

NBER WORKING PAPER SERIES

RAMSEY STRIKES BACK:
OPTIMAL COMMODITY TAXES AND REDISTRIBUTION IN THE PRESENCE OF SALIENCE EFFECTS

Hunt Allcott
Benjamin Lockwood
Dmitry Taubinsky

Working Paper 24233
<http://www.nber.org/papers/w24233>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2018

We thank Keith Ericson for comments. We are grateful to the Sloan Foundation for grant funding. 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. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2018 by Hunt Allcott, Benjamin Lockwood, and Dmitry Taubinsky. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Ramsey Strikes Back: Optimal Commodity Taxes and Redistribution in the Presence of Salience Effects

Hunt Allcott, Benjamin Lockwood, and Dmitry Taubinsky

NBER Working Paper No. 24233

January 2018

JEL No. H21,H23

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 essentially 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.

Hunt Allcott
Department of Economics
New York University
19 W. 4th Street, 6th Floor
New York, NY 10012
and NBER
hunt.allcott@nyu.edu

Dmitry Taubinsky
University of California, Berkeley
Department of Economics
530 Evans Hall #3880
Berkeley, CA 94720-3880
and NBER
dmitry.taubinsky@berkeley.edu

Benjamin Lockwood
The Wharton School
University of Pennsylvania
1400 Steinberg-Dietrich Hall
3620 Locust Walk
Philadelphia, PA 19104
and NBER
ben.lockwood@wharton.upenn.edu

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 strong. Chetty, Looney and Kroft (2009) and Taubinsky and Rees-Jones (forthcoming) find 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 behavioral responses than direct income tax reforms.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

¹This 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.

²See Strumpf (2011) for evidence that the introduction of Medicaid 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 the effect energy efficiency subsidies.

in the context of non-salient commodity taxes. Specifically, we show that the optimal commodity tax follows the Diamond-Ramsey formula, but scaled down by the degree of inattention.

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 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 $w \in R$, 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 $T(z)$.³ Consumers use their net income to choose a consumption bundle (c_1, c_2) , which is sold at before-tax prices (p_1, p_2) and are subject to additional linear (ad valorem) commodity taxes $t = (t_1, t_2)$. 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

$$\int (p_1 t_1 c_1(w) + p_2 t_2 c_2(w) + T(z(w))) dF(w) \geq R,$$

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

$$(c_1(w), c_2(w), l(w)) = \arg \max_{\{c_1, c_2, l\}} \{U(c_1, c_2, l; w)\},$$

subject to the budget constraint

$$p_1(1+t_1)c_1(w) + p_2(1+t_2)c_2(w) \leq wl(w) - T(wl(w)).$$

Implicit in this optimization is a strong assumption: when choosing labor supply l , individuals fully account for the effect of commodity taxes on the returns to labor. We relax that assumption to 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 is motivated by the sparsity-based bounded rationality model of Gabaix (2014). The framework of Gabaix (2014) would predict that people pay relatively less attention to changes in taxes or prices on commodities which constitute a small budget share.

Consistent with that reasoning, we suppose consumers correctly perceive the tax-inclusive price of c_1 —perhaps because c_1 is a composite good that constitutes a large share of expenditures—but misperceive the tax-inclusive price of c_2 when making labor supply decisions. Specifically, consumers mistakenly believe that the ratio of the (total) price of c_2 to the price of c_1 will be $(1 - \theta)\hat{r} + \theta \frac{p_2(1+t_2)}{p_1(1+t_1)}$, where \hat{r} is a “mental default” for the price ratio. The attention parameter θ captures the extent to which consumers’ labor supply is sensitive to

³Throughout the paper, labor supply decisions should be interpreted broadly to include decisions about human capital acquisition, career choice, and extensive margin decisions such as retirement.

variations in prices generated by the commodity tax on c_2 .

To focus our attention on the implications of commodity tax salience for *labor supply* decisions, we assume consumers correctly perceive commodity taxes at the time the bundle (c_1, c_2) is chosen. This contrasts with the form of non-salience studied in Chetty, Looney and Kroft (2009) and Taubinsky and Rees-Jones (forthcoming), wherein consumers are inattentive to commodity taxes at the time of purchase. Our framework thus accommodates settings where, e.g., taxes are not included in posted prices, but consumers research them prior to commodity purchases.⁴ More generally, this setup can represent any situation where t_2 is fully salient at the time of purchase because it is included in the posted price, but some other mental default anchors attention at the time labor supply decisions are made.

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 representation satisfies the two conditions—weak separability of consumption (taken together) and labor in the utility function, and homogeneous subutility of consumption—which give rise to the Atkinson-Stiglitz theorem. In particular, 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 effects, rather than consumption preferences varying with earnings ability w . The Atkinson-Stiglitz theorem holds 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).⁵ This assumption sharpens the consequences of salience effects, as it implies that any deviation from uniform commodity taxation is due to salience effects.

We assume the income tax does not generate bunching (in which multiple ability levels earn the same income), so that there is a one-to-one mapping between ability and income. As a result, we can write choice variables as a function of (endogenous) income z at the optimum, rather than unobservable w .

Following Diamond (1975), we let $\alpha(z)$ denote the social marginal utility from 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 equal to the multiplier on the government budget constraint at the optimum. We let \bar{c}_2 denote the average consumption of c_2 , and we let $\xi = -\frac{d\bar{c}_2}{dp_2} \cdot \frac{p_2}{\bar{c}_2}$ denote the (aggregate) price elasticity of demand for c_2 , and let ξ^c denote the compensated elasticity. Finally, we let $\hat{c}_2(z)$ denote the level of c_2 consumption that a z -earner anticipates when setting labor supply. We let $\hat{\theta}(z) = \theta \frac{\hat{c}_2(z)}{c_2(z)}$ denote the income-effect salience of changes in t_2 ; that is, $\hat{\theta}$ is the ratio of perceived to actual effects on a consumers’ wealth. (The term $\hat{c}_2(z)$ can be written in terms of actual c_2 consumption and the inattention parameters θ and \hat{p}_2 —see the Appendix for that derivation.)

We now generalize the Atkinson-Stiglitz theorem to characterize the optimal commodity tax structure as a function of salience θ .

Proposition 1. *Any optimal tax system must satisfy*

$$\frac{t_2}{1+t_2} - \frac{t_1}{1+t_1} = \frac{1}{\lambda \xi^c} \frac{E[(1-\hat{\theta}(z))(\lambda - \alpha(z))c_2(z)]}{\bar{c}_2},$$

⁴In this case, a natural mental default is the mistaken believe that the tax t_1 applies to commodity c_2 . In this case, $\hat{r} = \frac{p_2(1+t_1)}{p_1(1+t_1)}$, and the tax on c_2 is misperceived to be $(1-\theta)t_1 + \theta t_2$.

⁵Lockwood and Taubinsky (2017) further show that for 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 is identical to full nonsalience ($\theta = 0$) in this model.

and, in particular, the social optimum can be implemented with t_1 and t_2 satisfying

$$t_1 = 0$$

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

When agents are fully rational ($\theta = 1$), Proposition 1 reproduces 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 no commodity taxes.

However, the uniform commodity taxation result breaks down when consumers are not fully attentive to the commodity tax t_2 . In this case, the social optimum is implemented with a tax t_2 that follows the “many-person Ramsey tax rule” of Diamond (1975), scaled by the degree of inattention ($1 - \hat{\theta}$). When $\theta = 0$, the Diamond (1975) Ramsey tax rule is exactly replicated.

The optimal t_2 can be written as a simpler approximation if two additional conditions hold: (1) income effects on labor supply are negligible (in which case $\lambda \approx E[\alpha(z)]$), and (2) $\hat{\theta}(z)$ is constant across the population. In this case, we can use the fact that $E[(1 - \hat{\theta}(z))(\lambda - \alpha(z))c_2(z)] \approx -(1 - \hat{\theta})Cov[\alpha(z), c_2(z)]$ to rewrite the expression for t_2 above as:

$$\frac{t_2}{1 + t_2} = \frac{1 - \hat{\theta}}{\xi^c} \left(\frac{-Cov[\frac{\alpha(z)}{\lambda}, c_2(z)]}{\bar{c}_2} \right).$$

Recalling that $\alpha(z)$ corresponds to the social marginal value of net income for a z -earner, in the conventional case $\alpha(\cdot)$ declines with z . This means that in the presence of income taxation and salience effects, the optimal commodity tax on c_2 satisfies a simple principle: it should be positive if c_2 is consumed more heavily by the rich (i.e., when t_2 is progressive) and it should be negative—a subsidy—if c_2 is consumed more heavily by the poor.

We now consider a useful extension to Proposition 1, which deals with the case of corrective commodity taxes. That is, we consider the case in which the government taxes c_2 because it generates an externality or “internality” (e.g., due to behavioral factors like present bias). We assume that while $\frac{\partial u}{\partial c_2}$ is consumers’ (perceived) marginal utility from consuming c_2 , the social marginal utility from consumption of c_2 is actually $\frac{\partial u}{\partial c_2} - \chi$, for some $\chi \in R$. For simplicity, we consider the special case where χ is homogeneous and constant.⁶

The optimal Pigovian tax rate on c_2 is simply $t_2 = \frac{\chi}{\lambda p_2}$. In the presence of income inequality, however, a common intuition is that the tax on c_2 should be adjusted away from this Pigovian benchmark depending on whether the tax is progressive or regressive. Yet a natural 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 $t_2 = \frac{\chi}{\lambda p_2}$, because the regressive burden of the commodity tax can be perfectly offset by making the income tax more progressive in a way that exactly preserves 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 to the income tax, a change in the income tax that exactly offsets the burden of t_2 will nevertheless generate changes in labor supply. We formalize our result in the proposition below.

Proposition 2. *The social optimum can be implemented with $t_1 = 0$ and t_2 satisfying*

⁶Lockwood and Taubinsky (2017) analyze the case in which χ may be heterogeneous across consumers and nonlinear in c_2 .

$$t_2 = \frac{\chi}{\lambda p_2} - \frac{1 + t_2}{\lambda \xi^c} \frac{E[(1 - \hat{\theta}(z))(\lambda - \alpha(z))\hat{c}_2(z)]}{\bar{c}_2}$$

The term $\frac{\chi}{\lambda p_2}$ is the “Pigovian correction”—it is simply the value that the tax would take on in the absence of any redistributive concerns. The second term in the formula represents regressivity costs, and comes from redistributive concerns. These concerns are immaterial when consumers are fully rational ($\theta = 1$). In the presence of salience effects, however, if the tax is regressive and $\theta < 1$, then the optimal t_2 lies below the Pigovian benchmark. When the tax is progressive and $\theta < 1$, the optimal t_2 lies above the Pigovian benchmark.

4 Conclusion

Together, the results in this paper serve as an exploration of the robustness of optimal tax results derived in the prevailing Mirrlees framework, which assumes 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 displaced earlier results about optimal tax structures in the Ramsey tradition.

This paper underscores, however, that the implicit assumption of perfect rationality is a strong one, perhaps especially so when tax instruments affect labor-supply incentives in opaque or nuanced ways. 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 that their optimal size follows two intuitive principles: taxes 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. These results provide quantitative and qualitative guidance for a range of fiscal policies including capital income taxation, consumption taxes, and in-kind transfers.

References

- Allcott, Hunt, Benjamin B. Lockwood, and Dmitry Taubinsky.** 2018. “Regressive Sin Taxes, with Applications to the Optimal Soda Tax.” *Working Paper*.
- Atkinson, Anthony B, and J.E. Stiglitz.** 1976. “Design of Tax Structure - Direct Versus Indirect Taxation.” *Journal of Public Economics*, 6: 55–75.
- Chetty, Raj, Adam Looney, and Kory Kroft.** 2009. “Salience and Taxation: Theory and Evidence.” *American Economic Review*, 99(4): 1145–1177.
- Chetty, Raj, John N Friedman, Soren Leth-Petersen, Torben Nielsen, and Tore Olsen.** 2014. “Active vs. Passive Decisions and Crowd-Out in Retirement Savings Accounts: Evidence from Denmark.” *Quarterly Journal of Economics*, 139(3): 1141–1219.

- Diamond, Peter.** 1975. “A Many-Person Ramsey Tax Rule.” *Journal of Public Economics*, 4: 335–342.
- Gabaix, Xavier.** 2014. “A Sparsity Based Model of Bounded Rationality.” *Quarterly Journal of Economics*, 129: 1661–1710.
- Gallagher, Kelly Sims, and Erich Muehlegger.** 2011. “Giving green to get green? Incentives and consumer adoption of hybrid vehicle technology.” *Journal of Environmental Economics and Management*, 61(1): 1–15.
- Jacobs, Bas, and Robin Boadway.** 2014. “Optimal linear commodity taxation under optimal non-linear income taxation.” *Journal of Public Economics*, 117: 201–210.
- Lockwood, Benjamin B., and Dmitry Taubinsky.** 2017. “Regressive Sin Taxes.” *Working Paper*.
- Monks, James.** 2004. “An empirical examination of the impact of college financial aid on family savings.” *National Tax Journal*, 189–207.
- Ramsey, Frank P.** 1927. “A Contribution to the Theory of Taxation.” *The Economic Journal*, 37(145): 47–61.
- Rees-Jones, Alex, and Dmitry Taubinsky.** 2018. “Taxing Humans: Pitfalls of the Mechanism Design Approach and Potential Resolutions.” *Tax Policy and the Economy*, 1.
- Saez, Emmanuel.** 2002. “The desirability of commodity taxation under non-linear income taxation and heterogeneous tastes.” *Journal of Public Economics*, 83(2): 217–230.
- Strumpf, Erin.** 2011. “Medicaid’s effect on single women’s labor supply: Evidence from the introduction of Medicaid.” *Journal of Health Economics*, 30(3): 531–548.
- Taubinsky, Dmitry, and Alex Rees-Jones.** forthcoming. “Attention Variation and Welfare: Theory and Evidence from a Tax Salience Experiment.” *Review of Economic Studies*.

A Mathematical Appendix

Formal definition of $\alpha(z)$

An increase dy in a consumer’s (after-tax) income has the following three effects:

1. By the envelope theorem, the utility impact of an increase in net income equal to dy is $dy \cdot \frac{u_1(c_1, c_2)}{p_1(1+t_1)}$. (We denote partial derivatives using the notation $u_1 := \frac{\partial u}{\partial c_1}$, etc., throughout.)
2. This generates fiscal externalities equal to $dy \frac{dc_1}{dy} p_1 t_1 + dy \frac{dc_2}{dy} p_2 t_2$.
3. By changing consumption of c_2 , this also alters the quantity of externalities (or internalities) produced, by $-dy \frac{dc_2}{dy} \chi$.

The net effect is thus

$$\alpha(z) = \frac{u_1}{p_1(1+t_1)} + \frac{dc_1}{dy} p_1 t_1 + \frac{dc_2}{dy} (p_2 t_2 - \chi).$$

Proof of Propositions 1 and 2

Let $(\hat{c}_1(z), \hat{c}_2(z))$ denote the consumption bundle that a z -earner anticipates consuming when setting labor supply, while $(c_1(z), c_2(z))$ denotes the bundle they will actually choose.

For this proof, we consider the following joint perturbation of commodity and income taxes, which preserves labor supply choices: the commodity tax t_2 is raised by dt while the income tax is reduced by $dt \cdot \theta \hat{c}_2(z) p_2$ at each income z . At the time of labor supply choice, consumers perceive the commodity tax increase to be θdt , a reform which, on its own, reduces their anticipated utility from each possible choice of z by $dt \cdot \theta p_2 \hat{c}_2(z) u_1(\hat{c}_1(z), \hat{c}_2(z))$. By construction, the income tax reduction also *raises* anticipated utility by $dt \cdot \theta p_2 \hat{c}_2(z) u_1(\hat{c}_1(z), \hat{c}_2(z))$. Together, these reforms have offsetting effects on consumers' anticipated utility from each possible choice of z , so the joint reform does not alter earnings decisions. This design therefore simplifies the characterization of the total welfare effect of this reform, which can be decomposed into the following effects for each z -earner.

1. The reform mechanically raises revenue

$$dt \cdot (c_2(z) - \theta \hat{c}_2(z)) p_2 = dt \cdot (1 - \hat{\theta}) c_2(z)$$

from each consumer, at a marginal value of public funds equal to λ .

2. The social value of the resulting mechanical change in income for each consumer (including the resulting fiscal externalities through changes in c_1 and c_2 due to income effects) is

$$-dt \cdot \alpha(z) (c_2(z) - \theta \hat{c}_2(z)) p_2 = -dt \alpha(z) (1 - \hat{\theta}(z)) c_2(z) p_2$$

3. The commodity tax change also generates a substitution effect from c_2 to c_1 . Since the resulting change in income has mostly been compensated through the income tax reform (and the remaining true income change has been handled through the income effects in (2)) we can write this effect in terms of the compensated elasticity of demand for c_2 . Specifically, the change in c_2 consumption due to a compensated tax change of dt is $d\bar{c}_2 = -dt \cdot \xi^c \frac{\bar{c}_2}{1+t_2}$. Correspondingly, the change in c_1 consumption from this readjustment is $d\bar{c}_1 = -d\bar{c}_2 \frac{p_2(1+t_2)}{p_1(1+t_1)}$. Therefore the total impact of fiscal externalities from this adjustment is equal to

$$d\bar{c}_1 p_1 t_1 + d\bar{c}_2 p_2 t_2 = dt \cdot \xi^c \bar{c}_2 p_2 \left(\frac{t_1}{1+t_1} - \frac{t_2}{1+t_2} \right),$$

weighted by the marginal value of public funds λ .

4. Finally, this substitution from c_2 to c_1 also alters externalities. This generates a welfare change of $-d\bar{c}_2 \chi = dt \cdot \xi^c \frac{\bar{c}_2}{1+t_2} \chi$.

Under the optimal policy, the sum of first-order effects (1)–(4) must equal zero:

$$\lambda \xi^c \bar{c}_2 p_2 \left(\frac{t_1}{1+t_1} - \frac{t_2}{1+t_2} \right) + \frac{\bar{c}_2 \xi^c}{(1+t_2)} \chi + (1-\theta) p_2 E[(1-\hat{\theta}(z))(\lambda - \alpha(z)) c_2(z)] = 0$$

and thus

$$t_2 - t_1 \left(\frac{1 + t_2}{1 + t_1} \right) = \frac{\chi}{\lambda p_2} + \frac{1 + t_2}{\lambda \xi^c} \frac{E[(1 - \hat{\theta}(z))(\lambda - \alpha(z))c_2(z)]}{\bar{c}_2}$$

at the optimal policy.

To complete the proof, we show that if an optimum can be implemented with some pair of commodity tax rates (t_1^*, t_2^*) , then it can also be implemented with the pair of tax rates (t_1^{**}, t_2^{**}) satisfying $t_1^{**} = 0$. To that end, let T^* be the optimal income tax given the tax rates (t_1^*, t_2^*) . Now consider $t_1^{**} = 0$, $t_2^{**} = \frac{t_2^* - t_1^*}{1 + t_2^*}$, and $T^{**}(z) = \frac{t_1^* z}{1 + t_1^*} + \frac{T^*(z)}{1 + t_1^*}$. Note that $p_2(1 + t_2^{**})(1 + t_1^*) = p_2(1 + t_2^*)$ and $(z - T^{**}(z))(1 + t_1^*) = z - T^*(z)$. Thus, compared to the (t_1^*, t_2^*, T^*) regime, the $(t_1^{**}, t_2^{**}, T^{**})$ regime simply multiplies both the after-tax prices and consumers' after-tax incomes by $\frac{1}{1 + t_1^*}$. Because both of the after-tax prices are multiplied by the constant $A = \frac{1}{1 + t_1^*}$, the perceived after-tax price of c_2 is multiplied by this constant as well, as the perceived price will be given by

$$\theta A p_2(1 + t_2) + (1 - \theta) A p_1(1 + t_1) \hat{r} = A [\theta p_2(1 + t_2) + (1 - \theta) p_1(1 + t_1) \hat{r}]$$

Thus, the set of consumption bundles available to a z -earner is identical under the two different tax regimes, and thus the equilibrium allocation will be identical under the two tax regimes.

Writing anticipated consumption in terms of actual consumption

The optimal tax depends on (mistaken) anticipated consumption $\hat{c}_2(z)$, as is evident in the preceding proof. However, since the difference between anticipated consumption $\hat{c}_2(z)$ and actual consumption $c_2(z)$ is effectively driven by an unanticipated price change from $(1 - \theta)\hat{p}_2 + \theta p_2(1 + t_2)$ to $p_2(1 + t_2)$ —a difference of $(1 - \theta)(p_2(1 + t_2) - \hat{p}_2)$. Thus the resulting change in c_2 consumption (written in terms of the demand elasticity) as a share of actual c_2 consumption, satisfies the following expression:

$$\frac{c_2(z) - \hat{c}_2(z)}{c_2(z)} \approx (1 - \theta) \left(1 - \frac{\hat{p}_2}{p_2(1 + t_2)} \right) \xi(z),$$

where $\xi(z)$ is the income-conditional (uncompensated) demand elasticity for c_2 . (This is an approximation because this effective price change need not be infinitesimal change.) Thus

$$1 - \hat{\theta}(z) \approx (1 - \theta) + \theta(1 - \theta) \left(1 - \frac{\hat{p}_2}{p_2(1 + t_2)} \right) \xi(z).$$