

NBER WORKING PAPER SERIES

A TAX-BASED ESTIMATE OF THE ELASTICITY
OF INTERTEMPORAL SUBSTITUTION

Jonathan Gruber

Working Paper 11945
<http://www.nber.org/papers/w11945>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2006

I am grateful to David Seif for exceptional research assistance and to Daron Acemoglu, Olivier Blanchard, Gary Engelhardt, Erzo Luttmer, Jonathan Parker, and Jim Poterba for helpful comments and conversations. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

©2006 by Jonathan Gruber. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

A Tax-Based Estimate of the Elasticity of Intertemporal Substitution
Jonathan Gruber
NBER Working Paper No. 11945
January 2006
JEL No. H2, E2, E6

ABSTRACT

One of the most important behavioral parameters in macroeconomics is the elasticity of intertemporal substitution (EIS). Starting with the seminal work of Hall (1978), researchers have used an Euler equation framework to estimate the EIS, relating the growth rate of consumption to the after-tax interest rate facing consumers. This large literature has, however, produced very mixed results, perhaps due to an important limitation: the impact of the interest rate on consumption or savings is identified by time series movements in interest rates. Yet the factors that cause time series movements in interest rates may themselves be correlated with consumption or savings decisions. I address this problem by using variation across individuals in the capital income tax rate. Conditional on observable characteristics of individuals, tax rate movements cause exogenous shifts in the after-tax interest rate. Using data on total non-durable consumption from the Consumer Expenditure Survey over two decades, I estimate a surprisingly high EIS of 2. This finding is robust to a variety of specification checks.

Jonathan Gruber
MIT
Department of Economics
E52-355
50 Memorial Drive
Cambridge, MA 02142-1347
and NBER
gruberj@mit.edu

One of the most important behavioral parameters in macroeconomics is the elasticity of intertemporal substitution (EIS). Within the standard life-cycle model on which much academic and policy analysis is based, the EIS summarizes the extent to which changes in the price of consumption between periods, the after-tax interest rate, affects the consumption growth across periods. Starting with the seminal work of Hall (1978), researchers have used an Euler equation framework to estimate the EIS, relating the growth rate of consumption to the after-tax interest rate facing consumers. This large literature has, however, produced very mixed results. Time series estimates of the EIS have ranged from Hall's (1988) finding of no effect of the rate of return on consumption growth, to Mulligan's (2002) finding of more than one. Micro-data work has also produced a variety of estimates as well, ranging from Dynan's (1993) estimate of 0.1 to Blundell, Browning and Meghir's (1994) estimates of 0.64 to 1.17.

The estimates of the EIS that have emerged from this literature suffer from an important limitation: the impact of the interest rate on consumption or savings is identified by time series movements in interest rates. Yet the factors that cause time series movements in interest rates may themselves be correlated with consumption or savings decisions. Indeed, the interest rate and consumption/savings decisions are the joint equilibrium outcome of the capital market, making it difficult to interpret causally the effect of one on the other. Instrumenting in a time series context by using lagged interest rates doesn't solve this problem if there are any slowly evolving technology or taste shocks that affect capital market equilibrium.

In other regressions of quantities on prices, these types of time series problems are solved by introducing exogenous cross-sectional variation in the price. In the case of consumption/savings, this is difficult, since the variation in interest rates across individuals is

unobserved. Some individuals will have higher rate of return savings opportunities than others, but those differences are typically unknown to the econometrician.

Yet, in the context of savings, there is a joint solution to both of the problems identified above: instrumenting the interest rate using cross-sectional and time series variation in capital taxation. The tax rate on capital income varies in observable, and exogenous, ways across individuals in a cross-section and over time, providing identification for the EIS.

Several previous papers have tried to use income taxation to identify the EIS (Zeldes, 1989; Lawrance, 1991; Shea, 1995).¹ Yet these papers have suffered from three important problems. The first is imprecision: once pure time series variation is removed by including year dummies, the estimates of the EIS are very imprecise. The second is identification: the factors that determine tax rates, primarily income, are likely also correlated with consumption growth. Finally, these papers have focused only on a particular component of consumption, food, and researchers such as Attanasio and Weber (1995) have convincingly argued that the results of Euler equation estimation can be biased by using such a limited measure of consumption.

In this paper, I move beyond the previous literature by using exogenous tax rate variation across individuals to cleanly identify the EIS. I do so by using the most comprehensive available source of consumption information, the Consumer Expenditure Survey (CEX). The CEX has repeated cross-sections of information on consumption and income over a long period during which there has been significant cross-sectional and time series capital income tax variation. In addition, the CEX provides a short panel across individuals which has been used successfully to test for the predictive ability of the life-cycle consumption model in other contexts, such as the

¹In addition, Mulligan (2002) uses capital tax variation in the time series to help identify

impact of anticipated and unanticipated tax rebates on consumption behavior. This panel has also been used to estimate the EIS, by Vissing-Jorgenson (2002), but not using capital income taxation variation.

An obvious concern with such an approach is that tax rates are a function of income, and income is a function of savings decisions past and present. To surmount this problem, I create “simulated” tax rates that are based only on exogenous characteristics such as education, age, and sex. I then control in a very detailed and flexible way for these exogenous characteristics in the regression specification, so that the effect of taxes is identified only by changes in the tax system over the 1980 to 2001 period.

The results of this analysis are striking. The estimated EIS is very large, larger than most estimates from the previous literature, and this estimate is robust to a wide variety of specification checks. For some specifications, the EIS from this tax-based approach is actually very similar to that from time series identification, but the tax-based approach is much more robust to variations in the empirical specification.

The paper proceeds as follows. In Part I, I provide some background on the EIS. Part II discusses the data and empirical strategy that I will employ in my effort to improve upon the past literature. Part III presents the results, and Part IV concludes with implications of the findings.

his model, as discussed below, but not variation in capital taxes across individuals.

Part I: Background

The Euler Equation and the EIS

Since the seminal work of Hall (1978), a large number of articles have recognized that, in the standard life cycle framework, the EIS can be obtained by estimating the first order conditions from the utility maximization problem, or the Euler Equation. Following the excellent exposition in Browning and Lusardi (1996), consider a household with a within-period utility function denoted by $v(C, \mathbf{Z})$, where C is consumption and \mathbf{Z} is a vector of factors that modify the impact of consumption on utility, such as household demographics, health status, or labor supply. Assuming intertemporal additivity, the optimal allocation between periods t and $t+1$ is described by the Euler Equation:

$$v_c(C_t, \mathbf{Z}_t) = E_t [\beta (1 + r_t) v_c(C_{t+1}, \mathbf{Z}_{t+1})]$$

where r is the real *after-tax* interest rate between the two periods, and β is the agent's discount rate. That is, individuals try to keep the discounted marginal utility of consumption constant across periods.

Most articles in this literature assume the isoelastic, or Constant Relative Risk Aversion (CRRA) form of the utility function:

$$V(C, \mathbf{Z}) = (1 / 1 - \gamma(\mathbf{Z})) * (C / \alpha(\mathbf{Z}))^{1 - \gamma(\mathbf{Z})}$$

where the parameter $\gamma(\mathbf{Z})$ is the coefficient of relative risk aversion, which by construction is independent of lifetime wealth but not necessarily independent of demographics. The function $\alpha(\mathbf{Z})$ is an equivalence scale of sorts; it shows how various demographic factors, such as family size or ill health, impact the utility of consumption.

Using this utility function in the Euler equation, we obtain:

$$\beta (1 + r_t) [\alpha(\mathbf{Z}_{t+1}) / \alpha(\mathbf{Z}_t)]^{\gamma(\mathbf{Z}) - 1} [C_{t+1} / C_t]^{-\gamma(\mathbf{Z})} = 1 + e_{t+1}$$

where e_{t+1} is a white noise error with $E_t(e_{t+1}) = 0$ and expected variance of σ_{t+1}^2 . Taking \mathbf{Z} to be a scalar, and parameterizing $\alpha(\mathbf{Z}) = \exp(\theta\mathbf{Z})$, and taking a second-order log-linearization, we now have the standard linearized Euler equation:

$$\Delta \ln C_{t+1} = \phi \ln(\beta) + \phi r_t + \lambda \Delta \mathbf{Z}_{t+1} + 0.5\phi \sigma_{t+1}^2 + u_{t+1}$$

where $\phi = 1/\gamma(\mathbf{Z})$, the inverse of the coefficient of relative risk aversion, $\lambda = [\theta(\gamma - 1)/\gamma]$, and u_{t+1} is a white noise error term. Thus, consumption growth is a function of four factors. The first is impatience: the lower β equates to higher impatience, so that consumption is large early in life and growing more slowly over time. The second is the change in the relative price of consumption across periods, the interest rate: the parameter ϕ is the proportional change in consumption for a one percent change in the price of consumption, or the EIS. The third is the change in demographics; for example, a new child in the family may raise consumption growth. The final term, the variance of consumption shocks, captures the precautionary motive for saving: as there is a higher variance of consumption shocks, individuals will save more, consuming less today and more in the future, for a higher rate of consumption growth.

Following Hall (1978), there is a large literature which estimates the Euler equation in an attempt to test the key predictions of this framework. This literature includes both aggregate time series studies and micro-data studies. Typically, this literature has *not* had as its primary focus estimating the magnitude of the EIS, but rather testing for the structural assumptions of the separable life cycle model. For example, Hall emphasized that predictable income growth (based on information available at time t) should not affect the growth rate of consumption. As reviewed in Browning and Lusardi (1996), many articles have tested this proposition, with mixed

results. Yet neither Browning and Lusardi's review, nor most of the articles in this literature, spend much time discussing the estimates of the EIS - perhaps for the reasons discussed below.

Time Series Evidence on the EIS

Estimates from the time series approach to estimating the EIS vary. At one end, Hall (1988) assessed whether consumption growth over five year intervals in the U.S. corresponded negatively to the real interest rate during those same periods. He found no such correspondence, and he rejects earlier studies which found such a correspondence. He concludes that there is no evidence for a sizeable elasticity of intertemporal substitution. The most recent time series approach is that of Mulligan (2002), who argues that the previous time series studies suffered from using a particular financial rate of return, and not the overall after-tax return to capital investments. Using the after-tax rate of return to capital, Mulligan estimates a very large elasticity of intertemporal substitution of over 2.

The wide variability in these time series studies likely reflects the difficulty of disentangling the impact of the interest rate from other factors that affect savings and consumption decisions. Any time series factors that change taste for consumption will affect both consumption decisions and, in equilibrium, the interest rate. Many studies attempt to overcome this problem by instrumenting current interest rates with lagged interest rates. But this only solves the problem if time series shocks to the interest rate are sharp and dissipate quickly. Any slowly evolving factors that affect the tastes for consumption will have multi-period effects on both consumption and interest rates and bias the estimates in this context.²

²Mulligan goes the farthest in attempting to overcome these problems, using in his

Microdata Estimates: Two Basic Strategies

A large micro-data literature has also emerged to estimate the EIS. As noted earlier, the purpose of these articles is not to estimate the EIS per se, but rather to test the assumptions underlying the additive utility function that motivates Euler equation estimation. Perhaps as a result, this literature provides a broad range of estimates for the EIS, from 0.1 to over 1.

These studies differ along a number of dimensions, but it is useful to categorize the studies into two types. The first is studies that use panel data on food consumption from the Panel Study of Income Dynamics (PSID) to estimate the Euler equations (e.g. e.g. Zeldes, 1989; Lawrance, 1991; Shea, 1995). These studies identify the effect of the after-tax interest rate by using variation in capital tax rates across individuals. As Lawrance (1991) highlights, such models must include year dummies for the identification to come from tax rate variation and not from time series factors which are also correlated with consumption. Unfortunately, in most of these studies, once time series dummies are included the EIS is not well identified, leaving the authors unable to rule out estimates of the EIS of either zero or well above one.

An additional problem with these studies is that they focus on only one component of consumption, food. As Attanasio and Weber (1995) highlight, non-separability between food and other types of consumption can seriously bias estimates of the Euler equation based solely on food consumption; Browning and Lusardi (1996) state that “Most widely used additive separability assumptions are dubious at best” (p. 1827).

instrument set the tax wedge on capital. Unfortunately, however, his models using just the tax rate are not precise, and he does not estimate a statistically significant elasticity of intertemporal substitution using this instrument set. Moreover, time series movements in the tax wedge may

The second strand of literature in this area therefore takes a different approach, using surveys on total consumption expenditures from the U.S. (the CEX) and the U.K. (the Family Expenditure Survey, FES) (e.g. Attanasio and Weber, 1993 and 1995; Blundell, Browning and Meghir, 1994; Attanasio and Browning, 1995). These articles use the consumption data to construct “synthetic” panels of birth cohorts, and then to model how consumption evolves over time for these birth cohorts. These articles typically produce fairly large and precise estimates of the EIS, generally around 0.75.

While this literature has the advantage of focusing on total consumption, there is an offsetting disadvantage: the identification of the interest rate effect on consumption growth once again comes purely from time series variation in the interest rate (instrumenting the rate with the twice lagged interest rate). Moving to a micro-data context does not solve the problems with this identification strategy.

This point is highlighted by the provocatively titled paper by Carroll (1997). Carroll notes that estimates of the linearized Euler equation make two important assumptions: they assume that both the discount rate term and the consumption variance term are constant, or at least do not vary systematically across agents in a way which is correlated with other parameters of interest (such as the interest rate). Carroll criticizes these assumptions, particularly the latter. As he highlights, the omission of the consumption variance term can bias estimates of the EIS: a higher interest rate will lead to higher wealth holdings, which reduces the variance of consumption, lowering the growth rate of consumption. This offsets the EIS effect (a higher interest rate leads to a faster growth rate of consumption), so that the estimated EIS is biased

also not be exogenous, for example if government policy reflects tastes for consumption.

downward. Blundell et al. (1994) present a macro-economic version of this critique: if the macroeconomic policies that cause the rise in the interest rate also trigger a change in the variance of consumption, it will similarly lead to biased EIS estimates.

While most previous papers recognize this issue, they typically assume that it is dealt with through instrumenting the interest rate with lagged interest rates or using tax variation, but this approach does not surmount the problem raised earlier: lagged interest rates may not be independent of the conditional variance of consumption. This point is illustrated in the simulation analysis of Guvenen (2003). He builds a simulated macroeconomy, and then estimates a macroeconomic euler equation using lags as instruments. He then reestimates the model including as well the conditional variance, and his estimated EIS doubles, suggesting that lagged variables are not independent of the conditional variance of consumption.

Two of the Euler equation articles more directly deal with this concern.³ The first is Zeldes (1989), who includes in his estimates family fixed effects to capture the unobserved consumption variance across families. These fixed effects will help with the problem, but will not fix it if changes in income (his instrument for the after-tax interest rate) are associated with changes in the variance in consumption, as seems possible.

The second is Dynan (1993), using the quarterly changes in consumption available in the Consumer Expenditure Survey (CEX), the source of consumption changes used for my analysis. She explicitly includes the squared growth rate in consumption in her regression, and uses a variety of instruments for consumption variance: education, occupation, industry, number of

³This is in addition to the very large literature that focuses explicitly on estimating the impact of uncertainty, or insurance against uncertainty, on consumption, such as Carroll (1994); see Attanasio (1998) for a review.

earners in the household, and capital income (as a proxy for financial assets). Doing so, she obtains very small and imprecise estimates of the EIS. These instruments, however, seem unlikely to be independent of other factors determining consumption growth, such as patience (e.g. more patient and less patient individuals will obtain different levels of education and choose different occupations).

Part II: Data and Empirical Strategy

Data

My main source of data for this analysis is the Consumer Expenditure Survey (CEX). The CEX is a nationally representative quarterly survey of households in the U.S. that gathers up to four quarters of data on the entire bundle of consumption expenditures by households, as well as income and demographic information. This is the most comprehensive source of micro-data on consumption patterns in the U.S. The sample varies from 5-10,000 households per year, and data are available on a continuous basis since 1980.

The CEX surveys each household in the sample for up to five quarters. The first interview is used to gather baseline information on household demographics and income. Each of the subsequent quarterly interviews collects data on a roster of household consumption categories. As well, the fifth and final interview collects a new round of data on income. Surveys between the second and fifth track changes in household demographics. I use data for the CEX interviews from 1980 through 2001, which covers the last quarter of 1979 through the first quarter of 2002. I only use those households who report consumption information in all four interviews.

Non-durable consumption is defined following Autor, Katz and Kearney (2003), and includes expenditures on: food and alcohol; apparel; fuel, utilities and public services; household operations; non-durable vehicle expenses; entertainment; non-durable shelter expenses; and all other non-durable goods and services. This category excludes two types of consumption goods from total consumption. The first is durable expenditures, primarily durable spending on home and vehicles. The other is transfer spending, which includes insurance expenditures, charitable gifts, and other transfers to others. The first of these categories is quite large, and the line between durable and non-durable expenditures is a vague one; I therefore consider the sensitivity of the results to the consumption definition below. Consumption is deflated by a month-specific CPI to 2002.

I match information on the after-tax rate of return to these data. I use two measures of the rate of return. I first follow the traditional EIS literature in using the t-bill rate, measured over 9 months to reflect the periodicity of consumption changes in my data. For the after-tax t-bill rate, I multiply by one minus the marginal tax rate on interest income.

Recent work has emphasized, however, that this may not be the right rate to use for pricing consumption decisions over time. One line of argument is presented by Mulligan (2002), who argues that intertemporal decisions will reflect the overall productivity of capital in the economy. Another is presented by Vissing-Jorgenson (2002), who argues that with limited asset participation individuals will face different intertemporal prices. To account for this concern, I also use returns (and tax rates) specific to the savings vehicles used by different income groups over time. Specifically, I compute the income-group specific distribution of asset holdings, and take a weighted average of the after-tax rates of return to each type of asset, assuming that the

marginal rate of return equals the average rate of return on that income group's portfolio.

I do so using the Survey of Consumer Finances (SCF), the best data source for information on the asset holdings of U.S. families. I use the comparable data that is provided in the surveys in 1989, 1992, 1995, 1998 and 2001 to compute the asset mix of portfolio holdings for families by income decile in the U.S.⁴ I then match that asset mix by income decile back to the CEX survey data, assuming asset mixes are constant before 1989, then evolving according to the asset mix data in each of the SCF surveys over time. I divide assets in 8 categories: corporate bonds held in taxable accounts; government bonds held in taxable accounts; municipal bonds; bonds held in tax-deferred (retirement account) forms; equity held in taxable forms; equity held in tax-deferred forms; cash held in taxable accounts; cash held in tax-deferred accounts. For each category, I then assign a corresponding rate of return: the corporate bond rate of return for corporate bond holdings; the T-bill rate for federal government bond holdings; the municipal bond rate of return for municipal bond holdings; the rate of return on certificates of deposit for cash holdings; and the dividend and capital gains rates of returns for equity holdings⁵ All rates of return are deflated by the consumer price index for the relevant time period.

To turn these pre-tax rates of return into after-tax rates of return, I assign each household a tax rate using the NBER's TAXSIM model. I use the household's total income over the year

⁴Comparable asset holdings data are not available in the earlier SCF surveys.

⁵I assume tax-deferred bond holdings are corporate and use the corporate bond rate of return. For all assets except equity, I use the rate of return as of the month of the first interview (for EIS regression) or as of the beginning of the year (for consumption/savings regression). For equity, I use the rate of return for the period ending with the first interview (or first month), under the assumption that individuals expect that same rate of return over the coming period. I treat inflation the same way, assuming that past inflation proxies for inflation over the coming period. I discuss sensitivity to this treatment below.

that precedes the first interview, as well as information on their family structure and state of residence at the first interview to compute their marginal tax rate on different forms of capital income: interest income, dividend income, and capital gains income. The tax rates on dividend and interest income are very highly correlated, but not perfectly due to different state-level treatment of these forms of income in some state income taxes. The tax rate on capital gains differs substantially due to tax preferences to capital gains. Moreover, the effective tax rate on capital gains will be much lower than the statutory tax rate on capital gains, since those gains are only taxed on realization (or not at all if passed on at death). Absent a convincing estimate of the effective capital gains tax rate, I assume that the effective rate is 50% as large as the statutory rate

I then combine the relevant tax rate with the relevant asset's return. For taxable bond and cash accounts, I use the tax on interest income; for tax-deferred bonds/cash and municipal bonds, I assume no income taxation. For taxable equity, I divide the return into its capital gains and dividend components, and assign the relevant tax rate to each.

Empirical Framework

I use these data to estimate a standard log-linearized Euler equation:

$$(1) \quad GC_{it,t+1} = \alpha + \beta ATRATE_{it} + \mathbf{X}_{it} \delta + \Delta \mathbf{Z}_{it,t+1} \eta + \epsilon$$

where GC is the growth rate of consumption for household i from period t to $t+1$; ATRATE is the after-tax interest rate that applies to savings between period t and $t+1$; \mathbf{X} is a vector of baseline characteristics of the household in period t ; and $\Delta \mathbf{Z}$ is a vector of demographic changes between periods t and $t+1$. The \mathbf{X} vector contains: dummies for 5 year age categories for the

head and (if present) spouse; dummies for education of head and spouse (HS dropout, HS grad with no college; some college; college grad); dummies for the number of children in the household; dummies for the number of elderly in the household; and dummies for the number of non-elderly adults in the household. The ΔZ vector contains the change in the number of children, adults, and elderly in the household, and the change in the marital status of the head. This is similar to the approach of Dynan (1993) and Vissing-Jorgenson (2002) using these CEX data. In addition, recent papers by Parker (1999), Souleles (1999), and Johnson, Parker and Souleles (2004) use a similar approach to assess the response of consumption to anticipated and unanticipated tax refunds.

The discussion above highlighted three potential problems with this regression. First, much of the variation in after-tax rates of return is time series variation which may not be a legitimate source of identification of the EIS or the interest elasticity. Second, cross-sectional variation in both taxes and the pre-tax rate of return (in models where that variation is included) is driven largely by income differences. Since income is mechanically a function of tastes for consumption through capital income, this can bias the estimated effects of capital taxes on consumption. Finally, even non-capital income differences are likely correlated with omitted determinants of both the level and growth rate of consumption, such as consumption growth uncertainty.

I deal with the first of these problems through an instrumental variables strategy: instrumenting the after-tax rate of return with the tax rate on interest income, while controlling

for a full set of year dummies.⁶ By doing so, I use only the variation in tax rates across individuals to identify the model, and not time series movements in interest rates.

To deal with the second problem, that income that is endogenously determined by the taste for consumption, I use in my instrumental variables strategy not actual income, but *predicted* income. I first predict for each household their income, based on marital status interacted with sex of the household head; education of the household head and (if present) spouse; age dummies for head and spouse; race of head; and dummies for each state of residence. To ensure sufficient sample size, I estimate such a model on three year rolling panels of the CEX data, and apply the predictions to the middle year. These prediction equations perform fairly well, with R-squareds in the range of 0.42. I then use that predicted income to compute the various tax rates that go into computing the after-tax interest rate. In this way, I have a tax rate measure which is independent of unobserved tastes for consumption.

This tax rate remains, however, a function of observed factors which may also be correlated with the taste for consumption. In particular, the tax rate is most importantly a function of predicted income, tax filing status (joint, single, or head of household), state of residence and year, and interactions of these factors. I therefore include in the \mathbf{X} vector a set of controls for each these factors. A particular concern is that controlling for linear income, or even a fairly restrictive non-linear form, is not sufficient, as the non-linearities in the tax schedule may correspond to non-linearities in tastes for consumption. I therefore include a full set of 100 dummies in the regression for each point in the income distribution, so that identification is not coming from cross-sectional differences in income. I also interact these 100 dummies with a

⁶I also tried a set of instruments that includes as well the tax rate on capital gains and on

dummy for each of the three tax filing statuses, to allow for different effects of income on taxes for different types of families that face different tax schedules. I also include a full set of state fixed effects in the model to capture any fixed state differences in tastes for consumption. The model is therefore identified only by changes in state taxes that deviate from the national trend, or changes in national taxes that have differential impacts along the income distribution. In the estimation below, I consider further refinements of this identification strategy.

To summarize, this strategy deals with the concerns raised about previous papers in the EIS literature. I use total non-durable consumption, to address concerns about non-separability of consumption. I instrument for the interest rate and include a full set of year dummies, to address identification concerns about the co-movement of interest rates and consumption growth. Identification comes from a tax rate based on predicted, not actual income, to avoid endogeneity of income to the factors that determine consumption growth. I then control in a very detailed way for predicted income so that identification comes only from tax changes, not from income differences.

There is one aspect of criticism of the previous EIS literature that I cannot fully address, however, which is the omission of the variance of consumption. It is possible that tax changes not only impact the after-tax rate of return, but also the anticipated variance of consumption. In particular, a rise in the tax rate will in general lower the variance of consumption, since the government absorbs a higher share of income risk. Thus, higher tax rates lead to both slower consumption growth by lowering the after-tax rate of return, and slower consumption growth by lowering the degree of consumption uncertainty. As a result, there will be an upward bias to my

dividend income, and the results are very similar.

estimate of the EIS: it will incorporate both the effect of taxes on the rate of return and the effect of taxes on the variance of consumption growth. With only one instrument, the tax rate, I cannot identify both effects. I suggest an approach to addressing this issue below.

Table 1 presents the means of the data. The mean level of non-durable consumption in the data is \$25,268, and consumption growth over a nine-month period averages 1.1%.

Part III: Results

Basic Results

Table 2 presents the results of estimating equation (1). The regression include all the control variables described above, but I show here only the coefficient of interest, that on the interest rate.

The first column shows the result of OLS estimation on the after-tax T-bill rate. This model follows previous work in excluding year dummies and only including a linear time trend. The estimated coefficient here is wrong-signed, statistically significant and rather large, implying an EIS of -0.55. In the second column I show the results of OLS estimation where the after-tax T-bill rate is replaced with the broader after-tax rate of return measure described above. In this case, the EIS estimate is right-signed, significant, but small, implying an EIS of only 0.1.

The next row shows the results of estimating the model where the interest rate is instrumented by its twice-lagged value, and once again a linear year trend is included rather than year dummies. Using the after-tax T-bill rate in the first column, I estimate a very large and significant EIS of 2.6. This is well above the rate estimated in the earlier cross-sectional literature, and at the high end of the time series literature as well. On the other hand, moving to

the broader after-tax rate of return measure in the second column, the estimate is a much smaller 0.33. So the result from this method is sensitive to the definition of the rate of return.

The last column shows the results using my preferred identification strategy, instrumenting the after-tax interest rate by the capital income tax rate, and thereby using variation in tax rates to identify the EIS (while including year dummies to control for general changes in taxes and consumption growth over time). Using the after-tax T-bill rate in the measure of the interest rate, the estimated EIS is 2.03, which is fairly similar to that obtained using the time series approach, but is very different from OLS. Unlike the time-series approach in the second row, however, using this tax-based identification strategy yields a very similar estimate when using the after-tax rate of return measure in the second column, which is a slightly larger 2.24.

Summing up, these results suggest that there is a very large EIS. When instrumenting by lagged interest rates, the EIS using the after-tax T-bill rate is over 2. When the most reliable tax instrument is used, the EIS is above 2 regardless of the measure of interest rate. This large estimate is higher than most cross-sectional estimates, and is in line with Mulligan's (2002) time series estimates that use total returns to capital are partly identified by capital income tax variation. The estimates are fairly imprecise, however; while I can rule out OLS estimates, I can't rule out EIS estimates in the range of 0.75 estimated in earlier work.

Specification Checks

This large EIS estimate is striking, and bears further scrutiny. In this section, I subject the results to that scrutiny.

Sensitivity to Consumption/Savings Definition

The first variation I consider is the measurement of consumption. The measure of non-durable consumption used here excludes primarily expenditures on home and vehicles, but it also includes other consumer durables such as entertainment goods. These “quasi-durables” may not properly be viewed as savings, so it is important to assess the sensitivity of the results to their treatment. I therefore recreate consumption measures by shifting these quasi-durables into the non-durables category. On average, this results in an increase in non-durable consumption of \$1658.

The results of doing so are shown in the second row of Table 3; the first row replicates the baseline results from Table 2 for comparison. There are three panels of two columns in this Table: each panel represents a different estimation strategy (OLS, instrumenting by lagged interest rates, or instrumenting by tax rates), and each column within the panel represents a different measure of the interest rate. Each coefficient in the table is from a separate regression.

The second row of the Table shows that this change in definition has little effect on the OLS or lagged IV estimates. The tax IV estimates do shrink somewhat, but remain significant and close to 2.

The other difficult definitional issue is the treatment of transfers, such as insurance spending. These categories might properly be viewed as non-durable consumption as well, depending on the utility flow from purchasing insurance or transferring to others. In the next row of Table 3, I include transfers in the non-durable consumption measure used for defining consumption growth. The effects here are relatively similar to the previous row, and consistent

with the basic findings of the first row.

Finally, another concern with these data are outlying observations for consumption growth. There are a number of observations with extreme positive or negative consumption growth rates over this nine-month period, and very large movements are very unlikely to be related to the after-tax interest rate. I therefore assess the sensitivity to censoring any observations with a consumption growth rate below -100% or above 100%, setting these extreme values to either -100% or 100%. As the fourth row of the Table shows, this has relatively little effect on the results.

Construction of the After-Tax Interest Rate

Constructing the after-tax interest rates used here, particularly the overall rate of return measure, requires some judgement calls. Most importantly, for assigning inflation rates and rates of return to equities, I use past returns as a predictor of future returns, rather than assuming that individuals have perfect foresight. The alternative may be a natural assumption.

The fifth row of Table 3 shows the results from using this alternative construction. In the first column of each panel, which uses the after-tax T-bill rate, the only change here is a move from using the previous 9 months inflation rate to the actual realized 9 month inflation rate; in the second column, I also use such a forward looking approach for calculating the rate of return on equities in my overall weighted average rate of return. This change has dramatic effects on the OLS, and in particular the lagged interest rate IV results. The OLS result in the first column get much more negative. The lagged IV results become very peculiar, with an enormous positive EIS for the after-tax T-bill rate, and an enormous negative EIS for the after-tax rate of return.

For the tax IV, however, the results are very robust to this change in definition.

Omitted Income-Group Specific Trends

As noted earlier, a key advantage of my approach over the previous literature is that I am able to control for general changes over time in consumption/savings that might be correlated with interest rate movements. But I still use tax changes that differ along the income distribution for identification of the EIS. This implies that I am assuming that differential changes over time in tastes for consumption/savings by income group are not correlated with tax rate changes for those same income groups (e.g. that higher income groups are not becoming more patient at the same time their taxes are being cut). This seems a reasonable assumption, but it bears testing.

I do so in the next row of Table 4. In this row, I include in the model not only 297 dummies for each income group*filing status, but also another 297 interactions of these dummies with a linear time trend. This allows for evolving tastes for consumption-savings across income groups that might be correlated with evolutions in capital income tax rates.

Adding these regressors has virtually no impact on the OLS or lag IV estimates. The effect on the tax IV estimates are also fairly modest: the after-tax T-bill estimate goes down somewhat, and the after-tax rate of return estimate goes up somewhat. The estimates also become less precise, but remain statistically significant. Thus, the estimated EIS appears robust to trends by income group over time.

Controlling for Labor Income Tax Rates

As discussed above, one key limitation of most previous approaches to estimating the EIS

is the exclusion of the variance of consumption. This is also potentially problematic in my context. If a rise in taxes lowers the variance of consumption (effectively through government risk-sharing), it will slow the rate of consumption growth, leading to an upward bias to the EIS. With only one instrument, the capital income tax rate, I cannot identify both the rate of return and consumption variance effects of tax changes.

I can, however, address this concern by noting that what drives the variance in after-tax income, and therefore consumption, is mostly variation in *labor* income tax rates, while what matters for identifying the EIS is *capital* income tax rates. Therefore, to the extent that there is independent variation in capital income tax rates, I can include in the model the labor income tax rate to capture the general effect of taxation on consumption growth, and use the remaining independent variation in capital income tax rates to identify the EIS. This control is not perfect, in that some of the taxation effects on the variance of consumption arise through capital taxation. But since most of the variation for most taxpayers arises through labor taxation, if the results are robust to including the labor income tax rate, it suggests that omitted variance terms are not important.

I include the labor income tax rate in the final row of Table 3. As one would expect, this has very little effect on the OLS and lag IV estimates. Strikingly, however, it also does not much effect the tax IV estimates either. The standard error of these estimates rises, but the estimated effects are very similar. This is quite a strong test, since it says that conditional on labor taxation, capital taxation in particular leads to slower consumption growth.

Part IV: Discussion and Conclusions

The elasticity of intertemporal substitution is one of the central parameters necessary for evaluating standard models of consumption behavior. Yet, despite the importance of this parameter, there has been to date relatively little reliable evidence on its magnitude. In this paper, I present a tax-based approach to estimating the EIS that addresses important problems with the previous literature. Throughout a variety of specifications, the results are consistent and clear: the EIS is large, with a value around 2.

What are the implications of this large value? One potential implication is for the effects of tax policy on savings. In the canonical life cycle model of Summers (1981), such a large EIS would imply an enormous impact of tax policy on savings: even with an EIS of 0.33, Summers estimated that the interest elasticity of savings was almost three. But even within such a model, the impact of the EIS will depend critically on the other parameters that determine savings behavior, in particular the time preference rate, as highlighted by Evans (1983). For example, the implied savings elasticity for an EIS of 0.33 varies from 0.23 to 2.97 as Evans varies the discount rate, productivity growth rate, and population growth rate in simulations of a life cycle savings model. Thus, we cannot necessarily infer a large interest elasticity of savings from this large EIS, although for most parameter values that would be the result.

More problematically, there is a growing consensus that the life cycle model may not be the right description of savings behavior for much of the U.S. population. Elmendorf (1986) simulates the interest elasticity in a variety of alternative models of the savings decision, such as with and without bequests, with and without “savings targets”, and with and without short planning horizons. Unsurprisingly, the implied interest elasticity varies widely across these

models, including negative interest elasticities in some cases. If the life-cycle model is not the right one, it is unclear how the EIS estimated from this Euler equation framework translates to an implied effect of interest rate changes on savings.

Second, in the standard life-cycle model developed earlier, the EIS is the inverse of the coefficient of relative risk aversion (CRRA). My findings therefore imply a very small CRRA of roughly 0.5. This is at odds with much of the prior literature estimating the CRRA. But this may once again reflect the limitations of the structural model relating the EIS to the CRRA; in other models, such as Epstein-Zin (1989) preferences, this direct inverse relationship does not hold. To the extent that both my estimates and those of the CRRA literature are credible, this may offer support for such an alternative formulation.

Thus, this paper provides a useful component to the larger evaluation of the determinants of savings in the U.S. But such a richer evaluation is required to draw strong conclusions about the effect of interest rates and tax policies on savings decisions and risk-taking.

References

- Attanasio, Orazio and Browning, Martin (1995). "Consumption Over the Life Cycle and Over the Business Cycle," *American Economic Review*, 85, 1118-1137.
- Attanasio, Orazio and Weber, Guglielmo (1993). "Consumption Growth, the Interest Rate, and Aggregation," *Review of Economic Studies*, 60, 631-649.
- Attanasio, Orazio and Weber, Guglielmo (1995). "Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey," *Journal of Political Economy*, 103, 1121-1157.
- Autor, David, Lawrence Katz and Melissa Kearney (2003). "U.S. Wage and Consumption Inequality Trends: A Re-Assessment of Basic Facts and Alternative Explanations," mimeo, MIT.
- Blundell, Richard, Martin Browning and Costas Meghir (1994). "Consumer Demand and the Life-Cycle Allocation of Household Expenditures," *Review of Economic Studies*, 61, 57-80.
- Browning, Martin and AnnaMaria Lusardi (1996). "Household Saving: Micro Theories and Micro Facts," *Journal of Economic Literature*, 34, 1797-1855.
- Carroll, Christopher (1997). "Death to the Log-Linearized Euler Equation! (And Very Poor Health to the Second-Order Approximation)," NBER Working Paper #6298.
- Dynan, Karen (1993). "How Prudent Are Consumers?," *Journal of Political Economy*, 101, 1104-1113.
- Elmendorf, Douglas (1996). "The Effect of Interest-Rate Changes on Household Saving and Consumption: A Survey," mimeo, Federal Reserve Board.
- Epstein, Larry, and Stanley Zin (1989). "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica*, 57, 937-969.
- Evans, Owen (1983). "Tax Policy, the Interest Elasticity of Saving, and Capital Accumulation: Numerical Analysis of Theoretical Models," *American Economic Review*, 73, 398-410.
- Guvenon, Fatih (2001). "Mismeasurement of the Elasticity of Intertemporal Substitution The Role of Limited Stock Market Participation," mimeo, Rochester University, 2001.
- Hall, Robert (1978). "Stochastic Implications of the Life Cycle - Permanent Income Hypothesis:

- Theory and Evidence,” *Journal of Political Economy*, 86, 971-987.
- Hall, Robert (1988). “Intertemporal Substitution in Consumption,” *Journal of Political Economy*, 96, 339-57.
- Johnson, David, Jonathan Parker, and Nicholas Souleles (2004). “Household Expenditure and the Income Tax Rebates of 2001,” mimeo, August.
- Lawrance, Emily (1991). “Poverty and the Rate of Time Preference: Evidence from Panel Data,” *Journal of Political Economy*, 99, 54-77.
- Mulligan, Casey (2002). “Capital, Interest, and Aggregate Intertemporal Substitution,” NBER Working Paper #9373.
- Parker, Jonathan (1999). “The Reaction of Household Consumption to Predictable Changes in Social Security Taxes,” *American Economic Review*, 89, 957-973.
- Shea, John (1995). “Union Contracts and the Life-Cycle/Permanent-Income Hypothesis,” *American Economic Review*, 85, 186-200.
- Souleles, Nicholas (1999). “The Response of Household Consumption to Income Tax Refunds,” *American Economic Review*, 89, 947-958.
- Summers, Lawrence (1981). “Capital Taxation and Accumulation in a Life Cycle Growth Model,” *American Economic Review*, 71, 533-544.
- Vissing-Jorgenson, Annette (2002). “Limited Asset Market Participation and the Elasticity of Intertemporal Substitution,” *Journal of Political Economy*, 110, 825-853.
- Zeldes, Stephen (1989). “Consumption and Liquidity Constraints: An Empirical Investigation,” *Journal of Political Economy*, 97, 305-346.

Table 1: Means of the CEX Data		
Variable	Mean	Standard Deviation
Non-Durable Consumption	25,268	15,306
Consumption Growth	0.011	0.394
After-Tax Interest Rate (9 month)	0.039	0.049
Capital Income Marginal Tax Rate	0.236	0.094

Notes: Means and standard deviations of variables described in text.

Table 2: Base Case Estimates		
	After-tax T-Bill Rate	After-tax Rate of Return
OLS, no year dummies	-0.551 (0.116)	0.105 (0.032)
Lag IV, no year dummies	2.616 (0.490)	0.328 (0.130)
Tax IV	2.032 (0.796)	2.239 (0.894)
Number Obs	66314	66208

Notes: Estimates from models such as equation (1) in text. Each cell represents the estimated EIS from a separate model: first column uses after-tax T-bill rate, while second column uses weighted average after-tax rate of return. Standard errors in parentheses.

Table 3: Alternative Specifications						
	OLS		Lag IV		Tax IV	
	T-Bill	Rate of Return	T-Bill	Rate of Return	T-Bill	Rate of Return
Base Case	-0.551 (0.116)	0.105 (0.0320)	2.616 (0.490)	0.328 (0.130)	2.032 (0.796)	2.239 (0.894)
Include Quasi-Durables	-0.613 (0.121)	0.101 (0.034)	2.368 (0.513)	0.320 (0.136)	1.647 (0.835)	1.792 (0.931)
Include Transfers	-0.604 (0.105)	0.069 (0.032)	2.187 (0.450)	0.236 (0.129)	1.539 (0.788)	1.706 (0.886)
Censor Outliers	-0.288 (0.108)	0.119 (0.030)	2.614 (0.456)	0.325 (0.121)	1.809 (0.748)	1.993 (0.838)
Use Forward-Looking Rate	-0.869 (0.119)	0.072 (0.028)	7.657 (1.342)	-3.173 (1.060)	1.738 (0.720)	1.916 (0.778)
Add Trends	-0.312 (0.109)	0.118 (0.030)	2.801 (0.489)	0.327 (0.123)	1.766 (0.922)	2.353 (1.254)
Control for Labor Income Tax Rate	-0.546 (0.116)	0.107 (0.032)	2.345 (0.470)	0.322 (0.130)	1.733 (0.982)	2.359 (1.324)
Year Dummies?	No	No	Yes	No	No	Yes

Notes: Estimates from models such as equation (1) in text. Each cell represents the estimated EIS from a separate model. First panel shows models estimated by OLS; second panel shows models instrumenting the interest rate with the 24 month lagged rate; and the third panel shows models instrumenting with the tax rate. Within each panel, first column uses after-tax T-bill rate, while second column uses weighted average after-tax rate of return. Standard errors in parentheses.