

Taxation and the Earnings of Husbands and Wives: Evidence from Sweden

Alexander M. Gelber¹

The Wharton School, University of Pennsylvania, and NBER

September 2011

Abstract

This paper examines the response of husbands' and wives' earnings to a tax reform in which husbands' and wives' tax rates changed independently, allowing me to examine the effect of both spouses' incentives on each spouse's behavior. I compare the results to those of more simplified econometric models that are used in the typical setting in which such independent variation is not available. Using administrative panel data on approximately 11% of the married Swedish population, I analyze the impact of the large Swedish tax reform of 1990-1. I find that in response to a compensated fall in one spouse's tax rate, that spouse's earned income rises, and the other spouse's earned income also rises. A standard econometric specification, in which one spouse reacts to the other spouse's income as if it were unearned income, yields biased coefficient estimates. Uncompensated elasticities of earned income with respect to the fraction of income kept after taxes are over-estimated by a factor of more than three, and income effects are of the wrong sign. A second common specification, in which overall family income is related to the family's tax rate and income, also yields substantially over-estimated own compensated and uncompensated elasticities. Standard econometric approaches may substantially mis-estimate earnings responses to taxation.

¹Email: agelber@nber.org. For helpful guidance, I am grateful to Lawrence Katz, David Cutler, Martin Feldstein, Caroline Hoxby, Richard Freeman, Hanley Chiang, and three anonymous referees, as well as to others too numerous to name individually. Per-Anders Edin deserves special thanks for his generosity in helping to obtain data. I thank Håkan Björk, Hans Grönquist, Håkan Selin, and Marcus Vingren for graciously answering questions regarding the data. An earlier draft of this paper circulated under the title "Taxation and Family Labor Supply." NBER, the Uppsala Center for Fiscal Studies, Uppsala University, IFAU, and NIA Grant Number T32-AG00186 provided generous financial support. Part of this paper was completed while visiting at Uppsala University for a project with Per-Anders Edin. All errors are my own.

I. Introduction

Standard investigations of the impact of taxation on taxable income typically relate the family's overall taxable income to a measure of the family's tax rate.² In a parallel literature on the effect of tax rates on labor supply, it is standard to relate a spouse's labor supply decision to his or her own tax rate.³ It has been typical to assume that an individual's labor supply responds to the income of his or her spouse as it would respond to unearned income, following a long tradition beginning with Jacob Mincer (1962). In this paper, I relax these restrictions by examining how independent variation in both spouses' tax rates impacts each spouse's earnings decision. A richer econometric model allowing for such independent variation shows that the simplifications inherent in the standard approaches may lead to strongly biased results.

Swedish tax reforms in the early 1990s represent a particularly promising setting for studying these issues, for three primary reasons. First, Sweden has individual taxation, meaning that an individual's marginal tax rate on earned income depends only on his or her own income. When the Swedish tax schedule changes, husbands and wives face different changes in their marginal tax rates, and the relative size of these changes differs across households, allowing me to identify the response of one spouse's income to the other spouse's incentives.⁴ In the U.S., by contrast, married couples are almost always taxed jointly on the sum of their incomes, implying that husbands and wives face the same marginal tax rate. Second, I use the Longitudinal Individual Data for Sweden (LINDA), which contain information on the separate income of each spouse, unlike the IRS-Michigan-NBER Tax Panel on the U.S., which measures married couples' taxable income at the family level and does not allow investigation of how each spouse's income responds to incentives. The LINDA data are a panel of detailed administrative data on the labor force activity, government program participation, demographic characteristics, and other relevant features of approximately 11% of the Swedish population, which allows me to estimate parameters precisely, including cross responses.

²This large literature includes Lindsey (1987), Feldstein (1995), Navratil (1995), Auten and Carroll (1997), Sammartino and Weiner (1997), Saez (1999), Goolsbee (1999), Moffitt and Wilhelm (2000), Saez (2004), Kopczuk (2005), Giertz (2007), Saez (2010), and Singleton (2011). See Joel Slemrod, Saez, and Seth Giertz (2009) for a review.

³Examples include Jerry Hausman (1981), Nada Eissa (1995), and Richard Blundell, Alan Duncan, and Costas Meghir (1998).

⁴I use "own response," "own elasticity," or "own effect" to refer to the reaction to one's own tax rate or income, and "cross response," "cross elasticity," or "cross effect" to refer to the reaction to the tax rate or income of one's spouse.

Third, from 1989 to 1991, the top marginal income tax rate in Sweden decreased from 76% to 51%, with substantial but smaller decreases in other tax brackets. This represents an opportunity to examine labor supply responses to large exogenous changes in incentives.⁵

With a specification allowing for cross responses, I estimate a rich set of parameters, including own and cross income and substitution effects for both husbands and wives. Standard econometric models, in which one's own income or hours worked is assumed to respond to spousal income as it responds to unearned income, cannot reflect the possibility that the spouse's tax rate could have both income and substitution effects on one's earnings. The literature examining the response of families' taxable income to the marginal tax rates they face leaves open the question of how husbands' and wives' decisions separately contribute to families' aggregate responses, and implicitly assumes that the family can be treated as if it reacted to one (family) unearned income.

The results show that husbands and wives react to each other's marginal tax rates and unearned incomes, as well as to their own. My central estimates show compensated elasticities of individuals' earned income with respect to their own net-of-tax share of .25 and .49 for husbands and wives, respectively. (The net-of-tax share is defined as one minus the marginal tax rate.) It is noteworthy that the female earned income elasticity is higher than the elasticity for men, even in a country known for its relative gender equality and high female labor force participation rate. Compensated cross elasticities are .048 and .051, respectively, implying that in response to a compensated decrease in one spouse's tax rate, both spouses' earned incomes rise. Elasticities of earned income with respect to own unearned income are large (-.074 and -.056 for husbands and wives, respectively) and precisely estimated. I estimate elasticities of own earned income with respect to spouses' unearned income of -.0041 for husbands and -.018 for wives. When the dependent variable is a measure of taxable labor income, calculated by subtracting tax deductions from labor income, the elasticities are similar.

I find that standard econometric models, which make various simplifications, may yield very different results. I first assume that one spouse's earned income responds to the other spouse's

⁵Raj Chetty (2010) argues that large reforms give much more precise parameter estimates, and Chetty, John Friedman, Luigi Pistaferri, and Tore Olsen (2010) argue that large reforms may give estimates closer to long-run labor supply elasticities. Indeed, Chetty (2010) argues that because my paper examines a large reform in Sweden, I perform much more precise parameter estimates than other papers on income responses to taxation.

income as it reacts to unearned income. This yields an estimate of the income effect that is large and of the wrong sign. When a husband's marginal tax rate falls, he works more, and the estimates imply that his wife works more, as well. This induces a spurious positive correlation between the change in the measure of the husband's unearned income (which includes his wife's income) and the change in the husband's own earnings. Thus, the estimated coefficient on the husband's unearned income, which represents the income effect on his earnings, is overly positive.⁶ For both husbands and wives, this specification also produces an estimate of the uncompensated earnings elasticity that is biased upward by a factor of more than three, as well as an over-estimate of the compensated elasticity. A standard assumption from the parallel literature on the elasticity of taxable income is that the family's overall earnings reacts to the net-of-tax share and overall family unearned income. This specification likewise produces divergent results, with the compensated and uncompensated own elasticity substantially over-estimated and the uncompensated cross elasticity substantially under-estimated. Other specifications frequently run in the literature, including specifications omitting unearned income, also produce substantially biased results. Because these lessons are methodological—reflecting whether standard approaches produce correct results—I argue that the results are relevant to our understanding of these methods in contexts both inside and outside Sweden.

Since I estimate own and cross uncompensated and compensated effects, I am able to perform two separate tests of a unitary model of family taxpaying decisions. The unitary model is defined by the feature that the family can be characterized as maximizing a single utility function. I reject a unitary model based on violations of the "income pooling condition," which implies in this context that a married individual's pre-tax earnings should react equally to an increase in that individual's unearned income as it reacts to an increase in the unearned income of his or her spouse.⁷ The unitary model also predicts that the Slutsky matrix should be symmetric: the compensated response of the husband's pre-tax earnings to the net-of-tax rate of the wife is predicted to be equal to the compensated response of the wife's pre-tax earnings to the net-of-tax rate of the husband.

⁶Analogous reasoning implies that the estimated coefficient on the wife's unearned income should be overly positive.

⁷More generally, the income pooling condition states that a married individual's consumption of any good should react equally to an increase in that individual's unearned income as it reacts to an increase in the unearned income of his or her spouse. See Shelly Lundberg, Robert Pollak, and Terence Wales (1997) for an alternative test of income pooling in a developed country.

I cannot reject Slutsky symmetry at conventional significance levels.

The paper proceeds as follows. Section II discusses the empirical specification. Section III describes the policy environment and other relevant features of the Swedish economy in the period under consideration. Section IV describes the data. Section V presents the empirical results and relates them to a unitary model of family taxpaying decisions. Section VI concludes.

II. Empirical Model

A. Basic Framework

The primary goal of the paper is to estimate the causal effect of husbands' and wives' net-of-tax rates and unearned incomes on their earnings. Following Gruber and Saez (2002) and the literature cited therein, my empirical model relates the change in an individual's log real pre-tax earnings to the change in the individual's log net-of-tax rate and the change in the individual's log real after-tax unearned income. Relating the logs of the variables yields coefficient estimates interpretable as elasticities. In addition, as Gruber and Saez write (p. 10), "For large tax changes, it is perhaps more natural to use a log-log specification that is also closer to previous studies' specifications." It is illuminating to follow a specification that has been estimated in previous literature, as the results under alternative specifications can then be compared in a straightforward way to what has been estimated in previous work.

I extend the Gruber and Saez (2002) specification to allow each spouse's earnings to depend not only on his or her own tax rate and unearned income, but also on the tax rate and unearned income of the other spouse. This represents a natural extension, since one spouse's tax rate and unearned income are likely to influence the earnings of the other spouse if spouses' decisions interact. The log of a spouse's earned income, $\ln(E_{it}^s)$, is specified as a function of the log of that individual's net-of-tax share (i.e. the log of one minus that individual's marginal tax rate), $\ln(1 - \tau_{it}^s)$, the log of the other spouse's net-of-tax share, $\ln(1 - \tau_{it}^{-s})$, the log of the individual's own income, $\ln(Y_{it}^s)$, and the log of the other spouse's income, $\ln(Y_{it}^{-s})$. Here the superscript $s \in \{h, w\}$ represents the individual in question, whereas $-s$ denotes that individual's spouse, and h and w refer to the husband and wife, respectively. i indexes couples, and t represents the time period.

To remove individual-level fixed effects that may be correlated with the tax and income variables of interest, the model will be estimated in first differences, again following Gruber and Saez (2002):

$$\begin{aligned}\Delta \ln(E_{it}^h) &= \beta_0^h + \beta_1^h \Delta \ln(1 - \tau_{it}^h) + \beta_2^h \Delta \ln(1 - \tau_{it}^w) + \beta_3^h \Delta \ln(Y_{it}^h) + \beta_4^h \Delta \ln(Y_{it}^w) \\ &\quad + X_{iT}^h \beta_h^h + X_{iT}^w \beta_w^h + \vartheta_t^h + \varepsilon_{it}^h\end{aligned}\tag{1}$$

$$\begin{aligned}\Delta \ln(E_{it}^w) &= \beta_0^w + \beta_1^w \Delta \ln(1 - \tau_{it}^w) + \beta_2^w \Delta \ln(1 - \tau_{it}^h) + \beta_3^w \Delta \ln(Y_{it}^w) + \beta_4^w \Delta \ln(Y_{it}^h) \\ &\quad + X_{iT}^w \beta_w^w + X_{iT}^h \beta_h^w + \vartheta_t^w + \varepsilon_{it}^w\end{aligned}\tag{2}$$

where $\Delta \ln(Z_t)$ represents the change from $t - 1$ to t in the log of Z . (I use “base year” to indicate $t - 1$, the initial year in each pair of years over which the first difference is taken, and “final year” to refer to t , the last year in each pair of years over which the first difference is taken.) The subscript t still appears in the empirical model since multiple first differences will be used. Time dummies ϑ_t^s control for economy-wide earned income growth specific to each period over which the first difference is taken. ε_{it}^h and ε_{it}^w are error terms. X_{iT}^h and X_{iT}^w represent other variables—age, age squared, education, region, number of children, industry, occupation, and sometimes interactions of the covariates—that control for other factors that could influence changes in earned income. The control variables bear the subscript T , which refers to an initial period prior to the earliest observation of t .

B. Discussion of Specification

The dependent variable, the change in the log of real earned income, may best reflect the welfare consequences of taxation among the variables commonly observed in tax datasets, as discussed in Chetty (2009).⁸ The response of earned income is straightforward to examine because in the data, the definitions of several types of capital income changed from before to after the Tax Reform of 1991. In some regressions, I also examine how a measure of taxable labor income, formed by subtracting a set of deductions from earned income, responds to the net-of-tax share.⁹

⁸I also follow the suggestion of Chetty (2009) and explore the response of measures of non-wage compensation to the tax rate.

⁹Feldstein (1999) develops a measure of the deadweight loss of taxation in terms of the elasticity of taxable labor income with respect to the net-of-tax share, but the empirical literature has focused on the elasticity of taxable

The specification relates earnings to the net-of-tax share, which is the accepted standard specification in the tax literature. Indeed, the literature on taxable labor income responses to taxation has essentially exclusively adopted this specification, including Lindsey (1987), Feldstein (1995), Navratil (1995), Auten and Carroll (1997), Sammartino and Weiner (1997), Saez (1999), Goolsbee (1999), Moffitt and Wilhelm (2000), Saez (2004), Kopczuk (2005), Hansson (2007), Giertz (2007), Holmlund and Söderstrom (2008), Heim (2009), Saez (2010), Chetty, Friedman, Olsen and Pistaferri (2011), and Singleton (2011).¹⁰ However, it is important to note that in all of this literature, including in my paper, the incidence of a tax cut could also affect the estimates. Standard economic theory predicts that the incidence of a tax cut should be shared between employees and employers. In particular, in a framework in which workers are offered an hourly wage and choose their hours worked, a tax cut should lead to a fall in the pre-tax wage (assuming that labor supply is not completely inelastic and labor demand is not infinitely elastic) and a rise in the after-tax wage. Changes in earnings will therefore reflect both changes in hours worked and changes in the pre-tax hourly wage. Differences in earnings elasticities across groups, such as men and women, could reflect both differences in the labor supply or demand elasticities across these groups.¹¹ I follow the theoretical and empirical framework in Gruber and Saez (2002) and the literature cited above in abstracting from these issues by relating earnings to net-of-tax shares and virtual incomes. This is often defended on the grounds that empirical work tends to find no evidence of an impact of taxation on the pre-tax wage, and finds instead that taxes affect the after-tax wage, consistent with a model in which labor supply is much more inelastic than labor demand (e.g. Anderson and Meyer 2000). Feldstein (1995) and Blomquist and Selin (2009) add a further twist to this literature, arguing that the pre-tax wage may change because individuals put forth more effort in response to tax cuts.

The specification adopted in all of the literature cited above, as well as in my paper, implicitly assumes that changes in hourly wages are uncorrelated with changes in taxes (since changes in hourly wages could affect earnings). If they were correlated, this could lead to biased and inconsistent coefficients on the tax rate. However, I am able to address this limitation further than

income (including capital income) with respect to the net-of-tax share. By investigating the elasticity of taxable labor income, I estimate a parameter that more closely corresponds to Feldstein's (1999) model.

¹⁰Saez, Slemrod, and Giertz (2009) review papers that rely on this specification.

¹¹In a model in which employers and employees bargain, the incidence of a tax cut could also reflect employers' and employees' bargaining power.

much of this previous literature has, because I have information on various demographics that are likely to be correlated with hourly wage changes over time. In particular, I am able to control for factors such as age, education, region, number of children, two-digit industry, and two-digit occupation. In the data used in most previous literature, most of these demographic variables have not been available, including occupation and industry. The additional controls prove to make little difference to the results.

When observations of earnings E are zero, the log of earnings is undefined. This implies that when earnings are zero, with the log-log specification, observations in which earnings are zero in either the base year or the final year would be set to missing. This would in turn imply that the sample would be selected according to values of the dependent variable: only observations of the dependent variable for which earnings are non-zero in both the base year and the final year would be included in the regressions. Selecting the sample according to values of the dependent variable in general leads to biased and inconsistent estimates of coefficients. Thus, I find a way of including these observations in the regressions by adding one to earnings and then logging earnings. This implies that zeroes of earnings are not treated as missing.¹² The dependent variable for spouse s in couple i is therefore $\ln[(1 + E_{it}^s)/(1 + E_{it-1}^s)]$, and the notation in (1) and (2) can be considered shorthand for this expression. As I show in the Appendix, the results are generally insensitive to other choices, such as adding .5 before taking the log. In order to address the concern that logging the dependent variable could have an effect on the results particularly at low incomes, in one specification, I exclude from the regressions those who exited the labor force. When I log earnings in the Appendix table (without adding anything to earnings before logging earnings), and the sample is therefore limited to those with positive earnings in both periods, the results are again similar to the basic results. Because the prediction of Slutsky symmetry of the unitary model of family responses to taxation (discussed further below) only holds when both spouses participate in the labor market, I exclude couples from my main regressions in those pairs of years in which at least one member of the couple does not participate in the labor market in the base year. The measure of income used as an independent variable is “virtual income,” which represents the intersection of the individual’s extended budget segment in consumption-effort space with the

¹²Other literature has adopted this technique for including zeroes in the estimates (e.g. Kuziemko and Werker 2006).

Y-axis. The construction of virtual income is discussed at greater length in Appendix I.¹³

In the main regressions, I consider two sets of one-year differences, which are pooled in the regressions: one from 1989-1990, and the other from 1990-1991. These are the years of the tax reform. This strategy will identify a short-term effect of the changes in the tax schedule.¹⁴ The main source of exogenous variation is that in TR91, marginal tax rates were reduced much more for those at the top of the income distribution than for those at the bottom. This generates very large exogenous variation across households and time in the net-of-tax shares of husbands relative to their wives. For example, suppose that in Couple 1, the wife is in the lowest tax bracket, and the husband is in the highest tax bracket (both before and after the reform). In Couple 2, both the husband and wife are in the highest tax bracket (both before and after the reform). Those in the highest tax bracket receive a large cut in their marginal tax rate, whereas those in the lowest bracket receive a small cut. Therefore, due to the tax reform, the net-of-tax share of the husband relative to that of the wife increases in Couple 1 but stays constant in Couple 2. Thus, I can effectively relate the changes over time in the relative earnings of the husbands and wives in the two couples, to the changes over time and couples in their relative net-of-tax shares (and to the changes in virtual incomes associated with these tax changes and any simultaneous changes in capital taxation).¹⁵

C. Instruments

The actual marginal tax rate that an individual faces is potentially endogenous. For example, if an individual responds to an increase in his or her own marginal tax rate by decreasing his or her earned income, and marginal tax rates are progressive, then an OLS estimate of the effect of the net-of-tax share on earned income will be biased downward. Thus, it is typical to instrument for the net-of-tax share with a so-called “simulated instrument.” This instrument is constructed by calculating the change in the net-of-tax share that would have occurred if the individual had maintained the behavior he or she exhibited in the initial period (Gruber and Saez 2002). The

¹³Gary Burtless and Hausman (1978) explain virtual income and why it is the appropriate income measure for estimating income effects in the presence of a nonlinear budget set.

¹⁴Gruber and Saez (2002) find relatively similar elasticities at 1-year, 2-year, and 3-year intervals.

¹⁵My regressions in fact allow for more flexibility than a specification that literally related the relative earnings of the spouses to their relative net-of-tax shares, because I run separate regressions for husbands and wives and enter each spouse’s net-of-tax share separately in each regression.

intuitive notion that underlies this procedure is that the change in the tax schedule is exogenous to individuals' initial behavior, so the value of this instrument will not be affected by the endogenous response to the new tax schedule.

In particular, the instrument is constructed by projecting final year taxable income to be base year taxable income for spouse s in couple i , Z_{it-1}^s , multiplied by the growth of mean taxable income per taxpayer in the sample, $(1 + g)$. Letting \hat{Z}_{it}^s be projected taxable income, I set $\hat{Z}_{it}^s = (1 + g)Z_{it-1}^s$. Suppose that the net-of-tax share (as a function of taxable income) before the tax change is given by $T_{t-1}(\cdot)$ and the net-of-tax share after the tax change is given by $T_t(\cdot)$. I use $T_t(\hat{Z}_{it}^s) - T_{t-1}(Z_{it-1}^s)$ to instrument for $T_t(Z_{it}^s) - T_{t-1}(Z_{it-1}^s)$. In the regressions relating to the extensive margin, the average after-tax share is instrumented analogously. Because virtual income for spouse s in couple i in year t , $Y_{it}^{s,v}(\cdot)$, varies according to which budget segment the individual locates on, it is a function of actual taxable income.¹⁶ Thus, virtual income is also potentially endogenous. I construct a simulated instrument for the actual change in virtual income, by predicting the change in virtual income that would have occurred, if the individual had projected taxable income \hat{Z}_{it}^s in the final period. In other words, I use $Y_{it}^{s,v}(\hat{Z}_{it}^s) - Y_{it-1}^{s,v}(Z_{it-1}^s)$ as an instrument for $Y_{it}^{s,v}(Z_{it}^s) - Y_{it-1}^{s,v}(Z_{it-1}^s)$.¹⁷

Following previous literature (e.g. Hausman 1981; Blomquist and Hansson-Brusewitz 1990; Blomquist and Selin 2009), my empirical specification effectively assumes that the response to actual income shocks and virtual income shocks is the same. Variation in virtual income therefore comes from three sources. First, virtual income is influenced by changes in actual pre-tax unearned income. Second, virtual income is influenced by changes in the average tax rate on capital income, since this influences after-tax unearned income. Third, virtual income is influenced by changes in the marginal tax rate on labor income.

¹⁶ $Y_{it}^{s,v}(\cdot)$ is subscripted by i because it also depends on capital income and government transfers, which vary by individual.

¹⁷Since each spouse's tax rate on capital income was potentially different prior to the reform (because each spouse's capital income was taxed separately), this created an incentive for couples to avoid taxes by allocating capital income to the lower-taxed spouse. However, this does not affect my estimates because I instrument for the actual change in virtual income using the change that would have been expected on the basis of the different components of pre-reform virtual income. The estimation procedure therefore effectively throws away any variation coming from individuals' endogenous responses to the new tax schedule, and therefore throws away any variation relating to re-allocation of capital income. It is also worth noting that capital income has been taxed at a flat rate of 30% since the 1991 reform, thus eliminating any incentive for couples to re-allocate their assets to the lower-taxed spouse.

However, variation in *predicted* changes in virtual income, which serves as the instrument for virtual income, comes from only two sources. First, predicted virtual income is influenced by changes in the tax rate on capital income, since this influences after-tax predicted unearned income. Changes in tax policy toward capital income (discussed below) therefore drive variation in predicted virtual income. Second, predicted virtual income is influenced by changes in policy affecting the marginal tax rate on labor income, because this affects the slope of the budget segment in question and therefore its intersection with the y-axis. It is important to note that these changes in the marginal tax rate on labor income also have an effect on earnings that operates through the resulting change in the net-of-tax share on labor income, which in the empirical model has a direct effect on earnings. As in previous literature (e.g. Hausman 1981), changes in predicted virtual income arising from changes in the marginal tax rate on labor income are therefore correlated with changes in the net-of-tax share on labor income. As in this previous literature, therefore, to the extent that variation in predicted virtual income is driven by changes in the marginal tax rate on labor income, the coefficient on (instrumented) virtual income will be separately identified from the coefficient on the (instrumented) net-of-tax share owing to functional form restrictions.

However, it is important to note that since changes in predicted virtual income are also influenced by changes in capital income tax policy, part of the variation in predicted virtual income is not mechanically related to variation in the predicted net-of-tax share on labor income. In addition, it is worth noting that percentage changes in predicted virtual income are not influenced by changes in actual income, since predicted virtual income is calculated using lagged virtual income.

D. Controlling for the Evolution of the Income Distribution

In their regressions relating taxable income to the net-of-tax share and an income effect, Gruber and Saez (2002) control for a ten-piece spline in the log of base year real income. Since the size of the tax change is correlated with income, it may be difficult empirically to tease apart variation in base-year income from variation in the change in marginal tax rates. Indeed, Gruber and Saez (2002) write that using rich controls for base-year income “may destroy identification. This problem is especially acute when the size of the tax rate change is directly correlated with the income level as in the TRA of 1986...In practice, rich controls for base year income make it very difficult to separately identify income and substitution effects with only one tax change.”

(pp. 11-12). Because I examine only one tax reform, over-controlling for base year income is a major cause for concern. Given the correlation between base year income and the change in the marginal tax rate, the regression results may be highly sensitive to mis-specification, for example of the functional form with which base year income enters.

To address this issue, I calibrate the evolution of the income distribution using a period in which no major tax change occurs, and I assume that absent the tax change, the income distribution would have evolved similarly during the period of the change. I then relate the remaining variation in earned income to exogenous variation in the marginal tax rate, controlling for a rich set of covariates that can capture effects unique to the period of the tax change. I begin this procedure by performing the following regression during a period in which the change in the tax schedule is negligible:

$$\begin{aligned} \Delta \ln(E_{it}^s) = & \xi_0^s + f[\ln(E_{it-1}^s)]\xi_{E,s}^s + f[\ln(E_{it-1}^{-s})]\xi_{E,-s}^s + f[\ln(Z_{it-1}^s)]\xi_{Z,s}^s \\ & + f[\ln(Z_{it-1}^{-s})]\xi_{Z,-s}^s + X_{iT}^s \xi_s^s + X_{iT}^{-s} \xi_{-s}^s + v_{it}^s \end{aligned} \quad (3)$$

Here f is a ten-piece spline in lagged log real income, s denotes the spouse in question, and $-s$ denotes the other spouse. I use ten-piece splines in one's own lagged log real earned income, one's spouse's lagged log real earned income, one's own lagged log real taxable income, and one's spouse's lagged log real taxable income. I include a ten-piece spline in lagged log real taxable income because in the main regressions of interest, changes in log real earned income will be related to changes in marginal tax rates. Marginal tax rates are computed based on taxable income, so controlling for lagged log real taxable income addresses possible mean reversion relating to taxable income. The knots of the spline are placed at deciles of the income distribution. $\xi_{E,h}^h$, $\xi_{E,w}^h$, $\xi_{Z,h}^h$, $\xi_{Z,w}^h$, $\xi_{E,w}^w$, $\xi_{E,h}^w$, $\xi_{Z,w}^w$, and $\xi_{Z,h}^w$ represent vectors of coefficients on these splines.

These regressions yield an estimated set of coefficients $\hat{\xi}_{E,h}^h$, $\hat{\xi}_{E,w}^h$, $\hat{\xi}_{Z,h}^h$, $\hat{\xi}_{Z,w}^h$, $\hat{\xi}_{E,w}^w$, $\hat{\xi}_{E,h}^w$, $\hat{\xi}_{Z,w}^w$, and $\hat{\xi}_{Z,h}^w$, which collectively calibrate how income evolves in the absence of a tax change. In the later period that spans the tax change, I use these estimated coefficients to partial out the predicted effect of base year income, thus creating residual changes in the log of real earned income, $\tilde{\Delta} \ln(E_{it}^s)$,

for the wife and the husband:

$$\begin{aligned} \tilde{\Delta} \ln(E_{it}^s) &= \Delta \ln(E_{it}^s) - f[\ln(E_{it-1}^s)] \hat{\xi}_{E,s}^s - f[\ln(E_{it-1}^{-s})] \hat{\xi}_{E,-s}^s \\ &\quad - f[\ln(Z_{it-1}^s)] \hat{\xi}_{Z,s}^s - f[\ln(Z_{it-1}^{-s})] \hat{\xi}_{Z,-s}^s \end{aligned} \quad (4)$$

These residuals represent the remaining variation in the change in earned income, with the predicted effect of lagged income removed. I now modify equations (1) and (2), relating the residuals to the independent variables:

$$\begin{aligned} \tilde{\Delta} \ln(E_{it}^s) &= \beta_0^s + \beta_1^s \Delta \ln(1 - \tau_{it}^s) + \beta_2^s \Delta \ln(1 - \tau_{it}^{-s}) \\ &\quad + \beta_3^s \Delta \ln(Y_{it}^s) + \beta_4^s \Delta \ln(Y_{it}^{-s}) + X_{iT}^s \beta_s^s + X_{iT}^{-s} \beta_{-s}^s + \vartheta_t^s + \varepsilon_{it}^s \end{aligned} \quad (5)$$

I instrument for tax rates and virtual incomes using the simulated instruments described earlier.

The procedure described in this section is conceptually similar to a “triple difference” strategy, in which the differences across couples over time are contrasted between a period of no policy change and a period of a policy change. The assumption is that the influence of all of the factors that are unique to the period spanning the tax change can be removed with the controls. I control extensively for occupation, industry, region, education, and several other demographic variables. The evidence is consistent with the contention that this procedure removes the true effect of lagged income and business cycle effects, since adding more extensive controls makes little difference to the estimated coefficients of interest.¹⁸

E. Implications of the Unitary Model for the Parameter Estimates

The unitary model of family decision-making is defined by the feature that the family’s behavior can be characterized as maximizing a single utility function.¹⁹ This yields two central predictions

¹⁸My procedure also bears a conceptual resemblance to the empirical strategy of Lawrence Lindsey (1987). Lindsey predicts how much taxable income should exist in each part of the income distribution, absent the tax change. The difference between the actual amount of taxable income in each part of the distribution and the predicted amount is then attributed to the effect of taxation. My procedure performs a similar comparison, but differs from the Lindsey strategy by employing panel data, rather than repeated cross sections.

¹⁹It is important to emphasize that the paper’s empirical specification is not explicitly derived from the unitary model (whose validity is in fact called into question by the empirical results). The empirical specification instead builds on a long tradition of other work in the taxable income literature that has adopted similar specifications, as cited above.

(Blundell and MaCurdy 1999). First, income pooling: the husband’s (wife’s) earnings should react equally to a change in his (her) own unearned income as to a change in his wife’s (her husband’s) unearned income. This condition holds because the family does not distinguish between the unearned income of the husband and the unearned income of the wife in making its labor effort and consumption decisions; rather, the household acts as a single agent that pools its unearned income and reacts to it the same whether its source is the husband or wife. Second, Slutsky symmetry: The compensated response of the husband’s earnings to the wife’s net-of-tax wage should be equal to the compensated response of the wife’s earnings to the husband’s net-of-tax wage. Standard consumer demand theory implies this condition. The family has a single utility function, so the Slutsky matrix for the consumption of goods, including the effort of the husband and wife, must be symmetric about the diagonal.²⁰ Since the empirical model is specified in terms of elasticities, I transform the coefficient estimates to relate them to the predictions of the unitary model. For individuals at the sample means of income, income pooling implies:

$$\beta_3^h \bar{Y}^w = \beta_4^h \bar{Y}^h \tag{6}$$

and

$$\beta_3^w \bar{Y}^h = \beta_4^w \bar{Y}^w \tag{7}$$

where bars above the income variables represent their sample mean values.²¹

To test Slutsky symmetry, begin by recalling the Slutsky relation in the context of earned

²⁰The traditional unitary model applies in the context of hours worked responses to net-of-tax wages. The Slutsky symmetry prediction stated above—that the compensated response of the husband’s earnings to the wife’s net-of-tax wage should be equal to the compensated response of the wife’s earnings to the husband’s net-of-tax wage—follows from the Slutsky symmetry of the traditional unitary model if in the empirical estimation, changes in hourly wages across individuals are uncorrelated with changes in net-of-tax shares. If this is true, then the elasticity of hours worked with respect to the net-of-tax hourly wage will be the same as the elasticity of earnings with respect to the net-of-tax share. Similarly, if changes in hourly wages are uncorrelated with changes in virtual unearned income, then the estimated elasticity of hours worked with respect to unearned income will be the same as the estimated elasticity of earnings with respect to unearned income. In this case, all of the predictions of the traditional unitary model will carry over to the context of the relationship between earnings and the net-of-tax share. This prediction also follows from a framework similar to one adopted in Gruber and Saez (2002) or Saez (2010), in which individuals derive utility from pre-tax earnings as well as (after-tax) consumption. If family utility is defined over the pre-tax earnings of the husband and wife, as well as the family’s after-tax consumption, and regularity conditions are satisfied, then Slutsky symmetry and income pooling, as explicated below, should hold.

²¹In my tests, I use the sample mean values from 1989, before the tax change. I also test these predictions for individuals at other points in the income distribution and find similar results.

income:

$$\frac{\partial E^h}{\partial(1 - \tau^w)} \Big|_u = \frac{\partial E^h}{\partial(1 - \tau^w)} - E^w \frac{\partial E^h}{\partial Y^w} \quad (8)$$

$$\frac{\partial E^w}{\partial(1 - \tau^h)} \Big|_u = \frac{\partial E^w}{\partial(1 - \tau^h)} - E^h \frac{\partial E^w}{\partial Y^h} \quad (9)$$

After performing transformations to express the elasticity estimates as marginal effects, the following equality is implied by Slutsky symmetry, evaluated at the sample means of the variables:

$$\beta_2^h \frac{\bar{E}^h}{1 - \bar{\tau}^w} - \bar{E}^w \beta_4^h \frac{\bar{E}^h}{\bar{Y}^w} = \beta_2^w \frac{\bar{E}^w}{1 - \bar{\tau}^h} - \bar{E}^h \beta_4^w \frac{\bar{E}^w}{\bar{Y}^h} \quad (10)$$

III. The Tax Reform of 1991

The Tax Reform of 1991 changed income tax rates dramatically, as the top marginal tax rate fell from 76% to 51%.²² TR91 revised several other aspects of the tax system, including the VAT and corporate taxes. The period considered in this paper includes two tax reductions, from 1989-1990 and 1990-1991, the latter of which was substantially larger. Table 1 shows the tax schedule for the national Swedish government, called the "state tax schedule," in 1989 and 1991. Marginal tax rates fell substantially for those at the top of the income distribution but fell little for those at the bottom. Before TR91, the state tax schedule was comprised of two different schedules, the basic schedule and the additional schedule. Basic taxable income differed from additional taxable income because a number of deductions could be taken on the basic schedule that could not be taken on the additional schedule. The total state marginal tax rate was calculated by summing the basic marginal tax rate and the additional marginal tax rate. Starting in 1991, the distinction between basic and additional taxable income was eliminated, and income was taxed according to a single state tax schedule.

Prior to 1991, Sweden had a global tax system, under which earned income and capital income were taxed at the same marginal tax rate, calculated on the basis of an individual's earned income,

²²A detailed description and analysis of TR91 can be found in Jonas Agell, Peter Englund, and Jan Södersten (1998). This section and the next also often draw on the description of TR91 in Martin Ljunge and Kelly Ragan (2005). Many important features of the reform had been anticipated since 1987, when a commission began to plan the reform. Åsa Hansson (2007) also examines the Swedish Tax Reform of 1991, focusing on the response of individuals' earned income to taxation and assuming that one spouse reacts to the other's income as if it were unearned income. Bertil Holmlund and Martin Söderstrom (2008) examine reforms following the 1991 reform.

taxable government transfers, capital income, and deductions. Starting in 1991, Sweden changed to a dual tax system, under which the marginal tax rate on earned income is computed only based on earned income (and taxable government transfers and deductions), and capital income is taxed at the flat rate of 30%. These changes in the taxation of unearned income provide sizeable exogenous variation in after-tax unearned income, thus aiding in the identification of income effects. In the Swedish system, each spouse has their own separate assets and capital income. The simulated instrument approach described above isolates the effect of a labor income tax policy reform by calculating the change in the net-of-tax share that would have occurred if the individual had maintained the behavior he or she exhibited in the initial period. In much the same way, the simulated instrument for each spouse's change in after-tax unearned (virtual) income will isolate the effect of the changes in capital tax policy on the earnings of each spouse by calculating the change in after-tax unearned income that would have occurred if the individual had maintained the behavior he or she exhibited in the initial period, thereby effectively treating each spouse's pre-reform capital income as given.

The reform also broadened the tax base, to make up for the revenue lost due to the tax cuts. For example, before 1991, nominal interest expenses were fully deductible against typically high marginal income tax rates, whereas after the reform, they were deductible against the lower capital income tax rate of 30%. Due to such broadening of the base, deductions and exclusions fell as a share of total income. The reform was designed to be almost revenue-neutral. The Swedish Ministry of Finance (1991) projected that 89.1 billion Swedish Kronor (SEK) would be lost due to the tax cuts, and that SEK 8.2 billion would be lost due to increased spending planned for 1991. However, the projections indicated that SEK 95.1 billion would be recouped through the combination of base broadening (SEK 79.6 billion), dynamic gains from increased economic activity in response to the tax cuts (SEK 5.0 billion), and increases in other revenues (SEK 10.4 billion).

The total marginal tax rate is calculated as the sum of the local, municipal and state tax rates. The mean value of the sum of local and municipal rates is 31% (both before and after the reform), with a minimum of 27% and a maximum of 33% over all the years examined. It is possible to construct an alternative measure of the marginal tax rate that includes the phase-outs and phase-ins of the basic deduction and of various transfers (such as a housing-related transfer).

Ultimately, how much individuals respond to such incentives is an empirical question, and the results are similar when other measures of the marginal tax rate are employed.

I use the evolution of the income distribution from 1988 to 1989 to calibrate how the income distribution develops. I use these years because they are temporally adjacent to the TR91 changes in 1990-1991, and thus it is likely that the nature of the evolution of the income distribution did not change much by the 1989-1991 period; because there were no other substantial policy changes in 1988-1989; and because these years are before TR91, so individuals are not still reacting to the policy changes. 1986 or 1987 could in principle also be candidates for years on which to estimate the evolution of the income distribution, even though they are further removed from 1990-1991. However, in 1986 and 1987, the measure of income in the LINDA data includes *both* earned income and government transfers, so there is no way to separately measure earnings without including government transfers. Sweden has large government transfers (such as unemployment compensation, sick leave, parental leave, or welfare), and government transfers comprise a large fraction of income. The period immediately following TR91 also has important limitations because of a number of major policies likely to strongly affect labor force participation were enacted in 1992. Under Sweden's 1992 Equal Opportunity Act, employers must try to obtain a well-balanced sex distribution in various jobs and must make it easier for workers to combine work and family (Blau, Ferber, and Winkler 2005; Melkas and Anker 1998). In addition, in 1992 important legislation was enacted that cut funding for public daycare (Mahon 1997). Both of these policy shifts are likely to have substantially affected labor force participation differentially at different income levels. Finally, individuals may still have been reacting to the major changes in tax policy in the period following 1991, as adjustment costs may prevent immediate adjustment to policy changes (e.g. Holmlund and Söderstrom 2008 on Sweden or Chetty 2010). However, it is important to note that I obtain results in the same range when I also use 1991-2 to calibrate the evolution of the income distribution, rather than 1988-1989.²³

Like most empirical estimates, the estimates must be interpreted as local in the sense that they

²³Data from far before or after the period of TR91 are less likely to reflect the same earnings dynamics, as the autoregressive process characterizing the evolution of earnings may change substantially over time (e.g. Gottschalk and Moffitt 1994). Data from prior to 1986 are also likely to be confounded by other factors. A series of large tax cuts over the period from 1981 to 1986 are likely to have affected the evolution of the income distribution in ways that are not replicated in other years. The 1995 tax increase for upper-income individuals, as well as the 1994 changes in capital income taxation, are also likely to confound comparisons in later years.

may not generalize outside of the particular context studied. For example, although tax changes throughout the income distribution are included in the estimation, the tax cuts were largest near the top of the income distribution, and thus a large portion of the useful variation comes from tax cuts in this part of the income distribution. As another example, these estimates are largely based on tax reductions.²⁴ It is possible that tax reductions have a different impact on earnings than tax increases do, as it may be easier or more difficult to increase labor supply than to decrease it for a variety of reasons (see, for example, the discussion in Blomquist and Selin 2009).

Some features of the Swedish macroeconomic environment are shown in Figure 1. The dashed line represents real percentage GDP growth per capita. Sweden entered a recession in late 1990, with real per capita GDP growth rates of 1.0% and -1.1% in 1990 and 1991, respectively. The solid line shows the unemployment rate, which increased substantially during the recession.²⁵ To control for the influence of these macroeconomic factors, I control for a rich set of covariates, including dummies for 2-digit industry codes, 2-digit occupation codes, and other covariates interacted with year. Income effects could also have come not only from changes in capital income measured in the data, but also by the changes in wealth induced by the macroeconomic environment or by the capitalization of changes in the tax rules into asset prices. A particular source of concern is that housing prices in Sweden fell substantially around the time of the reform. When I include a measure of the imputed income from housing wealth in my measure of virtual income, or when I control for local housing price growth, I estimate similar income effects.

To understand the context in which the tax reform occurred, it is also worth noting various relevant features of the Swedish environment.²⁶ Completed fertility of the 1961 birth cohort is 2.03. The percentage of the population currently divorced in 2003 was 11.3%. 75.6% of Swedish women aged 15-64 participated in the labor market in 2002, and the male labor force participation rate was 79.4%. Swedish GDP per capita in 1990 was \$27,240. Finally, a relatively large fraction

²⁴The estimates are not entirely based on tax reductions because the tax base also changed in TR91, which resulted in a rise in the marginal (or average) tax rate for some individuals.

²⁵It is possible to argue that this macroeconomic turmoil could help me to uncover family labor supply responses. During a period of economic calm, couples may re-optimize their decisions infrequently, but in a period of turmoil, we may be able to observe these changes more readily and relate them to exogenous changes in tax policy. On the other hand, one could argue that in a weak labor market, families may not have labor supply choices available to them that they otherwise would have made. The overall impact on the parameter estimates is *a priori* unclear.

²⁶The statistics in this paragraph are drawn from Blau, Ferber, and Winkler (2005) and Betsey Stevenson and Justin Wolfers (2007).

of couples in Sweden cohabit rather than formally marrying, and the percentage married was only 45.2% in 2003. The sample of married Swedes is thus selected in certain ways, but it is *a priori* unclear whether and how this should affect the parameter estimates.

IV. The LINDA Data

I use the Longitudinal Individual Dataset for Sweden (LINDA), described in detail in Per-Anders Edin and Peter Fredriksson (2000). Based on the administrative records of the Swedish government, these data follow individuals and their families longitudinally. I examine yearly data from 1988 to 1991, inclusive. The data contain approximately 3.35% of the Swedish population, in addition to family members of these individuals. A random sample of 20% of the immigrants to Sweden and their families is also included. Prior to restricting the sample, the full data consist of approximately 950,000 sampled individuals per year, comprising approximately 11% of the Swedish population. After weighting by sample weights to correct for the over-sampling of immigrants, the full sample is cross-sectionally representative of the married Swedish population in any given year.

Gender, age, region of residence, occupation, industry, number of children, educational attainment, and other covariates are included in the data. In the regressions, the values of all of these control variables are taken from 1988. Most of these covariates are not available in the U.S. administrative data on tax returns, including the IRS-Michigan-NBER tax panel. My measure of earned income includes only wages paid from employers to employees (and excludes government transfers). I construct taxable labor income by subtracting certain deductions from earned income. During the period under consideration, the data do not contain a measure of hours worked. Further details about the data are contained in Appendix I.

I include in my main sample married individuals who are between 18 and 65 years old (inclusive), whose earned income in the base period is greater than zero, and whose spouses share all of these characteristics. 178,366 individuals fit these criteria, consisting of 89,183 husbands and the same number of wives. Summary statistics are shown in Table 2. The mean income in the sample is SEK 174,932 for husbands, as opposed to SEK 103,459 for wives.²⁷ Since men

²⁷In 2007 U.S. dollars, these amounts are equivalent to \$46,634 in mean earnings for husbands and \$27,580 for wives.

tend to have higher earnings and marginal tax rates are progressive, the mean net-of-tax share of husbands (.45) is somewhat lower than that of wives (.57). Since virtual income increases as the marginal tax rate increases (*ceteris paribus*), and since men have larger capital income than women, it makes sense that husbands have substantially higher virtual income on average (SEK 75,477 for husbands, as opposed to SEK 13,863 for wives).

Couples display positive assortative mating. In 1989, their earned incomes have a modest positive correlation of .22, and the correlation of their net-of-tax shares is .36.²⁸ Pooling the changes from 1989 to 1990 and from 1990 to 1991, the correlation between the changes in their earned incomes is .13, the correlation between the changes in their simulated net-of-tax shares is .27, and the correlation between the changes in their actual net-of-tax shares is .26. Appendix Table A1 shows the number of couples who experienced each type of tax change. A cell displays the number of couples in which the husband and wife received a given combination of simulated percent changes in their net-of-tax shares. For a given simulated change in one spouse’s net-of-tax share, there is substantial heterogeneity in the other spouse’s simulated change.

V. Empirical Results

A. Preliminary Evidence

Figure 2 shows the evolution of earnings from 1988-1989, 1989-1990, and 1990-1991, in panels a, b, and c respectively. The x-axis of the figure shows real taxable income in the base year (1988, 1989, and 1990 in a, b, and c respectively). The y-axis shows mean log real earnings growth among married individuals from the base year to the final year in each 10,000-SEK range of base year real taxable income (where the final year is 1989, 1990, and 1991 in a, b and c, respectively). The figure shows starkly different patterns in 1990-1991 than in the previous years. From 1988-1989, log real earnings growth is moderately negatively correlated with base year income, consistent with the mean reversion in income noted in other contexts (e.g. Gottschalk and Moffitt 1994 on the U.S.). From 1989-1990, when approximately 1/3 of the tax cuts from TR91 were enacted (with larger cuts for those with larger incomes), there is a fairly flat relationship between earnings growth and

²⁸If an increase in one spouse’s income causes a decrease in the other spouse’s earnings—as the empirics below show—then the observed positive correlation between spouse’s incomes would understate the positive correlation that would obtain absent the income effect of one spouse’s income on the other spouse’s labor supply.

base year income, consistent with the notion that the larger tax cuts for those with higher earnings spurred faster earnings growth relative to 1988-1989. From 1990-1991, however, a starkly different pattern is visible, with a substantial positive correlation between base year income and earnings growth. This strong positive correlation—and the reversal of the negative correlation in the period before the tax cuts—is again consistent with the notion that higher tax cuts for higher-income individuals spurred relatively faster earnings growth in that group.

Figure 3 shows that during the period of the tax reform, larger gains in earned income occurred in the parts of the income distribution that also experienced larger tax cuts, relative to the period without the tax reform. As described in Section III, my regressions effectively contrast the change in the income distribution from 1988 to 1989, to the changes from 1989 to 1990 and from 1990 to 1991, and relate these relative changes to those in marginal tax rates, controlling for other factors. Figure 3 graphically depicts these relationships.

On the x-axis of Figure 3 is real earned income in the base year (in SEK divided by 10,000). The squares represent the mean simulated change in the the log of the net-of-tax share from 1990 to 1991 in each 1990 income group within a 10,000-Kronor range, minus the mean simulated change in the log of the net-of-tax share from 1988 to 1989 in each 1988 income group.²⁹ The circles show the mean gain in the log of real earned income in each 1990 earned income group from 1990 to 1991, minus the mean gain in the log of real earned income in each 1988 earned income group from 1988 to 1989. Figure 3a shows the graph for husbands, and Figure 3b shows the graph for wives.³⁰ It is evident that there are much larger gains in log real earned income from 1990 to 1991 at the top of the income distribution, relative to the bottom of the income distribution, than from 1988 to 1989. For both husbands and wives, the line showing the gain in log real earned income and the line showing the simulated increase in the log net-of-tax share tend to grow quickly with base year real earned income until about SEK 170,000, after which both lines level off. It is notable that it is not simply the case that income growth and tax cuts are both higher at higher income

²⁹The simulated change in the net-of-tax share is not monotonically increasing in base year real earned income because the brackets before and after the tax change do not occur at exactly the same points in the income distribution, because the tax base changed from 1990 to 1991, and because the marginal tax rate is based on taxable income (rather than earned income).

³⁰There are few women in the high income ranges—between SEK 180,000 and SEK 250,000, there are only 306 women on average in each 10,000-Kronor range—so it is unsurprising to find substantial volatility in mean earnings growth in this range.

levels, but also that both level off in the same income range.

B. Comparison to Other Periods in Sweden

It is worth noting that inequality is only weakly related to the business cycle in Sweden. Using time series data, Anders Björklund (1991) calculates the correlation between income inequality and Swedish GDP growth using data on 1958 and 1960-73. I then use the coefficients he estimates using these older data, in combination with actual data on Swedish GDP growth during the 1990-91 period, to calculate that relative to the increase in inequality that one would have predicted on the basis of the estimates of Björklund (1991), the actual increase in inequality in 1990-91 is substantially larger. This is again consistent with the notion that the tax cuts spurred relatively greater earnings growth for high-income individuals than for low-income individuals. The available evidence also shows that inequality does not usually increase more in the first year of a recession than in subsequent recession years (which is apparent in the work of Kopczuk, Saez, and Jae Song 2010 on the U.S.). Moreover, the pattern of large relative income gains at the top of the income distribution from 1990 to 1991 survives when partialing out the effects of characteristics such as industry, occupation, education, age, and interactions of these variables. This suggests that the large relative decreases in marginal tax rates at the top of the income distribution help to drive the large increase in inequality in 1991 relative to the surrounding years.

Furthermore, the period of TR91 shows a notably different pattern of wage growth over the earnings distribution than other recessionary periods in Sweden. In particular, the period of TR91 shows substantially larger gains at the top of the income distribution relative to the bottom, relative to these other periods. In order to assess the relative gains at the top and bottom of the earnings distribution, I examine how wages grew on average for individuals in high- and low-wage industries in the period of TR91, as well as in periods spanning the previous two recessions in Sweden. In particular, I examine 1975-1978, 1980-1983, and 1988-1991, periods that span the three recessions in Sweden observed in my data.³¹ I calculate mean real earnings for married individuals in each two-digit industry at the beginning of these periods (in 1975, 1980, and 1988, respectively). I

³¹The recessions themselves occurred within these periods in 1976-1978, 1981, and 1991. I examine three-year periods because in the earlier years in the data, I have data only for 1975, 1978, 1980, 1983, and 1985. In order to examine a comparable 3-year period around TR91, I examine the three years examined in the main analysis in the paper (1988-1991). These 3-year periods all featured similarly sluggish growth, with a mean real GDP growth rate of .43% in 1975-78, .93% in 1980-1983, and .90% in 1988-91.

then correlate mean real earnings by industry at the beginning of the period with mean log real earnings growth for married individuals over the period in question (1975-1978, 1980-1983, and 1988-91, respectively). The rationale behind this exercise is to assess whether individuals in high-wage industries were similarly affected relative to individuals in low-wage industries in all three recessions.

The results show that individuals in high-wage industries increased their earnings by substantially more relative to individuals in low-wage industries in the 1988-1991 period relative to the 1975-1978 or 1980-1983 periods. The correlation between mean initial earnings by industry and subsequent log real earnings growth of married individuals aged 18-65 is $-.106$ in 1975-1978 and is $-.120$ in 1980-1983; earnings growth was higher in low-wage industries than in high-wage industries in the two periods of sluggish growth prior to the 1988-1991 period. However, in the 1988-1991 period, this correlation reverses sign and reaches $.021$. Thus, the 1988-1991 period stands out as different from previous periods of sluggish growth because earnings growth was positively correlated with mean industry earnings in this period. This is again consistent with the notion that tax cuts for individuals with high incomes (who tend to be concentrated in high-earnings industries) caused a larger increase in earnings in high-earnings industries than low-earnings industries, relative to the experience in previous periods of sluggish growth in Sweden.

C. Basic Results

I anchor the evolution of the income distribution by regressing the change in the log of real earned income from 1988 to 1989 on a ten-piece spline in own and spousal 1988 log real earned income and 1988 log real taxable income, as well as control variables. The coefficients on the different pieces of the spline are significantly different from each other, indicating that rich controls for lagged income are warranted. In the main regressions, I instrument for four independent variables. The first-stage regressions show F-statistics ranging from just over 6,000 to nearly 10,000. The R-squared ranges from $.47$ to $.56$. When a given variable is the dependent variable, its predicted value enters highly significantly, with a coefficient between $.8$ and $.9$.

Table 3 shows the basic regression results. In columns 1 and 2, I perform regression (5) for the husband and wife. The controls include age, age squared, number of children below 8, as

well as dummies for nine possible levels of education, 24 Swedish regions, and year. I estimate own uncompensated elasticities of .17 for husbands and .25 for wives. The estimates are precise, with standard errors of .021 and .026, respectively. These point estimates are lower than the uncompensated elasticity estimated in Gruber and Saez (around .4) for the population as a whole, but comparable to the uncompensated elasticities estimated in Gruber and Saez for the low- and middle-income populations (.18 and .12, respectively), whose mean incomes are more closely comparable to the sample in Sweden. I find substantial and precisely-estimated own income elasticities of -.074 for husbands and -.056 for wives, similar to those in Sören Blomquist and Håkan Selin (2007). The negative sign is consistent with the presumption that leisure is a normal good. Consistent with the typical finding that women's labor supply is more elastic than men's, wives' own uncompensated elasticity is significantly higher than husbands' ($p < .01$).

Given these parameter estimates, it is possible to calculate compensated own and cross elasticities, using the Slutsky equation and the transformation from elasticities into effects at the sample means. These are shown in the bottom section of the table. The compensated own elasticity is .25 for husbands and .49 for wives (significantly different from each other, and from zero, at the 1% significance level). Compensated cross elasticities, .048 and .051 for husbands and wives, respectively, are also substantial. Both are significantly different from zero ($p < .05$). As one would expect, these are smaller than the compensated elasticities with respect to one's own net-of-tax share. Interestingly, husbands and wives have similar uncompensated cross elasticities and similar compensated cross elasticities. The uncompensated cross elasticities are not significantly different from zero at conventional levels. Cross income elasticities, -.0041 for husbands and -.018 for wives, are also substantial and significantly different from zero ($p < .01$).

It is theoretically ambiguous whether one's earnings should rise or fall when the spouse's tax rate rises. Parallel with the literature on labor supply, in which the leisure of husbands and wives could be complementary or substitutable, the "effort" of one spouse (reflected in pre-tax earnings) could be complementary or substitutable with that of the other. The results show complementarity: as the net-of-tax share of one's spouse rises, one's own earnings rise. A number of factors could lead to complementarity. If one's spouse takes more leisure time, it may be more enjoyable (provide higher marginal utility) to take more leisure time oneself. Complementarity is

also consistent with several forms of social interactions, such as spouses imitating one another.

D. Robustness and Extensions

Columns 3 and 4 add further controls for the 2-digit occupation and the 2-digit industry of both the husband and wife, and interactions of all of the control variables with the year dummies. This controls still further for business cycle factors unique to each year of the tax change. These regressions show broadly similar estimates, with slightly smaller elasticities. The results are also similar when I add more interactions of these variables to the regression, such as interactions of occupation and industry with education, age, or region.

In Columns 5 and 6, taxable labor income is the dependent variable. The compensated own elasticity of taxable labor income with respect to the net-of-tax share is .22 for husbands and .35 for wives, and the compensated cross elasticities are .061 and .051, respectively. Interestingly, the elasticity estimates are similar to those in the regressions in which earned income is the dependent variable. This may relate to the fact that in Sweden, the deductions available both before and after the reform may not have been particularly elastic. For example, one of the major deductions was for costs associated with commuting to work, and these choices are typically found to be relatively inelastic with respect to price in the short run.

In Columns 1 and 2 of Appendix Table 2, I take the sample from Table 3, but in order to focus more directly on the intensive margin, I exclude those individuals in couples in which at least one member does not participate in the labor market in the final period.³² Since this involves selecting the sample on the basis of an outcome variable, the results should be interpreted with caution. Because I exclude labor market non-participants, I can allow the dependent variable to be the residuals of $\ln(E_t^s/E_{t-1}^s)$ (rather than the residuals of $\ln[(1 + E_t^s)/(1 + E_{t-1}^s)]$, as in Table 3). When either dependent variable is used, the results are nearly identical. The estimated coefficients are usually about half as large as the estimated coefficients in the comparable specification in Columns 1 and 2 of Table 3. As shown in Columns 3 and 4 of Appendix Table 2, the results are generally insensitive to other choices, such as adding .5 to earnings before taking the log. Moreover, the results are similar when applying the inverse hyperbolic sine transformation to the data, which

³²The results are similar when I take the sample from Table 3 and instead exclude those individuals who do not participate in the labor market in the final period.

approximates the log transformation and is defined for values of earnings that are zero (e.g. Karen Pence 2006). In Columns 5 and 6 of Appendix Table 2, I control for mean percentage housing price growth by county and year. The results are very similar to the basic results in Columns 1 and 2 of Table 3.

One potential concern is that families who received tax cuts of different magnitudes are systematically different from one another. Typically, the more a husband's income exceeds his wife's, the larger his tax cut will be relative to hers. If shocks to couples were correlated with how much husbands earned relative to their wives, then the coefficient estimates could be biased. To address this concern, I examine couples in which the usual positive correlation between spouses' relative incomes and their relative simulated changes in tax rates is reversed. Specifically, I examine couples in which the spouse that had the larger income prior to the tax cut received the smaller simulated tax cut. This occurred in 28% of the sample. There were two primary reasons this occurred. First, as shown in Table 1, there are certain income ranges in which those with higher incomes received smaller tax cuts. Second, the tax base shifted in TR91, implying that the size of one's tax cut depended not only on the size of one's income, but also on the composition of that income; those making income from sources that received lighter tax treatment in TR91 tended to get larger tax cuts. When I run the specification corresponding to Columns 1 and 2 of Table 3 on only the 28% subsample, my estimates are .14, .024, -.057, -.0038, .23, .019, -.054, and -.013 for the uncompensated own elasticity for husbands, the uncompensated cross elasticity for husbands, the own income effect for husbands, the cross income effect for husbands, the uncompensated own elasticity for wives, the uncompensated cross elasticity for wives, the own income effect for wives, and the cross income effect for wives, respectively. When I run the same regressions on the remaining 72% of couples, the point estimates are only slightly larger and all statistically indistinguishable from those in the 28% subsample: .18, .035, -.084, -.0046, .27, .026, -.066, and -.025, respectively. As a second check on the results, it is also possible to replace each spouse with a "placebo" spouse, who is chosen randomly from among all other married individuals of the same gender as the real spouse. I then run the same regressions as in the main regressions, but with the values of the independent variables of the true spouse replaced with the values for the placebo spouse. As expected, the results show no significant response to the placebo spouse's incentives.

Chetty (2009) suggests that the response of measures of non-wage compensation to taxation should be examined in addition to measures of wage compensation for determining the deadweight cost of taxation. Measures of non-wage compensation show no significant response to taxation in my data. In particular, I examine the only non-wage perk measured in the data—company car benefits—that is consistently measured throughout the sample period.³³ When I regress the change in husbands’ log real company car benefits on the instrumented change in both spouses’ log net-of-tax share and log real virtual incomes (using the same empirical strategy as above), I estimate coefficients (standard errors) of .22 (.25), .19 (.14), .19 (.19), and .10 (.11) on own net-of-tax share, spouse’s net-of-tax share, own virtual income, and spouse’s virtual income, respectively; for wives, these results are .18 (.18), .14 (.14), .12 (.13), and .04 (.10), respectively. These insignificant responses are suggestive of the conclusion that changes in earned income represent changes in labor supply, either through changes in hours worked or changes in effort per hour worked that is reflected in the pre-tax wage, rather than changes in the form of compensation. This finding is also relevant to the interpretation of the main results. Spouses’ earnings in principle could respond to own net-of-tax share because the spouses’ form of compensation changes, rather than because the spouse’s leisure shifts. For example, in principle it is possible that when one’s own tax rate rises, the form of spouses’ compensation shifts away from non-wage and toward wage compensation. This seems unlikely, and indeed the empirical results bear out that there is no significant response. However, the large confidence intervals on the estimates and incompleteness of the data on non-wage benefits prevent firm conclusions on this issue.

E. Implications

My results can be compared with the predictions of the unitary model of family earnings responses to taxation discussed above. The own income effect for husbands is significantly different from the cross income effect for husbands ($p < .001$), and the own income effect for wives is significantly different from the cross income effect for wives ($p < .001$). These inequalities represent a violation of the income pooling prediction of the unitary model, under which own and cross income effects should be equal. Income pooling is also rejected at the 1% significance level when I test

³³Changes in the taxation of many non-wage benefits in TR91 implied that many benefits are measured inconsistently before and after the reform in the data.

whether it holds at quantiles of the income distribution other than the mean of income.³⁴ I cannot reject the unitary model's prediction of Slutsky symmetry—the equality in equation (10)—at conventional significance levels ($p=.19$). Income pooling is also violated, and Slutsky symmetry is not violated, when I evaluate the implied effects at values of the independent and dependent variables other than the sample means.

Using these elasticity estimates, it is possible to calculate counterfactual GDP growth in the absence of the tax cuts. I assume a Cobb-Douglas production function in capital and labor, and a labor income share of output of .73 (Bentolila and Saint-Paul 2003). Under these assumptions, if the individuals affected had not received a tax cut, GDP growth would have risen by 0.2% in 1990 (rather than rising by 1% in 1990) and would have fallen by 3.5% in 1991 (rather than falling by 1.1%). These are in the same range as GDP growth rates in Finland at the time, where real GDP grew by .5% in 1990 and fell by 6% (more than Sweden) in 1991.

F. Comparison with Standard Empirical Specifications

Tables 4-6 show the results of frequently-estimated econometric specifications, which yield quite different results than those in Table 3. In Columns 1 and 2 of Table 4, I estimate a standard specification in which married individuals treat their spouses' income as unearned income. This specification imposes a restriction—the coefficient on one's own change in log real virtual income is restricted to be the same as the coefficient on the change in the log real income of one's spouse—and is therefore *a priori* undesirable. Own uncompensated elasticities are .57 and .78 for husbands and wives, respectively. These are much greater than the estimates of .17 and .25 in Columns 1 and 2 of Table 3. The income elasticities, .24 and .19 for husbands and wives, are very large and of the wrong sign.³⁵ The implied compensated own-elasticities are .46 and .62 for husbands and wives, respectively, which are also higher than the compensated elasticities in Table 3 (.25 and .49, respectively). Very similar biases occur in all of the estimates in Table 4 when the dependent variable is taxable labor income.

³⁴In particular, I have tested income pooling at each decile of the distribution of income and reject income pooling at the 1% significance level at each of these points.

³⁵In principle, it is possible that leisure is an inferior good, which would be consistent with the positive coefficient on the change in log real unearned income. However, it appears implausible that leisure would be so strongly inferior, and the coefficient restriction associated with this specification makes it independently undesirable.

If the leisure of husbands and the leisure of wives are complementary, then we would expect the standard specification to yield income elasticities that are more positive than the true income elasticities.³⁶ Suppose, for example, that a wife receives a tax cut, leading her earned income to rise. Under complementarity, the tax cut for this wife also leads her husband to earn more. Thus, a rise in the husband's earned income is correlated with a rise in the wife's earned income. Under the standard specification, the change in the husband's earned income contributes to the change in the measure of the wife's unearned income. Thus, the wife's unearned income (which includes the husband's income in this specification) tends to rise when the wife's earned income rises, which contributes to a positive coefficient on the wife's unearned income. Given that income elasticities are overly positive, it also makes sense that uncompensated elasticities are over-estimated. *Ceteris paribus*, a rise in the net-of-tax share will decrease virtual income, since the intersection of the extended budget segment with the y-axis falls when the net-of-tax share rises. In the specification of Tables 3 and 4, the change in one's own log net-of-tax share is therefore negatively correlated with the change in one's log real virtual income. Yet when spousal income is included along with own virtual income in the measure of unearned income, as in the standard specification, this negative correlation is dulled. Thus, in the specification in Columns 1 and 2 of Table 4, the coefficient on the change in one's own log net-of-tax share picks up some of the variation that is actually attributable to the change in one's log real virtual income.

To explore the factors responsible for these results, Columns 3 and 4 run the traditional specification, but with own (instrumented) virtual income entered separately from (uninstrumented) spousal income. This replicates the specification in Columns 1 and 2, but without the restriction imposed that own virtual income is summed with spousal income in calculating the measure of own unearned income. The coefficient on spousal income is positive and large, whereas the coefficient on own virtual income is negative and large. This is consistent with the assertion that the positive correlation between the change in spousal income and the change in one's own income induces overly positive income effects in Columns 1 and 2. I reject the hypothesis that the coefficient on own virtual income is the same as the coefficient on spousal income ($p < .001$).³⁷

³⁶In the presence of substitutable spousal leisure, the estimated income elasticities should be overly negative. Here I use "leisure" as shorthand to describe the absence of effort, as reflected in pre-tax earnings.

³⁷I also reject the hypothesis that these coefficients are equal when I run a specification in which each of these variables enters the regression linearly (rather than the logarithmic specification in Table 6).

Uncompensated own elasticities and the coefficients on own log virtual income are similar to those in the specification in Table 3. As argued above, it appears that the own uncompensated elasticity is over-estimated in Columns 1 and 2 because when spousal income is included along with own virtual income in the measure of unearned income, the own uncompensated elasticity term picks up variation actually attributable to own virtual income.

To shed more light on the factors driving the results and evaluate to what extent other common specifications may be biased, Table 5 shows the results of omitting various key right-hand-side variables. Columns 1 and 2 omit spouse income and show that the coefficient on own NTS and own income are nearly unchanged. As expected due to the strong negative correlation of spouse virtual income and spouse NTS, the coefficient on spouse NTS is much more positive than in the Table 3 specification. Columns 3 and 4 omit the spouse terms entirely from the regressions. Interestingly, this makes little difference to the estimates of the own elasticities, presumably due to the relatively low correlation between changes in spouse tax rates shown in Appendix Table 1. Columns 5 and 6 include only own and spouse NTS as independent variables of interest. As expected due to the negative correlation of virtual income and NTS, the coefficients are substantially larger than in Table 3. The own NTS coefficients are smaller than in Table 4 Columns 1 and 2, presumably because the income effect was so over-estimated in the Table 4 specification. Columns 7 and 8 estimate the results with only own NTS as an independent variable, a familiar specification from much of the literature. As before, I find little difference in the elasticities from the specification in Columns 5 and 6 that also includes spouse NTS. Evidently ignoring the spouse in an empirical specification makes little difference, but constraining responses to own and spouse variables makes a large difference.

Table 6 investigates specifications often employed in the literature on the elasticity of taxable income. In particular, in this literature it is common to include a measure of overall family income as an independent variable (e.g. Gruber and Saez 2002), disallowing the possibility of separate effects of each spouse's income on earnings. In Columns 1 and 2, I investigate the consequences of this specification. The change in log family virtual income is the change in the log of husband virtual income plus wife virtual income, and this is instrumented using the predicted change. The own elasticities are much larger than in the basic specification, whereas the reaction to the spouse's

net-of-tax share (NTS) is negative rather than positive. It makes sense that the uncompensated own NTS elasticity is over-estimated: the own NTS is negatively correlated with own virtual income, and the estimate of the effect of virtual income is less negative in this specification than the estimated effect of own virtual income in Table 3. The coefficient on own NTS therefore picks up some of the variation that is picked up by own virtual income in Table 3. Likewise, it makes sense that the effect of the spouse's NTS is overly negative: the effect of spouse's income is now estimated to be more negative than in Table 3 (since spouse's income and own income are lumped together), and since spouse's virtual income is highly negatively correlated with spouse's NTS, the coefficient on the latter is now more negative than in the Table 3 specification.

In the U.S., we typically have information on the sum of spouses' earnings, which is then related to the family's marginal tax rate and income. In the Swedish context, spouses have different marginal tax rates; because the family lacks one marginal tax rate, it does not make sense to relate overall family income to the family's marginal tax rate. To address this issue, it is possible to investigate the results in Sweden when the simulated changes in each spouse's marginal labor income tax rate are equal, as in the U.S. context. In Columns 3-5 of Table 6, I regress earnings on husband's net-of-tax share, husband income, and wife income, using husband's earnings, wife's earnings, and the sum of husband's and wife's earnings as the dependent variables. As expected, when the sum of husband and wife income is the dependent variable, the elasticity with respect to each independent variable is the sum of the elasticities for husbands and wives separately. The regressions with husband's and wife's income as the separate dependent variables are informative because they allow us to disaggregate the overall elasticity into the separate contributing parts. When I regress earnings on the change in the NTS and the change in family income in Columns 6-8, I estimate similar uncompensated elasticities but a substantially higher income elasticity than in Columns 3-5. The higher income elasticity makes sense, since I sum together husband and wife income before taking the log; if smaller percentage changes in virtual income are associated with similar percentage changes in family earned income, the elasticity of family earned income with respect to virtual income is accordingly higher.

VI. Discussion and Conclusion

This paper looks inside the family to uncover rich aspects of spouses' earnings responses to

taxation. I argue that this may lead to new conclusions about the effects of labor income taxation that were obscured in earlier analyses that estimated more simplified econometric models. The results reveal that individuals respond substantially to their spouses' incentives, with sizeable compensated cross elasticities and cross income effects. A customary specification, which treats spousal income as unearned income, produces income effects that are wrong-signed and large, and considerable bias results in the estimates of the uncompensated and compensated own elasticities. A specification that assumes that individuals respond to overall family taxable income—precluding the possibility that the response to each spouse's unearned income is different—may also lead to substantial biases. Relatedly, specifications that relate overall family earnings to the family's tax rate obscure the separate contribution of each spouse's earnings decision to overall family earnings.

While it has been recognized that these simplified specifications are not fully justified, it has not been recognized that these simplifications may lead to substantial biases in the resulting parameter estimates. The reasons for these substantial differences, relating to the simplifications inherent in the customary specifications, stand independent of features of the particular Swedish context in which I estimate the results. The more general lesson, applicable outside of the Swedish context, is that standard specifications simplify a richer family decision and may therefore lead us astray. The results suggest that at the least, we should place a premium on gathering data in other contexts on the separate earnings of each spouse, and relating each spouse's earnings separately to measures of both spouses' incentives (including separate measures of the virtual income of each spouse).

These results have implications for models of family taxation. At a basic level, the paper uncovers a reaction to spousal incentives, suggesting that treating individuals without considering the incentives of their spouses may be an unwarranted simplification. The unitary model of family labor income responses to taxation is rejected by the income pooling test, with own income effects much larger than cross income effects, but I cannot reject Slutsky symmetry. If bargaining power depends on the family distribution of income but not on the tax rate, this could rationalize these patterns in the data (Martin Browning, Pierre-André Chiappori, and Valérie Lechene 2006). This is consistent with a model in which a spouse's outside option in the marriage market is influenced by their income but not directly by their hourly wage, and thus their bargaining power within the marriage depends on their income but not their hourly wage. This makes sense if potential mates

are attracted by the possibility of high income. Overall, the paper suggests a picture of the family in which wives' labor supply is more elastic than husbands', spousal leisure is complementary, and bargaining power depends on income but not wages.

Ideally it would be desirable to use these tax policy changes to examine other models of family taxpaying behavior, for example by testing of efficiency of intra-household allocations as in Browning and Chiappori (1998). However, the LINDA data are not suited to this task: Browning and Chiappori show that one must observe demands for at least five commodities in order to test efficiency, and the LINDA data do not contain information on consumption. Feldstein (1999) derives a measure of the deadweight cost of taxing labor income under the assumption of a unitary decision-maker. Likewise, the literature on the optimal taxation of the family has typically assumed a unitary decision-maker (Michael Boskin and Eytan Sheshinski 1983; Henrik Kleven, Claus Kreiner, and Saez 2009; Louis Kaplow 2008; Patricia Apps and Ray Rees 2009). This paper suggests that such models could be enriched to account for non-unitary behavior.

Some recent literature has departed from the assumption of a unitary decision-maker in modeling optimal tax policy or the welfare consequences of tax reforms (e.g. Alesina, Ichino, and Karabarbounis 2011; Apps and Rees 1988; Apps and Rees 1999; Craig Brett 1998; Chiappori 1992; Olivier Donni 2003; Elisabeth Gugl 2009). My results are supportive of a key prediction of the Alesina, Ichino, and Karabarbounis (2011) model—that women's labor supply should be more elastic than men's—even in Sweden where gender differences in labor market outcomes are arguably less dramatic than elsewhere. Further empirical work to validate the predictions of their model in Sweden—such as predictions relating to the ratio of home work to market work—would be valuable.

In the model of Chiappori (1992), under some assumptions about the nature of the Pareto-efficient collective household process for sharing resources, the parameters of the household bargaining process can be estimated from only the demand for leisure of the husband and wife (as they relate to household unearned income and the wage rates of the husband and wife). These parameters can then be used to estimate the welfare consequences of family taxation. If one were willing to impose this structure on the distribution of family resources, a useful extension of the work in this paper would be to estimate the parameters of this model. The non-unitary Apps and

Rees (1999) model implies that if male and female uncompensated own and spouse elasticities are all equal, then in a “traditional” household in which the ratio of male to female labor supply is above the ratio of average male to average female labor supply, a tax reform in which the wife’s rate is increased and the husband’s rate is decreased represents a potential Pareto improvement in traditional households. The results of my paper suggest that this case does not hold, as own and cross elasticities are not found to be equal across genders.

As noted above and in Alesina, Ichino, and Karabarbounis (2011), it is noteworthy that the compensated earnings elasticity of married women is higher than that of married men even in Sweden. The higher female elasticity has been noted in many other contexts (e.g. Blundell and MaCurdy 1999). The reasons for these differences are not fully clear, although it is worth noting that the estimated *effect* of taxes on earnings—as opposed to the elasticity—is more similar across genders (e.g. Blomquist and Urban Hansson-Brusewitz 1990). The difference in calculated elasticities across genders therefore arises in part because the baseline level of labor supply is lower among women than among men, leading mechanically to a higher calculated elasticity among women. In the sample in this paper, mean earnings among men (SEK 174,932) is 69% higher than mean earnings among women (SEK 103,459), and the compensated earnings elasticity among women (0.49) is 96% higher than that among men (0.25). This functional form issue can therefore account for most, but not all, of the difference in elasticities (relative to effect sizes). Other work has found that the female labor supply elasticity may decline as female labor force participation rises (Blau and Kahn 2007), perhaps in part because of this functional form issue, and perhaps because the intensive margin of labor supply may be less elastic than the extensive margin.³⁸

The results are also noteworthy insofar as they may reflect changes in labor supply. Taxable labor income is sometimes thought to reflect a broader measure of labor supply than hours worked does (Feldstein 1995, 1999). For example, effort per hour worked should influence earned income by increasing the marginal product and thus the hourly wage. Moreover, Chetty (2009) argues that taxable earnings responses could provide a better measure of welfare than does taxable income because earned income responses may be associated with smaller transfer costs. Measures of hours worked are also subject to substantial measurement error. While earned income will also reflect

³⁸However, as I discuss further below, I use earnings as the dependent variable, so the results are not directly comparable with those of studies that use hours worked as the dependent variable.

changes in the form of compensation—for example, taxation might affect the mix of compensation between fringe benefits and wage compensation—the results also show no significant response of non-wage compensation to the independent variables.

To the extent that these results reflect changes in labor supply, these results may relate to a long line of literature on family labor supply (e.g. Orley Ashenfelter and James Heckman 1974; Paul Devereux 2004; Blau and Lawrence Kahn 2007). This literature has often been limited by the difficulty of finding exogenous variation in the incentives of the spouse. Interpreted as labor supply responses, the estimates suggest that the leisure of one spouse is complementary with the leisure of another (e.g. Jennifer Hunt 1998; Alan Gustman and Thomas Steinmeier 2000; Daniel Hamermesh 2002). In this context, the results suggest that the customary labor supply specification, in which labor supply is related to wages or the net-of-tax wage, and one spouse's income is treated as the unearned income of the other spouse, can produce substantially biased estimates. Nonetheless, my estimates are of earnings responses, not labor supply responses, and thus cannot directly yield information about the labor supply elasticity. As noted above, earnings could be impacted by several factors other than hours worked. Thus, the empirical results yield a composite of all of these responses, including hours worked responses and also changes in the form of compensation, tax avoidance, and tax evasion. The potential implications for labor supply, therefore, are ultimately speculative. In order to address the relationship between hours worked and earnings responses conclusively, we must await a different administrative dataset in which data on hours worked and earnings can be linked.

The ability to examine the effects of separate changes in each spouse's marginal tax rate has allowed me to go beyond the limitations of standard specifications, and the difference in the resulting estimates is large. Several European countries have individual taxation and have made available administrative micro-data on the income of each spouse. This suggests the possibility of future work in such contexts, which could add to the picture of family earnings decisions emerging from this paper.

Appendix I. Additional Data Description

Education dummies are dummies for nine categories measuring highest school attainment. Industry and occupation are defined at the 2-digit level. Occasional missing values of these covariates are represented by dummies indicating missing values. For the vast majority of households considered in the regressions, both spouses' earnings are positive in both base periods (i.e. both 1989 and 1990). However, a number of households have positive earnings for both household members in the base period in one of these years but not in the other. Observations for these individuals are included in the regressions only when the income of both household members is positive in the base year; otherwise, the dependent variable is a missing value. 178,366 individuals are in households in which both spouses have positive earnings in at least one of the years examined.

In 1991, Sweden switched from a global tax system, under which the marginal tax rate on earned income depends on the sum of earned income, capital income, and taxable government transfers (minus deductions), to a dual tax system, under which the marginal tax rate on earned income is computed only based on earned income (and deductions and taxable government transfers), and capital income is taxed at a flat rate. This implies that the proper way to calculate virtual income is different in 1991 than it was before 1991. Prior to 1991, virtual income is calculated by computing the intersection of the individual's extended budget segment with the y-axis in taxable income-consumption space, and adding the value of untaxed transfers. Predicted virtual income in 1990 is calculated by inflating the value of taxable income in 1990 by the mean per-person growth in taxable income of individuals in the sample, calculating the virtual income associated with this predicted budget segment, and adding this amount to the predicted value of untaxed transfers (calculated by inflating 1989 untaxed transfers by the mean per-person growth in untaxed transfers from 1989 to 1990 of individuals in the sample).

In 1991, virtual income is computed by adding three quantities: the intersection with the y-axis of the individual's extended budget segment in pre-tax taxable labor income-consumption space, the after-tax value of capital income, and the value of untaxed government transfers. (Here taxable labor income is taken to include government transfers.) Because of the change in the tax base, in constructing the instrument for the marginal tax rate for 1991, I project 1991 taxable labor income by multiplying each individual's 1990 taxable labor income by the mean per-individual growth in taxable labor income of individuals in the sample from 1990 to 1991. I calculate predicted virtual income in 1991 by determining what virtual income would have been in 1991 if an individual had the projected taxable labor income in 1991, as well as the projected values of capital income and untaxed transfers (calculated by inflating the values of capital income and untaxed transfers from 1990 by the mean growth from 1990 to 1991 in the per capita values of these variables of individuals in the sample). Like all income variables, virtual income is always represented in real terms.

When it enters as a dependent variable in my regressions, I construct taxable labor income by subtracting deductions from earned income. The deductions in question do not include deductions for interest payments or capital losses. To form a consistent measure of deductions, I exclude those that were available only before or only after 1991. When I subtract deductions from earned income, the result is occasionally negative. (Because the sample excludes labor market non-participants, earned income minus deductions is negative for only .28% of the sample.) Since I examine the change in the log of real taxable labor income, and the log of zero or of a negative number is undefined, I set the values of real taxable labor income equal to 1 for these individuals in the years in which it is negative. The results are

insensitive to this choice. Before 1991, certain deductions could be claimed only against the basic tax schedule. However, all of the deductions included in my measure of deductions prior to 1991 could be claimed against both the basic schedule and against the additional schedule. Thus, their marginal tax price was equal to the net-of-tax share associated with earned income, so a specification that relates my measure of taxable labor income to this net-of-tax share is appropriate.

References

- Agell, Jonas, Peter Englund, and Jan Södersten. 1998. *Incentives and Redistribution in the Welfare State: The Swedish Tax Reform*. London: Macmillan.
- Alesina, Alberto, Andrea Ichino, and Loukas Karabarbounis. 2011. "Gender Based Taxation and the Organization of the Family." *American Economic Journal: Economic Policy* 3(2): 1-40.
- Anderson, Patricia, and Bruce Meyer. 2000. "The Effects of the Unemployment Insurance Payroll Tax on Wages, Employment, Claims, and Denials." *Journal of Public Economics*, 78(1-2): 81-106.
- Apps, Patricia, and Ray Rees. 1988. "Taxation and the Household." *Journal of Public Economics* 35(3): 355-69.
- Apps, Patricia, and Ray Rees. 1999. "On the Taxation of Trade within and between Households." *Journal of Public Economics* 73(2): 241-63.
- Apps, Patricia, and Ray Rees. 2009. *Public Economics and the Household*. Cambridge: Cambridge University Press.
- Auten, Gerald and Robert Carroll. 1999. "The Effect of Income Taxes on Household Behavior." *Review of Economics and Statistics*, 81(4): 681-693.
- Ashenfelter, Orley and James Heckman. 1974. "The Estimation of Income and Substitution Effects in a Model of Family Labor Supply." *Econometrica*, 42: 73-85.
- Bentolila, Samuel and Saint-Paul, Gilles. 2003. "Explaining Movements in the Labor Share." *B.E. Press Contributions to Macroeconomics*: Vol. 3: Iss. 1, Article 9.
- Björklund, Anders 1991. "Unemployment and Income Distribution: Time Series Evidence from Sweden." *Scandinavian Journal of Economics*, 93: 457-465.
- Blau, Francine, Marianne Ferber, and Anne Winkler. 2005. *The Economics of Women, Men, and Work*. Englewood Cliffs, NJ: Prentice Hall.
- Blau, Francine and Lawrence Kahn. 2007. "Changes in the Labor Supply Behavior of Married Women: 1980-2000." *Journal of Labor Economics*, 25: 393-438.
- Blomquist, Sören and Urban Hansson-Brusewitz. 1990. "The Effect of Taxes on Male and Female Labor Supply in Sweden." *Journal of Human Resources*, 25(3): 317-357.
- Blomquist, Sören and Håkan Selin. 2009. "Hourly Wage Rate and Earnings Responsiveness to Changes in Marginal Tax Rates." Uppsala University mimeo.
- Blundell, Richard, Alan Duncan, and Costas Meghir. 1998. "Estimating Labor Supply Responses Using Tax Reforms." *Econometrica*, 68: 827-861.

Blundell, Richard and Thomas MaCurdy. 1999. "Labor Supply: A Review of Alternative Approaches." In *Handbook of Labor Economics*, vol. 3A, ed. O. Ashenfelter and D. Card, 1560–1695. Amsterdam: Elsevier.

Boskin, Michael, and Eytan Sheshinski. 1983. "Optimal tax treatment of the family: Married couples." *Journal of Public Economics*, 20: 281–297.

Brett, Craig. 1998. "Tax Reform and Collective Family Decision-Making." *Journal of Public Economics*, 70(3): 425–40.

Browning, Martin, Pierre-André Chiappori, and Valérie Lechene. 2006. "Collective and Unitary Models: A Clarification." *Review of Economics of the Household*, 4: 5–14.

Browning, Martin, and Pierre-André Chiappori. 1998. "Efficient Intra-Household Allocations: A General Characterization and Empirical Tests." *Econometrica* 66: 1241-1278.

Burtless, Gary, and Jerry Hausman. 1978. "The Effect of Taxation on Labor Supply: Evaluating the Gary Negative Income Tax Experiment." *Journal of Political Economy*, 86: 1103–1130.

Chetty, Raj 2010. "Bounds on Elasticities with Optimization Frictions: A Synthesis of Micro and Macro Evidence on Labor Supply." NBER Working Paper 15616.

Chetty, Raj 2009. "Is the Taxable Income Elasticity Sufficient to Calculate Deadweight Loss? The Implications of Evasion and Avoidance." *American Economic Journal: Economic Policy*, 1(2). 31–52.

Chetty, Raj, John Friedman, Tore Olsen, and Luigi Pistaferri. 2011. "Adjustment Costs, Firm Responses, and Labor Supply Elasticities: Evidence from Danish Tax Records." *Quarterly Journal of Economics* 126(2): 749-804.

Chiappori, Pierre-André. "Collective Labor Supply and Welfare." *Journal of Political Economy* 100(3): 437-467.

Devereux, Paul 2004. "Changes in Relative Wages and Family Labor Supply." *Journal of Human Resources*, 39: 696–722.

Donni, Olivier. 2003. "Collective Household Labor Supply: Non-Participation and Income Taxation." *Journal of Public Economics* 87: 1179–1198.

Edin, Per-Anders and Peter Fredriksson. 2000. "LINDA – Longitudinal Individual Data for Sweden." Department of Economics, Uppsala University Working Paper 2000:19.

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

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: 551–572.

Feldstein, Martin 1999. "Tax Avoidance and the Deadweight Loss of the Income Tax." *Review of Economics and Statistics*, 81: 674–680.

Giertz, Seth. 2007. "The Elasticity of Taxable Income over the 1980s and 1990s." *National Tax Journal*, 60(4): 743-768.

Goolsbee, Austan 1999. "Evidence on the High-Income Layer Curve from Six Decades of Tax Reform." *Brookings Papers on Economic Activity*, 2: 1–47.

Gottschalk, Peter, and Robert Moffitt. 1994. "Growth in Earnings Instability in the U.S. Labor Market". *Brookings Papers on Economic Activity* Vol. 2: 217-254.

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

Gugl, Elisabeth. 2009. "Income Splitting, Specialization, and Intro-family Distribution." *Canadian Journal of Economics* 42(3): 1050–71.

Gustman, Alan, and Thomas Steinmeier. 2000. "Retirement in Dual-Career Families: A Structural Model." *Journal of Labor Economics*, 18: 503–545.

Hamermesh, Daniel 2002. "Timing, Togetherness and Time Windfalls." *Journal of Population Economics*, 15: 601–623.

Hansson, Åsa 2007. "Taxpayers' Responsiveness to Tax Rate Changes and Implications for the Cost of Taxation in Sweden." *International Tax and Public Finance*, 14: 563–582.

Heim, Bradley. 2009. "The Effect of Recent Tax Changes on Taxable Income: Evidence from a New Panel of Tax Returns." *Journal of Policy Analysis and Management*, 9(1), 147-163.

Holmlund, Bertil, and Martin Söderstrom. 2008. "Estimating Dynamic Income Responses to Tax Reforms: Swedish Evidence." Institute for Labour Market Policy Evaluation, IFAU Working Paper 2008:28.

Hunt, Jennifer 1998. "Hours Reductions as Work-Sharing." *Brookings Papers on Economic Activity*, 1: 339–381.

Kaplow, Louis 2008. *The Theory of Taxation and Public Economics*. Princeton: Princeton University Press.

Kleven, Henrik, Claus Kreiner and Emmanuel Saez. 2009. "The Optimal Income Taxation of Couples." *Econometrica* 77: 537-560.

Kopczuk, Wojciech 2005. "Tax Bases, Tax Rates and the Elasticity of Reported Income." *Journal of Public Economics*, 89(11-12): 2093–2119.

Kopczuk, Wojciech, Emmanuel Saez, and Jae Song. 2010. "Earnings Inequality and Mobility in the United States: Evidence from Social Security Data since 1937." *Quarterly Journal of Economics*, 125: 91-128.

Kuziemko, Ilyana, and Eric D. Werker. 2006. "How Much Is a Seat on the Security Council Worth? Foreign Aid and Bribery at the United Nations." *Journal of Political Economy* 114(5): 905-930.

Lindsey, Lawrence 1987. "Individual Tax Payer Response to Tax Cuts: 1982-1984, with Implications for the Revenue Maximizing Tax Rate." *Journal of Public Economics*, 33: 173–206.

Ljunge, Martin and Kelly Ragan. 2005. "Labor Supply and the Tax Reform of the Century." University of Chicago mimeo.

Lundberg, Shelly, Robert Pollak and Terence Wales. 1997. "Do Husbands and Wives Pool their Resources? Evidence from the U.K. Child Tax Credit." *Journal of Human Resources*, 32: 463–480.

Mahon, Paule Rianne. 1997. "Child Care in Canada and Sweden: Policy and Politics." *Social Politics* 4(3): 382-418.

Mas-Colell, Andreu, Michael Whinston, and Jerry Green. 1995. *Microeconomic Theory*. Oxford: Oxford University Press.

Melkas, Helina, and Richard Anker. 1998. *Gender Equality and Occupational Segregation in Nordic Labor Markets*. Geneva: ILO.

Mincer, Jacob. 1962. "Labor Force Participation of Married Women." In *Aspects of Labor Economics*, ed. H.G. Lewis. Princeton: Princeton University Press.

Moffitt, R., Wilhelm, M., 2000. "Taxation and the labor supply decisions of the affluent." In: Slemrod, J. (Ed.), *Does Atlas Shrug? The Economic Consequences of Taxing the Rich*. Cambridge University Press, Cambridge.

Pence, Karen. 2006. "The Role of Wealth Transformations: An Application to Estimating the Effect of Tax Incentives on Saving." *The B.E. Journal of Economic Analysis and Policy* 5(1). Available at http://works.bepress.com/karen_pence/1.

Saez, E., 1999. "The Effect of Marginal Tax Rates on Income: A Panel Study of 'Bracket Creep'" *Journal of Public Economics* 87: 1231-1258

Saez, Emmanuel. 2003. "The Effect of Marginal Tax Rates on Income: A Panel Study of 'Bracket Creep'." *Journal of Public Economics*, 87: 1231–1258.

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

Saez, Emmanuel. 2010. "Do Tax Filers Bunch at Kink Points?" *American Economic Journal: Economic Policy*, 2(3): 180-212.

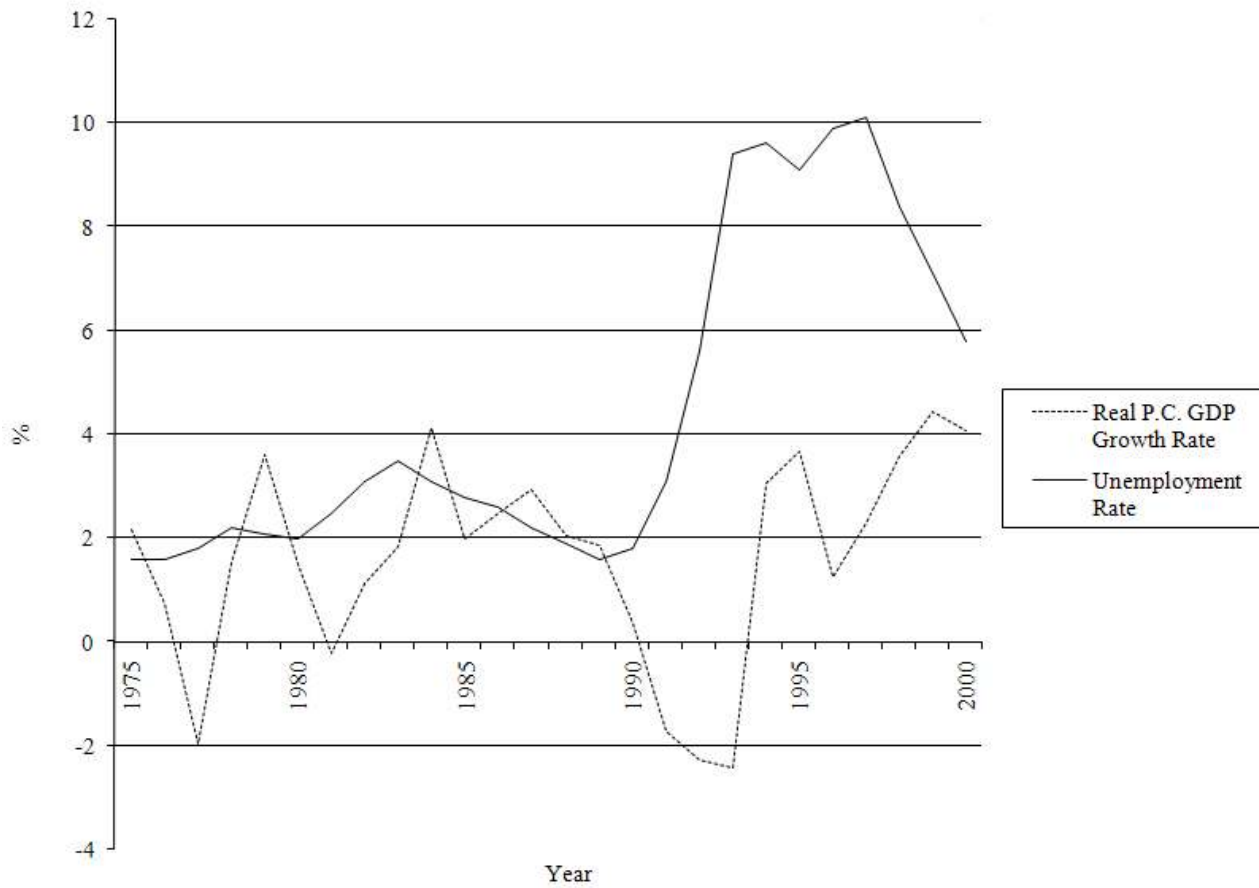
Sammartino, F., Weiner, D. 1997. "Recent evidence on taxpayers' response to the rate increases in the 1990's." *National Tax Journal* 50(3): 683–705.

Singleton, Perry. 2011. “The Effect of Taxes on Taxable Earnings: Evidence from the 2001 and Related US Federal Tax Acts.” *National Tax Journal* 64(2): 323-352.

Stevenson, Betsey and Justin Wolfers. 2007. “Marriage and Divorce: Changes and their Driving Forces.” *Journal of Economic Perspectives*, 21: 27–52.

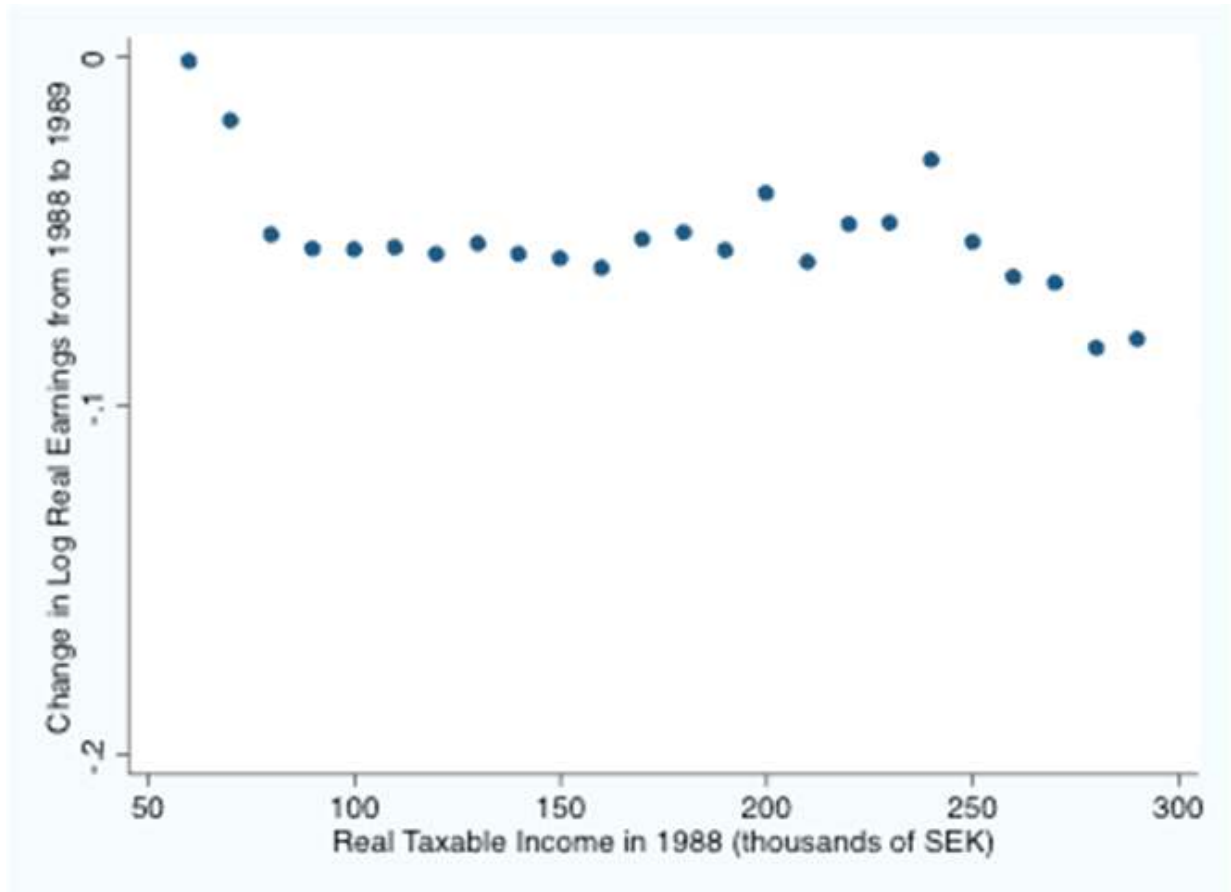
Swedish Ministry of Finance. 1991. The Swedish Tax Reform of 1991. Stockholm: Norstedts Tryckeri AB.

Figure 1. Macroeconomic Variables in Sweden, 1975-2000



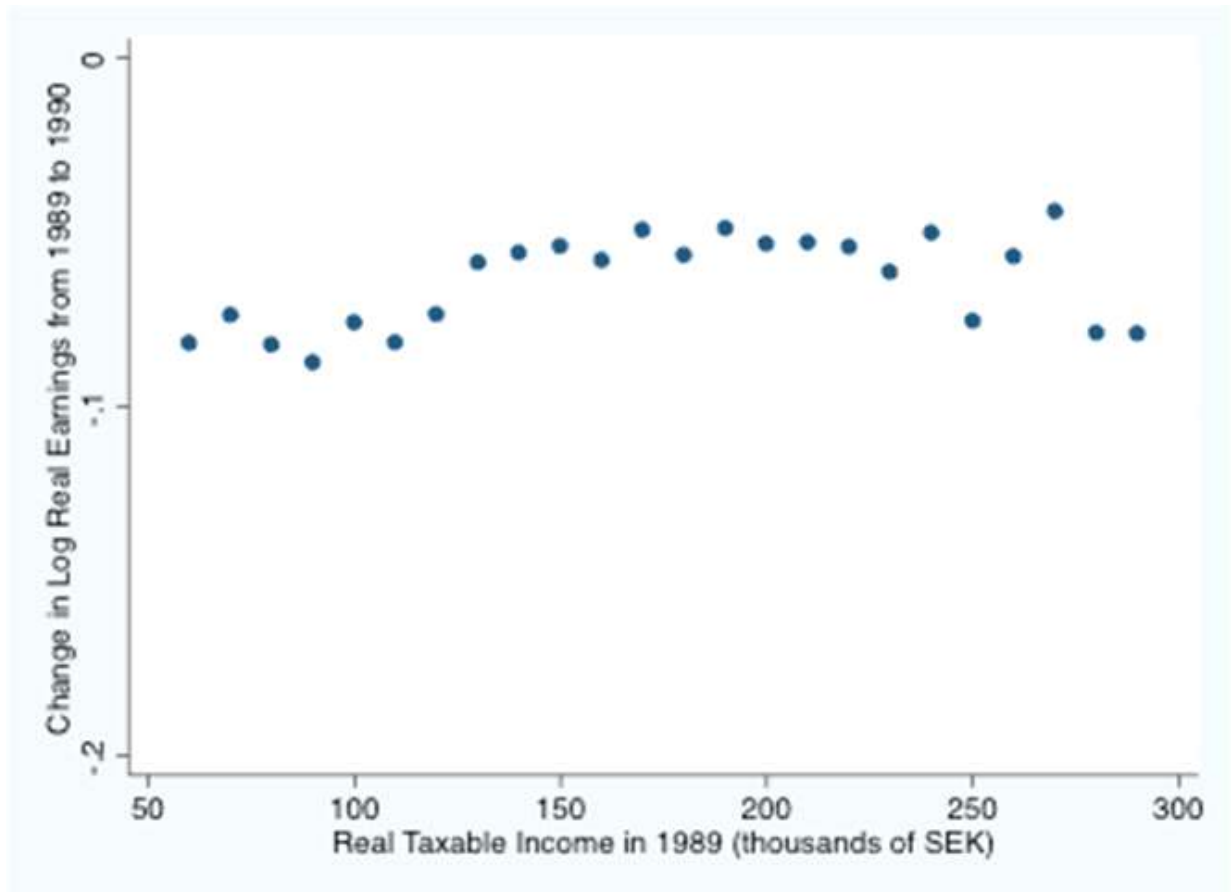
Source: U.S. Bureau of Labor Statistics, Foreign Labor Statistics. The dashed line shows the yearly growth rate of real GDP per capita in Sweden in each year from 1975 to 2000. The solid line shows the unemployment rate in Sweden in each year.

Figure 2a. Change in Log Real Earnings from 1988 to 1989 vs. Real Taxable Income in 1988



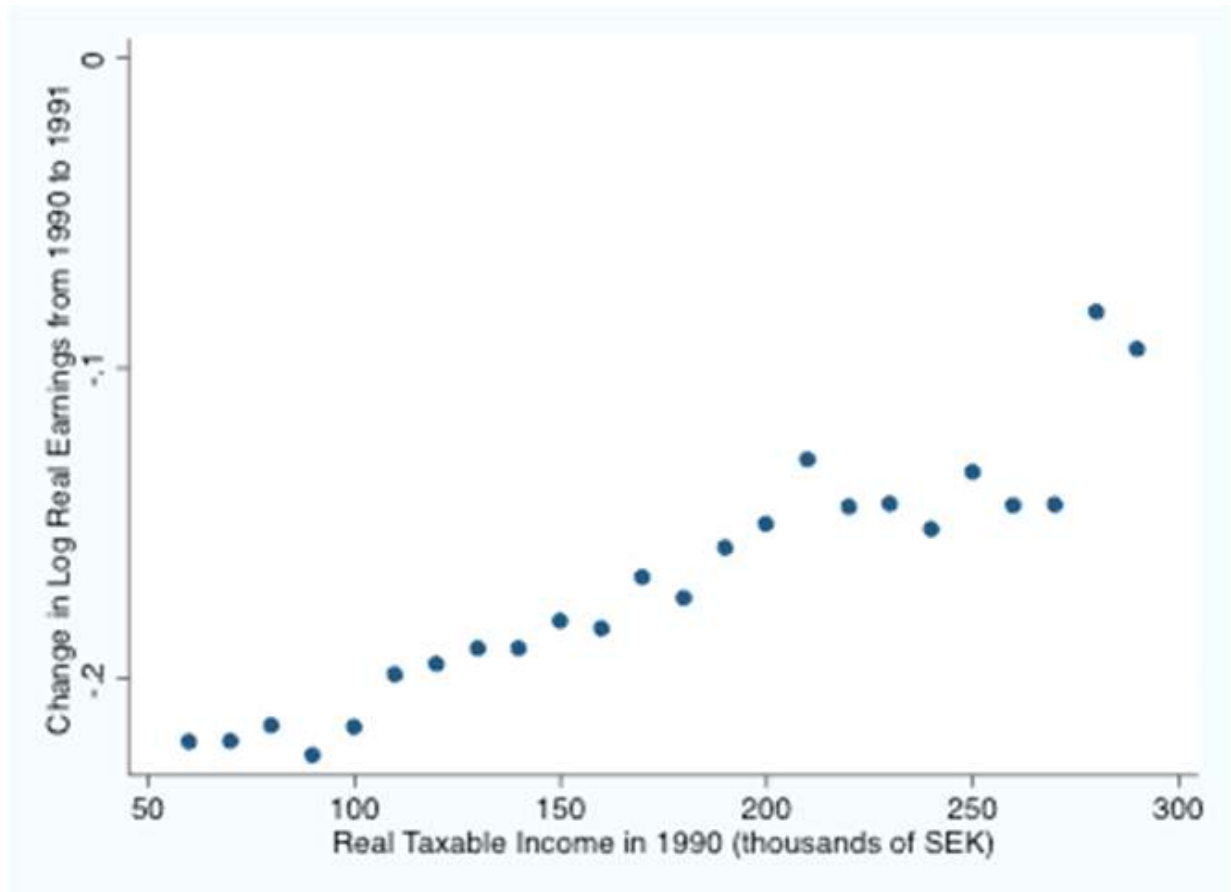
Notes: The figure shows the change in log real earnings from 1988 to 1989 on the y-axis, and real taxable income in 1988 (expressed in thousands of SEK) on the x-axis. Within each SEK 10,000 range of real taxable income in the base year, the figure shows the mean change in log real earnings from 1988 to 1989. The figure shows a negative correlation between initial taxable income and subsequent earnings growth in this period.

Figure 2b. Change in Log Real Earnings from 1989 to 1990 vs. Real Taxable Income in 1989



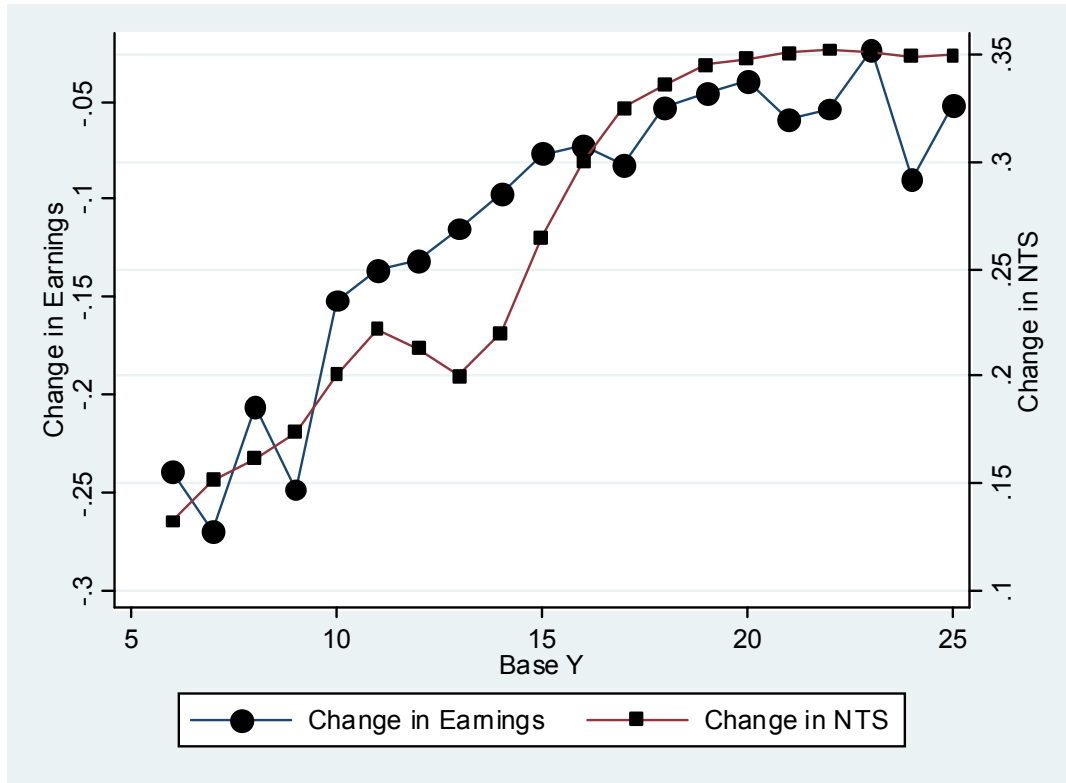
Notes: The figure shows the change in log real earnings from 1989 to 1990 on the y-axis, and real taxable income in 1989 (expressed in thousands of SEK) on the x-axis. Within each SEK 10,000 range of real taxable income in the base year, the figure shows the mean change in log real earnings from 1989 to 1990. The figure shows little correlation between initial taxable income and subsequent earnings growth in this period.

Figure 2c. Change in Log Real Earnings from 1990 to 1991 vs. Real Taxable Income in 1990



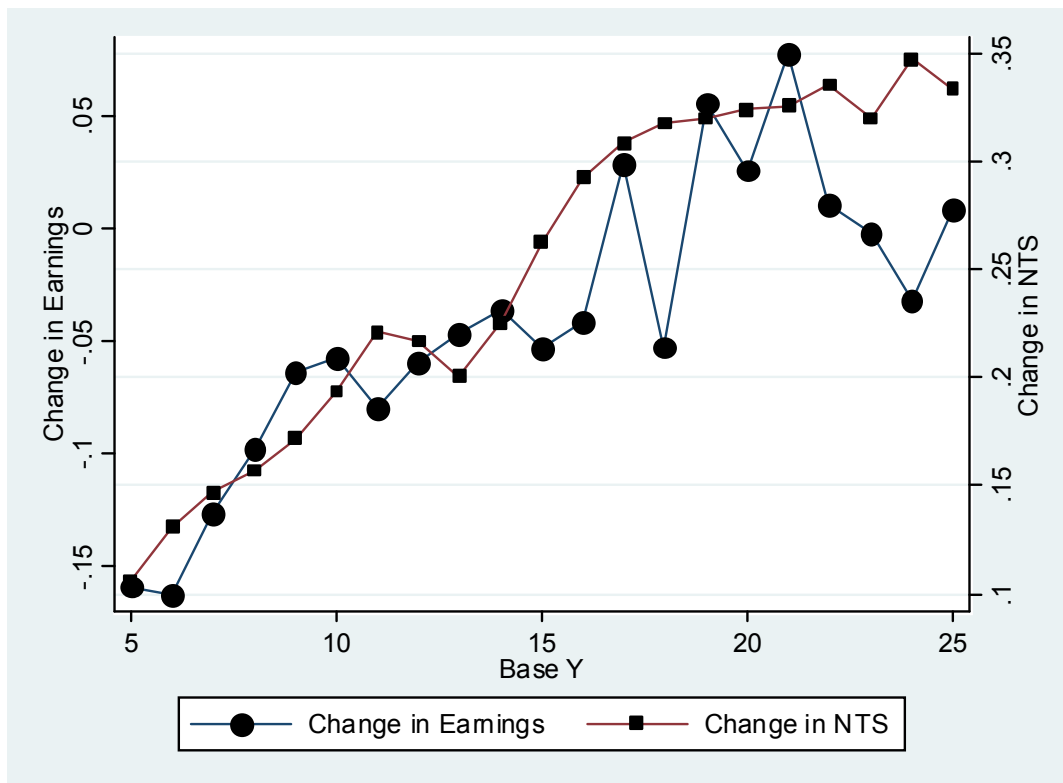
Notes: The figure shows the change in log real earnings from 1990 to 1991 on the y-axis, and real taxable income in 1990 (expressed in thousands of SEK) on the x-axis. Within each SEK 10,000 range of real taxable income in the base year, the figure shows the mean change in log real earnings from 1990 to 1991. The figure shows a substantial positive correlation between initial taxable income and subsequent earnings growth in this period.

Figure 3a. Changes in Earnings and Changes in Net-of-Tax Shares of Husbands, by Base Year Income Group



Notes: The figure shows that larger gains in real earned income took place in the same parts of the income distribution that experienced larger cuts in marginal tax rates, and that both rise until approximately the same point in the income distribution, at which point both lines level off. On the x-axis is real earned income (in Swedish Kronor) in the base year, divided by 10,000. The squares represent the mean simulated change in the log of the net-of-tax share (NTS) from 1990 to 1991 in each base-year income group (within a 10,000-Kronor range), minus the mean simulated change in the log of the net-of-tax share from 1988 to 1989 in each base-year income group. The mean simulated change in the log of the net-of-tax share is not monotonically increasing in base year real earned income because the brackets before and after the tax change do not occur at exactly the same points in the income distribution, because the tax base changed from 1990 to 1991, and because the marginal tax rate is based on taxable income (not earned income). The circles show the mean gain in the log of real earned income from 1990 to 1991 in each base year earned income group, minus the mean gain in the log of real earned income from 1988 to 1989 in each base year earned income group.

Figure 3b. Changes in Earnings and Changes in Net-of-Tax Shares of Wives, by Base Year Income Group



Notes: The figure shows that larger gains in real earned income took place in the same parts of the income distribution that experienced larger cuts in marginal tax rates, and that both rise until approximately the same point in the income distribution, at which point both lines level off. On the x-axis is real earned income (in Swedish Kronor (SEK)) in the base year, divided by 10,000. The squares represent the mean simulated change in the log of the net-of-tax share (NTS) from 1990 to 1991 in each base-year income group (within a 10,000-SEK range), minus the mean simulated change in the log of the net-of-tax share from 1988 to 1989 in each base-year income group. The mean simulated change in the log of the net-of-tax share is not monotonically increasing in base year real earned income because the brackets before and after the tax change do not occur at exactly the same points in the income distribution, because the tax base changed from 1990 to 1991, and because the marginal tax rate is based on taxable income (not earned income). The circles show the mean gain in the log of real earned income from 1990 to 1991 in each base year earned income group, minus the mean gain in the log of real earned income from 1988 to 1989 in each base year earned income group. It is unsurprising that in the higher income ranges, wives' mean income gains exhibit substantial volatility, since between SEK 180,000 and SEK 250,000, on average there are only 306 women in each 10,000-Kronor range.

Table 1. Marginal Tax Rates by Income and Year

1989 Schedule			1991 Schedule		
Bracket Start	Bracket End	MTR	Bracket Start	Bracket End	MTR
0	70,000	.36	0	149,284	.31
70,000	140,000	.51	149,284	—	.51
140,000	190,000	.65			
190,000	—	.76			

Source: Statistics Sweden. "MTR" refers to the marginal tax rate. The marginal tax rate is calculated by summing the Swedish state marginal tax rate with the average sum of local and municipal marginal tax rates (31% both before and after the reform). All amounts shown in the table are in real 1989 Swedish Kronor (SEK). In nominal terms, the end of the first bracket in 1991 was SEK 180,300. In 1989, an individual's tax liability was the sum of his or her liabilities on two different tax schedules, the basic tax schedule and the additional tax schedule. "Additional taxable income" refers to the measure of taxable income on the basis of which the liability on the additional tax schedule was calculated; "basic taxable income" refers to the measure of taxable income on the basis of which the liability on the basic tax schedule was calculated. Additional taxable income differed from basic taxable income because one could claim more deductions on the basic schedule than on the additional schedule. The additional schedule applied to individuals whose additional taxable income was above SEK 140,000. The tax schedule shown above for 1989 assumes that basic taxable income is equal to additional taxable income. The tax base also shifted in a number of ways from 1989 to 1991. For example, before 1991, the marginal tax rate on earned income was calculated as a function of both capital and labor income, whereas starting in 1991, capital income became irrelevant to the calculation of the marginal tax rate on earned income. "—" indicates that the bracket continues at all higher levels of income.

Table 2. Summary Statistics

	Husbands		Wives	
	Mean	SD	Mean	SD
Earnings	174,932	84,644	103,459	49,778
Taxable Labor Income	155,283	77,684	93,831	46,267
Age	44.0	9.16	41.29	9.11
Children < 18	1.40	1.18	1.40	1.18
Net-of-Tax Share	.45	.10	.57	.088
Virtual Income	75,477	20,092	13,863	13,561

Source: LINDA data. The sample contains 178,366 individuals, of whom 89,183 are husbands and 89,183 are wives. The sample includes only married Swedes who are not self-employed, do not hold shares in a closely held corporation, are between 18 and 65 years old (inclusive), whose earnings are greater than zero in 1989 or 1990, and whose spouses share these characteristics. The values of the variables are from 1988 and are expressed in 1988 SEK. The net-of-tax share is defined as one minus the marginal tax rate. Taxable labor income is calculated by subtracting certain deductions from earned income, as described in Appendix I.

Table 3. IV Regressions of the Change in Log Real Earnings or Log Real Taxable Labor Income on the Change in both Spouses' Log Net-of-Tax Shares and Log Real Virtual Incomes

	Earned Inc.				Taxable Labor Inc.	
	(1) H	(2) W	(3) H	(4) W	(5) H	(6) W
Δ Own	.17	.25	.15	.22	.17	.19
NTS	(.02)***	(.03)***	(.02)***	(.03)***	(.02)***	(.03)***
Δ Spouse	.03	.03	.03	.02	.05	.04
NTS	(.03)	(.02)	(.03)	(.02)	(.03)	(.02)***
Δ Own Income	-.07	-.06	-.06	-.04	-.05	-.05
	(.005)***	(.002)***	(.005)***	(.002)***	(.005)***	(.004)***
Δ Spouse Income	-.004	-.02	-.0007	-.01	-.0007	-.02
	(.002)**	(.004)***	(.002)	(.004)***	(.002)*	(.002)***
Add'l. Controls?	No	No	Yes	Yes	Yes	Yes
N	89,183	89,183	89,183	89,183	89,183	89,183
Compensated Own Elasticity	.25	.49	.21	.38	.22	.37
	(.02)***	(.02)***	(.02)***	(.02)***	(.03)***	(.04)***
Compensated Cross Elasticity	.05	.05	.04	.04	.06	.06
	(.02)**	(.02)**	(.02)*	(.02)**	(.03)**	(.03)***

Notes: The dependent variable in Columns 1-4 is the residuals of $\ln[(1+E_t)/(1+E_{t-1})]$. E is earned income. The residuals are calculated by partialing out the predicted effect of lagged income, as described in the text. " Δ Own NTS" is $\ln[(1-MTR_t)/(1-MTR_{t-1})]$, where MTR is one's own marginal tax rate. " Δ Own Income" is $\ln[(1+VI_t)/(1+VI_{t-1})]$, where VI is one's own virtual income. " Δ Spouse NTS" and " Δ Spouse Income" are the analogs. The sample includes married Swedes who are not self-employed, do not hold shares in a closely held corporation, are between 18 and 65 years old (inclusive), whose earnings are positive in 1989 or 1990, and whose spouses share these characteristics. The years examined are 1989-90 and 1990-91. All regressions control for year dummies and the 1988 values of age, age squared, number of children, and dummies for education level and region. "Add'l. Controls" means that 2-digit industry and occupation dummies for both spouses and interactions of all of the controls with year dummies are included. "H" and "W" denote regressions for husbands and wives, respectively. "N" is the total number of individuals in the regressions. In Columns 5 and 6, the dependent variable is the residuals of taxable labor income. Standard errors clustered by individual are in parentheses. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

Table 4. Comparison with Other Specifications: IV Regressions of Husbands' and Wives' Change in Log Real Earnings on the Instrumented Change in Own Log Net-of-Tax Share and the Instrumented Change in a measure of Log Real Unearned Income or the Instrumented Change in Spouse's Log Net-of-Tax Share

	(1) H	(2) W	(3) H	(4) W
Δ Own NTS	.57	.78	.17	.24
	(.02)***	(.03)***	(.02)***	(.03)***
Δ Virtual Inc.	.24	.19		
+ Δ Spouse Inc.	(.01)***	(.01)***		
Δ Virtual Inc.			-.07	-.06
			(.005)***	(.002)***
Δ Spouse Inc.			.11	.12
			(.005)***	(.006)***
N	89,183	89,183	89,183	89,183

Notes: The dependent variable is the residuals of $\ln[(1+E_t)/(1+E_{t-1})]$, where E denotes earnings and the subscript denotes the year. The residuals are calculated by partialing out the predicted effect of lagged income on the true value of $\ln[(1+E_t)/(1+E_{t-1})]$, using the 1988-9 evolution of the income distribution to determine the coefficients, as described in the text. " Δ Own NTS" is $\ln[(1-MTR_t)/(1-MTR_{t-1})]$, where MTR refers to one's own marginal tax rate and the subscript refers to the year. " Δ Virtual Inc.+ Δ Spouse Inc." refers to $\ln[(1+(VI+SI)_t)/(1+(VI+SI)_{t-1})]$, where VI is the individual's own virtual income and the subscript refers to the year, and SI is actual spousal income in the year in question. " Δ Virtual Inc." refers to $\ln[(1+VI_t)/(1+VI_{t-1})]$, and " Δ Spouse Inc." is defined similarly. The sample includes only married Swedes who are not self-employed, do not hold shares in a closely held corporation, are between 18 and 65 years old (inclusive), whose earnings are greater than zero in 1989 or 1990, and whose spouses share these characteristics. The years included in the regressions are 1989-1990 and 1990-1991, when the tax changes occurred. These years are pooled in the regressions. All regressions control for year dummies and the 1988 values of age, age squared, number of children, dummies for nine education levels, and dummies for 24 regions. "H" refers to regressions for husbands, and "W" refers to regressions for wives. "N" refers to the total number of individuals included in the regressions, the vast majority of whom appear in both 1989-1990 and 1990-1991. Standard errors clustered by individual are in parentheses. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

Table 5. Comparison with Other Specifications: IV Regressions of Husbands' and Wives' Change in Log Real Earnings on the Instrumented Change in Own Log Net-of-Tax Share, the Instrumented Change in Spouse's Log Net-of-Tax Share, and/or the Instrumented Change in Log Real Own Unearned Income

	(1) H	(2) W	(3) H	(4) W	(5) H	(6) W	(7) H	(8) W
Δ Own	.17	.24	.18	.25	.42	.71	.42	.71
NTS	(.02) ^{***}	(.03) ^{***}	(.02) ^{***}	(.03) ^{***}	(.02) ^{***}	(.03) ^{***}	(.02) ^{***}	(.03) ^{***}
Δ Spouse	.06	.08			.08	.08		
NTS	(.03) ^{***}	(.02) ^{***}			(.03) ^{***}	(.02) ^{***}		
Δ Own	-.07	-.06	-.07	-.06				
Income	(.005) ^{***}	(.002) ^{***}	(.005) ^{***}	(.002) ^{***}				
N	89,183	89,183	89,183	89,183	89,183	89,183	89,183	89,183

Notes: The dependent variable is the residuals of $\ln[(1+E_t)/(1+E_{t-1})]$. E is earned income. The residuals are calculated by partialing out the predicted effect of lagged income, as described in the text. " Δ Own NTS" is $\ln[(1-MTR_t)/(1-MTR_{t-1})]$, where MTR is one's own marginal tax rate. " Δ Own Income" is $\ln[(1+VI_t)/(1+VI_{t-1})]$, where VI is one's own virtual income. " Δ Spouse Income" is the analog. The sample includes married Swedes who are not self-employed, do not hold shares in a closely held corporation, are between 18 and 65 years old (inclusive), whose earnings are positive in 1989 or 1990, and whose spouses share these characteristics. The years examined are 1989-90 and 1990-91. All regressions control for year dummies and the 1988 values of age, age squared, number of children, and dummies for education level and region. "H" and "W" denote regressions for husbands and wives, respectively. "N" is the total number of individuals in the regressions. Standard errors clustered by individual are in parentheses. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

Table 6. Comparison with Other Specifications: IV Regressions of Husbands' or Wives' Change in Log Real Earnings, or the Sum of Changes in Husbands' or Wives' in Log Real Earnings, on the Instrumented Change in Spouses' or Family Log Net-of-Tax Share, the Instrumented Change in Husband Log Real Unearned Income, the Instrumented Change in Wife Log Real Unearned Income, or the Instrumented Change in Log Real Family Unearned Income

	(1) H	(2) W	(3) H	(4) W	(5) Sum	(6) H	(7) W	(8) Sum
Δ Own	.30	.59						
NTS	(.02)***	(.03)***						
Δ Spouse	-.02	-.05						
NTS	(.03)	(.02)***						
Δ Family			.25	.11	.36	.24	.12	.36
NTS			(.07)***	(.06)**	(.10)***	(.07)***	(.06)**	(.10)***
Δ Husb			-.08	-.04	-.12			
Income			(.02)***	(.01)***	(.02)***			
Δ Wife			-.02	-.05	-.07			
Income			(.02)	(.01)***	(.02)***			
Δ Fam.	-.05	-.07				-.10	-.09	-.19
Income	(.006)***	(.006)***				(.01)***	(.008)***	(.01)***
N	89,183	89,183	23,759	23,759	23,759	23,759	23,759	23,759

Notes: "H" and "W" denote regressions for husbands and wives, respectively, in which the dependent variable is the residuals of $\ln[(1+E_t)/(1+E_{t-1})]$, where E denotes the unearned income of husbands or wives, respectively. In the columns labeled "Sum," the dependent variable is the change in the sum of the residuals of the log earned income of the husband and wife summed. Columns 1 and 2 are based on the full sample, and Columns 3-8 limit the sample to those observations in which have the same simulated changes in their NTS, which I call " Δ Family NTS." The residuals are calculated by partialing out the predicted effect of lagged income, as described in the text. " Δ Husb NTS" is $\ln[(1-MTR_t)/(1-MTR_{t-1})]$, where MTR is one's own marginal tax rate. " Δ Husb Income" is $\ln[(1+VI_t)/(1+VI_{t-1})]$, where VI is one's own virtual income. " Δ Wife NTS" and " Δ Wife Income" are the analogs. " Δ Family Income" refers to $\ln[(VI_{ht}+VI_{wt})/(VI_{ht-1}+VI_{wt-1})]$, where VI refers to virtual income. The sample includes married Swedes who are not self-employed, do not hold shares in a closely held corporation, are between 18 and 65 years old (inclusive), whose earnings are positive in 1989 or 1990, and whose spouses share these characteristics. The years examined are 1989-90 and 1990-91. All regressions control for year dummies and the 1988 values of age, age squared, number of children, and dummies for education level and region. "N" is the total number of individuals in the regressions. Standard errors clustered by individual are in parentheses. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

Appendix Table A1. Number of Spouses with each Possible Combination of Simulated Percent Changes in Net-of-Tax Shares

		Changes for Wives											
		-23	-6	-2	8	26	29	33	38	77	81	82	156
Changes for Husbands	-23	0	2	1	71	1	8	34	0	0	0	0	0
	-6	0	89	0	906	0	199	1,470	0	0	28	9	3
	-2	1	2	10	245	1	13	164	0	0	3	3	2
	8	5	58	18	4,582	1	256	2,182	3	1	56	10	12
	26	0	0	4	61	2	7	33	0	0	4	1	1
	29	1	5	17	9,553	0	3,107	17,221	0	0	428	193	28
	33	2	249	2	9,066	1	1,508	13,592	1	0	213	84	23
	38	0	0	0	13	1	0	8	2	0	0	0	0
	77	0	1	0	4	0	0	0	0	0	0	0	0
	81	0	502	7	5,500	3	2,873	10,064	1	1	1,047	359	115
	82	0	2	0	205	0	33	332	0	0	18	57	4
	156	0	20	0	36	0	8	45	0	0	6	15	32

Notes: The number in each cell represents the number of couples in the sample in which the husband and wife had a given combination of simulated percent changes in their net-of-tax shares from 1989 to 1991. The heading of each column shows the simulated percent change in wives' net-of-tax share; the heading of each row shows the simulated percent change in husbands' net-of-tax share. The simulated percent change in the net-of-tax share is calculated as the percent change in the net-of-tax share that an individual would have experienced from 1989 to 1991 if his or her income had grown at the economy-wide growth rate of per-person income. The table shows 87,649 observations, which is 1,534 observations smaller than the sample size in the regressions in Table 3 (89,183). This is because some additional individuals enter the sample in 1990, because both of the members of their household have positive earnings in 1990, but not in 1989. There are 12 possible percentage changes in the net-of-tax share because in 1989, the total marginal tax rate was the sum of the marginal tax rate on the basic schedule and the marginal tax rate on the additional schedule, and two individuals with the same marginal tax rate on one of these schedules could have a different marginal tax rates on the other schedule. For simplicity, I have assumed that everyone faces the same total local and municipal tax rate of 31% (which is the mean sum of these tax rates both before and after the reform).

Appendix Table A2. Alternative Specification: IV Regressions of the Change in Log Real Earnings on the Instrumented Change in both Spouses' Net-of-Tax Shares and the Instrumented Change in both Spouses' Real Virtual Incomes

	(1) H	(2) W	(3) H	(4) W	(5) H	(6) W
Δ Own	.10	.13	.18	.29	.18	.25
NTS	(.01)***	(.01)***	(.02)***	(.03)***	(.02)***	(.03)
Δ Spouse	.02	.02	.03	.03	.03	.02
NTS	(.02)	(.01)	(.03)	(.02)	(.03)	(.02)
Δ Own	-.03	-.03	-.08	-.06	-.07	-.06
Income	(.003)***	(.002)***	(.01)***	(.005)***	(.005)***	(.002)***
Δ Spouse	.0002	-.01	-.003	-.02	-.004	-.02
Income	(.001)	(.002)***	(.002)	(.005)***	(.002)**	(.004)***
N	86,276	86,276	89,183	89,183	89,193	89,193

Notes: In Columns 1 and 2, I include only observations on individuals in couples in which both members have positive earnings in both the base year and the final year. Since couples are selected on the basis of an outcome variable, the results should be interpreted with caution. The dependent variable is the residuals of $\ln(E_t/E_{t-1})$, where E is earned income. The residuals are calculated by partialing out the predicted effect of lagged income from $\ln(E_t/E_{t-1})$, as described in the text. The dependent variable in Columns 3-4 is the residuals of $\ln[(.5+E_t)/(.5+E_{t-1})]$. Δ Own NTS" is $\ln[(1-MTR_t)/(1-MTR_{t-1})]$, where MTR is one's own marginal tax rate, and " Δ Own Income" refers to $\ln[(1+VI_t)/(1+VI_{t-1})]$, where VI is one's own virtual income. " Δ Spouse NTS" and " Δ Spouse Income" are the analogs. Columns 5 and 6 adopt the specification of Columns 1 and 2 of Table 3 but additionally control for mean percentage housing price growth in the county and year. The sample includes only married Swedes who are not self-employed, do not hold shares in a closely held corporation, are between 18 and 65 years old (inclusive), and whose spouses share these characteristics. The years examined are 1989-90 and 1990-91. All regressions control for year dummies and the 1988 values of age, age squared, number of children, dummies for nine education levels, and dummies for 24 regions. "H" refers to regressions for husbands, and "W" to those for wives. "N" is the total number of individuals included in the regressions. Standard errors clustered by individual are in parentheses. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.