The Taxable Income Elasticity
and the Implications of Tax Evasion for Deadweight Loss

Jon Bakija, April 2011

I. The "Taxable Income Elasticity" Literature

Traditionally, the microeconometric literature on how taxes affect behavior focused in isolation on particular decisions that might be influenced by tax rates, such as decisions about how many hours to work. Labor supply is an important economic decision, and evidence on how it responds to taxation tells us something important about what the deadweight loss from taxation is likely to be. But taxes change incentives to engage in all sorts of behavior, and *any* of these changes in behavior might be sources of deadweight loss. For example, the incentive effects of high tax rates might cause people to do any of the following: work fewer hours; not work at all; save less; invest less; change their occupation to one that pays less but offers better non-taxable amenities (such as pleasant working conditions); be less likely to take the risk of becoming an entrepreneur; take labor compensation as non-taxable fringe benefits instead of as taxable cash wages; engage in more legal tax "avoidance" (e.g., taking advantage of loopholes in the tax law, choosing tax deductible forms of consumption over taxable forms, etc.); switch to an occupation where tax evasion is easier; move to the informal sector; or otherwise engage in more tax evasion. In the mid-1990s, Martin Feldstein (1995, 1999) started an extensive new literature in public economics with the following insight: we can potentially infer the deadweight loss arising from *all* behavioral responses to taxation by observing how taxable income (i.e., income that is fully taxable and is reported to the tax authority) changes in response to changes in marginal tax rates. If that is true, then empirical estimates of the elasticity of taxable income with respect to marginal tax rates (the "taxable income elasticity" for short) can be very informative about the deadweight loss from taxation.

Figure 1 below helps illustrate Feldstein's basic point, which is laid out in more detail in Feldstein (1999). Consider a person's decision about how much taxable income to earn and report to the tax authorities. In the absence of taxation, the marginal benefit of $1 of taxable income to the person would be $1. The marginal cost of earning $1 of taxable income is simply the marginal benefit of the next alternative that person must give up in order to get the $1 of taxable income.

To begin with a simple example, suppose the only alternative to earning taxable income is to take leisure. If, for example, the person has a wage of $10 per hour, then the alternative to earning $1 of taxable income is to take 6 minutes of leisure time.[1] In Figure 1, at all points along the marginal cost curve to the left of Q^*, 6 minutes of leisure are

[1] If your wage is $10 per hour, it takes 60 minutes / $10 = 6 minutes to earn $1.

worth less than $1 to the person. At all points along the marginal cost curve to the right of Q^* in Figure 1, 6 minutes of leisure are worth more than $1 to the person. So in the absence of taxes, the person chooses to earn a taxable income of Q^*, at which point the marginal benefit of one more dollar of taxable income is just equal to the marginal cost of earning that dollar (that is, the 6 minutes of leisure the person has to give up to get the dollar is worth just a dollar). The person would not go beyond Q^*, because beyond that point, the leisure given up is more valuable to the person, at the margin, than the additional $1 of taxable income.

Figure 1 – The deadweight loss from a change in taxable income in response to taxation

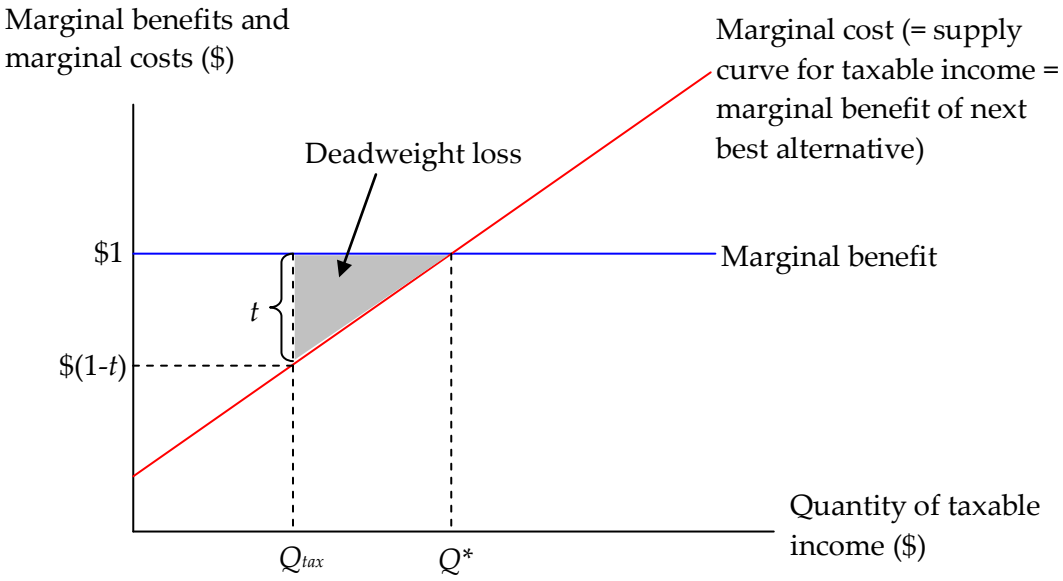

Now suppose the government imposes an income tax at a tax rate of t, say, 30%. In that case, the marginal benefit of earning $1 of taxable income is only $(1-$t$), or $0.70. When the tax is in effect, the person will only choose to earn taxable income up to the point where the marginal cost of earning that income (or in other words, the marginal benefit of the next best alternative) is $(1-$t$). So in Figure 1, when the tax is in effect, the person chooses to earn an amount of taxable income equal to Q_{tax}. For example, if the only alternative to earning taxable income is once again to take leisure, and the hourly wage before taxes is $10 per hour, the person will choose leisure over work as soon as 6 minutes of leisure is worth $0.70, which happens when taxable income is at Q_{tax}. The deadweight loss caused by taxation is then the area of the gray shaded triangle in Figure 1 above. It represents the loss of economic surplus, in dollars, arising from the fact that the tax induces people to switch away from something that in the absence of taxes was worth more to them (the taxable income) and towards something (in this example, leisure) that was would have been worth less to them in the absence of taxes.

When the problem is framed in this way, it is easy to see how the responsiveness of overall taxable income to marginal tax rates could tell us about deadweight loss arising from *all* forms of behavior which change in response to taxation. All alternatives to earning taxable income can be thought of as involving some cost that makes them imperfect substitutes for taxable income. The height of the marginal *cost* curve at each point in Figure 1 can be thought of as the marginal *benefit* of the next best alternative to earning that particular dollar of taxable income, whatever that alternative might be. Consider some examples. One alternative to earning $1 of taxable income might be to take the $1 of labor compensation as a non-taxable fringe benefit instead of as a taxable cash wage. An employer should be indifferent between providing you with $1 of taxable cash wages or $1 of non-taxable fringe benefits, because they each cost the employer $1. But given the choice, an employee with a 30% marginal tax rate would prefer that the employer provide a fringe benefit that costs the employer $1 as long as the marginal benefit of that fringe benefit to the employee is greater than $0.70. So in Figure 1, the height of the marginal cost curve at a particular point might represent the marginal value of a non-taxable fringe benefit that costs the employer $1 to provide, if that is the next best alternative to earning that particular dollar of taxable income. When the 30% tax is imposed, some of the foregone taxable income between Q_{tax} and Q^* represents switching from work to leisure, and some represents switching from cash wages to non-taxable fringe benefits. Both are sources of deadweight loss, in both cases the person is sacrificing taxable income that is worth a dollar to society (including both the benefit to the taxpayer and the value of the tax revenue going to the government) in order to substitute towards something that is worth less than a dollar.

Consider another example. Suppose that one alternative to earning and reporting taxable income is to pay a fancy tax lawyer to help you take advantage of a complicated tax shelter that enables you to receive income without having to pay tax on it. This is an example of "tax avoidance." Tax avoidance reduces the amount of taxable income you report to the tax authorities. The marginal benefit of shielding $1 from taxation in this way is less than $1, however, because it involves some real resource costs (for example, the opportunity cost of the time of the tax lawyer). So in Figure 1, the height of the marginal cost curve at a particular point might represent the marginal value of earning $1 but shielding it from taxation through a tax shelter, if that is the next best alternative to that particular $1 of taxable income. In our example, the person will choose to engage in the tax shelter as long as the net benefit to him or her from doing so ($1 less the associated costs, such as paying the tax lawyer) is greater than $0.70. Once again, this is a source of deadweight loss, as the person is switching from something that is worth $1 to society, towards something that is worth less than $1 because of the costs of avoiding taxes such as the opportunity cost of the lawyer's time.

Now consider tax evasion. Suppose that when tax rates go up, people report less taxable income because they are engaging in more tax evasion. This could be a source of deadweight loss from taxation in exactly the same way as in our previous examples.

Evading taxes involves some costs to the taxpayer. For example, in order to successfully hide income from the tax authority, you might need to operate in cash instead of using banks. This has real costs in terms of reduced convenience and security, reduced access to credit, having to buy suitcases to carry your cash around and a safe to keep it secure, etc. You might need to switch jobs to make it easier to hide income from the tax authorities. For example, hiding income from the tax authorities might require switching from a job working for a large employer that pays you more (before taxes), to working as a self-employed person in the informal sector at lower pay. The latter is not subject to withholding and information reporting, making it easier to evade taxes, so you might choose to switch to the informal sector even if it gives you a lower pre-tax income. So the marginal benefit of earning $1 but hiding it from the tax authorities is less than $1, because of the costs associated with hiding it. Thus the height of the marginal cost curve at a particular point in Figure 1 might represent the marginal benefit of earning $1 of income but hiding it from the tax authorities (evading). When the marginal tax rate is 30%, the person would choose to hide a dollar from the tax authorities if the marginal cost of hiding the dollar is less than $0.30, because then earning the dollar and paying the costs of evasion brings higher marginal benefit than does earning the dollar and reporting it to the tax authorities.

The marginal cost curve in Figure 1 can be thought of more generally as the supply curve for taxable income. The more elastic is this supply curve, the larger is the deadweight loss from taxation. Our analysis above suggests that if we can figure out how elastic the supply of taxable income is with respect to marginal tax rates, we'd be able to get a comprehensive measure of the deadweight loss caused by taxation that includes the effects of *all* behavioral responses to taxation. This insight has led to a vast literature trying to estimate the elasticity of taxable income with respect to marginal tax rates. Many such studies involve a difference-in-differences design, exploiting situations where marginal tax rates changed for one group but not another, and examining how taxable incomes changed over time for the "treatment group" (the group experiencing the change in marginal tax rates), compared to the "control group" (the group that did not experience a change in marginal tax rates). Usually, due to the nature of the tax reforms that have happened in recent years, high-income people have been the treatment group, and middle-income people have been the control group. One of the biggest challenges in this literature has been the question of how to distinguish the effects of changing marginal tax rates from the effects of other factors that affect incomes and which are changing in different ways over time for high-income people compared to middle-income people. For instance, Feldstein (1995) found that during the 1980s, taxable incomes of the rich went up much more than did the taxable incomes of the middle class, while marginal tax rates for the rich were reduced much more than were tax rates for the middle class. This could suggest a large elasticity of taxable income with respect to marginal tax rates, but it might also be explained by all sorts of other factors that were causing incomes to change differently over time at different points in the income distribution, such as globalization, skill-biased technical change, or

developments in financial markets (for example, a stock market boom has a big impact on pay of top executives because much of their pay comes in the form of stock options). Saez, Slemrod, and Giertz (2009) offer a comprehensive, critical review of this literature, and argue that studies to date have not adequately addressed these and other concerns, so that there is still a great deal of uncertainty about the magnitude of the taxable income elasticity.

II. Does it matter whether the response to taxation is "real" or evasion?

Gorodnichenko, Martinez-Vasquez, and Peter (2009) fit squarely in this tradition of empirical work aiming to estimate the elasticity of taxable income. They use a difference-in-differences approach comparing changes in income over time for high-income Russians who experienced a big cut in marginal tax rates, relative to changes in income over time for middle-income Russians who experienced little change in marginal rates. A major innovation in their approach is to come up with an indirect way of inferring how much of the response of income to marginal income tax rates represents changes in "real" economic activity, and how much represents changes in tax evasion. The analysis above suggested that either sort of response to taxation (evasion or real) would have similar implications for the deadweight loss from taxation. However, Gorodnichenko *et al*., take into account a more recent insight from Chetty (2009), which suggests that tax evasion and changes in real economic behavior might have different implications for deadweight loss. Chetty's basic point is that part of the marginal cost to the taxpayer of evading taxes is not a real resource cost to the taxpayer, but rather a transfer to other agents in the economy. In particular, part of the cost of evading taxes is that the evader is exposed to the risk of being caught and paying a fine. So for example, suppose that hiding a particular dollar of income from the tax authorities involves two costs: $0.15 of real resource costs (e.g., buying a suitcase to hold cash that would not otherwise be needed), and $0.05 of expected fines. (The expected fine is just the fine, times the probability of being caught). If the marginal tax rate is 30%, the person will choose to go ahead and hide that dollar from the tax authorities, incurring the $0.20 cost. But in this case, the deadweight loss is not the $0.20 difference between the $1 value of the taxable income and the $0.80 value of the income if hidden from the tax authorities. Rather, the deadweight loss is only $0.15, because that is the value of wasted resources. The $0.05 expected fine just represents a transfer. It makes the tax evader worse off, but it makes the recipient of the fine better off by the same amount. For instance, the proceeds from fines might be used to keep taxes on other people lower, and they benefit in aggregate by $0.05.

Figure 2 – The deadweight loss from a change in taxable income in response to taxation, when the response is due to tax evasion

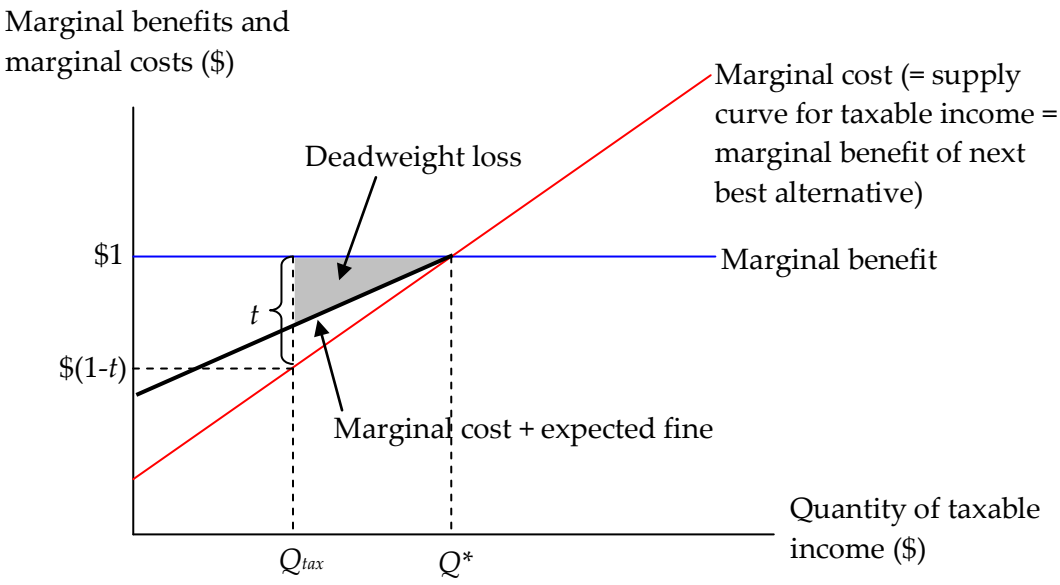

The implications for deadweight loss are illustrated in Figure 2 above. Consider a simple example where all of the change in taxable income in response to marginal tax rates represents a change in tax evasion. Remember that the marginal cost of reporting taxable income is the same thing as the marginal benefit of the next best alternative to reporting taxable income, which in this case is earning the income but hiding it from the tax authorities (evading). The marginal benefit of $1 hidden from the tax authorities is less than $1 both because of real resource costs of hiding income, and because of expected fines. If there's a 30% marginal tax rate, the person will evade as long as the cost of hiding one more dollar is less than $0.30. That leads the person to hide the income between Q_{tax} and Q^* from the tax authorities. The deadweight loss caused by each dollar of hidden income is the difference between the $1 of value that the dollar would have produced for society if it had been reported to the tax authorities (that is, the $0.70 after-tax benefit to the taxpayer, plus the $0.30 that goes to whoever benefits from the tax revenue), and the value to society that it produces for society if hidden from the tax authorities ($1, minus the cost to the taxpayer of hiding the income, plus the expected fine). In the diagram above, the thick black line represents the marginal cost of earning taxable income (or in other words, the marginal benefit of a $1 of hidden income for the taxpayer), plus the expected fine associated with hiding that dollar. In this case, the deadweight loss is the gray shaded triangle. The white triangle beneath it represents the expected fines associated with the tax evasion that actually occurs. That's a loss to the taxpayer who evades, but a gain to whoever benefits from the revenue from the fines.

Gorodnichenko *et al*. incorporate Chetty's insight when they discuss the implications of their evidence for the deadweight loss arising from personal income taxation, on p. 544. Under the assumption that tax evasion involves only real resource costs, and no transfers, their evidence suggests that the deadweight loss from personal income taxation in Russia is equal to about 13.1 percent of the revenue raised by the tax. On the other hand, if they assume that expected fines are the *only* cost to the taxpayer of evading taxes, so that all of the cost of tax evasion to the taxpayer represents a transfer to others, then their evidence implies that the deadweight loss from personal income taxation represents only 3.9 percent of the revenue collected by the tax. The large difference (13.1 percent versus 3.9 percent) reflects the fact that they estimate that a very large share of the response of income to changing marginal tax rates in Russia is due to changes in tax evasion.

An important question about the Russian flat tax reform that is not directly addressed in Gorodnichenko *et al*. is: exactly what were the consequences of the tax reform for tax revenues? It is at least possible that a reduction in tax rates could increase compliance with taxation so much that revenue actually goes up. In that case, the reform would actually be a Pareto improvement – the reform would have made the people who got the tax cut better off, without hurting anyone else, because tax revenues actually went up. We can see that personal income tax revenue in Russia did go up after the reform (*Gorodnichenko* et al. Table 1), but this is not convincing evidence because we need a counterfactual – other things were changing over time in Russia that also affected incomes and tax revenues, so we need to know what would have happened in the absence of reform. In the Gorodnichenko *et al*. study, the counterfactual is provided by observing changes in income among middle income people who did not experience large changes in tax rates. Their difference-in-differences estimates imply that taxable income went up by 19 percent more over time for the high-income "treatment" group that experienced the tax cut, compared to the middle-income "control" group that did not experience the tax cut (Table 10). To get a sense of the implications of this behavioral response for tax revenues, here is a rough "back-of-the-envelope" calculation. Consider someone who earns and reports 100,000 Rubles of taxable income before the reform. In the time before the reform, the tax was 13% of taxable income below 50,000 rubles, and 21% of taxable income above 50,000 Rubles, so this person would have paid 13%*50,000 + 21%*50,000 = 17,000 rubles in tax. Gorodnichenko's estimates suggest that, holding other influences on income constant, the flat tax reform caused such a person to increase taxable income by 19%, to 119,000 rubles, all of which would be taxed at a flat rate of 13%. This produces tax revenue of 13%*119,000 Rubles = 15,470 rubles. This implies that even with the increased reporting of taxable income, the reform caused tax revenues collected from this person to decline by 17,000 – 15,470 = 1,530 rubles. So this suggests there is a tradeoff here. The evidence suggests that the tax reform reduced tax evasion and deadweight loss, so there were benefits in terms of both efficiency and equity (i.e., better horizontal equity between evaders and non-evaders). But it also reduced tax revenues collected from high-income people relative to what would have been collected

in the absence of the reform, which is costly in terms of equity (depending on one's philosophy of distributive justice) and ability to finance public goods. One way to put this in perspective is to figure out how much revenue would have been lost if there were no behavioral response. If we imposed a 13% tax rate on an taxable income of 100,000 rubles, it would have brought in 13,000 rubles in tax revenue, a decline of 4,000 rubles relative to the pre-reform progressive tax system. The behavioral response caused revenues for this person to decline by only 1,530. So the behavioral response to the tax reform caused tax revenue to decline by only 1,530 / 4,000 = 38% of the decline that would have happened had there been no behavioral response. This suggests that the reform did cost us some revenue, but not nearly as much as would have happened in the absence of the huge increase in income that largely reflected improved compliance.

While the evidence Gorodnichenko *et al*. does suggest that the Russian flat tax involved a tradeoff (better efficiency and compliance, in exchange for reduced tax revenue from high income people), a preliminary analysis of a Polish flat tax reform by Kopczuk (2010) suggests that the Polish reform may have been very close to a Pareto improvement, making those who got the tax cut better off while costing little or no revenue.

Whether the behavioral response to taxation represents tax evasion or changes in "real" economic behavior can matter for other reasons as well. For example, Slemrod (1994) points out an important reason why the policy implications might differ. If the reason people reduce their taxable incomes when marginal tax rates go up is that they are changing their real economic behavior (for example, working less), there is not much the government can do to avoid the deadweight loss from taxation except to lower marginal tax rates. So if this is what is going on, it suggests that the deadweight loss associated with taxation is unavoidable, and therefore we might want to reduce how much government tries to do. For example, the deadweight loss would imply that redistribution is economically costly, that there is an unavoidable "leak in Okun's bucket," which reduces the optimal amount of redistribution (at least in a Utilitarian framework). If, however, the reason that people reduce their incomes when marginal tax rates go up is that they are engaging in more tax evasion or tax evasion, then the government may have another alternative to reducing marginal tax rates that avoids some of the deadweight loss – it could choose to do a better job of enforcing taxation, and reform the tax code to reduce opportunities for tax avoidance. In other words, maybe it could fix the leak in Okun's bucket, for example by improving tax administration. Improving tax administration involves real resource costs, so it does not *necessarily* enhance social welfare to do so, but it might. This also points out that the elasticity of taxable income is not some immutable behavioral parameter – rather, it can be changed by public policy. A tax system that is well-enforced, with a broad base and few loopholes, deductions, and other opportunities for tax avoidance, will tend to involve a lower elasticity of taxable income. In that case, raising marginal tax rates to pay for public goods or redistribution would involve less deadweight loss, compared to

doing so in the context of a tax system with a narrow base and lots of easy opportunities for evasion and avoidance.

References

Chetty, Raj. 2009. "Is the Taxable Income Elasticity Sufficient to Calculate Deadweight Loss? The Implications of Evasion and Avoidance." *American Economic Journal: Economic Policy* 1, no. 2: 31-52.

Feldstein, Martin. 1995. "The Effect of Marginal Tax Rates on Taxable Income: A Panel Study of the 1986 Tax Reform Act." *Journal of Political Economy*. 103, no. 3: 551-572.

Feldstein, Martin. 1999. "Tax Avoidance and the Deadweight Loss of the Income Tax." *Review of Economics and Statistics*, 81(4): 674-680.

Gorodnichenko, Yuriy, Jorge Martinez-Vazquez, and Klara Sabirianova Peter. 2009. "Myth and Reality of Flat Tax Reform: Micro Estimates of Tax Evasion Response and Welfare Effects in Russia." *Journal of Political Economy* 117, no. 3.

Kopczuk, Wojciech. 2010. "Polish Business Flat Tax and its Effect on Reported Incomes: a Pareto Improving Tax Reform?" Working Paper, Columbia University <http://econ-www.mit.edu/files/5502>.

Saez, Emmanuel, Joel B. Slemrod, and Seth H. Giertz. 2010. "The Elasticity of Taxable Income with Respect to Marginal Tax Rates: A Critical Review." Forthcoming, *Journal of Economic Literature*. (Working paper version available at http://elsa.berkeley.edu/~saez/saez-slemrod-giertzJEL10final.pdf).

Slemrod, Joel. 1994. "Fixing the Leak in Okun's Bucket: Optimal Tax Progressivity When Avoidance Can Be Controlled." *Journal of Public Economics*. 55, no. 1: 41-51.

Slemrod, Joel. 2001. "A General Model of the Behavioral Response to Taxation." *International Tax and Public Finance* 8, no. 2: 119-128.