2009

A Response to Seth H. Giertz

Daniel Feenberg
National Bureau of Economic Research, feenberg@nber.org

Published in TAX POLICY LESSONS FROM THE 2000s,
ed. Alan D. Viard (Washington, DC: AEI Press, 2009).

6

A Response to Seth H. Giertz

Daniel Feenberg

There are two aspects of Seth Giertz's excellent chapter that I want to talk about. One is slightly technical; I want to try to provide some explanation for why estimating elasticity of taxable income (ETI) is so difficult. I think this difficulty is unappreciated by nonspecialists, who are quick to latch onto a favorite estimate without understanding the weaknesses in the estimation. The other aspect is a bit more philosophical and addresses the different functions of the partial equilibrium analysis done here and the general equilibrium work done a few years back in the macro group at the Congressional Budget Office (CBO). Perhaps surprisingly, I strongly endorse the partial equilibrium approach taken here for the comparison of tax reforms.

There is a section of the chapter called "Issues That Complicate Estimation," but the complex solutions offered by the works Giertz cites may be raising standard errors more than they are reducing bias.

Figure 6-1 below shows the share of adjusted gross income going to the top half percent of taxpayers from 1960 to 2004; the data are from the public use files of the IRS's Statistics of Income division. Jim Poterba and I looked at an earlier version of this figure a decade ago and saw constancy before 1981 and after 1987, with a one-time jump for a transition. Figure 6-2 excludes capital gains—a more legitimate approach, in my opinion, but one that shows much the same thing. We thought the likeliest explanation for change from 1981 to 1987 was the series of Reagan tax cuts; these were sharp for the very well off and led to an increased willingness to realize taxable income within that class. We didn't think that such a quick change in the economy could be related to deunionization, globalization, or skill-biased technical change, because those are things that occur slowly.

FIGURE 6-1

SHARE OF AGI RECEIVED BY TOP 0.5 PERCENT OF AGI RECIPIENTS

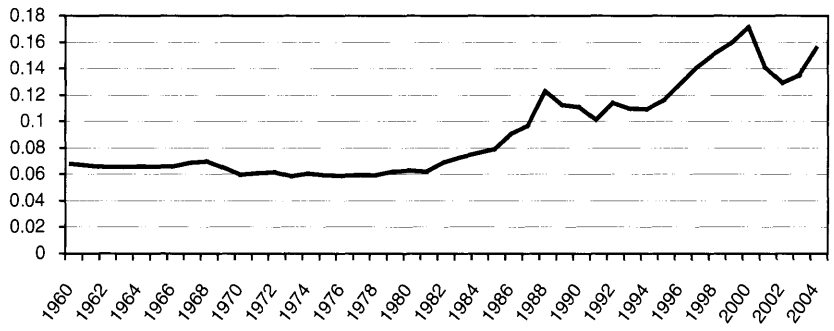

SOURCES: IRS-SOI and TAXSIM.

FIGURE 6-2

SHARE OF INCOME RECEIVED BY TOP 0.5 PERCENT OF TAXPAYERS, RANKED BY NON-GAIN INCOME

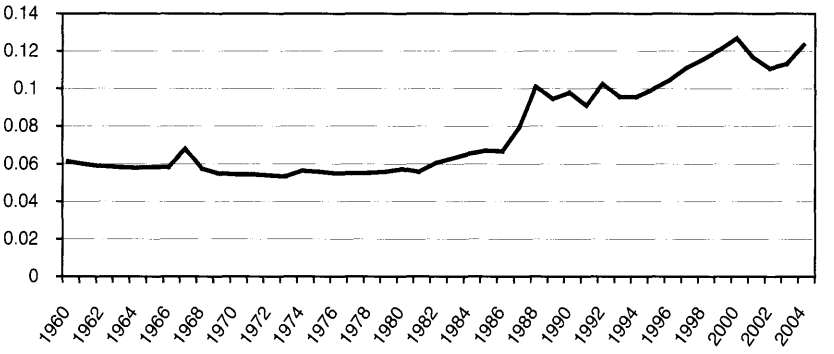

SOURCES: IRS-SOI and TAXSIM.

Of course, resting as it did on only two effective observations, our argument was not airtight. A variety of authors went ahead to use individual panel data; these are surveyed in Giertz's chapter. Hundreds of thousands or even millions of tax records should be more informative than one graph, but in fact they produce a disturbingly wide range of results.

Consider the problem of estimating the elasticity of taxable income using data from 2001 to 2005. There was a low-income group with no change in tax rate, a middle-income group with a small reduction in rates, and a high-income group with a larger reduction. It is certainly possible to treat this panel of taxpayers as a natural experiment. Regress the change in log income on the change in log net-of-tax share and you have a ready-made estimate of the elasticity of taxable income with respect to the net-of-tax share. Giertz mentions some of the problems with this regression, but I want to discuss the inadequacy of the available solutions.

Mean Reversion

Mean reversion is something that comes up in tax-price regressions but isn't much noticed in other contexts. This is not because mean reversion isn't universal—it largely is—but because it doesn't cause bias or inconsistency in the analysis of random cross-sections. But the studies cited here are mostly nonrandom panels. Those characteristics make mean reversion a problem. Consider a typical panel with 100 percent of base-year taxpayers of very high income, and one in ten thousand taxpayers with a modest income. Then the sample includes all the taxpayers going from high to low, but only one in ten thousand of those going from low to high. Even without stratification, there will still be a tendency for high-income taxpayers to be headed down, independent of the change in tax, and low-income taxpayers to be headed up. Given the correlation of income and change in rates, this will tend to bias the coefficient on the change in rates.

Moffitt and Wilhelm (2000) suggest controlling for base-period income. This makes some sense, intuitively. If high base-period income signals a likely decline in income, then adding base-period income as an explanatory variable can absorb the bias. This approach works if mean reversion is an AR-1 process with a coefficient that is constant across incomes and has

meaningful independence from the change in tax rates. If mean reversion is AR-2, then using more early years could serve to control for reversion bias, but as far as I am aware this has never been done, nor has anyone tested for the structure of mean reversion.

Mean reversion may vary across incomes, so recent authors have followed Gruber and Saez (2002) in including more general functional forms, up to and including a ten-piece spline function of income. If a single (log) linear term in income steals variance and raises the standard errors, then a flexible function of income is far more problematic. With a ten-piece spline, it is difficult to imagine that there is much independent variation in tax rates left to measure. Now in the papers Giertz cites there are some other sources of variation—changes in state taxes, differences in itemized deductions or the alternative minimum tax (AMT), etc.—that do provide some non-income-related variation. It has become the norm in economics papers to be very explicit about what is controlled for, and let the reader figure out for himself what is left over to serve for identification. In these papers, very little is left over, and it isn't obvious that the leftover variance is more independent than what is controlled for. Giertz points out that various authors' best estimates are widely spread from 0 to perhaps 1, with large standard errors. With income so partialed out, these regressions are not strong evidence that the ETI is 0. A long panel, or several concatenated panels, could ameliorate this situation—it would have periods of stability and tax rate changes in both directions. Several authors have used this approach, but it hasn't achieved really tight results, and some authors have used the long time period to make separate ETI estimates for each reform—which rather defeats the purpose.

Inequality

Another problem that Giertz brings up is the potential effect of a long-term trend in income inequality. After a decade of stability, income inequality started to grow again in 1997 without the benefit of a tax cut. A variety of authors have argued that there is long-term trend growth in inequality independent of taxes, and that this should be controlled for in regression estimates of ETI. The cure seems to be roughly the same as the cure for mean reversion: add a measure of the taxpayer's place in the income distribution

FIGURE 6-3
SHARE OF WAGES ACCRUING TO THE TOP 0.5 PERCENT OF EARNERS AND INCOME FROM NONQUALIFIED STOCK OPTIONS AS A SHARE OF WAGES

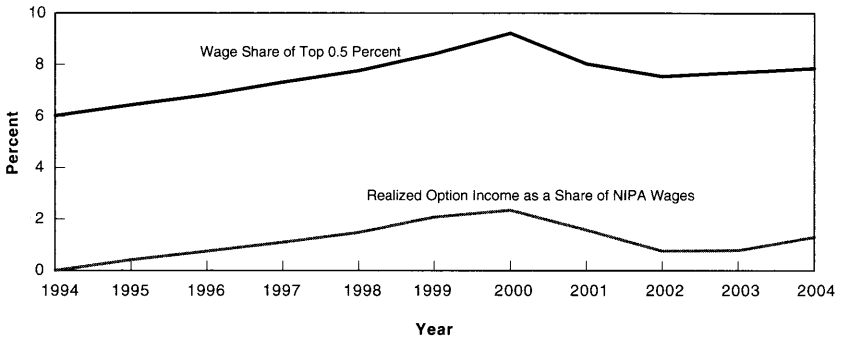

SOURCE: Congressional Budget Office using data from Scott Jaquette, Matthew Knittel, and Karl Russo, Recent Trends in Stock Options, OTA Working Paper 89 (Washington, D.C.: Department of the Treasury, Office of Tax Analysis, March 2003).
NOTES: Nonqualified stock options can be granted in unlimited amounts and, for purposes of the individual income tax, are treated the same as wages once they are exercised and the stock is purchased. NIPA = National Income and Product Accounts.

as an explanatory variable. Giertz mentions that you can't compensate for two sources of bias with one variable, which is true, but a ten-piece spline adds ten variables, so that shouldn't be a problem.

Personally, I don't think the spline is doing much that is good. My impression was always that the post-1997 increase in inequality was something new, and probably related to executive stock options. Most stock options are nonqualified and are taxed as wages to the recipient. This is the correct tax treatment, since they are deducted as wages from corporate income, but for the purpose of measuring inequality, they should really be looked at as capital gains. Data on stock options are not collected by the IRS, but the CBO (2008) produced this very interesting figure (reproduced here as figure 6-3). You can see that stock options increased by a startling amount just during the period of increased inequality, even to reproducing the downturn in 2001. So it is possible that the upturn in inequality is mostly an artifact of an overheated stock market, which can hardly become

a long-term trend. If this is the case, the spline is merely absorbing variance that could usefully be pinning down the ETI. I should add that if you control for changes in income inequality, nothing of the evidence from the first figure contributes to the result.

What are we to conclude about the best choice of a value for ETI? If we don't know the true value, is it reasonable to use zero? I don't think so. I think Giertz is right to consider a range.

Partial versus General

Giertz takes a range of estimates for ETI and considers how they affect the forecast of revenue for several plausible tax reforms. He does this in the simplest way possible—multiplying the ETI by the relative change in marginal tax rates by the base-period income. Obviously this is partial equilibrium, and doesn't account for changes in relative prices or changes in individual and government savings that might result. The CBO macro group (CBO 2004) takes a full general equilibrium approach and provides forecasts for all major macroeconomic variables. These forecasts depend on assumptions about how taxpayers, bondholders, congressmen, and foreign countries will respond to the tax change and the resulting change in the deficit. A variety of assumptions yields a wide variety of forecasts; the common theme is that behavioral considerations—if they have any significance at all—involve increased effort by taxpayers preparing for future tax increases.

The partial equilibrium answer is not a forecast of the future, but more like a price list. We can think of a menu of possible tax and expenditure changes, each with a partial equilibrium cost estimate, and let the legislature pick a budget-consistent set of choices from that menu. If the legislature respects a budget-balance constraint, the general equilibrium considerations are minimal. This approach probably makes more sense than pricing every possible combination of expenditure and financing methods. An exception would be if there were significant interactions that make the cost of one program dependent on the cost of another. While there are always such interactions, are they significant enough to justify the complexity of the general equilibrium results? Are they even as large as the discrepancies arising from stacking order issues? Those don't seem to bother anyone.

Consider a pure spending program. When a new bridge is proposed, the cost is summarized by the quantity of resources times the price of resources. The funding source isn't considered. Should it be? It is just as significant as in a tax proposal of similar magnitude.

There is a good analogy to the benefits of money economies over barter economies. The existence of money prices reduces the need to find a coincidence of wants. The existence of revenue scores plays a similar role in simplifying government budget planning. The fact that the revenue score might not take into account financing decisions is not a mark against it, but of course not an excuse to ignore the effect of the overall deficit either.

Just to show how open-minded I am, I can make an argument in favor of general equilibrium analysis. Fifty years ago London was planning the Victoria subway line, and did a cost-benefit analysis showing that the benefits exceeded the costs. However, the construction was then financed with an increase in subway fares that so depressed patronage that the benefit of the new line was wiped out. I suppose if the inefficient financing had been known in advance, the line might not have been built. Of course, in the long run, the line was justified even with the inefficient financing, but that is a different story.

According to most people with whom I have discussed the general equilibrium analysis, that analysis proves that behavioral effects don't matter; but that isn't a fair summary of the implications for making tax policy. Using the methodology of Giertz's chapter, which is conventional in a large literature on the behavioral effects of income taxation, one can make interesting comparisons of tax reforms. For instance, Jim Poterba and I compared changing the top rate to changing the AMT. We found that changing the AMT had few behavioral consequences, because marginal rates were not much affected, while changing the top rate had significant behavioral effects that changed the revenue estimate substantially. The general equilibrium analysis is really about the effects of deficit policy, not tax policy.

In summary, I think Giertz's chapter is very well done, and I am pleased to see that, within the government, well-trained and thoughtful economists are writing papers for the open literature on these topics. It is a far distance from scholarship like this to the usual secretive alchemy of revenue scoring.

References

Congressional Budget Office. 2004. *Macroeconomic analysis of a 10 percent cut in income tax rates.* Washington, DC: CBO.

_____. 2008. *Sources of the growth and decline in individual income tax revenues since 1994.* Washington, DC: CBO.

Gruber, Jonathan, and Emmanuel Saez. 2002. The elasticity of taxable income: Evidence and implications. *Journal of Public Economics* 84, no. 1: 1–32.

Moffitt, Robert, and Mark Wilhelm. 2000. Taxation and labor supply decisions of the affluent. In *Does Atlas shrug? The economic consequences of taxing the rich*, ed. Joel Slemrod, 193–234. Cambridge, MA: Russell Sage Foundation and Harvard University Press.