Economica (1996) **63**, 495–512

The Intertemporal Elasticity of Substitution: An Exploration using a US Panel of State Data

By PAUL BEAUDRY* and ERIC VAN WINCOOP†

* *University of British Columbia, CIAR and NBER.* † *Boston University*

Final version received 5 December 1994

This paper uses state-level consumption data to estimate the intertemporal elasticity of substitution of consumption (IES). In contrast to the results of Hall (1988) and Campbell and Mankiw (1989), we provide evidence indicating that the IES is significantly different from zero and probably close to one. Since inference about the IES in the context of the standard Euler equation is problematic as a result of mis-specification bias, we cast most of our discussion in the context of the framework developed by Campbell and Mankiw. This modifies the Euler equation in that a fraction of agents simply consume their income. The use of panel data to examine the relationship between interest rates and consumption growth has two advantages. First, we achieve a significant increase in precision, which in particular allows us to rule out a zero IES. Second, we can use the panel aspect of the data to bypass asset return measurement problems. In particular, we identify a common time component in expected consumption growth that is associated with movements in interest rates when the IES is positive.

INTRODUCTION

The intertemporal elasticity of substitution of consumption (IES) plays a key role in innumerable policy evaluations. Much of the literature that focuses on this parameter is based on the estimation of the first-order conditions associated with intertemporal optimization (the Euler equation approach). However, as has been emphasized by, among others, Mankiw (1981), Summers (1982), Hansen and Singelton (1983) and Mankiw *et al.* (1985), most of the existing estimates are problematic, since they rely on specifications where the over-identifying restrictions are rejected by the data. Therefore, without any firm knowledge of the extent to which such mis-specification biases estimates of the IES, we can infer nothing from the standard Euler equation formulation with respect to whether consumption is or is not sensitive to interest rates.

In an attempt to offer a characterization or consumption behaviour that is more consistent with the data, Campbell and Mankiw have explored a modification to the permanent income hypothesis in which a fraction of consumers follow the rule of thumb consisting of consuming their income.[1] Under this alternative specification, Campbell and Mankiw (1989) find that an IES of zero cannot be rejected by the data. Their results clearly indicate that the mis-specification of the standard consumption–Euler equation may result in substantially misleading inferences with respect to the relationship between consumption growth and the interest rate.[2] In fact, their results suggest that the main reason why expected consumption growth appears correlated with interest rates is that interest rates are correlated with expected income growth. While the Campbell–Mankiw specification may lack theoretical foundation, their setup is attractive in that over-identifying restrictions are not rejected.

Therefore their findings cast serious doubt on the usefulness of the vast literature that postulates that consumption is strongly affected by interest rates.

The goal of this paper is to use a panel of state-level data for the United States to re-evaluate whether or not consumption is sensitive to the interest rate. In particular, we want to explore the issue in a manner that is robust to possible mis-specifications inherent to the standard consumption–Euler equation. Hereby we cast most of our discussion within the framework developed by Campbell and Mankiw, since we believe it offers an important benchmark against which results can be compared.[3]

The use of panel data to evaluate the relationship between interest rates and consumption growth is attractive for two reasons. First, using state-level data provides more precise estimates of the IES than using nationwide aggregate data. The gain in precision is of particular importance when we adopt the specification of consumption behaviour that includes the possibility of rule-of-thumb consumers. We show that, when using aggregate US data, estimates of the IES are very imprecise when rule-of-thumb consumers are permitted. For example, using nondurable goods as the measure of consumption, almost any value between 0 and 1·5 cannot be rejected.[4]

Second, we can use the panel aspect of the data to bypass asset return measurement problems.[5] To this effect, we identify a common time component in expected consumption growth across states that is associated with interest rate movements when the IES is positive. This allows us to evaluate whether changes in expected asset returns affect consumption growth, without using any particular measure of asset returns. The approach has the additional advantage that it provides us with a means of examining whether the documented rejections of the over-identifying restrictions in the Euler equation are a result of asset return measurement problems.

The main results of the paper can be summarized as follows. First, in contrast to the recent results from aggregate data, using state-level consumption data we strongly reject a zero IES. When we include rule-of-thumb consumers, we find the IES to be significantly different from 0 with point estimates close to 1. Second, we find that there is a significant time-varying common component in expected consumption growth across states, after controlling for expected income growth. Moreover, the common time component is quite similar to the expected real return on Treasury bills. These results confirm that consumption growth is affected by interest rates (the IES is positive). Finally, we find that the presence of rule-of-thumb consumers is a simple, but robust, interpretation of the failure of the permanent income hypothesis, which cannot be explained away by interest rate mismeasurement.

The paper is structured as follows. In Section I we review the Euler equation approach to estimating the IES as well as Campbell and Mankiw's modification which introduces rule-of-thumb consumers. In Section II we use aggregate US data to document the failures of the standard Euler equation approach and to indicate how the introduction of rule-of-thumb consumers leads to very imprecise estimates of the IES. Section III describes our state-level data-set, which is subsequently used to examine the relationship between consumption growth and interest rates through a set of panel regressions. Section IV gathers concluding comments.

I. EMPIRICAL SPECIFICATIONS FOR CONSUMPTION BEHAVIOUR

Much of the literature that examines the IES is based on estimating the first-order conditions associated with the following maximization.

$$\max_{\{C_t\}_{t=0}^{\infty}} \quad E_0\left[\sum_{t=0}^{\infty}(1+\beta)^{-t}\frac{C_t^{1-\sigma}}{1-\sigma}\right]$$

$$\text{s.t.} \qquad A_t+C_t=Y_t+A_{t-1}(1+r_t).$$

In this maximization C_t represents consumption, Y_t stands for labour income, A_t represents assets and r_t denotes the real return of an asset held from period $t-1$ to period t. Under the assumption that asset returns have a log-normal distribution, the Euler equation derived from this maximization leads to the characterization of consumption growth given by equation (1). In (1), g_t^c is the growth rate in consumption from $t-1$ to t, E_{t-1} is the conditional expectation operator based on information available at time $t-1$, ε_t is an expectational error that is uncorrelated with $t-1$ information, and k_1 represents a covariance term which is assumed to be constant over the sample period:

$$(1) \qquad g_t^c=k_1+\frac{1}{\sigma}E_{t-1}(r_t)+\varepsilon_t.$$

A consistent estimate of $1/\sigma$, which is the IES, can be obtained by replacing $E_{t-1}(r_t)$ by the *ex post* real return on an asset and then estimating the equation by instrumental variables. Under rational expectations, any variables dated $t-1$ or earlier are admissible as instruments. However, as emphasized by Hall (1988), instruments dated $t-1$ may be problematic since the variables used for analysis are all time aggregates. Therefore, in all that follows, we use only instruments that are dated $t-2$ or earlier.

If, instead of assuming that all consumers in the economy choose consumption in order to satisfy equation (1), we follow Campbell and Mankiw and assume that a fraction of consumption is done by 'rule-of-thumb' consumers who simply consume their present income, then the appropriate specification for consumption growth is given by equation (2). In (2), λ represents the fraction of consumption that is done by rule-of-thumb consumers, g_t^y represents the growth rate in income and $\theta=(1-\lambda)/\sigma$:

$$(2) \qquad g_t^c=k_1+\theta E_{t-1}(r_t)+\lambda E_{t-1}(g_t^y)+\varepsilon_t.$$

Consistent estimators of θ and λ in equation (2) can again be obtained by instrumental variables once the expected return on the asset and the expected growth rate of income are replaced by their realized values. As instruments, any variable dated $t-2$ or earlier should be admissible.

Equation (2) can alternatively be derived while maintaining the assumption that all consumers are intertemporal maximizers but dropping the assumption that changes in asset positions are without cost. For example, if we assume that consumer's intertemporal utility function is given by the equation below, where $V(\cdot)$ represents the utility cost of changing assets, then a log-linear approximation to the corresponding Euler equation results in an equation identical to (2). The only difference is the interpretation of λ. Under an adjustment cost formulation, λ reflects the cost of portfolio adjustment relative to the gains in consumption smoothing. Nevertheless, the relationship between the IES and the parameters in (2) remains IES$=\theta/(1-\lambda)$:

$$\sum_{t=0}^{\infty}(1+\beta)^{-t}\{U(C_t)-V[A_t-(1+r_t)A_{t-1}]\},$$

$$U''(\cdot)<0 \quad \text{and} \quad V''(\cdot)>0.$$

Throughout the paper we will refer to (2) as the rule-of-thumb specification, even though there are obviously alternative interpretations, as the above example suggests.

II. Estimates of the Intertemporal Elasticity of Substitution using Aggregate US Data

In this section we estimate equations (1) and (2) using aggregate US data. The main objective is to document the difficulties associated with using economy-wide data to infer values for the IES. We also determine whether results are sensitive to sample period and consumption measures, since this will be important for assessing the applicability of our results where we use only state-level data on nondurable goods consumption over the period 1978–91.

The estimations are based on quarterly aggregate US data from 1953(I) to 1991(I). Campbell and Mankiw (1987, 1989, 1990) also start their sample period in 1953(I). Some earlier studies, such as Hall (1978, 1988) and Flavin (1981), use a sample period starting in 1947. By starting the period in 1953 we avoid the Korean war, and also circumvent an extreme outlier in the last quarter of 1950.[6]

Details on data sources and construction of variables can be found in the Appendix. Here we will provide only a brief description. Four types of consumption variables are considered: deseasonalized expenditures on non-durables (*ND*), non-deseasonalized expenditures on non-durables (*ND-NS*), deseasonalized expenditures on non-durables plus services (*NDS*), and non-deseasonilized retail sales of non-durables. The latter is included since it corresponds to the data on consumption that are available at the state level. Moreover, since the state-level retail sales data are available only since 1978, we also provide estimates of the IES for the sub-period 1978(I)–1991(I). We use two measures of aggregate retail sales data: one for the country as a whole (*Retail*) and one that corresponds to the aggregate consumption of the 19 states in our panel (*Retail-19*).[7] Data for *ND* and *NDS* are available at constant prices, while the other series are divided by the quarterly nondurables consumer price index. All consumption variables are divided by the population. We use two types of nominal return variables: the three-month Treasury bill (T-bill) rate and the stock return on the Standard and Poor 500. The *ex post* real return on Treasury bills is calculated by taking the average three-month after-tax T-bill interest rate and subtracting the quarterly inflation rate, which depends on the choice of the consumption variable. The real stock return is the quarterly after-tax dividend return, plus capital gains, minus the quarterly inflation rate. For now we will assume a constant tax rate of 30% on interest and dividend. However, we will also report some results based on time-varying estimates of the effective marginal tax rate.[8]

In order to be consistent with the earlier work by Campbell and Mankiw (1989), we use the same measure for income in the estimation of equation (2) when *ND* or *NDS* is used as the consumption measure. The income variable is per capita disposable personal income at constant prices, a series published by Citibase. Since this measure is not available at the state level, in order to make our results comparable with those in the panel regressions of the next section, we use a different measure for real per capita income when either *Retail*, *Retail-19* or *ND-NS* is used as the consumption measure. In these cases, nominal personal income is divided by the overall CPI deflator and population.

TABLE 1

ESTIMATION RESULTS FOR EULER EQUATION: AGGREGATE US DATA

Sample (1)	Consumption (2)	Return (3)	Instruments (4)	Lags (5)	IES (s.e.) (6)	p-value OID test (7)
1953(I)–1991(I)	*ND*	T-bill	T-bill	3	0·375 (0·124)	0·001
	ND	T-bill	T-bill	4	0·370 (0·124)	0·002
	ND	T-bill	T-bill, g^c	4	0·334 (0·108)	0·002
	ND	Stock	Stock, g^c	3	0·035 (0·037)	0·115
	NDS	T-bill	T-bill	3	0·255 (0·118)	0·017
	NDS	T-bill	T-bill	4	0·245 (0·118)	0·021
	NDS	T-bill	T-bill, g^c	4	0·229 (0·111)	0·005
	NDS	Stock	Stock, g^c	3	−0·009 (0·027)	0·114
1978(I)–1991(I)	*ND*	T-bill	T-bill	3	0·486 (0·158)	0·026
	ND	T-bill	T-bill, g^c	3	0·358 (0·142)	0·001
	ND	Stock	Stock, g^c	3	0·091 (0·068)	0·126
	ND-NS	T-bill	T-bill	3	0·494 (0·382)	0·034
	ND-NS	T-bill	T-bill, g^c	3	0·348 (0·364)	0·019
	ND-NS	Stock	Stock, g^c	3	0·026 (0·101)	0·059
	Retail	T-bill	T-bill	3	0·423 (0·222)	0·034
	Retail	T-bill	T-bill, g^c	3	0·339 (0·201)	0·081
	Retail	Stock	Stock, g^c	3	0·069 (0·051)	0·300
	Retail-19	T-bill	T-bill	3	0·509 (0.248)	0·125
	Retail-19	T-bill	T-bill, g^c	3	0·399 (0·223)	0·119
	Retail-19	Stock	Stock, g^c	3	0·104 (0·059)	0·459
	NDS	T-bill	T-bill	3	0·379 (0·164)	0·056
	NDS	T-bill	T-bill, g^c	3	0·344 (0·160)	0·011
	NDS	Stock	Stock, g^c	3	0·037 (0·040)	0·022

Notes: The table shows the results from estimating the Euler equation $g_t^c = S_j + \text{IES } E(r_t)$, using aggregate US data. Here S_j is a season dummy when the consumption measure is not de-seasonalized, and a constant otherwise. Column (2) shows the type of consumption measure: *ND* = (de-seasonalized) nondurables, *ND-NS* = non-de-seasonalized nondurables, *NDS* = nondurables plus services, *Retail* = aggregate US retail sales (non-de-seasonalized), *Retail-19* = aggregate retail sales over 19 US states. Column (3) shows the type of asset return measure used. Columns (4) and (5) show the type of instruments and the number of lags used (e.g. 'lag = 3' means that instruments of 2, 3 and 4 quarters lagged are used). Column (6) reports the estimated intertemporal elasticity of substitution and its standard error. The final column provides the p-value of the $\chi^2(k-1)$ over-identification (OID) test, where k is the number of instruments.

Table 1 presents estimates derived from estimating equation (1). The column 'Lags' indicates how many lags are used for each instrument. For example, three lags means that instruments at $t-2$, $t-3$, and $t-4$ are used. The first point to note is that there is a significant difference between the results based on stock returns and those on Treasury bills. Like Hall (1988), we find that the results based on stock returns always show estimates of the IES that are very close to zero, with small standard errors. However, results based on stock returns are potentially questionable since stock returns are hard to predict. In particular, Nelson and Startz (1990) have shown that when instruments are poor, the small sample properties of IV estimates are unreliable. Therefore, we consider the results based on T-bill returns to be more dependable.

The results in Table 1 obtained when using the T-bill rate as the asset variable almost uniformly suggest an IES that is greater than 0. The point estimates for the long sample are mostly in a range between 0·25 and 0·4, with standard errors near 0·12. Using nondurables consumption and a T-bill return, for the sample period 1947–83 Hall (1988) found a small and insignificant IES. The difference can be attributed to the outlier in the last quarter of 1950 that was mentioned above. Remarkably, we find that, while the results are practically the same as those in Table 1 if we start the sample in 1951(I), starting the sample in 1950(IV) or earlier leads to an IES that is insignificant. When using the Hall sample (1947–83), the point estimate for the IES drops to 0·034 with a standard error of 0·130 when *ND* is used as our measure of consumption and four lags of the T-bill rate and consumption growth are used as instruments. A similar pattern emerges using other measures of consumption and other choices of instrument.[9]

The last column in Table 1 shows the *p*-value associated with the $\chi^2(k-1)$ Lagrange Multiplier test of over-identifying (OID) restrictions, where k is the number of instruments.[10] When the T-bill rate is used as the measure of asset return, the overidentifying restrictions are rejected for the 1953(I)–1991(I) sample. The test is obviously weaker for the shorter sample for 1978(I)–1991(I), although even then we still find rejections of the OID restrictions in most cases.

Table 2 presents results corresponding to the estimation of equation (2) where rule-of-thumb consumers are introduced. Here the return variable is always based on the T-bill rate. The estimate of λ, which represents the fraction of rule-of-thumb consumers, is always positive and significant.[11] This suggests that the standard Euler equation without rule-of-thumb consumers is based on a mis-specified model. The last column shows the *p*-value associated with a $\chi^2(k-2)$ Lagrange Multiplier test of OID restrictions. This provides some evidence in favour of the permanent-income-cum-rule-of-thumb specification. We are never able to reject the OID restrictions at the 5% level, even for the long sample 1953(I)–1991(I).

It is especially important to note that in Table 2 the estimates of the IES are insignificantly different from 0 in both sample periods. An IES of 0 cannot be rejected if one is ready to accept the possibility of rule-of-thumb consumers, which is the conclusion reached by Campbell and Mankiw (1989).[12] However, the opposing view of relatively high values for the IES cannot be rejected either. For example, based on the entire sample period and based on the measure of consumption that includes only nondurable goods, the IES could take on any value between 0 and 1·5. In the case where services are included in consumption, our estimates of the IES suggest values between 0 and 0·6.

Table 3 reports some results when the marginal tax rate on interest income is allowed to vary over time instead of being fixed at 30%. Annual estimates of the effective marginal tax rate on interest income are taken from Gravelle (1994). Tax rates are assumed to be constant within a year. The results in Table 3 are very similar to those found in Tables 1 and 2, even though tax rates vary from 23% to 36% in both samples.

TABLE 2
ESTIMATION RESULTS FOR CAMPBELL–MANKIW SPECIFICATION: AGGREGATE US DATA

Consumption (1)	Instruments (2)	Lags (3)	θ (s.e.) (4)	λ (s.e.) (5)	IES (s.e.) (6)	p-value OID test (7)
(a) 1953(I)–1991(I)						
ND	T-bill, g^c	3	0·211 (0·122)	0·763 (0·252)	0·889 (0·959)	0·074
ND	T-bill, g^c	4	0·222 (0·112)	0·717 (0·203)	0·785 (0·590)	0·162
ND	T-bill, g^c, g^y	3	0·215 (0·108)	0·596 (0·196)	0·522 (0·319)	0·083
NDS	T-bill, g^c	3	0·122 (0·107)	0·463 (0·139)	0·228 (0·197)	0·206
NDS	T-bill, g^c	4	0·134 (0·105)	0·469 (0·112)	0·251 (0·192)	0·417
NDS	T-bill, g^c, g^y	3	0·116 (0·102)	0·388 (0·111)	0·190 (0·162)	0·363
(b) 1978(I)–1991(I)						
ND	T-bill, g^c	4	0·231 (0·164)	0·848 (0·274)	1·522 (2·696)	0·786
ND	T-bill, g^c	3	0·311 (0·204)	1·008 (0·353)	−37 (−1569)	0·925
ND	T-bill, g^c, g^y	3	0·295 (0·184)	0·885 (0·203)	2·598 (6·843)	0·874
ND-NS	T-bill, g^c	3	−0·168 (0·429)	1·272 (0·459)	0·615 (−1·462)	0·510
ND-NS	T-bill, g^c	4	−0·417 (0·464)	1·387 (0·660)	1·076 (−1·487)	0·867
ND-NS	T-bill, g^c, g^y	3	−0·141 (0·415)	1·262 (0·448)	0·537 (−1·460)	0·840
Retail	T-bill, g^c	3	0·215 (0·196)	0·636 (0·242)	0·589 (0·584)	0·525
Retail	T-bill, g^c	4	0·202 (0·173)	0·631 (0·234)	0·548 (0·535)	0·746
Retail	T-bill, g^c, g^y	3	0·219 (0·190)	0·592 (0·236)	0·536 (0·236)	0·328
Retail-19	T-bill, g^c	3	0·271 (0·219)	0·659 (0·271)	0·793 (0·784)	0·552
Retail-19	T-bill, g^c	4	0·210 (0·194)	0·672 (0·262)	0·641 (0·713)	0·666
Retail-19	T-bill, g^c, g^y	3	0·280 (0·213)	0·601 (0·265)	0·701 (0·607)	0·286
NDS	T-bill, g^c	3	0·288 (0·182)	0·524 (0·180)	0·604 (0·423)	0·595
NDS	T-bill, g^c	4	0·223 (0·153)	0·392 (0·125)	0·367 (0·248)	0·234
NDS	T-bill, g^c, g^y	3	0·274 (0·166)	0·439 (0·155)	0·488 (0·309)	0·551

Notes: The table shows the results from estimating the Campbell–Mankiw specification $g_t^c = S_j + \theta E(r_t) + \lambda E(g_t^y)$, using aggregate US data. Here S_j is a season dummy when the consumption measure is not de-seasonalized, and a constant otherwise. The T-bill rate is used as the asset return measure. Column (1) shows the type of consumption measure: *ND* = (de-seasonalized) nondurables, *ND-NS* = non-de-seasonalized nondurables, *NDS* = nondurables plus services, *Retail* = aggregate US retail sales (non-de-seasonalized), *Retail-19* = aggregate retail sales over 19 US states. Columns (2) and (3) show the type of instruments and the number of lags used. The next three columns report parameter estimates, with standard errors in brackets. The final column provides the p-value of the $\chi^2(k-2)$ over-identification test, where k is the number of instruments.

In Tables 1, 2 and 3 we find that the estimates of the IES are always higher when based on nondurables consumption as opposed to the consumption of nondurables plus services. This may be explained by the fact that services include medical expenditures and a flow measure for housing services. For one, most medical services are insured, and insurance companies provide few incentives for consumers to substitute intertemporally. Therefore estimating a low IES for services may reflect more the difficulty of providing incentives within medical insurance policies than the actual willingness of consumers to substitute. Second, the *imputed* cost of housing is the largest item in services, and since its measurment does not include expected capital gains or losses, it may also be contributing to the lower estimates of the IES. In effect, since the quantity of housing services is almost fixed in the short run, any desire to substitute intertemporally will have to be reflected in the cost of housing. Consequently, any mis-measurement of the actual cost of housing could greatly impair any evidence of short-run intertemporal substitution, and therefore

TABLE 3

ESTIMATION RESULTS FOR AGGREGATE US DATA WITH A TIME-VARIABLE TAX RATE

Consumption	Instruments	Lags	IES (s.e.)	p-value OID test
Table 1 results with time-variable tax rate				
(a) 1953(I)–1991(I)				
ND	T-bill, g_c	4	0·340 (0·109)	0·003
NDS	T-bill, g_c	4	0·237 (0·108)	0·005
(b) 1978(I)–1991(I)				
ND	T-bill, g_c	3	0·361 (0·143)	0·002
NDS	T-bill, g_c	3	0·339 (0·152)	0·016
Retail-19	T-bill, g_c	3	0·409 (0·233)	0·141

Table 2 results with time-variable tax rate
(a) 1953(I)–1991(I)

Consumption	Instruments	Lags	θ (s.e.)	λ (s.e.)	IES (s.e.)	p-value OID test
ND	T-bill, g^c	3	0·206 (0·123)	0·749 (0·251)	0·820 (0·835)	0·067
NDS	T-bill, g^c	3	0·123 (0·105)	0·453 (0·140)	0·225 (0·187)	0·192
(b) 1978(I)–1991(I)						
ND	T-bill, g^c	3	0·298 (0·204)	0·991 (0·349)	34·33 (1375)	0·923
NDS	T-bill, g^c	3	0·274 (0·170)	0·499 (0·176)	0·546 (0·367)	0·559
Retail-19	T-bill, g^c	3	0·268 (0·230)	0·639 (0·275)	0·741 (0·730)	0·539

Notes: This table shows some of the results of Tables 1 and 2 when the after-tax return on Treasury bills is measured using a time-variable effective tax rate obtained from Gravelle (1994). The symbols are the same as in the previous two tables.

implies that estimating the IES using these measures for services may lead to unreliable estimates.

III. ESTIMATING THE IES WITH STATE-LEVEL DATA

The results of the previous section indicate that estimates of the IES that are based on the standard Euler equation specification for consumption (equation (1)) are generally significantly different from 0 with point estimates mostly in a range between 0·25 and 0·4. However, the fact that expected income growth significantly helps predict consumption growth in equation (2) signals an important problem with the standard specification of consumption Euler equations. If we accept the introduction of rule-of-thumb consumers as a reasonable modification to the theory, we find that aggregate US data provide very little information regarding the value of the IES. Therefore, based on aggregate US data, we believe that it is problematic to infer much regarding the value of the IES.

In this section we use state-level consumption data in view of improving our knowledge of the IES. We begin by showing that, conditional on accepting the structural interpretation of equation (2), the data strongly suggest that the IES (for nondurable goods) is greater than 0·5 and most probably close to 1. In order to assess the relevance of this result, we further exploit the panel aspect of the data to examine whether the permanent-income-cum-rule-of-thumb interpretation of the data is robust. In particular, we examine whether the presence of rule-of-thumb consumers may be an artefact caused by the mis-measurement of interest rates and whether the T-bill rate is in fact an appropriate measure of asset returns.

The consumption data we use are state-level retail sales data for nondurable goods. These are unpublished data available from the Bureau of the Census. Our data-set includes non-adjusted quarterly data from 19 states over the period 1978(I)–1991(I). The states are: California, Florida, Illinois, Indiana, Louisiana, Maryland, Massachusetts, Michigan, Minnesota, Missouri, New Jersey, New York, North Carolina, Ohio, Pennsylvania, Tennessee, Texas, Virginia and Wisconsin. For all other states, either these data are not available or they are available only since 1987(I). Consumption is deflated using regional consumer price indices for nondurable goods and is divided by state population.[13] For each of these states, we use personal income by state of residence as our measure of income. These data are obtained from the Bureau of Economic Analysis. State income is deflated by the consumer price index for all goods, and divided by state population. The *ex post* state asset returns are the average three-month US T-bill rate minus the holding period regional inflation rate. The Appendix provides further details on the data used.

TABLE 4

ESTIMATION CAMPBELL–MANKIW SPECIFICATION: 19 STATES PANEL DATA, 1978(I)–1991(I)

Instruments (1)	Lags (2)	θ (s.e.) (3)	λ (s.e.) (4)	IES (s.e.) (5)	p-value OID test (6)
T-bill, g^c	4	0·392 (0·115)	0·629 (0·095)	1·057 (0·370)	1·000
T-bill, g^c	3	0·469 (0·127)	0·599 (0·098)	1·171 (0·376)	1·000
T-bill, g^c	2	0·642 (0·143)	0·338 (0·130)	0·970 (0·242)	0·998
T-bill, g^c, g^y	3	0·476 (0·123)	0·597 (0·093)	1·182 (0·362)	0·974
T-bill, g^c, g^{gnp}	3	0·352 (0·111)	0·558 (0·094)	0·797 (0·263)	0·985
T-bill, g^c, g^{inv}	3	0·311 (0·108)	0·519 (0·092)	0·646 (0·225)	0·996
T-bill, g^c, g^{gov}	3	0·513 (0·125)	0·496 (0·092)	1·018 (0·280)	1·000
T-bill, g^y	3	0·381 (0·131)	0·675 (0·113)	1·171 (0·488)	0·328
T-bill, g^{gnp}	3	0·336 (0·117)	0·573 (0·111)	0·786 (0·284)	0·113
T-bill, g^{inv}	3	0·304 (0·120)	0·540 (0·114)	0·662 (0·252)	0·230
T-bill, g^{gov}	3	0·436 (0·137)	0·583 (0·112)	1·044 (0·365)	0·801
T-bill, g^c	3	0·465 (0·131)	0·600 (0·099)	1·162 (0·382)	1·000

Notes: The table reports results from estimating the Campbell–Mankiw specification based on a 19-state panel data-set. The estimated equation is $g^c_{it} = S_{i,j} + \theta E(r_t - \pi_{it}) + \lambda E(g^y_{it})$. Here $S_{i,j}$ is a state-specific season dummy. The state consumption measure is non-de-seasonalized nondurables retail sales. State income is personal income, deflated by the consumer price index. The asset return variable is the T-bill rates minus regional inflation (using a nondurables consumption price index). The instruments in column (1) are all aggregate variables. The final column provides the p-value of a $\chi^2(19 \times k - 2)$ over-identification test, where k is the number of instruments. The final row reports the results when a time-variable tax rate is used to compute the after-tax return on T-bills.

Table 4 presents results from estimating equation (2) using the state-level data. The instruments are a full set of state dummies interacted with different combinations of the real return on Treasury bills, and the growth of aggregate US personal income, aggregate US consumption of nondurable goods (retail sales data), aggregate US private investment and total government spending on goods and services. Note that, by allowing the instrument set to be comprised of aggregate variables interacted with state-level dummies, we are in fact allowing predicted income growth to differ between states. The variance–covariance structure of the error term allows for heteroscedasticity across states and for arbitrary contemporaneous covariances between states. The serial correlation in the error term was found to be insignificant and therefore was set to zero. Since the state-level consumption data are not de-seasonalized, we estimate equation (2) with seasonal dummies that are allowed to vary between states.

The results in Table 4 are all quite similar. The estimates of the IES cluster around 1 with a standard error of approximately 0·3, clearly rejected an IES of 0. Furthermore, we cannot reject the over-identifying restrictions.[14] Compared with the results using aggregate data, the most important change is in terms of precision. The improved precision comes from two sources. First, the cross-state variation in expected income growth greatly improves the precision in λ, which directly translates into more precision for the IES. Second, the estimation efficiency is improved by allowing for an arbitrary contemporaneous variance–covariance structure of the residuals across states. Potentially, the cross-state variation in real returns could also be contributing to the improved precision arising from differences in expected inflation. However, this variation is rather insignificant in our data.[15]

The results from Table 4 depend on our particular measure of asset returns. In principle, both the Euler equation and the Campbell–Mankiw specification should hold for all assets. However, asset return measurement problems may enter in various ways. First, not all agents (other than rule-of-thumb consumers) necessarily have access to all asset markets. Therefore a different marginal rate of return may be relevant for different consumers. Such consumer heterogeneity may warrant the use of an average rate of return on different assets instead of the rate on only one asset. Second, because of transaction costs, it is unclear how to determine the relevant marginal return for the representative consumer. Finally, without an exact knowledge of the representative consumer's decision periods, uncertainty about the appropriate time aggregation of the interest rates will be present.

Note that each of these problems can lead to a measurement error that is not uncorrelated over time and therefore is not necessarily cured by the use of instrumental variables. It is quite possible that the apparent superiority of the rule-of-thumb specification is a result of an imprecise measure of the appropriate asset return. Since our estimates of the IES depend on the validity of the Campbell–Mankiw specification, it is worthwhile investigating this possibility. One of the advantages of panel data is that it is feasible to examine whether the perceived superiority of the rule-of-thumb specification and our finding of a significantly positive IES, are a reflection of asset return measurement problems.

Examining biases arising from inappropriate measures of interest rates

In order to understand the biases that can be caused by an improper choice of interest rate, let us consider the state-level counterpart to equation (1), given by (3). Since the retail sales data are not de-seasonalized, (3) explicitly includes seasonal effects. In this equation, subscript i denotes the state, the variables $S_{i,j}$ represent state-level seasonal effects ($j = 1, \ldots, 4$) and Ω_{t-1} represents the aggregate information set at time $t-1$; the term $\pi_{i,t} - \pi_t$ is the differential between state-level inflation and aggregate inflation and reflects the possibility of different real interest rates between states.[16] Even though information sets may differ among states, it is always possible to express all expectations with respect to aggregate variables as long as the aggregate information is assumed to be available in each state:

$$g^c_{i,t} = S_{i,j} + \frac{1}{\sigma} E(r_t/\Omega_{t-1}) - \frac{1}{\sigma} E(\pi_{i,t} - \pi_t/\Omega_{t-1}) + \varepsilon_{i,t}. \tag{3}$$

If we assume that we know how to measure the relevant *ex post* real interest rate r_t, then we can estimate (3) by instrumental variables. This has been our maintained hypothesis to date. However, if our choice of interest rate is inappropriate, this could lead to rejections of the over-identifying restrictions and a bias in favour of finding rule-of-thumb consumers. To see this, let the conditional expectation of the appropriate real interest rate be given by equation (4) and that of an inappropriate interest rate, denoted $\tilde{r}_t$, be given by (5):

$$E(r_t/\Omega_{t-1}) = \Omega'_{t-1}\gamma_1, \tag{4}$$

$$E(\tilde{r}_t/\Omega_{t-1}) = \Omega'_{t-1}\gamma_2. \tag{5}$$

Using (4) and (5), we can rewrite (3) as in (6):

$$g^c_{i,t} = S_{i,j} + \frac{1}{\sigma} E_{t-1}(\tilde{r}_t) - \frac{1}{\sigma} E_{t-1}(\pi_{i,t} - \pi_t) + \varepsilon_{i,t} + \frac{1}{\sigma} \Omega'_{t-1}(\gamma_1 - \gamma_2). \tag{6}$$

If $\gamma_1 \neq \gamma_2$, then the error term in equation (6) includes all variables in the information set. Consequently, an inappropriate choice of asset return could be the case of the rejection of the over-identifying restrictions and could be the reason for the better fit of the rule-of-thumb specification.

In order to examine this possibility, it is useful to estimate equation (6) by replacing $E(r_t/\Omega_{t-1})$ by its unrestricted form given in (5) and including expected income growth as an additional regressor. On the one hand, if the failure of the standard Euler equations for consumption is due to a mismeasurement of the interest rate, then allowing for expected state income growth in this less restrictive specification should lead to an insignificant coefficient on income. On the other hand, even if λ is significant and the rule-of-thumb specification is appropriate, an inappropriate choice of interest rates could lead to a bias in the estimates of the fraction of the population that are rule-of-thumb consumers. Consequently, replacing the expected real interest rate by its unrestricted reduced form should allow us to determine the biases in the estimates of λ that can be due to the difficulty in measuring interest rates.

Table 5 presents the results associated with the estimation of equation (7) for different aggregate information sets Ω_{t-2}. Equation (7) modifies (3) in the following manner: (a) the expected interest rate is replaced by its unrestricted reduced form to avoid the need of identifying the proper measure for asset returns; (b) expected income growth is included to test for the presence of rule-of-thumb consumers; (c) the information set is restricted to include only information as of time $t-2$ in order to avoid problems caused by time aggregation; and (d) the expected inflation differential is set to zero. The last modification is adopted only because we could not reject the fact that the expected inflation differentials between states are 0 based on aggregate information.

TABLE 5

IDENTIFICATION OF A COMMON TIME EFFECT IN CONSUMPTION GROWTH FOR 19 STATES, 1978(I)–1991(I)

Instruments (1)	Lags (2)	λ (s.e.) (3)	p-value (time effect) (4)	p-value (T-bill = time effect) (5)	p-value (OID test) (6)
T-bill, g^c	4	0·449 (0·114)	0·004	0·069	1·000
T-bill, g^c	3	0·417 (0·118)	0·001	0·096	1·000
T-bill, g^c	2	0·267 (0·155)	0·000	0·020	0·994
T-bill, g^c, g^y	3	0·486 (0·111)	0·000	0·001	0·951
T-bill, g^c, g^{gnp}	3	0·334 (0·113)	0·002	0·016	0·959
T-bill, g^c, g^{inv}	3	0·335 (0·111)	0·002	0·007	0·984
T-bill, g^c, g^{gov}	3	0·369 (0·111)	0·000	0·043	1·000
T-bill, g^y	3	0·596 (0·132)	0·000	0·000	0·242
T-bill, g^{gnp}	3	0·475 (0·127)	0·010	0·068	0·070
T-bill, g^{inv}	3	0·521 (0·128)	0·006	0·013	0·190
T-bill, g^{gov}	3	0·555 (0·125)	0·001	0·068	0·762

Notes: The table shows results from estimating the equation $g^c_{it} = S_{i,j} + [(1-\lambda)/\sigma]\gamma'_1\Omega_{t-2} + \lambda E(g^y_{it}/\Omega_{t-2})$. Here Ω_{t-2} contains the aggregate variables that are used as instruments, and $S_{i,j}$ is a state-specific season dummy. Column (4) reports the p-value associated with the $\chi^2(k)$ test that there is no common time component in expected consumption growth ($[(1-\lambda)/\sigma]\gamma_1 = 0$). Here k is the number of instruments. Column (5) provides the p-value of the $\chi^2(k-1)$ test that the common time component is proportional to the expected T-bill rate. The final column reports the p-value of a $\chi^2(19 \times k - (k+1))$ test of over-identifying restrictions.

Nevertheless, none of our results are sensitive to the inclusion of this differential and, correspondingly, when it is included its estimation coefficient is always close to 0 and very imprecisely estimated.[17] Our estimation of (7) again allows for cross-state correlations in the error terms and for cross-state heteroscedasticity:[18]

$$(7) \qquad g^c_{i,t} = S_{i,j} + \frac{1-\lambda}{\sigma}\gamma'_1\Omega_{t-2} + \lambda E(g^y_{i,t}/\Omega_{t-2}) + \varepsilon_{i,t}.$$

It is important to note that the estimation of equation (7) requires panel data. With only time-series data, it is impossible to identify simultaneously an unrestricted time effect and the effect of expected income growth resulting from collinearity. With state-level data, the cross-sectional variation in expected income growth allows us to identify these two components, since for each state expected income growth is a different linear combination of the variables in the information set. In order words, there is state-specific time variation in predicted income growth owing to the difference in the predictability of state income. Note that this identification is possible because the instruments are aggregate variables interacted with state-level dummies.[19]

The third column in Table 5 presents estimates of λ, while the fourth column reports the p-value associated with the chi-square statistic for the test that the vector $[(1-\lambda)/\sigma]\gamma_1 = 0$. If the interest rate is predictable ($\gamma \neq 0$) and not all consumers are of the rule-of-thumb type ($\lambda < 1$), then this test corresponds to the null hypothesis that IES = 0. The fifth column is a chi-square statistic associated with the hypothesis that, conditional on expected income growth, aggregate variables affect expected consumption growth only in proportion to the linear combination that predicts the T-bill return; that is, it is a test of whether the T-bill return is an appropriate measure of the interest rate faced by optimizing consumers. The final column shows the p-value of the chi-square test of over-identifying restrictions.[20]

The p-values reported in column (4) indicate that there is a significant common time effect, $[(1-\lambda)/\sigma]\gamma_1'\Omega_{t-2}$, driving expected state-level consumption growth, even after controlling for expected income growth. This estimated time effect should reflect the common cross-state substitution pattern that is induced by changes in the intertemporal price of consumption. The results therefore provide strong evidence of a positive IES.

Table 5 also shows that mis-measurement of interest rates cannot account for the rejection of the standard Euler equation specification. In fact, the estimates of λ in all cases are quite significant, and the over-identifying restrictions are again not rejected. The average of the point estimates for λ is 0·44. This is only 0·11 below the average of the estimates reported in Table 4, when the real return on T-bills was assumed to be an appropriate measure of the interest rate. It may suggest a slight mismeasurement of the interest rate, although the difference is not statistically significant. Together, these results confirm the difficulties associates with the pure permanent income hypothesis and suggest that the rule-of-thumb specification may possibly be a better description of aggregate consumer behaviour.

Since our point estimates for the IES were based on the choice of the expected T-bill rate as the asset return measure, it is of interest to know whether the common time effect is fully captured by movements in the expected T-bill return. We therefore test whether equation (7) admits the restrictions imposed by setting $[(1-\lambda)/\sigma]\gamma_1'\Omega_{t-2}$ proportional to $E(r_t/\Omega_{t-2})$, where r_t is the real T-bill return. The p-values of this test, reported in column (5) of Table 5, show that at the 1% significant level we reject this null hypothesis in only 3 of the 11 cases; at the 5% level, we reject in 7 cases that the expected T-bill rate is the appropriate measure of asset return.

In order to provide a visual background to these findings, Figures 1 and 2 plot both the yearly average of this estimated common time effect ($[(1-\lambda)/\sigma]\gamma_1'\Omega_{t-2}$) and the yearly average of the expected real return on three-month Treasury bills. The unobserved expected return variable should be a linear function of the estimated common time effect and therefore this time effect has been scaled to be comparable with the expected real return on Treasury bills. The difference between the two figures is the set of variables used as instruments. Figure 1 corresponds to the specification in row 4 of Table 5, where the instruments are the real return on Treasury bills, US consumption growth and US income growth. This figure has been chosen as a representation of one of the worst fits within the table (the p-value in column (5) is 0·001). Figure 2 corresponds to the specification in row 2, where the instruments are the real return on Treasury bills and US consumption growth. This figure has been chosen as one of the best fits (the p-value is 0·096). Even though our formal test corresponding to the case of Figure 1 rejects the fact that the T-bill rate is the correct asset return measure, we find that the estimated common time effect does in fact move quite closely with the expected return on Treasury bills. This is suggestive of a common response across states to changes in expected asset returns. The similarity of the expected T-bill rate and common time component is even more striking in Figure 2, where the expected T-bill return explains over 90% of the variance in the common time effect.

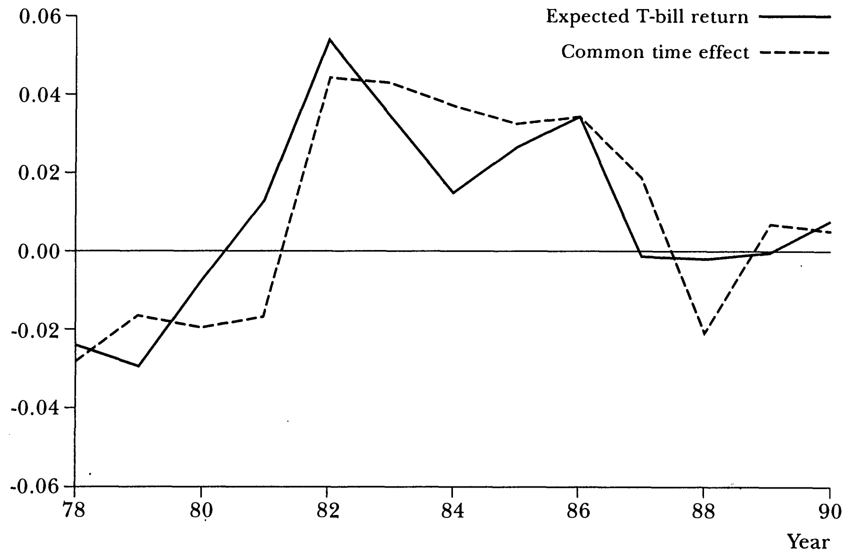

FIGURE 1

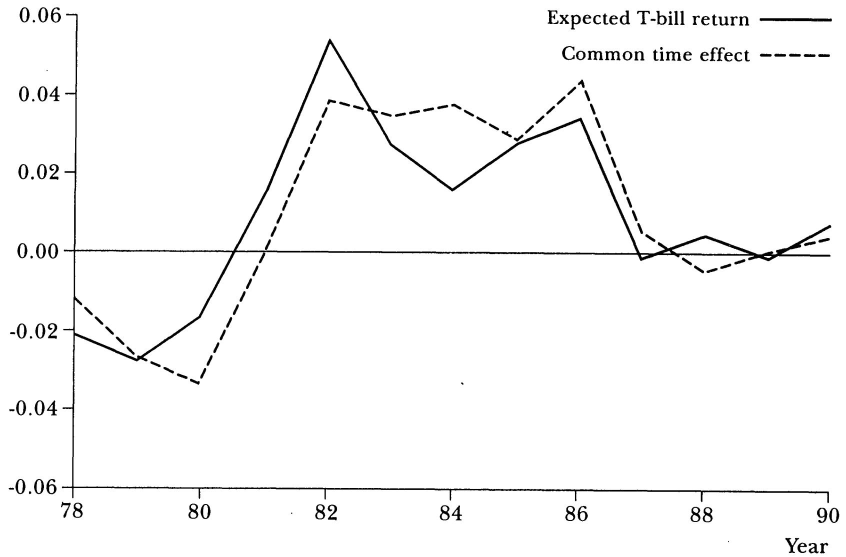

FIGURE 2

The most telling aspects of Figures 1 and 2 come from the observations in the late 1970s and early 1980s. In the late 1970s, income growth was positive but the common time effect in expected consumption growth is observed below mean potentially because of a low interest rate. In contrast, in 1982 income growth was weak, but high expected interest rates seem to drive expected consumption growth high above its mean.

The similarity between the expected T-bill rate and the common time effect provides further support for our point estimates of the IES based on the expected T-bill rate. But at a minimum, the observation that there is a significant common time effect in state-level expected consumption growth is substantial evidence in favour of the hypothesis that the IES is different from 0.

IV. Conclusion

The aim of this paper is to provide insight with regard to the value of the intertemporal elasticity of substitution of consumption. As we have shown, aggregate data for the entire US economy are rather uninformative on this subject. In contrast, using a panel of state-level data, we find that the IES for nondurables consumption is significantly different from 0, and probably close to 1. In further support of a positive relationship between interest rates and expected consumption growth, we identify a significant common time component in expected consumption growth across states. Moreover, we document a close relationship between this common time component and the expected T-bill return. Together, these results provide considerable evidence of the willingness of consumers to substitute intertemporally in response to interest rate movements.

DATA APPENDIX

Data sources

1. Citibase: Retrieval codes are provided in brackets. From this source we obtain quarterly de-seasonalized personal consumption of nondurable goods at constant prices (*GCN*), quarterly de-seasonalized personal consumption of services at constant prices (*GCS*), disposable income per capita at 1987 prices (*GYDPCQ*), three-month T-bill rate (*FYGN*3), Standard and Poors' composite stock price index (*FSPCOM*), Standard and Poors' dividend yield on common stock composite (*FSDXP*), quarterly disposable personal income (*GYD*), quarterly GNP at 1987 prices (*GNPQ*), quarterly gross private domestic investment at 1987 prices (*GPIQ*), quarterly government consumption (*GGEQ*), quarterly nondurables consumption price deflator (*GDCN*), quarterly services price deflator (*GDCS*), monthly CPI-U index for all items (*PZUNEW*), and monthly population (*POP*).

2. Survey of Current Business, Department of Commerce: From this source we obtain nominal quarterly non-seasonally adjusted personal consumption of nondurables. Below it is referred to as *NDNSNOM*.

3. Monthly Retail Survey, Bureau of the Census: From this source we obtain nominal monthly non-seasonally adjusted retail sales of nondurables, both the aggregate and the data for the 19 states mentioned in the text. These data are unpublished. We refer to this series as *RET* below.

4. Bureau of Economic Analysis (quarterly personal income diskette): From this source we obtain non-deseasonalized quarterly nominal personal state income data, referred to as *NSY* below. We also obtain annual state population data from this source, referred to as *N* below.

5. CPI Detailed Report, Bureau of Labor Statistics: From this source we obtain a quarterly regional price index for nondurable consumption goods. The data are available for four large regions. Refer to this as *PND*.

Computation of Variables used in the Regressions

To the extent that monthly data are used, they are first aggregated to the quarterly frequency. The annual state population data are converted to the quarterly frequency through interpolation.

1. Left-hand-side consumption variable: $ND = GCN/POP$; $ND-NS = NDNSNOM/(N * GDCN)$; $NDS = (GCN + GCS)/POP$; $Retail = RET/(N * GDCN)$; $Retail_i$ (panel regressions) $= RET_i/(N_i * PND_i)$.

2. *Right-hand-side variables:* (i) *Income*: When the left-hand-side variable is *not* retail sales, income is equal to *GYDPCQ*; otherwise, income is $NSY/(POP * PZUNEW)$ for aggregate data, and $NSY_i/(N_i * PZUNEW)$ for panel data.

(ii) *Asset returns.* After-tax T-bill return $= (1-\tau)FYGN3 - \pi$. The tax rate τ is either set at 0·3, or is the time variable rate obtained from Gravelle (1994). Since the data from Gravelle are annual, the tax rate is assumed to be the same in all quarters within a year. The inflation rate $\pi = \pi_{GDCN}$ (inflation rate in the *GDCN* index) when *ND*, $ND - NS$, or aggregate retail sales is the left-hand-side consumption measure; $\pi = 0{\cdot}57\pi_{GDCS} + 0{\cdot}43\pi_{GDCN}$ when *NDS* is the consumption measure; $\pi_i = \pi_{PND_i}$, in the panel regressions. The after-tax stock return is computed as follows. Monthly dividends are obtained by multiplying the yield series *FSDXP* with the stock price index *FSPCOM*. The stock return from quarter t to $t+1$ is then computed as the average after-tax dividend in quarter $t+1$, divided by the average stock price index in t, plus the relative change in the average stock price index from quarter t to $t+1$.

3. *Instruments.* To the extent that (the growth rate of) consumption, asset returns or (the growth rate of) personal income are used as instruments, they are simply based on lags of the variables discussed above. When *NDS*, *ND*, *ND-NS*, or *Retail* are used as the left-hand-side consumption measure, their lagged growth rates are used when g_c is an instrument. When *Retail-19* or panel data are used, lagged growth rates of *Retail* are used when g_c is an instrument. In the panel regressions, additional instruments are (the growth rate of) per capita real GNP $= GNPQ/POP$, per capita real government consumption $= GCEQ/POP$, and per capita real private investment $= GPIQ/POP$.

ACKNOWLEDGMENTS

The authors would like to thank John Geweke, Sam Kortum, Larry Kotlikoff, Kevin Lang, Annemaria Lusardi, Jeff Miron and seminar participants at the Federal Reserve Bank of Minneapolis and IGIER, Milan, for useful comments and discussions.

NOTES

1. They find that about 45% of all consumption is done by such rule-of-thumb consumers, with a standard error of about 20%. This is a little higher, but consistent, with studies by Hall and Mishkin (1982), Mariger (1986) and Jappelli (1990), who, using different methods, all find that about 20% of US households are liquidity-constrained.
2. Campbell and Mankiw (1989) find estimates of the IES for the optimizing consumers close to 0·2. By redoing their regressions, we find that the corresponding standard error is also 0·2.
3. Papers that have adopted this specification are Campbell and Mankiw (1987, 1989, 1990, 1991), as well as Bayoumi and Klein (1994), Lusardi (1991), Shea (1993) and van Wincoop (1995).
4. Our approach of using regional panel data also has some advantages over the use of micro-level panel data, as in Runkle (1991) and Zeldes (1989). In particular, PSID data suffer from significant measurement error and a consumption measure that includes only food.
5. As shown by Christiano (1989), adding measurement error to the interest rate can lead to low and insignificant point estimates of the IES, which might account for the results by Hall (1988) and Campbell and Mankiw (1989). In Section III below we discuss various factors that may give rise to mismeasurement of asset returns.
6. In the last quarter of 1950 deseasonalized consumption of non-durables rose by 0·2% annually, leading to significant inflation and a real return on Treasury bills of −16% annually for that quarter.
7. The instruments used are always based on national data, and not on the sum of the 19 states. This is because retail sales data for 1976 and 1977 are available nationally, but not for individual states.
8. We focus mainly on the results with a tax rate set at 30% in order to be more easily comparable with previous studies.
9. When we estimate the reverse normalization of the Euler equation, which corresponds to regressing asset returns on expected consumption growth, we find estimates of 1/IES close to 1. These estimates are similar to the results of Hansen and Singleton (1983) and Summers (1982), who also estimate 1/IES instead of the IES. The difference in results coming from a choice of normalization is an indication of mis-specification.

10. The $\chi^2(k-1)$ test statistic is T times the R^2 from the regression of the residuals derived from estimating (1) by IV on the instruments. This test is appropriate in our context since we do not find any evidence of autocorrelation in error terms from the IV regression.
11. We examined whether non-separabilities between either consumption and leisure, present consumption and past consumption or the consumption of durable and nondurable goods could explain this observation. We did not find any evidence in favour of such a hypothesis.
12. Our results using aggregate data are very similar to those presented by Campbell and Mankiw. For example, based on the sample period 1953–86 and on the consumption of nondurable goods and services, we estimate the IES to be approximately 0·2 with a standard error of 0·2.
13. State-level price indices are not available. Therefore we use regional price indices from the Bureau of the Census. They are available for four large regions (the West, Midwest, Northeast and South).
14. Since we are interacting state-level dummies with aggregate variables, we have a total of $19\times k$ instruments, where k is the number of aggregate instruments. Consequently, we have $19\times k-2$ over identifying restrictions that can be tested.
15. One other possible means of obtaining more precise estimates of the IES would be to include state-specific variables in the instrument set. However, we found that such additions had virtually no effects on our results. For example, when we use three lags of the T-bill rate, g^c and g_i^c as instruments, our estimates are $\hat{\theta}=0{\cdot}505$ (0·143), $\hat{\lambda}=0{\cdot}561$ (0·082) $\widehat{\text{IES}}=1{\cdot}150$ (0·346),. Both point estimates and standard errors are close to the numbers reported in Table 4 with three lags of the Treasury bill and g^c as instruments. The advantage of using only aggregate instruments in Table 4 is that it allows for better comparison with the results of Table 5.
16. Note that the estimates presented in Table 4 were obtained by imposing that the coefficient on $E_{t-1}(r_t)$ be equal to the coefficient on $-E_{t-1}(\pi_{i,t}-\pi_t)$.
17. The estimates presented in Table 4 are also insensitive to adopting the hypothesis that the expected inflation differentials across regions are zero.
18. We found no evidence of serial correlation of the error term.
19. In order to see more formally that equation (7) is well identified, consider an example with two states. Let $E(g_{it}^y/\Omega_{t-2})=\Omega_{t-2}'\alpha_i$, where α_i is a $k\times 1$ vector (k is the number of aggregate variables in Ω). Let g^c be the vector that stacks the time-series of consumption growth in the two states. Then, ignoring the season dummies, we need to estimate the regression $g^c=X\beta+\varepsilon$, where

$$\beta=\begin{pmatrix}\dfrac{1-\lambda}{\sigma_\lambda}\gamma_1\end{pmatrix} \quad\text{and}\quad X=\begin{pmatrix}\Omega & 0\\ 0 & \Omega\end{pmatrix}\begin{pmatrix}I & \alpha_1\\ I & \alpha_2\end{pmatrix}.$$

Here Ω is the $T\times k$ matrix of aggregate variables and I is the $k\times k$ identity matrix. For the regression to be well identified, we need the matrix $\left(\begin{smallmatrix}I & \alpha_1\\ I & \alpha_2\end{smallmatrix}\right)$ to have rank $k+1$, which equals the number of parameters to be estimated. This is clearly the case as long as $\alpha_1\neq\alpha_2$, so that expected income growth is a different linear combination of aggregate variables in the two states.
20. The number of over-identification restrictions is $(19\times k)-(k+1)$, since we are estimating $k+1$ parameters.

REFERENCES

Bayoumi, T. and Klein, M. W. (1994). A provincial view of capital mobility. Working Paper, International Monetary Fund.

Campbell, J. and Mankiw, G. (1987). Permanent income, current income, and consumption. NBER Working Paper no. 2436.

—— and —— (1989). Consumption, income, and interest rates: reinterpreting the time series evidence. In O. Blanchard and S. Fischer (eds), *NBER Macroeconomic Annual 1989*, pp. 185–216. Cambridge, Mass.: NBER.

—— and —— (1990). Permanent income, current income, and consumption. *Journal of Business and Economic Statistics*, **8**, 265–79.

—— and —— (1991). The response of consumption to income. *European Economic Review*, **35**, 723–67.

Christiano, L. J. (1989). Comment. In O. Blanchard and S. Fischer (eds), *NBER Macroeconomic Annual 1989*, pp. 216–33. Cambridge, Mass.: NBER.

Flavin, M. A. (1981). The adjustment of income to changing expectations about future income. *Journal of Political Economy*, **89**, 974–1009.

Gravelle, J. (1994). *The Economic Effects of Taxing Capital and Income*. Cambridge, Mass.: MIT, Press.

HALL, R. E. (1978). Stochastic implications of the life-cycle–permanent income hypothesis: theory and evidence. *Journal of Political Economy*, **86**, 971–87.

—— (1988). Intertemporal substitution in consumption. *Journal of Political Economy*, **96**, 339–57.

—— and MISHKIN, F. S. (1982). The sensitivity of consumption to transitory income: estimates from panel data on households. *Econometrica*, **50**, 461–81.

HANSEN, L. P. and SINGLETON, K. J. (1982). Generalized instrumental variables estimation of nonlinear rational expectation models. *Econometrica*, **50**, 1269–86.

—— and —— (1983). Stochastic consumption, risk aversion, and the temporal behaviour of asset returns. *Journal of Political Economy*, **91**, 249–65.

JAPPELLI, T. (1990). Who is credit-constrained in the US economy? *Quarterly Journal of Economics*, **105**, 219–34.

LUSARDI, A. (1991). Euler equations in micro data: merging data from two samples. Working Paper, Dartmouth College.

MANKIW, N. G. (1981). The permanent income hypothesis and the real interest rate. *Economics Letters*, **7**, 307–11.

—— ROTEMBERG, J. J., and SUMMERS, L. H. (1985). Intertemporal substitution in macroeconomics. *Quarterly Journal of Economics*, **100**, 225–51.

MARIGER, R. P. (1986). *Consumption Behavior and the Effects of Fiscal Policies.* Cambridge, Mass.: Harvard University Press.

NELSON, C. R. and STARTZ, R. (1990). The distribution of the instrumental variables estimator and its t-ratio when the instrument is a poor one. *Journal of Business*, **63**, 125–40.

RUNKLE, D. E. (1991). Liquidity constraints and the permanent income hypothesis: evidence from panel data. *Journal of Monetary Economics* **27**, 73–98.

SHEA, J. (1993). Myopia, liquidity constraints, and aggregate consumption: a simple test. Working Paper, University of Wisconsin.

SUMMERS, L. H. (1981). Capital taxation and accumulation in a life-cycle growth model. *American Economic Review*, **71**, 533–44.

—— (1982). Tax policy, the rate of return, and savings. Working Paper no. 995, NBER.

—— (1984). The after-tax rate of return affects private savings. *American Economic Review, Papers and Proceedings*, **74**, 249–53.

VAN WINCOOP, E. (1995). Regional risksharing. *European Economic Review*, **37**, 1545–67.

WORKING, H. (1960). Note on the correlation of first differences of averages in a random chain. *Econometrica*, **29**, 916–18.

ZELDES, S. P. (1989). Consumption and liquidity constraints: an empirical investigation. *Journal of Political Economy*, **97**, 305–46.