the total saving rate and include a linear income term as a regressor. We obtain similar results for other measures of saving or when we include a quadratic income term.

The first two columns of panel A in Table 4 present results using either consumption or food consumption as an instrument. To measure consumption, we follow Dynan *et al.* (2004) and use a weighted average of food at home, food away from home, rental payments and imputed housing flows. The estimated relationship between income and saving rates becomes steeper and the negative impact of reference income in the lower tail of the income distribution becomes stronger. Next, in column 3, we use average labour income as an instrument for permanent income. In order to avoid problems of reverse causality, this instrument excludes any source of capital income that could be correlated with past saving decisions. Furthermore, this measure also excludes labour income from the current period. The signs and statistical significance of our coefficients of interest remain unchanged. Finally, the last column of this table reports results using years of schooling of the household

<sup>24</sup> In our model a mean-preserving spread in the distribution of income increases inequality in saving rates only through a compositional effect (since rich households become richer and poor households poorer). This is simply a consequence of the fact that this spread does not affect the level of reference income, which is simply average income. In the data changes in inequality most likely affect reference income. This change in reference income is behind the direct effect of inequality on saving.

<sup>25</sup> Results are similar when we use other measures of inequality, particularly we reproduce our baseline using the Gini coefficient, the 90/10 percentile ratio and the standard deviation of the log of labour income. The picture that emerges from these extensions is the same: Inequality steepens the gradient between saving rates and income.



TABLE 4  
*Saving rates, permanent income, reference group and inequality*

	(1)	(2)	(3)	(4)
	<i>Consumption</i>	<i>Food consumption</i>	<i>Avg labour income</i>	<i>Education</i>
<i>Panel A: Reference group</i>				
ln(permanent income)	0.148*** (0.015)	0.119*** (0.019)	0.070*** (0.006)	0.115*** (0.005)
ln(state income)	-1.745*** (0.258)	-1.676*** (0.279)	-1.572*** (0.211)	-1.508*** (0.194)
ln(permanent income) × ln(state inc.)	0.164*** (0.025)	0.157*** (0.027)	0.151*** (0.020)	0.143*** (0.019)
Pseudo-R <sup>2</sup>	0.008	0.008	0.010	0.009
F-excluded instruments	2,753	760	10,492	5,088
N	20,578	20,195	21,360	24,036
<i>Panel B: Inequality</i>				
ln(permanent income)	0.132*** (0.018)	0.103*** (0.022)	0.038*** (0.010)	0.090*** (0.010)
Inequality	-0.149 (0.127)	-0.227* (0.127)	-0.339*** (0.118)	-0.254*** (0.104)
ln(permanent income) × inequality	0.013 (0.012)	0.020* (0.012)	0.030*** (0.011)	0.023*** (0.009)
Pseudo-R <sup>2</sup>	0.009	0.008	0.010	0.008
F-excluded instruments	2,574	781	10,522	5,722
N	20,578	20,195	21,360	24,036

*Notes:* The column heading indicates the instrument used in the regression. As a measure of saving we use change in wealth. Perm. income is a measure of permanent income estimated, using an OLS regression, including the instrument and our set of controls. State income is the state level average of this measure. The regression includes demeaned log state income. Inequality is measured as the coefficient of variation of disposable income per region an year. All regressions also include controls for age, age squared, household size, head's education and year and state dummies which are not reported. Clustered standard errors are reported in parentheses. \*\*\*Significant at less than 1%; \*\*significant at 5%; \*significant at 10%.

head as an instrument.<sup>26</sup> Again the same picture emerges, saving rates increase with income, decrease with reference income, and more so for low-income households.

Because of the interaction, it is hard to read marginal effects off Table 4 directly. Therefore, the marginal effects of average state income at different percentiles of the income distribution are depicted in Figure 2a. The marginal effect for the lowest percentile, -0.35, indicates that for households in this percentile of the income distribution, an increase in reference income of 10% reduces their saving rate by 0.035 percentage points. Although small, this effect is far from negligible, particularly if one considers that the median saving rate for households around that percentile is zero. This negative effect disappears between the 30th and 40th percentiles depending on the choice of instrument. As we pointed out in the previous section, this may be a consequence of the fact that average state income is not an appropriate reference level for richer households. In Figure 2b, we repeat the previous exercise restricting the sample to those households with income above the mean and recal-

<sup>26</sup> In this case we exclude years of schooling from the set of controls.

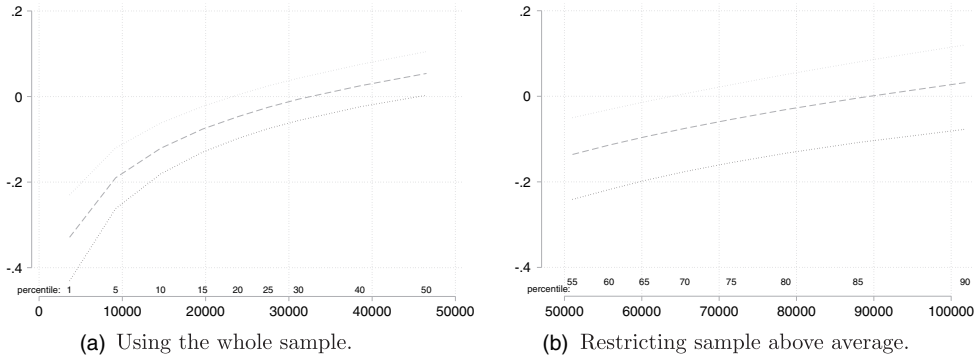


Figure 2. Marginal effect of reference group on saving rates

Notes: (a) Marginal effects of state income on saving rates. The marginal effects are computed using results from Table 4. Permanent income is estimated, using average labour income as instrument. All saving and disposable income measures are given in 2007 dollars. (b) Marginal effects of state income on saving rates restricting the sample to those households with income above average income in their state of residence. The marginal effects are computed using results from Table 13 in Appendix S3. Permanent income is estimated using average labour income as instrument.

culating reference income as the average income of this restricted sample. The sign, size, and significance of the coefficients remains unchanged. Nonetheless, the marginal effect of reference income now remains negative for households up to the 90th percentile of the (whole) income distribution, suggesting that increases in reference income reduce saving rates for most households.

Panel B in Table 4 explores the effects of income inequality on saving rates. The results are similar to those reported in the previous section. After controlling for permanent income, inequality has a negative impact on saving rates. This impact decreases with permanent income. Figure 3 illustrates the relationship between income and the saving rate implied by the coefficients reported in column 3 of panel B in Table 4. The solid line represents the income-saving relationship in a state with low inequality (Oregon) and the broken line that in a state with high inequality (Texas). The gradient between saving rates and income is positive and becomes steeper in states with high inequality. The effect of inequality on saving rates is substantial. For instance, if we take a Texan at the median of the US income distribution, his total saving rate is 5 percentage points or 40% lower than that of a household with the same level of income living in Oregon. Notice that inequality reduces saving rates for most of the income distribution.

## Further analysis

### *Sample composition and reverse causality*

In this section, we explore the robustness of our results to differences in the sample composition. We recalculate our measures of reference income and inequality for each sample. In the analysis that follows, we will use the average labour income as an instrument for permanent income and we will focus on total saving – change in wealth – as our measure of saving. Nonetheless, results are similar for the other measures of saving or when we use alternative instruments for permanent income. Results from these exercises are summa-

TABLE 5

*Saving rates, permanent income, reference group and inequality – Robustness checks*

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Baseline	Excl. sav. rate 5%	Income >10,000 <100,000	Income >35,000	Labour income	Housing Redistribution prices	Housing
<i>Panel A: Reference group</i>							
ln(perm.inc.)	0.070*** (0.006)	0.068*** (0.007)	0.067*** (0.008)	0.105*** (0.008)	0.085*** (0.008)	0.069*** (0.006)	0.071*** (0.007)
ln(state inc)	-1.572*** (0.211)	-1.549*** (0.200)	-3.787*** (0.576)	-2.420*** (0.413)	-0.637*** (0.083)	-1.532*** (0.202)	-1.492*** (0.247)
ln(perm.inc.) × (state inc)	0.151*** (0.020)	0.150*** (0.019)	0.370*** (0.055)	0.225*** (0.038)	0.062*** (0.008)	0.147*** (0.019)	0.141*** (0.024)
Governor index						0.006 (0.004)	
Housing prices							0.001 (0.0001)
Pseudo-R <sup>2</sup>	0.010	0.038	0.007	0.011	0.010	0.010	0.010
N	21,360	19,472	17,886	19,472	21,360	21,076	19,542
<i>Panel B: Inequality</i>							
ln(perm.inc.)	0.038*** (0.010)	0.039*** (0.010)	0.030 (0.020)	0.136*** (0.035)	0.023** (0.012)	0.036*** (0.011)	0.045*** (0.013)
Inequality	-0.339*** (0.118)	-0.305*** (0.115)	-0.578*** (0.215)	-0.199 (0.342)	-0.457*** (0.123)	-0.359*** (0.126)	-0.257** (0.115)
ln(perm.inc.) × (inequality)	0.030*** (0.011)	0.027** (0.011)	0.053*** (0.020)	0.016 (0.031)	0.042*** (0.012)	0.032*** (0.012)	0.023** (0.011)
Governor index						0.007 (0.005)	
Housing prices							0.001*** (0.000)
Pseudo-R <sup>2</sup>	0.010	0.038	0.006	0.010	0.010	0.010	0.010
N	21,360	19,472	17,886	14,180	21,360	21,076	19,542

*Notes:* As a measure of saving we use change in wealth. Perm. income is a measure of permanent income estimated using an OLS regression, including the instrument and our set of controls. State income is the state level average of this measure. The regression includes demeaned log state income. Inequality is measured as the coefficient of variation of disposable income per region an year. All regressions also include controls for age, age squared, household size, head's education and year and state dummies which are not reported. In columns 2–4, the column heading indicates the way the sample has been altered. In column 2, 5% has been eliminated from both sides of the saving rate distribution. In column 5, the income of the reference group and the coefficient of variation of income (inequality measure) have been computed using only labour income per state and year. In column 6, the regression includes the governor index, which goes from 0, the state governor was a Democrat during the period in which the saving decision took place, to 1, the state governor was a Republican during the entire period. In column 7, the regression includes a measure of real estate prices at the state level. Clustered standard errors are reported in parentheses. \*\*\*Significant at less than 1%; \*\*significant at 5%; \*significant at 10%.

rized in Table 5, where, for ease of comparison, column 1 reproduces our results from the baseline sample.

Excluding some observations at the top and bottom of the distribution of saving rates, we obtain some indication of the sensitivity of our results to extreme values. In column 2, we use a sample that drops households with saving rates at the top and bottom 5% of the

distribution. Not surprisingly, since we are using median regression, results are unaffected by this or similar trimmings.

Columns 3 and 4 present results after restricting the sample to households with different levels of permanent income. On the one hand, excluding households with permanent income below a certain threshold, \$10,000 per year, we ensure that our results are not driven by the behaviour of the very poor whose saving rates may simply reflect the presence of a subsistence level of consumption. On the other hand, several authors have pointed out difficulties in rationalizing the saving behaviour of very rich households in terms of a standard life-cycle model, see for instance Carroll (2000). Following a path that goes back, at least, to Weber (1958), Carroll (2000) argues that the pursuit of wealth for its own sake may be behind the saving choices of the very rich. Along these lines, excluding high-income households from our sample we ensure that our results are not driven by what Carroll (2000) refers to as the ‘spirit of capitalism’. In both cases, there is always a fraction of the income distribution that saves less as average income or inequality increase. Results in column 3 suggest that our previous estimates were not driven by the behaviour of very poor or very rich households.

Trimming our sample by income is also useful to infer whether the presence of asset-based means-tested government programs affects our results. These programs distort the saving incentives for marginally eligible households. In our context, this would be particularly important if the eligibility criteria or coverage change over time and this change is somehow correlated with inequality at the state level. By restricting our sample to households with permanent income over \$35,000, well above the eligibility criteria for most of these programs, we work with a set of households whose saving decisions should not be affected by the presence of these plans.<sup>27</sup> Column 4 in Table 5 report the results of this exercise. The estimates remain similar to those reported using the full sample. In view of this, we conclude that the presence of asset-based means tested government programmes is unlikely to be driving our results.

Additionally, we have also explored the effects of dropping from our sample the states with highest (Texas, New York and California) and lowest (Oregon, Washington and Wisconsin) inequality. Once more, the basic results remain unchanged.

One might think that differences in inequality across states may result from differences in the distribution of capital income that arise from different saving patterns across the income distribution. If this was the case, past differences in saving would be at the root of current differences in inequality. To show that our results are robust to this specific type of reverse causality, we repeat our basic regression constructing our measures of inequality using only labour income. Results, reported in column 5 of Table 5, are consistent with our baseline estimates.<sup>28</sup>

<sup>27</sup> Although the eligibility criteria of these programs is often stated in terms of income and wealth, we restrict our sample only in terms of permanent income. By doing so, we assume that permanent income is highly correlated with wealth. Finally, it is worth noticing that these programs tend to have a stronger impact on the saving decisions of older households which are already excluded from our baseline sample.

<sup>28</sup> Finally, since using the PSID to obtain state-level income or inequality measures could be problematic due to the non-representativeness of the PSID at state level, we use a different data source for alternative measures. State-level reference income from the Current Population Survey (CPS) is reported by the Bureau of Labour Statistics. We obtain state-level measures of inequality from Frank (2014). The main results are not altered by the use of these alternative measures.

*Does redistribution matter?*

There are two, somewhat opposite, concerns related to the degree of redistribution at the state level. First, along the lines suggested by Alesina and Rodrik (1994), states with more inequality may engage in more redistribution of goods and services for which households typically save, for instance health care or higher education. This may reduce saving in that state, particularly for poor households since health costs represent a higher fraction of their income. In this case inequality has an effect on saving rates but it does so through a channel different from the one illustrated in our theoretical model. Inequality affects saving through redistribution rather than through interpersonal comparisons. A second, and probably less relevant, concern arises if more equal states are so as a result of redistribution, this redistribution mainly takes the form of consumption goods and services, and these goods are necessities, i.e. they have a low income elasticity of demand. One can think of food stamps or public transportation as goods that satisfy these requirements. If this is the case, redistribution will reduce more than proportionally the consumption needs of poor households providing them with additional room for saving. As a result, low-income households may save more in more equal states.

This second concern arises if differences in inequality do simply reflect differences in redistribution at the state level, i.e. more equal states are so because they redistribute more. Since we have data on pretax labour income and income after taxes and transfers, we can test the extent to which redistribution drives the patterns of inequality across states.<sup>29</sup> We regress the estimated coefficient of variation of after-tax income on the coefficient of variation of pre-tax income, a time trend and regional dummies. The adjusted- $R^2$  of this regression is very high, 0.95, and the coefficient on the coefficient of variation of pretax income is 0.83. This suggests that pretax levels of inequality and state fixed characteristics account for almost all the variation in inequality in the data.

The first concern requires that more unequal states engage in more redistribution. To explore this hypothesis, we construct an index variable that indicates the fraction of time that the state governor was a Republican during the period in which the savings decision took place. This variable goes from 0, the state governor was a Democrat during the entire period, to 1, the state governor was a Republican during the entire period. As a first approximation, this variable should be a good proxy for the degree of generosity in the public provision of goods and services at the state level.

The correlation between regional inequality and our governor index is positive (10%), indicating that on average, states governed by Republicans have higher inequality. The scenario described in the first concern thus appears questionable. In addition, column 6 in Table 5 shows that when the governor index is included in our regressions, our results remain unchanged. Moreover, this index is only weakly significant, both on its own and when interacted with individual income. To the extent that differences in state governor capture differences in the generosity of the welfare state, our estimates of the effect of inequality on the saving-income relation do not seem to be driven by differences in redistribution at the state level.

<sup>29</sup> Notice that the denominator of any of our saving rates is income after taxes and transfers. This ensures that we do not overstate the saving rates of low-income households in more redistributive states inducing a spurious correlation between saving rates and inequality.

### *Housing prices*

It is possible that income inequality leads to increases in the prices of certain goods, particularly those that are locally in fixed supply. One can think of housing as a prime example. For instance, focusing on tight housing markets, Matlack and Vigdor (2008) find evidence consistent with this hypothesis. To the extent that this happens, higher housing prices would make residents of more unequal states poorer in real terms. As a consequence, they will tend to save less in line with the results of Dynan *et al.* (2004).

In column 7 of Table 5 we use a measure of real estate prices at the state level to test if the effect of inequality on savings results from differences in price levels. We include state level housing prices in our two standard specifications. These results show that higher house prices are associated with higher saving. This is only natural, as increases in real-estate prices directly feed into higher saving rates through capital gains on housing. The effects of reference group income and inequality become slightly smaller but remain significant.<sup>30</sup>

### *An alternative reference group*

Our analysis above has used average state income as income of the reference group, and therefore considered all residents of a state as a household's group of reference. It is not obvious what exactly constitutes a household's group of reference; all US citizens, all residents in the same state, or narrower groups? We therefore also provide results for three other definitions of reference group, one broader and two more narrow than our baseline. Results are shown in Table 6. Column 1 reproduces our baseline results. Column 2 shows results with average US income as the reference income. In column 3 we define the reference group as households in the same state of residence, the same level of education and the same occupation than the household of interest. Column 4 further restricts this last reference group by race. It is evident that results remain qualitatively the same and remain strongly significant. For national income, coefficients change qualitatively. We attribute this to the fact that in this setting, it is not possible to include time dummies. Results for the narrower reference groups, however, change quantitatively very little compared to our main results.

We also compute results for a placebo setting: We randomly assign to each household in our sample the reference income from a different state, not its own state. Results for this setting are shown in the last column. The random reference income is not significant. We therefore conclude that reference income robustly affects household saving, for different measures of the reference group.

### *The marginal propensity to save under upward-looking comparisons*

In this section, we follow the procedure outlined by Dynan *et al.* (2004) to estimate the marginal propensity to save (MPS). In line with our simple model of interpersonal comparisons we have used average state income as reference income for all the household in each state. Nonetheless, Veblen (1899), Duesenberry (1949) and Frank (2007) eloquently argue that the behaviour of successful individuals or groups sets the standard for the rest of the

<sup>30</sup> Similarly, Bertrand and Morse (2016) argue that it is unlikely that inequality affects consumption through home prices. They find that households consume a larger share of their income in more unequal places independently of whether they are homeowners (who could finance consumption by loans against home equity) or renters.

TABLE 6  
*Saving rates and permanent income – using different reference groups*

<i>Reference group</i>	(1)	(2)	(3)	(4)	(5)
	<i>Baseline (year, region)</i>	<i>Broader (US)</i>	<i>Narrower (year, region, occup, educ)</i>	<i>Narrower (year, region, occup, educ, race)</i>	<i>Placebo (random region)</i>
ln(perm. income)	0.070*** (0.006)	0.059*** (0.006)	0.108*** (0.009)	0.085*** (0.008)	0.068*** (0.007)
ln(reference group income)	−1.572*** (0.211)	−4.250*** (0.467)	−0.654*** (0.086)	−0.637*** (0.083)	0.222 (0.171)
ln(perm. income) × ln(reference group income)	0.151*** (0.020)	0.421*** (0.045)	0.062*** (0.008)	0.062*** (0.008)	−0.021 (0.016)
Pseudo-R <sup>2</sup>	0.010	0.010	0.010	0.010	0.010
<i>N</i>	21,360	21,360	18,326	18,141	21,360

*Notes:* The column heading indicates the reference group considered. Race consists in four categories (white, black, hispanic, other), education consists in three categories (Less than high school, high school, and more than high school), and occupation in four categories (group 1: management, professional, technical, financial, sales and public security occupations, group 2: administrative support and retail sales occupations, group 3: low-skill service occupations, group 4: production, transport and construction occupations). As a measure of saving we use change in wealth. Permanent income is a measure of permanent income estimated, using an OLS regression, including the instrument and our set of controls. State income is the state level average of this measure. The regression includes demeaned log state income. Inequality is measured as the coefficient of variation of disposable income per region an year. All regressions also include controls for age, age squared, household size, head's education and year and state dummies which are not reported. All regressions except the one using average national income as a reference group also control for year fixed effects. Clustered standard errors are reported in parentheses. \*\*\*Significant at less than 1%; \*\*significant at 5%; \*significant at 10%.

community. Ferrer-i Carbonell (2005) provides convincing microeconomic evidence on the importance of upward-looking comparisons as a determinant of subjective well-being. In view of this evidence, a more realistic setup would allow for individuals to compare themselves with those who are ranked above them in the income distribution.<sup>31</sup> Under upward-looking comparisons, state average income is only a sensible reference point for those households in the lower part of the income distribution. Along these lines in this section, we consider a different measure of reference income. We classify households in a given state and year into permanent income quintiles. We then estimate the reference income of each household as the average income of those quintiles above him. We do that for all the individuals except those in the top quintile, for whom we measure reference income as the average income of their own quintile.<sup>32</sup>

We exploit the panel dimension of our dataset and regress changes in total saving over time for a given household on changes in its income and changes in reference income. The coefficients associated with the income quintile dummies are our measures of the quintile specific MPS. Table 7 summarizes the results of this exercise. In both specifications the coefficients are statistically different from zero and, except for the bottom two quintiles,

<sup>31</sup> We will explore the theoretical implications of such an environment in the next section.

<sup>32</sup> Results excluding the top quintile are similar. Another possibility is to use average income of individuals with higher income as the reference for individuals in the highest income group. Although the results are unaffected, this specification suffers from important problems of multicollinearity between variation in households' level of income and variation in income of the reference group.

TABLE 7  
*Marginal propensity to save*

	<i>Change in saving</i> $\Delta(\Delta Wealth)$	
	(1)	(2)
Quantile 1	0.040 (0.024)	0.049 (0.027)
Quantile 2	0.064 (0.016)	0.067 (0.018)
Quantile 3	0.069 (0.009)	0.078* (0.011)
Quantile 4	0.139* (0.008)	0.111* (0.010)
Quantile 5	0.327* (0.002)	0.328* (0.003)
$\Delta$ income reference group		-0.062*** (0.008)
Pseudo- $R^2$	0.007	0.007
$N$	16,798	14,236

*Notes:* This table shows estimates of marginal propensity to save for each quintile obtained from changes in saving over time. The dependent variable is the change in household saving between periods divided by the number of years in the period. Reported in the table are marginal effects of changes in household income between periods by income quintile. They are computed from the coefficient on the interaction between the change in income and the income quintile indicators. Column 2 shows the results of the specification in which we included the change in income of the reference group, estimated as the average income of the quintiles above the quintile to which the individual belongs for a given space and a given year. Note that for quintile 5, we use the average income of the quintile as income of the reference group. Also included in the regression (coefficients not reported in the table) are the sets of indicator variables for permanent income, age, a quadratic term in age, the number of persons in the household, education and year and regional dummies. Standard errors are reported in parentheses. \*The coefficient is significantly greater than that for the previous quintile, on the basis of a one-sided 5% test.

different from each other. In the second column we include the change in income of the reference group as a regressor. Its coefficient is negative and statistically significant. To illustrate the implications of our result, consider again Figure 1b. Between 1984 and 2007, for households in the lower half of the income distribution, average income increased by \$2,018 and average saving decreased by \$1,844. As discussed in the introduction, a MPS that increases with income is difficult to reconcile with these patterns. Using the estimates in column 1 that abstract from the impact of changes in reference income, one would predict that savings of the average household with income below the median should have increased by \$129 over this period. Results in column 2 allow us to control for the



evolution of reference income. Since this measure increased by \$13,284 for this household, our estimates predict a drop in savings of \$688.<sup>33</sup>

#### IV. Aggregate saving and income inequality

In the previous model, we characterized the savings behaviour of an heterogeneous-agent economy under the restrictive assumption that the reference group was common and equal to the average household in the economy. In this section, we relax this assumption and revisit the relationship between aggregate savings and inequality. For the sake of illustration, we restrict interpersonal comparisons to the first period of life setting  $\zeta = 0$  and assume there are only two homogeneous income groups, denoted by the superscripts  $H$  (rich) and  $L$  (poor), with their wages given by  $w^H$  and  $w^L$  respectively, with  $w^H > w^L$ . The population is evenly distributed between these two groups. In line with the evidence discussed in the previous section, we assume that the reference group of the rich households is made up only of rich households while the reference group of poor households is composed of a weighted average of poor and rich households, with  $\rho$  being the weight placed on poor households. As a result, reference consumption levels for the two groups are given respectively by

$$\bar{g}_t^H = \bar{c}_t^H \quad \text{and} \quad \bar{\sigma}_{t+1}^H = \bar{d}_{t+1}^H \tag{10}$$

$$\bar{g}_t^L = \rho \bar{c}_t^L + (1 - \rho) \bar{c}_t^H \quad \text{and} \quad \bar{\sigma}_{t+1}^L = \rho \bar{d}_{t+1}^L + (1 - \rho) \bar{d}_{t+1}^H \tag{11}$$

The levels of saving for a rich and a poor household are given by<sup>34</sup>

$$s_t^H = (\beta(1 - \gamma)/(1 + \beta(1 - \gamma)))y_t^H, \tag{12}$$

$$s_t^L = \frac{\beta(1 - \gamma\rho)}{1 + \beta(1 - \gamma\rho)}(y_t^L - \frac{\beta\gamma(1 - \rho)}{1 + \beta(1 - \gamma)}y_t^H). \tag{13}$$

Although this model has similar implications for individual saving than the previous one, its aggregate implications differ. The following proposition summarizes the impact of income inequality on aggregate savings in the presence of upward-looking comparisons.

*Proposition 4.* An increase in income inequality, a mean preserving spread in the distribution of wages, leads to a decrease in the aggregate (average) saving rate.

*Proof.* Defining the share of total income received by the rich households by  $y_s^H \equiv \frac{y^H}{y^H + y^L}$ , we combine equations (12) and (13) to derive the aggregate (average) saving rate as,

$$s_{\text{rate}}^{\text{agg}} = y_s^H s_{\text{rate}}^H + (1 - y_s^H) s_{\text{rate}}^L = \frac{\beta}{1 + \beta(1 - \gamma\rho)} \left( 1 - \gamma\rho - \frac{\beta\gamma(1 - \rho) + \beta\gamma(1 - \gamma\rho)(1 - \gamma)}{(1 + \beta(1 - \gamma))} y_s^H \right) y_s^H \tag{14}$$

<sup>33</sup> This presupposes that the household with average income in the bottom 50% of the income distribution is in the second quintile of the distribution. For this household, reference income therefore is the average income of the top three quintiles.

<sup>34</sup> See Appendix S1 for a detailed derivation.

which is decreasing in the share of income received by rich households,  $y_s^H$ , and therefore in inequality.

In this framework, where upward-looking interpersonal comparisons matter, a mean preserving spread in the distribution of wages reduces the aggregate level of saving and therefore the aggregate saving rate. Although an increase in the wage of rich households leads to a proportional increase in their level of saving, leaving their saving rate unchanged, it is clear from equation (13) that this increase leads to a decrease not only in the level of saving of the poor but also in their saving rate. This decrease in saving of poor households is larger than the increase in saving of the rich.

*Remark to Proposition 4.* A distributionally neutral increase in income, i.e. an increase in income that leaves unchanged the shares of aggregate income earned by rich and poor households, leaves the aggregate saving rate unchanged.

It is evident from equation (14) that the aggregate saving rate remains unchanged as long as the share of income of rich households remains constant. This implies this last remark.

## V. Empirical analysis: inequality and aggregate saving

Proposition 4 states that an increase in inequality should reduce the aggregate saving rate. Our previous results suggest that this relationship holds in US microdata. For instance, Figure 3 indicates that individuals below the 85th percentile of the permanent income distribution save less in more unequal states. In this section, we test again Proposition 4 using panel data on six OECD countries. We estimate a simple saving specification that includes a measure of income inequality together with other standard determinants of saving. Our theory predicts that the coefficient on inequality should be negative.

In contrast to previous cross-country studies on the determinants of aggregate saving that focus on measures of saving derived from the current account and therefore include elements of public or corporate saving (Edwards, 1996; Schmidt-Hebbel and Serven, 2000; Li and Zou, 2004; Leigh and Posso, 2009), we will use a measure of the personal saving rate, from several OECD yearbooks, as our dependent variable. This measure is closer to our theory, which focuses on the determinants of household saving and abstracts from other components of national saving. Reasonably long series for the personal saving rate are not available for many countries, so we are forced to restrict our sample to Canada, France, Germany, Japan, the UK and the US. Although the number of countries in our sample is not large, the time dimension is, covering the years 1955–2007. Over this period, our sample accounts for roughly 40% of world GDP.

Our measure of inequality is the fraction of total income accruing to the top five per cent of the income distribution.<sup>35</sup> We obtain this measure from Alvaredo *et al.* (2013). This variable is available at a yearly frequency, except for Germany, where it is measured on a tri-annual basis. We follow Schmidt-Hebbel and Serven (2000) and include the age

<sup>35</sup> We have estimated the same specification using other conventional indicators of inequality like the Gini coefficient or the ratio of the income shares of the richest 5% and the bottom 95%, with similar results.

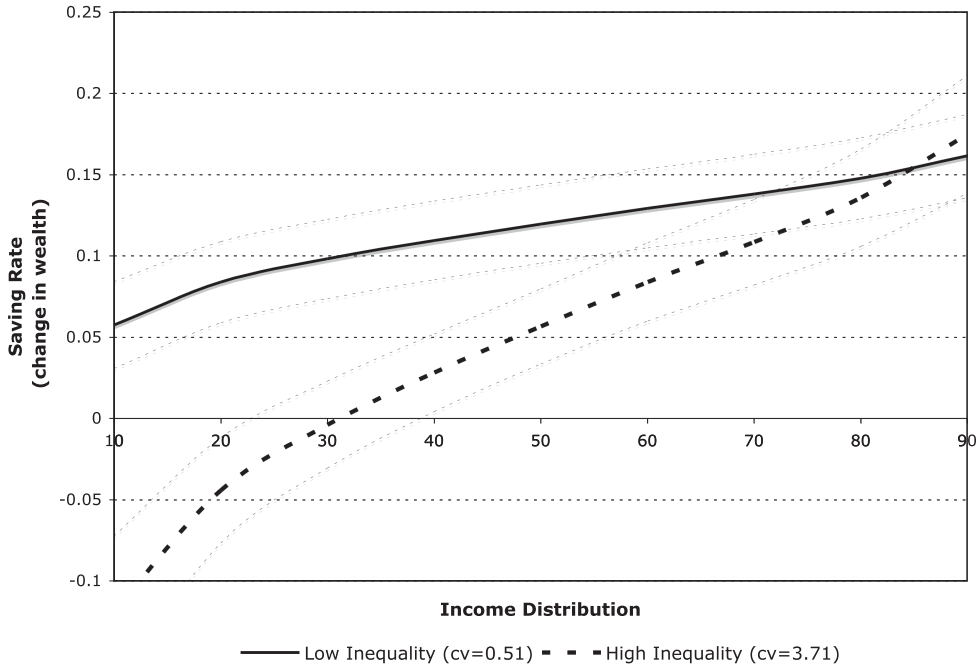


Figure 3. The saving rate by income percentile in states with different levels of inequality  
*Notes:* Percentiles of the permanent income distribution on the  $x$ -axis. Saving rates are predicted values from the regression in column 4 in Table 4, for 2005, for a 45-year old head of a household with four members and 12 years of education. Percentiles of the permanent income distribution used in the prediction refer to the entire sample. Therefore, the level of income at each point is the same for both groups of states. The values of the coefficient of variation correspond to the minimum and maximum in the sample across regions. (Note: low being Oregon and high being Texas.)

dependency ratio, the rate of growth of real GDP per capita, and the real interest rate as additional regressors.<sup>36</sup>

Results are reported in Table 8. Column 1 shows a negative association between aggregate saving rates and income inequality, controlling for other determinants of saving. This relationship remains robust to the inclusion of country-fixed effects, as shown in column 2. In both specifications, the coefficient on inequality is negative and significant. The results from column 2 suggest that a percentage point increase in the share of income accruing to the top 5% of the income distribution is associated with a reduction in the personal saving rate of 0.26 percentage points. The evolution of the share of income of the top 5%, reported in Figure 1, together with this estimate suggests that around half of the decrease in the US personal saving rate (3.38 percentage points) can be associated with the process of income concentration that began more than three decades ago.<sup>37</sup>

<sup>36</sup> Notice that these authors include the level of per capita GDP and its square to capture the nonlinear relationship between income and savings. This is particularly important when the variation in income through time or across countries is substantial. This is not the case in our sample. Nonetheless, the sign and significance of our coefficient of interest remains unchanged when we replace the rate of growth of per capita GDP by its level.

<sup>37</sup> The remaining coefficients also have the expected signs. The dependency ratio has a negative effect on the aggregate saving rate while in most specifications the real interest rate and the GDP growth rate have a positive effect.

TABLE 8  
Aggregate saving rates and income inequality

	Fixed effects					Fixed effects financial development	
	Pooled OLS baseline		Time dummies	5-years average	GMM		5-years average
	(1)	(2)	(3)	(4)	(5)	(6)	
Inequality	-0.297*** (0.066)	-0.262*** (0.074)	-0.162** (0.080)	-0.282* (0.158)	-0.032** (0.015)	-0.119** (0.055)	-0.596** (0.293)
Age dependency ratio	-0.311*** (0.109)	-0.410*** (0.120)	-0.237 (0.144)	-0.157 (0.294)	-0.054 (0.030)	-0.065** (0.130)	-0.705*** (0.225)
GDP growth	0.226** (0.095)	0.215** (0.093)	0.192* (0.098)	0.800** (0.315)	-0.166*** (0.046)	0.333* (0.197)	0.189 (0.119)
Real interest rates	0.513*** (0.072)	0.530*** (0.072)	0.566*** (0.109)	0.566*** (0.149)	0.003 (0.115)	0.077 (0.034)	0.192 (0.161)
Financial development							0.088*** (0.026)
Inequality × fin. development							-0.003** (0.001)
Wald test ( <i>P</i> -value)	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Sargan test ( <i>P</i> -value)				0.792	0.230		
AR(2) ( <i>P</i> -value)				0.187	0.651		
Observations	216	216	216	41	214	39	86

*Notes:* The dependent variable is the country-level household saving rate. Source: OECD yearbooks (1975, 1980, 1983, 1988, 1991 and 2001). The first independent variable captures inequality and is the share of income received by the top 5% households. Source: Alvaredo *et al.* (2013). The age dependency ratio is the population aged 65 and over as share of total population. Source: Groningen Growth and Development Centre, Total Economy Database, <http://www.rug.nl/research/ggdc/data/total-economy-database->. GDP growth is real per capita GDP growth from the World Bank National Accounts. The real interest rate is interest on bonds minus inflation from the OECD. The measure of financial development is stock market capitalization over GDP. Source: Beck and Demirguc-Kunt (2009). The countries included in the sample are Canada, France, Germany, Japan, the UK and the US. The sample consists of annual observations covering 1954–2007, except for column 3 in which we use five-years averages for all variables. Standard errors are reported in parentheses. \*\*\*Significant at less than 1%; \*\*significant at 5%; \*significant at 10%.

In the next two columns, we make sure that our results are not due to unique events in certain periods or business cycles. We do so by including decade dummies (column 3) and by using five-year averages of our yearly data (column 4). In both cases, the negative relationship between saving and inequality remains significant.<sup>38</sup>

At the aggregate level, the growth rate of per capita GDP and the interest rate are jointly determined with the saving rate. As a result, endogeneity becomes a concern. To deal with this issue, we estimate the model using Blundell and Bond's (1998) system GMM estimator with one lag. To further reduce endogeneity concerns, we use five-year averaged

<sup>38</sup> Furthermore, we estimate the model dropping one country at a time to verify that our results are not driven by the behaviour of any individual country. The basic results remain unchanged.

data to extend the length of the lag. Columns 5 and 6 show that the coefficient on inequality remains negative and significant.

Finally, it is worth noticing that the timing of income in our theoretical model is equivalent to the assumption that credit markets are perfect. In this sense, one should expect the negative effect of inequality on aggregate saving through interpersonal comparisons to be reduced by the presence of credit market imperfections. These imperfections would prevent low income households from borrowing against their future income to emulate the consumption of richer households. To allow for this relationship, we control for a measure of financial development and its interaction with inequality in column 7. Among the several popular measures of financial development in the literature, we choose stock market capitalization because this measure, unlike ratios of credit or deposits to GDP, is less directly related to the saving rate. Results are consistent with the predictions of the model. The relationship between inequality and saving remains negative, and more so for higher levels of financial development.<sup>39</sup>

Despite limited availability of data – which limits the set of econometric techniques that can be used – we can thus establish a negative relationship between inequality and saving not only in individual-level data, but also in aggregate data. We see these results as promising, with more data on private saving, we could probe the robustness of these results with more detail and take the analysis to the level of detail of our microdata analysis.

## VI. Conclusions

We have presented two models, four propositions and some statistical evidence on the importance of interpersonal comparisons and its connection with the positive gradient between saving rates and income documented by Dynan *et al.* (2004). The intuition behind our theoretical model is simple. When households care about relative consumption, the comparison of their consumption with that of their peers, the level of individual saving is determined by the interaction of two opposing forces. First, saving increases in the household's lifetime income, as in the permanent income hypothesis. Second, as a result of interpersonal comparisons, saving decreases in the lifetime income of the reference group as in the relative income hypothesis.

In addition, our model suggests that the negative effect of reference income on saving rates is stronger for low-income households, that increases in income inequality translate into increases in inequality of saving rates, and that, in the presence of upward-looking interpersonal comparisons, changes in income inequality affect aggregate saving.

We use more than two decades of PSID data and a long, although thin, panel of OECD countries to test these implications. The data supports the basic predictions of our theory. Like Dynan *et al.* (2004), we find that saving rates increase in permanent income. Approximating income for the reference group by average state income, we find that, after controlling for own income, saving rates decrease in comparison income and they do so more for poor households. When we include a measure of income inequality, we find that those regions or periods with higher inequality are characterized not only by a more

<sup>39</sup> Unfortunately, we are limited to using panel fixed effects with yearly data in this regression due to limited data availability. The World Bank's measure of stock market capitalization only covers the years 1989–2007. Taking five-year averages would leave us with only two observations per country.

unequal distribution of saving rates but also by lower saving rates for most of the income distribution. These results do not seem to be driven by measurement error, reverse causality, differences in redistribution across time or states, or variation in real estate prices. Finally, although preliminary, our cross-country results suggest that the distribution of income may have important aggregate implications.

In their seminal paper, Dynan *et al.* (2004) conclude that ‘the rich do, indeed, save more’ or, equivalently, that non-rich households save less. Our analysis suggests one of the reasons for these lower rates of saving, the importance of interpersonal comparisons.

*Final Manuscript Received: February 2018*

## References

- Abel, A. (1990). ‘Asset prices under habit formation and catching up with the joneses’, *American Economic Review*, Vol. 80, pp. 38–42.
- Abel, A. (2005). ‘Optimal taxation when consumers have endogenous benchmark levels of consumption’, *Review of Economic Studies*, Vol. 72, pp. 21–42.
- Alesina, A. and Rodrik, D. (1994). ‘Distributive politics and economic growth’, *The Quarterly Journal of Economics*, Vol. 109, pp. 465–490.
- Alonso-Carrera, J., Caballe, J. and Raurich, X. (2008). ‘Estate taxes, consumption externalities, and altruism’, *Journal of Public Economics*, Vol. 92, pp. 1751–1764.
- Alvaredo, F., Atkinson, A., Piketty, T. and Saez, E. (2013). ‘The top 1 percent in international and historical perspective’, *Journal of Economic Perspectives*, Vol. 27, pp. 3–20.
- Alvarez-Cuadrado, F. and Long, N. V. (2011). ‘The relative income hypothesis’, *Journal of Economic Dynamics and Control*, Vol. 35, pp. 1489–1501.
- Alvarez-Cuadrado, F. and Long, N. V. (2012). ‘Envy and inequality’, *Scandinavian Journal of Economics*, Vol. 114, pp. 949–973.
- Atkeson, A. and Ogaki, M. (1996). ‘Wealth-varying intertemporal elasticities of substitution: evidence from panel and aggregate data’, *Journal of Monetary Economics*, Vol. 38, pp. 507–534.
- Atkeson, A. and Ogaki, M. (1997). ‘Rate of time preference, intertemporal elasticity of substitution, and level of wealth’, *The Review of Economics and Statistics*, Vol. 79, pp. 564–572.
- Attanasio, O. and Browning, M. (1995). ‘Consumption over the life cycle and over the business cycle’, *American Economic Review*, Vol. 85, pp. 1118–1137.
- Attanasio, O., Hurst, E. and Pistaferri, L. (2012). *The Evolution of Income, Consumption, and Leisure Inequality in The US, 1980–2010*, NBER Working Papers No. 17982, National Bureau of Economic Research.
- Autor, D. H., Katz, L. F. and Kearney, M. S. (2008). ‘Trends in US wage inequality: revising the revisionists’, *Review of Economics and Statistics*, Vol. 90, pp. 300–323.
- Beck, T. and Demirguc-Kunt, A. (2009). *Financial Institutions and Markets Across Countries and over Time: Data and Analysis*, World Bank Policy Research Working Paper No. 4943.
- Becker, G. S. and N. Tomes (1986). ‘Human Capital and the Rise and Fall of Families’, *Journal of Labor Economics*, Vol. 4, pp. 1–39.
- Bertrand, M. and Morse, A. (2016). ‘Trickle-down consumption’, *The Review of Economics and Statistics*, Vol. 9, pp. 863–879.
- Blundell, R. W. and Bond, S. (1998). ‘Initial conditions and moment restrictions in dynamic panel data models’, *Journal of Econometrics*, Vol. 87, pp. 115–143.
- Boskin, M. J. (1978). ‘Taxation, saving and the rate of interest’, *Journal of Political Economy*, Vol. 76, pp. 3–27.
- Bosworth, B. and Anders, S. (2008). *Savings and Wealth Accumulation in the PSID, 1984–2005*, CRR Working Paper No. 2008-2.
- Browning, M. and Leth-Petersen, S. (2003). ‘Imputing consumption from income and wealth information’, *The Economic Journal*, Vol. 113, pp. F282–F301.

- Carroll, C. D. (2000). *Portfolios of the Rich*, NBER Working Papers No. 7826, National Bureau of Economic Research.
- Carroll, C. D. and Weil, D. N. (1994). 'Saving and growth: a reinterpretation', *Carnegie-Rochester Conference Series on Public Policy*, Vol. 40, pp. 133–119.
- Carroll, C. D., Overland, J. and Weil, D. N. (2000). 'Saving and growth with habit formation', *American Economic Review*, Vol. 90, pp. 341–355.
- Caselli, F. and Ventura, J. (2000). 'A representative consumer theory of distribution', *American Economic Review*, Vol. 90, pp. 909–926.
- Charles, K. K., Hurst, E. and Roussanov, N. (2009). 'Conspicuous consumption and race', *The Quarterly Journal of Economics*, Vol. 124, pp. 425–467.
- Chatterjee, S. (1994). 'Transitional dynamics and the distribution of wealth in a neoclassical growth model', *Journal of Public Economics*, Vol. 54, pp. 97–119.
- Clark, A., Frijters, P. and Shields, M. (2008). 'Relative income, happiness, and utility: an explanation for the Easterlin paradox and other puzzles', *Journal of Economic Literature*, Vol. 46, pp. 95–144.
- Coibion, O., Gorodnichenko, Y., Kudlyak, M. and Mondragon, J. (2014). *Does Greater Inequality Lead to More Household Borrowing? New Evidence from Household Data*, NBER Working Papers No. 19850, National Bureau of Economic Research.
- Coleman, J. (1961). *The Adolescent Society*, Free Press, New York.
- Corneo, G. and Jeanne, O. (1998). 'Social organization, status, and savings behavior', *Journal of Public Economics*, Vol. 70, pp. 37–51.
- Corsaro, W. and Eder, D. (1990). 'Children's peer cultures', *Annual Review of Sociology*, Vol. 16, pp. 197–220.
- Cronqvist, H. and Siegel, S. (2015). 'The origins of savings behavior', *Journal of Political Economy*, Vol. 123, pp. 123–169.
- Cushing, M. J. (2005). 'Net marginal social security tax rates over the life-cycle', *National Tax Journal*, Vol. 58, pp. 227–246.
- De Giorgi, G., Frederiksen, A. and Pistaferri, L. (2016). *Consumption Network Effects*, NBER Working Papers No. 22357, National Bureau of Economic Research.
- De Nardi, M. and Fella, G. (2017). *Saving and Wealth Inequality*, CEPR Discussion Paper No. DP11746. <https://ssrn.com/abstract=2896036>.
- De Nardi, M., French, E. and Jones, J. B. (2010). 'Why do the elderly save? The role of medical expenses', *Journal of Political Economy*, Vol. 118, pp. 39–37.
- Duesenberry, J. (1949). *Income, Saving and the Theory of Consumer Behavior*, Harvard University Press, Cambridge, MA.
- Dynan, K., Skinner, J. and Zeldes, S. (2004). 'Do the rich save more?', *Journal of Political Economy*, Vol. 112, pp. 397–444.
- Easterlin, R. (1995). 'Will raising the incomes of all increase the happiness of all?', *Journal of Economic Behavior and Organization*, Vol. 27, pp. 35–47.
- Easterlin, R. (2010). 'Does economic growth improve the human lot? Some empirical evidence', in David P. A. and Reder M. W. (eds), *Nations and Households in Economic Growth: Essays in Honour of Moses Abramovitz*, New York: Academic Press, pp. 89–125.
- Edwards, S. (1996). 'Why are Latin Americas savings rates so low? An International comparative analysis', *Journal of Development Economics*, Vol. 51, pp. 5–44.
- Feldstein, M. and Samwick, A. (1992). *Social Security Rules and Marginal Tax Rates*, NBER Working Papers No. 3962, National Bureau of Economic Research.
- Ferrer-i Carbonell, A. (2005). 'Income and well-being: an empirical analysis of the comparison income effect', *Journal of Public Economics*, Vol. 89, pp. 997–1019.
- Frank, R. (1985). *Choosing the Right Pond: Human Behavior and the Quest for Status*, Oxford University Press, New York and Oxford.
- Frank, R. (2007). *Falling Behind: How Income Inequality Harms the Middle Class*, University of California Press, Berkeley, CA.
- Frank, M. W. (2014). 'A new state-level panel of annual inequality measures over the period 1916–2005', *Journal of Business Strategies*, Vol. 31, pp. 241–263.
- Friedman, M. (1957). *A Theory of the Consumption Function*, Princeton University Press, Princeton, NJ.

- Gali, J. (1994). 'Keeping up with the joneses: consumption externalities, portfolio choice, and asset prices', *Journal of Money, Credit, and Banking*, Vol. 26, pp. 1–8.
- Huggett, M. and Ventura, G. (2000). 'Understanding why high income households save more than low income households', *Journal of Monetary Economics*, Vol. 45, pp. 361–397.
- Hurst, E. and Lusardi, A. (2004). 'Liquidity constraints, household wealth, and entrepreneurship', *Journal of Political Economy*, Vol. 112, pp. 319–347.
- Imbens, G. and Angrist, J. D. (1995). 'Identification and estimation of local average treatment effects', *Econometrica*, Vol. 62, pp. 467–475.
- Katz, L. F. and Autor, D. H. (1999). 'Changes in the wage structure and earnings inequality', in Ashenfelter O. and Card D. (eds), *Handbook of Labor Economics*, Vol. 3, Amsterdam: Elsevier, chapter 26, pp. 1463–1555.
- Kopecky, K. and Koreshkova, T. (2010). *The Impact of Medical and Nursing Home Expenses and Social Insurance*, Federal Reserve Bank of Atlanta Working Paper No. 2010-19.
- Kuznets, S. (1942). *Uses of National Income in Peace and War*, National Bureau of Economic Research, New York.
- Leigh, A. and Posso, A. (2009). 'Top incomes and national savings', *Review of Income and Wealth*, Vol. 55, pp. 57–74.
- Levine, A., Frank, R. H. and Dijk, O. (2014). 'Expenditure cascades', *Review of Behavioral Economics*, Vol. 1, pp. 55–73.
- Li, H. and Zou, H.-f. (2004). 'Savings and income distribution', *Annals of Economics and Finance, Society for AEF*, Vol. 5, pp. 245–270.
- Liu, W.-F. and Turnovsky, S. (2005). 'Consumption externalities, production externalities, and long-run macroeconomic efficiency', *Journal of Public Economics*, Vol. 89, pp. 1097–1129.
- Lusardi, A. (1999). 'Information, expectations, and savings for retirement', in Aaron H. (ed), *Behavioral Dimensions of Retirement Economics*, Washington, D.C.: Brookings Institution Press and Russell Sage Foundation, pp. 81–115.
- Luttmer, E. F. (2005). 'Neighbors as negatives: relative earnings and well-being', *Journal of Economics*, Vol. 120, pp. 963–1002.
- Matlack, J. L. and Vigdor, J. L. (2008). 'Do rising tides lift all prices? Income inequality and housing affordability', *Journal of Housing Economics*, Vol. 17, pp. 212–224.
- Mayer, T. (1966). 'The propensity to consume permanent income', *American Economic Review*, Vol. 56, pp. 1158–1177.
- Mayer, T. (1972). *Permanent Income, Wealth, and Consumption: A Critique of the Permanent Income Theory, the Lif-Cycle Hypothesis, and Related Theories*, Berkeley University Press, California, CA.
- Modigliani, F. (1970). 'The life-cycle hypothesis and intercountry differences in the saving ratio', in Eltis M. F. S. W. A. and Wolfe J. N. (eds), *Induction, Growth, and Trade: Essays in Honour of Sir Roy Harrod*, Oxford: Oxford University Press, pp. 197–225.
- Modigliani, F. and Brumberg, R. H. (1954). 'Utility analysis and the consumption function: an interpretation of cross-section data', in Kurihara K. K. (ed), *Post-Keynesian Economics*, New Brunswick, NJ: Rutgers University Press, pp. 388–436.
- Oswald, A. J. (1997). 'Happiness and economic performance', *The Economic Journal*, Vol. 107, pp. 1815–1831.
- Piketty, T. and Saez, E. (2003). 'Income inequality in the United States, 1913–1998', *The Quarterly Journal of Economics*, Vol. 118, pp. 1–39.
- Quadrini, V. (1999). 'Growth, learning and redistributive policies', *Journal of Public Economics*, Vol. 74, pp. 263–297.
- Runkle, D. E. (1991). 'Liquidity constraints and the permanent-income hypothesis: evidence from panel data', *Journal of Monetary Economics*, Vol. 27, pp. 73–98.
- Schmidt-Hebbel, K. and Serven, L. (2000). 'Does income inequality raise aggregate saving?', *Journal of Development Economics*, Vol. 61, pp. 417–446.
- Simmons, R. and Blyth, D. (1987). *Moving Into Adolescence: The Impact of Pubertal Change and School Context*, Aldine Press, New York.
- Veblen, T. (1899). *The Theory of the Leisure Class: An Economic Study of Institutions*, Modern Library, New York.
- Weber, M. (1958). *The Protestant Ethic and the Spirit of Capitalism*, Routledge Classics, London, UK.



Wendner, R. (2014). 'Ramsey, Pigou, heterogeneous agents, and non-atmospheric consumption externalities', *Journal of Public Economic Theory*, Vol. 16, pp. 491–521.

## **Supporting Information**

Additional supporting information may be found in the online version of this article:

**Appendix S1.** Details on the solution of the model in section II.

**Appendix S2.** Details on the data sources and sample restrictions.

**Appendix S3.** Additional results.