

The Choice of the Personal Income Tax Base

Roger H. Gordon and Wojciech Kopczuk
UCSD and Columbia University

The Suntory Centre
Suntory and Toyota International Centres for
Economics and Related Disciplines
London School of Economics and Political Science
Houghton Street
London WC2A 2AE
Tel: (020) 7955 6674

Abstract

Starting with Mirrlees (1971) and Vickrey (1945), the optimal tax literature has studied the design of a personal income tax. The ideal would be to tax earnings ability. Earnings ability is unobservable for tax purposes, however, and 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. 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, when constructing a tax base that is most attractive on distributional grounds.

To begin with, we find that the current tax base does slightly less well than the far simpler tax base equal just to a couple's joint earnings. In accordance with current practice, we find that the optimal tax base should include capital income, at least to some degree. In contrast to current practice, property tax payments and mortgage interest payments should not be deductible, since these deductions are costly on equity and presumably on efficiency grounds. We also find that joint filing and separate filing by a couple have similar consequences on equity grounds.

Public Economics Programme

The Public Economics Programme was established in 2009. It is located within the Suntory and Toyota International Centres for Economics and Related Disciplines (STICERD) at the London School of Economics and Political Science. The programme is directed by Frank Cowell and Henrik Kleven. The Discussion Paper series is available free of charge. To subscribe to the PEP paper series, or for further information on the work of the Programme, please contact our Research Secretary, Leila Alberici on:

Telephone: UK+20 7955 6674

Fax: UK+20 7955 6951

Email: l.alberici@lse.ac.uk

Web site: <http://sticerd.lse.ac.uk/PEP>

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 (wage rate) 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 (wage rate times hours of work), so that 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.

Are there any welfare gains, though, from including in the tax base not just labor income but also 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 based on an optimal income tax framework, without adding subsidiary objectives or externalities?

Atkinson and Stiglitz (1976) derive conditions under which the ideal tax base should not make use of information about consumption of other goods, and include just labor income.² One key condition is that consumption goods not differ in the degree to which they are complements to leisure. A second condition is that the amount an individual consumes is not correlated with the individual's marginal utility of income, conditional on earnings. It is this second condition that we test.

Are there consumption choices that reveal information about ability beyond what is revealed by reported earnings? The tagging argument of Akerlof (1978) in the context of welfare programs presumes so; for example Besley and Coate (1992) argue that providing low-quality in-kind rather than cash transfers help reveal who among low earners has 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.

Are there observables that help detect high vs. low wage rates among individuals with the same labor income? The first objective of this paper is to explore empirically using PSID data to what degree the observable variables that are commonly part of the income tax base help forecast earnings ability (wage rates) among those with any given level of labor income. The variables we

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

²For more recent attempts to explicate this result, see Saez (2002), Laroque (2005), and Kaplow (2006).

focus on are interest and dividend income, the labor income of one's spouse, and expenditures on mortgage interest and property taxes.

Why might it be plausible that individuals with the same labor income but different underlying abilities choose different consumption patterns? Possible reasons include:

a) Among those with any given level of labor income, those with higher earnings ability may save more.³ Higher income from savings then signals higher earnings ability.

b) Conditional on labor income, individuals with higher ability may own a larger house, and so have a larger mortgage, perhaps due to sorting across locations based on ability/education.

c) Holding labor income fixed, those with higher wage rates have more leisure, so have higher potential labor income to help in the event of income shocks. This may make it easier to invest in riskier securities, suggesting that higher wage-rate individuals will have more dividend income, holding labor income fixed.

d) Among those with any given level of labor income, those with higher earnings ability are likely to value education for their kids more highly, so will live in communities with higher public school spending and therefore higher property taxes.

e) The labor supply of secondary earners is very heterogeneous, with some out of the labor force and others working full time. Despite the sensitivity of the participation decision of secondary earners to their wage rate, their observed earnings are only weakly linked to wage rates. Due to assortative mating, however, the earnings ability of spouses are highly correlated. The earnings of the primary earner can then provide valuable supplementary information about the earnings ability of the secondary earner, even conditional on that individual's reported earnings.

f) Within a couple, there may be some specialization in market work vs. work within the household, so that the higher wage-rate spouse increases labor supply while the lower-wage spouse focuses more on non-market activities. As a result, a low-earning spouse may signal high hours for the primary earner, so a lower wage rate than would otherwise have been expected given observed labor income.

Our results show that all of the additional observations we examine help in forecasting wage rates, even after controlling for labor income. Property tax payments are particularly useful. Capital income also helps forecast wage rates.

The second main objective of the paper is to solve for how this information should be used in the design of the tax base, trading off any resulting equity gains vs. efficiency costs. Finally, the paper reexamines the data to provide evidence needed in choosing the tax base suggested by the theory.

The outline of the paper is as follows. We begin in section 1 with a general discussion of an "ability" tax. In particular, to what degree should any of the underlying differences across individuals in addition to differences in earnings ability affect the ideal tax base? The analysis in the rest of this paper follows the past literature in presuming that the ideal tax is an "ability" tax, but this

³This could be the case, for example, if those with high ability tend to discount future utility less severely, a possibility focused on in Diamond and Spinnewijn (2010).

section at least highlights questions that might be raised about this characterization of the ideal tax.

Section 2 provides some preliminary empirical evidence on the extent to which observable information in addition to own labor income helps forecast an individual's earnings ability, as measured by the wage rate, using data from the PSID.

Section 3 then generalizes the standard optimal tax model to allow for multiple sources of information about the individual, and derives expressions characterizing the optimal tax base. Section 4 makes use of these expressions and the PSID data first to test whether the existing tax base makes optimal use of available information and then to approximate the tax base that would be most attractive on equity grounds. Section 5 then discusses the resulting equity vs. efficiency trade-off, and draws several conclusions about tax reforms that might be attractive on both equity and efficiency grounds. 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. Finally, our results imply similar equity consequences of joint and separate taxation of couples.

1 Is an "ability" tax the ideal tax?

In order to assess the appropriate tax base, the past literature has focused on a utilitarian objection function, where welfare depends on the sum of individual utilities plus some function of tax revenue. Individual utilities are commonly assumed to differ only due to differences in earnings ability. Those with higher ability are presumed to have a lower marginal utility of income.⁴

The optimal tax literature then presumes that "ability" cannot be monitored for tax purposes. It instead focuses on labor income as the tax base: labor income is the key observable variable linked to earnings ability, since it equals the individual's wage rate times hours of work. The optimal tax then trades off equity gains with efficiency costs from discouraging labor supply.

Of course, individuals differ in behavior for many reasons beyond differences in earnings ability. Tastes for consumption vs. leisure, tastes for consuming now vs. in the future, tastes for housing and local public services vs. other goods, all can differ among individuals with the same labor income. Given these multiple differences across individuals, how compelling is it to assume that marginal utilities of income differ across individuals only due to differences in earnings ability, the implicit assumption in the past optimal tax literature? Do differences in tastes convey further information about marginal utility of income, conditional on earnings ability?

Consider, for example, an additively separable utility function: $U_h(C_h) - V_h(L_h)$, where L_h denotes labor supply for individual h , $C_h = w_h L_h$ measures her consumption, while w_h is her hourly wage rate (her "earnings ability"). In equilibrium, each individual supplies labor until $w_h U'_h = V'_h$.

⁴The shape of this function is key in any quantitative discussion of the optimal tax rate schedule, although qualitative results need not depend on it.

Given w_h , chosen hours can be higher because the utility cost of time at work, V'_h , is lower evaluated at any particular L_h , or because U'_h is higher at any given value of C_h . In the first case, the equilibrium marginal utility of income, U'_h , is lower, whereas in the latter case the equilibrium U'_h is higher. Variation in stamina (V'_h) or in needs (U'_h) both yield the same observable outcomes, so cannot be distinguished based on data for C_h , L_h , and w_h .

The optimal tax literature usually assumes that U'_h differs across people simply due to differences in w_h , so that U'_h is the same regardless of chosen hours of work. This in effect assumes that whether the variation in behavior is due to variation in needs or in stamina is unobserved and equally likely. If, for any given wage rate, hours in fact vary across people primarily because of variation in stamina, then those with higher hours have a lower equilibrium value of U'_h , justifying on equity grounds a tax on those who work more. Conversely, hours could vary across individuals due to variation in needs, so that those working more should be treated more leniently. In either case, there are other dimensions beyond ability affecting the marginal utility of income (differences in “needs” and “stamina”), raising a question of how to combine these indicators to define the underlying marginal utility of each individual.⁵

Similarly, those who choose to save more could do so because the marginal utility of consuming now is lower, or conversely because the marginal utility of consuming later is higher. In the first case, those who save more have lower marginal utility of income, given “earnings ability,” and conversely.

By focusing on “earnings ability,” we view each of these two alternative scenarios as “equally likely”. Formally, we assume that observed labor supply, savings, or other choices reveal no information about an individual’s marginal utility of income, given their wage rate. An individual’s marginal utility of income at the observed allocation is simply assumed to depend on their wage rate.⁶

Similarly, are there equity grounds for taxing an individual differently depending on her marital status, or the “earnings ability” of her spouse, given her own “earnings ability”? Marrying, or marrying a higher wage spouse, could in principle reflect a different dimension of ability: attractiveness as a possible mate. Conversely, marrying (or marrying a more “able” spouse) can just reflect a willingness to put up with a lot in order to have access to the extra income. Even if we presume the same utility level for individuals who make different choices, the individual with a higher earning spouse may have a lower marginal utility of income. Whenever an individual has an unusually high-earning spouse, though, the spouse has an unusually low-earning mate. While this argument suggests treating the low-earnings spouse as having a lower marginal utility of income, given “earnings ability,” it simultaneously suggests treating the high-earning spouse as having a higher marginal utility of income. On net, assuming that the overall “ability” of the couple is unaffected by the nature of the

⁵One observable reason for having less “stamina” could be poorer health, providing a rationale for treating those with poorer health more favorably under the tax law. Health problems, though, can also directly affect U'_h .

⁶This normalization of utility functions is parallel to but not identical to the approach used in Weinzierl (2009). The difference is that Weinzierl instead assumes that marginal utilities of income are equal for those with the same wage rate at the *laissez-faire* allocation (without taxes), implying that marginal utilities differ at the observed allocation among those with the same earnings ability, due to the differential effects on utility of taxes (and presumably of the associated changes in government spending and in market prices) arising from heterogeneous preferences.

match is a not unreasonable choice.

What impact do children have on an individual's earnings ability? From one perspective, the decision to have children is analogous to the decision to allocate time/resources to other activities, and in these other cases we have assumed equal marginal utility of income regardless of the choices made. Even if we maintain that assumption here, though, children can affect social welfare directly, and not simply through their effect on the utility of the parents. Transfer programs in fact often seem to direct resources towards poor children in particular, and not just towards poor families, e.g. aid to families with dependent children, EITC (where benefits are heavily dependent on the presence of children), SCHIP, and even Medicaid (where most of the expenditures are on children). In order to avoid dealing with children in the analysis, our assumption will be that the current treatment of children in the law handles appropriately the implications of children for ability to pay. We take these provisions as given, and focus on the rest of the tax base.

Following the past literature, we therefore stick with the assumption that marginal utilities differ across people at the observed allocation due just to differences in earnings ability. Given this ability, some have chosen to marry and some not, some work longer hours and others shorter hours, and some consume now and others consume later. Regardless, all are assigned the same marginal utility of income if they have the same "earnings ability."

A second key assumption we make is that "earnings ability" can be approximated by an individual's wage rate. There can be many reasons, though, why "ability" differs from observed wage rates. To begin with, some individuals choose more pleasant jobs (being a professor?) rather than higher paying but less pleasant jobs (being an investment banker?). The observed wage rate then does not fully characterize the range of possible jobs the individual could have taken.

Education and on-the-job training also affect wage rates. If we take into account the foregone income and the out-of-pocket costs of education, though, then we would still measure the degree to which some individuals have better options than others.

Observed wage rates also leave out nonwage but still monetary compensation. We have no information, for example, about employer pension contributions or compensation through options or stock. In addition, the labor-contracting literature suggests that wage rates can differ period by period from earnings ability, even if their present values correspond. While there is a long tradition in labor economics of using the market wage rate as a measure of earnings ability, we do recognize the possible omissions.

2 Forecasting wage rates based on observables

Building on Atkinson and Stiglitz (1976), Saez (2002) shows that one key assumption needed for excluding consumption choices from the tax base is that these choices do not help forecast the individual's marginal utility of income, after controlling for labor income. This is testable if wage rates are a reasonable proxy for ability and marginal utility depends on ability only. The objective of this section is then to see to what degree observable information now part of the income tax base

helps forecast wage rates even controlling for labor income.⁷

Our data were taken from the Panel Study of Income Dynamics. We included all individuals between ages 18 and 65 observed in any year between 1968 and 2001 who were not self-employed during that year.^{8,9,10} We restricted the sample to the Survey Research Center subsample, in order to have a random representative sample.

Let Z_{ht} denote reported wage and salary income of individual h at time t . For w_{ht} , we normally used Z_{ht} divided by reported hours of work during the year. When reported hours of work were low, however, we worried that this measure could be quite noisy, and of course it is undefined if hours equal zero. For cases with hours worked in the year below 500, we imputed wage rates based on a regression forecasting $\log(\text{wage rate})$ as a function of individual fixed effects, year dummies, and age dummies, estimated separately by gender.^{11,12}

In the empirical work we focus on 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. In particular, we focus on: spouse's labor income, interest income, dividends, mortgage payments, and property tax payments. Since these are in current use, we face no need to assess whether they can be observed for tax purposes.¹³ 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 relationship between this list of observables and wage rates, we restricted our sample 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 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.

⁷Building on the approach introduced in this paper, Pirtillä and Suoniemi (2010) pursue analysis similar to the one in this section using Finnish data.

⁸We 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.

⁹For the self-employed, observed earnings are a noisy and misleading measure of earnings ability, since earnings can show up as dividends, capital gains, or royalty payments as well as wage income.

¹⁰We use the data from 1968 and 2001 to construct our variables but in the actual empirical specifications we rely on observations between 1994 and 2001 because interest and dividend income cannot be observed separately before 1994.

¹¹For those first entering the labor force after age 18, we discounted their initial observed wage rate by 10% per year to construct wage rates back to age 18, on the presumption that the return to education/maturity was roughly 10% per year.

¹²Wage and salary income of course is only one form of compensation. Unfortunately, we have no data on fringe benefits such as employer-financed pension contributions, or other non-wage forms of compensation.

¹³Other 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 (2010) 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.

Under U.S. tax law, married individuals normally file a joint tax return though, for example, in Canada and the U.K. they each file a separate return. We will ultimately examine whether separate or joint filing provides a better tax base. Let us first presume separate filing, and forecast the wage rate of each married individual as a function of own labor income and couple's earnings from financial assets and their joint expenses for property tax and mortgage interest payments.

In these initial tests, we estimate nonlinear regressions of the form

$$w_{ht} = f(Z_{ht} + X_{ht}\alpha) + \tilde{\epsilon}_{ht} \quad (1)$$

Here, w_{ht} is the wage rate of individual h in year t , Z_{ht} is this individual's labor income that year, X_{ht} is a vector of other information about this individual, and $\tilde{\epsilon}_{ht}$ is an iid error term.

The relationship between wage rates and observables can be highly nonlinear. To allow for this, we introduce an unconstrained function $f(\cdot)$ relating the observable information to the wage rate.¹⁴ The estimation procedure is semi-parametric minimum least squares. In particular, we use a kernel estimator for the function $f(\cdot)$, with the bandwidth chosen following the procedure recommended in Härdle et al. (1993). The estimates for the function $f(\cdot)$ will be of direct interest, since they describe the strength of the local link between the tax base and the wage rate. Under this specification, we expect substantial heteroskedasticity in the error term. To deal with this, we use a type of GLS procedure also suggested by Härdle et al. (1993). In the first stage, we estimate equation (1) as written. We then forecast the square of the resulting residuals as a function of the square of the estimated value of $f(Z_{ht} + X_{ht}\alpha)$. Finally, we weight observations by the inverse of the forecasted standard deviation for $\tilde{\epsilon}_{ht}$ in a second stage regression. Reported results come from this second-stage regression.

The function $B_{ht} \equiv Z_{ht} + X_{ht}\alpha$ characterizes the tax base that most closely approximates an individual's earnings ability. In column 1 of Table 1, we report estimates for α when forecasting the wage rate for each married individual as a function of own labor income, the couple's income from interest and dividends, and the couple's joint expenses for property tax and mortgage interest payments. If the assumptions in Atkinson and Stiglitz (1976) are right, the coefficients α should all be insignificantly different from zero. In contrast, we find that all of the coefficients of these supplementary variables are statistically significantly different from zero. The most striking coefficient is that on property tax payments. One interpretation is that, controlling for own labor income and mortgage payments, those with a higher wage rate demand more local public services, presumably better schooling for their kids. Those with higher wage rates (controlling for own labor income) also have more interest and dividend income.¹⁵ Note that when the forecasted wage rate is higher, forecasted hours of work are lower.

Those with higher mortgage payments also seem to have higher wage rates, controlling for the other variables. Perhaps individuals tend to sort across communities based on ability, with more expensive houses in higher ability communities.

¹⁴This function is normalized so that the coefficient on Z_{ht} by construction equals one.

¹⁵Perhaps, those with higher wage rates give more weight to future utility, so save more even controlling for current labor income. This is the setting examined by Golosov et al. (2010) and by Diamond and Spinnewijn (2010). See Banks and Diamond (2010) for a summary of the empirical evidence consistent with this pattern.

In column 2, we report comparable estimates for the special case in which we constrain the function $f(\cdot)$ to be linear, so that the specification is equivalent to ordinary least squares. The coefficients are similar to those in column 1. Inspection of the estimated function $f(\cdot)$ reveals that, for most of the earnings distribution, it in fact is quite close to being linear.¹⁶

In columns 3 and 4 of Table 1, we report equivalent results for two subsamples. In column 3 (4), the sample consists of the higher (lower) earning individual within each couple. Hours of work vary much less among the higher earning spouses, so that own labor income alone does a much better job of forecasting wage rates. Since there is little unexplained variation in the wage rates among the higher earnings spouses, it is not surprising that the other variables have small coefficients. Only interest income still enters strongly, while the sign of the coefficients on mortgage and property tax payments are now negative. Plausibly, the higher earner in a couple that wants better housing or better schooling for their kids works longer hours in order to pay the required extra property taxes, implying a lower wage rate than otherwise expected given observed labor income.

In the results for lower earning spouses reported in column 4, however, all of the variables help in forecasting wage rates. Property tax payments in particular play a major role. Perhaps the secondary earners most concerned about education also choose low or even zero hours of work, in order to care for the kids. In order to check on the role played by spouses who were out of the labor force, we reestimated this specification on the subsample of lower-earning spouses who had positive hours of work. In these results (available on request), the coefficients are each roughly half as large, suggesting that an important role for these variables is in forecasting the wage rate for those who are out of the labor force.

Column 5 reports equivalent results presuming joint rather than separate filing. Here the dependent variable is the average wage rate of husband and wife, while labor income equals their combined earnings and the other variables equal their joint income from capital and their joint expenses. Again, the key additional indicators of use in forecasting wage rates are interest income and property tax payments.

Column 6 tests whether the labor income of the secondary earner should enter with a different weight when trying to forecast average wage rates of the couple. Given that the lower earning spouse presumably works fewer hours, we should expect a higher coefficient on this spouse's labor income, since the coefficient converting wH to w should be proportional to $1/H$. In contrast, our estimated coefficient for the labor income of the secondary earner is $1 - .41 = .59$, compared to a weight of one on the labor income of the primary earner. A natural explanation for this result is that hours of work are more random among secondary earners, making this indicator for the wage rate much less informative.¹⁷

¹⁶Figure 1 provides a graph of this function for one of the other specifications.

¹⁷For example, if hours satisfy $H = \hat{H}(1 + \tilde{\eta})$ for some *i.i.d.* random noise $\tilde{\eta}$, then the expected coefficient on wH is $1/[H(1 + \sigma^2)]$. When we constrained the sample to couples where both spouses report positive hours, the estimated weight on the labor income of the secondary earner is .83. While still less than one, suggesting that hours of work are more variable for secondary earners than for primary earners even among those with positive hours, this effect is now much less strong.

How does the tax base suggested by these results correspond to current law? The results in column 5 of Table 1 imply that the weight on interest income is 1.7 times the weight on own earnings. Interest income enters more strongly in the tax base than under current law, even though the only reason for its presence is to help in forecasting unobserved ability. The estimated weight on dividend income in this specification is 0.18, suggesting the inclusion of dividend income but at a more favorable rate, as under current law.

The suggested treatment of mortgage and property tax payments differs dramatically, however, from current law. Rather than allowing these payments as a deduction, the results instead suggest imposing additional taxes on households based on these expenditures, particularly property tax payments. Note, though, that property tax payments are not currently deductible for households that do not itemize or for those who are subject to the AMT. The 2005 President’s Advisory Panel on Federal Tax Reform also proposed eliminating the deductibility of property tax payments, and severely restricting the deductibility of mortgage payments, moving the tax code towards the tax base suggested by our initial results.¹⁸

All of these regressions report results forecasting wage rates in a year using various information from that year in use in the tax code. From a lifecycle perspective, however, the individual’s “ability to pay” should depend on average wage rates over time and not just the wage rate in any one year. When wage rates fluctuate over time, for example, current wage rates and current earnings would be more closely linked than they are from a lifecycle perspective, leaving less potential role for other indicators.¹⁹ To test the sensitivity of our results, we rerun the regressions from Table 1 using instead as a dependent variable a measure of the individual’s “permanent” wage rate.²⁰ Results are reported in Table 1a. As expected, the R^2 of the regression falls, since current earnings predict “permanent” wage rates less well than they predict current wage rates, in the process resulting in an increased role for the remaining indicators of ability. In particular, the role of interest income, dividends and property taxes (all of which are potential proxies for permanent income) is strengthened, while the effect on the coefficient on mortgage payments (which are likely to fluctuate significantly over time and are less cleanly related to permanent income) is small.

Our leap, though, from regression results to the desired tax base is very casual. The aim of the next section is to think through this link more carefully.

The estimates for the function $f(\cdot)$ can also be revealing. In Figure 1, we graph our estimates for the forecasted average wage rate for a couple as a function of just the couple’s average labor income. Here, we find that the forecasted average wage rate is roughly proportional to average labor income for those with average earnings above about \$50,000. Reported earnings are much less informative

¹⁸We do not model alternative rationales for these deductions such as potential externalities from house ownership and local government expenditures.

¹⁹As an extreme example, if all individuals have the same age-wage rate profile but are observed at different ages, then current earnings help forecast current wage rates, but have no use in forecasting an individual’s average wage rate over her lifetime.

²⁰Specifically, we used as the dependent variable the forecasted wage rate at age 40, standardized to the earnings levels in year 2001, for that individual estimated based on a regression of $\log(\text{wages})$ against a fixed effect, age dummies, and year dummies (all gender-specific).

about wage rates, however, among couples with lower earnings. If income for those earning below \$50,000 is not very informative about earnings ability, then any redistribution *within* this poorer segment of the population provides limited equity gains while still generating important efficiency costs.

3 Theory

3.1 Initial derivation

When there are multiple indicators that are jointly useful in forecasting earnings ability, how should the tax base be designed? To assess this, we begin with an analysis of a single individual. Consider a one-period economy consisting of H single individuals. Individual h has wage rate w_h , labor income Z_h , a vector of consumption C_h , and some other observable indicators X_h that can potentially enter the tax base.

Tastes for leisure and consumption vary arbitrarily across individuals. The utility function for individual h is denoted $U(C_h, X_h, Z_h; w_h, h)$, where h enters the utility function directly to capture any heterogeneity in individual tastes. Here, labor income serves as a proxy for leisure, i.e. labor (the complement of leisure) equals Z_h/w_h .

Tax payments are denoted by $T(B_h)$, where the tax base B_h is some function of observables. Assume in particular that $B_h \equiv Z_h + X_h\alpha$, to capture the typical structure of the income tax base.²¹

We assume that the marginal utility of income for an individual, denoted by U_{hY} , simply depends on the individual's wage rate at the observed market equilibrium, regardless of the hours of work and the consumption bundle they choose. We assume that marginal utility is a declining function of the wage rate.

To see more explicitly the assumed normalization of marginal utility, suppose that the utility function $U(C, X, Z; w, h)$ is some general cardinal representation of preferences for each individual. At the particular date when we entertain the possibility of a reform, marginal utility of income is given by $g(w, h) \equiv U_Y(C(w, h), X(w, h), Z(w, h); w, h)$, where C , X , and Z denote the individual's choices given the environment at that particular date. On equity grounds, we assume that the marginal utility of income at that date depends just on the wage rate, so can be expressed by some function $U_Y(w)$. Consequently, the normalized marginal utility of each individual equals $U_{hY}(C, X, Z; \tilde{w}) \equiv \frac{U_Y(w)}{g(w, h)} U_Y(C, X, Z; \tilde{w}, h)$.²²

The government's objective function is assumed to equal: $\max \sum_h U_h + W(\sum_h T(B_h))$, where $W(\cdot)$ is a concave function that reflects the aggregate utility that individuals receive from government

²¹Following actual tax law, we assume that the tax function does not vary with the age or other characteristics of the individual. We also assume that the extra indicators X_h enter linearly in the tax schedule. See Golosov, Tsyvinski, and Weinzierl (2009) for an analysis of taxation with heterogeneous discount rates where the additional taxes can be a general function of X_h and Z_h

²²While unneeded in the following discussion, one may assume as well that the level of utility is also the same for individuals who have the same wage rate.

expenditures and/or the utility that officials extract through their control over tax revenue.²³ Given this objective function, the government has an incentive to collect revenue more so from individuals with lower marginal utility of income, so higher wage rates. It does not observe wage rates, though, so must rely on observable information. Any distributional gains depend on the average marginal utility of income among those with a higher vs. a lower value of the available tax base. If the tax base were entirely uninformative about wage rates, then these marginal utilities would be the same at all values of the tax base. The more informative is the tax base about the unobserved wage rates, the larger the difference in the average marginal utilities of income among those with a higher vs. a lower value of the tax base, and so the larger the potential distributional gains from basing taxes on this tax base.

The resulting first-order condition for α (the parameter defining the tax base) is:

$$\sum_h (1 - U_{hY})T'(B_h)X_h = - \sum_h T'(B_h) \left(\frac{\partial Z_h}{\partial \alpha} + \alpha \frac{\partial X_h}{\partial \alpha} \right) \quad (2)$$

Here, the left-hand side measures the equity gain from transferring a given amount of revenue from individuals to the government, where individual 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 utility of private funds to each individual h , U_{hY} . The right-hand side measures the marginal excess burden created by this increase in tax rates arising from any drop in either element in the tax base.

To shed more light on the left-hand side, we group individuals based on their initial tax base. With some abuse of notation, group B contains all N_B individuals with tax base equal to B . Restating this first-order condition,

$$\sum_B N_B T'(B) E[(1 - U_{hY})X_h | B] = - \sum_h T'(B_h) \left(\frac{\partial Z_h}{\partial \alpha} + \alpha \frac{\partial X_h}{\partial \alpha} \right) \quad (3)$$

Consider next an alternative tax change that alters the tax schedule T , holding α fixed, so as to raise just as much extra revenue from those in group B as occurs with the initial reform.²⁴ The resulting increase in tax payments for each individual in group B must then equal $T'(B)E[X_h|B]$. The change in welfare resulting from this alternative tax change equals:

$$\sum_B N_B T'(B) (1 - E[U_{hY}|B]) E[X_h|B] = - \sum_h T'(B_h) \left(\frac{\partial Z_h}{\partial n} + \alpha \frac{\partial X_h}{\partial n} \right), \quad (4)$$

where “ n ” denotes this alternative reform. If the initial tax structure is optimal, then any marginal change from the existing optimal rate schedule must have no marginal effect on welfare.

The question is then how expressions (3) and (4) compare. At the optimal tax base, both expressions equal zero, and the difference then equals zero. Consider first the difference between

²³Without loss of generality, we normalize utilities so that $W' = 1$ at the existing policy.

²⁴A similar approach has been used recently by Laroque (2005) and Kaplow (2006) to show that under separability there is always a Pareto improving reform that eliminates differential taxation of bases other than labor income.

the left-hand sides. This difference (capturing the gains from redistribution) equals:

$$-\sum_B N_B T'(B) E[(X_h - E[X_h|B])U_{hY}|B] = -\sum_B N_B T'(B) \text{cov}(U_{hY}, X_h|B). \quad (5)$$

This is nonzero to the extent that information about X_h helps forecast U_{hY} among those with the same B . Recall that $U_{hY} = U_Y(w_h)$ depends simply on the individual's wage rate. Therefore, this expression implies an equity gain, supporting including X in the tax base to the extent that X helps in discriminating between low and high wage-rate individuals who have the same B . If X_h provides no information about wage rates, given B , then the covariance will be zero.

For example, assume that B represents labor income, while X_h is some form of consumption. If consumption is simply a function of labor income, and 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 be nonzero and there are equity grounds for including X_h in the tax base.

Consider next the difference between the right-hand sides in expression (4), measuring the excess burden from this compensated change in the tax base. Consider for illustration the special case focused on in Atkinson and Stiglitz (1976), where consumption is weakly separable from leisure for each individual. If a particular good comprises $\theta\%$ of an individual's consumption, then a tax at rate $\alpha T'$ on this good has the same marginal effect on individual utility as a tax on labor income at rate $\alpha\theta T'$, regardless of that individual's wage rate. When θ varies across individuals, the equivalent increase in the labor tax rate varies across individuals. If those with high θ on average have the same labor supply elasticity as those with low θ , however, then the commodity tax change has the same impact on aggregate labor income, $\sum_h Z_h$, as a change in the tax rate on labor income of $\alpha E(\theta)T'$.²⁵

The efficiency effects of the combined tax changes arising from changes in X_h can then be expressed by

$$-\sum_h T'(B_h)\alpha \left(\frac{\partial X_h}{\partial \alpha} - E(X_h|B) \frac{\partial X_h}{\partial I} \right) = -\sum_h T'(B_h)\alpha \frac{\partial X_h}{\partial \alpha} |_c. \quad (6)$$

Here, the expression in parentheses corresponds to the sum of the compensated price effects on X_h , as measured by the second expression in this equation, if the income effects $\partial X_h/\partial I$ are not correlated with X_h , given B , again corresponding to Assumption 2 in Saez (2002). Under these assumptions, therefore, the marginal excess burden equals zero when $\alpha = 0$, increases as α differs from zero, and more so the more responsive is X_h to its price. The same argument applies to the efficiency terms involving Z_h .

Our paper does not attempt to estimate these efficiency effects. Instead, we focus on the size of the redistributive effects of tax reforms, as measured by expression (5). When these are nonzero, the optimal value of α is also nonzero and would be set to balance distributional gains and efficiency losses.

²⁵This assumption corresponds to Assumption 2 in Saez (2002).

Figure 2 provides a possible graph of both the redistributive gains and the efficiency losses generated by changes in α . Here, $\hat{\alpha}$ represents the value of α that maximizes the redistributive gains, so where expression (5) equals zero (this is the value that we estimate in the paper). Under the separability assumptions in Atkinson and Stiglitz (1976), efficiency is maximized at $\alpha = 0$. In general, the optimal value of α that accounts for both efficiency and equity considerations is α^* , a value between zero and $\hat{\alpha}$.²⁶

3.2 Derivation for Married Couples Filing Jointly

Several further issues arise when considering the tax base for a tax on married couples, assuming joint filing.²⁷ To begin with, only the couple's joint income from interest, 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.²⁸

An important new question is how any changes in tax liabilities are shared between the two spouses. For simplicity, we assume that any tax change is equally divided between the two spouses.²⁹ Decreasing the taxes on any married individual by a dollar then increases the utility of the couple by $.5(U_{pY} + U_{s(p)Y}) \equiv \bar{U}_{pY}$, where $s(p)$ denote the spouse of any given primary earner p and where \bar{U}_{pY} is defined to equal the average marginal utility of income in the couple.

With joint taxation, the earnings of the secondary earner can potentially be treated differently from that of the primary earner. If so, assume that the tax base for the couple under joint filing equals $B_p = Z_p + \gamma Z_{s(p)} + \alpha X_p$. For any $\gamma \neq 1$, the marginal tax rate varies between high and low earners within a couple.

The equivalent to expression (5) used to characterize the optimal choice of α , now summed over couples (primary earners) rather than over individuals becomes:

$$-\sum_B N_B T'(B) E[\bar{U}_{pY}(X_p - E(X_p|B))|B] \quad (7)$$

If earnings of the secondary earner are treated differently, then we have an additional expression for the optimal γ :

$$\sum_B N_B T'(B) E[\bar{U}_{hp}(Z_{s(p)} - E(Z_{s(p)}))|B] \quad (8)$$

Under the tax base that best approximates an ability tax, one or both of these expressions would equal zero, depending on whether secondary earners can be taxed differently.

²⁶With multiple dimensions of α , we instead solve for the vector of marginal distributional benefits for changes in the vector α . The marginal excess burden term then takes the form $-\Omega\alpha$ for some matrix Ω . In the special case in which the matrix Ω is diagonal, the marginal distributional gain at the optimum must be larger when the marginal excess burden is larger.

²⁷In order to decompose this analysis from that on single individuals, we hold fixed the tax base for single individuals when considering possible changes to the tax base for married couples.

²⁸In community property states in the U.S., such changes in the nominal owner have no legal implications for divorce settlements, though they can in other states.

²⁹See the Appendix for a derivation justifying this assumption when the couple engages in Nash bargaining, where each has a fall-back position equal to the utility he or she would have as a single individual, and where their bargaining power is equal.

3.3 Married Couples: Joint vs. Separate Filing

A number of countries tax married individuals separately. The U.S. provides this as an option as well, though relatively few married individuals make use of it. Can a case be made on equity grounds for joint vs. separate filings?

We derived the optimal tax base on equity grounds under joint filing in the previous section. What would the optimal tax base be instead, assuming separate filing? Here we assume the same tax schedule for both primary and secondary earners, so need to assess the welfare effects on both spouses from modifying the tax base.³⁰ Again we assume that the combined tax change is shared equally between the two spouses.

With separate filing, there is also a policy choice about how to deal with joint assets and liabilities. One approach is to base the law on legal ownership and liability for each asset, giving the couple substantial flexibility to reassign ownership to the spouse facing the lower marginal tax rate, in order to minimize taxes. The approach we presume instead is that the tax law assigns half of their joint assets and liabilities to each spouse. To facilitate the comparison with other specifications, we always include joint assets and liabilities as regressors (including a half would simply scale the coefficient).

Consider then the effects of a marginal change in α on a couple, taking into account the effects on the tax liabilities of both spouses. The resulting term equivalent to that on the left-hand side of equation (1) is:

$$\sum_p (1 - \bar{U}_{pY}) [T'(B_p)X_p + T'(B_{s(p)})X_p] = \sum_h (1 - \bar{U}_{hY}) T'(B_h)X_h \quad (9)$$

Here, the first expression is summed over primary earners, so over p , while the second term is summed over individuals, h . Note that this second expression is identical to the term on the left-hand side of equation 1, except that marginal utility is for the couple rather than for the individual. With this reinterpretation, we can then follow the previous derivation, implying that the optimal tax base on distributional grounds satisfies

$$- \sum_B N_B T'(B) E [\bar{U}_{hY}(X_h - E[X_h|B])|B] = 0 \quad (10)$$

The question then is how to compare the welfare resulting from use of separate vs. joint filing. Consider the welfare effect of a shift from joint filing to separate filing, with a compensating transfer designed to leave the expected taxes paid by those with any given joint-filing tax base unchanged. The resulting impact on welfare equals

$$\sum \bar{U}_{pY} (T^S(B_p) + T^S(B_{s(p)}) - E(\sum_s T^S|B_p^j)) \quad (11)$$

Here, T^S denotes the tax schedule for an individual assuming separate filing, B_p ($B_{s(p)}$) is the optimal tax base estimated above under separate filing for the primary (secondary) earner, B_p^j is

³⁰We again hold fixed the tax base for single individuals when considering possible changes to the tax base for married couples filing separately.

the tax base for the couple under joint filing, while $E(\sum_s T^S | B_p^j)$ measures the expected tax payments by the couple under separate filing, conditional on B_p^j .

3.4 Lifecycle 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 is then measured based on their wage rate in that year. These derivations ignore, though, that individuals can smooth consumption through borrowing and lending, so that their marginal utility of income should depend on their earnings ability throughout their life. In addition, they are affected by a change in the tax law throughout their life, and not just in one year.

How would the prior results change if we take into account these lifecycle considerations? The first question is how to characterize the marginal utility of an individual, given the set of wage rates they face at all different ages: $U_Y(w_1, \dots, w_T)$. In general, individuals 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 earnings 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: $\sum_t w_t(1+r)^{-t} \equiv w^*$. 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 using earnings levels in 2001 derived from a wage regression that includes fixed effects and a full set of year and age-specific dummies, estimated separately by gender.³¹

If the value of α were to change, then the cumulative welfare impact of this change over the life of the individual equals: $\sum_h(1 - U_{Ch}(w^*))\{\sum_t T'(B_{ht})X_{ht}(1+r)^{-t}\} = \sum_t(1+r)^{-t}[\sum_h(1 - U_{Ch}(w^*))T'(B_{ht})X_{ht}]$. As long as the composition of the population is stable over time, we can evaluate the expression in brackets at any given date in order to evaluate the equity implications of a tax reform. The only change from our prior expressions is then in the use of a lifetime rather than an annual wage rate for each individual.

³¹There 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 impose risk-bearing costs. By ignoring the variation in wage rates, we implicitly presume that these two effects are on average offsetting. We also ignore possible liquidity constraints, which when binding imply that the marginal utility of consumption depends on current but not future wage rates. Our static results are instead appropriate if individuals are systematically liquidity constrained.

4 Estimation given theory

Our initial empirical work provided suggestive evidence about the possible value of various observations about individuals in addition to their labor income in the design of the tax base. The value of these other indicators in forecasting wage rates can easily vary though across individuals. Based on the theory (as summarized for example by expression (7)), we find that any additional ability to forecast wage rates is more important for those individuals with a higher marginal utility of income facing a higher marginal tax rate.

Our aim in this section is first to test whether the optimal value of α differs from current law, and then to approximate the value of α that is most attractive on equity grounds, given the theory. We do this first in a static setting, and then in a lifecycle context. We evaluate the gains from use of additional indicators assuming either that married couples file jointly or file separately. Given these results, we can then compare the equity gains from joint vs. separate filing.

According to the theory, the optimal tax base for a married couple filing jointly should satisfy equation 7.³² This equation can be reexpressed as follows:

$$\begin{aligned}
 & - \sum_B N_B T'(B) E[\bar{U}_{pY}(X_p - E(X_p|B))|B] = \\
 & - \sum_B N_B T'(B) E(X_p \bar{U}_{pY}|B) + \sum_B N_B T'(B) E(X_p|B) E(\bar{U}_{pY}|B) = \tag{12} \\
 & - \sum_p T'(B_p) X_p \bar{U}_{pY} + \sum_p T'(B_p) X_p E(\bar{U}_{pY}|B)
 \end{aligned}$$

The resulting first term can immediately be calculated, once we have functional forms for $T(\cdot)$ and \bar{U}_Y . Evaluating the second term requires evaluating $E(\bar{U}_{hY}|B)$.³³ Conditional on this expression, the second term can also immediately be calculated.

Deriving the equivalent expression for the optimal tax base for married couples filing separately, the only change in the expression is the replacement of the tax rates T' for married couples filing jointly with the tax rates for these individuals filing separately.

The key next steps in estimation are choosing functional forms for $T'(\cdot)$ and \bar{U}_Y . For $T'(\cdot)$ we simply use the observed tax rate schedule for couples filing jointly in 1995. For married individuals filing separately, we used the tax schedule for single individuals in 1995.³⁴

Our normative assumption is that $U_{hY} = u(w_h)$. In our base case, we assume that $u(w_h) = 1/\max(w_h, \$5)$, consistent with a utility function equaling $w_h^{1-\gamma}/(1-\gamma)$ with $\gamma = 1$, 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.

³²For now, we ignore the possibility of taxing secondary earnings differently.

³³Note that this expression does not vary across the first-order conditions for each element in α , which provides an important computational advantage.

³⁴Unfortunately, the NBER Taxsim program did not include a procedure for married couples filing separately, so we approximated their tax schedule with that for single individuals.

Estimation is done on the same sample from the PSID used in section 2, and implemented in R, relying on the *np* package from Hayfield and Racine (2008) for implementation of semiparametric and nonparametric components.

We construct an empirical estimate of $E[\bar{U}_{hY}|B]$ by a nonparametric (kernel) regression of \bar{U}_{hY} on $Z + \alpha X$. We select the bandwidth by least-squares cross validation.³⁵ 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 considered modifications to this tax base, changing the weight on each of the available indicators, calculated based on equation (7), and perhaps equation (8). The existing tax base in the U.S. in that year includes the sum of joint earnings plus dividend and interest income, minus mortgage payments and property tax payments among those who itemize.³⁶ 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.³⁷

Results are reported in columns (1) and (2) of Table 2. 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. However, it would maintain interest income in the tax base, 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.

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 $1 + 1.35 = 2.35$). This higher weight compensates for the lower hours worked by the secondary earner. Ignoring any randomness in labor income, given wage rates, this would be consistent with the weighted average hours worked by the secondary earner (with those with a low wage rate being

³⁵Because 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.

³⁶We constructed a one-dimensional proxy for the tax base by mapping the 1995 tax liability that accounts for available information (including state information, children, AMT etc., as described in the appendix) into the corresponding AGI level for a married couple with no kids (and using the actual AGI in cases when tax liability was zero).

³⁷Our results are only approximations since we do not control for possible general equilibrium changes in prices or wage rates due to a change in the tax law and do not control for changes in the marginal utility of income for each individual caused by changes in their tax liabilities. These figures also do not take into account any efficiency costs/benefits arising from a change in the tax base.

weighted heavily) being only 1/2.35 times the average hours worked by the primary earner.³⁸ When spousal earnings receive more weight, the other indicators also become more important. Each help in forecasting hours of work for the secondary earner, controlling for spousal earnings. When these earnings receive more weight, the corrections need to receive more weight as well.

Note that the qualitative results are broadly the same as we found in Table 1, with interest income and property tax payments being particularly useful in forecasting ability, after controlling for earnings. This is true even though those with low wage rates facing high marginal tax rates receive much more weight, and even though our starting point is the existing tax base rather than joint earnings.

Does the existing tax base even do as well as a tax base that simply equals the couples joint labor income? To test this, we use an analogue of equation (11) to 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, holding expected taxes unchanged at each value of the actual tax base.³⁹ In order to convert this utility gain to a dollar figure, we divide through by the sum of the welfare weights. The resulting figure then is a weighted average of the change in tax liabilities, weighting by the marginal utility of income. Here we find a weighted average gain of \$173: a tax base limited to joint earnings does slightly better on equity grounds than does the actual tax base.

We next tested to see to what extent further equity gains are feasible, starting from a tax base equal to joint earnings, by making use of the available indicators. Here, we are implicitly testing the separability assumption in Atkinson and Stiglitz (1976) that consumption decisions convey no information about ability other than what is already contained in reported earnings. Statistically significant coefficients for particular variables reject separability for those variables, indicating that such variables contain usable information about ability. The specific coefficient estimates again provide an approximation to the optimal tax base.

Estimation results for married couples filing jointly are reported in columns (3) and (4) of Table 2. Results are broadly similar to those implied by the coefficients in columns (1) and (2). 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.

When we allow for a separate tax rate on the labor income of primary vs. secondary earners in column (4), we now find that the weight on the labor income of the secondary earner should be 2.5 times that on the primary earner. The other coefficients again increase, roughly in proportion to

³⁸ 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.

³⁹ In calculating the alternative tax payments, we use the existing tax schedule in 1995 and the couple's joint earnings.

the change in the weight on the income of the secondary earner.

Estimation results for the optimal tax base for married couples filing separately are reported in column (5) of Table 2. We now find yet more weight given to the supplementary sources of information. These results apparently are driven by the extra weight being put on approximating well the tax base for secondary earners with low earnings, but also low hours of work, who are not as poor as they would appear based on reported labor income.

In column 6 of Table 2, we consider a somewhat counterintuitive tax system, with separate filing but the possibility of making taxable income depend on the spouse's income. Interestingly, this specification finds a very significant role for spousal income and eliminates the importance of mortgages and property taxes. In particular, in forecasting the wage rate of the secondary earner, we find that the head's income does a better job than mortgage and property tax payments in forecasting hours of work (as can be seen by comparing R^2 in columns 3 and 4).

Table 3 provides some robustness checks. In columns 1-2 of Table 3, we reestimate the results in columns 3-4 of Table 2, but now assume that the minimum effective wage rate is \$3, rather than \$5. Remarkably little changes. In columns 3-4, we assume a less concave utility function ($\gamma = .5$), while in columns 5-6 we assume a more concave utility function ($\gamma = 2$). The reestimates of column 3 of Table 2 result in very small changes, so this specification is very robust. The reestimates of column 4 of Table 2 prove to be much more sensitive. The more weight that is put on individuals with lower wage rates, the higher the estimated weight on spousal income, and on most of the other indicator variables. This pattern is entirely consistent with our prior explanation that lower-wage-rate individuals have lower hours than do those with higher wage rates, and these lower wage rate individuals receive more weight in the reestimation results in Table 3.

Does separate filing do a better job than joint filing in approximating an ability tax? To conduct this test, we make use of the same procedure used in comparing a tax on joint earnings with a tax based on the current tax base. When we compare the joint filing specification characterized by the results in column 1 of Table 3 to the separate filing specification in column 3, there is a gain from separate filing corresponding to just \$51.65. When we compare the joint filing in specification 2 to the separate filing in specification 4, the gain changes to \$113.75: In both cases the difference is small suggesting that both separate and joint filing do equally well on equity grounds.

We next examine the sensitivity of our results to lifecycle considerations. Even with the PSID, we observe individuals over only part of their lifecycle, with the observed age range varying by individual. To proceed, we forecasted each individual's wage rate at age 40, based on the earnings levels in year 1995, using a regression of $\log(\text{wages})$ against an individual fixed effect, age dummies, and year dummies, with the regression estimated separately for each gender. We then used this estimated wage rate instead of the individual's wage rate in any given year in the various formulas used to calculate the optimal tax base. Results for joint filing, comparable to those in columns (3) and (4) of Table 2, are reported in Table 4. Here, we find that more (and in some cases much more) weight is given to each of the indicators than was found in Table 2. If annual earnings fluctuate due to random variation in wage rates, 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 have seen in the data, we re-estimate the lifetime specification allowing 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.⁴⁰ The resulting estimates for the tax base, using the permanent component for each of the variables, are reported in columns 3 and 4 of Table 4. In column 3, income of the lower earner is given equal weight, while it can have a different weight in column 4. As expected, comparing the estimates in columns 3-4 to those in columns 1-2, more of the variation in lifetime wage rates can be explained by lifetime earnings than by annual earnings, leaving less room for the various indicators to help in forecasting lifetime wage rates. The coefficient on property tax payments in particular falls substantially, suggesting that this indicator is large particularly in years when labor supply is low.⁴¹ In contrast, dividend receipts now receive more weight, suggesting that over the course of a lifetime those with high dividends work fewer hours. The smaller coefficient in the annual data could arise from random variation in dividend receipts, or from a tendency to own more stock at ages with high labor supply. Interest income in contrast tends to be high in years with low hours of work, but has little association with hours over the course of a lifetime.

Note also that the overall quality of the fit in columns 3-4 is much higher than in columns 1-2 of Table 4, implying important equity gains from lifetime averaging.

5 Equity vs. Efficiency Gains

In our empirical work, we focused on estimating $\hat{\alpha}$, 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.

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:

a) 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

⁴⁰Since 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.

⁴¹This could occur, for example, if school-aged kids induce a couple to move to a community with good schools (and high property taxes), but also to spend more time away from work caring for the kids.

earners and primary earners can be interpreted as providing a rough balancing between equity and efficiency considerations.

b) 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 savings, focusing solely on efficiency considerations. The recent literature suggests a positive tax on savings on efficiency grounds.⁴² 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 nontrivial role in the tax base. The current deductibility of mortgage interest cannot be justified on equity 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. These equity costs raise questions, though, about the current tax treatment.⁴³

We have shown in the paper how to estimate the implications of equity considerations on the choice of the tax base. Rigorously incorporating efficiency considerations is a natural extension.

⁴²One strand of this literature, reviewed in Golosov et al. (2006), shows that a positive tax rate on capital income makes it easier to redistribute between high ability and low ability individuals, by making it less attractive for high ability individuals to save to finance consumption during periods with low future earnings. 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.

⁴³Perhaps the current deductibility of property tax payments serves as another mechanism to aid children, by inducing parents to spend more on their education.

A 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, 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, they face a joint budget constraint equal to

$$C_{HM} + C_{WM} = w_H(1 - L_{HM}) + w_W(1 - L_{WM}) - T_M \quad (13)$$

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

$$\left(U^H(C_{HM}, L_{HM}) - U^H(C_{HS}, L_{HS}) \right) \left(U^W(C_{WM}, L_{WM}) - U^W(C_{WS}, L_{WS}) \right) \quad (14)$$

where we implicitly assume here that the two spouses have equal bargaining power. To simplify the resulting algebra, we take a first-order approximation to each change in utility, e.g. $U^{HM} - U^{HS} = U_C^{HM} \Delta C_H + U_L^{HM} \Delta L_H$.

The resulting first-order conditions can be combined to show that labor supply is allocated efficiently between the two spouses:

$$\frac{U_L^{HM}}{U_C^{HM}} = \frac{U_L^{WM}}{U_C^{WM}} \frac{w_H^n}{w_W^n} \quad (15)$$

Here, w^n denotes the net-of-tax wage rate on marginal labor supply. Unless, total resources change, there is no impact of marriage then on labor supplies.

Given this result, the first-order conditions for the consumption of each spouse can be expressed by

$$\Delta C_W = \Delta C_H + \left(\Delta L_H \frac{w_H^n}{w_W^n} - \Delta L_W \right) \frac{U_L^{WM}}{U_C^{WM}} \quad (16)$$

If each couple spends the same fraction of any extra resources they receive on consumption vs. leisure, then we infer that $\Delta C_W = \Delta C_H$.

B 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 and trust income (disaggregated when

possible), mortgage payments, property taxes, marginal tax rates, total federal and state tax liability. 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. We often weight observations by either the average family labor income or the average own labor income depending on whether family-level or individual observations are used. Unless otherwise indicated, we adjust dollar variables for changes in the price level using the CPI-U index, using 2001 as the baseline.

Labor income. We construct full labor income by adding up wages and salaries, bonuses, overtime payments, tips, commission income, professional practice income, miscellaneous labor income and labor part of business/farm income. The last of these variables is imputed in 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. This is one reason why we usually exclude individuals owning a business.

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 wife and we do so.

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.

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 tax if the latter is not greater than mortgage payments.

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.

Taxes. We compute tax liability by relying on the TAXSIM calculator available at <http://www.nber.org/taxsim>. Our calculation accounts for: labor income of both spouses, unemployment insurance, dividends and interest income (separately starting with 1993, jointly and treated as interest income prior to 1993). We account for the number of children for the purpose of computing exemptions/child tax credit. We assign filing status of married filing jointly (to those married), single (to those single without children) or head of households (single with children). Because our analysis is performed for those between ages of 25 and 60, age exemptions are not an issue. We use TAXSIM to compute federal and state tax liability and marginal tax rates, while accounting for the information listed above. We use federal marginal tax rates unless otherwise indicated (although state taxes do affect federal taxes through their effect on itemized deductions).

We use marginal tax rates for weighting in the optimal tax base formulae. The 1995 tax rates are obtained by first indexing all the dollar variables to 1995 values using the Social Security Administration wage index and then applying the 1995 law. For specifications with separate filing, we use the hypothetical tax rate that would arise if the individual were single. We allocate half of the children to each parent, with the odd child going to the mother. Own unemployment income, if any, is used. Half of capital income is allocated to each spouse. We multiply property and mortgage interest payments by 2/3 to rudimentarily account for returns to scale in these kinds of payments occurring within a family. Finally, we assign the marital status of single or head of household, depending on the presence of children, and apply the current law.

References

- Akerlof, George A.**, “The Economics of “ Tagging” as Applied to Optimal Income Tax, Welfare Programs, and Manpower Planning,” *American Economic Review*, March 1978, *68* (1), 8–19.
- Alesina, Alberto, Andrea Ichino, and Loukas Karabarbounis**, “Gender based taxation and the division of family chores,” Working Paper 13638, National Bureau of Economic Research November 2007.
- Atkinson, Anthony B. and Joseph E. Stiglitz**, “The Design of Tax Structure: Direct versus Indirect Taxation,” *Journal of Public Economics*, July-August 1976, *6* (1-2), 55–75.
- Banks, James and Peter A. Diamond**, “The Base for Direct Taxation,” in James A. Mirrlees, Stuart Adam, Timothy Besley, Richard Blundell, Steve Bond, Robert Chote, Malcolm Gammie, Paul Johnson, Gareth D. Myles, and James Poterba, eds., *Dimensions of Tax Design: the Mirrlees Review*, Oxford University Press, 2010, chapter 6.
- Besley, Timothy and Stephen Coate**, “Workfare versus Welfare Incentive Arguments for Work Requirements in Poverty-Alleviation Programs,” *American Economic Review*, March 1992, *82* (1), 249–61.
- Blomquist, Soren and Vidar Christiansen**, “The Role of Prices for Excludable Public Goods,” *International Tax and Public Finance*, January 2005, *12* (1), 61–79.
- Browning, Martin and Costas Meghir**, “The Effects of Male and Female Labor Supply on Commodity Demands,” *Econometrica*, July 1991, *59* (4), 925–51.
- Conesa, Juan Carlos, Sagiri Kitao, and Dirk Krueger**, “Taxing Capital? Not a Bad Idea After All!,” *American Economic Review*, March 2009, *99* (1), 25–48.
- Diamond, Peter and Johannes Spinnewijn**, “Capital Income Taxes with Heterogeneous Discount Rates,” January 2010. MIT and LSE, mimeo.
- Golosov, Mikhail, Aleh Tsyvinski, and Ivan Werning**, “New Dynamic Public Finance: A User’s Guide,” *NBER Macroeconomics Annual*, 2006, 317–388.
- , – , and **Matthew Weinzierl**, “Preference Heterogeneity and Optimal Commodity Taxation,” 2010. Yale and Harvard, mimeo.
- Gordon, Roger H.**, “Taxation of Interest Income,” *International Tax and Public Finance*, January 2004, *11* (1), 5–15.
- Gruber, Jonathan and Emmanuel Saez**, “The Elasticity of Taxable Income: Evidence and Implications,” *Journal of Public Economics*, April 2002, *84* (1), 1–32.
- Härdle, Wolfgang, Peter Hall, and Hidehiko Ichimura**, “Optimal Smoothing in Single-Index Models,” *The Annals of Statistics*, 1993, *21* (1), 157–78.

- Hayfield, Tristen and Jeffrey S. Racine**, “Nonparametric Econometrics: The np Package,” *Journal of Statistical Software*, 2008, *27* (5).
- Kaplow, Louis**, “On the Undesirability of Commodity Taxation Even When Income Taxation is Not Optimal,” *Journal of Public Economics*, August 2006, *90* (6-7), 1235–50.
- Kopczuk, Wojciech**, “Redistribution when Avoidance Behavior is Heterogeneous,” *Journal of Public Economics*, July 2001, *81* (1), 51–71.
- Laroque, Guy R.**, “Indirect Taxation is Superfluous under Separability and Taste Homogeneity: A Simple Proof,” *Economics Letters*, April 2005, *87* (1), 141–44.
- Mankiw, N. Gregory and Matthew Weinzierl**, “The Optimal Taxation of Height: A Case Study of Utilitarian Income Redistribution,” *American Economic Journal: Economic Policy*, February 2010, *2* (1), 155–76.
- Mirrlees, James A.**, “An Exploration in the Theory of Optimum Income Taxation,” *Review of Economic Studies*, April 1971, *38* (114), 175–208.
- Pirtillä, Jukka and Ilpo Suoniemi**, “Public Provision, Commodity Demand and Hours of Work: An Empirical Analysis,” Working Paper 3000, CESifo March 2010.
- Saez, Emmanuel**, “Using Elasticities to Derive Optimal Income Tax Rates,” *Review of Economic Studies*, January 2001, *68* (1), 205–29.
- , “The Desirability of Commodity Taxation under Non-Linear Income Taxation and Heterogeneous Tastes,” *Journal of Public Economics*, 2002, *83* (2), 217–320.
- Vickrey, William**, “Measuring Marginal Utility by Reactions to Risk,” *Econometrica*, October 1945, *13* (4), 319–333.
- Weinzierl, Matthew**, “Incorporating Preference Heterogeneity into Optimal Tax Models: De Gustibus non est Taxandum,” May 2009. Harvard University, mimeo.
- , “The Surprising Power of Age-Dependent Taxes,” 2010. Harvard University, mimeo.

Table 1: Predicting wages

Variables	Married Individuals		Married	Married	Married Couples	
	(1)	Linear (2)	Head (3)	Spouse (4)	(5)	(6)
Spousal income						-0.408 (0.018)
Dividends	0.359 (0.032)	0.378 (0.042)	0.052 (0.011)	0.423 (0.023)	0.176 (0.053)	0.483 (0.059)
Interest	1.173 (0.039)	1.493 (0.080)	0.559 (0.018)	1.893 (0.094)	1.700 (0.070)	1.702 (0.067)
Mortgage	0.418 (0.022)	0.432 (0.026)	-0.034 (0.011)	0.600 (0.045)	0.223 (0.041)	0.019 (0.037)
Property Tax	2.522 (0.058)	3.056 (0.105)	-0.191 (0.045)	5.126 (0.155)	1.381 (0.043)	1.961 (0.137)
Bandwidth	2344.369		721.082	3290.666	2700.804	3455.407
R^2	0.650	0.629	0.820	0.341	0.747	0.763
N	22736	22736	11351	11379	11396	11390

Table 1a: Predicting permanent wages

Variables	Married Individuals		Married	Married	Married Couples	
	(1)	Linear (2)	Head (3)	Spouse (4)	(5)	(6)
Spousal income						-0.257 (0.021)
Dividends	0.751 (0.018)	0.825 (0.094)	0.601 (0.069)	1.251 (0.112)	0.851 (0.107)	1.357 (0.033)
Interest	3.188 (0.038)	4.270 (0.220)	3.226 (0.147)	7.579 (0.379)	7.597 (0.212)	4.573 (0.032)
Mortgage	0.031 (0.013)	0.187 (0.042)	-0.153 (0.046)	0.362 (0.084)	0.116 (0.076)	0.374 (0.047)
Property Tax	4.053 (0.033)	6.262 (0.208)	2.742 (0.199)	9.770 (0.435)	6.419 (0.318)	8.350 (0.076)
Bandwidth	1579.078		3676.866	6220.167	7333.920	4655.394
R^2	0.508	0.501	0.625	0.254	0.552	0.567
N	22726	22726	11331	11368	11381	11391

Table 2: Optimal tax base

Variables	Using tax base		Joint Filing		Separate filing	
	(1)	(2)	(3)	(4)	(5)	(6)
Spousal Income		1.353 (0.167)		1.513 (0.191)		3.492 (0.041)
Dividends	-0.739 (0.139)	-0.575 (0.240)	0.049 (0.093)	0.224 (0.212)	0.305 (0.068)	0.419 (0.028)
Interest	0.149 (0.387)	0.787 (0.196)	1.055 (0.071)	1.884 (0.344)	3.415 (0.312)	1.561 (0.244)
Mortgage	0.883 (0.114)	1.394 (0.228)	0.184 (0.106)	0.707 (0.217)	2.204 (0.414)	-0.185 (0.085)
Property Tax	1.883 (0.399)	3.939 (0.883)	1.451 (0.361)	3.826 (1.058)	9.223 (1.205)	0.358 (0.404)
Bandwidth	3926.355	5138.061	2398.916	6270.476	2153.311	550.600
R^2	0.581	0.575	0.588	0.580	0.384	0.430
OLS R^2	0.440	0.484	0.453	0.477	0.291	0.404
N	11357	11369	11355	11368	22681	22747

Table 3: Sensitivity of the optimal tax base

Variables	Minimum wage = \$3		$\gamma = 0.5$		$\gamma = 2$	
	(1)	(2)	(3)	(4)	(5)	(6)
Spousal income		1.521 (0.247)		0.873 (0.105)		2.700 (0.538)
Dividends	-0.028 (0.164)	0.115 (0.228)	0.110 (0.154)	0.243 (0.193)	-0.079 (0.130)	0.153 (0.398)
Interest	1.112 (0.356)	1.413 (0.343)	0.913 (0.187)	1.435 (0.247)	1.559 (0.278)	2.277 (0.440)
Mortgage	0.111 (0.150)	0.328 (0.252)	0.276 (0.106)	0.526 (0.147)	0.191 (0.154)	0.927 (0.503)
Property Tax	1.112 (0.683)	2.755 (1.456)	1.565 (0.342)	2.949 (0.776)	1.206 (0.713)	4.502 (2.003)
Bandwidth	4012.444	3099.947	4076.863	5899.339	4055.656	6847.871
R^2	0.485	0.492	0.656	0.647	0.460	0.468
OLS R^2	0.358	0.391	0.548	0.567	0.297	0.338
N	11355	11368	11363	11364	11355	11397

Table 4: Optimal tax base

Variables	Fixed effect wage		All fixed effects	
	(1)	(2)	(3)	(4)
Spousal Income		3.343 (1.646)		1.376 (0.384)
Dividends	1.797 (0.114)	4.018 (1.780)	4.314 (0.144)	6.657 (1.239)
Interest	5.673 (1.883)	12.398 (11.935)	0.204 (0.647)	2.892 (3.062)
Mortgage	0.354 (0.270)	1.878 (1.655)	0.386 (0.264)	0.678 (0.418)
Property Tax	9.329 (1.746)	24.196 (12.506)	1.275 (0.837)	2.283 (1.081)
Bandwidth	5778.778	10467.022	5408.576	5911.344
R^2	0.434	0.419	0.666	0.688
OLS R^2	0.328	0.362	0.532	0.564
N	11340	11330	2668	2666

Figure 1: Prediction of the average wage rate of a couple based on joint labor income




Figure 2: Marginal effects of α : Distributional gains vs efficiency losses

