



The choice of the personal income tax base[☆]

Roger H. Gordon^{a,c}, Wojciech Kopczuk^{b,c,*}

^a University of California, San Diego, United States

^b Columbia University, 420 West 118th Street, MC 3308, New York, NY 10027, United States

^c NBER, United States



ARTICLE INFO

Article history:

Received 30 September 2013

Received in revised form 30 June 2014

Accepted 1 July 2014

Available online 16 July 2014

Keywords:

Income tax

Tax base

Optimal taxation

ABSTRACT

Starting with Vickrey (1945) and Mirrlees (1971), the optimal tax literature has studied the design of a personal income tax. The assumed ideal would be to tax earnings ability. Earnings ability is unobservable for tax purposes, however. Past papers have focused instead on designing a tax on labor income.

Existing tax bases, though, depend on a broader range of information about each individual than just labor income. In principle, this supplementary information can help in designing a tax that has more attractive distributional properties, by more closely approximating an ability tax. The objective of this paper is to lay out theoretically and estimate empirically how to make best use of available information about each individual in addition to earnings, in a setting where the first-best tax would be an ability tax. The theory lays out an equity/efficiency trade off when choosing the tax base. In the empirical work, we find the tax base that is best on equity grounds alone.

We find that the choice to tax couples based on their joint income, and the inclusion of dividends, interest income, and dependents' deduction in the tax base in roughly their current form can be rationalized simply based on their value in better approximating an ability tax, without any need for supplementary motivations for these provisions. However, the inclusion of mortgage and property tax payments in the list of itemized deductions cannot be defended on these grounds.

© 2014 Elsevier B.V. All rights reserved.

1. Introduction

The choice of a personal income tax consists first of the choice of a tax base and then the choice of a tax rate schedule. The past literature on the optimal design of the income tax, as exemplified by Vickrey (1945) and Mirrlees (1971),¹ presumes that the ideal tax base is the earnings ability of each individual, since this is the only characteristic that is assumed to differ across people. In practice, however, earnings ability cannot be monitored for tax purposes. A close observable proxy for earnings ability is labor income. The initial optimal tax literature presumed that labor income is the natural choice of a tax base and then derived the optimal rate schedule given this tax base.

While the correlation between labor income and ability (proxied by wage rates) is high (around 0.87), differences between the two certainly remain. Fig. 1 compares the average wage rate for each married couple in the PSID data with their joint labor income,² including a

non-parametric estimate for the curve characterizing the joint relationship that minimizes least-squared deviations. Particularly among those with lower levels of labor income, the relationship between labor income and wage rates is remarkably weak.

These limitations to the use of labor income as a proxy for ability in the tax code at least raises the question about whether there would be welfare gains from extending the tax base to include as well other observable information about individuals. Actual tax bases certainly include information beyond labor income, such as interest, dividend, and capital gains income. In addition, by taxing couples as a unit rather than taxing each spouse separately, the labor income of one's spouse affects one's own tax rate. Mortgage and property tax payments are allowed as deductions for those who itemize. The tax base is certainly more complicated than labor income. To what degree can these additional elements in the tax base be explained simply by the value of these other indicators in coming up with a tax base that better approximates ability, without taking into account any subsidiary objectives or externalities?

Atkinson and Stiglitz (1976) derive conditions under which the ideal tax base should include just labor income, and not make use of information about consumption of other goods.³ Saez (2002) then extended these results to a setting where individuals have heterogeneous tastes

[☆] Kopczuk gratefully acknowledges the financial support from the Research Council of Norway, grant 217139/H20.

* Corresponding author at: Columbia University, 420 West 118th Street, MC 3308, New York, NY 10027, United States.

E-mail addresses: rogordon@ucsd.edu (R.H. Gordon), wojciech.kopczuk@columbia.edu (W. Kopczuk).

¹ For recent work in this tradition, see Saez (2001) and Gruber and Saez (2002).

² See Appendix C for details about how the data are constructed.

³ For more recent attempts to explicate this result, see Laroque (2005), and Kaplow (2006).




Fig. 1. Prediction of the average wage rate of a couple based on joint labor income.

and not just heterogeneous abilities. His paper derives three assumptions under which the ideal tax base is just labor income, even when consumption patterns are observable for tax purposes. The first assumption is that the marginal utility of income is uncorrelated with consumption of each good among individuals with the same labor income. Given this assumption, there are no equity gains from modifying the tax base.⁴

We examine both theoretically and empirically how the optimal tax base changes on equity grounds when this assumption does not hold.⁵ In particular, we derive an expression for the potential equity gains from varying the tax treatment of individuals with the same labor income and then proceed to estimate the tax base that does best on equity grounds using PSID data.

Of course, ours is not the first paper to consider how observable information about individuals beyond their labor income may provide useful information about their ability. For example [Besley and Coate \(1992\)](#) argue that providing low-quality in-kind rather than cash transfers helps reveal who among low earners have low earnings ability, on the presumption that only those with low earnings ability are in fact willing to consume low quality goods. Similarly, [Blomquist and Christiansen \(2005\)](#) argue that users of excludable public goods should be charged a price different from marginal cost to the degree that demand depends on earnings ability. [Kopczuk \(2001\)](#) argues that tax avoidance should be facilitated if the low skilled can avoid taxes more easily than the high skilled, conditional on labor income. Closer to the choice of income tax base, [Gordon \(2004\)](#) argues that income from savings (dividends or interest income) should be part of the tax base to the degree to which those with high ability save more (or in different forms) than those with low ability, among those with the same labor income.

Our first objective, undertaken in [Section 1](#), is to reexamine the theory for the optimal tax structure, allowing for multiple observed characteristics for married couples that are each correlated with the unobserved earning abilities of the two spouses. As in Mirrlees, but applied to couples rather than single individuals, we begin with the assumption

that the first-best would be to tax couples based on their earning abilities, implemented by assuming that the social marginal utility of income for a couple under the current allocation depends solely on some function of the earning ability of the two spouses. While earning ability is not observable, many other attributes of the couple can be observed, including but not limited to the labor income of each spouse. How should this set of indicators best be used in the design of the tax base, trading off any resulting equity gains vs. efficiency costs? Here, we find that indicators beyond labor income should be included in the tax base to the extent that they are correlated with the marginal utility of income to the couple, conditional on the couple's observed labor income.⁶ These equity gains from better approximating an ability tax, though, must be traded off with the efficiency implications arising from a changing pattern of distortions to the couple's economic decisions.

Our second objective is to make use of PSID data to estimate the degree to which various indicators now in use under the personal income tax are correlated with a couple's marginal utility of income, conditional on the existing tax base. We then make use of the first-order conditions from the theory to approximate the optimal tax base from an equity perspective, given these correlation patterns.⁷ Our sample is restricted to married couples from ages 18 to 65. The indicators we consider are: dividend income, interest income, mortgage payments, property tax payments, and perhaps number of children. We also consider whether there are equity (and not just efficiency) reasons for imposing different tax rates on the labor income of primary vs. secondary earners within a couple.

Our basic results, in [Section 2](#), focus on the optimal tax base for married couples under a personal income tax in which each couple is taxed jointly based on their combined labor income along with information about the couple's joint income from interest and dividends and joint expenses for mortgage interest payments and property tax payments.

⁴ The first assumption in [Saez \(2002\)](#) is that these correlations will all be zero.

⁵ Our focus is limited, though, to the equity effects of changes to the tax base. We made no attempt to measure any resulting impact on economic efficiency. Measuring the equity and efficiency effects requires completely different approaches: The equity effects rely on cross-sectional non-causal evidence on the relationship of ability and observable attributes of a couple, whereas the efficiency effects depend on causal impacts of tax rates. In general, absent specific assumptions about heterogeneity and preferences, there is no link between the magnitudes of the two effects. A large literature estimates the latter effect and in this work we focus on the novel former one.

⁴ The other two assumptions assure that there are no efficiency gains from any deviation from a tax base equal to labor income.

⁵ Two more recent papers by [Diamond and Spinnewijn \(2011\)](#) and [Golosov et al. \(2013\)](#) incorporate specific forms of taste heterogeneity and analyze their implication for optimal capital taxation theoretically and in calibrations.

We also allow for a separate tax rate on the labor income of the lower earning spouse.

The resulting empirical estimates suggest that on equity grounds capital income should be included in the tax base with roughly the weights seen under current law: Those with the same labor income but higher capital income tend to have higher wage rates. However, we also find that, conditional on labor income, those with higher property tax payments also tend to have higher wage rates, suggesting that the tax base should be higher for those making larger property tax payments, contrary to the current income tax law.⁸ Mortgage payments plays little role in the optimal tax base, again contrary to their current deductibility. We also find, when we allow for a separate tax rate on the lower earning spouse, that this tax rate should exceed the tax rate on the labor income of the higher earning spouse: For any given labor income, the lower earning spouse tends to have a higher wage rate, implying a lower marginal utility of income.

We then consider in Section 3 several extensions of these initial results, in each case reworking the theory and redoing the estimation. Our first extension considers the optimal tax base under separate taxation of each spouse based on that individual's labor income along with information about their joint income from capital and their joint expenses for mortgage payments and property tax payments.⁹ Under separate filing, we find that the optimal tax base would give much more weight to the supplementary indicators that we consider, far beyond what is done under existing law in countries using separate filing. This is because the relationships between own earnings and household earning ability for primary and secondary earners are substantially different so that when the two groups are pooled together: own earnings do a poor job explaining household ability in the full sample even though they perform well for each group separately. As the result, the additional information about household level variables is very useful. However, even with much greater use of these supplementary indicators than is seen in practice, separate filing does much worse than joint filing on equity grounds.

So far, our results do not raise the issue of how children should enter the analysis, on the implicit assumption that policies towards children mainly involve transfer and government expenditure programs rather than the income tax. Our second extension explores how results change if we add children to the analysis, again re-deriving the theory and redoing the empirical work. In particular, we allow children to enter separately from the parents into the social welfare function, and also include information about the number of children in the family as an additional indicator in the tax base. Here, the main change is the finding that couples should be given a large deduction for each child if extra expenditures within the family are entirely on public goods within the family, but a deduction per child roughly in line with the current law if extra expenditures are spent primarily on private goods.

All of these results assume that the marginal utility of income for a couple depends solely on *that year's* wage rates, implicitly assuming that couples are credit constrained. Our third extension assumes instead that couples can borrow and lend without constraint throughout their life, implying that their marginal utility of income depends on lifetime opportunities, and not just on their wage rate in a given year. Here, we consider both an annual income tax based on information from that year as well as a tax on lifetime income. When the tax base depends on just annual income, as under current law, but the marginal utility of income depends on lifetime opportunities, the weights on capital income and property tax expenses increase: These indicators in any given year are more closely linked to average wage rates over the lifetime, controlling for that year's observed labor incomes, than they

were to that year's wage rates. If the tax base were instead to depend on average values over the lifetime for each of these variables, however, then the optimal weights on each variable return to the qualitative sizes they had in the earlier cases.

Section 4 then discusses more informally how efficiency considerations would modify our conclusions. In several settings, there would appear to be both equity and efficiency gains from particular changes to the existing tax base. In particular, our estimates suggest that capital income should be included to some degree in the tax base on equity grounds, complementing findings in studies such as Conesa et al. (2009) that it should be included on efficiency grounds. In addition, we find that reducing the deductibility of property tax and mortgage payments can be justified on both equity and efficiency grounds, ignoring any externalities from these expenditures.

Section 5 then provides a brief summary of the key findings of the paper.

2. Theory

The aim of this section is to derive conditions characterizing the optimal tax base in a one-period setting where the ideal tax base for each household h would depend solely on the earning ability of the two spouses, w_h and $w_{s(h)}$, where h identifies the higher-earning individual within each couple and $s(h)$ identifies the spouse of this individual. Given the inability to observe this tax base, the government instead must make use of a variety of other observable characteristics of couples, including their joint labor income Z_h but also possibly including some other vector X_h of observable characteristics of the couple.

While our analysis focuses on couples in order to explicitly introduce household considerations that will be of interest in our subsequent empirical implementation, the model nests the more standard (in the literature) case of redistribution across individuals.¹⁰

Our analysis builds on the existing optimal tax models dating back to the work of Mirrlees (1971). In particular, we begin by assuming:

Assumption 1. The government's objective function is to maximize $\sum_h U_h + W(R)$, where U_h measures the utility of household h ,¹¹ while $W(\cdot)$ is an increasing but concave function of overall tax revenue R that captures the welfare benefits from the expenditures financed by R .¹²

The utility $W(R)$ from tax revenue captures the sum of utilities that couples receive from government expenditures, assumed separable from the utility they receive from private goods or leisure, but it can also include any additional effects of tax revenue on the welfare of the decision-maker.

Consistent with the Mirrlees (1971) framework, we assume that the ideal tax system, at least for marginal interventions, would link tax payments to each couple's joint earnings ability. In particular, we assume:

Assumption 2. The marginal utility of income for each couple under the existing allocation satisfies $U_{hY} = V(w_h, w_{s(h)})$, where U_{hY} denotes the couple's marginal utility of income under the current allocation while $V(\cdot, \cdot)$ is an increasing concave function of the earning ability of the two spouses.

If the marginal utility of income depended on more than the two wage rates, then the first-best marginal intervention would make use of this extra information.

⁸ The contrast with the current law looks much smaller, however, if one treats the current law as including the existing property tax as well as the existing income tax.

⁹ Only the couple's joint income from interest and dividends, and their joint mortgage and property tax payments are observable in practice, given the flexibility couples have in changing the name of the owner of any given asset for tax purposes.

¹⁰ To shift to an analysis of individuals, we could set $w_{s(h)} = 0$, and then treat each individual as a separate household.

¹¹ We could easily generalize this measure of social welfare to be some function $G(U)$ of the vector of utilities received by these households, but chose not to complicate the notation.

¹² Without loss of generality, we normalize utilities so that $W' = 1$ at the existing policy.

If we were to assume homogeneous tastes, then when the marginal utilities of income of couples with the same earning abilities are equated at one market allocation, they would remain equal at all other market allocations. When tastes are heterogeneous, in contrast, the decision to make these interpersonal comparisons at the current market allocation becomes a more restrictive assumption. Given this normalization, marginal utilities of income for those with the same earning ability will in general no longer be equal at other market allocations.

Our choice to normalize utility at the existing allocation raises no new complications to the welfare analysis of tax reforms, though, if this normalization remains unchanged in the future. If utilities would be renormalized at some future date, however, when tax reform discussions again occur, then tax policies that looked optimal based on the initial normalization would in general no longer look optimal once utilities are renormalized. Policies would then iterate towards a steady state in which the existing policy used in the normalization is itself the optimal policy. One interpretation of our results is that we provide a test for whether the existing policies are such a fixed point.

Another implication of Assumption 2 is that we ignore life-cycle considerations at this point, an issue we return to in Section 4.3. Focusing on earning abilities in one year would be appropriate to the extent that couples are credit constrained, as can commonly arise earlier in life when individuals face rising wage rates over time and would like to but cannot borrow against future earnings.¹³

Following Mirrlees (1971), we also assume that wage rates (and hours of work, H_h) are not observable for tax purposes, even though joint labor income $Z_h = w_h H_h + w_{s(h)} H_{s(h)}$ and the labor income of each spouse are observable. In general, tax payments would then be denoted by some general function of observables, $T(Z_h, X_h)$. To facilitate our empirical work, and to use a structure that more closely matches the structure of actual tax bases, we simplify by assuming:

Assumption 3. The tax schedule takes the form $T(B_h)$, where $B_h \equiv Z_h + X_h \alpha$.

Here, B_h is a scalar representing the tax base, which is assumed to be a linear function of observables,¹⁴ while $T(\cdot)$ is a flexible tax rate schedule.

What then is the impact of any perturbation to the vector α on the government's objective function $\sum_h U_h + W(\sum_h T(Z_h + X_h \alpha))$, for any given tax-rate schedule $T(\cdot)$? We find that:

Theorem 1. Under assumptions 1–3, the welfare impact of a perturbation to any α_i , compensated on average for those at each value of the overall tax base by an appropriate adjustment to the tax-rate schedule, equals the sum of an equity term and an efficiency term.

The equity term equals

$$-\sum_B N_B T'(B) \text{cov}(U_{hy}, X_{hi}|B). \tag{1}$$

The efficiency term equals

$$\sum_h T'(B_h) \left(\frac{\partial Z_h}{\partial \alpha} \Big|_{\bar{c}} + \frac{\partial X_h}{\partial \alpha} \Big|_{\bar{c}} \alpha \right), \tag{2}$$

¹³ While in the cross section, earnings peak when workers are roughly in their mid-40s, time-series evidence in contrast normally shows rising earnings throughout the working career.

¹⁴ See Golosov et al. (2013) for an analysis of taxation with heterogeneous discount rates where taxes can be a general function of X_h and Z_h .

where the notation $\Big|_{\bar{c}}$ indicates that we measure the combined effects on behavior of the perturbation in α_i and the changes in the tax-rate schedule that compensate on average within each tax bracket for this perturbation.¹⁵

Proof. See Appendix A.

We find that the equity impact of a compensated increase in α_i is non-zero to the extent that X_{hi} covaries with U_{hy} , controlling for B . With a non-zero covariance, X_{hi} helps to discriminate between individuals with different marginal utilities of income who have the same value for their tax base B .

For example, assume that B represents labor income, while X_i is some form of consumption. If consumption is simply a function of labor income, as would be the case under the Atkinson and Stiglitz (1976) assumptions, and therefore does not vary across individuals with different wage rates but the same labor income, then the covariance will be zero. If tastes vary with wage rates, even given labor income, however, then the covariance will in general be non-zero and there will be equity grounds for including X_i in the tax base.

Expression (2), measuring the efficiency gain from any increase in some α_i , takes the standard form of $\sum_j t_j \Delta x_j$. By the envelope condition, changes in behavior have no direct effect on individual utilities, so that without any externalities the only net efficiency effect arises from the impact of behavioral responses on government revenue.

To find the tax base that is ideal on equity grounds alone, we need to choose the vector α that sets expression (1) equal to zero for all i . The objective of the next two sections will be to solve empirically for this vector. We ignore, though, any impact of such a tax reform on economic efficiency.

While our empirical focus is on the equity rather than the efficiency effects of any change in the tax base, we did try to make further theoretical progress characterizing the efficiency implications of a change in α . To shed additional insights, we make use of the following assumptions:

Assumption 4. The utility function for each individual is weakly separable between leisure and consumption, so that $U_h(L_h, C_h) = \tilde{U}_h(L_h, f_h(C_h))$, where C_h is the vector of consumption goods, while $f_h(\cdot)$ is a scalar measure of the quality of this consumption basket. The functions $f_h(\cdot)$ and $\tilde{U}_h(\cdot, \cdot)$ can vary by individual, however.

Assumption 5. Among those with any given tax base B , the impact of price and income changes on labor supply and demands for each good are uncorrelated with the level of demands X_h .

Assumption 6. When X_i is positively (negatively) correlated with the marginal utility of income, conditional on B , the average fraction of B spent on X_i is a decreasing (increasing) function of reported labor income.

Assumption 4 is a version of the weak separability assumption in Atkinson and Stiglitz (1976), but without imposing the assumption of homogeneous tastes. Its role is to assure that goods are equally substitutable with leisure, individual by individual, avoiding this efficiency reason for differential tax rates by good.

Assumption 5 assures that neither lump-sum redistribution nor zero-sum variation in the marginal tax rate on labor income within a group of people with the same tax base will affect overall labor income or overall demand for each of the indicators within that tax bracket. With no aggregate changes in behavior within each tax bracket, there are no efficiency consequences of these tax perturbations.

The tax modifications we will be examining do in fact alter marginal tax rates on labor income on average within a tax bracket. Assumption 6 assures that effective tax rates on labor income go down (up) when consumption of some X_i should be taxed (subsidized), so that efficiency

¹⁵ This notion of compensation is analogous to the one that has been used recently by Laroque (2005) and Kaplow (2006) to show that under separability there is always a Pareto improving reform that shifts the tax base towards simply labor income.

effects due to changes in labor supply reinforce the equity effects we focus on. The required assumption to get this result seems eminently plausible: If those with high wage rates tend to spend a higher fraction of their income on X_i , then it is natural to assume that the fraction of income spent on X_i is higher for those with higher earnings (who on average should have higher wage rates).

With these additional assumptions, we find:

Theorem 2. Under Assumptions 1, 3, and 4–6, when expression (1) is positive (negative), the efficiency loss from a compensated increase in any α_i is smaller (larger) than

$$-\sum_h T'(B_h) \left(\sum_j \frac{\partial X_{hj}}{\partial \alpha_i} \Big|_c \alpha_j \right). \tag{3}$$

Proof. See Appendix B.

Here, the notation $|_c$ reflects compensation to restore utility of the couple. Note that expression (3) omits any changes in labor supply, or the implications of changes in labor income on the vector X_i . The proof in Appendix B shows that these terms lessen (increase) the net efficiency loss when expression (1) is positive (negative), so that the remaining expression overstates (understates) the efficiency loss from an increase in any α_i .

An immediate implication of Theorem 2 is:

Corollary 1. When the equity impact of an increase in any α_i is positive (negative), when evaluated at $\alpha = 0$, we can conclude that the overall welfare impact of this tax perturbation is also positive (negative).

Proof. We know from Theorem 1 that the equity change from a perturbation of α is measured by expression (1). Hence, expression (1) is positive (negative) whenever the desired perturbation is positive (negative). Starting from $\alpha = 0$, Theorem 2 then shows that these equity-improving perturbations also raise efficiency, and therefore raise overall welfare.

Given Corollary 1, we can therefore be confident that at least some intervention is justified whenever there are equity gains from use of some indicator, X_i , in the tax base: whenever equity improves from a marginal intervention, starting from $\alpha = 0$, so does efficiency.

More generally, whatever the efficiency gain, a measure of the equity effect given by formula (1) is a necessary empirical piece of information for the purpose of evaluating the desirability of a potential reform.

3. Estimation given theory

Our aim in this section is to approximate the value of α that is most attractive on equity grounds, given the theory, and in particular to test the degree to which the optimal value of α differs from current law.

We therefore start with the actual tax base under current U.S. law, and consider modifications to this base, setting expression (1) equal to zero for all i , to solve for the optimal modifications to this tax base. Our first test will be to see if we can reject statistically that at the optimum $\alpha = 0$.

Our data were taken from the Panel Study of Income Dynamics. In specifications that we estimate, we include all individuals between ages 18 and 65 observed in any year between 1994 and 2001 who were not self-employed during that year. In some cases, we use information from other years (1968–1993) for imputations: see Appendix D for details.^{16,17} We restricted the sample to the Survey Research Center subsample, in order to have a random representative sample.

¹⁶ We chose 18 as an initial age so as to include the foregone earnings due to a college education. Time spent in school is viewed as uncompensated “work” rather than leisure, so earns a zero wage rate.

¹⁷ For the self-employed, observed earnings are a noisy and misleading measure of earning ability, since earnings can show up as dividends, capital gains, or royalty payments as well as wage income.

Z_{ht} measures the joint labor income of the h 'th couple at time t . For w_{ht} ($w_{s(h)t}$), we normally used the labor income of that individual divided by reported hours of work during the year, with imputations for those with low hours: see Appendix D for details.¹⁸

In the empirical work we set the vector X_h equal to a short list of income sources and expenditures for each individual that are in current use in the personal tax base, in addition to own labor income.¹⁹ In particular, we focus on: interest income, dividends, mortgage payments, and property tax payments. In addition, we allowed for the possibility that the labor income of the secondary (lower labor income) earner enters with a different weight than the labor income of the primary earner. Since all of these financial variables enter the current tax base, we face no need to assess whether they can be observed for tax purposes. These variables are all reported directly in the PSID. All monetary variables were converted to real 2001 dollars using the consumer price deflator.

In estimating the value of α at which expression (1) equals zero for all i , we restricted our study to married couples. The sample of single individuals is particularly heterogeneous, consisting largely of the very young and the very old, along with some divorcees and never-married middle-aged individuals, so it can yield results that vary substantially over time as sample composition changes due to changes for example in life expectancy and divorce rates. We defined an individual as “married” if he/she had been married at any point during the year.

We will ultimately examine whether separate or joint filing provides a better tax base. In this section, we assume joint filing, and then consider separate filing in Section 4.1.²⁰

To ease the computational challenge of setting expression (1) to zero for all i , note that this equation can be re-expressed as follows:

$$\begin{aligned} &-\sum_B N_B T'(B) E[U_{hY}(X_h - E(X_h|B))|B] = \\ &-\sum_B N_B T'(B) E(X_h U_{hY}|B) + \sum_B N_B T'(B) E(X_h|B) E(U_{hY}|B) = \\ &-\sum_h T'(B_h) X_h U_{hY} + \sum_h T'(B_h) X_h E(U_{hY}|B_h) \end{aligned} \tag{4}$$

where the last line simply switches to aggregating over households rather than over the tax base levels. The resulting first term is straight-forward to calculate, once we have functional forms for $T(\cdot)$ and U_{hY} . Evaluating the second term requires evaluating $E(U_{hY}|B)$.²¹ Conditional on this expression, the second term is also easy to calculate.

The key next steps in estimation are choosing functional forms for $T(\cdot)$ and U_{hY} . For $T(\cdot)$ we simply use the observed tax rate schedule for married couples in 1995.

For the marginal utility of income for a couple, U_{hY} , all we have assumed so far is that this value is a general function $V(w_h, w_{s(h)})$ of the wage rates of the two spouses, when evaluated at the current allocation. To proceed further, we make the simplifying assumption of additive

¹⁸ Wage and salary income of course is only one form of compensation. Unfortunately, we have no data on fringe benefits such as the value of health insurance or employer-financed pension contributions, or other non-wage forms of compensation.

¹⁹ Other indicators not in current use in the tax code, may also be helpful in forecasting wage rates, given labor income. Examples could be age, gender, race, ethnic background, region of residence, or even height. (See Mankiw and Weinzierl (2010) for an analysis of the optimal taxation of height, Alesina et al. (2007) for an analysis of gender-dependent taxes, and Weinzierl (2011) and Best and Kleven (2013) for an analysis of age-dependent taxes.) Since these indicators are not in current use for tax purposes, though, we do not examine their potential role. Perhaps policies that rely on them are deemed horizontally inequitable and hence are not a part of the policy toolkit for reasons not explicitly accounted for by the utilitarian framework.

²⁰ Throughout, we hold fixed the tax treatment of single individuals. At least with Nash bargaining over the division of joint income within a couple, changes in the utility arising from the outside option of each spouse due to changes in the tax treatment of single individuals would affect the division of income within the couple, and therefore affect the social welfare derived from couples.

²¹ Note that this expression does not vary across the first-order conditions for each element in α , which provides an important computational advantage.

separability for this expression: $V(w_h, w_{s(h)}) = .5(v(w_h) + v(w_{s(h)}))$. Intuitively, we are taking the average of the marginal utilities of income of the two spouses, assuming that any dollar transfer to the couple is split equally between the two spouses.²²

In our base case, we then assume that $v(w_h) = 1/\max(w_h, \$5)$, consistent with a utility function equaling $w_h^{1-\gamma}/(1-\gamma)$ with $\gamma = 1$,²³ and assuming that social-safety-net programs assure a minimum standard of living equivalent to what would be faced by someone with a wage rate equal to \$5. As robustness checks, we also report some results for alternative utility functions (other values for γ) and also for a minimum effective wage rate of \$3.

Estimation is implemented in R, relying on the np package from (Hayfield and Racine, 2008) for implementation of semi-parametric components. We construct an empirical estimate of $E[U_{HY}|B]$ by a non-parametric (kernel) regression of U_{HY} on $Z + X\alpha$. We select the bandwidth by least-squares cross validation.²⁴ We then search for the values of the vector α that simultaneously satisfy the above first-order conditions for each element of α . The number of moment conditions is the same as the number of elements of α and hence the optimal value is just identified.

Our first aim is to test whether the existing tax law is consistent with the theory. To do this, we take the actual tax base under U.S. tax law for married couples filing jointly in 1995 as our starting point, and consider modifications to this tax base (so that $\alpha = 0$ yields the current tax base). If the existing tax base makes best use of available information, then there should be no gain from introducing non-zero weights on any of the indicator variables we focus on. Through estimating the optimal weights on these indicators, we approximate the tax base that does best on equity grounds.

The existing tax base (taxable income) in the U.S. in that year equals the sum of joint earnings plus dividend and interest income, minus mortgage payments and property tax payments among those who itemize, and also depends on a few other indicators such as state of residence: see Appendix D for details.

If we cannot reject that $\alpha = 0$ at the optimum, then we can conclude that the current law is (statistically) consistent with the forecasts from our model, and conversely. To the extent that $\alpha \neq 0$, then our inference for the optimal values of α is only approximate. In particular, we do not control for possible general equilibrium changes in prices or wage rates due to a change in the tax law, do not control for changes in the marginal utility of income for each individual caused by changes in their tax liabilities, do not control for changes in job choice (and therefore in observed wage rate) resulting from changes in the tax law, and do not control for the impact of behavioral responses on the covariance terms in expression (1) that would be present when non-marginal changes in policy are implemented. These figures also ignore the efficiency-cost terms.

The results are reported in columns 1 and 2 of Table 1. In column 1, we find that each of the available indicators (except for interest income) is statistically significant, indicating that each is correlated with the couple's marginal utility of income even after controlling for the existing tax base. The results imply that the tax base that does best on equity grounds would virtually eliminate dividends from the existing tax base,²⁵ arguably consistent with their preferential treatment in the US since 2003. However, it would maintain interest income in the tax

²² See Appendix C for a simple model of Nash bargaining within a couple that implies such an equal division. (The key assumption is equal bargaining power between the two spouses.)

²³ Note that $\log(w)$ is the limiting case of this general specification as $\gamma \rightarrow 1$.

²⁴ Because least-squares cross validation is computationally intensive, we implement it by optimizing at a given bandwidth, performing cross validation at the optimum, and iterating until the bandwidth converges. In practice, results are not very sensitive to the choice of bandwidth.

²⁵ The large negative coefficient indicates that the weight of dividends in the tax base should be reduced from one to 0.26.

Table 1
Optimal tax base – joint filing.

Variables	(1)	(2)	(3)	(4)
Initial tax base	Actual tax base		Joint labor income	
Secondary earnings		1.374 (0.159)		1.338 (0.164)
Dividends	−0.900 (0.175)	−0.757 (0.238)	0.041 (0.162)	0.196 (0.230)
Interest	0.240 (0.374)	0.804 (0.336)	1.005 (0.129)	1.630 (1.493)
Mortgage	1.054 (0.167)	1.414 (0.234)	0.166 (0.127)	0.455 (0.207)
Property tax	2.301 (0.569)	3.978 (0.737)	1.540 (0.485)	3.253 (0.887)
Bandwidth	4699.669	5862.370	2545.317	6316.676
R ²	0.584	0.581	0.588	0.582
OLS R ²	0.454	0.490	0.453	0.477
N	11,359	11,366	11,355	11,368

Notes: GMM estimates of the tax base (α) setting equity term (formula (1)) to zero – optimal tax base on equity grounds. The potential tax base has the form of $Z + \alpha \cdot X$. Z is the actual tax base in specifications 1 and 2 (see Appendix D for details) and it is joint labor earnings in specifications 3 and 4. The vector X consists of variables listed in the table. The marginal utility of a couple is given by $U_{HY} = .5 \cdot \frac{1}{\max\{w_h, \$5\}} + .5 \cdot \frac{1}{\max\{w_{s(h)}, \$5\}}$. Moment conditions evaluated using Formula 4, with estimates of $E[U_{HY} | B]$ obtained by a non-parametric regression of U_{HY} on $Z + X\alpha$ using the np package (Hayfield and Racine, 2008) with bandwidth selected by least-squares cross-validation. Household sample (married couples only). Standard errors are clustered at the household level.

base (coefficient indistinguishable from zero), even though the only aim in the design of the tax base is to approximate an “ability tax.” In addition, we find that the itemized deductions for mortgage payments and property tax payments are not appropriate on equity grounds. Property tax payments in particular serve as a useful proxy for ability. Three of the four indicators we considered are statistically significant implying a strong rejection of the hypothesis that the current law approximates the optimal tax base.

When we allow spousal earnings to have a different weight than earnings of the primary earner, we find that spousal earnings should receive more than twice as much weight (a weight of $1 + 1.37 = 2.37$). This higher weight apparently compensates for the lower hours worked by the secondary earner: To impose the same tax liabilities on a primary vs. a secondary worker facing the same wage rate, we would need to multiply the reported income of the secondary earner by the ratio of the hours worked by the primary worker compared with that of the secondary worker. Ignoring any randomness in hours worked, the estimated coefficient would be consistent with the weighted average hours worked by the secondary earner (with those with a low wage rate being weighted more heavily) being only 1/2.37 times the average hours worked by the primary earner.²⁶ When spousal earnings receive more weight, the other indicators also become more important. Each helps in forecasting hours of work for the secondary earner, controlling for spousal earnings. When spousal earnings receive more weight, the corrections receive more weight as well.

Does the existing tax base even do as well as a tax base that simply equals the couple's joint labor income? To test this, we estimate the welfare effect of a shift from the actual tax base towards one where taxes are based on just the couple's joint earnings. Specifically, we set X_h equal to the difference between the couple's joint earnings and their taxable income. If $\hat{\alpha} = 0$, then the current law works well, whereas if $\hat{\alpha} = 1$, then a tax base limited to joint earnings does best. When we estimated such a specification, we found that $\hat{\alpha} = 0.965[0.106]$, implying that any deviations in the current tax base from joint earnings

²⁶ According to the data, the ratio of the average hours worked by secondary earners to those of primary earners is 0.60. The estimation procedure, though, gives much more weight to those with low wage rates, and these individuals have much lower hours.

provide no help in better approximating the earning ability of a couple. This is consistent with our estimates in column 1 of Table 1: compared to an earnings base, the actual tax base appears to account only for interest income in the right way, while it puts too much weight on dividends (compared to an earnings base) and it deviates from an earnings base in the wrong direction by allowing for the deductibility of property taxes and mortgage payments.

We next tested to see to what extent further equity gains are feasible, starting from a tax base equal to the couple's joint labor income (rather than the actual tax base as in column 1), by making use of the available indicators. Here, we are implicitly testing the assumption that consumption decisions convey no information about ability beyond what is already contained in reported earnings. Statistically significant coefficients for particular variables indicate that such variables contain usable information about ability. The specific coefficient estimates again provide an approximation to the optimal tax base (relative to taxation of joint labor income).

Estimation results are reported in columns 3 and 4 of Table 1. The results are quite similar to those implied by the coefficients in columns 1 and 2.²⁷ We find in column 3 that on equity grounds dividends should not be included in the tax base, while the coefficient on interest income is virtually equal to one, consistent with the current tax law. Mortgage payments provide little or no added information in forecasting ability, so should not enter into the tax base without some other justification. As before, we find that property tax payments not only should not be allowed as a tax deduction but instead should be taxable, since those who make higher property tax payments tend to be more able, for any given value of joint earnings. Note though that an individual's actual combined tax payments under the property tax as well as the personal income tax equal $T(B) + P$, where P denotes property tax payments. Our forecast that the optimal tax structure equals $T(B + \alpha P)$ approximates the actual tax schedule of $T(B) + P$ when $T'\alpha \approx 1$ as it roughly does given our estimates.

Table 2 provides some robustness checks. In column 1 of Table 2, we re-estimate the results in column 3 of Table 1, but now assume that the minimum effective wage rate is \$3, rather than \$5. In column 2, we assume a less concave utility function ($\gamma = .5$), while in column 3 we assume a more concave utility function ($\gamma = 2$). These re-estimates yield small changes from the results in Table 1. Only the coefficient on interest income changes significantly as we put more weight on individuals with lower wage rates.

4. Extensions

In this section, we consider three directions for extending the prior results. In each case, we revise the theoretical derivation as needed, and then redo the empirical work.

First, we examine separate taxation of each spouse compared to joint taxes on the couple. Second, we look more closely at how the presence of children in the household affects the perceived marginal utility of income for the household. Finally, we bring in life-cycle considerations, shifting from our initial implicit assumption of universal credit constraints to the opposite extreme of unrestricted ability to borrow and lend. Here, we first see how results change when the marginal utility of income depends on an average of wage rates over the life-cycle. Next we consider a tax base that uses information about lifetime income rather than annual income.

²⁷ If the actual tax base were simply the sum of components that we account for, the only difference between columns 1 and 3 (columns 2 and 4) would be due to different reference points – the tax base in one case, each individual's labor income in the other – so that coefficients should differ by one. This is an accurate description for interest and dividend income, but not for mortgage and property tax payments given that only a fraction of people itemize. The fact that the actual tax base accounts for state taxation and family size is another difference.

Table 2
Sensitivity of the optimal tax base.

Variables	Minimum wage = \$3	$\gamma = 0.5$	$\gamma = 2$
	(1)	(2)	(3)
Initial tax base	Joint labor income		
Dividends	−0.046 (0.168)	0.110 (0.206)	−0.104 (0.151)
Interest	1.116 (0.297)	0.899 (0.218)	1.609 (0.048)
Mortgage	0.106 (0.159)	0.227 (0.154)	0.139 (0.173)
Property tax	1.375 (0.689)	1.640 (0.534)	1.375 (0.600)
Bandwidth	4072.761	3923.981	4006.723
R ²	0.485	0.657	0.461
OLS R ²	0.358	0.548	0.297
N	11,355	11,363	11,355

Notes: GMM estimates of the tax base (α) setting equity term (formula (1)) to zero – optimal tax base on equity grounds. The potential tax base has the form of $Z + \alpha \cdot X$. Z is joint labor earnings. The vector X consists of variables listed in the table. The marginal utility of a couple is given by $U_{h,y} = .5 \cdot \max\{w_h, M\}^{-\gamma} + .5 \cdot \max\{w_{s(h)}, M\}^{-\gamma}$, where: $M = \$3$ and $\gamma = 1$ in column 1; $M = \$5$ and $\gamma = 0.5$ in column 2; and $M = \$5$ and $\gamma = 2$ in column 3. Moment conditions evaluated using Formula 4, with estimates of $E[U_{h,y} | B]$ obtained by a non-parametric regression of $U_{h,y}$ on $Z + X\alpha$ using the np package (Hayfield and Racine, 2008) with bandwidth selected by least-squares cross-validation. Household sample (married couples only). Standard errors are clustered at the household level.

4.1. Separate taxation

With separate taxation, each spouse is taxed separately. We set the tax base equal to that individual's labor income, and then include in the X_h vector half of the couple's joint income from interest and dividends and half of their joint expenses for mortgage payments and property tax payments.²⁸

In this section, we will estimate the optimal tax base on equity grounds when separate taxation of each spouse is used. We continue to assume that the marginal utility of income for each spouse equals $U_{h,y} = .5(v(w_h) + v(w_{s(h)}))$, consistent with any change in the couple's combined tax liabilities being split equally between the two spouses.²⁹

The estimates for the optimal α under separate taxation appear in column 1 of Table 3. The only change from column 3 of Table 1 is that we replace joint labor income with the labor income of a given spouse and then include each spouse as a separate observation. If there were perfect assortative matching, with equal labor incomes of the two spouses within any given household, then coefficients (and the R^2) would remain entirely unchanged.³⁰ To begin with, though, we find a sharp fall in the R^2 compared to Table 1. To compensate for this fall in the information provided by the labor income of each spouse for the couple's marginal utility of income, the other indicator variables become much more important components of the tax base.

In order to understand these findings, in the following two columns we re-estimate the same specification on the subsample of primary and secondary earners. For each of those groups alone, the results are much closer to those in Table 1 although the weight on all coefficients increases somewhat. Additional indicators are more useful for primary earners. Fig. 2 illustrates why pooling the two groups results in such a dramatic difference in results. Within each group separately, the relationship between couples' marginal utility and own earnings is very

²⁸ Here, we presume that any attempt under the tax law to link financial income flows or consumption expenditures to a particular spouse is fruitless, given the flexibility a couple has to reassign income and/or expenses to the spouse with the more favorable tax rate.

²⁹ To test the importance of this assumption, we also ran the same specification but assuming that any tax change for an individual goes entirely to that individual. The qualitative story is very similar to our results in column 3 of Table 1.

³⁰ Since the separate observations for each spouse within a couple have the same marginal utility of income, these observations are certainly not statistically independent. We therefore cluster at the level of the couple.

Table 3
Optimal tax base – separate filing.

Variables	Everyone	Primary earners	Secondary earners
	(1)	(2)	(3)
Initial tax base	Individual labor income		
Dividends	1.683 (0.351)	0.283 (0.081)	0.344 (0.201)
Interest	10.299 (1.787)	2.669 (0.421)	1.054 (0.495)
Mortgage	4.307 (1.428)	1.062 (0.406)	0.434 (0.182)
Property tax	21.167 (1.976)	4.361 (1.642)	2.477 (0.652)
Bandwidth	2241.064	2361.804	1596.999
R ²	0.386	0.483	0.422
OLS R ²	0.291	0.334	0.392
N	22,681	11,308	11,363

Notes: GMM estimates of the tax base (α) setting equity term (formula (1)) to zero – optimal tax base on equity grounds. The potential tax base has the form of $Z + \alpha \cdot X$. Z is individual labor earnings. The vector X consists of variables listed in the table. The marginal utility of a couple is given by $U_{HY} = .5 \cdot \frac{1}{\max\{w_h, SS\}} + .5 \cdot \frac{1}{\max\{w_{sh}, SS\}}$. Moment conditions evaluated using Formula 4, with estimates of $E[U_{HY} | B]$ obtained by a non-parametric regression of U_{HY} on $Z + X\alpha$ using the np package (Hayfield and Racine, 2008) with bandwidth selected by least-squares cross-validation. Married couples only, individual observations. In specification 1 both spouses are included. In specification 2, only primary earners (defined as the spouse with higher earnings are included) and in specification 3 only secondary earners (defined as the spouse with lower earnings are included). Standard errors are clustered at the household level (i.e., allow for arbitrary correlation in error terms over time and across spouses).

strong. However, this relationship is different for primary and secondary earners so that when the two groups are pooled together, own earnings become much more poorly associated with the marginal utility of income of a couple, especially at lower levels of income. As a result, other indicators become much more important relative to own earnings.³¹

One response, while formally maintaining separate taxation would be to include the labor income of the other spouse as a separate indicator in the tax base. Given the symmetry of the two observations within each household, the optimal weight on this indicator will be 1.0, and we return to the results in column 3 of Table 1.

Note that a number of countries (including Canada and the U.K.) make use of separate taxation, with no role for the labor income of the spouse in the tax base. We find that the resulting tax base does a poor job of approximating the marginal utility of income of a couple, even with much more aggressive use of other sources of information than is seen in practice. As Fig. 2 demonstrates, this is driven by the difference in the relationship of earnings and marginal utility of a couple for primary vs secondary earners. This is of course specific to this sample and this particular period. If earnings/marginal utility profiles of primary and secondary earners were similar, earnings would continue to be informative about ability even when the two groups are pooled. One might expect that the extent of labor force participation of women and the extent of assortative mating would affect this difference, suggesting that these factors could potentially explain why different countries pick different systems.³² These more general considerations notwithstanding, our results indicate that in the U.S. as of the 1990s, joint taxation seemed to be called for on equity grounds.

³¹ The large coefficients in column 1 simply reflect that own earnings have little explanatory power – the coefficients that we report are analogous to normalizing by the coefficient on own earnings if we estimated this specification by OLS.

³² In some cases, countries switched some aspects of their tax/benefit system from joint to separate – for example (Persson, 2004) analyzes the implications of such a reform of the Swedish Social Security system.

4.2. Children

Various new issues arise once we focus on the presence of children in the household. One easy extension is simply to include information about the number of children in the vector of indicators, X_h .

These updated results appear in column 1 of Table 4. Here, we replicate our specification from column 3 in Table 1 except that we now include number of children as an additional indicator. While the other coefficients do not substantively change, the added coefficient suggests a \$1060 deduction for each child, even though the presence of children does not affect the assumed marginal utility of income for the family. (In comparison, the size of the deduction in 1995 per child was \$2500.) Apparently, those with children tend to work longer hours, so that their labor income slightly overstates their earning ability.³³

We also consider whether the existing marginal utility of income of the household should differ across households with the same earnings potential but with a different number of children. If an ability tax were to remain the first-best outcome, then the answer would be no. So far, utility functions were normalized so as to equate the marginal utility of income among those with the same earning ability, regardless of the decisions the couple makes about what goods to buy or how much to work. The same reasoning suggests that utility functions should be normalized so that the marginal utility of income of the couple in the existing allocation should be the same regardless of their choices on family size.

However, the government's objective function plausibly includes not only the utility of the couple (which itself includes the utility they receive through their children) but also the utility of the children. To the degree to which children enter the objective function separately, the government could well want to redistribute towards families away from childless couples. Many government programs engage in such redistribution, including the EITC program, food stamps, the SCHIP program, school lunch programs, and even the public finance of education. Here, we consider how the tax structure could be affected by a concern to redistribute towards families with children.

In particular, we now define the objective function to be $\sum_h U_h + a \sum_c U_c$, where c captures the sum over the utility of children.

Previously, we implicitly set $a = 0$. Now, we instead set $a = 1$. In the process, we capture any of a variety of rationales that justify a welfare gain from redistribution from childless couples to parents with dependent children among those with the same earning ability.

The next question is how to measure the marginal utility of income of each child. Many things could affect this marginal utility, including not only the family's overall income, but also the time the parents have available to spend with their children, the number of other children in the household, and even the allocation of the family's income across goods that benefit the parents vs. the children.³⁴

For simplicity, we summarize the marginal utility of each child by the value of U_{HY} for their parents. In the process, for example, we implicitly assume that a choice by the parents to spend less time with their children in order to increase the income available to the household has no net effect on the marginal utility of the children: the extra goods simply substitute for the extra time.

If any dollar drop in a couple's tax payments leads to a dollar increase in "public goods" within the family, then the contribution to social

³³ This occurs in spite of the stylized view that the secondary earner reduces hours of work in order to take care of the children. Apparently, the primary earner increases hours by enough that earnings on net increase slightly with children.

³⁴ For example, the kids may benefit particularly from those expenditures that give the children access to better schools or the opportunity to grow up in a better neighborhood. Such considerations could potentially explain a more favorable treatment of property tax payments (largely fees for public schools) or mortgage payments (an indirect measure of the quality of neighborhood chosen by the parents). We did not see any clear way to capture the empirical size of such concerns. Interpretation of our findings should then keep these omissions in mind.




Fig. 2. Prediction of the average marginal utility of a couple based on each spouse's individual labor income.

welfare equals $(2 + N)U_{hY}$, where N measures the number of children. If instead, any extra income is divided in some manner among family members to use for added private consumption for each individual, then the contribution to social welfare of a dollar transfer to a given family equals $\sum_i d_i U_{hY} = U_{hY}$, where d_i denotes the fraction of the extra dollar going to family member i . If the funds are entirely spent on private goods, then the resulting objective function is the same as we used in the previous sections.

In column 2, we examine the opposite polar case in which the funds are entirely spent on public goods within the family, with the resulting marginal utility of the children as well as that of the parents adding to social welfare. Everything else equal, there are now welfare gains from redistribution towards larger families, since the resulting public goods are shared by more people. The main change in our estimates is that the optimal deduction per child grows dramatically, to \$29,750. While this figure is much higher than the deduction allowed at the time per dependent of \$2500, it looks much more plausible if one includes EITC, Medicaid (and perhaps other transfer programs) as part of the existing tax law. For example, the estimated \$59,500 deduction for two children results in a net dollar gain of roughly \$8200 over a law with the existing \$2500 deduction per child for those in the bottom 15% tax bracket. But lower-income couples with two kids would have been eligible at the time for up to roughly \$4000 from the EITC program, while dollar-equivalent benefits from the Medicaid program were likely larger than the remaining \$4200. Existing law, though, targets transfers much more heavily to poor families than is feasible in the tax structures we consider.

We also find that the weight on interest income becomes much larger than under current law, perhaps because childless couples have more savings.

4.3. Life-cycle considerations

All of the above derivations focus on a static setting in which individuals are observed and taxed in one year only. The ability of an individual has been measured based on their wage rate in that year. These derivations ignore, though, that individuals can potentially smooth consumption through borrowing and lending. With unrestricted borrowing and lending, their marginal utility of income should depend on their earnings ability throughout their life. In addition, utility is then affected by

a change in the tax law throughout their life, and not just in that one year.

How would the prior results change if we take into account such life-cycle considerations? The first question is how to characterize the marginal utility of income for a couple, given the set of wage rates they face at all different ages. In general, couples will differ in the weights they would place on wage rates at different ages, depending on the time pattern of their desired labor supplies. As before, we seek to develop a measure of marginal utility of income that depends on earning ability but not on tastes for how much to work, when to work, or what to consume. To summarize the vector of annual wage rates by a scalar measure of “ability”, a natural approach would be to use the present value of these wage rates: $\sum_t w_{ht}(1+r)^{-t} \equiv w_h^*$ ³⁵

Given that we do not observe individuals over their entire career, we instead use for the lifetime wage rate the forecasted wage rate evaluated at age 40, derived from a wage regression that includes fixed effects and a full set of year and age-specific dummies, estimated separately by gender.

If the value of α were to change, then the cumulative welfare impact of this change over the life of the individual equals:

$$\sum_h \left(1 - V(w_h^*, w_{s(h)}^*) \right) \left\{ \sum_t T'(B_{ht})(X_{ht} - E(X_{ht}|B_h))(1+r)^{-t} \right\} = \sum_t (1+r)^{-t} \left\{ \sum_h \left(1 - V(w_h, w_{s(h)}) \right) T'(B_{ht})(X_{ht} - E(X_{ht}|B_h)) \right\}. \tag{5}$$

As long as the composition of the population is stable over time, with the population growing at rate r , then the expression in braces in Eq. (5) can be estimated using cross-sectional data in any given year.³⁶ The only change from our prior expressions is then in the use of a lifetime rather

³⁵ There are of course many questions that can be raised here. Shouldn't different weights be applied to wage rates at different ages, for example because labor supplies are systematically different at different ages? Here, we confine our sample to individuals in prime-earning years, making this issue of secondary importance. If individuals can borrow and lend over time and have full information, then fluctuations in wage rates across years are beneficial, since labor supply can be reallocated from low-wage-rate to high-wage-rate years. Without full information, though, these fluctuations (and particularly those at older ages) impose risk-bearing costs. By ignoring the variation in wage rates, we implicitly presume that these two effects are on average offsetting.

³⁶ Here, in place of the discounting of taxes paid at older ages, we implicitly underweight the old due to population growth.

Table 4
Optimal tax base.

Variables	Control for children (1)	$(2 + \#Children) \cdot u'$ (2)	Fixed effect wage (3)	Fixed effect wage, permanent earnings (4)	All fixed effects (5)
Initial tax base	Annual joint labor income			Permanent joint labor income	
Dividends	0.048 (0.128)	−0.049 (0.103)	1.830 (0.605)	0.414 (0.110)	3.655 (0.677)
Interest	1.061 (0.128)	1.737 (0.082)	5.211 (1.314)	1.201 (0.242)	0.627 (0.704)
Mortgage	0.216 (0.140)	0.114 (0.147)	0.120 (0.293)	−0.154 (0.075)	0.327 (0.201)
Property tax	1.633 (0.528)	1.194 (0.646)	8.648 (1.976)	1.481 (0.439)	1.230 (0.671)
Children * \$10 K	−0.106 (0.046)	−2.975 (0.159)			
Bandwidth	3333.854	6247.566	5728.941	5491.711	5254.481
R ²	0.589	0.744	0.436	0.639	0.671
OLS R ²	0.453	0.621	0.326	0.428	0.510
N	11,360	11,375	11,355	11,395	2692

Notes: GMM estimates of the tax base (α) setting equity term (formula (1)) to zero – optimal tax base on equity grounds. The potential tax base has the form of $Z + \alpha \cdot X$. Z is household labor earnings in specifications 1, 2 and 3 and it is permanent household labor earnings in specifications 4 and 5. The vector X consists of variables listed in the table; in specifications 1–4, these are annual values while in specification 5 they are estimated permanent components. The marginal utility of a couple is given by $U_{HY} = .5 \cdot \frac{1}{\max\{w_h, .55\}} + .5 \cdot \frac{1}{\max\{w_{in}, .55\}}$, where w_h is the annual value in columns 1 and 2, and a fixed effect estimate capturing an individual's average wage rate over the lifetime in columns 3–5. Moment conditions evaluated using Formula 4, with estimates of $E[U_{HY} | B]$ obtained by a non-parametric regression of U_{HY} on $Z + X\alpha$ using the np package (Hayfield and Racine, 2008) with bandwidth selected by least-squares cross-validation except for specification (4) that has bandwidth assigned as an average of bandwidths for specifications (3) and (5) (all coefficient estimates are within 0.03 of the reported ones for any choice of bandwidths between 5000 and 6000). Standard cross-validation for specification (4) would generate very small bandwidths due to multiple observations at any observed level of permanent income that dominates variation in the data. Married couples only. Standard errors are clustered at the household level.

than an annual wage rate for each individual when constructing a measure of the marginal utility of income.

In column 3 of Table 3, we report our re-estimates of the specifications from column 3 of Table 1, returning to our implicit initial assumption that families spend any extra income entirely on private goods ($a = 0$). Now, however, we measure the marginal utility of income using our fixed effect estimate for the wage rate of each individual at age 40, rather than the individual's wage rate in each year. Here, we find that more (and in some cases much more) weight is given to each of the indicators than was found in our initial results. Given that annual earnings fluctuate over time due to transitory variation in an individual's wage rate, annual earnings are less effective at forecasting average wage rates than in forecasting that year's wage rate. This leaves more room for other sources of information when average wage rates are used in measuring the marginal utility of income.

One possible reason why these various indicators can help in forecasting wage rates is that they reflect decisions based on the individual's permanent earnings rather than earnings in that year. To test to see to what degree this explains the patterns we see in the data, we re-estimate the lifetime specification while replacing the current earnings by average earnings.³⁷ The results are in column 4: as expected, comparing the estimates in columns 3 and 4, more of the variation in the marginal utility of income, as measured using lifetime wage rates, can be explained by lifetime earnings than by annual earnings, leaving less room for the various indicators. The coefficients on property tax payments and interest income in particular both fall substantially. Apparently, these two indicators are particularly good proxies for lifetime earnings, giving them an important role in the tax base when lifetime earnings are not controlled for directly.

In column 5 we allow full income averaging over an individual's life, so that the tax base is average earnings plus a weighted sum of the average values of each of the indicators over the individual's life. The

main difference between this specification and that in column 4 is that dividend receipts now receive substantial weight, suggesting that over the course of a lifetime those with high dividends work fewer hours. There could be various explanations for the smaller coefficient on dividends when annual rather than lifetime data on dividends are used for each individual. One could be random variation in reported dividend receipts in any given year that average out over a lifetime. Another could be receipt of stocks as a bequest at some age, with prior hours chosen in anticipation of this bequest.

5. Equity vs. efficiency gains

In our empirical work, we focused on estimating the tax base that would be optimal taking into account solely distributional considerations. The choice of tax base clearly has important effects on economic efficiency as well. We have made no attempt to estimate the marginal excess burden of any change in the tax base. As we argued before, the evaluation of efficiency gains would constitute a totally different study, imposing very different data requirements (variation in tax incentives) and requiring very different controls to yield causal estimates.

Even without any *quantitative* information about the size of efficiency effects, however, we can at least take into account the sign of these efficiency effects. Consider each of the additional sources of information in turn:

- Spousal income*: On equity grounds, our results suggest that spousal labor income should be weighted more heavily than is the labor income of the primary earner, reflecting lower average hours of work for secondary earners. Yet on efficiency grounds, the labor income of secondary earners should be taxed more lightly, given the greater sensitivity of labor supply to the net-of-tax wage rate among secondary earners. Having the same weight on the labor income of secondary earners and primary earners can be interpreted as providing a rough balancing between equity and efficiency considerations.
- Interest and dividend income*: Here, we find that under the ideal tax base focusing solely on distributional considerations, both dividends and interest income would be included in the tax base, with the weights varying depending on the degree of income averaging. There is a large literature assessing the optimal tax treatment of

³⁷ Since we do not have data for an individual's entire life, we forecast the permanent component of earnings and of each indicator using the same approach we used in forecasting the permanent component in wage rates: we regress the log of each indicator against a fixed effect, age dummies, and year dummies, separately by gender. The forecasted value for someone aged 40 in year 1995 is then used to measure the permanent component of each variable.

savings, focusing solely on efficiency considerations. The recent literature suggests a positive tax on savings on efficiency grounds.³⁸ As a result, there appear to be reasons to include income from capital in the tax base on both equity and efficiency grounds.

- c) *Mortgage interest payments*: Here, on equity grounds, we forecast that these payments should not play any non-trivial role in the tax base. At the same time, many writers have also argued for eliminating the implicit subsidy to owner-occupied housing on efficiency grounds.
- d) *Property tax payments*: Our results indicate that property tax payments provide valuable information in forecasting wage rates, even given labor income. On equity grounds, not only should these expenditures not be deductible, but they should instead add to the tax base. On efficiency grounds, in contrast, economists often argue against the inclusion of property tax payments (and other state and local tax payments) as an itemized deduction, finding little evidence for positive externalities from state and local expenditures. Taking both efficiency and equity considerations into account, property taxes should still add at least some to the tax base. While this does not occur under our current income tax, it does occur indirectly through the local government taxation of property values.

Finally, apart from efficiency and redistributive effects, there may be additional reasons for some tax preferences, such as the deductibility of mortgage and property tax payments for itemizers. One potential reason discussed above is that these expenditures are more valuable for the children in a family than are most other forms of expenditure by the family. To the extent that the government wants to redirect expenditure patterns within the household towards those expenditures of particular value to children, then a more favorable treatment of mortgage interest payments and property tax payments would be appropriate, perhaps explaining any deviations between our forecasts and actual policies.

6. Conclusions

The aim of this paper has been to derive theoretically and then to estimate empirically how to make best use of information about household consumption patterns as well as labor income when the government is attempting to redistribute across couples with different earning abilities, but cannot observe earning ability directly.³⁹

The past theoretical work, as derived most generally in Saez (2002), focuses on a set of assumptions under which the optimal tax base is simply labor income. The first key assumption needed to reach this result is that among individuals with the same labor income a person's marginal utility of income is uncorrelated with that person's consumption of any given commodity. Under this assumption, any information about an individual's consumption patterns provides no help in forecasting that person's earning ability, given their labor income.

Our theoretical results examine what happens to the optimal tax base when this first assumption in Saez (2002) is relaxed. Our theoretical analysis works out the equity/efficiency trade-offs when constructing a tax base that maximizes the sum of utilities. On equity grounds, any specific piece of information about a household should be included in the tax base if it is correlated with the earning abilities of the two spouses, holding fixed the rest of the tax base.

Our empirical work then makes use of the PSID data to solve numerically for the tax base that is best on equity grounds, given the observed

patterns of correlation between various observables now included in the tax base and a plausible measure of the marginal utility of income for a couple as a function of the couple's earning abilities.⁴⁰

Here, we find that the implied tax base in many ways looks like the types of tax bases seen in practice. In particular, we find that the optimal tax treatment of interest and dividend income that looks best on equity grounds is quite close to current practice, without any supplementary desire to redistribute from those with large assets to those with low assets. We allow number of children to serve as a possible additional term in the tax base, and find an estimated treatment that is again rather close to current practice, under the assumption that transfers to a family are largely spent on private rather than public goods within the family.

We also find that joint taxation of the couple does far better on equity grounds than separate taxation of each spouse. In particular, in a pooled sample of primary and secondary earners, individual earnings is a poor indicator of household earnings ability even though it does reasonably well if applied to subsamples of primary and secondary earners in isolation. As a result, without a separate tax schedule for primary and secondary earners within a couple, other household-level information becomes useful.

In contrast, our results cannot easily explain the current deductibility of mortgage and property tax payments, without appealing to other motivations for these policies.⁴¹ When we allow there to be a separate tax rate on the labor income of the secondary earner and the primary earner within a family, we find another inconsistency with current practice. Our results forecast a much higher tax rate on the labor income of the secondary earner, reflecting the fact that the earning ability of a secondary earner is likely to be much higher than that of a primary earner with the same reported labor income. Perhaps the equal tax treatment we see under the current law reflects the larger presumed efficiency costs of taxing the labor income of secondary earners, given the evidence that secondary earners have much more elastic labor supply.

The specific figures we forecast for the tax base, though, are very sensitive to particular assumptions we make regarding the ability of individuals to smooth consumption over time. The above results hold when individuals have no ability to smooth income over time, or when they can not only smooth consumption over time but also smooth over time the reported figures for their taxable income. If individuals can smooth consumption but have no ability to smooth taxable income, then we find that the tax law should (on equity grounds) impose much more weight on capital income, and make taxes more heavily dependent on observed property tax payments than is seen in practice.

Appendix A. Proof of Theorem 1

We begin by differentiating the social objective function with respect to α to obtain the following welfare impact of this tax perturbation:

$$\sum_h (1 - U_{hY}) T'(B_h) X_h + \sum_h T'(B_h) \left(\frac{\partial Z_h}{\partial \alpha} + \frac{\partial X_h}{\partial \alpha} \alpha \right). \quad (A1)$$

Here, the first term measures the equity gain from transferring a given amount of revenue from households to the government, where household h pays an additional amount proportional to $T'(B_h) X_h$. The resulting net benefit depends on the difference between the marginal value of public funds, by normalization equal to one, and the marginal

³⁸ One strand of this literature, reviewed in Golosov et al. (2006), shows that discouraging saving mitigates efficiency consequences of redistribution between high ability and low ability individuals. A second strand, exemplified by Conesa et al. (2009), argues that taxation of capital income serves as an indirect way to impose age-dependent taxes, an issue we ignore.

³⁹ If the ideal is to redistribute based on earning ability, then there is an implicit assumption that the marginal utility of income at the existing allocation depends solely on earning ability. Otherwise any other information affecting the marginal utility of income should affect the ideal tax. We incorporate this assumption into our analysis.

⁴⁰ In order to focus on these neglected equity issues, we ignore in the empirical work any implications of changes to the tax base on economic efficiency.

⁴¹ One example of other motivations would be a desire to shift expenditures within the family towards those expenditures of particular value to children within the family, specifically spending on schools (largely through the property tax) and expenditures to pay for living in a better neighborhood (proxied by mortgage payments). Any other notion of externalities associated with housing or state/local government spending could naturally provide such arguments.

utility of private funds to each household h , U_{hY} . The second term measures the marginal impact on efficiency created by this increase in tax rates arising from any change in either element in the tax base.

To shed more light on the first term, we group households based on their initial tax base. With some abuse of notation, group B contains all N_B individuals with tax base equal to B . The first term can then be re-expressed as:

$$\sum_B N_B T'(B) E[(1 - U_{hY}) X_h | B]. \tag{A2}$$

An increase in any α_i implies a net increase in tax burden (assuming $X_{hi} > 0$), but one that varies by tax bracket. To compensate individuals on average within each tax bracket for this tax increase, we alter the tax schedule T . The resulting drop in tax payments for each individual in group B must then equal $T'(B) E[X_h | B]$. We denote this compensating policy intervention by n . The change in welfare resulting from this drop in the schedule of tax rates equals:

$$-\sum_B N_B T'(B) (1 - E[U_{hY} | B]) E[X_h | B] - \sum_h T'(B_h) \left(\frac{\partial Z_h}{\partial n} + \frac{\partial X_h}{\partial n} \alpha \right). \tag{A3}$$

Next, add the first term in expression (A3) to expression (A2) to measure the net distributional effects of this compensated increase in α . The resulting expression equals:

$$-\sum_B N_B T'(B) E[(X_h - E[X_h | B]) U_{hY} | B] = -\sum_B N_B T'(B) \text{cov}(U_{hY}, X_h | B). \tag{A4}$$

The efficiency impact of these combined tax changes equals the sum of the second terms in Eqs. (A1) and (A3). This is directly measured by expression (2) in Theorem 1.

Appendix B. Proof of Theorem 2

In the proof for Theorem 1, we found that the efficiency impact of the compensated change in α_i equals the sum of the second terms in Eqs. (A3) and (A1).

Consider first the effects of this combined set of tax changes on labor supply. To simplify the notation, let $\theta_{hi} \equiv X_{hi}/B_h$ denote the fraction of income spent on the i 'th indicator, and let $\theta_i^B \equiv E(\theta_{hi} | B)$ denote the expected value of this fraction for those in any given tax bracket B . Also, denote the net return to extra work by w^n . The efficiency costs due to changes in labor supply caused by a perturbation in α_i can then be written as:

$$\sum_h T'(B) \left(-\frac{\partial Z_h}{\partial w^n} \theta_{hi} + \frac{\partial Z_h}{\partial l} \theta_{hi} + \frac{\partial Z_h}{\partial l} (\theta_i^B - \theta_{hi}) + \left(\frac{\partial Z_h}{\partial w^n} \right) \Big|_c \theta_i^B + \left(\frac{\partial Z_h}{\partial w^n} \right) \Big|_c B \frac{\partial \theta_i^B}{\partial Z_h} \right). \tag{B1}$$

The first term in expression (B1) captures the labor supply effects of the drop in the real wage rate caused by the increase in α_i . The weak separability in the utility function (Assumption 4) implies that a surtax on good i of dt has the same effects on labor supply of individual h as an uncompensated increase in the tax rate on labor income of $\theta_{hi} dt$.

The sum of the second and third terms captures the income effect on labor income of the compensating transfer, where l denotes income. The sum of the first and second terms then measures the compensated effect on labor supply of an increase in α_i . The third term equals zero, given Assumption 5, which assumes that the income elasticity is uncorrelated with demand X_{hi} .

The substitution effect of the compensation depends on the change in the compensating transfer as labor income changes: $\partial \theta_i^B / \partial Z$, which equals $\theta_i^B + B \partial \theta_i^B / \partial Z$. Any resulting increase in the transfer with income is implicitly a cut in the worker's marginal tax rate. Here, we are measuring a compensated price effect, since the change in the marginal

tax rate is local to B , while the infra-marginal changes are reflected in the sum of the second and third terms.

Note that the compensated increase in the wage rate captured in the fourth term equals the expected value of the compensated cut in the wage rate captured by the sum of the first and second terms. There is no net effect on labor supply from these three terms together if the price elasticities of labor supply are uncorrelated with θ_{hi} , conditional on B , as assumed in Assumption 5.

The only term left in Eq. (B1) is the fifth term. By Assumption 6, the fraction of expenditures grows (shrinks) with labor income for goods where the optimal α_i is positive (negative). This efficiency term then has the same sign as the equity term, implying that equity effects and those efficiency effects arising from a change in labor supply both push in the same direction.

What about the efficiency terms arising from changes in the X_{hj} ? Using the same steps as above, we can express the efficiency impact of the compensated increase in α_i by:

$$-\sum_{h,j} \alpha_j T'(B_h) \left[\left(\frac{\partial X_{hj}}{\partial \alpha_i} \right) \Big|_c + \frac{\partial X_{hj}}{\partial l} \frac{dZ_h^n}{d\alpha_i} \right]. \tag{B2}$$

As in the derivation of the impact of the change in α_i on labor supply, the first term here captures the effects of a compensated change in the price of the i 'th indicator on each of the X_{hj} .

The remaining term captures the impact of the resulting change in net-of-tax labor income on demand for each of the X_{hj} . If $\alpha_j > 0$, we concluded above that net-of-tax labor income increases, and conversely. This second term therefore increases (reduces) efficiency when expression (1) is positive (negative). We therefore conclude that the efficiency impact of an increase in α_i is less (more) severe than would be measured by the first term in Eq. (B2) when expression (1) is positive (negative), due both to the change in labor supply captured by the fifth term in Eq. (B1) and to the impact of this change in labor supply on demand for the indicators, as measured by the second term in Eq. (B2).

Appendix C. Behavior of married couples

Assume that married couples choose their joint labor supplies and joint consumption levels based on Nash bargaining over the division of their joint resources. Under Nash bargaining,⁴² each spouse has as a fall-back position the utility he or she would have as a single individual.

Denote the utility function of the husband by $U^H(C_H, L_H)$, where C_H is consumption and L_H is leisure. Similarly, let the utility function of the wife be $U^W(C_W, L_W)$. If single, consumption and leisure satisfy $C_{HS} = w_H(1 - L_{HS}) - T_{HS}$, and similarly for the wife, where T_{HS} denotes tax payments when single. As a married couple, their budget constraint equals:

$$C_{HM} + C_{WM} = w_H(1 - L_{HM}) + w_W(1 - L_{WM}) - T_M. \tag{C1}$$

Here, T_M measures their combined tax payments as a couple, either under joint filing or separate filing.

Under Nash bargaining, they choose their consumption levels and labor supplies to maximize

$$\left(U^H(C_{HM}, L_{HM}) - U^H(C_{HS}, L_{HS}) \right) \left(U^W(C_{WM}, L_{WM}) - U^W(C_{WS}, L_{WS}) \right) \tag{C2}$$

subject to the budget constraint (C1). To simplify the resulting algebra, we take a first-order approximation to each change in utility, e.g. $U^{HM} - U^{HS} = U_C^{HM} \Delta C_H + U_L^{HM} \Delta L_H$.

⁴² Here we assume that the two spouses have equal bargaining power.

The first-order conditions for the labor supplies of each spouse remain $U_C^H w_H^n = U_L^H$ and $U_C^W w_W^n = U_L^W$, where w^n denotes the net-of-tax wage rate. Substituting, the Nash bargain maximizes:

$$U_C^H U_C^W (\Delta C_H + w_H^n \Delta L_H) (\Delta C_W + w_W^n \Delta L_W). \quad (C3)$$

In response to any change in the tax law, say a change in α , the budget constraint implies:

$$\frac{\partial C_H}{\partial \alpha} + w_H^n \frac{\partial L_H}{\partial \alpha} + \frac{\partial C_W}{\partial \alpha} + w_W^n \frac{\partial L_W}{\partial \alpha} = -\frac{\partial T_M}{\partial \alpha}. \quad (C4)$$

We can rewrite this Nash bargaining problem as maximizing AB such that $A + B = x$, yielding a solution of $A = B$. We then infer that:

$$\frac{\partial C_H}{\partial \alpha} + w_H^n \frac{\partial L_H}{\partial \alpha} = \frac{\partial C_W}{\partial \alpha} + w_W^n \frac{\partial L_W}{\partial \alpha}. \quad (C5)$$

Each side of Eq. (C5) measures the dollar equivalent impact of the policy change on the utility of each spouse, implying that they share equally in the tax change.

Appendix D. Data

We rely on PSID data for 1968–1997, 1999 and 2001. We extract the following categories of variables (further details follow): demographic variables (gender, age, marital status, number of children, state of residence), hours, labor income, business and farm income, business and farm assets, unemployment income, mortgage payments, property taxes, dividend, interest income and other categories of capital income when available. We rely on this information to construct the following variables used in the analysis: wages, labor income, dividend, interest income, mortgage payments, property taxes, marginal tax rates, and total federal and state tax liabilities. Because in some cases questions in PSID change over time, we devoted considerable effort to come up with definitions that are consistent over time. We discuss our procedure for each variable in more detail below.

Generally, we restrict the sample to the Survey Research Center sample which is a random equal probability sample of the U.S. population. We exclude the so called Survey of Economic Opportunities sample that was not consistently covered throughout the whole period (changes occurred in 1997) and is not representative of the U.S. population as a whole. We also exclude the Latino sample that was followed for a short period of time only. We use in our analysis data for heads and wives, and do not rely on information about a person when she was a dependent. We do not use PSID sampling weights anywhere. Unless otherwise indicated, we adjust dollar variables for changes in the price level using the CPI-U index, using 2001 as the baseline.

We use the data between 1968 and 2001 to construct our variables – in particular, to measure permanent ability or current ability when current wage information is unreliable – but in the actual empirical specifications we only include observations between 1994 and 2001 because interest and dividend income cannot be observed separately before 1994.

D.1. Labor income

We construct full labor income by adding up wages and salaries, bonuses, overtime payments, tips, commission income, professional practice income and miscellaneous labor income. We do not include individuals owning a business because the labor part of business/farm income is imputed in the PSID with rules regarding spousal shares and labor/capital division varying somewhat over time, but with not enough information to adjust to any standard definition.

D.2. Asset income

PSID has always included questions about some forms of capital income but due to changes in how the questions were asked and changes in the detail of breakdown of asset income into various categories, we can only construct a consistent definition starting in 1984. For 1984–1992, we construct a single variable “dividend, interest and trust income” for both heads and wives. This variable specifically excludes other categories of asset income that cause definitional problems prior to 1984 such as rental income, alimony, business and farm income, “market gardening” and “roomers and boarders.” It is constructed separately for heads and wives but in practice we use it aggregated to the family level. Starting in 1993, we can actually separate dividend, interest and trust income for both heads and wives and we do so.

D.3. Wage rate

While a wage rate measure is already available in the PSID, our labor income measure does not exactly coincide with the PSID measure and therefore we calculated the wage rate by dividing our definition of labor income by the hours reported in the PSID. We do so for those who work more than 500 hours in a given year. For those who work at most 500 hours, we impute the wage rate by first regressing log wages when working more than 500 hours on the full set of year and age dummies and individual fixed effects, for men and women separately, and then using fitted values from this regression (accounting for the individual fixed effect) in years when a person works 500 hours or less. Although we use data starting with 1984 for most of the analysis, the wage imputations and lifetime wage calculations use wage rates starting from 1968.

D.4. Deductions

Information about mortgage payments and property tax payments is systematically available in the PSID starting in 1984. In some cases mortgage payments include property taxes. In those cases, we reduce mortgage payments by separately reported property taxes if the latter is not greater than the mortgage payments.

D.5. Other issues

State of residence is not available in years 1994–1997. In these cases, we assign state of residence from the nearest year among 1992, 1993, 1999, 2001 for which state of residence is available.

We use the term “secondary income” to denote the spouse with lower income.

D.6. Taxes

In estimating the equity term based on formula (4) we use federal marginal tax rate schedule as of 1995 for childless married couples filing jointly applied to the appropriate tax base indexed to 1995 using Social Security average wage index. For specifications 1 and 2 in Table 1, we apply this marginal tax rate schedule to the tax base under the actual law (described below). For all other specifications, we apply it to labor income measure used in the particular specification.

In order to construct taxable income for specifications 1 and 2 in Table 1, we rely on the TAXSIM calculator available at <http://www.nber.org/taxsim>. We calculate the federal taxable income for each individual by accounting for labor income of both spouses, unemployment insurance, dividends and interest income, property taxes, mortgage deduction and state of residence. (We do not account for the presence of children in our initial specifications, deferring to Section 4.2 a discussion of the role of children.) We define our measure of taxable income by adding the standard deduction and personal exemption (\$11,550) to the TAXSIM-defined measure. For those with zero TAXSIM taxable

income (453 out of 11,422 observations), we use AGI instead (except for 43 cases, AGI is smaller than \$11,550).

References

- Alesina, Alberto, Ichino, Andrea, Karabarbounis, Loukas, 2007. Gender based taxation and the division of family chores. Working Paper 13638 National Bureau of Economic Research (November).
- Atkinson, Anthony B., Stiglitz, Joseph E., 1976. The design of tax structure: direct versus indirect taxation. *J. Public Econ.* 6 (1–2), 55–75 (July–August).
- Besley, Timothy, Coate, Stephen, 1992. Workfare versus welfare incentive arguments for work requirements in poverty-alleviation programs. *Am. Econ. Rev.* 82 (1), 249–261 (March).
- Best, Michael, Henrik Kleven, Optimal income taxation with career effects of work effort, February 2013. LSE, mimeo.
- Blomquist, Soren, Christiansen, Vidar, 2005. The role of prices for excludable public goods. *Int. Tax Public Financ.* 12 (1), 61–79 (January).
- Conesa, Juan Carlos, Kitao, Sagiri, Krueger, Dirk, 2009. Taxing capital? Not a bad idea after all! *Am. Econ. Rev.* 99 (1), 25–48 (March).
- Diamond, Peter, Spinnewijn, Johannes, 2011. Capital income taxes with heterogeneous discount rates. *Am. Econ. J. Econ. Policy* 3 (4), 52–76.
- Golosov, Mikhail, Tsyvinski, Aleh, Werning, Ivan, 2006. New dynamic public finance: a user's guide. NBER Macroeconomics Annual, pp. 317–388.
- Golosov, Mikhail, Troshkin, Maxim, Tsyvinski, Aleh, Weinzierl, Matthew, 2013. Preference heterogeneity and optimal commodity taxation. *J. Public Econ.* 97 (1), 160–175.
- Gordon, Roger H., 2004. Taxation of interest income. *Int. Tax Public Financ.* 11 (1), 5–15 (January).
- Gruber, Jonathan, Saez, Emmanuel, 2002. The elasticity of taxable income: evidence and implications. *J. Public Econ.* 84 (1), 1–32 (April).
- Hayfield, Tristen, Racine, Jeffrey S., 2008. Nonparametric econometrics: the np package. *J. Stat. Softw.* 27 (5).
- Kaplow, Louis, 2006. On the undesirability of commodity taxation even when income taxation is not optimal. *J. Public Econ.* 90 (6–7), 1235–1250 (August).
- Kopczuk, Wojciech, 2001. Redistribution when avoidance behavior is heterogeneous. *J. Public Econ.* 81 (1), 51–71 (July).
- Laroque, Guy, 2005. Indirect taxation is superfluous under separability and taste homogeneity: a simple proof. *Econ. Lett.* 87 (1), 141–144 (April).
- Mankiw, N. Gregory, Weinzierl, Matthew, 2010. The optimal taxation of height: a case study of utilitarian income redistribution. *Am. Econ. J. Econ. Policy* 2 (1), 155–176 (February).
- Mirrlees, James A., 1971. An exploration in the theory of optimum income taxation. *Rev. Econ. Stud.* 38 (114), 175–208 (April).
- Persson, Petra, Social insurance and the marriage market, February 2004. Stanford University, mimeo.
- Saez, Emmanuel, 2001. Using elasticities to derive optimal income tax rates. *Rev. Econ. Stud.* 68 (1), 205–229 (January).
- Saez, Emmanuel, 2002. The desirability of commodity taxation under non-linear income taxation and heterogeneous tastes. *J. Public Econ.* 83 (2), 217–320.
- Vickrey, William, 1945. Measuring marginal utility by reactions to risk. *Econometrica* 13 (4), 319–333 (October).
- Weinzierl, Matthew, 2011. The surprising power of age-dependent taxes. *Rev. Econ. Stud.* 78 (4), 1490–1518 (October).