Optimal Tax Preferences*

Ryan Nunn†

November 3, 2014

Abstract

A key goal of the tax system is to raise the desired quantity of revenue in the least economically costly manner. However, policymakers sometimes wish to use the tax system to encourage certain types of economic behavior via tax preferences. This paper demonstrates that tax preferences structured as deductions from taxable income are not ideal tax policy. Those tax preferences could be restructured to achieve all existing objectives while lowering the excess burden of income taxation. Rather than offering a deduction for the expense of a tax-favored activity, the optimal tax system would instead feature a labor tax rate that is a decreasing function of that activity. Such a system encourages taxpayers to use the preference, increasing their consumption of the tax-favored good, but at the same time offers a general reduction in their tax rate. This reduction would lower the marginal distortions associated with consumption of all other taxed goods.

In fact, this proposal allows for an arbitrary level of tax-preferred consumption without the loss of revenue implied by a standard tax expenditure. Worker productivity heterogeneity weakens this result, but only slightly: when the model (using a simple, linear tax function) is calibrated to the case of itemized charitable giving, the "tax revenue advantage" of the optimal tax preference, relative to the normal deduction, falls by about 20%.

JEL codes: H21, H24

*The views in this paper are solely the responsibility of the author and should not be interpreted as reflecting the views of the US Treasury Department or of any other person associated with the US Treasury Department. Thanks to Aditya Aladangady, Adam Cole, Gabriel Ehrlich, Laura Kawano, Shanthi Ramnath, Collin Raymond, and seminar participants at the US Treasury Department for their comments. All errors are the author's own.

†Office of Economic Policy, US Department of the Treasury. Email: ryan.nunn@treasury.gov.

1 Introduction

Any labor income tax will raise revenue and generate deadweight loss; in other words, even if a labor tax is fully refunded to workers, their welfare will be diminished by the tax. In devising optimal instruments and schedules of income taxation, it is necessary to find the least disruptive and costly tax system that nonetheless raises a given quantity of revenue. As more revenue is raised, deadweight loss increases: moreover, it is generally the case that the marginal excess burden of taxation rises with the tax rate. Given this reality, tax preferences are problematic, as they entail a give-away of revenue without a corresponding reduction in the general rate of taxation. These preferences may be justified in light of the positive spillovers associated with a given activity (e.g., charitable donations). But it remains the case that these deductions necessitate the "narrowing" of the taxable income base and an increase in the rate of taxation, all else equal, with an implication of higher marginal deadweight loss.

However, tax preferences could be restructured to avoid this problem. Instead of allowing taxpayers to deduct the expense of their tax-preferred consumption, a worker's labor tax rate, though constant with respect to labor income, could be made a decreasing function of tax-preferred consumption. As with the existing system, this proposal would provide a tax incentive to consume the favored good. Unlike the existing system, an optimal tax rate function would lower the marginal tax rate on all labor income, raising more tax revenue for a given level of deadweight loss.

The intuition for this result is straightforward. A normal tax deduction lowers the relative after-tax price of the favored consumption. In the case of the charitable donation deduction, for instance, a worker is encouraged to view charitable giving as a less expensive good than it would be without the deduction. The deduction does not change the after-tax price of "ordinary" (non-charitable) consumption relative to leisure, however. If, instead, the tax preference were structured as a reduction (proportional to charitable spending) in the labor income tax rate, the after-tax price of ordinary consumption would fall relative to leisure. Because taxes on labor income also distort that relative price, making leisure excessively attractive, this price shift yields an improvement in social welfare.[1]

First, it is shown in a general setting that an optimal tax rate function obviates the need to give away tax revenue in order to encourage consumption of a particular good; i.e., any combination of labor supply and tax-preferred consumption that is obtainable with a normal tax expenditure is also obtainable with a tax rate function, but without the need to narrow the tax base and lose revenue. Equivalently, when each tax preference regime yields the same level of tax revenue, the traditional system requires a higher tax rate on labor income and increased deadweight loss. This logic generalizes to any situation in which fiscal activities can be "outsourced" to taxpayers (by contrast to goods, like national defense, that cannot be purchased by taxpayers). Direct purchase

[1] It is important to note that the disadvantage of the normal tax expenditure, relative to the proposed alternative, does *not* depend on the fact that the typical expenditure is restricted to be a subsidy at the worker's labor tax rate. For any desired implicit subsidy of the tax-favored good, a tax rate function of the type described in this paper will be a preferable alternative.

of a good, paid for by an income tax, is similarly dominated by the proposed alternative of a tax rate function.

Finally, a specific implementation of the model, incorporating functional form assumptions and worker heterogeneity, is calibrated to the case of itemized charitable giving. This allows a numerical calculation of the optimal tax preference's advantage relative to a traditional deduction when a single tax rate function must be applied to workers of varying types. With productivity heterogeneity calibrated to the level observed in the U.S., the optimal tax preference still increases tax revenue by about 80% of the original tax expenditure while maintaining social welfare and charitable donations at the same level induced by the original tax preference. Intuitively, high productivity workers are led to excessively consume the tax-preferred good, while low-productivity workers consume it insufficiently, from a social planner's perspective. Worker preference heterogeneity, by contrast, does not affect the results.

2 Related Literature

While a voluminous theoretical and empirical literature has examined the optimality of the income tax schedule, relatively few efforts have been made to characterize the optimality of tax expenditures.[2] Saez (2004) embeds tax expenditures in a framework with positive spillovers, direct purchase by government (and consequent partial crowding-out of private activity in a model that accommodates "warm glow" preferences), and earnings heterogeneity. He derives expressions for the optimal linear income tax rate and tax expenditures under a variety of different assumptions, then conducts numerical simulations that depict the range of optimal tax expenditure subsidies and income tax rates. For most of the paper, Saez allows the tax preference to be a subsidy of arbitrary size; in other words, the expenditure is not constrained to be proportional to the worker's tax rate, as is generally the case in U.S. tax law. However, he does not consider the possibility of a tax preference similar to that discussed in this paper, in which the labor tax rate itself is a function of tax-preferred consumption.

Importantly, Saez characterizes opponents of tax expenditures as claiming that "as tax expenditures are likely to be much more responsive to taxation than labor supply, they point out that allowing tax expenditures may both reduce the size of the tax base and increase significantly the elasticity of taxable income, thus increasing significantly the total deadweight burden from the income tax." (Saez, 2004, p. 2658) This is precisely the concern that motivates this paper. Indeed, the tradeoff between encouragement of a socially-beneficial activity and erosion of the tax base can be eliminated if typical tax expenditures (i.e., deductions from taxable income) are replaced with tax rate functions.

A number of empirical papers have investigated the consequences of specific tax preferences (e.g., the mortgage interest and charitable deductions).[3] This work examines the aforementioned assumption that the elasticity of tax-preferred consumption is higher than that of taxable income.

[2]For important examples of the former, see Mirrlees (1971) and Diamond (1998).
[3]See Rosen (1985) and Auten et al. (2002).

That assumption will likely depend on the particular tax preference considered. In the case of charitable giving, Auten finds that its elasticity with respect to the after-tax price of giving ranges from -0.8 to -1.3. Dunsky and Follain (2000) find an elasticity of about -1 for mortgage debt. These can be compared with the overall taxable income elasticity of -0.4 reported in Gruber and Saez (2002). As Gruber and Saez point out, this figure is itself inflated in magnitude by the existence of tax preferences; the tax response of actual labor income is likely lower.

3 Model

3.1 Standard Tax Expenditures

What follows is a simple setting in which it will be demonstrated that the existing regime of tax expenditures is strictly dominated by one in which the income tax rate declines linearly with consumption of the tax-favored good.

Consider a single worker who maximizes utility over labor supplied L, ordinary consumption c, and a separate good H.

$$U = c + \phi(L) + \eta(H),$$

where $\phi' < 0$, $\phi'' < 0$, $\eta' > 0$, and $\eta'' < 0$. The worker is subject to a budget constraint

$$p_H H + c \leq (1 - \tau)wL + R,$$

where τ is the constant labor tax rate, p_H is the relative price of H, and w is the constant pre-tax wage. The price of ordinary consumption is normalized to one. Tax revenue $R = \tau wL$ is refunded lump-sum to the worker.

Now introduce a tax preference for H. This takes the form of a deduction of expenditure on H from taxable labor income. The worker's lifetime budget constraint is now:

$$p_H \left(1 - \tau\right) H + c \leq (1 - \tau)wL + R,$$

where $R = \tau(wL - p_H H)$.

Note that while this deduction reduces the after-tax price of tax-preferred consumption by a factor $(1-\tau)$, the relative price of ordinary consumption c and leisure $(1-L)$ remains $(1-\tau)w$ with or without the tax preference. Intuitively, this means that the tax preference will not ameliorate any of the deadweight loss associated with a tax-distorted consumption-leisure choice. In the present quasilinear setting, with no income effects, it is further the case that $(1 - L^*)$ will be unaffected by the deduction permitted for H.

Below, set up the worker's problem and take first-order conditions.

$$\begin{aligned}
\mathcal{L} &: c + \phi(L) + \eta(H) + \lambda\left([1-\tau]wL + R - c - p_H H(1-\tau)\right) \\
\mathcal{L}_c &: 1 - \lambda = 0 \\
\mathcal{L}_L &: \phi'(L) + \lambda(1-\tau)w = 0 \\
\mathcal{L}_H &: \eta'(H) - \lambda p_H(1-\tau) = 0
\end{aligned}$$

Collectively, this implies that the marginal disutility of labor $\phi'(L^*) = -(1-\tau)w$ and the marginal utility of preferred consumption $\eta'(H^*) = p_H(1-\tau)$. Tax revenue R, net of the deduction, is $\tau wL - \tau p_H H$.

3.2 Tax Rate Functions

Rather than allow a deduction equal to $\tau p_H H$, this proposal would implement a wage tax rate that is constant with respect to labor income, but decreasing in tax-preferred consumption. Let the tax rate be $\tau(H)$, where $\tau'(H) < 0$ and $\tau''(H) \geq 0$. As before, the worker problem and its associated first-order conditions are first presented.

$$\begin{aligned}
\mathcal{L} &: c + \phi(L) + \eta(H) + \lambda\left([1-\tau(H)]wL + R - c - p_H H\right) \\
\mathcal{L}_c &: 1 - \lambda = 0 \\
\mathcal{L}_L &: \phi'(L) + \lambda(1-\tau(H))w = 0 \\
\mathcal{L}_H &: \eta'(H) - \lambda\tau'(H)wL - \lambda p_H = 0
\end{aligned}$$

Tax revenue is now $R = \tau(H)wL$. It will be convenient to distinguish the standard regime by a subscript S and the proposed regime by F.

Proposition 1 *For any $\tau_S > 0$ and $H_S^* > 0$, there exists a tax rate function $\tau_F(H)$ such that $H_F^* = H_S^*$, $c_F^* = c_S^*$, $L_F^* = L_S^*$, and $R_F^* = R_S^* + \tau_S p_H H_S^*$. This implies that worker choices (and by implication, social welfare) are identical across the two regimes, but tax revenue is strictly higher in the proposed regime ($R_F > R_S$).*[4]

Proof. From the first-order conditions on ordinary consumption c, $\lambda_S = \lambda_F = 1$. In equilibrium, if labor supply is chosen to be equal across regimes ($L_F^* = L_S^*$), then from the first-order conditions on L it is apparent that $\tau_F(H^*) = \tau_S$, so that the level of labor income tax rates are equal across regimes.

[4] Throughout the paper, the analysis is restricted to tax rates that are both constant with respect to income and linear in consumption of the tax-favored good.

Now consider the first-order conditions pertaining to H for each regime:

$$\eta'(H_S) - p_H(1-\tau_S) = 0$$
$$\eta'(H_F) - \tau_F'(H)wL_F - p_H = 0$$

If tax-preferred consumption is chosen to be equal across regimes ($H_F^* = H_S^*$), then $p_H(1-\tau_S) = \tau_F'(H)wL_F + p_H$. This condition defines the first derivative of the proposed tax function $\tau_F(\cdot)$. To see that both $H_F^* = H_S^*$ and $\tau_F = \tau_S$ may hold simultaneously, let the previous condition hold at the equilibrium value of L, i.e., $p_H(1-\tau_S) = \tau_F'(H)wL_F^* + p_H$. Quasilinear utility makes this step trivial.

As mentioned previously, tax revenue is fully refunded to the worker. Therefore, conditional on choices of L and H, any change in tax revenue collected will not affect c. This implies that $c_F^* = c_S^*$. The worker's welfare (identical to social welfare in this setting) is therefore identical across regimes. Finally, the conditions on derivatives of the tax rate, $\tau_F' < 0$ and $\tau_F'' \geq 0$, ensure that the budget set is convex and that the equilibrium is unique.[5]

By stipulation, tax revenue in the standard regime is $R_S = \tau_S(wL_S^* - p_H H_S^*)$, and tax revenue in the proposed regime is $R_F = \tau_F w L_F^*$. Given the identies established above, $R_F^* = R_S^* + \tau_S p_H H_S^*$. Therefore, the proposed tax system strictly dominates the existing system, as it manages to raise increased revenue for a given level of deadweight loss and tax-preferred consumption. Conversely, one could hold tax revenue constant across regimes, and the proposed regime would then imply higher social welfare. ■

Corollary 2 *Given a government tax revenue target $\mathcal{R}$, there exists a tax function $\tau_F(H)$ such that $H_F^* = H_S^*$, $c_{\tau_F=0}^* \geq c_F^* > c_S^*$, $L_{\tau_F=0}^* \geq L_F^* > L_S^*$, and $R_F^* = R_S^* = \mathcal{R}$.[6] Social welfare is therefore strictly higher in the proposed regime.*

Proof. Intuitively, Proposition 1 established that the two regimes, despite imposing identical labor income tax rates in equilibrium, delivered different levels of tax revenue. An appropriately-chosen reduction in the tax rate under the proposed regime, then, can deliver identical tax revenue and strictly higher social welfare.

As before, begin with the first-order conditions pertaining to H for each regime:

$$\eta'(H_S) - p_H(1-\tau_S) = 0$$
$$\eta'(H_F) - \tau_F'(H)wL_F - p_H = 0.$$

Also as previously, require that $H^* = H_F^* = H_S^*$, indicating that the two regimes induce the same level of tax-preferred consumption, and that $\eta'(H_S) = \eta'(H_F)$. Further impose the restriction that tax revenue be equal across regimes, with $\mathcal{R} = \tau_S(wL_S^* - p_H H^*) = \tau_F(H^*)wL_F^*$.

[5] Given that the tax rate function is only defined in its level and first derivative, there are infinitely many $\tau(\cdot)$ functions consistent with this equilibrium.

[6] Restrict consideration to equilibria for which tax revenue is increasing in the tax rate, i.e., $\frac{\partial(\tau w L)}{\partial \tau} > 0$.

[I plan to revise the rest of this section to avoid using a particular tax rate function.]

It now suffices to show that the corollary holds for a particular tax rate function: $\tau_F = t(1 + ap_H H)$.

Then find the a, t that solve the equations

$$p_H(1-\tau_S) = wp_H L_F ta + p_H$$
$$\tau_S(wL_S^* - p_H H^*) = t(1 + ap_H H^*)wL_F^*.$$

A solution is given by

$$t = \frac{\tau_S L_S^*}{L_F^*}$$
$$a = \frac{-1}{wL_S^*},$$

with $t > 0$ if τ_S, L_S^*, and L_F^* are all strictly positive, and $a < 0$ if $L_S^* > 0$. It remains to be shown that $L_{\tau=0}^* \geq L_F^* > L_S^*$, which would imply both $c_{\tau=0}^* \geq c_F^* > c_S^*$ and that social welfare is strictly higher in the proposed regime than the standard regime. First, from differentiation of the first-order condition $\phi'(L) = -(1-\tau)w$ and the assumption that $\phi'' < 0$, it is apparent that L^* is monotonically decreasing in the tax rate τ. This implies that $L_{\tau_F=0}^* \geq L_F^*$: a positive tax rate distorts the labor-leisure choice, reducing labor supply relative to the socially-optimal level.

Now consider the revenue requirement $\mathcal{R} = \tau_S(wL_S^* - p_H H^*) = \tau_F w L_F^*$. Recall that, previously, consideration was restricted to equilibria for which tax revenue is increasing in the tax rate, i.e., $\frac{\partial(\tau w L)}{\partial \tau} > 0$. Suppose that $\tau_S = \tau_F$, which would hold iff $L_S^* = L_F^*$. Then $R_S + \tau_S p_H H^* = R_F$ and $R_S < R_F$, consistent with Proposition 1. Raising R_S to meet the revenue target $\mathcal{R}$, and bring it into equality with R_F, necessitates an increase in τ_S, which in turn reduces labor supply such that $L_S^* < L_F^*$. Finally, we have the identity $L_{\tau_F=0}^* \geq L_F^* > L_S^*$. Because labor supply L_F and ordinary consumption c_F are strictly closer to their socially-optimal levels, with all other variables identical across regimes, it follows that social welfare is strictly higher in the proposed regime than in the standard regime. ■

4 Direct Purchase by a Fiscal Authority

An alternative to either a tax deduction or a decreasing tax rate function is direct purchase of a good by the government. In the environment developed previously, direct purchase is relatively simple. If the government would like to increase consumption of H beyond the level that would obtain in a private equilibrium, it must purchase the entire quantity of H. Let H_g be the amount purchased by the government and H_p be the amount purchased by the worker, with $H = H_g + H_p$. The worker takes H_g as given when choosing H_P.

$$U = c + \phi(L) + \eta(H_g + H_p)$$

The first-order condition determining the choice of H_g is $\eta'(H_g + H_p) - p_H = 0$. If H^* is the total quantity that would be chosen in an equilibrium with $H_g = 0$, inspection of the first-order condition makes it clear that $H_p^* + H_g = H^* \;\; \forall H_g \leq H^*$. In other words, private purchase of H is perfectly crowded out by government purchase.[7]

Proposition 3 *For a given level of tax-preferred consumption $\overline{H}$, there exist an optimal tax rate function that implements $\overline{H}$ with strictly less deadweight loss than is generated by direct government purchase.*

Proof. Assume, for simplicity, that the fiscal authority has no need of revenue for purposes other than subsidization of H. Then it follows that direct purchase would require a tax rate $\tau = \frac{p_H \overline{H}}{wL(\tau)} > 0$, where $\overline{H} > H^*$ is the desired quantity of the good and $L(\tau)$ is the quantity of labor supplied at a given tax rate. Given that $\tau > 0$, deadweight loss would be strictly positive.

Now consider the alternative of a declining tax rate function. It is a straightforward implication of proposition 1 that $\tau_F(H) = 0$ exists such that the worker chooses $\overline{H}$, and the marginal tax rate on labor is zero at the desired quantity of H. This alternative dominates the direct purchase option, which generates a strictly positive deadweight loss. Intuitively, direct purchase requires that the government raise revenue with a distortionary tax, then procure the optimal quantity of the good. A decreasing tax rate function, however, can encourage H consumption without actually raising revenue or distorting the relative price of leisure and ordinary consumption. Deadweight loss in this instance is therefore zero. ■

5 Numerical Results for the Optimal Tax Function

Return now to the decentralized problem, but with a particular specification for the representative worker's preferences and for the tax rate function. For now, the assumption of a single worker type is maintained. Results presented here will be useful for later comparison with the cases of preference and productivity heterogeneity. Let the worker's utility be given by:

$$U = c + \alpha \log(l) + \beta \log(H),$$

where $l = 1 - L$ is leisure and tax revenue is refunded to the worker. The tax rate function $\tau(\cdot)$ is $\tau(H) = t(1 + ap_H H)$, where $t \geq 0$, $a \leq 0$. Parameters of the model are given in Table 1. The model targets the share of itemized charitable giving in total reported income, which is about 2.0% in recent years.[8] The parameters a and t of the tax function, $\tau(H) = t(1 + ap_H H)$, are chosen to maximize excess tax revenue, $R_F - R_S$, while holding social welfare and H consumption constant across regimes.[9]

[7] This result depends entirely on the nature of preferences for the tax-preferred good; if the representative worker received "warm-glow" utility from spending on H, rather than the more standard utility derived from consumption of H, crowd-out would not be perfect. However, the argument that follows does not depend on this detail.

[8] Internal Revenue Service, Statistics of Income.

[9] Any combination of parameters $\{w, \alpha, \beta\}$ that yields a particular expenditure share for the tax-favored good will generate the same excess tax revenue.

For the calibration of Table 1, optimal $\hat{a} = -1.5$ and $\hat{t} = 0.26$. Given selection of equilibrium H, this implies that the proposed regime lowers the worker's income tax rate by 0.5 percentage points relative to a choice of $H = 0$. Put differently, the increase in H resulting from the proposed regime, relative to an economy with no tax preference, lowers the worker's income tax rate by 0.1 percentage points. These figures may seem small, but it is important to keep in mind that the model is calibrated to a low (2%) fraction of income constituted by H consumption.

The tax revenue savings generated by the proposed regime, relative to the standard tax preference, are quite large despite the relatively small variation in tax rate. Holding social welfare constant and setting a and t at their optimal levels, tax revenue is higher by an amount equal to the full amount of the original tax expenditure in the standard regime, as indicated in Proposition 1. Table 2 gives results for various values of a that are consistent with equal worker utility and equal H^* for several distinct values of τ. The first line associated with each level of τ gives results for the optimal values of a (t is suppressed, as it is implied by the combination of τ, a, and equilibrium H^*). The next line, for each τ, shows the values associated with an a value that is 0.1 higher. In all cases, this increase in a raises equilibrium consumption of the tax-favored good, but lowers the revenue advantage of the proposed tax regime relative to the standard regime.

6 Heterogeneity

6.1 Preference Heterogeneity

The model presented above assumes worker homogeneity; it may be important, however, to allow for heterogeneity in workers' utility over the tax-preferred good. Both the standard and proposed regimes provide a tax advantage that is increasing in the consumption of H. The extent to which the preference is exploited depends on β, the coefficient in the utility function

$$U = c + \alpha \log(l) + \beta \log(H),$$

and it may be that heterogeneity in β affects the two tax preference regimes differently. For simplicity, assume that the social planner weights all workers equally.

However, it turns out not to be the case that a mean-preserving change in the distribution of β will differentially affect the desirability of the two regimes. Solving the model with β heterogeneity is relatively simple. Because the model is solved in partial equilibrium, with no interactions between workers, it suffices to solve separately for multiple values of β, then calculate aggregate equilibrium quantities as weighted averages of individual worker quantities, with weights given by a probability density function over β. There exists a unique (t, a) combination that sets aggregate welfare equal across standard and proposed regimes, aggregate H consumption equal across regimes, and maximizes aggregate tax revenue in the proposed regime. Given this (t, a) combination, aggregate quantities of H, C, l, and R are all identical under different densities (both symmetric and non-symmetric) that have the same mean. Thus, the argument for a declining tax rate function is unaffected by preference heterogeneity.

This is not to say that preference heterogeneity presents no difficulties for tax preferences more generally, though. "Horizontal equity", the principle that similarly-situated taxpayers (in terms of income and wealth) ought to bear similar tax burdens, may be violated by a tax preference of any kind. Under the current system, those with a particular preference for expensive housing experience a smaller tax burden than those with a weaker preference, holding income constant. This inequity would persist in the proposed regime, as workers with a relatively strong preference for H would face lower tax rates. Also, I assume that the social planner only cares about aggregate H, which may be a poor restriction.

6.2 Wage Heterogeneity

Dispersion in wages also requires investigation. Using the same strategy from above, I aggregate across workers earning various wages. Interestingly, wage dispersion does substantially affect the relative performance of the proposed tax function, narrowing the gap between the regimes. The intuition for this is as follows. The social planner is constrained to choose a tax rate function (i.e., t and a) such that *aggregate* welfare and tax-preferred consumption are equal across regimes. This may entail a tax function that is suboptimal for a particular type of worker. Indeed, the highest wage types are poorly-served by a tax function that is optimal for the aggregate. They are induced to increase their consumption of H by a large amount, and obtain lower utility than they would under the standard regime. The social planner, in turn, would prefer to tailor the tax rate function to each worker type, rather than meet the aggregate H constraint by inducing such large tax-preferred consumption by high types. As the social planner can only provide one tax function to all workers, this is impossible.

In the 2012 Current Population Survey outgoing rotation group sample, the coefficient of variation of reported hourly wages is 0.80. Assuming a normal distribution of wages with this coefficient of variation, the tax revenue advantage of the proposed regime is diminished by roughly 20%. In other words, the increase in tax revenue, relative to the equivalent standard regime, that is induced by an optimal tax rate function is only 1.6% rather than 2.0%. Table 3 compares the two regimes for various values of wage dispersion.

[This result is specific to the particular tax rate function chosen. A function with $\tau_F'' > 0$ should mitigate this problem. Will work on this shortly.]

7 SOI Calibration

The models developed previously are necessarily simplified in a variety of ways. These simplifications facilitate general statements about optimal taxation, but may affect the accuracy of quantitative predictions. In this section, the model developed previously is applied to data from the public-use Statistics of Income file. This data identifies a less restrictive version of the model in which intensity of preference for the tax-advantaged good and worker productivity have an arbitrary correlation.

[Still to do.]

8 Conclusion and Future Work

This paper demonstrates that tax preferences structured as reductions in labor income tax rates are an important part of optimal tax policy. This type of tax preference would diminish the marginal distortions associated with consumption of all goods, providing both an incentive to consume the preferred good and a reduction in deadweight loss. The efficiency gains from this innovation are potentially large, though heterogeneity in wages can both mitigate the welfare benefit and introduce "horizontal equity" concerns.

A model of worker behavior demonstrates the optimal tax preference's advantage relative to a normal deduction. For a given tax revenue target, the optimal tax preference is able to generate lower deadweight loss while holding constant the level of tax-preferred consumption, eliminating the tradeoff between raising tax revenue and encouraging consumption of a preferred good. However, wage heterogeneity of the magnitude observed in recent U.S. data re-establishes such a tradeoff, albeit one that is conducted on substantially better terms than in the standard regime.

One interesting question not answered in this paper concerns the implementation of a tax rate function when the social planner wants to encourage the consumption of multiple goods. In principle, the tax rate function could be entirely separable in the multiple goods. In a quasilinear setting, this will likely be optimal; however, practical considerations may militate against a tax function of such complexity. Instead, it could be preferable to aggregate consumption across tax-preferred goods and allow the labor tax rate to take only that aggregate as an argument. In this environment, heterogeneity in either preferences or wages may have important effects on the optimal schedule.

References

Auten, G., Sieg, H., Clotfelter, C., 2002. Charitable giving, income, and taxes: An analysis of panel data. The American Economic Review 92, 371–382.

Diamond, P., 1998. Optimal income taxation: An example with a u-shaped pattern of optimal marginal tax rates. The American Economic Review 88, 83–95.

Dunsky, R., Follain, J., 2000. Tax-induced portfolio reshuffling: The case of the mortgage interest deduction. Real Estate Economics 28, 683–718.

Gruber, J., Saez, E., 2002. The elasticity of taxable income: Evidence and implications. Journal of Public Economics 84, 1–32.

Mirrlees, J., 1971. An exploration in the theory of optimum income taxation. Review of Economic Studies 38, 175–208.

Rosen, H., 1985. Housing behavior and the experimental housing allowance program: What have we learned? NBER Working Paper 1161.

Saez, E., 2004. The optimal treatment of tax expenditures. Journal of Public Economics 88, 2657–2684.

Table 1: Model parameters and targets

Externally-calibrated		*Target*
τ	0.25	Total U.S. wage tax rate
w	1	Relative price of charitable
p_H	1	contributions = 1
Internally-calibrated		*Target*
α	0.25	Expenditure share of
β	0.01	charitable donations

Table 2: Tax Revenue for Standard and Alternative Regimes

Tax parameters	H^*	Standard Revenue	τ Function Revenue	Revenue increase (%)
		$\boldsymbol{\tau = 0.15}$		
$a = -1.39$	0.0118	0.1041	0.1059	1.68%
		$\boldsymbol{\tau = 0.20}$		
$a = -1.43$	0.0125	0.1350	0.1375	1.84%
		$\boldsymbol{\tau = 0.25}$		
$a = -1.47$	0.0133	0.1633	0.1667	2.02%
		$\boldsymbol{\tau = 0.30}$		
$a = -1.52$	0.0143	0.1886	0.1929	2.25%
		$\boldsymbol{\tau = 0.35}$		
$a = -1.59$	0.0154	0.2100	0.2154	2.53%

Notes: For any given value of τ, a is chosen to maximize the increase in revenue over the standard regime, subject to a constraint that H and worker utility be equal across the regimes. "Revenue increase" denotes the percentage increase in tax revenues generated by the alternative regime relative to the standard. "Standard" refers to the existing regime with full deductibility of tax-preferred consumption; "τ Function" refers to the proposed regime with a tax rate function given by $\tau = t(1 + a \cdot p_H H)$. Values of t are suppressed, but are determined by the values of τ and a stipulated. Workers are homogenous in preferences and wages.

Table 3: Tax Revenue as a Function of Wage Heterogeneity

Tax parameters	H^*	Standard Revenue	τ Function Revenue	Revenue increase (%)
		$\boldsymbol{\sigma_w = 0}$		
$a = -1.47$	0.0133	0.1633	0.1667	2.02%
		$\boldsymbol{\sigma_w = 0.4\overline{w}}$		
$a = -1.41$	0.0133	0.1633	0.1661	1.69%
		$\boldsymbol{\sigma_w = 0.8\overline{w}}$		
$a = -1.40$	0.0133	0.1633	0.1660	1.64%
		$\boldsymbol{\sigma_w = 1.2\overline{w}}$		
$a = -1.39$	0.0133	0.1633	0.1660	1.63%

Notes: The labor tax rate $\tau = 0.25$ throughout this table. For any given value of σ_w, a is chosen to maximize the increase in revenue over the standard regime, subject to a constraint that H and worker utility be equal across the regimes. "Revenue increase" denotes the percentage increase in tax revenues generated by the alternative regime relative to the standard. "Standard" refers to the existing regime with full deductibility of tax-preferred consumption; "τ Function" refers to the proposed regime with a tax rate function given by $\tau = t(1 + a \cdot p_H H)$. Values of t are suppressed, but are determined by the values of τ and a stipulated. Workers are homogenous in preferences but heterogenous in wages.