WORKING
PAPER

Income Taxes, Compensating Differentials, and Occupational Choice

How Taxes Distort the Wage-Amenity Decision

DAVID POWELL AND HUI SHAN

WR-705-1

December 2010

This paper series made possible by the NIA funded RAND Center for the Study of Aging (P30AG012815) and the NICHD funded RAND Population Research Center (R24HD050906).

RAND LABOR AND POPULATION

Income Taxes, Compensating Differentials, and Occupational Choice: How Taxes Distort the Wage-Amenity Decision

David Powell*
Hui Shan†

December, 2010

Abstract

The link between taxes and occupational choices is central for understanding the welfare impacts of income taxes. Just as taxes distort the labor-leisure decision, they may also distort the wage-amenity decision. Yet, there have been few studies on the full response along this margin. When tax rates increase, workers favor jobs with lower wages and more amenities. We introduce a two-step methodology which uses compensating differentials to characterize the tax elasticity of occupational choice. We estimate a significant compensated elasticity of 0.03, implying that a 10% increase in the net-of-tax rate causes workers to change to a 0.3% higher wage job.

Keywords: Income Taxes, Occupational Choice, Compensating Differentials
JEL classification: H24, H31, J24

*RAND, dpowell@rand.org
†Federal Reserve Board of Governors, hui.shan@frb.gov. We are grateful to Gerald Auten, David Autor, Neil Bhutta, Tonja Bowen Bishop, Tom Garrett, Jon Gruber, Jerry Hausman, Amanda Pallais, Jim Poterba, Nirupama Rao, Claudia Sahm, and Chris Smith for their comments, advice, and support. We thank Dan Feenberg and Inna Shapiro for their help with NBER's TAXSIM program and Donald Bruce for graciously sharing his PSID code. This research was supported by the National Institute on Aging, Grant Number P01-AG05842. The findings and conclusions expressed are solely those of the authors and do not represent the views of the Federal Reserve System or the National Institute on Aging.

1 Introduction

It is well-known that taxes can distort individual labor supply decisions. Following the dramatic declines in marginal income tax rates in the 1980s, a sizeable literature exploited the tax schedule changes to study the price elasticity of labor supply. Numerous papers have estimated the impact of income taxes on the number of hours worked and labor force participation decisions. Economists have also noted that income taxes may affect many other aspects of labor supply decisions. For example, Feldstein (1997) highlights the importance of understanding other facets of labor supply:

> The relevant distortion to labor supply is not only the effect of tax rates on participation rates and hours but also their effect on education, occupational choice, effort, location, and all of the other aspects of behavior that affect the short-run and long-run productivity and income of the individual. Unfortunately, we still know very little about how taxes affect labor supply defined in this broad way. (page 209)

More than a decade later, researchers have made limited progress on understanding the effect of income taxes on labor supply beyond working hours and labor force participation decisions.

In this paper, we study the effect of income taxes on occupational choices using the 1981-1997 Panel Study of Income Dynamics (PSID). In a simplified equilibrium model, every job entails two types of compensation to workers: the pecuniary compensation that is subject to income taxes (which we call wages) and the non-pecuniary compensation that is non-taxable (which we call amenities). These amenities are very broadly defined, including qualities such as difficulty of the job, convenience of the hours, collegiality of the working environment, provision of generous health insurance, on-the-job safety, etc. When the marginal tax rate (τ) increases, the net-of-tax rate ($1-\tau$) is lowered, meaning that the worker keeps less of each dollar that he earns. As a result, a high tax rate diminishes the benefit to high wage occupations and induces the worker to choose a job with more amenities.

Compared to hours worked or labor force participation, the occupational-choice facet of labor supply is difficult to study empirically for a number of reasons. While working hours and participation status are observable to econometricians, job-specific amenities often vary along multiple dimensions and many of them are not observed in the data (e.g., difficulty of the job, convenience of the hours, and collegiality of the working environment). Furthermore, classification of occupations in survey data contains substantial measurement error. When using an indicator for job changes as the dependent variable in regression models, occupation misclassification cannot be considered classical measurement error and may bias estimates.

We introduce a two-step estimation procedure to study the relationship between income taxes and occupational choices. This procedure recognizes that compensating differentials can be used to measure all disamenities associated with a job. Assuming a worker receives a fixed amount of total compensation in a competitive labor market, a job with a lot of disamenities must pay higher compensating differentials to attract workers. Therefore, the choice over jobs with different amenities is equivalent to the choice over jobs with different compensating differentials. In the first step of the two-step procedure, we estimate the compensating differential for each occupation in each year which will be used as the dependent variable in the second step. Using compensating differentials to study occupational choices has two advantages. First, it allows us to measure total amenities without having to observe each specific type of amenities that a worker receives. Second, because the estimated compensating differentials are continuous variables, we can model occupation misclassifications in the data as classical measurement error. Classical measurement error in the dependent variable does not bias estimates in linear models.

In the second step of the two-step procedure, we relate the difference in the compensating differentials of the new occupation and the old occupation to the change in a worker's marginal net-of-tax rate $(1 - \tau)$. Because a worker's marginal net-of-tax rate is endogenous to compensating differentials, we use an instrumental variable (IV) strategy. The standard instruments found in the tax literature such as Gruber and Saez (2002), Eissa and Hoynes (2004), and Auten et al. (2008) calculate the predicted tax rate change by holding the base period income constant and only using tax schedule changes to generate exogenous variation in the change in the tax rate. As explained later in the paper, however, the error term in our second-step specification contains compensating differential heterogeneity which is correlated with the base period labor income. This correlation is potentially problematic so we purge the predicted tax change variation resulting purely from differences in the base period labor income using a "simulated IV" approach introduced by Currie and Gruber (1996a, 1996b) and Cutler and Gruber (1996). By ridding the variation in the base year labor income, our identification strategy relies on the differential impact that tax schedule changes have on workers with different initial tax-related characteristics such as spousal earnings and capital income.

Our empirical model includes occupation-year fixed effects, flexibly controlling for tax-driven general equilibrium effects that may change wages and the provision of amenities in an occupation. Because we use a two-step procedure, we adjust for the estimation error in the first step when estimating the second-step equation. We also apply the multi-way clustering method introduced by Cameron et al. (2006) to account for correlations between observations of the same worker over time and between observations of workers in the same occupation. Our analysis sample includes both workers who switch jobs from the first period to the second period and workers who do not.

Basic theory suggests that, all else equal, a higher marginal tax rate induces the worker to choose a job that pays less in wages but more in amenities. A naïve OLS model produces an estimate with the wrong sign: Higher tax rates lead workers to choose higher wage jobs. This finding confirms the endogenous relationship between wages and marginal tax rates. Our IV estimates, however, have the expected sign: Higher tax rates induce workers to choose jobs with lower wages and presumably higher amenities. Therefore, the findings in this paper show that occupational choice is indeed a component of income tax distortion. Although statistically significant, the magnitude of our estimates turns out to be economically modest. Our preferred estimates suggest a compensated elasticity of 0.03, implying that a 10 percent increase in the net-of-tax rate causes workers, on average, to switch to an occupation paying 0.3 percent higher wages. Interestingly, we estimate similar elasticities for both men and women and we find no evidence that younger workers are more responsive to tax changes than older workers in choosing occupations with different wages and amenities.

By construction, our empirical strategy captures any amenity that varies across jobs. For example, some jobs are riskier than others, and some jobs provide more generous health insurance plans than others. Our estimates reflect the worker's choice over on-the-job safety and the generosity of employer-provided health insurance. However, if a worker decides to work harder when he faces a lower marginal tax rate regardless of which job he chooses, this effort effect will not be captured by our estimates. As shown in Albouy (2009), federal taxes are levied on nominal wages which discourage workers from living in cities with above-average wages and high cost-of-living. Our paper does not explicitly model location choices, but to the extent some jobs are concentrated in certain geographic areas, local amenities associated with these jobs are captured in our estimates.

The wage-amenity choice with respect to taxes is a critical component of the individual labor supply response to taxes. By distorting occupational choices, income taxes can lead to inefficiently allocated workers with potentially long-term economic consequences. A previous literature has studied the link between income taxes and specific amenities. For example, Gruber and Lettau (2004) investigate the tax effect on health insurance provision by employers, and Powell (2009) examines the wage responses to tax changes of jobs with different injury and mortality risks. However, no study has estimated the overall tax elasticity of the wage-amenity tradeoff. We view this paper as a direct complement to the existing literature.

The rest of this paper proceeds as follows. In section 2, we review the existing research. We then present a theoretical framework on the effect of income taxes on the demand for wage and amenities in section 3. In section 4, we discuss the data used in this paper. Section 5 introduces our two-step methodology and explains our identification strategy. Section 6 presents the estimation results, and section 7 concludes.

2 Previous Research

There is a large existing literature on taxes and labor supply. This literature is summarized by Hausman (1985) and Blundell and MaCurdy (1999) and focuses primarily on the impact of taxes on number of hours worked and labor force participation rate. The consensus of the literature is that women appear to change working hours and labor force participation status in response to taxes, whereas prime-age men do not. While working hours and labor force participation are important aspects of labor supply, there are other margins where taxes could play an important role. For example, Feldstein (1995) points out:

> [These] studies focus on labor force participation and hours because those are the aspects of labor supply that are easily measured. In actual practice, individuals can vary their labor supply in the short run by changing how hard they work and in the long run by their location and the types of jobs that they accept. (page 553)

Our paper contributes to the literature by examining the occupational choice aspect of labor supply.

A related set of studies has looked at the effect of taxes on self-employment and entrepreneurship. Gentry and Hubbard (2002) study the impact of tax progressivity on the decision to become an "entrepreneur." They find evidence suggesting that a progressive tax schedule with imperfect loss offsets discourages entry to entrepreneurship. Bruce (2000, 2002) investigates the link between tax rates and transitions into and out of self-employment. He finds higher marginal tax rates actually increase the probability of self-employment. He interprets this counter-intuitive finding as evidence of tax evasion among the self-employed.

Few studies have examined the relationship between taxes and occupational choice. Gentry and Hubbard (2004) look at the effect of tax rates and tax progressivity on changing to a self-reported "better" job. They argue that a more progressive tax system reduces the return to job search and discourages upward job mobility. They find that both higher tax rates and increased tax progressivity decrease the probability that a head of household will move to a better job in the coming year.

In contrast to specifications where the dependent variable is an indicator for job changes, our specification uses compensating differentials as the dependent variable. A sizable literature in the field of labor economics estimates compensating differentials associated with various job characteristics. For example, Duncan and Stafford (1980) establish the presence of compensating wage differentials for union workers. Garen (1988) estimates wage premia for risk of fatality and injury. Kostiuk (1990) shows a positive wage premium for shift work. However, no studies have used compensating differentials to examine the link between incomes taxes and occupational choices. While an indicator for job changes as the dependent variable may capture whether taxes matter to occupational choices, using compensating differentials as the dependent variable allows us to investigate both the direction of the tax effect on occupational choices (i.e., a higher tax rate causes the worker to choose a lower wage job) and the magnitude of the effect.

This paper is also closely related to the literature on the elasticity of taxable income. Feldstein (1999) argues that the compensated elasticity of taxable income is the central parameter needed to calculate the deadweight loss associated with income taxes. Given its importance, many studies have estimated the elasticity of taxable income using different data and empirical strategies.[1] We use a specification that is similar to the one found in Gruber and Saez (2002) and Auten et al. (2008). Our parameter of interest is, theoretically, a component of the overall tax elasticity, and we specify this relationship in section 6. We believe that understanding the magnitude of the components is important because the component elasticities help us understand the mechanisms through which the

[1] See Saez et al. (2009) for a survey of the literature.

tax distortions occur. The economic consequences can be very different if workers respond to taxes by changing the number of hours they work than if they respond by changing occupations.

3 Theoretical Framework

In this section, we illustrate the intuition behind our empirical test using a simple model. Assume that a worker chooses from a continuum of job options. Each job offers a combination of wages (w) that are subject to income tax and amenities (n) that are non-taxable. For any given worker at a given time, higher amenities imply relatively lower wages ($w'(n) < 0$), all else equal. Let y denote the worker's other income, and $T[z]$ the total tax liability given total taxable income z. The worker maximizes his utility over consumption (c) and amenities (n) subject to his budget constraint.

$$\max_{c,n} U(c,n)$$

$$\text{s.t. } c = w(n) + y - T[w(n) + y]$$

The first order condition of this maximization problem can be expressed as

$$\text{FOC: } w'(n)(1 - T') = -\frac{U_n}{U_c} \tag{1}$$

Assuming that $w'(n)$ is unchanged by tax changes for the worker, equation (1) suggests that when the marginal tax rate (T') increases, the left hand side of the equation increases, implying that $\frac{U_n}{U_c}$ decreases. Under the standard assumption that the utility function is concave, the demand for n must increase relative to the demand for c. Thus, the relative demand for wage compensation decreases when the marginal tax rate increases.[2]

[2] In Appendix A, we also consider models where the worker chooses occupation and working hours simultaneously. The tradeoff between wage and amenities remains the same when the hours decision is included, and these models also predict that workers will demand more non-taxable compensation when the marginal tax rate increases.

In practice, tax changes may induce changes to $w'(n)$. For example, higher taxes may cause firms to increase the provision of non-taxable compensations. This paper focuses specifically on the individual occupational decision and does not explicitly model the firm-level decisions concerning wages and amenity provision. Our empirical strategy, however, will flexibly account for such general equilibrium effects.

The simple model described above offers two key insights. First, the model shows that the marginal net-of-tax rate, $1 - T'$ or $1 - \tau$, is the price of amenities faced by the worker. Therefore, it is the marginal tax rate, not the average tax rate, that is the relevant tax measure in studying the wage-amenity tradeoff. Second, the model suggests that for any given after-tax income I, we expect a higher price of amenities reduces the demand for amenities ($\frac{\partial n}{\partial(1-\tau)}|_I < 0$). Since we do not observe n in the data, our empirical strategy focuses on the relationship between the marginal net-of-tax rate and the choice of wage. Because $w'(n) < 0$, $\frac{\partial n}{\partial(1-\tau)}|_I < 0$ implies $\frac{\partial w}{\partial(1-\tau)}|_I > 0$. When a worker's marginal net-of-tax rate increases, we expect to see him move to a higher wage job with fewer amenities, all else equal.[3] This model prediction guides our empirical specification where we study the relationship between the change in the individual's wage and the change in the tax rate, holding after-tax income constant.

4 Data Description

We use the Panel Study of Income Dynamics (PSID), a longitudinal data set containing household- and individual-level variables on a wide range of topics. Starting with the 1981 data, the PSID provides consistent occupation and industry codes using the 1970 3-digit Census coding. After the 1997 survey, the PSID becomes a biennial survey and

[3]In practice, the magnitude of the tax effect on occupational choices depends on factors such as whether the tax change is perceived as temporary or permanent, the extent to which employers would respond to tax changes by adjusting the provision of amenities, and how costly it is for workers to change jobs. As a result, one has to resort to actual data to gauge the size of the potential effect, which makes our empirical analysis more interesting.

we can no longer observe individuals' annual income in a continuous manner.[4] Our final data set, therefore, includes the years 1981-1997.

There are three sets of variables in the PSID that we use in our analysis. The first set includes income and family composition variables. We use these variables as inputs to the NBER's TAXSIM program to estimate tax liabilities and marginal tax rates. Specifically, the TAXSIM calculation of marginal tax rates is based on a very small finite difference in income (1 cent), and the TAXSIM calculator is widely used in the elasticity of taxable income literature.[5] Butrica and Burkhauser (1997) show that the tax rates and tax liabilities calculated by NBER's TAXSIM are similar to tax burden values supplied by the PSID staff from 1980 through 1991, the last year the PSID staff provided such information. Second, the PSID provides detailed information on labor supply. We derive the hourly wage rate for each worker in each year by taking the ratio of the worker's total labor income to his total hours worked. This hourly wage rate measure is used in the estimation of compensating differentials.

Third, the PSID contains 3-digit occupation and industry codes from the 1970 Census coding system. We use the occupation and industry codes that refer to each person's main job. One concern with the 3-digit coding system is that it may be too detailed and have a great deal of misclassifications. To minimize these errors while still capturing meaningful transitions between jobs, we categorize workers by the 2-digit occupation codes interacted with the 1-digit industry codes.[6] We believe that it is potentially interesting to distinguish the same occupation by industry. For example, accountants in the finance industry and accountants in public administration may receive very different amenities. In addition, this type of job transition may be relatively easy to make, providing a rich source of occupational changes in response to taxes. In the rest of this paper,

[4] The income measures in the PSID refer to the previous year so after 1997, we cannot determine a person's marginal tax rate and occupation in the same year. For example, the 1999 PSID tells us the worker's occupation as of 1999 but we have no way of knowing his 1999 income because the 1999 PSID only reports his 1998 income and the 2001 PSID only reports his 2000 income.

[5] See Feenberg and Coutts (1993) for more details about the TAXSIM calculator.

[6] We follow Kambourov and Manovskii (2009) in constructing the 2-digit occupation codes.

we refer to occupation-industry combinations simply as "occupations."

We limit our data set to workers between the ages of 25 and 55. We also exclude the self-employed from the sample because they may face a very different set of amenities than other workers in the same occupation. In our main specification, we follow Feldstein (1995) and Gruber and Saez (2002) and focus on behavior changes over a three-year time interval. This time interval seems appropriate because it is long enough for workers to respond to tax changes and switch occupations and it is also short enough for the tax change to stay relevant. We try alternative time intervals as robustness checks. In all analysis, we use sample weights to obtain nationally-representative estimates.

Our final analysis sample has 44,062 worker-year observations. Table 1 shows the summary statistics of our sample. On average, we observe 40.7 percent of the respondents change their occupations within the next three years. We are not the first to find such a large occupation change rate in the PSID. Kambourov and Manovskii (2009) provide an extensive discussion on this issue. The possibility of misclassification is high. If one constructs an indicator for job changes using these occupation categories directly, the indicator variable would contain substantial measurement error and such measurement error cannot be treated as classical because of the nature of binary variables. Our empirical strategy avoids this problem by looking at the change in compensating differentials rather than the change in these occupation categories. Figure 1 shows the occupation change rate across years by age group. It is reassuring that the younger workers appear to have higher occupational change rates than older workers. Table 1 also shows that the average hourly wage is $16.36 in 1997 dollars in the analysis sample and the average household income is $53,997. The average tax liability – the sum of federal, state, and half of the FICA taxes – faced by PSID respondents is $13,528, and the average marginal tax rate is 34.2 percent.

Our sample has 174 unique occupations.[7] The median size of an occupation is

[7]There are 9 unique 1-digit industry codes and 24 unique 2-digit occupation codes in our sample. However, not all possible occupation-industry combinations exist and we observe 175 unique occupation-industry combinations in our sample.

97 observations; the mean is 410 observations. Panel A of Table 2 lists the most frequent occupation changes observed in our data. Many of the listings occur twice with the original job and the new job reversed. It is possible that these observed transitions are driven by occupation misclassifications in the data. On the other hand, the original jobs and new jobs look similar. It is possible that workers freely move across similar jobs and these observed transitions are real. Panel B of Table 2 lists the most frequent transitions as a percentage of the original occupation population.

In Figure 2, we graph the sample-average marginal tax rate and occupation change rate together. We are not suggesting that any relationship from this graph should be interpreted as causal. The overall occupation change rate during the sample period remains relatively stable, but the calculated average marginal tax rate declined significantly. The Economic Recovery Tax Act of 1981 (ERTA1981) and the Tax Reform Act of 1986 (TRA1986) are the major tax changes taking place during the sample period. For example, ERTA1981 decreased marginal rates from 70 percent to 50 percent, and TRA86 cut the top federal income tax rate further to 28 percent. In addition, the Earned Income Tax Credit (EITC) was significantly expanded in 1987. The Omnibus Budget Reconciliation Act of 1990 and 1993 subsequently increased the top federal income tax rate. It is reassuring that the average marginal tax rate shown in Figure 2 reflects the major tax schedule changes. The many and significant changes introduced by these tax reforms provide useful variation to identify the effect of income taxes on occupational choices.

5 Estimation Strategy

5.1 The Setup of the Main Specification

The empirical question that we are interested in is whether workers move to higher wage jobs when net-of-tax rates increase. Suppose individual i works in occupation k at time

$t-1$ and faces a marginal tax rate of τ_{ikt-1}. At time t, the tax schedule changes and he chooses to work in occupation j. If we could observe individuals' counterfactual wages in occupations they are not working in, our ideal specification would be:

$$\ln w_{ijt} - \ln w_{ikt} = \beta_1 \ln\left(\frac{1-\tau_{ijt}}{1-\tau_{ikt-1}}\right) + \beta_2 \ln\left(\frac{z_{ikt-1} - T_t[z_{ikt-1}]}{z_{ikt-1} - T_{t-1}[z_{ikt-1}]}\right) + \gamma_{kt-1} + X'_{it-1}\Pi + (\nu_{ijt} - \nu_{ikt}) \tag{2}$$

where $T_t(z)$ is the tax liability under the tax schedule at time t for total pre-tax income z.

The left-hand side of equation (2) is the log difference between the wage rate that individual i receives in his new occupation j at time t (w_{ijt}) and the wage rate he would have received in his old occupation k at time t (w_{ikt}). We choose to use this variable as the dependent variable rather than the log difference of the two observed wage rates ($\ln w_{ijt} - \ln w_{ikt-1}$) for two reasons. First, note that

$$\ln w_{ijt} - \ln w_{ikt} = (\ln w_{ijt} - \ln w_{ikt-1}) - (\ln w_{ikt} - \ln w_{ikt-1}).$$

When the tax schedule changes, the worker may receive a different wage rate even if he stays in the same occupation (i.e., the second term $\ln w_{ikt} - \ln w_{ikt-1} \neq 0$). For example, tax decreases may cause the worker to work harder and receive a higher wage. By focusing on wage differences in the same time period, we net out wage changes over time that have nothing to do with occupational choices. For workers who do not change jobs from time $t-1$ to t (i.e., $j = k$), our dependent variable equals zero. Second, the choice of dependent variable makes the economic interpretation of equation (2) intuitive. After the tax change, the worker faces a menu of jobs with different wage-amenity combinations which may have shifted from time $t-1$. Given the change in his marginal tax rate, the worker re-optimizes along the wage-amenity margin and chooses from the current menu of wage-amenity combinations.

The right-hand side of equation (2) is similar to the Gruber and Saez (2002) specification and separately estimates the substitution and income effects.[8] By separately estimating the income effect, we can interpret the coefficient on the marginal net-of-tax rate variable (β_1) as a compensated elasticity. When the marginal net-of-tax rate increases, we expect workers to move to higher wage jobs. An increase in after-tax income should cause workers to move to lower wage jobs due to an increased demand for amenities. Therefore, we expect $\beta_1 > 0, \beta_2 < 0$.

The right-hand side of equation (2) also includes occupation-year fixed effects (γ_{kt-1}). These fixed effects flexibly control for any natural job transitions at each point of time. There may be occupation-specific shocks, including tax-induced general equilibrium effects on each occupation, which influence the flows into and out of a specific job. The occupation-year fixed effects absorb these confounding factors and help us identify the true effect of income taxes on occupational choices. In addition, X_{it-1} is a vector of control variables at time $t-1$, including race, sex, marital status, education, number of dependents, job tenure, job tenure squared, and age group fixed effects.

Although equation (2) has many nice features, it cannot be estimated directly in practice because we do not observe w_{ikt}, the wage that individual i would have earned at time t in occupation k had he not switched to occupation j. To circumvent this problem, we use the compensating differentials of these occupations instead of the wages themselves.

5.2 Estimating Compensating Differentials

The wage received by individual i in occupation j at time t can be decomposed into two parts.

$$\ln w_{ijt} = \alpha_{it} + \phi_{it}(n_{ijt})$$

[8] Our income effect variable is slightly different than the Gruber-Saez variable. Our specification holds the taxable income at time $t-1$ fixed, whereas the Gruber-Saez specification lets taxable income change between periods $t-1$ and t. The income effect is the response to the budget constraint shift and we want the specification to reflect this insight. We hold z constant so that households are responding only to the budget constraint shift.

α_{it} is the return to the individual's ability and effort at time t. It represents the wage earned by the individual regardless of occupation. $\phi_{it}(n_{ijt})$ represents the tradeoff between wages and amenities faced by the individual. A higher n_{ijt} means that the worker receives more amenities. $\phi_{it}(n_{ijt})$ can be interpreted as the individual-specific compensating differential.

Let ϕ_{jt} represent the average compensating differential for occupation j at time t. We rewrite the wage equation as

$$\ln w_{ijt} = \alpha_{it} + \phi_{jt} + \mu_{ijt} \tag{3}$$

where $\mu_{ijt} \equiv \phi_{it}(n_{ijt}) - \phi_{jt}$ is individual heterogeneity in the compensating differential. Note that equation (3) places no real restrictions on the wage function and is essentially tautological. The purpose is simply to divide each person's wage into separate components: return to skill/effort, mean compensating differential for the occupation, and individual heterogeneity. This heterogeneity in the compensating differential may exist for two reasons. First, the worker may receive a different level of amenities ($n_{ijt} \neq n_{jt}$). Second, the worker may face a different price function for amenities. Put differently, a worker may get more amenities than the average worker in an occupation and the worker may get paid a higher wage for a given level of amenities. Our wage model allows for both types of individual heterogeneity in the compensating differential within an occupation.

Using the compensating differentials in place of wages, we plug equation (3) into equation (2) and get

$$\phi_{jt} - \phi_{kt} = \beta_1 \ln\left(\frac{1-\tau_{ijt}}{1-\tau_{ikt-1}}\right) + \beta_2 \ln\left(\frac{z_{ikt-1} - T_t[z_{ikt-1}]}{z_{ikt-1} - T_{t-1}[z_{ikt-1}]}\right) + \gamma_{kt-1} + X'_{it-1}\Pi + (\epsilon_{ijt} - \epsilon_{ikt}) \tag{4}$$

$\phi_{jt} - \phi_{kt}$ is the difference in the compensating differentials between the old and new occupations of individual i at time t, reflecting differences in amenity levels.[9] The individual ability term α_{it} conveniently drops out in this equation. If we have measures of ϕ_{jt} and ϕ_{kt}, we can estimate our parameter of interest β_1. It is worth emphasizing that the residual ϵ terms include both the ν terms in equation (2) (true random error) and the μ terms in equation (3) (individual heterogeneity in compensating differentials). When constructing our instruments, which will be discussed in the next section, we must make sure that they are orthogonal to the μ terms in the residual.

Equation (3) cannot be estimated because there is only one observation per person-year and α_{it} is not identified. Instead, we use $\alpha_{it} = \alpha_i + V_{it}'\delta$ and rewrite equation (3) as

$$\ln w_{ijt} = \alpha_i + \phi_{jt} + V_{it}'\delta + \mu_{ijt} \tag{5}$$

where V includes job tenure, job tenure squared, and a vector of age group fixed effects.[10] Because the α_{it} term drops out when we take the log difference of two wage rates ($\ln w_{ijt} - \ln w_{ikt}$), the form of α_{it} is largely irrelevant. Identifying α_{it} would certainly help the precision of our estimates in our first step, but it is unnecessary for obtaining consistent estimates of the tax elasticities. We test the sensitivity of our results to the specification of α_{it} in a robustness check later in this paper and show that the exact specification does not, in fact, seem to matter.

The compensating differentials (ϕ_{jt}) capture any job characteristics that vary across occupations. Because our empirical strategy compares the compensating differential of the old job to that of the new job in studying the wage-amenity tradeoff after tax changes, we cannot apply our empirical strategy to study the occupational choices of recent graduates who just enter the labor market as they do not have an old job. Neither can we apply our empirical strategy to workers who exit the labor market as they do not have a new job.

[9] Even though the left-hand side of equation (4) ($\phi_{jt} - \phi_{kt}$) may seem to only vary at the occupation-year level, it actually varies at the individual level because k and j index the occupations of individual i's choosing at times $t-1$ and t.

[10] Because equation (5) includes individual fixed effects, time-invariant individual characteristics such as gender, race, and education attainment are implicitly controlled for.

5.3 Instruments

The main specification shown in equation (4) cannot be estimated consistently using ordinary least squares (OLS) because the tax variables are endogenous. Tax rates and tax liabilities are functions of wages and, similarly, compensating differentials. Workers who switch to high wage occupations will see their taxes increase. To solve the endogeneity problem, we need to construct valid instruments for the tax variables.

The natural candidate to instrument for the tax change variable $\ln\left(\frac{1-\tau_{ijt}}{1-\tau_{ikt-1}}\right)$ is the Gruber-Saez instrument $\ln\left(\frac{1-\tau_{ikt}^{p}}{1-\tau_{ikt-1}}\right)$, where τ_{ikt}^{p} is the "predicted" marginal tax rate constructed by holding the worker's base period income and characteristics fixed and applying the new tax schedule. Recall that the residual term in equation (4) contains both a random error component and a component representing the individual heterogeneity in compensating differentials. The compensating differential heterogeneity term is correlated with labor income, and the Gruber-Saez instrument is constructed using base period characteristics including labor income. As a result, the residual term in our main specification may be correlated with the Gruber-Saez instrument.[11]

We construct a valid instrument by purging the variation in workers' base period labor income from the standard Gruber-Saez instrument. The idea is that the Gruber-Saez instrument includes two sources of variation: One is the variation in workers' base period labor income, and the other is the variation in workers' base period other income and tax-related characteristics. If we remove the former from the Gruber-Saez instrument, we obtain an instrument that varies across individuals only because the tax change affects workers differently based on their other income (e.g., spousal earnings and asset income) and tax-related characteristics (e.g., marital status and number of dependents) in the initial period.

[11] Another way to think about the difference between our main specification and the Gruber-Saez specification is that we do not observe individual-specific compensating differentials whereas Gruber and Saez (2002) observe their dependent variable at the individual level. We use estimates of occupation-specific compensating differentials as the dependent variable and move the individual heterogeneity term to the right hand side of the main specification. This rearrangement allows us to measure the dependent variable, but it causes the Gruber-Saez instrument to be problematic in our model.

The idea is best illustrated through an example. Say that worker A has a spouse who makes \$50,000 in earnings. The other workers in the same occupation-year have \$30,000 spousal earnings. The Gruber-Saez instrument tells us the predicted tax change for worker A given A's labor income, \$50,000 in spousal earnings, and the change in the tax schedule. We can also calculate worker A's predicted tax change given A's labor income, the spousal earnings of the other workers (\$30,000), and the change in the tax schedule. This is the counterfactual tax change that worker A would have experienced with his own initial labor income but the spousal earnings of the other workers in the same occupation-year. In other words, the only source of variation in this counterfactual tax change is the interaction between the tax schedule change and the worker's initial period labor income. The difference between the Gruber-Saez instrument and the counterfactual predicted tax change is, therefore, due entirely to the fact that worker A has \$50,000 spousal earnings in the initial period whereas other workers have \$30,000.

Specifically, our instrument is

$$\ln\left(\frac{1-\tau_{ikt}^{p}}{1-\tau_{ikt-1}}\right)-\ln\left(\frac{1-\tau_{ikt}^{s}}{1-\tau_{ikt-1}^{s}}\right)$$

where the first term is the Gruber-Saez instrument and the second term is the counterfactual predicted tax change. In practice, we borrow the idea from the simulated instrument literature introduced by Currie and Gruber (1996a, 1996b) and Cutler and Gruber (1996) and calculate the second term by "cycling" the other income and tax-related characteristics of every other worker in the same occupation-year and then taking the average. As illustrated in the example above, we subtract this term from the Gruber-Saez instrument to remove the part of predicted tax change that is driven by workers' own labor income and to retain the part of predicted tax change that comes from differences in initial characteristics (other than labor income). We also construct an instrument for the after-tax income change variable in a similar way.

The variation in our instrument comes from the fact that when the tax schedule changes, workers with different initial other income and tax-related characteristics may experience different changes in their marginal tax rates.[12] Our identification assumption, therefore, is that these initial other income and tax-related characteristics do not systematically affect workers' decisions to change jobs other than through the nonlinear interaction with the tax schedule change. Similar identification assumptions have been used in many previous studies. For example, researchers compare single women with and without children before and after EITC expansions to identify the effect of EITC on labor supply (see Eissa and Hoynes (2006) for a summary). Studies such as Eissa (1995) treat primary earners' income as exogenous in identifying the tax effect on the labor supply decision of secondary earners. Although identification assumptions cannot be directly tested, we estimate models allowing workers with different base period spousal income, asset income, and marital status to follow different trends in job switching as robustness checks later in the paper.

6 Estimation Results

6.1 Basic Results

The first step of our two-step methodology is to obtain compensating differentials for each occupation in each year. We estimate equation (5) applying the PSID sample weights and clustering the standard errors at the individual level. We also require an occupation-year to have at least 3 observations to be included in the sample. This restriction affects 40 occupations out of the total of 174 occupations in the sample.[13]

Our measure of compensating differentials is a summary statistic of all amenities associated with an occupation. Since we cannot create a full list of all amenities, it would

[12] The full set of other income and tax-related characteristics that enter the NBER's TAXSIM model and generate variation in our instruments include marital status, number of dependents, number of individuals over 65 years of age in the household, wage and salary income of spouse, dividend income, other property income, child care expenses, and number of dependents under age 17.

[13] We also tried requiring an occupation-year to have at least 5 observations to be included in the sample. The results are similar since different cutoffs only affect a small number of observations.

be impossible to know the true compensating differential function and to verify whether $\hat{\phi}_{jt}$ is the "right" compensating differential estimate. However, there are ways for us to indirectly check whether these estimates are sensible. We previously presented the most frequent occupation transitions observed in the sample in Table 2. Because of the apparent fluidity between some occupations, we would think that these occupations may have very similar compensating differentials. We find it comforting that this seems to, in fact, be the case. For example, the two most frequent transitions are between "Operatives (Manufacturing)" and "Craftsmen (Manufacturing)." "Operatives (Manufacturing)" is estimated to have the 88th highest average compensating differential while "Craftsmen (Manufacturing)" has the 89th highest. The similarity of these estimates suggests that these jobs are close substitutes.[14]

The second step of our two-step estimation procedure is to regress the change in predicted compensating differentials on tax changes. The dependent variable, the difference in compensating differentials, is estimated rather than observed. In the second step of our estimation, we use the variance-covariance matrix of the first-step regression to adjust the sample weights. Specifically, we weight each observation by the inverse of the standard error of the dependent variable. In addition, we also incorporate the PSID weights to ensure that our sample is nationally representative. Because our sample observations are not independent of each other, we need to adjust the estimated standard errors accordingly. The PSID is a panel data set, so we must adjust the standard errors by clustering at the individual level. Furthermore, we estimate compensating differentials for each occupation in each year, and these compensating differentials appear to be serially correlated. Therefore, we also adjust the standard errors by clustering at the occupation level. We use the multi-way clustering procedure introduced by Cameron et al. (2006) to account for clustering at both the individual and occupation levels.

In our model, we hypothesize that workers will choose higher wage jobs when their

[14] We also check the internal consistency of the compensating differential estimates by examining how correlated the compensating differentials of the same occupation are over time. We find very strong correlations. The results are shown in the working paper version of this paper.

net-of-tax rate increases and β_1 should be positive. We also hypothesize that a higher after-tax income will cause the worker to demand more amenities and lower wages and β_2 should be negative. Column (1) of Panel B in Table 3 presents the OLS estimates. The OLS estimate of β_1 is negative and statistically significant, and the OLS estimate of β_2 is positive and statistically significant. This finding is consistent with our suspicion that the change in the tax variables are endogenous. A worker who moves to a higher wage job will probably also face a higher tax rate because of his higher income. In addition, the OLS estimation results show that male, white, and more educated workers tend to move to higher wage jobs. The number of dependents and the job tenure in the old occupation do not seem to have a significant impact.

Panel A in Table 3 presents the first stage results for our main specification using the instruments described in the previous section. We report the relevant coefficients and Shea's Partial R^2 statistic which indicates the strength of the first stage. Note that our instruments strongly predict the endogenous variables.[15] Column (2) of Panel B in Table 3 presents the IV estimates. Unlike the OLS estimates, the IV estimate of β_1 has the expected sign. The estimated magnitude suggests that a 10 percent increase in the net-of-tax rate would cause individuals to move to an occupation with a wage that is 0.33 percent higher. This elasticity is economically modest, though statistically significant. The IV estimate of β_2, the effect of changes in after-tax income on occupation choices, is small and statistically insignificant. The coefficients on other control variables are similar to the OLS estimates.

6.2 Robustness Checks and Extensions

We identify the effect of income taxes on occupational choice by exploiting the nonlinear interaction between tax schedule changes and workers' initial other income and tax-related characteristics. A central assumption of our empirical strategy is that these initial

[15] The strength of the prediction of the after-tax income variable is a side effect of holding initial pre-tax income constant in the endogenous variable, as discussed in section 5.

other income and tax-related characteristics do not predict secular trends in occupational choices that correlate with tax changes. For example, our identification may fail if workers with higher spousal earnings experience a larger increase in marginal tax rates and they also tend to move to jobs with lower wages and higher amenities in those years. Although the identification assumption that the initial other income and tax-related characteristics are exogenous cannot be directly tested, we can estimate models with various controls of initial spousal income, asset income, and marital status to check the robustness of our results.

In columns (2) and (3) of Table 4, we follow Gruber and Saez (2002) by including linear controls and a more flexible spline in the second-step specification. The linear controls simply involve including asset and spousal income as control variables. The spline allows the effect to be non-linear. We include two separate 5-piece splines for both asset income and spousal income. The results are very similar to the main results shown in column (1). In column (4) of Table 4, we include the interaction between year fixed effects and spousal and asset income in the initial period. This specification allows workers with different base period spousal and asset income to follow different trends in each year. The results remain very similar to the main results. In column (5) of Table 4, we include the interaction between year fixed effects and an indicator for being married. This specification allows married workers to follow different trends than single workers in each year, and again, the results are unchanged. Overall, these robustness checks provide support to our identification assumption.

Our wage model makes the assumption that $\alpha_{it} = \alpha_i + V_{it}'\delta$. We argue that the exact specification of α_{it} is unimportant since this term differences out. We can test the validity of this argument by changing the specification of α_{it}. In column (6) of Table 4, we replicate our analysis using $\alpha_{it} = \alpha_i$ in equation (5). Thus, the compensating differentials are estimated in a regression with only individual fixed effects and occupation-year fixed effects. The IV estimates are very similar to our main results, suggesting that our model is robust to different specifications of α_{it}.

We also look at the effect of our weighting method on the estimates. Recall that we weight each observation by the inverse of the standard error of the estimated compensating differentials. In column (7) of Table 4, we do not use these weights obtained from the first step of the two-step estimation procedure to see if our estimates are driven by these weights. We do, however, still use the PSID sample weights. The results are little changed, although the standard errors are slightly bigger. This robustness check suggests that our weighting procedure, while improving the precision of our estimates, is not driving the main results as we would come to the same conclusions without the weighting.

It is well-known that male and female labor markets may have different dynamics. Hence, we study the male and female samples separately in Panel A of Table 5. Note that whenever we use a different sub-sample in this paper, we estimate the compensating differentials using only the relevant sub-sample. In other words, we estimate compensating differentials separately by gender here. Also, the tax instruments are formed using only the relevant sub-sample. The IV estimate of the key coefficient is higher for women than men, though this difference is not statistically significant. The results suggest that women and men are similarly responsive to tax changes when it comes to wage-amenity decisions.[16] The literature on the tax elasticity of working hours and labor force participation typically finds that women are more responsive than men. However, it is not clear that this result should extend to the wage-amenity margin. Men's occupation choice sets may have a higher variation in amenities than women's. For example, Hersch (1998) documents that jobs with high injury rates are male-dominated. This difference suggests that men may be actively choosing between safe and risky jobs while women are less likely to make a job decision along the safety margin, regardless of taxes.

We have so far focused on a 3-year interval length in studying job switching.

[16] Note that the coefficient of interest is larger for each sub-sample than the aggregate sample. Because compensating differentials and tax instruments are constructed differently in the sub-samples and because we do not include occupation-year-gender three-way interactions in the full-sample regression, it is not guaranteed that the full sample estimate will be between the male and female estimates.

Even though the 3-year interval length is our preferred specification because it allows the worker ample time to respond to taxes by searching and moving to a different occupation, we believe the adjustment time itself is of interest. In Panel B of Table 5, we present IV results using 1-, 2-, 3- and 4-year intervals. The estimate of β_1 is small and statistically insignificant in the 1-year interval specification. However, it is positive and statistically significant in the 2-, 3-, and 4-year specifications. This pattern suggests that the occupation adjustment is not immediate, but that by the second year, the full adjustment has occurred.

Figure 1 shows that younger workers change jobs more often than older workers. We might think the tax elasticity of occupational choices may also differ by age. In Panel C of Table 5, we cut the sample into "young" (age 25-34) and "old" (age 35-55) to explore the potential heterogeneity across age groups. As mentioned before, the compensating differentials and tax instruments are generated using only the sub-sample in question. We find evidence suggesting that older workers are more responsive to tax changes than younger workers, although the two are not statistically different. This finding suggests that our main results are not being driven by young workers, as some might expect. Even though young workers are more likely to change occupations, they do not seem to be more responsive to taxes when making these changes. This may be because older workers understand their marginal tax rates better than younger workers.

6.3 Further Discussion

A few features of our empirical strategy are worth mentioning. Compared to the natural tendency to use an indicator variable for "changing jobs" to study how taxes affect job turnover, our approach uses compensating differentials instead and we model the substantial occupation misclassification as classical measurement error in the dependent variable. For the measurement error in the estimated compensating differentials to bias our estimates, it has to be systematically correlated with both the change in the marginal tax rate experienced by the worker and the worker's subsequent decision to switch jobs. Because our main specification controls for occupation-year fixed effects, such correlations have to hold within an occupation-year to break our identification, which seems unlikely.

One reason why our estimated compensating differentials may vary across occupations is because workers in different occupations value the amenities differently. The value is partially a function of the preferences of the workers in that occupation. Workers are likely to have heterogeneous preferences for specific amenities, and they may sort into different occupations based on these preferences. For example, teachers may value on-the-job safety differently than engineers. We are not using the average valuation of on-the-job safety among all workers. Instead, we care about the empirical market valuation of amenities among workers in that specific occupation. This is what we estimate and use.

Moreover, the dependent variable in our main specification is not the difference in compensating differentials of two random occupations, but the difference for the new occupation and the old occupation chosen by individual i. Such "selection," however, should not introduce bias in our estimates. The individuals in an occupation are likely good fits for that occupation relative to the average worker in the labor market. But because we include occupation-year fixed effects and only compare workers within the same occupation in the same year, the individual wage component (μ_{ijt}) is the individual's "fit" relative to the other workers in that occupation. Therefore, there should be no systematic bias introduced at the occupation level.

Another benefit of our approach is that we explicitly estimate a parameter of economic interest. This point can best be illustrated by introducing an instructive framework. Denote I as the total income, which is the sum of labor income L and capital income K. Denote s_L and s_K as the share of labor income and capital income respectively. As discussed before, a vast literature has attempted to estimate the elasticity of taxable income with respect to the net-of-tax rate $\epsilon_{I,1-\tau}$. This aggregate elasticity is a weighted average of the labor income elasticity and capital income elasticity:

$$\epsilon_{I,1-\tau} = \epsilon_{L,1-\tau} s_L + \epsilon_{K,1-\tau} s_K$$

Labor income can be expressed as $L = wh$ where w is the hourly wage rate and is modeled by equation (3) and h is number of hours worked. We can write the labor income elasticity $\epsilon_{L,1-\tau}$ as

$$\begin{aligned} \frac{\partial \ln L}{\partial \ln(1-\tau)} &= \frac{\partial \ln w}{\partial \ln(1-\tau)} + \frac{\partial \ln h}{\partial \ln(1-\tau)} \\ &= \frac{\partial\left(\alpha + \phi\right)}{\partial \ln(1-\tau)} + \frac{\partial \ln h}{\partial \ln(1-\tau)} \\ &= \epsilon_{w,1-\tau|\phi} + \epsilon_{w,1-\tau|\alpha} + \epsilon_{h,1-\tau} \end{aligned}$$

The above equation indicates that the elasticity of labor income consists of three components. The first component is the elasticity of the individual-specific return (the elasticity of the wage holding the compensating differential constant). Workers may decide to work harder when the tax rate changes, so we can think of this term as an elasticity of "effort." The second component is the choice of compensating differentials or amenities (the elasticity of the wage holding effort and skill constant). This is the margin we study. The third component is the hours worked. By focusing on wages and compensating differentials, we are able to quantify the relative importance of $\epsilon_{w,1-\tau|\alpha}$ in the tax elasticity of labor income.

7 Conclusion

Taxes can affect occupational choice by distorting the return to taxable compensation (wages) relative to non-taxable compensation (amenities). In this paper, we study the effect of tax changes on wage-amenity decisions where amenities are defined in a broad and agnostic manner. Job-specific amenities include both observable characteristics such as health insurance provision and unobservable characteristics such as stress and workplace environment. We introduce a two-step estimation procedure and use compensating differentials as the summary statistic of all amenities associated with a job. We believe that this methodology offers a fruitful means of characterizing the amenity decision faced by workers and can be extended to other contexts where occupational choice or amenities play a significant role.

We find that when the net-of-tax rate increases, workers keep a higher fraction of his wage earnings and he moves to higher wage jobs, implicitly receiving fewer amenities. We estimate a statistically significant compensated elasticity of 0.03, suggesting that a 10 percent increase in the marginal net-of-tax rate leads workers to choose an occupation with a 0.3 percent higher wage. This paper contributes to the tax literature on several fronts. First, we look beyond working hours and labor force participation and provide evidence on other aspects of labor supply. Second, this paper complements existing studies by looking at all amenities associated with a job as opposed to one single dimension of amenities. In summary, we show that occupational choice is an important mechanism of distortion resulting from income taxes that has been generally ignored by the literature. Our paper focuses on this important margin and provides critical evidence for the understanding of how taxes impact economic welfare.

Although this paper shows that income taxes indeed affect occupational choices through the wage-amenity tradeoff, the field calls for more studies on the link between income taxes and occupational choices using other data sets and methods. Additionally, many other interesting and important research questions in this area awaits for more attention from researchers, including the effect of taxes changes on the provision of amenities by employers, the tax distortion on education attainment, and the effect of income taxes on workers' effort levels.

Appendix A: Models with Intensive Labor Supply

Adding the intensive labor supply decision (h = hours worked) does not change the FOC for the wage-amenity tradeoff in a meaningful way. We can model amenities in two different ways. First, we can think of each job as having a fixed level of amenities n. The worker maximizes a utility function which now contains hours worked. Previously, we used $w(n)$ to represent labor income. After adding the intensive margin of labor supply in the model, we think of labor income as equal to the wage $\omega(n)$ times hours worked h:

$$\max_{c,h,n} U(c,h,n)$$
$$\text{s.t. } c = \omega(n)h + y - T[w(n)h + y]$$

The first order conditions of this maximization problem can be expressed as

$$\text{FOC1: } \omega'(n)h(1-T') = -\frac{U_n}{U_c}$$
$$\text{FOC2: } \omega(n)(1-T') = -\frac{U_h}{U_c}$$

Note that the FOC regarding to the choice of amenities is essentially the same as the model shown in section 3 of the paper.

Alternatively, we can think of amenity consumption as proportional to the number of hours worked. For example, a safe working environment decreases fatality rates per hour. Each hour worked, then, is extra consumption of this safety. We can model amenities as nh instead of n:

$$\max_{c,h,n} U(c,h,nh)$$
$$\text{s.t. } c = \omega(n)h + y - T[w(n)h + y]$$

The first order condition of this maximization problem can be expressed as

$$\text{FOC1: } \omega'(n)(1-T') = -\frac{U_n}{U_c}$$
$$\text{FOC2: } \omega(n)(1-T') = -\frac{U_h + U_n n}{U_c}$$

The relevant FOC is, again, essentially the same.

References

Albouy, David. 2009. "The Unequal Geographic Burden of Federal Taxation." *Journal of Political Economy*, 117(4): 635-667.

Auten, Gerald, Robert Carroll, and Geoffrey Gee. 2008. "The 2001 and 2003 Tax Rate Reductions: An Overview and Estimate of the Taxable Income Response." *National Tax Journal*, 61(3): 345-364.

Blundell, Richard, and Thomas MaCurdy. 1999. "Labor supply: A Review of Alternative Approaches." In *Handbook of Labor Economics*, Vol 3, ed. Orley Ashenfelter and David Card, 1559-1695. New York: Elsevier.

Bruce, Donald. 2000. "Effects of the United States Tax System on Transitions into Self-Employment." *Labour Economics*, 7(5): 545-574.

Bruce, Donald. 2002. "Taxes and Entrepreneurship Endurance: Evidence from the Self-Employed." *National Tax Journal*, 54(1): 5-24.

Butrica, Barbara A., and Richard V. Burkhauser. 1997. "Estimating Federal Income Tax Burdens for Panel Study of Income Dynamics (PSID) Families Using the National Bureau of Economic Research TAXSIM Model." Syracuse University Center for Policy Research Aging Studies Program Paper 12.

Cameron, Colin A., Jonah B. Gelbach, and Douglas L. Miller. 2006. "Robust Inference with Multi-way Clustering." NBER Technical Working Paper 327.

Currie, Janet, and Jonathan Gruber. 1996. "Health Insurance Eligibility, Utilization of Medical Care, and Child Health." *Quarterly Journal of Economics*, 111(2): 431-466.

Currie, Janet, and Jonathan Gruber. 1996. "Saving Babies: The Efficacy and Cost of Recent Changes in the Medicaid Eligibility of Pregnant Women." *Journal of Political Economy*, 104(6): 1263-1296.

Cutler, David M, and Jonathan Gruber. 1996. "Does Public Insurance Crowd Out Private Insurance?" *Quarterly Journal of Economics*, 111(2): 391-430.

Duncan, Greg J., and Frank P. Stafford. 1980. "Do Union Members Receive Compensating Wage Differentials?" *American Economic Review*, 70(3): 355-371.

Eissa, Nada 1995. "Taxation and Labor Supply of Married Women: The Tax Reform Act of 1986 as a Natural Experiment." NBER Working Paper 5023.

Eissa, Nada, and Hilary Hoynes. 2004. "Taxes and the Labor Market Participation of Married Couples: The Earned Income Tax Credit." *Journal of Public Economics*, 88(9-10): 1931-1958.

Eissa, Nada, and Hilary Hoynes. 2006. "Behavioral Responses to Taxes: Lessons from the EITC and Labor Supply." In *Tax Policy and the Economy*, Vol 20, ed. James M. Poterba, 73-110. Chicago: The University of Chicago Press.

Feenberg, Daniel, and Elisabeth Coutts. 1993. "An Introduction to the TAXSIM Model." *Journal of Policy Analysis and Management*, 12(1): 189-194.

Feldstein, Martin. 1995. "The Effect of Marginal Tax Rates on Taxable Income: A Panel Study of the 1986 Tax Reform Act." *Journal of Political Economy*, 103(3): 551-572.

Feldstein, Martin. 1997. "How Big Should Government Be?" *National Tax Journal*, 50(2) 197-213.

Feldstein, Martin. 1999. "Tax Avoidance And The Deadweight Loss Of The Income Tax." *The Review of Economics and Statistics*, 81(4): 674-680.

Garen, John. 1988. "Compensating Wage Differentials and the Endogeneity of Job Riskiness." *Review of Economics and Statistics*, 70(1): 9-16.

Gentry, William M., and R. Glenn Hubbard. 2002. "Tax Policy and Entry into Entrepreneurship." Mimeograph, Columbia University.

Gentry, William M., and R. Glenn Hubbard. 2004. "The Effects of Progressive Income Taxation on Job Turnover." *Journal of Public Economics*, 88(11): 2301-2322.

Gruber, Jonathan, and Michael Lettau. 2004. "How Elastic is the Firm's Demand for Health Insurance?" *Journal of Public Economics*, 88(7-8): 1273-1293.

Gruber, Jonathan, and Emmanuel Saez. 2002. "The Elasticity of Taxable Income: Evidence and Implications." *Journal of Public Economics*, 84(1): 1-32.

Hausman, Jerry A. 1985. "Taxes and Labor Supply." In *Handbook of Public Economics*, Vol 1, ed. Alan J. Auerbach and Martin Feldstein, 213-263. New York: Elsevier.

Hersch, Joni. 1998. "Compensating Differentials for Gender-Specific Job Injury Risks." *American Economic Review*, 88(3): 598-627.

Kambourov, Gueorgui, and Iourii Manovskii. 2009. "Occupational Specificity of Human Capital." *International Economic Review*, 50(1): 63-115.

Kostiuk, Peter F. 1990. "Compensating Differentials for Shift Work." *Journal of Political Economy*, 98(5): 1054-1075.

Powell, David M. 2009. "Using Income Tax Changes to Identify the Value of a Statistical Life." PhD diss. MIT.

Saez, Emmanuel, Joel Slemrod, and Seth H. Giertz. 2009. "The Elasticity of Taxable Income with Respect to Marginal Tax Rates: A Critical Review." NBER Working Paper 15012.

Figure 1:

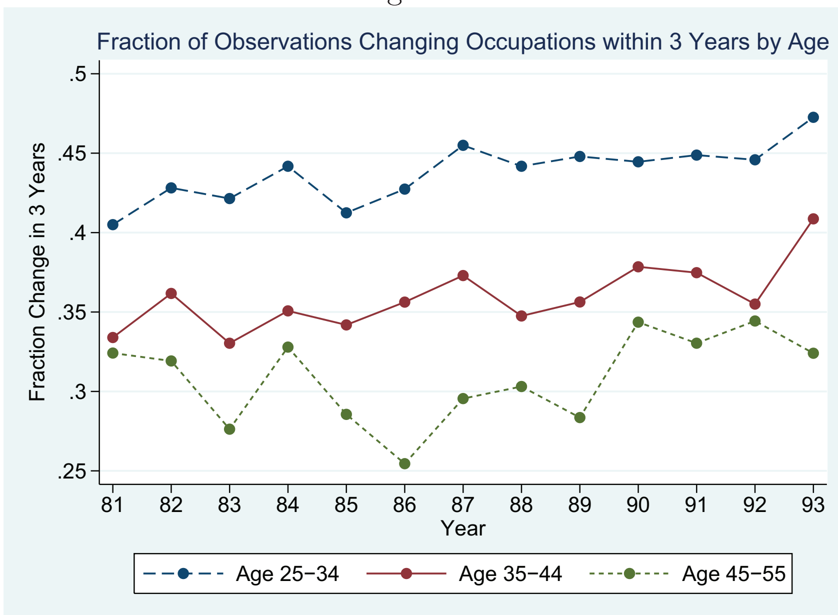

Figure 2:

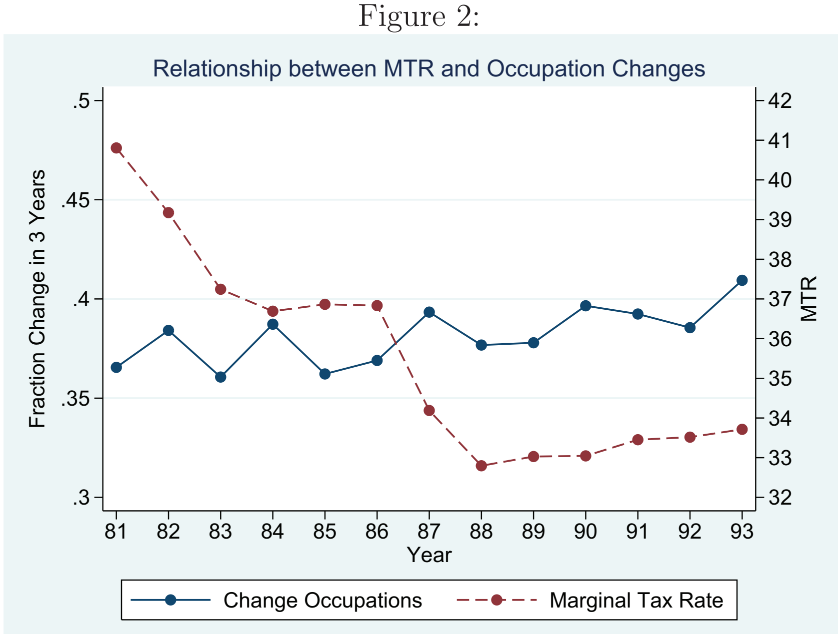

Table 1: Summary Statistics of the PSID Analysis Sample

	Mean	SD
Change Occupation within 3 Years (%)	40.7	49.1
Wage ($)	16.4	10.5
Age	35.8	7.2
Male (%)	53.4	49.9
Married (%)	75.1	43.2
Less than High School (%)	12.7	33.2
High School Graduates (%)	40.6	49.1
Some College (%)	23.2	42.2
College Graduates (%)	23.6	42.5
Total Income ($)	53,997	35,849
Marginal Tax Rate (%)	34.2	9.6
Tax Liability ($)	13,528	13,616

Note: N=44,062. Wage, total income, and tax liability are in 1997 dollars.

Table 2: Most Frequent Occupation (Industry) Changes

A. Most Frequent in Number of Observations

Original	New
Operatives (Manufacturing)	Craftsmen (Manufacturing)
Craftsmen (Manufacturing)	Operatives (Manufacturing)
Sales (Retail)	Managers (Retail)
Managers (Retail)	Sales (Retail)
Operatives (Manufacturing)	Unskilled Laborers (Manufacturing)
Unskilled Laborers (Manufacturing)	Operatives (Manufacturing)
Service Workers (Service)	Clerical (Service)
Clerical (Service)	Managers (Service)
Secretaries (Service)	Clerical (Service)
Service Workers (Service)	Other Medical (Service)

B. Most Frequent in Percentage of Workers in Original Occupation

Original	New
Unskilled Laborers (Manufacturing)	Operatives (Manufacturing)
Accountants (Finance)	Managers (Finance)
Sales (Services)	Managers (Services)
Sales (Manufacturing)	Sales (Retail)
Foremen (Construction)	Craftsmen (Manufacturing)

Note: Must have at least 40 observations in the original occupation to be considered in Panel B.

Table 3: Main Results

A. First Stage

	(1)	(2)
	$\Delta \ln(1-\tau)$	$\Delta \ln(z-T)$
MTR Instrument	0.463***	0.006**
	(0.019)	(0.002)
After-Tax Income Instrument	0.504***	0.882***
	(0.177)	(0.037)
Shea's R^2	0.0514	0.5727

B. OLS and IV Results

	(1)	(2)
	OLS	IV
$\Delta \ln(1-\tau)$	-0.016***	0.033**
	(0.003)	(0.015)
$\Delta \ln(z-T)$	0.072***	0.018
	(0.023)	(0.025)
Male	0.003**	0.003**
	(0.001)	(0.001)
White	0.004***	0.005***
	(0.001)	(0.001)
Married	0.002**	0.001
	(0.001)	(0.001)
Number of Dependents	-0.000	0.000
	(0.000)	(0.000)
High School Graduates	0.002	0.001
	(0.002)	(0.002)
Some College	0.006***	0.006***
	(0.002)	(0.002)
College Graduates	0.017***	0.016***
	(0.003)	(0.003)
Tenure in Old Job	0.003	0.002
	(0.002)	(0.002)
Tenure in Old Job Squared	-0.000	-0.000
	(0.001)	(0.001)
N	43,990	43,990

Note: The dependent variable is the difference in predicted compensating differentials between the old and new occupations at time t. Covariates included but not shown in this table are age group fixed effects and occupation-year fixed effects. Standard errors in parentheses are clustered by occupation and individual using the two-way clustering method. Significance levels: * 10%, ** 5%, *** 1%.

Table 4: Robustness Checks

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Main Results	Income Linear	Income Spline	Income* Year FE	Married* Year FE	Assume $\alpha_{it} = \alpha_i$	No Weights
$\Delta \ln(1-\tau)$	0.033**	0.033**	0.032**	0.038**	0.033**	0.032**	0.036**
	(0.015)	(0.016)	(0.016)	(0.017)	(0.015)	(0.015)	(0.018)
$\Delta \ln(z-T)$	0.018	0.018	0.018	0.034	0.015	0.027	0.007
	(0.025)	(0.024)	(0.024)	(0.028)	(0.025)	(0.025)	(0.035)
N	43,990	43,990	43,990	43,990	43,990	43,990	43,990

Note: The dependent variable is the difference in predicted compensating differentials between the old and new occupations at time t. Covariates included but not shown in this table are gender, race, education, job tenure, job tenure squared, number of dependents, marital status, age group fixed effects, and occupation-year fixed effects. Standard errors in parentheses are clustered by occupation and individual using the two-way clustering method. Significance levels: * 10%, ** 5%, *** 1%.

Table 5: Extensions

A. Heterogeneity by Gender

	(1) Male	(2) Female
$\Delta \ln(1-\tau)$	0.034*	0.054*
	(0.020)	(0.028)
$\Delta \ln(z-T)$	-0.172*	0.044
	(0.104)	(0.041)
N	23,229	19,996

B. Different Interval Length

	(1) 1-Year	(2) 2-Year	(3) 3-Year	(4) 4-Year
$\Delta \ln(1-\tau)$	-0.003	0.026*	0.033**	0.023**
	(0.012)	(0.014)	(0.015)	(0.011)
$\Delta \ln(z-T)$	-0.078	-0.065*	0.018	-0.006
	(0.056)	(0.039)	(0.025)	(0.027)
N	60,254	51,186	43,990	37,938

C. Heterogeneity by Age Group

	(1) Young	(2) Old
$\Delta \ln(1-\tau)$	0.021	0.044*
	(0.025)	(0.023)
$\Delta \ln(z-T)$	-0.019	-0.015
	(0.088)	(0.102)
N	21,034	22,159

Note: The dependent variable is the difference in predicted compensating differentials between the old and new occupations at time t. Covariates included but not shown in this table are gender, race, education, job tenure, job tenure squared, number of dependents, marital status, age group fixed effects, and occupation-year fixed effects. Standard errors in parentheses are clustered by occupation and individual using the two-way clustering method. Significance levels: * 10%, ** 5%, *** 1%.