COMMENT BY

WOJCIECH KOPCZUK Jesse Bricker, Alice Henriques, Jacob Krimmel, and John Sabelhaus have produced very careful estimates of the magnitude and trends (from 1989 to 2013) in top wealth and income shares in the United States, relying on data from the Survey of Consumer Finances (SCF). This is of course not a new question, and the existing estimates of the top 1 percent share and the like have been highly influential, both in the economic literature and in broader public discussions. This paper adds to the existing evidence by providing high-quality estimates and by reconciling discrepancies between different methods. The authors' key contribution is their estimation of the top wealth shares, a topic on which there has been recent controversy.

Before delving into the details of the paper, it is useful to comment on the broader question of why one might be interested in wealth inequality, and in the top shares in particular. The paper's opening paragraph signals one reason: There is much popular interest in this topic. I take as given that we may be interested in inequality—but why in wealth? Wealth is a much more complicated outcome than income. Income itself does not measure the inequality of well-being or opportunities, and it comingles them with decisions about skill acquisition, occupational choice, hours of work, effort, saving, and portfolio choice. Focusing on wealth shares has the same problems, and adds some. It is inherently linked to the life-cycle dynamics of wealth accumulation—it is an outcome of the income, transfer, spending, and investment decisions that individuals make up to a particular point in time when they happen to be observed. In the natural economic approach, wealth reflects potential consumption (including that done in the form of transfers to others). Correspondingly, it is related to lifetime resources—and it does have advantages over permanent income, in that it responds to intergenerational transfers. However, if

this is the objective of analyzing wealth inequality, then one should make clear how lifetime resources and wealth are related, and should at least account for age distribution—having the same amount of wealth means something very different at age 20 than at 65. Alternatively, one may be interested in the distribution of wealth itself rather than in learning from it about the distribution of consumption opportunities. The economic rationale for separating wealth from its consumption value is more speculative, but one can certainly consider the notions of political and personal power, and of control or status, that are tied to it. Arguably, the higher one goes in the distribution, the more important these issues become, providing some cover for focusing on the top wealth shares as they are, without a more carefully specified conceptual framework. This is not a complaint about this paper—measurement is important—but just a discussant's reminder that there is a considerable distance between what we can measure and the interpretation of what wealth inequality represents.

This paper provides estimates for both income and wealth, but its findings about wealth stand out as its key contribution. This is because estimates of the top wealth shares are much less settled than those of the top income shares, and there is substantial controversy about how they have evolved in recent years.[1] The paper provides estimates using the SCF, and it offers evidence that enables us to understand the sources of the differences between these estimates and the most prominent recent alternative: the capitalization approach offered by Emmanuel Saez and Gabriel Zucman (2016). The share of wealth of the top 0.1 percent, as estimated in the paper by Bricker and his colleagues, grew between 1989 to 2013 by about 4 percentage points—from a bit under 11 percent to close to 15 percent of aggregate wealth. In contrast, the estimated share of wealth of the same group, as analyzed by Saez and Zucman (2016), doubled from the similar level in 1989 to more than 20 percent in 2013. Both methods show that wealth concentration has

1 For an extensive discussion, see Kopczuk (2015).

increased, but the difference in trends is massive. And the temporal dynamics are also different; using the SCF approach, the top 0.1 percent share fluctuated somewhat but did not change much between 1995 and 2010, so the increase over the whole period is accounted for by changes between 1989 and 1995 and since 2010. In contrast, the capitalization approach shows relentless growth, with just a short break in about 2000.

With such a large difference in results, one would expect there to be a smoking gun as evidence for what is going on—and there is one here: Looking at the composition of assets of the top groups, the bulk of the discrepancy is due to the amount of fixed-income assets that both approaches yield. Both Saez and Zucman (2016) and I (Kopczuk 2015) have noted this discrepancy before, and this paper makes it clear that this is the mechanical source of the differences. How does it come about?

First, let us start with a potential problem on the SCF side. There is a discrepancy between capital income in the data from the Internal Revenue Service on which Saez and Zucman (2016) rely and what is observed in the SCF data. This could potentially mean that the SCF is not accurately capturing the very top of the distribution. This is certainly true in a narrow and obvious sense; the SCF explicitly excludes those individuals on the *Forbes* 400 list (to preserve confidentiality), but this particular issue is explicitly dealt with in wealth estimates by adding the estimated wealth of this group to the top shares.[2] There is an extensive and very informative discussion in the paper about the approach to and quality of sampling in the SCF that

2 Note, however, that this approach takes at face value the estimates of net worth reported in *Forbes* publications. There are reasons to be skeptical about precision here; these estimates sometimes mix the wealth of a whole family with an individual's wealth, and they may miss some components of net worth, in particular debt. Raub, Johnson, and Newcomb (2010) compared the *Forbes* estimates with estate tax reports for individuals who died while on the list and found that reported estates are only about 50 percent of the *Forbes* numbers. Though some of this may reflect tax avoidance, the magnitudes are substantially larger than existing evidence of the extent of tax avoidance (for a discussion, see Kopczuk 2013), suggesting that *Forbes* is likely to somewhat overestimate the net worth of these individuals. Hence, if anything, I suspect that the approach taken by the *Forbes* list leads to upward bias in estimated top shares.

compares presurvey income tax information for respondents and nonrespondents. This discussion indicates that the role of the sampling bias is limited, though it cannot prove it plays no role. In principle, it is still possible that even though respondents and nonrespondents are similar in prior years, their income trajectories could potentially diverge in the survey year (and perhaps be related to the reason for the difference in response behavior). This discussion is also limited to sampling for the 2013 SCF, leaving open the possibility of changes in the quality of the SCF's coverage. However, if by 2013 the survey is of a high quality, then the improved coverage of the top shares should strengthen rather than weaken the observed trend.

The paper's authors also note that the overall level of income of the top groups is consistent between the SCF data and the Internal Revenue Service data on which Saez and Zucman (2016) rely, and only its composition between capital income and other sources (primarily wages) differs. They speculate that the explanation may have to do with varying notions in the tax and survey data of what constitutes labor versus capital income, especially for business owners. I am quite sympathetic to this argument—as any public finance economist working on capital taxation knows, the line between labor and capital is inherently imprecise, and it is certainly possible that tax accounting differs from the common-language way of separating labor from capital. I also find persuasive the argument that the close match of the overall income concentration measures suggests differences in the classification of income rather than bias. Still, at the end of the day, there is a difference in capital income observed in the two sources, and this is clearly an important future research area for improving our understanding of the SCF's concepts and quality of sampling. Also, perhaps more can be done with the existing data to further explain which components of capital income are a problem and how these discrepancies evolved over time.

The alternative explanation for the discrepancy has to do with how capitalization estimates are constructed. In Saez and Zucman's (2016) capitalization approach, observed capital income must be multiplied by a capitalization factor in order to arrive at the underlying level of wealth. Thus, if unobserved asset worth *A* generates observed return *rA*, one needs to multiply *rA* by the capitalization factor, 1/*r*, to arrive at the original stock. If realized *r* were known, this would be an uncontroversial—and trivial—procedure. However, *r* varies over time, it varies on average across asset classes, and it varies across individual portfolios within an asset class. In a nutshell, Saez and Zucman's (2016) approach is to use aggregate information about flows and stocks by asset classes to construct average capitalization factors, while assuming that they do not vary across income distributions and providing a battery of approaches and outside data to test sensitivity. This procedure still allows for differences in rates of return across income groups, because their portfolio compositions might differ, but this can only be due to differences in portfolio composition across very broad asset classes, which include fixed income, equities, business assets, and housing—categories that match the limited level of detail observed in data on income tax returns.

For all this approach's reliance on microeconomic data, the capitalization factor for a particular asset class is a single number for a particular year, which is constructed on the basis of aggregate data. Any bias in this factor skews the estimated value of the whole asset class. Any bias in its trend generates a trend in the estimated value of the underlying asset. In an environment with a low rate of return, a seemingly small bias in the estimated rate of return has large consequences. The capitalization factor for taxable interest income used by Saez and Zucman (2016) for 2009 is 96.6, which corresponds to the estimated rate of return of 1.04 percent for the asset class that it reflects in the economy as a whole. Hypothetically, imagine that

we are underestimating the true rate of return by 1 percentage point. In this case, the true capitalization factor would be 50 (or, if it were instead an upward bias, it could be 2,500 . . .), so the assumed 96.6 capitalization factor would erroneously double the amount of wealth estimated in this particular asset category! In an environment with a higher rate of return, however, the implications of mismeasurement will be more benign. In the 1990s, the capitalization factor for taxable interest income was about 25. In that case, increasing the corresponding rate of return by 1 percentage point, from 0.04 to 0.05, would modify the capitalization factor to 20—still a bias, to be sure, but the value of the assets would be overestimated by 25 percent rather than 100 percent.

Moving beyond hypothetical situations, the paper's figure 9 shows that directly observed rates of return on some fixed-income assets (Treasuries, bonds) are higher than those implied by observed interest income on individual income tax returns, as analyzed by Saez and Zucman (2016), so that relying on them would translate into large differences (by a factor of 2 or more, by the end of the period) in capitalization factors. The paper's figure 10 then shows that reducing the capitalization factor for fixed-income assets brings the estimates of the SCF and Saez and Zucman (2016) much closer to each other, especially in the 2000s, when they track each other fairly closely.

The paper's authors suggest that the overestimation of the capitalization factor is the reason for the discrepancy in fixed-income estimates that constitutes the bulk of the difference. I have also suggested so in the past (Kopczuk 2015), and thus—not surprisingly—I concur. The key series for me are those capitalization factors that rely on the linked estate and (pre-death) income tax data; this approach constructs the rate of return that is specific to a high-net-worth population and, in particular, it reflects a wealthy-specific portfolio composition within asset

classes. One can still worry about the quality of information for the estate tax versus the income tax, the timing of when income and wealth are observed, and the representativeness of those who died for the whole wealthy group. However, the fact that it moves closely in sync with the Treasury rate and its growing discrepancy with the series assumed by Saez and Zucman (2016) over the 2000s strongly suggest the existence of a trending bias in their capitalization factor. If the capitalization factor based on the estate income rate of return was the approach used in the baseline figures of Saez and Zucman (2016) (rather than that reported in their figure B27b, on page 79 of their 385-page-long online appendix[3]), we would be left with an understanding of the remaining discrepancies in figure 10 rather than of the major differences in trends shown in figure 1.

Having said this, the remaining and interesting question is why the interest income observed on income tax returns would imply too low a rate of return. Let us assume that there are no problems with measuring the underlying aggregate stock of fixed-income assets. There are two main possibilities. One is that some interest income is not reported or that some fixed-income assets generate no interest income (my checking account!). The other possibility (which is closely linked) is that fixed-income assets are still a broad category that, in particular, includes checking accounts, savings accounts, certificates of deposit, and bonds. In practice, these different types of investments correspond to different rates of return, but Saez and Zucman's (2016) capitalization factor is based on the average rate of return for the whole class. The much lower implied capitalization factor, which is based on an income–estate link that is not far from the Treasury rate, suggests that the portfolios of the wealthy are tilted toward higher-yield assets (for example, bonds) relative to the general public's low-interest deposits. This would always result in bias; but in a world where the general public earns 3 percent and the top of the

3 The online appendix is found at http://eml.berkeley.edu/~saez/SaezZucman2016QJEAppendix.pdf.

distribution earns 5 percent, this bias is much smaller than in a world of 0 percent versus 2 percent earnings. If, for simplicity, each group had half the aggregate assets, we would be back to my original example, with average rates of 4 percent and 1 percent and a 1 percentage point difference between the average rate of return and the one that should be used for the wealthy population.

I am not aware of any outside evidence (other than the capitalization method) that would indicate that between 2000 and 2012, the top 0.1 percent did indeed rebalance their portfolios to increase their holdings of fixed-income assets from 21 to 43 percent of their net worth, as implied by the approach taken by Saez and Zucman (2016, table B5b). This finding is driven by declining fixed income, multiplied by strongly increasing capitalization factors. Given the issues with constructing the capitalization factors, I find the evidence in this paper that indicates no such rebalancing in the SCF much more plausible.

In conclusion, this very valuable paper provides timely and careful estimates of the top wealth shares and makes a persuasive argument for the source of the discrepancy between these results and those of Saez and Zucman (2016). This is not a mortal blow to the capitalization method; nor is it intended to be one. The two methods are certainly complementary, and one way of describing the discrepancy's source is that it is due to a particular implementation of the capitalization method rather than the method itself. Adjusting capitalization factors to match the portfolios of the rich is certainly a feasible task. However, the paper does highlight how the capitalization approach is very sensitive to hard-to-estimate capitalization parameters and how the assumption of the constant rate of return across income groups for broad asset classes is potentially problematic. This approach is also heavily based on tax reporting, with all its associated conceptual problems. Hence, I view it as a complement to approaches that are based

on observing wealth directly (such as surveys, the administrative data on wealth available in some countries, and estate tax data) rather than the preferred alternative. In the United States, the SCF remains the prime source of information for understanding wealth distribution.

REFERENCES FOR THE KOPCZUK COMMENT

Raub, Brian, Barry Johnson, and Joseph Newcomb. 2010. “A Comparison of Wealth Estimates for America's Wealthiest Decedents Using Tax Data and Data from the *Forbes* 400.” In *Proceedings of the 103rd Annual Conference on Taxation*. Washington: National Tax Association.

Kopczuk, Wojciech. 2013. “Taxation of Intergenerational Transfers and Wealth.” In *Handbook of Public Economics, Volume 5*, edited by Alan J. Auerbach, Raj Chetty, Martin Feldstein, and Emmanuel Saez. Amsterdam: North-Holland.

———. 2015. “What Do We Know about the Evolution of Top Wealth Shares in the United States?” *Journal of Economic Perspectives* 29, no. 1: 47–66.

Saez, Emmanuel, and Gabriel Zucman. 2016. “Wealth Inequality in the United States since 1913: Evidence from Capitalized Income Tax Data.” *Quarterly Journal of Economics* 131, no. 2: 519–78.