

The Welfare Cost of a Tax, in One Number

a toy model, after Chetty (2009)

Felix Muñoz-Garcia

October 2, 2026

Measuring the welfare cost of a tax looks as if it should demand a full model of the consumer: her utility, her income effects, her entire demand system. Chetty (2009) shows it does not. The deadweight loss of a tax turns on a single reduced-form number, the compensated elasticity of demand, and on nothing else about the consumer's preferences. This note builds the simplest version of that argument and shows that two consumers with very different market demands lose the same amount to the same tax, because their compensated demands agree.

1 The model

A consumer buys a good x at price p , with every other good represented by a numeraire. The government places a specific tax t on the good, raising its price to $p + t$ and collecting revenue $R = t x(p + t)$. The cost of the tax to the consumer is the compensating variation CV of Chapter 3, the income she would need to reach her pre-tax welfare. She hands part of that loss to the government as revenue; what is lost to no one is the *excess burden*,

$$EB = CV - R.$$

Measure the compensating variation along the Hicksian demand x^h , which holds welfare fixed, and the excess burden is the familiar triangle. For a small tax, to second order:

- The compensating variation is the area to the left of the Hicksian demand between p and $p + t$. Approximating x^h by a straight line with slope $\frac{\partial x^h}{\partial p}$, this area is a rectangle plus a triangle, $CV \approx t x + \frac{1}{2} t^2 \frac{\partial x^h}{\partial p}$ (the triangle term is negative, since compensated demand slopes down).
- Revenue, evaluated along the same demand, is $R = t x^h(p + t) \approx t \left(x + t \frac{\partial x^h}{\partial p} \right) = t x + t^2 \frac{\partial x^h}{\partial p}$.
- Subtracting, the rectangles $t x$ cancel and $EB = CV - R \approx \frac{1}{2} t^2 \frac{\partial x^h}{\partial p} - t^2 \frac{\partial x^h}{\partial p} = -\frac{1}{2} t^2 \frac{\partial x^h}{\partial p}$.

Finally, writing the slope in terms of the elasticity, $\frac{\partial x^h}{\partial p} = -E^c \frac{x}{p}$, gives

$$EB \approx -\frac{1}{2} t^2 \frac{\partial x^h}{\partial p} = \frac{1}{2} E^c \frac{x}{p} t^2, \quad \text{where } E^c \equiv -\frac{\partial x^h}{\partial p} \frac{p}{x}$$

is the compensated own-price elasticity. Everything the welfare cost needs is on the right: the tax t , the point (p, x) the analyst can read from data, and the one elasticity E^c .

2 Why one number is enough

The excess burden does not contain the slope of the utility function, the income elasticity, or the shape of demand away from (p, x) . It contains E^c and the tax. So any two consumers with the same compensated elasticity at (p, x) carry the same excess burden, however far apart their preferences lie elsewhere. The compensated elasticity is the only feature of preferences that the welfare cost depends on: the analyst need not recover the structural model behind demand, only this one number.

The market demand the analyst observes is not that number. The Slutsky equation of Chapter 2 splits the market (uncompensated) elasticity E into the compensated piece and an income piece. To see it, start from the Slutsky equation in derivatives, $\frac{\partial x}{\partial p} = \frac{\partial x^h}{\partial p} - x \frac{\partial x}{\partial M}$, where M is income, and multiply both sides by $-\frac{p}{x}$:

$$-\frac{\partial x}{\partial p} \frac{p}{x} = -\frac{\partial x^h}{\partial p} \frac{p}{x} + \frac{p x}{M} \cdot \frac{\partial x}{\partial M} \frac{M}{x}, \quad \text{that is,} \quad E = E^c + s \eta,$$

with $s = \frac{p x}{M}$ the budget share of the good and $\eta = \frac{\partial x}{\partial M} \frac{M}{x}$ its income elasticity. Only E^c belongs in the welfare cost; the income term $s \eta$ moves market demand but does not affect the excess burden. Reading welfare off the market elasticity therefore mixes in an income effect that carries no welfare weight.

A numerical example. Let $p = 10$, $x = 100$, and a tax $t = 1$. Fix the compensated elasticity at $E^c = 0.5$, so the excess burden is $EB = \frac{1}{2}(0.5) \frac{100}{10} (1)^2 = 2.5$, about one-fortieth of the revenue the tax raises. Now compare two consumers at this same point. The first has no income effect (quasilinear preferences); the second has a positive income effect, which by Slutsky makes her market demand twice as elastic. For Model 2, take a budget share $s = 0.5$ and an income elasticity $\eta = 1$: then $E = 0.5 + 0.5(1) = 1$; income is $M = \frac{p x}{s} = \frac{1,000}{0.5} = 2,000$; and the income effect is $\frac{\partial x}{\partial M} = \eta \frac{x}{M} = \frac{100}{2,000} = 0.05$. Their market demands could hardly differ more, yet the welfare cost is identical.

Consumer	Income effect $\partial x / \partial M$	Market E	Compensated E^c	Excess burden
Model 1 (no income effect)	0	0.50	0.50	2.5
Model 2 (income effect)	0.05	1.00	0.50	2.5

Table 1. Same compensated elasticity, same excess burden ($p = 10$, $x = 100$, $t = 1$; Model 2 has $s = 0.5$, $\eta = 1$).

Figure 1 draws the two consumers through the same point E_0 . The tax lifts the price to 11. Along the compensated demand the quantity falls by $E^c \frac{x}{p} t = 0.5(10)(1) = 5$, to 95, and the welfare cost is the shaded triangle of area $\frac{1}{2}(1)(5) = 2.5$. Model 2's more elastic market demand falls by $E \frac{x}{p} t = 1(10)(1) = 10$, all the way to 90, so a triangle drawn under it would double the loss.




Figure 1: The deadweight-loss triangle rests on the compensated demand: the tax raises the price from 10 to 11, the compensated quantity falls to 95, and the shaded area is 2.5. Model 2's market demand (income effects) falls to 90, but the welfare cost uses the shared compensated curve, not the market one.

Lesson: The welfare cost of a tax is a triangle under the compensated demand, so it is fixed by one measurable number, the compensated elasticity. An analyst does not need the whole utility function to find it, and should not use the market elasticity, which brings in an income effect that does not change the cost. A reduced-form statistic does the work of a full structural model. $\square$

References

Chetty, R. (2009) "Sufficient Statistics for Welfare Analysis: A Bridge Between Structural and Reduced-Form Methods." *Annual Review of Economics* 1(1): 451–488.