Statistics Norway's Open Research Repository - *SNORRe*

https://brage.bibsys.no/xmlui/handle/11250/177677

www.ssb.no

Heterogeneity in Returns to Wealth

and the Measurement of Wealth Inequality[1]

Andreas Fagereng
(Statistics Norway)

Luigi Guiso
(EIEF)

Davide Malacrino
(Stanford University)

Luigi Pistaferri
(Stanford University and NBER)

Abstract

Lacking a long time series on the assets of the very wealthy, Saez and Zucman (2015) use US tax records to obtain estimates of wealth holdings by capitalizing asset income from tax returns. They document marked upward trends in wealth concentration. We use data on tax returns and actual wealth holdings from tax records for the whole Norwegian population to test the robustness of the methodology. We document that measures of wealth based on the capitalization approach can lead to misleading conclusions about the level and the dynamics of wealth inequality if returns are heterogeneous and even moderately correlated with wealth.

JEL Classification: E13, E21, E24

Keywords: Wealth inequality, returns to wealth, heterogeneity in returns to wealth.

1. Introduction

Data on the stock of wealth of the very wealthy are rare. Yet, people at the top of the wealth distribution control a large share of the total assets in the economy. Variation in these shares can have first order effects on overall wealth inequality.[2]

Unfortunately, measuring wealth at the top and its evolution over time is difficult. Survey data are problematic. The wealthy are too few to be sampled and even oversampling some leaves out too much wealth. Furthermore, underreporting is notoriously a problem, it is increasing over time and, importantly, likely to be more relevant precisely among wealthier households (Meyer, Mok, and Sullivan, 2015). To bypass these problems, Saez and Zucman (2015) propose to estimate wealth holdings by capitalizing income from capital obtained from tax records. They apply this methodology to the US, documenting a marked increase in wealth inequality over the past 30 years. One key assumption behind the capitalization approach is that returns to wealth are homogeneous across households for broad asset categories. In this paper we use administrative Norwegian data, which have the unique feature of reporting both income from capital and actual wealth holdings for the whole population. This allows us to compare actual wealth inequality with wealth inequality predicted by the capitalization approach. The difference between the two reflects the role of heterogeneity in asset returns and their correlation with the level of wealth.

2. The capitalization method

Let w_{it} denote the wealth of individual i in year t and r_{it} the individual rate of return on wealth. Income from capital is $y_{it} = r_{it}w_{it}$, which is observed from tax records. Because wealth and

---

[2] More generally, if a small group (say the 0.5% of the population) owns a share S of the population wealth, the Gini index for the whole population is approximately equal to $G = G_r(1 - S) + S$ where G_r is the Gini inequality index of the remaining 99.5% of the population. Hence, swings in S can cause large changes in G.

individual returns are not observed, Saez and Zucman (2015) estimate the individual stock of wealth as: $\hat{w}_{it} = (y_{it}/\bar{r}_t) = (r_{it}w_{it}/\bar{r}_t)$, where $\bar{r}_t$ is a mean rate of return on wealth observed at time *t*. This methodology, first suggested by Giffen (1913), has undeniable advantages: because tax records cover the whole population, all people at the top of the distribution are included; also, measurement error (except for tax evasion) is likely to be small. Finally, because tax authorities store individual filing records, inequality measures can be constructed for long periods of time. Indeed, the US data go as far back as 1913. As Saez and Zucman acknowledge, however, the capitalization approach also has drawbacks. The one we focus on is heterogeneity in returns to wealth and potential correlation of returns with wealth.

The capitalization approach carries a measurement error $(\hat{w}_{it}/w_{it} = r_{it}/\bar{r}_t)$ whose importance depends on the extent of heterogeneity in returns to wealth and on the correlation between returns and wealth. Consider first the case of independence. Both the Gini coefficient of imputed wealth and the top wealth shares are likely to be higher than the corresponding statistics computed for actual wealth. Colombi (1990) shows analytically that if the distribution of net returns is lognormal and wealth is Pareto, then the Gini coefficient of $\hat{w}_{it}$, $G(\hat{w})$, exceeds the Gini coefficient of actual wealth, $G(w)$. Moreover, $G(\hat{w})$ is monotonically increasing in the standard deviation of individual returns to wealth. Similarly, Saez and Zucman (2015) show that returns heterogeneity overstates top shares by a factor $(Er_i^{\alpha})^{1/\alpha}/\bar{r}>1$ (by Jensen's inequality) if wealth is Pareto with shape parameter $\alpha > 1$. If one again assumes that net returns are lognormal, then this factor equals $e^{(\alpha-1)\sigma^2/2}$, which is similarly increasing in the standard deviation of returns σ. Hence, the higher the heterogeneity in returns, the greater the overstatement of wealth inequality and wealth concentration from the capitalization method.

To our knowledge, there are no analytical solutions in the more realistic case of lognormal gross returns and correlation between returns and wealth. In the Appendix we conduct Monte Carlo simulations to assess the direction and size of the biases. We assume that wealth is distributed Pareto with shape parameter α=1.3 (a value that fits Norwegian data) and gross returns are lognormally distributed in the cross-section. We first confirm that the bias in the Gini and the top wealth shares using the capitalization method increases with the standard deviation of returns in the case of independence (see Figure A1). The biases can be substantial. For example, if the standard deviation and mean of net returns are σ=0.04 and μ=0.03, respectively (as in the data, see next section), the bias in the Gini ($G(\widehat{w})/G(w)$) can be as high as 25%; for the top 5% wealth share, the bias is 15%. When we assume that returns and wealth are correlated, holding returns heterogeneity constant, our Monte Carlo simulations show that even a mild correlation between returns and wealth can generate significant gaps between the Gini and the top wealth shares from imputed and observed wealth. For example, if σ=0.04, μ=0.03 and the correlation coefficient between wealth and returns is 0.05, the Gini of imputed wealth is 35% higher than the one on actual wealth; the top 5% share based on imputed wealth is 35% higher than the one based on actual wealth (see Appendix, Figure A2).[3]

Needless to say, if the cross sectional variance of returns and/or the correlation between individual returns and wealth vary over time, inequality measures from imputed wealth may exhibit trends even when none are present in the true distribution of wealth.

---

[3] For moderate correlation, the overstatement is on average more pronounced for the top 1% and top 0.1% shares. Moreover, the bias in the Gini and the top wealth shares increases with the Pareto shape parameter α.

3. How large is the bias?

Assessing the size of this bias empirically requires observing actual wealth, capital income from tax records and the individual returns to wealth using the same population coverage. Our Norwegian data satisfy these requirement. As discussed in Fagereng, Guiso, Malacrino and Pistaferri (2016), because Norway has a wealth tax, Norwegians report all their assets, real and financial, to the tax authority, allowing us to obtain a measure of w_{it}. From the same source we also obtain capital income for various wealth categories. Data are available for 21 years, from 1993 to 2013. Hence we can construct imputed wealth as in Saez and Zucman (2015) and compare inequality measures from imputed wealth with those from actual wealth. Because we can compute individual returns (as $r_{it} = \frac{y_{it}}{(w_{it}+w_{it+1})/2}$), we can assess how heterogeneity in returns and their correlation with wealth affect the inequality measures of imputed wealth.[4] Fagereng, Guiso, Malacrino and Pistaferri (2016) document substantial heterogeneity in individual returns to wealth and provide a thorough analysis of the statistical properties of wealth returns.[5]

Panel A of Figure 1, taken from Fagereng et al. (2016), plots the cross-sectional mean, median and standard deviation of individual return on wealth for each of the 20 years covered by the Norwegian data. Panel B plots instead the median return for the 10[th] and 90[th] percentiles of the wealth distribution against time, a simple way to visualize the association between wealth and its returns.

---

[4]If the assets observed at the beginning of the calendar year were held until the end of the year, then dividing the income from capital by the beginning of period stock of wealth would provide the correct measure of the return to wealth. However, if assets are traded during the year, the income from capital will reflect the part earned over the holding period of the asset. Moreover, for tax purposes capital gains and losses are reported only when they are realized. To reduce the incidence of these errors, we do three corrections. First, we define returns as the ratio of income from capital and the average stock of wealth at the beginning and end of year. Second, we drop people with less than NOK 3000 in financial wealth (US $345). Third, we trim the distribution of returns in each year at the top and bottom 1%. These are conservative corrections that, if anything, reduce the extent of return heterogeneity.

[5]They show that not only returns are heterogeneous, but also that they have a persistent component both within and across generations. Persistent heterogeneity in returns to wealth is key to explain the long right tail in the wealth distribution (Benhabib, Bisin and Zhu, 2011).

Individual returns average around 3% (2.43% at the median) and show substantial heterogeneity over the sample period. The standard deviation of individual returns ranges between 2.5% in 2009 and 6.1% in 2005. Similarly, the range (not shown) between the 90[th] and the 10[th] percentile varies, depending on the year, between 3.2 and 7.5 percentage points. Interestingly, the extent of heterogeneity first increases and then trends downward.[6] Panel B shows that median returns increase with the household's position in the wealth distribution. For example, households in the 90[th] percentile have median returns that are often twice as large as those of households in the 10[th] percentile, and the correlation with wealth also varies over time.[7]

To summarize, the Norwegian data used by Fagereng et al. (2016) show that heterogeneity in returns and correlation of returns with wealth are a non-negligible feature of the data and both vary over time. Hence, measures of wealth inequality from capitalized tax returns can not only overstate true inequality, but also show dynamics that may differ significantly from those in actual data.

The potential importance of these biases is documented in Figure 2. In Panel A we plot the Gini coefficient of imputed and actual wealth. Panel B plots the top 5% share of imputed and actual wealth. The imputation follows closely the methodology used by Saez and Zucman (2015). In particular, we allow for heterogeneity in returns across asset classes, and impute individual wealth as $\widehat{w}_{it} = \sum_c \frac{y_{it}^c}{\bar{r}_{ct}}$ where c denotes the asset class c, y_{it}^c is the income for that asset class and $\bar{r}_{ct}$ the mean return (consistent with aggregate actual wealth). We define wealth as financial wealth and consider two assets – safe assets (deposits, bonds and outstanding claims) and risky assets (stocks traded on the Oslo Stock Exchange, shares in private businesses and mutual funds). We exclude

---

[6] Fagereng et al. (2016) show that heterogeneity is not due (only) to differences in average returns across asset classes. It remains significant even within asset classes. The standard deviation of returns on fixed income ranges between a minimum of 0.01 and a maximum of 0.04. That on risky assets (listed stocks and private business) is - not surprising - much higher (ranging between 0.06 and 0.17).

[7] Digging deeper in this correlation, Fagereng et al. (2016) show that it is mostly driven by the high returns at the very top of the wealth distribution.

housing wealth because housing data before 2010 are incomplete. As Figure 2 shows, the Gini coefficient on actual wealth shows mild fluctuations but no trend. The Gini implied by the capitalization method reveals no overall trend either; however, it overstates actual wealth inequality in all years. The average difference is 0.034 – much larger than the standard deviation of the true Gini over the 20 sample years (0.005). In some years the gap is substantial, as in 2004 (0.072). Furthermore, and perhaps most importantly, while the Gini on imputed wealth seems to track well the actual Gini at the beginning and particularly at the end of sample, the Gini coefficients of imputed and actual wealth show divergent dynamics in the middle of the sample period. While actual wealth inequality is slightly falling between 2001 and 2003, using imputed wealth would suggest increasing inequality.[8] Besides the Gini coefficients, the capitalization method overstates the top 5% wealth share in all years. The average difference between the top 5% share of imputed and actual wealth is 5.4 percentage points, but again in the middle of the sample period the overstatement is larger and the dynamics reversed. Over the whole period the actual top share is fairly flat; this absence of long-term dynamics is also reflected in the Gini of imputed wealth. The bias on top 1% and top 0.1% is milder, and sometimes negative, possibly because of the differential sensitivity of top shares to returns correlation (see Appendix and Figure A3).

4. What Drives the Dynamics of the Gap?

As argued in Section 2 both heterogeneity in returns and correlation between returns and wealth can overstate measured inequality from capitalized tax returns. Figure 1 shows that both features

---

[8] The 2006 shareholder income tax reform, which was announced as early as 2001, introduced a 28% flat tax rate on individual dividend income. This is probably responsible for the big discrepancy in the years prior to 2005. In fact the incentive to anticipate dividends has increased the correlation between returns and wealth magnifying the discrepancy between wealth estimated from tax returns and actual wealth. See Alstadsæter and Fjærli (2009).

are present in the data. Table 1 runs simple OLS regressions of the difference between $G(\hat{w})$ and $G(w)$, and of the difference between the top 5% share of imputed and actual wealth, against the standard deviation of individual returns and their correlation with wealth.

In columns (1) and (3) we control only for the standard deviation of returns, and find that both gaps increase with the extent of return heterogeneity. When we add the return correlation with wealth (columns (2) and (4)), we find that variation over time in the gap between imputed and actual Gini and the imputed and actual top 5% share are mostly sensitive to time variation in the correlation, while the standard deviation of returns turns statistically insignificant. Jointly, they explain around 82% of the total variation. We calculate that a one percentage point increase in the correlation coefficient increases the gap in the top 5% share by 1.3 percentage points and that in the Gini coefficient by almost 1 percentage point. We conclude that systematic heterogeneity of returns across the wealth distribution can generate large and time varying gaps between measures of inequality based on imputed and actual wealth.[9]

TABLE 1. EXPLAINING THE GAP BETWEEN IMPUTED AND ACTUAL INEQUALITY, 1994-2013

	$G(\hat{w}) - G(w)$		$S(\hat{w}) - S(w)$	
	(1)	(2)	(3)	(4)
St.dev.	0.81*	-0.15	1.55*	-0.24
	(0.44)	(0.24)	(0.83)	(0.46)
Correl.		0.69***		1.29***
		(0.09)		(0.17)
R^2	0.16	0.83	0.16	0.82

---

[9] We find similar evidence when looking at the difference between the top 1% and top 0.1% wealth shares (see Online Appendix, Table A1).

5. Conclusions

This paper has analyzed the effect of heterogeneity in asset returns and their correlation with the level of wealth on the discrepancies between measures of inequality based on actual wealth and those based on imputed measures of wealth obtained from capitalized tax returns, as in Saez and Zucman (2015). These factors may potentially lead to misleading conclusions about the level and the dynamics of wealth inequality.

REFERENCES

Alstadsæter, Annette, and Erik Fjærli. 2009. "Neutral Taxation of Shareholder Income? Corporate Responses to an Announced Dividend Tax." International Tax and Public Finance, 16(4): 571–604.

Benhabib, Jess, Alberto Bisin, and Shenghao Zhu. 2011. "The Distribution of Wealth and Fiscal Policy in Economies with Finitely Lived Agents", Econometrica, 79: 123–157.

Colombi, Roberto. 1990. "A New Model of Income Distribution: The Pareto-lognormal Distribution", in Income and Wealth Distribution, Inequality and Poverty, Dagum C. and Zenga M. (Eds.), Springer: Berlin.

Fagereng, Andreas, Luigi Guiso, Davide Malacrino, and Luigi Pistaferri. 2016. "Wealth returns dynamics and heterogeneity", Stanford University mimeo.

Giffen, Robert. 1913. *Statistics*. McMillan: London.

Meyer, Bruce D., Wallace K. C. Mok, and James X. Sullivan. 2015. "Household Surveys in Crisis", Journal of Economic Perspectives, 29(4): 199–226.

Saez, Emmanuel, and Gabriel Zucman. 2015. "Wealth Inequality in the United States since 1913: Evidence from Capitalized Income Tax data", Quarterly Journal of Economics, forthcoming.

FIGURE 1. TIME VARIATION IN RETURNS TO WEALTH, HETEROGENEITY
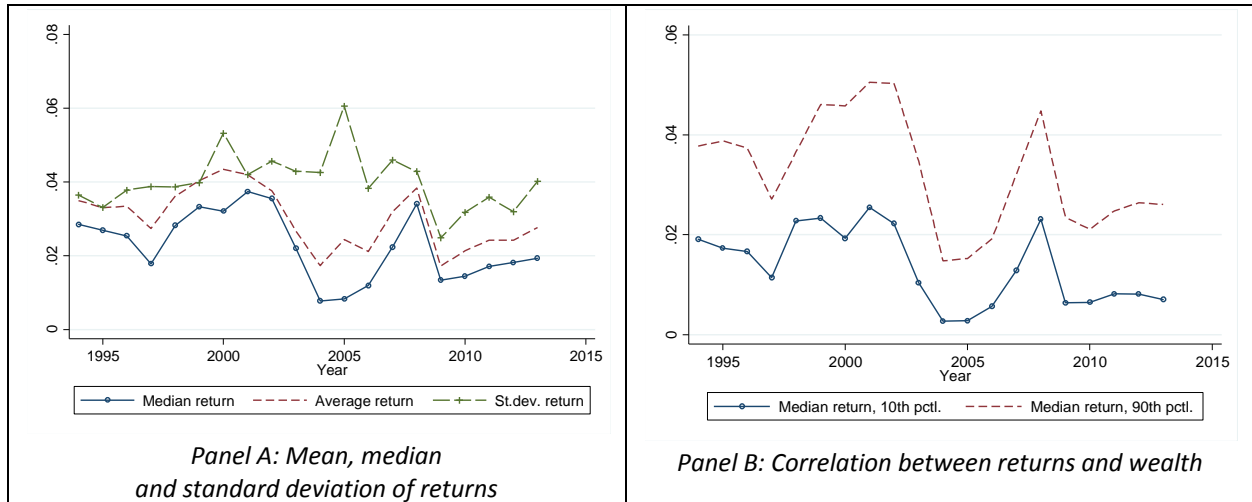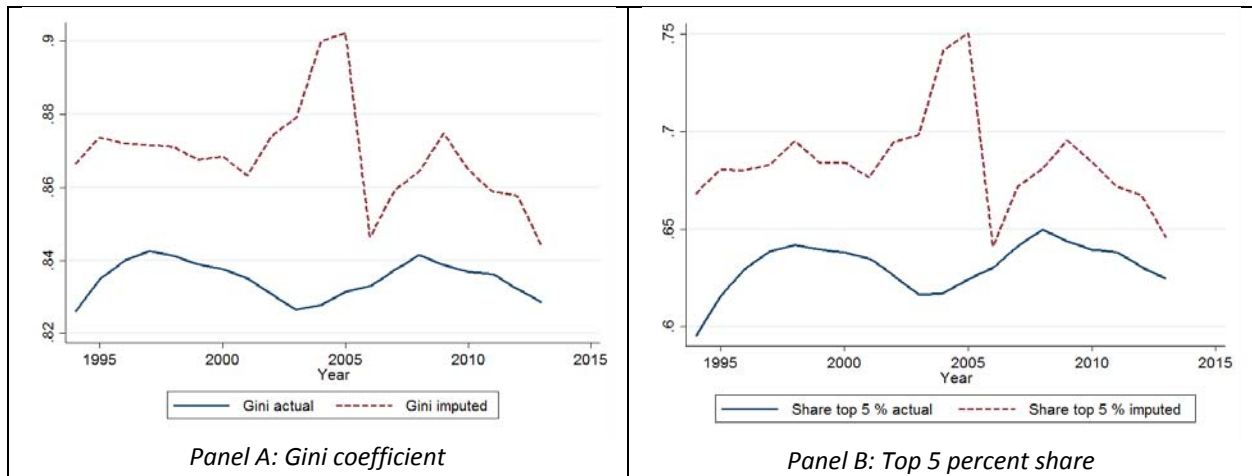AND WEALTH-RETURNS CORRELATION

Panel A: Mean, median
and standard deviation of returns

Panel B: Correlation between returns and wealth

FIGURE 2: THE GAP BETWEEN INEQUALITY OF ACTUAL AND IMPUTED WEALTH

Panel A: Gini coefficient

Panel B: Top 5 percent share

Appendix

This appendix complements the evidence shown in the text.

1. Simulations

Figure A1 and Figure A2 show the results of Monte Carlo simulations based on 200 replications and 100,000 individuals per sample. In Figure A1 we assume that wealth and returns are independent. Figure A2 relaxes this assumption. In both cases, we start by generating two standard normal (with correlation coefficient $\tilde{\rho}$, which we set equal to 0 in Figure A1). From these, we then generate a wealth distribution that is Pareto with shape parameter α and a distribution for the gross returns $(1 + r)$ that is lognormal. We assume $\alpha = 1.3$ (a value consistent with the data), and that the mean of net returns is 0.03. In Figure A1 we look at the bias induced by returns heterogeneity and hence run simulations for different values of the standard deviation of returns σ. In Figure A2 we set $\sigma = 0.04$ (a value consistent with the data) and run simulations for different values of the implied (median) correlation between returns and wealth, ρ.

The top left panel of the figures plots the ratio between the Gini coefficient using imputed wealth and the same statistics using actual wealth. The other three panels repeat the exercise for measures of wealth concentration (the top 5%, 1%, and 0.1% wealth shares). Imputed wealth is constructed replicating the procedure used by Saez and Zucman (2015). In particular, to avoid imputing a negative wealth value, whenever we draw a negative value of capital income we set it to zero before computing the capitalization factor.

The figures show the median, 5th and 95th percentile ratio of the 200 draws. Two interesting patterns emerge. First, the Gini ratio and the top 5%, 1% and 0.1% wealth share ratios are both greater than 1 and increasing with the level of heterogeneity in rates of return to wealth, even when rates of return and wealth are independent (Figure A1). Second, a positive correlation between returns and wealth widens the gap between the Gini measure on imputed and actual wealth, at least at non-negligible correlation levels. For instance, in the absence of correlation, the ratio $G(\hat{w})/G(w)$ is 1.26 for a standard deviation of returns σ=0.04 (Figure A1, top left panel). Holding the standard deviation constant, a positive correlation of returns and wealth of 0.05 results in a larger ratio (1.35) between imputed and actual Gini index (Figure A2, top left panel). The gap increases more if the correlation is just slightly higher at 0.08 (ratio 1.39). As the narrow confidence intervals show, the Gini coefficient is consistently overestimated by the simulated capitalization method either when returns are independent or when they are correlated with wealth.

As for the shares, simulations results depend both on the magnitude of the correlation and the standard deviation as well as on which top share we focus on. If the correlation is large enough, the capitalization method overstates inequality when measured by top shares. However, for low correlations, capitalization can *understate* inequality when measured by the very top shares such as the top 1% or 0.1%, even when capitalization overstates inequality measured by the Gini. This is clear from the fact that the confidence band widens considerably as we look at higher fractiles of the wealth distribution. For example, when the correlation between returns and wealth is 0.01, the Gini index on imputed wealth is overstated by 28% in median and the estimation interval is very contained. On the other hand, the top 0.1% wealth share of imputed wealth, while overstated at the median, can be short of the actual in more than 1 out of 20 cases.

In other words, summarizing inequality with top shares when the capitalization method is used can generate an upward or downward bias compared to the actual top share, depending on the degree of correlation between rates of returns and wealth. This ambiguity is absent if inequality is summarized by the more comprehensive Gini coefficient. Figure A3 shows that this property is present in our data. The two panels plot the top 1% and 0.1% shares of wealth from capitalized returns and true wealth. While the Gini measure and the top 5% share of capitalized tax returns always overstate their true counterpart, the top 1% and top 0.1% sometimes overstate and sometimes understate the corresponding actual shares.

2. Regression Evidence

Both heterogeneity in returns and correlation between returns and wealth can overstate measured inequality from capitalized tax returns. The discussion in the main text shows that both features are present in the data (see Figure 1). Table A1 complements the evidence presented in Table 1 by showing also the results of OLS regressions of the difference between the 1% and 0.1% share of imputed and actual wealth on the standard deviation of individual returns and the correlation between wealth and returns.

In columns (1), (3), (5) and (7) we control only for the standard deviation of returns, and find that all three gaps increase with the extent of return heterogeneity. However, in columns (2), (4), (6) and (8) we find that the gap between imputed and actual Gini and the imputed and actual top wealth shares are mostly sensitive to variation in the correlation between individual returns and wealth, while the effect of the standard deviation of returns turns statistically insignificant. Hence, as discussed in the main text, we conclude that it is the extent of systematic heterogeneity of returns across the wealth distribution that explains the gap between measures of inequality based on imputed and actual wealth.

Figure A1. Simulating the effect of return heterogeneity on the bias in inequality measures from capitalizing tax returns: independent returns

The Figure shows the results of a Monte Carlo simulation of the ratio between the Gini coefficient and three top wealth shares using imputed and actual wealth. The imputation assumes that true wealth is Pareto with shape parameter α =1.3 and the individual gross rate of return of wealth is distributed log normally and independently of wealth in the cross section with mean μ =1.03 and standard deviation σ. The figure shows the bias as we vary the value of the standard deviation of returns. The mean return is the average return observed in the data over the 1994-2013 period. Imputed wealth is computed by capitalizing the individual returns (computed as the product between the individual rate of returns and individual wealth) using the mean rate of return. To comply with the Saez and Zucman (2015) method we set at zero negative realizations of returns.

Figure A2. Simulating the effect of return heterogeneity on the bias in inequality measures from capitalizing tax returns: correlated returns

The Figure shows the results of a Monte Carlo simulation of the ratio between the Gini coefficient and three top wealth shares using imputed and actual wealth. The imputation assumes that true wealth is Pareto with shape parameter $\alpha =1.3$ and the individual gross rate of return of wealth is distributed log normally in the cross section with mean $\mu =1.03$ and standard deviation $\sigma = 0.04$, with median correlation with wealth equal to ρ. The figure shows the bias as we vary the value of the correlation parameter. The mean and standard deviation of returns are the average and the standard deviation of returns observed in the data over the 1994-2013 period. Imputed wealth is computed by capitalizing the individual returns (computed as the product between the individual rate of returns and individual wealth) using the mean rate of return. To comply with the Saez and Zucman (2015) method we set at zero negative realizations of returns.
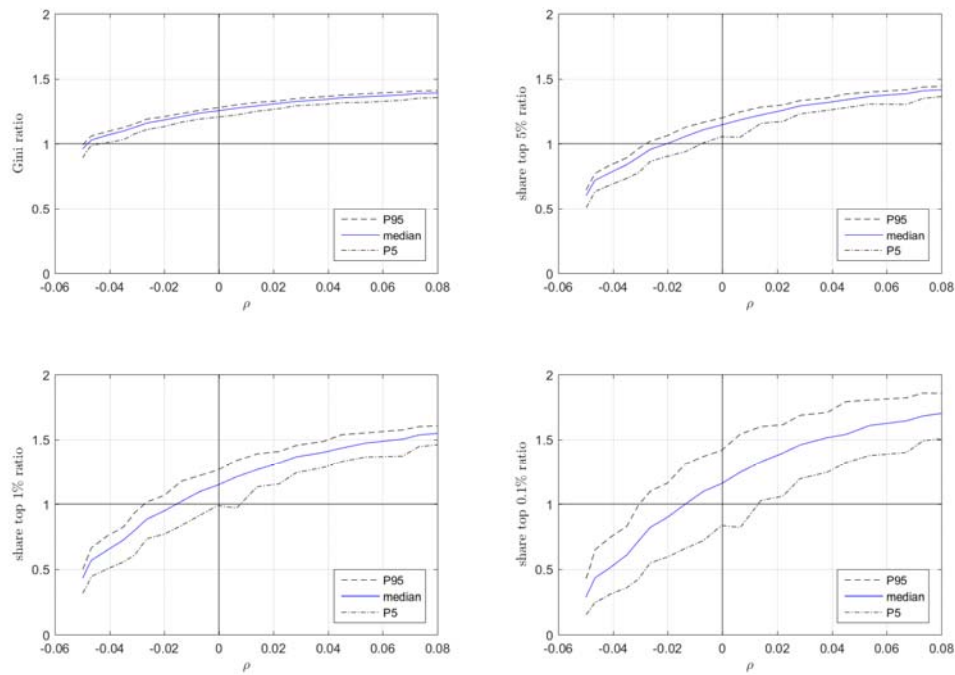
Figure A3. Shares of wealth to the top 1 and 0.1 percent of the population

The figure shows the pattern over time of the top 1% (top panel) and top 0.1% share (bottom panel) of the wealth estimated using the capitalization method and from the actual value of wealth.

Table A1. Explaining the gap between imputed and actual inequality

	$G(\hat{w}) - G(w)$		$S_5(\hat{w}) - S_5(w)$		$S_1(\hat{w}) - S_1(w)$		$S_{0.1}(\hat{w}) - S_{0.1}(w)$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
St.dev. returns	0.81*	-0.15	1.55*	-0.24	2.24	-0.50	2.45*	-0.39
	(0.44)	(0.24)	(0.83)	(0.46)	(1.30)	(0.78)	(1.37)	(0.86)
Correl. Returns/wealth		0.69***		1.29***		1.98***		2.06***
		(0.09)		(0.17)		(0.28)		(0.31)
Obs.	20	20	20	20	20	20	20	20
R^2	0.16	0.83	0.16	0.82	0.14	0.78	0.15	0.76