Ramsey strikes back: Optimal commodity taxes and redistribution in the presence of salience effects*

Hunt Allcott (NYU and NBER)

Ben Lockwood (Wharton and NBER)

Dmitry Taubinsky (UC Berkeley and NBER)

January 10, 2018

Abstract

An influential result in modern optimal tax theory, the Atkinson and Stiglitz (1976) theorem, holds that for a broad class of utility functions, all redistribution should be carried out through labor income taxation, rather than differential taxes on commodities or capital. An important requirement for that result is that commodity taxes are known and fully salient when consumers make income-determining choices. This paper allows for the possibility consumers may be inattentive to (or unaware of) some commodity taxes when making choices about income. We show that commodity taxes are useful for redistribution in this setting. In fact, the optimal commodity taxes follow the classic "many person Ramsey rule" (Diamond, 1975), scaled by the degree of inattention. As a result, to the extent that commodity taxes are not (fully) salient, goods should be taxed when they are less elastically consumed, and when they are consumed primarily by richer consumers. We extend this result to the setting of corrective taxes, and show how nonsalient corrective taxes should be adjusted for distributional reasons.

*The presentation of the model and derivations of the results follow Lockwood and Taubinsky (2017); see that paper for a generalization of the results to the case in which there is correlated preference heterogeneity. Allcott: New York University and NBER. hunt.allcott@nyu.edu. Taubinsky: Berkeley and NBER. dmitry.taubinsky@berkeley.edu. Lockwood: Wharton and NBER. ben.lockwood@wharton.upenn.edu. We thank Keith Ericson for comments. We are grateful to the Sloan Foundation for grant funding.

1 Introduction

The Atkinson and Stiglitz (1976) theorem—a pillar of modern tax theory—demonstrates that for a broad class of utility functions, all redistribution should be carried out through labor income taxation. That is, differential commodity taxes are suboptimal means of redistributing from rich to poor. This canonical result, which has become conventional wisdom in many modern public finance circles, stands in contrast to the widespread use of differential commodity taxes for redistributive purposes in practice. To cite a few examples, most states exempt groceries from sales tax, health insurance and education are heavily subsidized (often in an income-dependent manner), and capital income is subject to a progressive marginal rate schedule. This raises an obvious question: is the current tax system rife with suboptimal commodity taxes? Or alternatively, is there some feature of reality that the Atkinson Stiglitz model misses, but that policy makers (and perhaps common intuitions) take into account?

This paper relaxes a key assumption underlying the Atkinson-Stiglitz theorem: that all commodity taxes are fully salient when people make income-determining decisions, such as whether to attend college, what career to pursue, or how many hours to work each week.[1] According to the logic of the Atkinson-Stiglitz theorem, a tax on (say) some luxury good reduces the appeal of attaining high earnings—since one cannot purchase as much of that good—and thereby distorts labor supply in the same fashion as an income tax targeted at the high earners who consume that good. It is better to employ an income tax directly, which at least avoids distorting consumption choices. Key to this reasoning is the assumption that the commodity tax is fully salient when income-determining decisions are made.

A wave of recent empirical evidence suggests this full salience assumption may be too

[1] This paper is one of many which relax various assumptions underlying the Atkinson and Stiglitz (1976) model. Saez (2002) demonstrates that commodity taxes are useful to the extent that due to correlated preference heterogeneity, certain kinds of consumption patterns provide additional information about individuals' earnings ability. Jacobs and Boadway (2014) shows that if labor supply and commodity consumption are non-separable in the utility function, then commodities which boost labor supply should be taxed. To our knowledge, this is the first paper, together with Rees-Jones and Taubinsky (2018) who study a simple two-type model, which maintains the utility function restrictions of Atkinson and Stiglitz (1976) while relaxing the assumption of fully salient commodity taxes.

strong. Chetty et al. (2009) finds that taxes which aren't included in posted prices are not fully salient even at the time of purchase—calling into question whether they could be fully salient at the time of income-determining decisions such as choice of profession. Moreover, various subsidies often appear to generate more muted effects on labor supply than direct income taxes.[2] These results fit more broadly into a growing literature that demonstrates that individuals often do not re-optimize their choices in response to even substantial indirect changes in policy—see, for example, Chetty et al. (2014) on the insensitivity of savings decisions to subsidies.

In exploring the possibility that non-income taxes lack salience when income-determining decisions are made, this paper builds a perhaps unexpected bridge to an earlier optimal taxation literature in the tradition of Ramsey, 1927. The canonical Ramsey framework prescribes the well-known "inverse elasticity rule," that commodity taxes should be inversely proportional of the price-elasticity of demand of the good in question. The Diamond (1975) extension to heterogeneous income-earners shows that commodity taxes should also be focused more heavily on goods consumed primarily by the rich. Although this literature had once had a profound impact, its results have now largely been dismissed due to its add-hoc assumptions about the non-existence of nonlinear income taxation.

A key result of this paper, however, is that the canonical Ramsey-style formulas turn out to be relevant in the context of non-salient commodity taxes. We show that when a commodity tax is not at all salient for income-determining decisions, the formula for the optimal commodity tax is *identical* to the "many-person Ramsey" formula derived by Diamond (1975). More generally, the optimal commodity tax follows the Diamond-Ramsey formula, but scaled down by one minus its salience.

We then extend this result to corrective commodity taxes which target externalities or "internalities" (e.g., due to present bias or poor information). In the absence of salience

[2]See Strumpf (2011) for evidence that Medicaid introduction had little impact on labor supply, Monks (2004) for a discussion of the impact of income-dependent financial aid for college on household savings, and Gallagher and Muehlegger (2011) for energy efficiency subsidies.

effects, a simple extension of the Atkinson-Stiglitz theorem in this setting shows that the optimal commodity tax is Pigovian: it should be set to the marginal externality (or internality), regardless of whether consumption of the sin good is concentrated on rich or poor consumers. In the presence of salience effects, however, we show that the tax should be lower if the good is consumed by the poor, and higher if it is consumed by the rich.

2 Model

We consider individuals differentiated by earnings ability $w \in \mathbb{R}$, distributed according to a distribution F. Individuals choose a level of labor l, which generates earnings $z = wl$, and which is taxed according to the nonlinear income tax $T(z)$. Consumers use their net income to choose a consumption bundle (c_1, c_2), which are sold at before-tax prices (p_1, p_2) and are subject to additional linear commodity taxes $t = (t_1, t_2)$. For concreteness, we consider taxes that are included in the final posted price of the good. Each individual's budget constraint is $p_1(1+t_1)c_1 + p_2(1+t_2)c_2 \leq z - T(z)$. Individuals maximize $U(c_1, c_2, l; w))$.

In the classical formulation, the policymaker's problem is to maximize aggregate utility:

$$\max_{T,t} \int U(c_1(w), c_2(w), l(w); w) dF(w),$$

subject to the government's budget constraint

$$\int (p_1 t_1 c_1(w) + p_1 t_2 c_2(w) + T(z(w))) dF(w) \geq R,$$

where R is an exogenous revenue requirement, and to individual optimization:

$$(c_1(w), c_2(w), l(w)) \in \arg\max_{\{c_1,c_2,l\}} \{U(c_1, c_2, l; w) \mid p_1(1+t_1)c_1 + p_2(1+t_2)c_2 \leq wl - T(wl)\} \quad (1)$$

Implicit in equation (1) is the assumption that when choosing labor supply l, individuals fully account for the effect of commodity taxes on the returns to labor. We modify the standard formulation by relaxing the assumption that people always consider and can correctly compute how various commodity taxes—or differences in prices, more generally—determine the implicit marginal tax rate on their labor income. That is, we allow for the possibility that some commodity taxes may be under-internalized, or even ignored entirely, when labor supply decisions are made.

A possible micro-foundation for this possibility is motivated by the sparsity-based bounded rationality model of Gabaix (2014). Since expenditures on some dimensions of consumption are relatively small, while the total tax burden paid through the income tax is relatively high, the framework of Gabaix (2014) would predict that people pay less attention to changes in taxes or prices on particular commodities than they do to changes in the income tax.

Consistent with that reasoning, we suppose consumers correctly perceive the after-tax price $p_1(1+t_1)$ of c_1, perhaps because it is a composite good that constitutes a large share of expenditures, and hence whose prices and taxes are considered carefully by consumers. But we model consumers' perceived after-tax price of c_2 as given by $(1-\vartheta)\hat{p}+\vartheta p_2(1+t_2)$ when they are making a labor supply decision, where $\hat{p}$ is a "mental default" for the after-tax price. The attention parameter ϑ captures the extent to which consumers' labor supply is sensitive to variations in prices generated by the commodity tax on c_2. However, we assume that prices are understood when consumers choose how to spend their after-tax earnings $z-T(z)$ on c_1 and c_2.[3]

3 Optimal revenue-raising taxes

To simplify exposition, we assume that $U(c_1,c_2,l;w)=u(c_1,c_2)-\psi(l)$, where u and ψ are increasing, smooth, and (respectively) concave and convex. This formulation implies that individuals' earnings ability w is not related to their preferences for consumption. Thus, any variation in c_2 consumption across the income distribution is due to income/wealth

[3]We make this assumption because we assume that the taxes t_i are included in the posted prices. When taxes are not included in the posted prices immediately but are added at the register, they are not fully salient during shopping decisions (cites).

effects, rather than consumption preferences varying with earnings ability w. The Atkinson-Stiglitz theorem guarantees that under this condition, the optimal tax structure must levy a uniform tax rate on all commodities. In the absence of this assumption, the uniform commodity taxation result need not hold, as shown by Saez (2002) and Allcott, Lockwood and Taubinsky (2018).[4] Making this assumption allows us to draw out the consequences of salience effects most crisply, as it implies that any deviation from uniform commodity taxation is due to the salience effects.

We further make the simplifying assumption that income effects on labor supply are negligible. Gruber and Saez (2002) find small and insignificant income effects and Saez et al. (2012) review the empirical literature on labor supply elasticities and argue that "in the absence of compelling evidence about significant income effects in the case of overall reported income, it seems reasonable to consider the case with no income effects."

Following Diamond (1975), we let $\alpha(z)$ denote the social marginal utility of giving a z-earner one more unit of after-tax income (see the online appendix for a formal definition). We define λ to be the social marginal value of public funds, which is just the multiplier on the government budget constraint. We let $\bar{c}_2$ denote the average consumption of c_2, and we let $\xi = -\frac{d\bar{c}_2}{dp} \cdot \frac{p(1+t_2)}{\bar{c}_2}$ denote the (aggregate) price elasticity of demand for c_2, and let ξ^c denote the compensated elasticity.

We now characterize the optimal tax-structure, generalizing the seminal Atkinson-Stiglitz theorem. For brevity, we state results only for the optimal commodity taxes, relegating a characterization of the optimal income tax to the appendix.

Proposition 1. *Let $\tau = \frac{1+t_2}{1+t_1} - 1$ be percent difference in the tax rates on c_2 versus c_1. Then in any optimal tax system τ must satisfy*

$$\frac{\tau}{1+t_2} = (1-\vartheta)\frac{1}{\lambda\xi^c}\frac{E[(\lambda - \alpha(z))c_2(z)]}{\bar{c}_2}.$$

[4] Lockwood and Taubinsky (2017) further show that in the case in which all variation in c_2 consumption is driven by preference heterogeneity rather than income effects, the formula for the optimal t_2 corresponds to the formula in the case in which $\vartheta = 0$.

and the social optimum can be implemented with t_1 and t_2 satisfying

$$t_1 = 0$$

$$\frac{t_2}{1+t_2} = (1-\vartheta)\frac{1}{\lambda\xi^c}\frac{E[(\lambda-\alpha(z))c_2(z)]}{\bar{c}_2}$$

When agents are fully rational ($\vartheta = 1$) the result in proposition 1 is the Atkinson-Stiglitz theorem, giving the uniform commodity taxation result that $t_2 = t_1$ in any optimal tax system, and that the social optimum can be implemented with zero commodity taxation.

However, the uniform commodity taxation result breaks down whenever salience of the commodity tax on c_2 is incomplete. In fact, for $\vartheta < 1$, the social optimum is implemented with a tax t_2 that follows the "many-person Ramsey tax rule" of Diamond (1975), scaled by $(1-\vartheta)$. When $\vartheta = 0$, the Diamond (1975) Ramsey tax rule is exactly replicated. .

Because in our setting lump-sum taxation is available, $\lambda = E[\alpha(z)]$ (see the online appendix for a formal proof), and thus we have $E[(\lambda-\alpha(z))c_2(z)] = -Cov[\alpha(z), c_2(z)]$. Recalling that $\alpha(z)$ corresponds to the social marginal value of giving a unit of income to a $z-$earner, $\alpha(\cdot)$ should be declining with z. This means that in the presence of income taxation and salience effects, the commodity tax on c_2 follows a simple principle: it should be positive when the lower-income people consume less of c_2 (i.e., when the tax is progressive), and it should be negative when lower-income consumer more of c_2 (i.e., the tax is regressive, but a subsidy is progressive).

4 Optimal corrective taxes

We now turn to the case in which the government has a motive to tax c_2 because consumption of c_2 generates an externality or internality. We assume that while $\frac{\partial}{\partial c_2}u$ is consumers' (perceived) marginal utility from consuming c_2, the social marginal utility from consumption of c_2 is actually $\frac{\partial}{\partial c_2}u - \chi$, for some $\chi \in \mathbb{R}$. We make the simplifying assumption that χ is homogeneous and constant, though Lockwood and Taubinsky (2017) analyze the case in which χ may be heterogeneous across consumers and nonlinear in c_2.

The optimal Pigovian tax rate on c_2 would simply be $t_2 = \chi/p_2$. However in the presence of income inequality, a common intuition holds that the tax on c_2 should be adjusted away from this Pigovian benchmark depending on whether it is progressive or regressive. Yet an extension of the Atkinson-Stiglitz theorem states that this intuition is incorrect. Under our assumptions about U, even if high earners consume more c_2 than low earners, the optimal tax on c_2 is still $t_2 = \chi/p_2$. This is because in the absence of salience effects, the burden of the commodity tax can be perfectly offset by the income tax in such a way that consumers do not change their labor-supply decisions.

In the presence of salience effects, however, we show that the above logic breaks down. Because individuals react less to the labor-supply incentives induced by t_2 than by the income tax, any change in the income tax that offsets the burden of t_2 will generate changes in labor supply. We formalize our result in the proposition below.

Proposition 2. *The social optimum is implemented with t_1 and t_2 satisfying*

$$
\begin{aligned}
t_1 &= 0 \\
t_2 &= \underbrace{\frac{\chi}{p_2}}_{\textit{Pigovian correction}} \underbrace{- (1-\vartheta)\frac{1+t_2}{\lambda \xi^c}\frac{Cov\left[\alpha(z), c_2(z)\right]}{\bar{c}_2}}_{\textit{Regressivity costs}}
\end{aligned}
$$

The first term in the formula for t_2—the "Pigovian correction"—is simply the value that the tax would take on in the absence of any redistributive concerns. The second term in the formula comes out of redistributive concerns. These concerns are immaterial when consumers are fully rational ($\vartheta = 1$), which is a generalization of the Atkinson-Stiglitz theorem to the case of externalities/internalities with consumers not subject to salience effects. In the presence of salience effects, however, if the tax is regressive and $\vartheta < 1$, then its optimal value will be below that of the Pigovian benchmark. When the tax is progressive and $\vartheta < 1$, its optimal value will be above that of the Pigovian benchmark.

5 Conclusion

Together, the results in this paper serve as an exploration of the robustness of optimal tax results derived in the prevailing Mirrlees frameworks, which assume that distortions from taxation arise from asymmetric information about individuals' ability levels. These frameworks are appealing because they allow for non-trivial (and intellectually gratifying) optimal tax derivations that do not rely on ad-hoc assumptions such as the absence of nonlinear income taxes. Consequently, they have largely replaced earlier results about optimal tax structures in the Ramsey tradition.

This paper underscores, however, that the implicit assumption of perfect rationality is a strong and limiting one, perhaps especially so when tax instruments are mathematically shown to affect labor-supply incentives in nuanced ways that are not immediately intuitive. Plausible relaxations of the perfect rationality assumption can lead to optimal tax results that are strikingly similar to earlier results in the Ramsey literature, and which are in line with non-economists' intuitions about the distributional role of commodity taxes. Our paper shows that differential commodity taxes are useful when they are not fully salient, and their optimal size follows two intuitive principles: they should decrease in the price-elasticity of the taxed good, and they should increase in the extent to which they target goods more heavily consumed by the rich.

In addition to standard measures of elasticities and regressivity, our formulas highlight the need to measure salience bias for implementing the optimal tax system. Our results provide useful quantitative and qualitative guidance for the capital income taxation, consumption taxes, and in-kind transfers.

References

Atkinson, A. B. and J. Stiglitz (1976): "Design of Tax Structure - Direct Versus Indirect Taxation," *Journal of Public Economics*, 6, 55–75.

Chetty, R., J. N. Friedman, S. Leth-Petersen, T. Nielsen, and T. Olsen (2014): "Active vs. Passive Decisions and Crowd-Out in Retirement Savings Accounts: Evidence from Denmark," *Quarterly Journal of Economics*, 139, 1141–1219.

Chetty, R., A. Looney, and K. Kroft (2009): "Salience and Taxation: Theory and Evidence," *American Economic Review*, 99, 1145–1177.

Diamond, P. (1975): "A Many-Person Ramsey Tax Rule," *Journal of Public Economics*, 4, 335–342.

Gabaix, X. (2014): "A Sparsity Based Model of Bounded Rationality," *Quarterly Journal of Economics*, 129, 1661–1710.

Gallagher, K. S. and E. Muehlegger (2011): "Giving green to get green? Incentives and consumer adoption of hybrid vehicle technology," *Journal of Environmental Economics and Management*, 61, 1–15.

Gruber, J. and E. Saez (2002): "The elasticity of taxable income: evidence and implications," *Journal of public Economics*, 84, 1–32.

Jacobs, B. and R. Boadway (2014): "Optimal linear commodity taxation under optimal non-linear income taxation," *Journal of Public Economics*, 117, 201–210.

Lockwood, B. B. and D. Taubinsky (2017): "Regressive Sin Taxes," *Working Paper*.

Monks, J. (2004): "An empirical examination of the impact of college financial aid on family savings," *National Tax Journal*, 189–207.

RAMSEY, F. P. (1927): "A Contribution to the Theory of Taxation," *The Economic Journal*, 37, 47–61.

REES-JONES, A. AND D. TAUBINSKY (2018): "Taxing Humans: Pitfalls of the Mechanism Design Approach and Potential Resolutions," *Tax Policy and the Economy*, 1.

SAEZ, E. (2002): "The desirability of commodity taxation under non-linear income taxation and heterogeneous tastes," *Journal of Public Economics*, 83, 217–230.

SAEZ, E., J. SLEMROD, AND S. H. GIERTZ (2012): "The elasticity of taxable income with respect to marginal tax rates: A critical review," *Journal of Economic Literature*, 50, 3–50.

STRUMPF, E. (2011): "Medicaid's effect on single women's labor supply: Evidence from the introduction of Medicaid," *Journal of Health Economics*, 30, 531–548.