Feldstein, Martin, and Andrew Samwick. 1992. Social Security rules and marginal tax rates. *National Tax Journal* 45 (1): 1–22.

Fenge, Robert, Silke Uebelmesser, and Martin Werding. 2002. Second-best properties of implicit Social Security taxes: Theory and evidence. CESifo Working Paper no. 743.

Gruber, Jonathan, and David A. Wise. 1997. Social Security and retirement around the world. NBER Working Paper no. 6134. Cambridge, MA: National Bureau of Economic Research, August.

Shoven, John B. 2004. The impact of major improvement in life expectancy on the financing of Social Security, Medicare, and Medicaid. In *Coping with Methuselah: The impact of molecular biology on medicine and society,* ed. Henry J. Aaron and William B. Schwartz, 166–97. Washington DC: Brookings Institution.

Comment   Erzo F. P. Luttmer

This chapter tackles the important question of how the Social Security benefit formula can be adjusted so that it generates fewer incentives for individuals to retire early. Social Security provides retirement incentives when the additional Social Security taxes paid by postponing retirement for a year exceed the increase in the present value of future Social Security benefits from working this additional year. Goda, Shoven, and Slavov refer to this difference, when expressed as a fraction of earnings, as the implicit Social Security tax. Two features in the current Social Security law cause this implicit Social Security tax to be high for individuals with long careers. First, the current Social Security law bases benefits on the average of the thirty-five highest years of indexed earnings. Thus, current earnings will increase this average less for individuals who already have worked for thirty-five years than for individuals who have not yet worked thirty-five years because for the former group the current year's earnings crowd out a prior year's earnings in the benefit formula. Second, the progressivity of Social Security benefits depends on the average indexed earnings of the highest thirty-five years of earnings (including years with zero earnings) rather than basing this average only on those years with positive earnings. As a result, Social Security redistributes from workers with long careers to those with short careers even if these two groups have the same earnings per year worked. This redistribution further raises the implicit Social Security tax on those with longer careers.

Goda, Shoven, and Slavov analyze a reform proposal that would reduce the implicit early retirement incentives in the Social Security benefit rules. This reform would base benefits on the average of the forty highest years of

Erzo F. P. Luttmer is an associate professor at the John F. Kennedy School of Government, Harvard University, and a faculty research fellow of the National Bureau of Economic Research.

positive indexed earnings (i.e., excluding years with zero earnings) and reduce benefits pro rata based on the number of years with positive earnings for people who have fewer than forty years of positive earnings. For example, a person with twenty years of positive earnings would receive half the benefits of someone who has the same average earnings based on forty years of positive earnings. Finally, after forty years of positive earnings, one would become "paid-up," that is, exempt from paying any further Social Security taxes. Under this proposal, Social Security benefits and the Social Security tax rate are adjusted such that average benefits and average Social Security tax revenue are the same as under the current law.

Using data on a 1 percent random sample of Social Security beneficiaries in 2004 who started working after or during 1951 and claim benefits based on their own earnings record, Goda, Shoven, and Slavov carefully evaluate the reform proposal's impact on the Social Security incentive to retire as well as its distributional impact. They find that this reform would lead to a sharp reduction in retirement incentives: the implicit Social Security tax rate would fall substantially, typically by 4 to 7 percentage points, for men and women between the ages of fifty and seventy. As a result, Social Security would no longer provide this group of workers with an incentive to retire early. By design, this proposal would increase benefits for those with longer careers at the expense of those with shorter careers. However, the reform does not substantially affect the overall progressivity of the Social Security system. The proposed reform would reduce the average benefits of women relative to those of men, but this can be fixed with a minor adjustment.

Goda, Shoven, and Slavov analyze one of three implicit marginal Social Security taxes, namely the implicit Social Security tax "on postponing retirement by one year." In other words, it is the implicit tax on working this year when the counterfactual is retiring in the following year. This is probably the most plausible counterfactual for older workers. The second implicit Social Security tax is the tax on the extensive margin of working this year holding labor supply constant in the future years. In other words, it is the incentive to take one year off from working, for example, for child care or schooling reasons. This margin is probably the more relevant one for younger workers. Finally, there is the implicit Social Security tax on the intensive margin: the effect of earning one extra dollar on expected Social Security benefits net of taxes holding earnings in all other years constant.

Feldstein and Samwick (1992) also calculate implicit Social Security taxes and find, in apparent contradiction to Goda, Shoven, and Slavov's findings, that these tax rates fall with age. Goda, Shoven, and Slavov attribute this difference to the fact that Feldstein and Samwick only calculate implicit marginal Social Security tax rates for workers with at most thirty-five years of earnings. This, however, is not the reason why Feldstein and Samwick obtain different results. The difference in findings arises because

Feldstein and Samwick examine the implicit Social Security tax *on the intensive margin,* while Goda, Shoven, and Slavov examine the implicit Social Security tax on *postponing retirement.* The implicit Social Security tax on the intensive margin is lowest in those years included in the thirty-five highest years, whether or not they crowd out other years, and, therefore, this tax will remain relatively low for workers with a long work history as long as current earnings end up belonging to the thirty-five years of highest earnings. The Social Security system distorts labor supply decisions on both the intensive and extensive margins and may cause significant deadweight loss on each of these margins because of preexisting distortions. It would, therefore, be worthwhile to also examine how the proposed reform would affect these other two implicit marginal Social Security tax rates.

While Goda, Shoven, and Slavov present a compelling case that the reform proposal would improve incentives without having a major redistributive impact, it is less clear that the reform proposal is optimal. For example, the rules on the treatment of spouses imply high implicit Social Security taxes for individuals who will claim benefits based on their spouse's earnings. Perhaps the political viability of altering these rules is low, but it would be interesting to explore whether these rules can be adjusted to reduce implicit marginal Social Security tax rates with a distributional impact that is roughly neutral. More narrowly, the proposal currently analyzed by Goda, Shoven, and Slavov contains two parameters: the number of years of earnings that are included in the Social Security benefit formula and the number of years of earnings needed to reach the paid-up status. It would be relatively easy to analyze the impact of reforms that use different values for these parameters. For example, would retirement incentives be further reduced if both parameters were set at fifty?

A key component of the proposal analyzed by Goda, Shoven, and Slavov is that progressivity is based on average earnings in years with positive earnings rather than on average earnings regardless of whether earnings were positive. For practical purposes, the proposal defines positive earnings as earnings exceeding 5 percent of the earnings cap. This raises two issues. The first is practical. Holding lifetime income constant, the new proposal is more generous toward those who have more years with positive earnings. This creates incentives for individuals to shift earnings in order to attain this earnings threshold in each year. It would be useful to ascertain whether these incentives are strong enough that possibilities for gaming the system would become a serious concern. The second issue concerns the deeper theoretical point of whether income redistribution should be based on annual earnings or on lifetime earnings. As Liebman (2003) discusses, this issue is largely unresolved because it depends on how well each measure proxies for unobserved true ability. If, as seems plausible, both measures contain useful information about unobserved true ability, using a combination of both measures is optimal. Given that the majority of the

redistribution taking place through the tax and transfer system is based on annual (or sometimes even monthly) income, having some redistribution based on a measure of lifetime earnings, as currently is the case with Social Security, may well be optimal.

The chapter raises the empirical question of whether individuals understand the implicit retirement incentives from the Social Security system and whether they respond to them. The sharp break in retirement incentives induced by already having thirty-five years of positive earnings can be exploited to estimate this response. David Seif, Jeffrey Liebman, and I are currently analyzing this, and preliminary results indicate that the retirement hazard rate starts to increase sharply as soon as individuals have thirty-five years of positive earnings. This suggests that Goda, Shoven, and Slavov's concern about these implicit retirement incentives is pertinent and that the reform proposal will cause people to retire later.

Overall, this chapter makes an important contribution to the debate about Social Security reform because it makes a compelling case that a relatively straightforward and plausibly politically viable adjustment to the Social Security benefit formula can drastically reduce incentives from Social Security to retire early without major redistributive consequences. A reduction of the implicit Social Security tax will produce a first-order welfare gain because the implicit Social Security tax comes on top of other distortions, most notably from income taxation, that already encourage early retirement.

References

Feldstein, Martin S., and Andrew A. Samwick. 1992. Social Security rules and marginal tax rates. *National Tax Journal* 45:1–22.

Liebman, Jeffrey B. 2003. Should taxes be based on lifetime income? Vickrey taxation revisited. Harvard University. Unpublished Manuscript.