Journal of Public Economics 131 (2015) 115–123

Contents lists available at ScienceDirect

Journal of Public Economics

journal homepage: www.elsevier.com/locate/jpube

Optimal tax salience☆

Jacob Goldin

Stanford Law School, United States

ARTICLE INFO

Article history:
Received 11 September 2014
Received in revised form 27 August 2015
Accepted 11 September 2015
Available online 21 September 2015

Keywords:
Behavioral economics
Optimal taxation
Tax salience

ABSTRACT

Recent empirical work finds that consumers under-account for commodity taxes when the after-tax price is not prominent. I investigate how policymakers may utilize such "low-salience" taxes to promote welfare. The optimal combination of high- and low-salience taxes balances two competing effects: low-salience taxes dampen distortionary substitution but cause consumers to misallocate their budgets. Using a stylized model, I show the availability of taxes with differing salience provides an extra degree of freedom that can be used to implement the first-best welfare outcome. I characterize the optimal policy and derive a formula for incremental adjustments when the first-best is unattainable.

1. Introduction

Optimal commodity taxation is a classic subject in public finance. Most research studies how governments should levy taxes across distinct goods to promote social welfare when lump-sum taxes are unavailable. In contrast, questions relating to tax design have not received the same degree of theoretical attention.[1]

Recent empirical work suggests a need to reconsider this emphasis. A series of findings suggests that the salience of a tax has important effects on consumer behavior: the less prominent the after-tax price of a good, the less consumers respond to changes in the tax on that good.[2]

Such findings suggest an additional margin through which governments can shape the behavioral effects of a tax. Although policymakers typically lack perfect control over a tax's salience, they frequently face a choice between relying on high- and low-salience ways of raising revenue. For example, policymakers can manipulate the salience of a commodity tax by choosing whether to include the tax in the displayed price of the taxed good or to add it on at the register when the consumer completes her purchase. Because the former is more salient than the latter, the government can alter the tax's salience by adjusting the degree to which it relies on the two tax designs.[3]

This paper studies the optimal salience of commodity taxes: how should a benevolent government choose between high- and low-salience taxes on a particular good to raise revenue? The analysis highlights two distinct mechanisms through which tax salience affects consumers' well-being. On the one hand, low-salience taxes dampen the type of excess burden traditionally associated with distortionary taxation: because consumers are less prone to substitute away from goods subject to low-salience taxes, such taxes are less distortionary for a given amount of revenue raised. On the other hand, low-salience taxes drive taxpayers to make optimization errors, reducing welfare by causing consumers to misallocate income among consumption goods. The government's choice between high- and low-salience taxes trades off between these competing effects.

In the standard model, the presence of an untaxed good causes optimal policy to diverge from the first-best; commodity taxes generate excess burden by distorting consumption decisions for taxed and untaxed goods. In contrast, when the government can control the salience of a tax, that flexibility provides an additional degree of

☆ For helpful comments, I am grateful to Constantine Angyridis, Raj Chetty, Jonah Gelbach, Mikhail Golosov, Joseph Grundfest, Nikolaj Harmon, Tatiana Homonoff, Louis Kaplow, Alvin Klevorick, Nicholas Lawson, Yair Listokin, David Lee, Daniel Markovits, Alex Mas, Daniel Reck, and Dean Spears. All errors are my own.

[1] By "tax design", I mean policy decisions relating to characteristics of a tax that do not directly enter into consumers' budget constraints. Two exceptions are Slemrod and Kopcczuk (2002) and Krishna and Slemrod (2003).

[2] I use "salience" to refer to the prominence of the taxed good's tax-inclusive price. For example, an excise tax included in a good's posted price is "high-salience" even though consumers may not be able to identify how much of what they pay to the retailer is tax as opposed to the pre-tax price. For empirical research relating to commodity tax salience, see Chetty et al. (2009), Feldman and Ruffle (2015), Goldin and Homonoff (2013), and Bradley and Feldman (2015). Krishna and Slemrod (2003) and McCaffery (1994) review earlier evidence on how tax design affects behavior.

[3] Policymakers may also manipulate commodity tax salience by adopting tax-inclusive pricing regulations, which require retailers to include the full amount of consumption taxes in the prices displayed to consumers. Such regulations are common in Europe but are rare in the United States. Similarly, governments may require tax-inclusive pricing for a particular good. For example, the Federal Trade Commission requires airlines to include taxes and other fees in the initial price displayed to consumers. Policymakers may also shape salience in other contexts: road tolls can be collected manually by cash transfers or automatically through an EZ-Pass system (Finkelstein, 2009); property tax payments may be collected on their own or bundled into a monthly mortgage payment to an escrow account (Hayashi, 2014); and income tax payments may be collected from employees or automatically withheld (Jones, 2010).

http://dx.doi.org/10.1016/j.jpubeco.2015.09.005
0047-2727/

freedom. I show that when the government can utilize two taxes on a single good that differ in their salience, it can employ those taxes in combination to achieve the first-best welfare outcome — even when one of the available goods cannot be taxed. The key insight is that by adjusting the balance between high- and low-salience taxes, the government can maintain a given level of revenue while causing taxpayers to vary their consumption of the taxed and untaxed good. In this way, taxpayers can be induced to choose the same allocation they would choose under a lump-sum tax (even though that allocation is privately sub-optimal given the taxes that are actually in place).

I next turn to characterizing the optimal combination of high- and low-salience taxes. Solving the government's problem yields an intuitive formula for the optimal policy, which highlights the link between optimal salience and the nature of demand for the good being taxed. Notably, the formula implies that the optimal size of the low-salience tax is always non-zero. Although low-salience taxes drive consumers to make optimization errors, the welfare costs of those errors is second-order for small values of the tax. In contrast, even small values of a low-salience tax may raise substantial revenues, allowing the government to reduce distortionary high-salience taxes while still meeting its budget constraint.

In practice, adopting policies that are designed to exploit people's biases raises several important concerns. Although many of these, such as political transparency and credibility, are outside the scope of this paper, one that can undermine the results presented here is the possibility that taxpayers will become more attentive to low-salience taxes as the government increases its reliance on them — i.e., as the utility cost of neglecting the taxes grows larger. Before concluding, I consider an extension of the model to a setting in which the salience of a tax is endogenously related to the tax's size and derive conditions under which the first-best will be attainable. When the first-best is unattainable, I show how incremental adjustments in the balance between high- and low-salience taxes can still yield efficiency gains.

Despite the ubiquity of policy decisions that affect tax salience, the topic has received little theoretical attention. As Congdon et al. (2009) conclude in their review of the behavioral tax literature, "the theoretical literature has yet to yield the type of rules of thumb with respect to optimal tax salience that translate into practical policy recommendations". The research closest to the current analysis are Chetty et al. (2009), Chetty (2009), and Reck (2015). Those authors derive formulas for quantifying the excess burden of a tax that is less than fully salient but do not consider the implications of salience for optimal taxation. In addition, this paper is the first to consider the possibility of combining tax instruments that differ in their salience, and it is that possibility which drives the theoretical insights described here.

A number of influential papers have investigated how cognitive biases other than salience affect prescriptions for optimal tax policy (e.g., Liebman and Zeckhauser, 2004; O'Donoghue and Rabin, 2003). This literature evaluates the optimal level of a tax instrument conditional on taxpayers exhibiting an assumed behavioral bias. I build on this literature by studying a setting in which the government's choice of tax instrument controls the extent to which taxpayers exhibit the bias in the first place.

The remainder of the paper proceeds as follows. Section 2 develops the model and derives the main results — first graphically and then formally. Section 3 extends the model to account for the possibility that a tax's salience is endogenously related to its size. Section 4 concludes.

2. Model and results

Society is composed of a representative taxpayer who divides her income I between goods x and y. Production of x is characterized by constant returns to scale technology so that its pre-tax price is fixed at marginal cost p. Good y is the numeraire. The taxpayer's utility depends on consumption of x and y[4]:

$$U = U(x, y) \tag{1}$$

U is concave and smooth with respect to both goods.

The government's objective is to maximize the representative taxpayer's utility while raising revenue R_0.

2.1. First-best welfare outcome

Before turning to tax salience, it is helpful to characterize the first-best welfare outcome — i.e., what the government could achieve with access to a non-distorting tax. To derive this benchmark I will assume for purposes of this section that the government can levy a (fully-salient) lump-sum tax of size L.

When facing the lump-sum tax, the taxpayer's budget constraint is given by

$$p\,x + y = I - L \tag{2}$$

and her consumption satisfies the first-order condition associated with maximizing utility subject to this constraint:

$$U_x(x,\ y) = p\,U_y(x,\ y). \tag{3}$$

Because the revenue collected by a lump-sum tax of size L is simply L, the government's revenue constraint is satisfied if and only if

$$L = R_0. \tag{4}$$

Eqs. (2)–(4) pin down consumption under a lump-sum tax and hence characterize the first-best welfare outcome.

2.2. Tax salience

Having characterized the first-best, I assume now that the government lacks access to a lump-sum tax and can only raise revenue through commodity taxes on x. Good y (the numeraire) is untaxed. The government has at its disposal two tax designs that it can levy on purchases of x: a high-salience tax t_h and a low-salience tax t_l. Both t_h and t_l are unit taxes. The taxpayer's budget constraint takes the form:

$$y + (p + t_h + t_l)\,x = I. \tag{5}$$

2.2.1. Taxpayer behavior

Taking income as fixed, demand for x and y can be written as a function of the two taxes and the pre-tax price of x: $x = x(p, t_h, t_l)$ and $y = y(p, t_h, t_l)$. To capture the empirical findings described in the introduction, I assume that the extent to which a tax affects consumer demand depends on the tax's salience. As in Chetty et al. (2009), I

[4] Expressing utility as a function of consumption is standard in public finance models but implies that other factors (such as tax salience) do not affect welfare apart from their effect on consumption. For example, an agent would violate the assumption if she preferred facing a register tax to a posted tax on political grounds, perhaps because the amount going to the government is more transparent under the former than the latter. If low-salience taxes do generate direct welfare costs to consumers the results presented here will overstate the benefits of low-salience taxes. However, as Chetty et al. (2007) show, even relatively small cognitive costs generate substantial under-reaction to a tax; consequently, omitting such costs from the model may not be as misleading as would otherwise be the case. In addition, note that not all psychic cost models are ruled out: suppose that accounting for a low-salience tax is associated with some cognitive cost, but because of that cost, the consumer rationally chooses to ignore the tax. This agent's utility function can be described by Eq. (1) because given her decision-making strategy, she does not suffer any direct utility cost when confronted with the tax.