

# Effects of Permanent and Transitory Tax Changes in a Life-Cycle Labor Supply Model with Human Capital

By

Michael P. Keane

Nuffield Professor of Economics, University of Oxford  
Visiting Research Professor, CEPAR and UNSW

April 2014

Revised September 2014

**Abstract:** I examine the effect of labor income taxation in life-cycle models where work experience builds human capital. In this case, the wage no longer equals the opportunity cost of time – which is, instead, the wage *plus* returns to work experience. This has a number of interesting consequences: First, the data appear consistent with much larger labor supply elasticities than most prior work suggests. Second, contrary to conventional wisdom, permanent tax changes can have larger effects on current labor supply than temporary tax changes. Third, human capital dampens the response of young workers to transitory tax changes, while causing responses to both permanent and transitory tax changes to increase with age. Fourth, human capital amplifies the labor supply response to permanent tax changes in the long-run, as a permanent tax reduces the rate of human capital accumulation (reducing worker productivity).

**Acknowledgements:** This paper was presented as the Cowles Lecture at the 2011 North American Summer Meetings of the Econometric Society. Seminar participants at the University of Minnesota, the Australian National University, the University of Melbourne, the University of Zurich, the World Congress of the Econometric Society (Shanghai), the 2010 Australian Conference of Economists, the University of California at Berkeley, the 2011 SETA Conference (Melbourne), the Federal Reserve Bank of New York, the Prescott Center at ASU, Oxford, UCL, the Becker-Friedman Institute and the LSE provided useful comments, as did two editors and three anonymous referees. Gary Becker provided especially useful comments during my visit to Chicago in November 2013. I thank Susumu Imai for providing me with many of the simulation exercises reported here. This research has been supported by Australian Research Council grants FF0561843 and FL110100247, and by the AFTS Secretariat of the Australian Treasury. But the views expressed are entirely my own.

## I. Introduction

In this paper I examine the effects of taxes on labor supply, wages and assets in the context of a dynamic life-cycle model that incorporates both human capital and savings. Wages are endogenous because work experience builds human capital. As current work and consumption decisions affect the next period's state (i.e., human capital and savings), a full-solution dynamic programming (DP) approach is used to solve the model. I then use the model to simulate effects of permanent and transitory tax changes on life-cycle profiles of labor supply, wages and assets. As we'll see, the introduction of human capital leads to results that differ quite sharply from models that treat wages as exogenous.

The first paper to use a DP approach to model labor supply was the classic work by Eckstein and Wolpin (1989).<sup>1</sup> Their paper looked at work decisions by married women in the NLS Mature Women's cohort. The main focus was on how the decision to work today affects wages and tastes for work in the future. As a result, they did not construct labor supply elasticities. But I have calculated that their results imply married women's labor supply is highly elastic with respect to permanent changes in the wage rate (see Keane (2011) section 7.4.1). Large labor supply elasticities are common in the female labor supply literature.

In contrast to the literature on women, the male labor supply literature has mostly ignored human capital, focussing instead on savings as the main source of dynamics (see, e.g., MaCurdy (1981)). The reason for this difference in emphasis is that the prevalence of non-participation among women naturally raises the question of how their human capital depreciates when they are out of the labor force (see, e.g., Weiss and Gronau (1981), Eckstein and Wolpin (1989)). In sharp contrast to the literature on women, the male literature has typically found very small labor supply elasticities. I will argue that this difference largely stems from the decision to ignore human capital and treat wages as exogeneous.

The first paper to use a DP approach to structurally estimate a model of male labor supply in which work experience affects wages was Wolpin (1992). He also allowed work experience to affect job offer probabilities, thus unifying the labor supply and job search frameworks. Using data on male "terminal" high school graduates from the NLSY79, he concluded that "the wage return to experience is ... considerably smaller for blacks than for whites." Wolpin used the model to examine how changes in the technology mapping work experience to wages would affect outcomes. But, unfortunately for our purposes, he did not calculate labor supply elasticities.

---

<sup>1</sup> There is a distinction between labor supply models and job search models, such as Wolpin (1987). The latter study one aspect of labor supply, but do not address the allocation of work hours over the whole life-cycle.

The paper by Wolpin (1992) was an extraordinary achievement for its time. But as Wolpin noted, a lot of corners had to be cut to make estimation feasible. He concluded that “Methods for estimating dynamic stochastic models ... are still in a relatively undeveloped stage..., This enterprise, then, is part of a much larger long-term research agenda.” Two years later, Keane and Wolpin (1994) proposed an *approximate* DP solution method that made it feasible to structurally estimate much more complex (and, we hope, less stylized) dynamic models than was previously feasible. This led to a series of papers, such as Keane and Wolpin (1997, 2000, 2007, 2010) and Lee and Wolpin (2006, 2010), that estimate progressively more sophisticated dynamic models of labor supply that incorporate human capital accumulation.<sup>2</sup>

Despite these advances, there have been very few dynamic structural models of labor supply that incorporate *both* assets and human capital. As Eckstein and Wolpin (1989) note, it is difficult to model human capital and saving simultaneously, due to the very large size of the state space. And it remains difficult even given the speed of current computers. Thus, it is no accident that the literature on male labor supply has tending to emphasize human capital and ignore saving, while the literature on women does the reverse. Researchers have been forced for computational reasons to focus on either one source of dynamics or the other.

The first paper to structurally estimate a dynamic life-cycle model with both human capital and saving was Keane and Wolpin (2001). The focus of that paper was on the school and work decisions of young men – in particular, whether borrowing constraints lead to an inefficiently low level of educational attainment.<sup>3</sup> Our results implied that borrowing constraints have little effect on schooling. Unfortunately, given the focus on education, we did not calculate labor supply elasticities.

Imai and Keane (2004) was the first paper to use full solution methods to estimate a life-cycle labor supply model that includes both human capital investment and saving, along with a continuous choice of hours.<sup>4</sup> The paper focussed on estimating the Frisch elasticity,

---

<sup>2</sup> The literature that employs the approximate DP solution method developed in Keane and Wolpin (1994) is much larger than the set of 6 papers cited here. There have been many applications in other areas besides dynamic labor supply, and even outside of labor economics. Indeed, applications have spanned a large range of areas such as industrial organization, marketing, health economics, public economics, development economics, the economics of education, child development, and so on.

<sup>3</sup> I should also mention the important paper by van der Klaauw and Wolpin (2008), who were the first to consider retirement decisions in a model with both saving and human capital. As the focus is on retirement, their model begins at age 50. Thus, I would not classify it as a life-cycle labor supply model, although retirement is obviously a very important aspect of the life-cycle labor supply decision.

<sup>4</sup> As far as I am aware, the first paper to estimate a life-cycle model with human capital was Heckman (1976). The computing technology of that time did not permit estimation of a model where workers decide jointly on savings and human capital investment, especially not while also allowing for wage uncertainty and taste shocks. Thus, Heckman’s model is deterministic and only attempts to fit “typical” life cycle paths of wages and hours. The pioneering paper by Shaw (1989) was the first to estimate the Euler conditions of a life-cycle labor supply

and how it is affected by incorporating human capital into a standard life-cycle model similar to MaCurdy (1981). We simulated transitory tax changes to calculate the Frisch elasticity, and found it is modest at young ages but grows substantial at older ages. We did not consider implications of the model for permanent tax changes – i.e., Hicks and Marshall elasticities. This was an unfortunate omission, as these elasticities are more relevant for tax policy.

The present paper remedies this omission by examining the implications of the Imai and Keane (2004) model for *both* permanent and transitory tax changes.<sup>5</sup> The main findings with regard to how introduction of human capital alters the behavior of the standard life-cycle model with savings (e.g., MaCurdy (1981)) are as follows:

(i) Human capital *dampens* the response of workers to transitory tax changes. This effect is much stronger for young workers, so effects of transitory taxes grow with age.

(ii) Human capital *dampens* the short-run response of young and prime-age workers to permanent unanticipated tax changes (whether compensated or uncompensated). But the model implies very large responses for older workers (55+) who are nearing retirement.

(iii) The well-known elasticity inequality, Frisch > Hicks > Marshall, need not hold once we account for human capital. Permanent tax changes can have larger *current* effects on labor supply than transitory tax changes. Quantitatively, the Imai-Keane model implies that the Hicks elasticity does exceed the Frisch for young workers in their 20 and 30s. At those ages, compensated permanent tax changes have larger effects than transitory tax changes.

(iv) Due to income effects, uncompensated (Marshallian) responses to permanent tax changes remain modest when human capital is included in the model, *except* that:

(v) The Frisch, Hicks and Marshall elasticities converge and become quite large for older workers – e.g., at age 60 they are all roughly 2. This is because both human capital and income effects become small as one nears the end of the working life. The estimated utility function implies high willingness to substitute labor inter-temporally, and the forces that dampen this at younger ages (human capital and income effects) have weakened.

(vi) Human capital *amplifies* the response of workers to permanent tax changes in the long-run. It also amplifies the response of total lifetime labor supply. And it causes the effect of a permanent tax change to grow over time for each individual worker. These three results are perhaps the most important of the paper. The mechanism that drives these results is that

---

model with both savings and human capital. Her estimates implied large effects of work experience on wages. But as she did not do a full solution of the DP problem, she could not calculate labor supply elasticities.

<sup>5</sup> The Imai and Keane (2004) model provides a good fit to complete life-cycle paths of wages, hours and assets, so its quantitative predictions about tax effects have some credibility. As we will see in section II, it allows for a very flexible human capital production function, wage uncertainty, taste shocks, a 45-period working life, a bequest motive, etc. These are all needed to provide a good fit to life-cycle wage, hours and asset paths.

human capital creates a dynamic feedback loop from labor supply to wages. A tax that reduces labor supply today leads to less human capital tomorrow. This lowers labor supply even more tomorrow, which reduces human capital even further. And so on. The Imai-Keane model implies this “snowball” effect is very substantial.

In discussing how human capital alters the behavior of the standard life-cycle model, it is important to note that we are actually considering two separate effects: The first is how the introduction of the learning-by-doing mechanism affects the behavior of the model *ceteris paribus* – that is, holding utility parameters fixed. The second is how accounting for human capital affects the econometric estimates of the utility parameters:

In the standard life-cycle model without human capital, the intertemporal elasticity of substitution is a simple (decreasing) function of the curvature of the utility function in leisure. Imai and Keane (2004) showed that ignoring human capital causes econometric estimates to exaggerate the curvature of the utility function, leading to a false conclusion that people are quite unwilling to substitute leisure inter-temporally.<sup>6</sup> They went on to show that estimates of a model with human capital imply much less curvature than estimates of the standard model.

The intuition for these results is that the standard model (with exogenous wages) must use utility function parameters to rationalize why wages increase much more than hours over the early part of the life-cycle. The only way it can do this is to settle on estimates that imply the willingness to substitute inter-temporally is very small. But the Imai-Keane model, which nests the standard model, can explain why young workers have high hours despite low wages by settling on estimates that imply they have a strong incentive to invest in human capital. This frees up the utility function estimates, which then imply less curvature in leisure. Thus, there is a strong substitution effect, which leads directly to result (v).

In contrast, results (i)-(iii) and (vi) are due to the human capital mechanism itself, not the utility function estimates. Result (i) arises because, with human capital, the price of time is the wage plus the future returns to human capital investment. A temporary tax alters the current wage but not the future returns to investment. Thus, a temporary tax has a relatively small effect on the price of time for young workers, leading to a small labor supply response.

Result (ii) arises because permanent tax increases flatten the earnings profile, which flattens the hours profile. So the labor supply response is dampened in the short-run and amplified in the long-run. At older ages this effect becomes trivial, leaving a pure substitution effect. So unanticipated permanent tax increases have much larger effects on older workers.

---

<sup>6</sup> This implies the Frisch elasticity is quite small. And, as the Frisch is an upper bound on the Hicks and Marshall, it implies that all three elasticities are small.

Result (iii) arises because a permanent tax increase reduces *both* the current wage and the return to human capital investment, while a transitory tax only affects the current wage. Thus, a permanent tax increase has a larger effect on the current price of time.

Finally, result (vi), that the effect of a permanent tax increase on labor supply grows over time, is an obvious consequence of letting current labor supply affect future human capital. An interesting implication is that studies that focus on short-run effects of tax reforms may greatly understate long-run labor supply elasticities.

One key experiment sheds light on the dynamics of tax effects: I use the Imai-Keane model to simulate a permanent tax increase that takes effect at age 20 and lasts for the whole working life (i.e., to age 65). The compensated elasticity is about 0.6 in the 20s and 30s. But then it rises to 1.8 at age 50 and 4.0 at age 60. (Note that this is the effect of the tax on a particular worker at different points in his life). Averaged over the whole life, the Hicks elasticity implied by the Imai-Keane model is a substantial 1.3. Most of this comes from substantial hours reductions when workers are in their 50s and 60s.

To put these figures in perspective, the survey by Keane (2011) averages over 22 of the most prominent labor supply studies (using a wide range of empirical approaches) and finds an average Hicks elasticity of 0.31. For a subset of 14 prominent studies that use direct wage measures (or attempt to deal with the denominator bias that arises if one measures the wage as annual earnings over annual hours) the mean is 0.43. Thus this elasticity of 0.6 implied by the Imai-Keane model for workers in their 20s and 30s is relatively large, but not out of the ballpark of prior work. The elasticities at older ages, and the lifetime elasticity of 1.3, are large compared to most of the earlier results.

To give further context to these figures, note that dynamic life-cycle models of female labor supply generally imply much larger elasticities. For instance, Keane and Wolpin (2010) estimate a life-cycle model for women incorporating not only human capital but also marriage and fertility. The estimated elasticity of hours with respect to permanent tax changes is 2.8. This large value arises in part because women reduce fertility in response to an increase in the after-tax wage, a margin of adjustment not available to men.

To proceed, Section II gives a simple exposition of the model in Imai and Keane (2004). Section III discusses the data and estimation of the model. Section IV discusses the fit of the model and various efforts at validation. This is relevant for evaluating the credibility of the simulation results. Section V presents the simulations of the effects of permanent and transitory tax changes on labor supply, wages and assets. Section VI concludes with a summary of findings, discussion of limitations and suggestions for future research.

## II. A Simple Exposition of the Imai and Keane (2004) Life-Cycle Labor Supply Model

Since the pioneering work by MaCurdy (1981), the standard life-cycle model with saving and exogenous wages has strongly influenced economists' thinking about male labor supply. Letting  $t$  denote age and suppressing the individual subscript  $i$ , the period utility function in MaCurdy (1981) is given by:

$$(1) \quad U_t = \frac{\alpha_t}{1+\eta} C_t^{1+\eta} - \frac{\beta \varepsilon_{1t}}{1+\gamma} h_t^{1+\gamma} \quad t = 1, \dots, T \quad \eta \leq 0, \gamma \geq 0$$

Here  $C_t$  is consumption in period  $t$  and  $h_t$  is hours of labor supplied in period  $t$ . The parameters  $\alpha_t > 0$  and  $\beta \varepsilon_{1t} > 0$  are age specific "taste shifters" that allow the marginal utility of consumption and leisure, respectively, to vary with age. The taste shifter for leisure consists of both a deterministic component ( $\beta$ ) and a stochastic component ( $\varepsilon_{1t}$ ). The latter allows the model to capture idiosyncratic variation in hours of work from year-to-year.

The law of motion for assets ( $A_t$ ) is:

$$(2) \quad A_{t+1} = (1 + r)A_t + W_t h_t - C_t$$

where  $r$  is the interest rate and  $W_t$  is the wage rate. In the standard life-cycle model the wage path  $\{W_t\}_{t=1}^T$  is assumed to be exogenous. In most applied work the log wage is assumed to be a quadratic function of age, with the intercept (and sometimes the slope) shifted by the level of education plus period specific productivity shocks.

The worker's problem is to maximize the expected present value of utility over the planning horizon  $T$ :

$$(3) \quad \max_{C_t, h_t} E_\tau \left\{ \sum_{t=\tau}^T \delta^{t-\tau} U_t(C_t, h_t) \mid I_\tau \right\} \quad s. t. \quad A_{T+1} \geq 0$$

Here  $\delta$  is the discount factor and  $E_\tau \{ \cdot \mid I_\tau \}$  is the expectation conditional on the information set at time  $\tau$ . The information set would include, for example,  $A_\tau$ ,  $W_\tau$ ,  $\alpha_t$ ,  $\beta$ , and anything else that might to predict future consumption, hours and preferences.

An important feature of the standard model is that labor supply elasticities are simple functions of utility function parameters (see Keane and Rogerson (2012)). The Frisch or inter-temporal elasticity is simply  $1/\gamma$ . Note that the Frisch is defined as the elasticity with respect to a *transitory* wage change holding marginal utility of lifetime wealth fixed. Thus, it refers to a type of compensated wage change. However, it is usually reasonable to assume that transitory wage changes have a negligible effect on lifetime wealth (see Section IV).

The Hicks elasticity, which in the life-cycle framework is the compensated response to a *permanent* tax change, is simply  $1/(\gamma - \eta)$ , and the uncompensated or Marshallian response to a permanent tax change is  $(1+\eta)/(\gamma - \eta)$ . Keane and Rogerson (2011) derive these results formally, and show that the Hicks elasticity characterizes the response to a proportional (i.e., flat-rate) tax combined with lump sum transfer of the proceeds.

Thus the standard life-cycle model with saving (and exogenous wages) give the well-known elasticity inequality  $1/\gamma \geq 1/(\gamma - \eta) \geq (1+\eta)/(\gamma - \eta)$ , i.e., Frisch > Hicks > Marshall. Of course, the three elasticities are equal in the special case of no income effects ( $\eta=0$ ).

Otherwise, transitory tax changes will have larger (current) effects than permanent ones.

Imai and Keane (2004) extend the MaCurdy (1981) model to allow for endogenous wages. They use the same specification of preferences as MaCurdy, so as to keep their results as comparable to his as possible. Continuing to suppress the individual subscript  $i$ , the wage is assumed to equal the level of human capital  $K_t$  times the rental rate on human capital  $R$ :<sup>7</sup>

$$(4) \quad W_t = RK_t$$

The law of motion for human capital is given by the following three equations:

$$(5) \quad K_{t+1} = g(K_t, h_t, t)\varepsilon_{2,t+1}$$

$$(6) \quad g(K_t, h_t, t) = k_0 + (1 - d)K_t + G(K_t, h_t, t)$$

$$(7) \quad G(K_t, h_t, t) = \pi_0(1 + \pi_1 t)\{(K_t + \kappa_0)[(h_t + \kappa_1)^\theta - \lambda(h_t + \kappa_1)]\}$$

Equation (5) says that next period's human capital  $K_{t+1}$  depends on current human capital and hours of work, as well as age, and a productivity shock  $\varepsilon_{2,t+1}$  that is realized at  $t+1$ . In (6) the parameter  $k_0$  captures an initial endowment of skill,  $d$  is the depreciation rate, and  $G(\cdot)$  is the investment in human capital.

Equation (7) describes how investment in human capital depends on hours of work ( $h_t$ ), current human capital ( $K_t$ ) and age ( $t$ ). The term in square brackets captures the learning-by-doing mechanism, whereby labor supply generates more human capital. The parameters  $\theta$  and  $\lambda$  capture the curvature of the production function in hours of work. The parameter  $\kappa_1 > 0$  assures that the derivative of  $G(\cdot)$  with respect to  $h_t$  is bounded at  $h_t = 0$ . Imai and Keane argue there is strong evidence of complementarities between human capital and hours in producing

---

<sup>7</sup> In an equilibrium model such as Lee and Wolpin (2006, 2010) the rental rate  $R$  is determined in equilibrium. Imai and Keane (2004) adopt a partial equilibrium setting where  $R$  is fixed. Keane and Rogerson (2011) discuss the conditions necessary for rental rates to be independent of tax rates in this type of model.

additional human capital (see their Figure 2). To capture this they include the term  $(K_t + \kappa_0)$  in (7). If  $\kappa_0$  grows large and the scaling parameter  $\pi_0$  grows small the complementarity vanishes. Finally, the term  $(1 + \pi_1 t)$  in equation (7) allows for a pure age effect.

The model in (5)-(7) may look complicated, but it is a flexible specification that nests the standard model with exogenous wages. To see this, note that if  $\theta = 0$  and  $\lambda = 0$  in (7) then the learning-by-doing mechanism vanishes. In this case wages are exogenous, depending only on age, the lagged wage and a stochastic term.<sup>8</sup> Thus, the Imai-Keane model does not impose *a priori* that a learning-by-doing mechanism *must* account for any or all wage growth.

Following MaCurdy (1981), a large literature has estimated the Euler conditions implied by equations similar to (1)-(3). These papers consider different assumptions about the information set  $I_t$ , the taste shifters  $\alpha_t$  and  $\beta_t$  and the functional form for utility (see Keane (2011) for a detailed review). The standard procedure is to estimate the Euler conditions by GMM, using as instruments elements of the information set  $I_t$ . Given this approach, the terminal condition does not actually matter for the estimation.

However, Imai and Keane (2004) estimate their model by full information maximum likelihood, based on a full solution of a workers dynamic programming (DP) problem. The solution to the DP problem is obtained by back-solving from the terminal period. Hence, the specification of the terminal condition is important. The terminal period is set at  $T=65$ , after which all workers are assumed to stop working and enter retirement. They are assumed to value the assets  $A_{T+1}$  that they carry into retirement according to the terminal value function:

$$(8) \quad V(A_{T+1}) = \begin{cases} 3\ln(A_{T+1} + \phi) - 1 - 3\ln(\phi) & \text{if } A_{T+1} > 0 \\ ([A_{T+1} - \phi]/\phi)^3 & \text{if } A_{T+1} \leq 0 \end{cases}$$

where  $\phi > 0$  is a parameter that determines the marginal value of assets (at  $T + 1$ ). Higher values of  $\phi$  imply the marginal value of retirement assets is less (at any given asset level). This function was chosen because it is continuously differentiable and concave in assets. In particular, it is continuously differentiable at  $A_{T+1}=0$ . Intuitively, while this function strongly discourages negative assets it does not rule them out. [Note: (8) replaces  $A_{T+1} \geq 0$  in (3)].

Notably, even at  $T=65$  the model is dynamic and has an inter-temporal substitution motive. This is, by working more at  $T=65$  one can increase consumption in retirement. In fact, as we'll see in Section IV, labor supply elasticities peak at  $T=65$ .

---

<sup>8</sup> In the special case of no learning-by-doing, ignoring stochastic terms, and setting  $R=1$ , the model in (5)-(7) implies that wages start at  $W_1 = K_1 = k_0^*$ , and then follow the process  $W_{t+1} = K_{t+1} = k_0^* + (1-d)K_t + \pi_1^*$  where  $k_0^* \equiv k_0 + \pi_0 \kappa_0$  and  $\pi_1^* \equiv \pi_0 \kappa_0 \pi_1$ . Provided the depreciation rate satisfies  $0 < d < 1$ , wages will grow at a decelerating rate, and converge in the long-run to a steady state maximum of  $(k_0^* + \pi_1^*)/d$ . Introducing the stochastic terms creates fluctuations around this basic pattern.

Finally, the model is complete by assuming the stochastic terms  $\varepsilon_{1t}$  and  $\varepsilon_{2t}$  are distributed log normally with means equal to one:

$$(9) \quad \ln(\varepsilon_{it}) \sim iid N\left(-\frac{1}{2}\sigma_i^2, \sigma_i^2\right) \quad for \quad i = 1, \dots, 2$$

The model in (1)-(9) generates value functions of the form:

$$(10) \quad V_t(A_t, K_t, \varepsilon_{2t}) = \max_{C_t, h_t} \left\{ U_t(C_t, h_t, \varepsilon_{2t}) + \delta E_t \left\{ V_{t+1}(A_{t+1}, K_{t+1}, \varepsilon_{2,t+1}) | A_t, K_t, C_t, h_t \right\} \right\}$$

Using the first order conditions for hours and consumption implied by this model, we obtain the following form for the marginal rate of substitution between consumption and leisure:

$$(11) \quad MRS_{C,l} \equiv \frac{-\partial U_t / \partial h_t}{\partial U_t / \partial C_t} = W_t + \frac{1}{1+r} \frac{\partial E_t V_{t+1}(\cdot) / \partial K_{t+1}}{\partial E_t V_{t+1}(\cdot) / \partial A_{t+1}} \left[ \frac{\partial K_{t+1}}{\partial h_t} \right]$$

In the standard life cycle  $\partial K_{t+1} / \partial h_t = 0$ , and (11) simplifies to the familiar conditional that the MRS equals the wage rate. But in the model with human capital investment the wage rate is no longer equal to the opportunity cost of time (OCT). Instead, human capital creates a wedge between the wage and the OCT that I denote by  $hc_t$ :

$$(12) \quad hc_t \equiv \frac{1}{1+r} \frac{\partial E_t V_{t+1}(\cdot) / \partial K_{t+1}}{\partial E_t V_{t+1}(\cdot) / \partial A_{t+1}} \left[ \frac{\partial K_{t+1}}{\partial h_t} \right]$$

The term  $hc_t$  captures the marginal effect of hours at time  $t$  on the expected present value of lifetime (money metric) utility. Clearly,  $hc_t > 0$  if  $\partial K_{t+1} / \partial h_t > 0$ . It is also clear that  $hc_t$  will decline with age. In fact,  $hc_T = 0$  because at  $T+1$  a worker enters retirement and has no further use for human capital. But estimates of the Imai-Keane model indicate that  $hc_t$  is roughly as large as the wage itself when workers are young. This is because they have a long horizon over which to recoup the returns to human capital investment.

Thus, in the model with human capital, the OCT is much greater than the wage rate at young ages, but it converges (from above) to the wage rate at older ages. This means that the OCT grows (much) more slowly than the wage at young ages, but tracks the wage closely at older ages. This is how the Imai-Keane model can rationalize: (i) a high willingness to substitute labor inter-temporally with data showing that young workers have rapid wage growth but only modest hours growth, and (ii) a rising labor supply elasticity with respect to the current wage as people age.<sup>9</sup>

<sup>9</sup> See Keane, Todd and Wolpin (2010) or Keane (2009, 2011) for simple graphical illustrations of these points.

Finally, consider the effect of introducing a proportional tax on labor earnings. Then (2) is replaced by  $A_{t+1} = (1 + r)A_t + W_t(1 - \mu_t)h_t - C_t$  where  $\mu_t$  is the tax rate at time  $t$ . Define a change in  $\mu_t$  while holding  $\{\mu_\tau\}_{\tau=t+1}^T$  fixed as a “transitory” tax change, and a shift in the whole vector  $\{\mu_t, \dots, \mu_T\}$  as a “permanent” tax change. A key point is that a transitory tax increase reduces  $W_t(1-\mu_t)$ , the current period component of the OCT, but it does not reduce  $hc_t$ , the future component of the OCT. That is, the return to human capital investment is unaltered because as future taxes are unchanged. But a permanent tax increase reduces both  $W_t(1-\mu_t)$  and  $hc_t$ . Thus, a permanent tax has a larger effect on the OCT than a transitory tax.

As a result, it is possible for permanent tax changes to have a larger effect on *current* labor supply than transitory, and the Frisch elasticity may be smaller than the Hicks, or even the Marshall. Such a reversal is most likely for young workers, for whom the  $hc_t$  term is a much larger share of the OCT than for older workers.

Of course, a permanent uncompensated tax increase also has a larger income effect than a transitory tax. And a permanent compensated tax increase does not generate the same incentive for inter-temporal substitution as a transitory tax. So whether permanent taxes actually have a larger effect on current labor supply is an empirical question.<sup>10</sup>

### III. Estimation Issues, Data and Results

#### III.A. Estimation Issues

In order to take the model of Section II to the data it is necessary to consider worker heterogeneity, measurement error and missing data. Despite the fact that the errors are assumed *iid* in (9), the law of motion for human capital in (5)-(7) will generate persistence in wage shocks over time. For this reason Imai-Keane did not include time-invariant unobserved heterogeneity in the wage process. However, there is rich structure of observed heterogeneity in that the production function parameters  $k_0$  and  $d$  in (6) and  $\pi_0$ ,  $\pi_1$  and  $\theta$  in (7) are allowed to differ across four education levels (high-school dropout, high school, some college, and college graduate). This allows the model to capture the feature of the data that wage growth over the life-cycle tends to be greater for higher educated workers. The disutility of work parameter  $\beta$  in (1) is also allowed to vary across the four education groups. The taste for consumption parameter  $\alpha_t$  is a spline in age with kinks at 25 and 33.

---

<sup>10</sup> Using a two-period version of the model, Keane (2009) derives the condition for uncompensated permanent taxes to have larger effects on current labor supply than transitory taxes. It requires that the return to work experience be sufficiently large relative to the income effect. A calibration exercise shows that this condition can hold for plausible values of utility and wage function parameters. Keane (2009) also shows that with borrowing constraints (i.e., a static budget constraint) permanent tax effects must be larger than transitory.

Measurement error and missing data pose a very serious problem for dynamic structural models, because construction of value functions and choice probabilities require that one know the agents' true state. But with missing data or measurement error one does not know the true state. Furthermore, forming the likelihood by integrating over the distribution of the true state (conditional on observables) is typically completely infeasible.

Keane and Wolpin (2001) developed an estimation method that is very simple to use in such cases. The method involves assuming that *all* the endogenous variables in the model (in this case wages, hours and assets) are measured with error. This assumption is reasonable in almost all applied contexts. Then, one can simulate the likelihood using only *unconditional* simulations of choice and outcome histories.<sup>11</sup> Given a simulated history, the likelihood contribution for a person is the joint density of the set of measurement errors needed to reconcile their (error ridden) observed history with the (simulated) "true" history. Imai and Keane (2004) used this method to estimate the model.

As in MaCurdy (1981) and much of the male labor supply literature, the model assumes interior solutions for hours. Obviously, making school a choice and allowing for corner solutions would be valuable additions to the model. On the other hand, the present simpler formulation allows a direct assessment of how the standard life-cycle framework is affected by including human capital.

A limitation of the Imai-Keane model is that it does not incorporate the complexities of the US tax-transfer system, or changes in tax rules over time. It should be noted, however, that most all of the structural labor supply literature shares this problem. It is very difficult to model changes in tax rules over time, as one must model how agents form expectations of future rules.<sup>12</sup> However, in the next section I argue that, in the data used here, the variation of *individual level* after-tax wages due to changes in tax rules is quite trivial compared to either life-cycle variation or transitory fluctuations in pre-tax wages.

### **III.B. The NLSY79 Data**

The model was fit to white males from the NLSY79 born in Jan. 1957 to Dec. 1964. They were interviewed annually in the calendar years 1979 to 1994. To be in the sample a person must be at least 20 and have completed school. The modal respondent was born in Dec. 1960–Jan. 1961, was 20 in 1980 and 33 in 1994. The greatest age attained by any respondents is 36. The peak sample size is N=2439 at age 28.

---

<sup>11</sup> The use of unconditional simulation of histories circumvents the need to know the current state with certainty so as to simulate transitions or form transition probabilities.

<sup>12</sup> To my knowledge, the only work that tackles this problem is Keane and Wolpin (2007, 2010), who assume agents use a VAR to forecast future tax/transfer parameters. They focus on labor supply of low income women.

The life-cycle variation in wages is enormous. The mean (and median) wage roughly doubles from age 20 to 36, increasing by roughly 50% from 20 to 25, and by another 33% from 25 to 36. The Imai-Keane estimates imply that most observed transitory wage variation is measurement error, and that the true standard deviation of transitory (annual) wage shocks is 0.058. This is less than typical estimates that ignore measurement error, but it still implies that one standard deviation transitory (annual) wage shocks at the individual level are  $\pm 6\%$ .

According to the CBO,<sup>13</sup> during 1979-1994 the average tax rate in US was 21.1%. It fell from 22% in 1979-80 to 20.2% in 1981, and then slowly rose to 21.9% in 1994. Thus, the range of variation over the period was only 1.8 percentage points. Given that tax changes were very small relative to either life-cycle or idiosyncratic wage variation, I would argue that ignoring taxes is not a serious source of error in measuring after-tax wages for most individuals during the sample period.<sup>14</sup>

### III.C. Estimation Results

The key estimation results in the paper are: (1) the estimated value of the production function parameter  $\theta$  is 0.23 and highly significant, implying learning-by-doing is important, and (2) the estimates of the utility function curvature parameters are  $\gamma = 0.26$  and  $\eta = -0.74$ . The small value of  $\gamma$  implies that utility is only moderately concave in hours, implying that people are very willing to substitute labor inter-temporally. This estimate is very different from the typically much larger values of  $\gamma$  that have been obtained using the standard model with exogenous wages. In fact, in the standard model, this would imply a Frisch elasticity of  $1/\gamma = 3.8$ . However, as we discussed earlier, and as we will see in Section V, the human capital mechanism considerably dampens the response of workers to transitory wage changes until they reach their 50s.

The estimate of  $\eta$  is larger than in most prior studies, implying people are very willing to inter-temporally substitute consumption. But Keane and Wolpin (2001) obtained a similar value. They argue prior estimates of  $\eta$  were biased downwards (implying more curvature) because, in the data, youth do not borrow heavily against expected future wage growth to smooth consumption. The standard life-cycle model explains this via a high degree of prudence, which in CRRA utility is governed by  $\eta$ . But Keane and Wolpin (2001) explain the failure of youth to borrow by liquidity constraints. Rather than modeling liquidity constraints,

---

<sup>13</sup> See Congressional Budget Office, <http://www.cbo.gov/publication/44604>.

<sup>14</sup> The relative stability of the average tax rate masks substantial changes in the top marginal rate (which was cut from 50% to 28% by TRA 1986), and in the tax rate paid by the top 1% (which fell from 35.1% in 1979 to only 24.6% in 1986, and then rose to 34.8% in 1994). However, the top rate is only relevant for a tiny subset of the NLSY79 sample. The bulk of the sample faced a relatively stable tax rate over the sample period. Indeed, Chetty (2012) argues that only the top 1% would have suffered non-trivial utility losses by failing to adjust to TRA86.

the Imai-Keane model proxies for them in a “reduced-form” way by letting the marginal utility of consumption increase with age (via the age spline  $\alpha_t$ ). This may proxy for liquidity constraints, or factors like marriage/children that cause consumption to increase with age.

In terms of heterogeneity, the only production function parameter that differs substantially across the four education types is  $k_0$  in equation (6), which captures the initial skill endowment. Naturally, more educated workers have higher skill endowments at age 20. The estimate of  $\kappa_0$  in (7) is .04, while  $\pi_0$  ranges from .13 to .15 for the four types. Together, these values imply strong complementarity between human capital and labor supply in the production of additional human capital. The implication is that the human capital investment mechanism described in Section II will be more important for better educated workers.

#### **IV. Model Fit and Validation**

Wolpin (1996) argued that economists should make greater efforts to validate their models before using them to make forecasts or policy prescriptions. He returns to this theme in much subsequent work, both philosophically and in applications (see, e.g., Wolpin (2007), Todd and Wolpin (2006, 2008), Keane and Wolpin (1997, 2007, 2009) ). Two important criteria for assessing the credibility of a model are its in-sample fit and its out-of-sample predictive performance. Thus, before using the Imai-Keane model to predict effects of taxes, it is important to examine its performance under these criteria.

Imai and Keane (2004) provide extensive documentation on model fit. They show that the model provides a very good fit to the typical age profiles of wages, hours and assets (in sample). But they also look at more challenging aspects of fit such as transition rates and conditional correlations. One method they use was, to my knowledge, first used in Keane and Wolpin (1997). The idea is to run regressions on both the actual data and on simulated data from the model. This makes it simple to assess conditional correlations.

For example, in a regression of log wages on the lagged log wage, age and age<sup>2</sup>, the coefficient on the lagged log wage is 0.878 in the NLSY79 data and 0.823 in the simulated data. Thus, the model does an excellent job of capturing the persistence of wages, despite the fact there is no individual heterogeneity. For a similar regression using assets as the dependent variable, the lagged asset coefficients are 0.616 in the NLSY79 and 0.551 in the simulated data. So the model also captures the persistence in assets very well.

One area where the model is weak is that it does not generate enough persistence in observed hours (i.e., in a regression of log hours on lagged logged hours, the lag coefficient is 0.708 in the NLSY79 but only 0.413 in the simulated data). Imai and Keane (2004) show that

the model does generate a very high degree of persistence in actual hours, so the problem with the simulated data is that the variance of measurement error is too high. We suspect this is because of misspecification of the measurement error distribution.

An interesting question about the Imai-Keane model is whether it can replicate the results in the large prior literature based on MaCurdy (1981) that assumes exogenous wages. In other words, suppose the Imai-Keane model is the true data generating process. Would a researcher using methods like those in MaCurdy (1981) and Altonji (1986) to analyse the data conclude that the inter-temporal elasticity of substitution is small?

To address this question, Imai and Keane (2004) estimate Euler conditions based on (1)-(3) on both the simulated and NLSY79 data, using instruments that are typical of the prior literature (e.g., polynomials in age and education). In the specification that is closest to MaCurdy (1981), they obtain  $\gamma = 8.0$  in the NLSY79 data and  $\gamma = 14.3$  in the simulated data. These estimates imply the utility function is highly convex in hours, and give Frisch elasticities of 0.14 and 0.08, respectively. Thus, the simulated data looks like the actual data in the sense that when viewed through the lens of the standard model with exogenous wages it implies an extremely small Frisch elasticity, just like the bulk of the prior literature.

What explains the *prema facie* contradiction between Imai-Keane's small estimated value of  $\gamma$  and the large value obtained from their simulated data? It is the divergence between the OCT and the wage in their model (see eqns. (11)-(12)). In particular, they estimate that from age 20 to 36 the mean of the OCT increases by only 13%. In contrast, the mean wage increases 90% in the actual data, and 86% in the simulated data. Thus, the wage increases about 7 times faster than the OCT. Given these figures, Keane (2009) calculates that conventional methods of estimating  $(1/\gamma)$  will understate it by a factor of roughly 7.

As a final assessment of in-sample fit, Imai and Keane (2004) compare the raw correlation between log hours changes and log wage changes in the NLSY79 and the simulated data. Strikingly, these correlations are -0.231 and -0.293, respectively. Thus, despite the strong inter-temporal substitution mechanism implied by the model estimates, the simulated data still resembles the actual data in that it generates a *negative* correlation between wage changes and hours changes.

Next, I look at the out-of-sample fit of the model. First, consider the life-cycle path of hours. The model predicts that hours rise moderately through the 20s and 30s, peak in the 40s, and then decline sharply in the 50s and 60s. Also, the model accurately predicts the quantitative magnitude of the fall in hours at older ages – e.g., it predicts an hours decline

from ages 45-54 to 55-64 of 53%.<sup>15</sup> This is close to the 47% figure for this cohort projected by McGrattan and Rogerson (1998).<sup>16</sup>

Next, consider assets. As Imai and Keane (2004) note, the model fits the in-sample asset data well in the 20s and 30s. When it is used to forecast out-of-sample it predicts large asset accumulation between ages 40 and 60, and dis-saving afterwards. This asset path is similar to the actual pattern of savings and dis-savings discussed in Carroll (1997).

Finally, regarding wages, the model correctly predicts that wages rise much more sharply than hours for men in their 20s and 30s. Imai and Keane (2004) argue that this is the fundamental data pattern that forces estimates of the standard model with exogenous wages to imply that the Frisch elasticity is small. Wages then flatten out in the 40s and decline in the 50s and 60s, giving the familiar humped shape pattern over the life-cycle.

## **V. Simulations of the Imai-Keane Model: Permanent and Transitory Tax Effects**

As the Imai-Keane model provides a reasonably good fit to wage, hours and asset patterns – both in and out-of-sample – it seems credible to use it to predict labor supply responses. So in this section I use the model to simulate effects of various types of tax changes. Table 1 reports simulated effects of permanent and transitory tax increases. The effects of unanticipated transitory tax increases, displayed in the first column of the table, were already reported in Imai and Keane (2004). But the results in the next three columns (anticipated transitory, permanent uncompensated, permanent compensated) are new. Each column of Table 1 reports the response to a 5% tax increase of the specified type.

### **V.A. Effects of Transitory Tax Changes**

First, consider two columns labelled “transitory.” Here, the tax increase applies for only one year at the indicated age. The first column reports results for unanticipated tax changes. Labor supply elasticities with respect to unanticipated transitory tax changes are usually viewed as being good *approximations* to the Frisch, or marginal utility of lifetime wealth constant, elasticities. Of course, unanticipated transitory taxes do have small income effects, so the accuracy of this assumption is worth investigating. Thus, the second column of Table 1 reports responses to *anticipated* transitory tax changes. As they are anticipated they do not alter lifetime wealth, so this column gives pure Frisch effects.

---

<sup>15</sup> A limitation of the Imai-Keane model is it assumes interior solutions for hours, so it cannot generate complete retirement prior to age 65. But this limitation should not be exaggerated. In the 2008 CPS, 70% of men aged 55-64 still worked, and 52% of men aged 62-64 still worked (see Purcell (2009)).

<sup>16</sup> Overall, the Imai-Keane model predicts average weekly hours (for white males) of 44.4, 48.9, 43.4 and 19.9 at ages 25-34, 35-44, 45-54 and 55-64. This is similar to what McGrattan and Rogerson (1998) project for all men in this cohort (see their Table 8), but the Imai-Keane hours profile is shifted up due to exclusion of minorities.

As we can see by comparing columns 1 and 2 of Table 1, the assumption that transitory tax changes have negligible income effects is perfectly valid through age 35. At age 40 we begin to see a very slight divergence, and at age 60 the elasticity for unanticipated transitory changes is  $8.6/5 = 1.72$  and that for anticipated changes is  $9.8/5 = 1.96$ . Clearly, as a worker approaches retirement, the income effect of a transitory tax increase does become noticeable, but even at age 60 the divergence between columns 1 and 2 is modest.

A key result is that the effects of transitory taxes grow substantially with age. As we see in Table 1 column 2, at age 20 a temporary 5% tax increase reduces hours by 1.5%. This implies a Frisch elasticity of only 0.30. This is far smaller than one might expect, given that Imai-Keane estimate  $(1/\gamma) = 3.8$ . But at age 60 a 5% anticipated tax increase reduces hours by 9.8%, implying a much larger Frisch elasticity of 1.96.

The intuition for this result is clear from our discussion in Section II. Transitory taxes have relatively small effects at young ages because they only affect a part of the OCT; they affect the current after-tax wage, but not the return to human capital investment. Thus, the human capital mechanism dampens the response to transitory taxes for young workers.

But, as both Shaw (1989) and Imai-Keane note, human capital investment is not very important for people late in the life-cycle. For them, the wage is close to the opportunity cost of time.<sup>17</sup> Thus, the dampening effect of the human capital mechanism diminishes, causing the elasticity with respect to transitory taxes to increase sharply as workers age.

Interestingly, French (2005), in a study of retirement, finds large elasticities for 60 year olds in the PSID – despite ignoring human capital. In his model a different mechanism causes elasticities to increase with age: the extensive margin. Many older workers are close to indifferent between working and not working, so participation is very responsive to wages.

## **V.B. Effects of Permanent Tax Changes**

In this section I consider permanent tax simulations. These are more relevant for evaluating long-run changes in tax policy, or differences in tax rates across countries. To my knowledge, Imai-Keane is the only micro model that attempts to fit asset, hours and wage data over the whole working life.<sup>18</sup> Because it generates life-cycle paths from age 20 to 65, it can be used to simulate both short-run and long-run effects of permanent tax changes.

### **V.B.1. Short-Run Effects of Permanent Tax Changes**

The last two columns of Table 1 report effects of permanent 5% tax increases, both uncompensated and compensated. In each case, the tax increase occurs (unexpectedly) at the

---

<sup>17</sup> As a result, ignoring human capital should not cause much bias in labor supply models for older workers.

<sup>18</sup> Two other papers that fit assets, hours and wages are Van der Klauw and Wolpin (2008) and Keane and Wolpin (2001). The former simulates behavior of older workers, while the latter focuses on youth.

indicated age and lasts until age 65. The Table reports only the effect on current labor supply in the year the tax increase is first implemented. This may be called the short-run effect.

In the compensated case the proceeds of the tax are redistributed to the workers in lump sum form.<sup>19</sup> Comparing columns 3 and 4, we see that compensated effects are much larger than uncompensated – implying income effects are important – except at older ages.

A key result in Table 1 is that, for younger workers, (compensated) permanent tax increases have larger effects on current labor supply than do transitory tax increases. For instance, consider a 5% tax increase that takes place at age 25. If it is transitory, hours fall by 1.8%. But if it is permanent and proceeds are distributed lump sum, hours fall by 2.7%. So at age 25, the permanent tax effect is 50% greater. By the mid-30s permanent and transitory tax effects are roughly equal. Only in the 40s do transitory tax effects become somewhat larger.

These results confirm the intuition from Section II. Permanent tax changes can have larger effects on current hours than transitory changes because they affect both the wage and human capital terms in (11), while a transitory tax only hits the current wage. But, as workers age, returns to experience fall and the human capital effect becomes small relative to income effects. Then, transitory taxes begin to have larger current effects than permanent taxes.<sup>20</sup>

Another key result in Table 1 is that, starting at age 50, both compensated and uncompensated effects of permanent tax changes on current labor supply begin to grow quite rapidly. For instance, for workers aged 20 to 40, compensated effects of a 5% permanent tax increase are only -2.3 to -3.2%. But at ages 55 to 60 these effects grow to -7.2% and -10.5%. Uncompensated effects grow even more dramatically, from -0.7% at age 40 to -9.4% at 60.

Indeed, at age 60, the compensated and uncompensated elasticities with respect to unanticipated permanent tax changes become very close (2.1 vs. 1.9). Furthermore, these permanent elasticities are both close to the Frisch (anticipated transitory) elasticity of 1.96. Thus, Hicks, Marshall and Frisch elasticities all converge to about 2.0 by age 60.

Short-run responses to permanent unanticipated tax changes increase with age for two reasons. The first is that, as one nears the terminal period, income effects of permanent tax changes become small, because workers are about to enter retirement.<sup>21</sup> The second is that, at

---

<sup>19</sup> This is Slutsky compensation, which enables workers to achieve the original consumption level at the original hours level. In contrast, in a static model, Hicks compensation holds maximized utility fixed. In the dynamic case, the analogue is to hold the maximized value function fixed, as in Keane (2009). But this requires iterative solution for the value function. The two concepts are essentially equivalent for small tax changes.

<sup>20</sup> Note that uncompensated permanent tax effects never exceed transitory effects at younger ages; this outcome is theoretically possible (see Keane (2009)), but the income effect is too strong for it to occur in practice.

<sup>21</sup> If a worker experiences a permanent tax reduction at age 60, they will only get the benefit of the higher wage for a few years, while the extra earnings must be spread over the remainder of the working life and the entire length of the retirement period. So the income effect is small. Note: If there were no retirement period, we would have static model at  $T=65$ , so income effects would still be important.

young ages, a permanent tax increase reduces the rate of human capital investment, flattening the earnings profile. This dampens the short-run hours response, because labor is substituted toward earlier periods. But near the end of the working life this effect is no longer important.

### **V.B.2. Long-Run Effects of Permanent Tax Changes**

So far, I have only discussed effects of permanent tax changes on *current* period hours. From the point of view of tax policy it is more interesting to examine their long-run impact. This is the subject of Table 2. The table considers a permanent (compensated) 5% tax increase that takes effect at either age 25, 30 or 35, and that remains in effect for the rest of a person's working life. I report how this alters a person's labor supply at 5-year intervals from age 25 to 65. For instance, say a 5% tax increase goes into effect unexpectedly when the worker is 25. Then, at age 25, his hours are reduced by only 2.7%. But, at age 45 his hours are reduced by 5.1%, and at age 60 the reduction is 19.3%.

The striking finding here is that the effect of a permanent tax change grows substantially with age. This occurs for two reasons: First, as I've already noted, as workers get older, the after-tax wage makes up a larger fraction of the OCT, so a given tax has a larger direct effect. Second, a permanent tax hike produces a "snowball" effect: If a worker reduces his labor supply at time  $t$ , he will have less human capital at time  $t+1$ . This causes him to work even less at time  $t+1$ , leading to a lower wage at  $t+2$ , etc..

This "snowball" effect of taxes on pre-tax wages is also shown in Table 2. At first, tax effects on human capital are small, but they grow substantially with age. For instance, if a 5% tax increase is instituted when a worker is 25, then by age 40 his wage is reduced by 1.0%, but by age 55 his wage is reduced by 3.6%, and by age 65 the reduction is 11.6%. So in the long-run a permanent tax reduces the rate of human capital accumulation. This lowering of pre-tax wages creates an additional work disincentive, beyond the direct effect of the tax.

Thus, the human capital mechanism *amplifies* the effect of permanent tax changes in the long-run. Given that permanent tax effects grow over time, it is interesting to examine how a permanent tax change affects labor supply over the entire working life. The results are reported in Table 3. Here, I simulate the impact of a permanent 5% tax hike that starts at age 20 and lasts through age 65. If the tax revenue is discarded, total lifetime hours drop by 2%. If revenue is redistributed lump sum, total hours drop 6.6%. Thus, we have uncompensated and compensated lifetime elasticities of 0.4 and 1.3, respectively.

Notably, if the Imai-Keane parameter estimates are used in the standard life-cycle model with exogenous wages, *a la* MaCurdy (1981), the compensated elasticity would be  $1/(\gamma-\eta) = 1/(.262+.736) \approx 1.0$ . Thus, the human capital mechanism and the "snowball" effect

on wages combine to amplify the compensated elasticity by 30% (from 1.0 to 1.3).<sup>22</sup>

Another interesting result in Table 3 is that a person living in a higher tax regime chooses to reduce hours much more at older ages than at young ages. This means that a permanently higher tax rate shifts labor supply toward younger ages. Presumably this means a higher tax regime would induce more early retirement; but the present model assumes interior solutions, so it instead causes hours to drop to very low levels for older workers.

Finally, it is also interesting to examine how asset accumulation over the life-cycle responds to permanent tax changes. This is also reported in Table 2. The basic pattern is that, after a tax increase, savings increase at first but fall later. For example, given a 5% permanent tax increase at age 25, workers respond by increasing their assets by 26.3% at age 35, but end up with assets that are 3.8% lower at the age of retirement (age 65). In other words, a permanent tax increase reduces consumption in the short run by more than the amount of the tax, not only because labor supply falls, but also because the savings rate increases.<sup>23</sup>

The reason a permanent tax increase generates more saving in the short to medium run is precisely the "snowball" effect of the tax on wage growth described earlier. Given that a tax increase reduces labor supply, a worker knows that his rate of wage growth has been reduced. So the asset response pattern is as one would expect -- young workers consume less today if their perceived life-cycle wage path is flattened. This influence of the perceived life-cycle wage path on current consumption is a central issue in the papers by Domeij and Floden (2006) and Low (2005) – see Keane and Rogerson (2011, 2012).

### **V.B.3. Comparing the Model Predictions to Estimated Effects of Tax Reforms**

The finding that human capital amplifies the effect of permanent tax changes in the long-run has important implications for the growing literature that attempts to estimate labor supply elasticities by looking at responses to major tax reforms (see Saez et al (2011) or Keane (2011) for reviews). This literature adopts a difference-in-difference approach and generally focuses on short-run responses. The results presented here suggest that a short-run focus may cause one to seriously understate responses to tax changes.

With this point in mind, it is interesting to compare the short-run tax effects predicted by the model with results from the literature on tax reforms. As we saw in Table 1, the model implies that short-run compensated labor supply elasticities with respect to permanent tax changes range from 0.54 to 0.46 for workers aged 25 to 45. Elasticities for older and younger

---

<sup>22</sup> Similarly the Marshallian elasticity is amplified from  $(1+\eta)/(\gamma-\eta) \approx 0.25$  to 0.40.

<sup>23</sup> However, in Table 2 we see that the magnitude of the increase in savings following the tax increase, as well as the drop in assets at retirement, are less if the tax is implemented when the worker is older.

workers are larger. The aggregate elasticity (over all ages) is 0.70. Workers in their 50s and 60s account for half of the aggregate response. As a point of comparison, Chetty (2012) pools estimates from many existing studies, most using the short-run effects of tax reforms as the source of identification, and obtains a Hicks elasticity of 0.58.<sup>24</sup> The Imai-Keane prediction of 0.70 is a bit higher, but this may well be explained by the fact that many microeconomic studies screen out older men, who have the largest elasticities. Thus, the Imai-Keane model prediction is in the ballpark of the quasi-experimental evidence.

## VI. Conclusion

If human capital is added to the standard life-cycle labor supply model, the wage no longer equals the opportunity cost of time (OCT). Rather, the OCT is the wage *plus* returns to work experience. This has important implications for how workers respond to taxes, and for proper estimation of labor supply elasticities. In fact, given human capital, the data appear consistent with much larger labor supply elasticities than conventional wisdom suggests.

Another key implication is that the human capital mechanism dampens responses to transitory tax changes. This is because a transitory tax only affects one part of the OCT, the *current* after-tax wage, while leaving the return to human capital investment unaffected.

In contrast, permanent tax changes alter both the *current* wage and the *future* return to human capital investment. Thus, permanent tax changes have a larger effect on the OCT than transitory tax changes. As a result, it is possible for permanent tax changes to have larger effects on *current* labor supply than transitory tax changes – contrary to conventional wisdom based on models without human capital.

In order for permanent tax changes to have larger effects than transitory, the returns to work experience must be sufficiently large relative to income effects. Using the labor supply model of Imai and Keane (2004), which gives a good fit to life-cycle paths of wages, hours and assets, I found that compensated permanent tax changes do have larger effects than transitory for workers under 35 (for whom returns to work experience are large).<sup>25</sup>

As the Imai-Keane model generates complete life-cycle paths of worker behavior, it can also be used to simulate the long-run response to permanent changes in tax policy. The

---

<sup>24</sup> Chetty (2012)'s bounds imply a Hicks elasticity of 0.33 on the intensive margin and 0.25 on the extensive margin. As the Imai-Keane model does not distinguish between the intensive and extensive margins (i.e., it is fit to data on total annual hours), I take the sum (0.58) as the relevant point of comparison. To bound the intensive margin elasticity Chetty uses 15 studies, based on tax reforms in the US, UK, Iceland, Sweden and Denmark. The extensive margin elasticity is inferred from 11 studies that examine US, UK and Danish tax reforms, EITC expansions, and cross-country tax comparisons. It is beyond the scope of the present paper to comment on the quality of the various studies.

<sup>25</sup> Note that the income effect matters critically for the difference between a compensated permanent (Hicks) effect and a Frisch effect. This is because Frisch compensation is larger than Hicks (or Slutsky) compensation.

results indicate that human capital amplifies the response to permanent tax changes in the long-run. For example, consider a tax change that goes into effect at age 25 and lasts through age 65. The (compensated) elasticity of labor supply at age 25 is only 0.54. But the elasticity grows to 1.0 at age 45 and 3.9 at age 60.

The effect of permanent tax changes grows over time because of a “snowball” effect on human capital investment: A permanent tax increase at  $t$  leads to less labor supply at  $t$ , which lowers wages at  $t+1$ , further reducing labor supply at  $t+1$ , etc.. In fact, in the previous example, a 5% tax increase leads to a 7.5% reduction in human capital by age 60. This in turn causes labor supply to fall disproportionately at older ages.

If the effects of permanent tax changes grow substantially over time, it has important implications for the literature that attempts to estimate labor supply elasticities by looking at responses to major tax reforms (see Saez et al (2011) or Keane (2011) for reviews). This literature generally focuses on short-run responses. But the results presented here suggest that a short-run focus may cause one to seriously understate responses to tax reforms.

I find that human capital amplifies the response of lifetime labor supply to taxes. In a standard life-cycle model with no human capital, the preference parameters in the Imai-Keane model would imply a compensated (Hicks) elasticity of about 1.0. But simulation of their model generates a compensated elasticity of only 0.54 at age 25 rising to 3.9 at age 60, with an average over the whole working life of 1.32. Hence, the impact on lifetime hours is one-third greater than if the human capital mechanism were shut down. Thus, not only does human capital cause permanent taxes to have a larger effect in the long-run (and at older ages), it also amplifies their total impact on lifetime hours.

A more general point is that, in a model with human capital, changes in taxes cannot be viewed as a source of exogenous variation in after-tax wages for the purpose of identifying labor supply elasticities. This is because behavioral responses to tax changes feedback and alter the life-cycle wage path itself. Thus, elasticities estimated from quasi-experimental evidence on responses to exogenous changes in tax rates are not easily interpretable in terms of underlying preference parameters. Interpretable results can only be obtained via full structural estimation of the joint labor supply/human capital investment process.

The “snowball” effect of permanent tax changes also has important implications for the calculation of the welfare costs of taxation. Obviously, if human capital models generate larger labor supply elasticities than in the conventional labor supply literature, then these larger elasticity estimates will translate into larger estimates of the welfare losses from taxation of earnings. But there is a second and subtler reason that human capital matters, and

this involves *dynamics*: As I noted earlier, a permanent tax hike not only reduces labor supply, but also the rate of human capital accumulation, thus reducing worker productivity. This suggests that static models that focus on the effect of taxes *holding work experience fixed* are missing a key channel through which welfare costs of taxation arise.

Three limitations of the paper are worth noting: First, the Imai-Keane model assumes interior solutions for hours (as is true of MaCurdy (1981) and most of the male labor supply literature). The model correctly predicts that average hours fall by about 50% from ages 45-54 to ages 55-64. But it cannot generate that roughly 30% of males aged 55-64 do not work (see Purcell (2009)). Instead, it approximates this by having some workers reduce hours to low levels at ages 55-64. Also, the model imposes retirement after age 65, rather than treating it as a choice. Building corner solutions and retirement into the model is an important avenue for future research. However, evidence in Keane and Rogerson (2012) suggests that accounting for corner solutions and/or retirement leads to higher labor supply elasticities.<sup>26</sup> Hence, it seems unlikely that including these factors would alter the main message of the paper – i.e., that labor supply elasticities are larger than conventional wisdom suggests.

Second, the models presented here ignore schooling, and consider the behavior of workers conditional on their having entered the labor force. But as noted by Keane and Wolpin (2000, 2010), changes in the tax/transfer system that reduce rewards to working will also reduce educational attainment. So accounting for this additional channel would presumably magnify the long-run tax effects on human capital found here.

Third, I assume a particular investment mechanism (learning-by-doing) while others may be operative. A leading example is the on-the-job training (OJT) model of Ben-Porath (1967). However, as Becker (1962) stressed, OJT models are similar to learning models in key respects: In OJT models the observed wage (i.e., earnings/hours) differs from the OCT (productivity) because only a fraction of work hours are spent in production. The rest is spent learning. Learning time falls with age, so *wages grow more slowly than the OCT*. This is the exact same pattern that causes the degree of inter-temporal substitution to be underestimated if we ignore learning-by-doing. More generally, Becker (1962) argued that learning, OJT and other forms of human capital investment can and should be treated symmetrically.<sup>27</sup>

---

<sup>26</sup> My results imply taxes have much larger effects on labor supply at older ages. Thus, in a model with corner solutions and/or endogenous retirement, higher taxes would presumably generate more early retirement.

<sup>27</sup> For instance, Becker (1962, p. 35-36) states: "...an important virtue...in my concept of human capital is that learning—both on and off the job—is included along with training and schooling... In general...if one activity was said to require a given investment and to yield a given return, *another activity with the same net earning stream must be said to require the same investment and yield the same return, no matter how they differ in other respects.*" (emphasis added).

This paper is part of an emerging literature exploring mechanisms that may have caused prior work to understate labor supply elasticities. Besides human capital, other potentially important mechanisms include liquidity constraints (Domeij and Floden (2006)), uninsurable wage risk (Low and Maldoom (2004)), corner solutions (Rogerson and Wallenius (2007), French (2005), Kimmel and Kniesner (1998)) and fixed costs of adjustment (Chetty (2010)). An important task for future research is to assess the relevance of these mechanisms. Suffice it to say, while conventional wisdom says labor supply elasticities are small, more dissent from that position is emerging – see Keane and Rogerson (2011) for a detailed survey.

I'll end with some suggested directions for future research. The prediction of the human capital model that effects of permanent tax increases will “snowball” over time seems very difficult to test. It would require that we have data on two different cohorts who live under different tax rate regimes (but are otherwise similar), and that we can track these cohorts for many years. The problem is that tax regimes tend to be rather variable over time.<sup>28</sup>

However, several other predictions of the human capital model are easier to test, as they only require data from shorter time periods. One is that older workers will be much more responsive to tax changes than younger workers. [Existing evidence, such as French (2005), supports this, but more work is needed]. Second, in higher tax regimes we should see hours decline more sharply with age, which implies the transition to part-time work or retirement will occur at earlier ages. Third, young workers should be more responsive to tax changes that are perceived as long-lived than to tax changes perceived as temporary. Fourth, given the complementarity between skill and work experience in production of human capital, the human capital effects discussed here should be more pronounced for workers with more skill/education. This leads to several testable hypotheses.<sup>29</sup>

Finally, the tax changes considered in this paper – i.e., purely transitory/permanent, perfectly anticipated/unanticipated, perfectly compensated/uncompensated – are ideals or archetypes. They are meant to shed light on the properties of the life-cycle model with human capital in a well-defined idealized setting. Tax changes in the real world are much messier. Predicting effects of real world tax changes is complex as it requires one to take into account the complexities of real tax/transfer systems and to model expectations about changes to those systems. This is a formidable challenge for structural econometrics.

---

<sup>28</sup> Alternatively, one could compare labor supply across countries with persistently different tax rates, but this raises the difficult issue of controlling for other factors besides taxes that may lead to cross-country differences.

<sup>29</sup> For example, for college workers, elasticities should grow strongly with both age and with the time since a tax change was implemented. But for less educated workers, these age and time paths should be relatively flat.

## References

- Altonji, Joseph (1986). Intertemporal Substitution in Labor Supply: Evidence from Micro Data, *Journal of Political Economy*, 94:3, S176-S215.
- Becker, Gary (1962). Investment in Human Capital: A Theoretical Analysis, *Journal of Political Economy*, 70:5, part 2, 9-49.
- Ben-Porath, Yoram (1967). The Production of Human Capital and the Life Cycle of Earnings, *Journal of Political Economy*, 75 (August), 352-365.
- Chetty, Raj (2012). Bounds on Elasticities with Optimization Frictions: A Synthesis of Micro and Macro Evidence on Labor Supply, *Econometrica*, 80:3, 969-1018.
- Domeij, D. and M. Floden (2006). The Labor Supply Elasticity and Borrowing Constraints: Why Estimates are Biased, *Review of Economic Dynamics*, 9, 242-262.
- Eckstein, Z. and K.I. Wolpin. (1989). Dynamic Labor Force Participation of Married Women and Endogenous Wage Growth. *Review of Economic Studies*, 56: 375-90.
- French, E. (2005). The Effects of Health, Wealth and Wages on Labour Supply and Retirement Behaviour. *The Review of Economic Studies*, 72, 395-427.
- Heckman, J. (1976). Estimates of a Human Capital production Function Embedded in a Life-Cycle Model of Labor Supply, in N. Terlecky, ed., *Household Production and Consumption* (New York: Columbia University Press). 227-264.
- Imai, S. and M. Keane (2004). Intertemporal Labor Supply and Human Capital Accumulation. *International Economic Review*, 45:2, 601-642.
- Keane, Michael (2009). Income Taxation in Life-Cycle Model with Human Capital, CEPAR Working Paper #2011-17, Australian School of Business, UNSW.
- Keane, Michael (2011). Labor Supply and Taxes, *Journal of Economic Literature*, 49:4, 961-1075.
- Keane, M. and R. Rogerson (2012). Micro and Macro Labor Supply Elasticities: A Reassessment of Conventional Wisdom, *Journal of Economic Literature*, 50, 464-76.
- Keane, M. and R. Rogerson (2011) Reconciling Micro and Macro Labor Supply Elasticities: A Structural Perspective, NBER Working Paper No. 17430.
- Keane, M. P. and K. I. Wolpin (2010). The Role of Labor and Marriage Markets, Preference Heterogeneity, and the Welfare System in the Life Cycle Decisions of Black, Hispanic and White Women. *International Economic Review*, 51:3, 851-892.
- Keane, M. and K. Wolpin (2009). Empirical Applications of Discrete Choice Dynamic Programming Models. *Review of Economic Dynamics*, January, 1-22.
- Keane M. and K. Wolpin (2007). Exploring the Usefulness of a Non-Random Holdout Sample for Model Validation: Welfare Effects on Female Behavior, *International Economic Review*, 48:4 (November), p. 1351-78.

- Keane, M. and K. Wolpin (2001). Effect of Parental Transfers and Borrowing Constraints on Educational Attainment. *International Economic Review*, 42:4, 1051-1103.
- Keane, M. and K. Wolpin (2000). Equalizing Race Differences in School Attainment and Labor Market Success. *Journal of Labor Economics*, 18:4, 614-52.
- Keane, M. and K. Wolpin (1997). The Career Decisions of Young Men. *Journal of Political Economy*, 105:3.
- Keane, M. and K. Wolpin (1994). The Solution and Estimation of Discrete Choice Dynamic Programming Models by Simulation: Monte Carlo Evidence. *Review of Economics and Statistics*, 76:4, 648-72.
- Keane, M., P. Todd and K. Wolpin (2010). Structural Estimation of Behavioral Models. *Handbook of Labor Economics: Vol.4*, O. Ashenfelter and D. Card (eds.), Elsevier Science B.V., Chapter 4, p. 333-463.
- Kimmel J. and Knieser T.J. (1998). New evidence on labor supply: employment vs. hours elasticities by sex and marital status. *Journal of Monetary Economics*, 42, 289-301.
- Lee, Donghoon and K. Wolpin (2010). Accounting for Wage and Employment Changes in the U.S from 1968-2000: A Dynamic Model of the Labor Market Equilibrium,” *Journal of Econometrics*, May.
- Lee, Donghoon and K. Wolpin (2006). Intersectoral Labor Mobility and the Growth of the Service Sector, *Econometrica*, January.
- Low, Hamish and D. Maldoom (2004). Optimal Taxation, Prudence and Risk-Sharing, *Journal of Public Economics*, 88, 443-464.
- MaCurdy, T. (1981). An Empirical Model of Labor Supply in Life-cycle Setting. *Journal of Political Economy*, 89:6, 1059-1085.
- McGrattan, E. and R. Rogerson (1998). Changes in Hours Worked Since 1950. *Federal Reserve bank of Minneapolis Quarterly Review*, Winter, 2-19.
- Purcell, P. (2009). *Older Workers: Employment and Retirement Trends*, Congressional Research Service, Washington, D.C., Available at: [www.crs.gov](http://www.crs.gov).
- Rogerson, R. and J. Wallenius (2009). Micro and Macro Elasticities in a Life Cycle Model with Taxes, *Journal of Economic Theory*, 144, 2277-2292.
- Saez, E., J. Slemrod and S. Giertz (2012). The Elasticity of Taxable Income with Respect to Marginal Tax Rates: A Critical Review. *Journal of Economic Literature*, 50:1, 3-50.
- Shaw, Katherine (1989). Life-cycle Labor Supply with Human Capital Accumulation. *International Economic Review*, 30 (May), 431-56.
- Todd, Petra and K. Wolpin (2008). Ex-Ante Evaluation of Social Programs,” *Annals of Economics and Statistics*, 91/92 (July/December), 263-291.

- Todd, Petra and K. Wolpin (2006). Assessing the Impact of a School Subsidy Program in Mexico: Using a Social Experiment to Validate a Behavioral Model of Child Schooling and Fertility, *American Economic Review*, 96:5 (Dec.), 1384-1417.
- Van der Klaauw, W. and K.I. Wolpin (2008). Social Security and the Retirement and Savings Behavior of Low Income Households. *Journal of Econometrics*, 145, 21-42.
- Wolpin, Kenneth (2007). Ex Ante Policy Evaluation, Structural Estimation and Model Selection, *American Economic Review*, 97:2 (May), 48-52.
- Wolpin, Kenneth (1996). Public-Policy Uses of Discrete-Choice Dynamic Programming Models, *American Economic Review*, 96:2 (May), 427-432.
- Wolpin, Kenneth (1992). The Determinants of Black-White Differences in Early Employment Careers, *Journal of Political Economy*, 100:3 (June), 535-560.

**Table 1: Short-Run Labor Supply Responses to Transitory and Permanent Tax Increases**

Age	Transitory Tax		Permanent Tax (Unanticipated)	
	Unanticipated	Anticipated	Uncompensated	Compensated
	~Frisch	Frisch	Marshall	Hicks
20	<b>-1.5%</b>	<b>-1.5%</b>	-0.7%	<b>-3.2%</b>
25	<b>-1.8%</b>	<b>-1.8%</b>	-0.6%	<b>-2.7%</b>
30	<b>-2.2%</b>	<b>-2.2%</b>	-0.6%	<b>-2.4%</b>
35	-2.6%	-2.6%	-0.5%	-2.3%
40	-3.2%	-3.3%	-0.7%	-2.3%
45	-3.8%	-4.2%	-1.0%	-2.8%
50	-4.7%	-5.3%	-2.3%	-4.2%
55	-6.2%	-7.2%	-5.3%	-7.2%
60	-8.6%	-9.8%	-9.4%	-10.5%

Note: All figures are contemporaneous effects of a 5% tax increase. The “transitory” increase only applies for one year at the indicated age. The “permanent” tax increases take effect (unexpectedly) at the indicated age and last until age 65. In the “compensated” case the proceeds of the tax (in each year) are distributed back to agents in lump sum form.

**Table 2: Long-Run Effects of Permanent Tax Increases**

Age	Age 25 (unexpected)			Age 30 (unexpected)			Age 35 (unexpected)		
	Hours	Wage	Assets	Hours	Wage	Assets	Hours	Wage	Assets
25	-2.7%								
30	-2.9%	-0.4%	+19.8%	-2.4%					
35	-3.2%	-0.7%	+26.3%	-2.7%	-0.3%	+12.4%	-2.3%		
40	-3.8%	-1.0%	+14.5%	-3.3%	-0.6%	+8.5%	-2.7%	-0.2%	+3.2%
45	-5.1%	-1.3%	+6.9%	-4.4%	-0.9%	+4.3%	-3.8%	-0.5%	+1.9%
50	-7.9%	-2.0%	+2.6%	-7.0%	-1.4%	+1.5%	-6.2%	-1.0%	+0.5%
55	-13.3%	-3.6%	-0.4%	-12.2%	-2.9%	-0.8%	-11.0%	-2.3%	-1.2%
60	-19.3%	-7.5%	-3.0%	-18.4%	-6.6%	-3.0%	-17.4%	-5.8%	-3.0%
65	-29.2%	-11.6%	-3.8%	-28.1%	-10.7%	-3.6%	-26.9%	-9.7%	-3.5%

Note: The tax increase is 5%. It takes effect (unexpectedly) at the age indicated in the column headings (i.e., age 25, 30 or 35) and lasts until age 65. The proceeds of the tax (in each year) are distributed back to agents in lump sum form.

**Table 3: Lifetime Effects of a Permanent Tax Increase on Labor Supply**

Age	Uncompensated	Compensated
20	-0.7%	-3.2%
30	-0.7%	-3.3%
40	-0.9%	-4.2%
45	-1.2%	-5.7%
50	-2.1%	-8.7%
60	-9.1%	-20.0%
Lifetime Hours (20-65)	-2.0%	-6.6%

Note: This table compares the baseline simulation of the Imai-Keane (2004) model with an alternative scenario where the tax rate on earnings is increased by 5%. The increase is in effect from the first period (age 20) until the terminal period (age 65). The table reports both the uncompensated case and the case where the proceeds of the tax (in each year) are distributed back to agents in lump sum form.